Simple Forecasting Heuristics that Make us Smart: Evidence from Different Market Experiments

Mikhail Anufriev and Cars Hommes and Tomasz Makarewicz*

25 February 2016

Abstract

We study a model in which individual agents use simple linear first order price forecasting rules, adapting them to the complex evolving market environment with a smart Genetic Algorithm optimisation procedure. The novelties are: (1) a parsimonious experimental foundation of individual forecasting behaviour; (2) an explanation of individual and aggregate behaviour in four different experimental settings, (3) improved one-period and 50-period ahead forecasting of lab experiments, and (4) a characterisation of the mean, median and empirical distribution of forecasting heuristics. The median of the distribution of GA forecasting heuristics can be used in designing or validating simple Heuristic Switching Model.

JEL codes: C53, C63, C91, D03, D83, D84.

Keywords: Expectation Formation, Learning to Forecast Experiment, Genetic Algorithm Model of Individual Learning.

---

*Anufriev: University of Technology Sydney, Business School, Ultimo NSW 2007, Australia, Mikhail.Anufriev@uts.edu.au. Hommes: CeNDEF, University of Amsterdam, Valckenierstraat 65, 1018 XE Amsterdam, C.H.Hommes@uva.nl. Makarewicz: CeNDEF, University of Amsterdam, Valckenierstraat 65, 1018 XE Amsterdam, T.A.Makarewicz@uva.nl. The authors are thankful for stimulating discussions to Blake LeBaron and Murat Yildizoglu, as well as to all the participants of seminars at WEHIA 2015 and CEF 2015 conferences. We gratefully acknowledge the financial support from NWO (Dutch Science Foundation) Project No. 40611142 "Learning to Forecast with Evolutionary Models". The research leading to these results has also received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement Integrated Macro-Financial Modeling for Robust Policy Design (MACFINROBODS, grant no. 612796). Mikhail Anufriev acknowledges financial support from the Australian Research Council through Discovery Project (DP140103501).

1 Introduction

Expectations are a cornerstone of many dynamic economic models. In this paper we address the question of how individuals form expectations in a complex environment by fitting a novel learning model with explicit heterogeneity in expectations to the set of experiments reported in the literature. The traditional literature after Muth (1961), Lucas (1972) and others emphasises the Rational Expectations (RE) hypothesis, which states that the expectations of all agents have to be model consistent.[1] This hypothesis provides an elegant and universally applied solution concept for an economic expectations feedback system. Most economists would assume that agents are rational enough to avoid systematic errors and so RE remain a commonly used modelling tool to derive the aggregate market dynamics. Nevertheless, two problems with RE approach are widely recognised. First, the RE impose strict informational and computational assumptions. Second, they are at odds with many empirical studies. These shortcomings have already been highlighted in the early works of Pesaran (1987), Sargent (1993) and others, and recent evidence both from surveys[2] and controlled laboratory experiments[3] adds further support to the criticism of the RE. In particular, the outcomes of many Learning-to-Forecast experiments which we use in this paper contradict the RE hypothesis.

The empirical deficiencies of the RE benchmark signifies the importance of learning and bounded rationality. Recent years witnessed a surge of the adaptive learning literature in macroeconomics relaxing the strong informational assumptions underpinning RE.[4] This literature, whose

[1]See Wagener (2014) for a discussion of different versions of RE hypothesis and their theoretical and empirical obstacles.

[2]Perhaps, the most prominent recent example on failure of REs comes from the housing market in the US, which in the last decade exhibited first a boom and then a collapse. Case, Shiller, and Thompson (2012) conduct a survey of households' expectations about changes in their home value over the next years and reject the RE hypothesis. They conclude that people's expectations are consistent with trend-extrapolation and that people systematically misjudge the long-term value of their houses. Similar effects were observed with expectations before the previous housing bubble in the late 80's, see Goodman and Ittner (1992). In fact, economic history knows many similar examples of prolonged asset misvaluation, see, e.g., Reinhart and Rogoff (2009) and Kindleberger and Aliber (2011). Many studies use surveys of inflation expectations. For example, Malmendier and Nagel (2009) studies the responses in the Reuters/Michigan Survey of Consumers and find support for the backward looking, learning from experience model. Branch (2004) shows that the responses are consistent with a mixed model where non-rational expectations (such as naive or adaptive) have a high weight. A similar conclusion is reached in Nunes (2010) who uses, instead, the Survey of Professional Forecasters.

[3]Recent papers with a special focus on inflationary expectations include Adam (2007), Pfajfar and Žakelj (2014) and Assenza, Heemeijer, Hommes, and Massaro (2014). Despite differences in design of the experiment and underlying macroeconomic models, all of them reject the RE hypothesis. See Duffy (2014) for an overview of macroeconomic experiments.

[4]Non-learning streams of macroeconomic literature on bounded rationality include the rational inattention approach, see Sims (2010) for a comprehensive review, the rational or "near-rational" beliefs approach, see Woodford (2010) and Kurz and Motolese (2011), and the eductive approach of Guesnerie (2005). In the rational inattention literature agents do not react on all relevant information quickly but instead process information at some finite rate. Similarly to the adaptive learning models it induces sluggish behaviour which then can be translated into

early contributions are discussed in Evans and Honkapohja (2001), studies the emerging dynamics in the case when agents do not know the actual dynamic laws of the economy but only estimate their perceived model on observable data. The focus is on whether the RE equilibrium can emerge as the long-run outcome of such learning process, or the dynamics can become more complicated as in Bullard (1994) when it converges to a limit cycle. Adaptive learning and other streams of literature on bounded rationality in macroeconomics can reproduce some empirical regularities and are well-suited to address policy issues, but still face many empirical and theoretical obstacles as discussed, e.g., in Woodford (2013). Also their behavioural assumptions remain to be very demanding and the phenomenon of heterogeneity is often ignored. However, heterogeneity in expectations is one of the most recurring findings in surveys and experimental studies.[5]

Starting from the data-driven paper by Allen and Taylor (1990), heterogeneity in expectations is recognised as a driving force behind the bubbles and crashes in financial markets. Influential paper of Brock and Hommes (1997) proposed a new framework where agents entertain a set of possible forecasting models and select the model that performs best. A large and constantly growing body of literature uses this framework with heterogeneous and time-varying expectations. There are application to asset prices (Brock and Hommes, 1998; Anufriev and Panchenko, 2009; Branch and Evans, 2010), business cycle dynamics and monetary policy (Branch and McGough, 2010; Branch and McGough, 2011; Anufriev, Assenza, Hommes, and Massaro, 2013), dynamics in the housing market (Bolt, Demertzis, Diks, Hommes, and Van der Leij, 2014). Moreover, evidence for this predictor-selection learning model can be seen in survey data (Branch, 2004), estimated financial and macro models (Boswijk, Hommes, and Manzan, 2007; Cornea, Hommes, and Massaro, 2013), and, most importantly for this paper, in a series of experiments reviewed in Hommes (2011).[6]

This paper is inspired by and belongs to this line of research. More precisely, we address here an important question: *how exactly do people invent, reinforce and update their forecasting rules in a complex world?* We address the 'wilderness of bounded rationality' problem: there is a myriad of possible learning or other behavioural mechanisms with varied restrictions on human

sluggishness of economic variables. 'Near-Rational' expectations allow distortions of expectations with respect to the RE case within certain bounds. Eductive learning means that agents' expectations are consistent with the actual law of motion and *some* common knowledge assumption about expectations of others. All these threads of literature produce a richer set of equilibrium dynamics than RE.

[5] Heterogeneity in expectations is reported, for instance, in all the references in footnotes 2 and 3.

[6] There is a connection of this literature with experimental research in psychology (Tversky and Kahneman, 1974; Kahneman, 2011; Gigerenzer and Todd, 1999) and game theory (Erev and Roth, 1998; Camerer and Ho, 1999) suggesting that people rely on relatively simple behavioural rules in their decision making and that an important ingredient of their learning is reinforcement of the successful rules and forgetting the less successful. It is worth to notice that in the game theoretical studies there is also evidence of using more sophisticated belief-based learning, see, e.g., Feltovich (2000). However, in the experiments which we discuss in this paper there is not much space for belief learning, because the payoffs as well as the game-theoretical structure are not explicitly explained to the subjects. This would, actually, be the case in most real situations, where the law of motion of the market is unknown.

memory and computational capabilities. In the literature these range from simple linear models in the spirit of adaptive expectations in macroeconomics (Evans and Ramey, 2006), through models with switching between heterogeneous expectations (Brock and Hommes, 1997), to evolutionary learning mechanisms (Arifovic, Bullard, and Kostyshyna, 2013). Which of them shall be used? We propose a model that incorporates all these features and use the results of recent Learning-to-Forecast (LtF) experiments to fit the model to the data. Our findings are based on extensive simulations, but our model has a parsimonious and analytically tractable counterpart to it, namely the Heuristics Switching Model of Anufriev and Hommes (2012) and Brock and Hommes (1997).

The model which we present in this paper is a model of individual learning of forecasting heuristics based on Genetic Algorithms updating. Heterogeneous agents use linear first order price forecasting rules *with only two parameters*, adapting them to the current environment with a smart Genetic Algorithm optimisation procedure. In this sense agents use simple forecasting heuristics that make them smart (Gigerenzer and Todd, 1999). We fit our GA model to individual and aggregate data from LtF experiments to study the mean, the median and the distribution of individual forecasting heuristics in *four* different experimental laboratory settings.

Learning-to-Forecast experiments offer a simple laboratory testing ground for adaptive learning mechanisms (Lucas, 1986). These controlled experimental economies have a straightforward and unique fundamental equilibrium consistent with RE. As in real markets, subjects observe the realised prices and their own past individual predictions, but not the history of other subjects' predictions, and are not informed about the exact law of motion of the economy. The outcomes of many LtF laboratory experiments contradict the RE hypothesis, see the review in Hommes (2011). The experiments in Hommes, Sonnemans, Tuinstra, and van de Velden (2005), henceforth **HSTV05**, showed that the subjects can coordinate on oscillating and serially correlated time series, and that convergence to the fundamental equilibrium happens only under severe restrictions on the underlying law of motion. Further experiments in Heemeijer, Hommes, Sonnemans, and Tuinstra (2009), henceforth **HHST09**, and Bao, Hommes, Sonnemans, and Tuinstra (2012), henceforth **BHST12**, demonstrated that the expectations feedback structure plays crucial role. Negative feedback systems (i.e., where more optimistic forecasts lead to lower market prices, as in supply driven commodity markets) tend to generate convergence to the fundamental equilibrium rather easily, while positive feedback systems (i.e., where more optimistic forecasts lead to higher market prices, as in speculative asset markets) typically generate behaviour with the price oscillating around the fundamental equilibrium dynamics.

Anufriev and Hommes (2012) propose a parsimonious Heuristic Switching Model (HSM) to provide an explanation for different types of aggregate behaviour, including both convergence and oscillations, observed in the LtF experiments of **HSTV05**. The basic idea of the model is that agents have a small set of simple forecasting heuristics (rules of thumb, such as adaptive or trend extrapolating expectations) and gradually switch to relatively better performing rules as in Brock and Hommes (1997). Both in-sample and out-of-sample performance of the HSM is usually better than for the RE model and several other homogeneous and heterogeneous expectation models.

However, the HSM has some shortcomings. First, the small set of given heuristics cannot fully account for within treatment individual heterogeneity, observed in the experiments. Second, different experiments may require the HSM to utilise (for a better fit) different sets of heuristics with different parametrisation, see, e.g., Anufriev, Hommes, and Philipse (2013). It is unclear why the subjects would use only those particular forecasting heuristics and how they would learn them in the first place.

This paper addresses those weaknesses of the HSM by using Genetic Algorithms (GA). GA are a prominent tool in the economic literature to model individual learning (see, e.g., Sargent, 1993 and Dawid, 1996). From the very first economic application in Arifovic (1994), GA were used to model both the social and individual learning and to explain the results of experiments with human subjects. Areas of GA applications include the overlapping generation monetary economies (Arifovic, 1995), exchange rate volatility (Arifovic, 1996; Lux and Schornstein, 2005) and production level choices in a cobweb producers economy (Dawid and Kopel, 1998). Recently in a related paper Hommes and Lux (2013) investigate a model in which agents use GA to optimise a forecasting heuristic (instead of directly optimizing a prediction) and, much like the actual subjects in the LtF experiments, cannot observe each others behaviour or strategies. The authors replicate the distribution (mean, variance and first order auto-correlation) of the predictions and prices of the cobweb experiments by Hommes, Sonnemans, Tuinstra, and van de Velden (2007) and van de Velden (2001) (henceforth **HSTV07** and **V01**, respectively).

The main contribution of our paper is that it provides a general explanation of four different LtF experiments simultaneously and at different levels of aggregation. Agents forecast prices using a large set of heuristics from a simple, general class. The agents then *independently* use GA to update and select the heuristics based on their relative success. Agents thus learn to use simple heuristics that make them smart. This explicitly accounts for individual learning and endogenous heterogeneity observed in the experiments. Monte Carlo simulations of this model provide insight in the mean, median and the distribution of forecasting heuristics.

More specifically, this paper has four contributions. The first is that the computational model presented here is *empirically validated* and *parsimonious*. It is as parsimonious as a computational model can be because its heuristic space is based on a simple linear first order rule with *only two free parameters*. It is empirically validated because this simple rule is a mixture of adaptive and trend extrapolating heuristics, consistent with the individual forecasting behaviour estimated in **HSTV05** and **HHST09**. The second contribution is that we apply the same GA learning model to explain *four different* LtF experiments: (1) the simple, linear positive/negative feedback system with small shocks (**HHST09**); (2) the linear positive/negative price-expectations feedback system with unexpected large shocks to the fundamental price (**BHST12**); (3) a stable/unstable cobweb producers economy (**HSTV07**; **V01**), used also in the GA model of Hommes and Lux (2013); and (4) a non-linear positive feedback asset pricing economy, where the subjects are asked for two-period ahead predictions (**HSTV05**).

The third contribution of the paper is that our model is able to capture the dynamics at both the *aggregate* and the *individual* level for different experimental settings. The GA model replicates the long-run behaviour of the experimental prices, as well as the individual forecasting decisions. We are also the first to evaluate the out-of-sample predictive power of the model by means of a simple Sequential Monte Carlo technique. We find that depending on the experiment, our model is comparable or better than the HSM in terms of predicting both *prices* and *individual price forecasts* one period ahead. This is an important contribution to the literature on heterogeneous agent models, which usually focuses only on a model's fit to aggregate stylised facts.

Finally, the fourth contribution is that the Monte Carlo studies of the GA model enable us to characterise the emerging *median forecasting behaviour*, together with its corresponding confidence bounds, in various experimental settings. The GA simulations thus (1) provide a solid motivation for describing the LtF experimental dynamics in terms of simple 'stylised' heuristics, and (2) guide the specific choice of these heuristics for a particular experimental market. This yields natural *empirical micro-foundations* for heterogeneous expectations models such as HSM.

The paper is organised as follows. In Section 2 we present the set-up and findings of the LtF experiments and briefly discuss the HSM by Anufriev and Hommes (2012). In Section 3 we introduce our GA model and fit it to the experimental set-up of **HHST09**. Section 4 investigates three other experimental settings. Finally, the concluding Section 5 gives an overview of the results and suggestions for future research. The appendices contain GA simulation details and various robustness checks.

2 Learning to Forecast and Heuristic Switching

Consider an experimental market with I subjects indexed by $i \in \{1, \ldots, I\}$ who at each period t forecast the price of a certain good. The subjects are informed that they act as forecasting consultants for firms and are rewarded only for the accuracy of the predictions. The relationship between the prices and predictions is summarised by a law of motion of the form

$$p_t = F(p^e_{1,t}, \ldots, p^e_{I,t}) + \varepsilon_t, \tag{1}$$

where the realised price p_t is a function of all individual forecasts $p^e_{i,t}$ and a small white noise noise term ε_t. The function $F(\cdot)$ is obtained from the market clearing condition with aggregate supply and demand derived from optimal (i.e., profit/utility maximizing) choices of firms, consumers or investors, given the subjects' individual forecasts.[7] Define the fundamental price p^f as the steady state self-fulfilling (RE) prediction: $p^f = F\left(p^f, \ldots, p^f\right)$. In all examples below the fundamental price exists and is unique.

The subjects in the LtF experiments are endowed with limited information about the market. They are told that their predictions affect realised prices, but the feedback's description is only

[7]The LtF experiments focus only on the forecasting behaviour and abstract from other considerations (e.g., trading) by assuming that the subjects' actions are rational conditional on the submitted forecast. See Hommes (2011) for an in-depth discussion on the methodology of the LtF experiments.

qualitative. Subjects do not know exact number and nature of other participants and are not explicitly informed about the fundamental price.[8] The forecasts are submitted repeatedly for a number of periods and the experimental screen shows the past realised prices and the past predictions and earnings of the participant.

HHST09 study the subjects' behaviour conditional on whether the market is built upon negative or positive feedback. A typical example of positive feedback is a stock exchange: optimistic investors will buy more stock and due to increased demand the stock price will go up. In this sense the investor sentiments are self-fulfilling. Negative feedback arises in a supply driven market where producers face a lag in production. If they expect a high price, they will increase production and the market clearing price will go down. **HHST09** run two treatments with linear specifications of (1):

$$\text{(2)} \qquad \textbf{Negative feedback:} \quad p_t = p^f - \frac{20}{21}\left(\bar{p}^e_t - p^f\right) + \varepsilon_t \, ,$$

$$\text{(3)} \qquad \textbf{Positive feedback:} \quad p_t = p^f + \frac{20}{21}\left(\bar{p}^e_t - p^f\right) + \varepsilon_t \, ,$$

where $\bar{p}^e_t = \sum_{i=1}^{I} p^e_{i,t}/I$ is the average prediction of all individuals at period t and p^f is the fundamental price. The experiment run for 50 periods for groups of $I = 6$ subjects. The two treatments are symmetrically opposite, with the same fundamental price $p^f = 60$, and the dampening factors of the same absolute value but opposite signs.[9] Under homogeneous naive expectations (i.e., $\bar{p}^e_t = p_{t-1}$) the fundamental price for both treatments is a stable steady state of dynamics.

The aggregate price dynamics in the two feedback treatments were very different, see Figs. 1a and 1b (two lower panels are explained in Section 3). Under the negative feedback after a short volatile phase of $7-8$ periods, the price converged to the fundamental value $p^f = 60$, after which the subjects' forecasts coordinated on the fundamental price as well. In most of the positive feedback groups[10] persistent price oscillations arose where the price overshot and undershot p^f. In spite of the price oscillations the subjects' forecasts became close to each other after already $2-3$ periods and remained so until the end of the experiment. In positive feedback markets subjects' forecasts are thus strongly coordinated, but not on the fundamental price. It is the almost self-fulfilling character of the near-unit root positive feedback system that allows subjects to coordinate on trend following behaviour, which results in price oscillations (Hommes, 2013).

[8]The fundamental price can sometimes be inferred from the experimental instructions. For example, in the asset pricing (positive feedback) treatment in **HHST09** the fundamental price is equal to the present value of the future dividends, which is the ratio of the average dividend to the interest rate. Both variables were provided to the subjects, but most of the individual first period predictions were not at the fundamental.

[9]In an asset pricing market, the near unit root coefficient 20/21 arises from a realistic discount factor. To have symmetric treatments, the factor in the negative feedback was set to $-20/21$.

[10]There were 6 experimental groups for the negative feedback treatment with very similar price dynamics. Fig. 1a is a typical example. There were 7 experimental groups for the positive feedback treatment and in 4 of them price oscillated. Fig. 1b is a typical example for the oscillating group. Even when price converged (which happened for 3 groups), it did so only towards the end of the experiment.

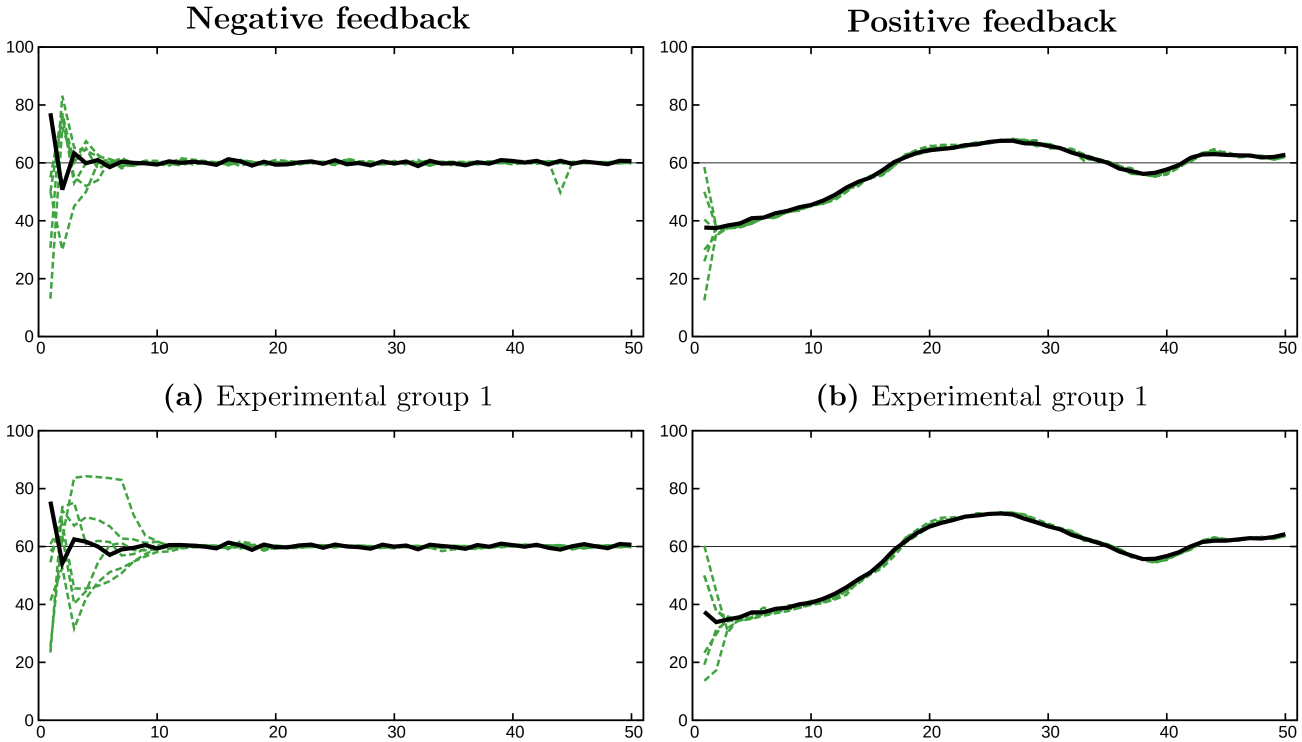

(a) Experimental group 1

(b) Experimental group 1

(c) Simulation of GA model

(d) Simulation of GA model

Figure 1: HHST09 experimental groups (*upper panels*) and sample 50-period ahead simulations of **GA-P1** model with random initial predictions (*lower panels*). Black thick line shows the price, green dashed lines show 6 individual predictions. The long-run dynamics of the model is close to the experiment both under negative (*left*) and positive (*right*) feedback.

HHST09 described the subjects' forecasting behaviour in the experiment with the first-order rule (FOR):

$$p^e_{i,t} = \alpha_1 p_{t-1} + \alpha_2 p^e_{i,t-1} + \alpha_3 p^f + \beta(p_{t-1} - p_{t-2}), \tag{4}$$

for $\alpha_1, \alpha_2, \alpha_3 \geqslant 0$, $\alpha_1 + \alpha_2 + \alpha_3 = 1$, $\beta \in [-1, 1]$. Rule (4) is an anchor and adjustment rule extrapolating a price change from an anchor given by a weighted average of the previous price and forecast, and the fundamental price p^f.[11] **HHST09** estimated this simple rule separately for each subject, fitting well the forecasting behaviour of around 60% of the individuals.

HHST09 found a significant variability in terms of the individual forecasting, within the same treatment, and even more so between treatments. The main difference appears to lie in the trend extrapolation, which is popular under positive feedback ($\beta > 0$) and disregarded under negative feedback ($\beta \approx 0$). This shows that a model with a homogeneous forecasting rule may explain one of the two treatments, but not both at the same time.

[11] Under RE, the FOR in (4) should be specified with $\alpha_1 = \beta = 0$, and subjects always predict the fundamental price, $p^e_{i,t} = p^f = 60$.

These findings led Anufriev, Hommes, and Philipse (2013) to investigate the Heuristic Switching Model (HSM), in which the subjects are endowed with two prediction heuristics

$$\textbf{adaptive expectations:} \quad p^e_{i,t} = \alpha p_{i,t-1} + (1-\alpha) p^e_{i,t-1} \quad \text{with } \alpha \in [0,1] \,,$$
$$\textbf{trend extrapolation:} \quad p^e_{i,t} = p_{i,t-1} + \beta (p_{t-1} - p_{t-2}) \quad \text{with } \beta \in [-1,1] \,.$$

Both heuristics are special cases of the first-order rule (4), and for the benchmark specification the values $\alpha = 0.75$ and $\beta = 1$ were used. The idea of the HSM model is that the subjects can at any time use any of the two heuristics, but tend to focus on the rule with a higher relative past performance. The dynamics of the HSM are similar to the experimental outcome. Under positive feedback agents quickly coordinate to use the trend extrapolation heuristic, leading to persistent price oscillations and thus self-confirming trend chasing predictions. In contrast, under negative feedback the trend extrapolation rule performs poorly and agents tend to switch to adaptive expectations. This does not allow a fast coordination, but eventually causes the price to converge to the fundamental price.

HSM captures the essence of the aggregate forecasting behaviour in the LtF experiment by successfully replicating the results of **HHST09** in both treatments in a stylised fashion. It leaves open, however, important questions about the origins of the forecasting heuristics. It does not say where do those particular rules come from and is silent about which rules (and how many of them) should be used in a more general setting. Moreover, the HSM cannot fully account for the within-treatment heterogeneity of predictions and hence does not explain the experiment at the individual level. To overcome these drawbacks, we will introduce a model with explicit individual heuristic-learning through Genetic Algorithms.

3 The Genetic Algorithms model

Genetic Algorithms (GA) form a class of numerical stochastic maximisation procedures that mimic the evolutionary operations with which DNA of biological organisms adapts to the environment. GA were introduced to solve 'hard' optimisation problems, which may involve non-continuities or high dimensionality with complicated interrelations between the arguments. They are flexible and efficient and so are often used in computer sciences and engineering. See Haupt and Haupt (2004) for technical discussion and Dawid (1996) for applications in economics.

A GA routine starts with a population of random trial solutions to the problem. Individual trial arguments are encoded as binary strings (strings of ones and zeros), so-called 'chromosomes'.[12] They are retained into the next iteration with a probability that increases with their relative functional value (performance or 'fitness'). This so-called **procreation** operator means that with each iteration, the trial arguments are likely to have a higher functional value, i.e., be

[12]We use a binary representation for the sake of parsimony. The real number variant of the GA requires additional parametrisation, such as distribution of the mutation changes.

'fitter'. On top of procreation, GA use three evolutionary operators that allow for an efficient search through the problem space: mutation, crossover and election, where the last operator was introduced in the economic literature in Arifovic (1995).

Mutation. At each iteration, every bit in each chromosome has a small probability to mutate, in which case it changes its value from zero to one and vice versa. The mutation operator utilises the binary representation of the arguments. A single change of one bit at the end of the chromosome leads to a minor, numerically insignificant change of the argument. But with the same probability a mutation of a bit at the beginning of the chromosome can occur, which changes the argument drastically. With this experimentation, GA can easily search through the whole parameter space and have a good chance of shifting from a local maximum towards the region containing the global maximum.

Crossover. Pairs of arguments can, with a predefined probability, exchange predefined parts of their respective binary strings. In practice, the crossover is set to exchange subset of the arguments. For example, if the objective function has two arguments, crossover would swap the first argument between pairs of trial arguments. This allows for experimentation in terms of different mixtures of arguments.

Election. This operator screens inefficient outcomes of the experimentation phase by transmitting the new chromosomes (selected from the old generation and treated with mutation and crossover) into the new generation only if their fitness is greater than that of the original chromosome. This ensures that once the routine finds the global solution, it will not diverge from it due to unnecessary experimentation.

These four operators have a straightforward economic interpretation for a situation in which the agents optimise their behavioural rules such as forecasting heuristics. The procreation means that – as in the case of HSM – people focus on better solutions (or heuristics). The mutation and crossover are experimentation with the heuristics' specifications, and finally the election ensures that the experimentation does not lead to suboptimal heuristics.[13]

3.1 Model specification

In GA model we follow the experimental information structure where the subjects did not have an access to the predictions and performances of other subjects and, therefore, could learn only individually. We populate the price-expectation feedback economy (1) by $I = 6$ GA agents. At the beginning of each period t an agent i submits the forecast using one of $H = 20$ specifications[14]

[13]An important additional condition for a GA routine is that it requires a predefined interval for each parameter. For the example with updating behavioural rules through GA it means that we confine them to some predetermined, finite grid of heuristics.

[14]Experimentation with $H = 10$, $H = 50$ and $H = 100$ yield similar results.

of a general linear forecasting rule. Different specifications available to this agent in the period t are indexed by h and the agent's forecast $p^e_{i,h,t}$ of price p_t conditional on picking specification h is described by

$$p^e_{i,h,t} = \alpha_{i,h} p_{t-1} + (1 - \alpha_{i,h}) p^e_{i,t-1} + \beta_{i,h}(p_{t-1} - p_{t-2}) \,, \tag{5}$$

where $p^e_{i,t-1}$ denotes the prediction of price p_{t-1} submitted by agent i in period $t-1$. Our model has *empirical validation* because rule (5) is a version of the general FOR estimated in **HHST09** on individual data.[15]

Heuristic (5) depends on two parameters, $\alpha_{i,h}$ (price weight) and $\beta_{i,h}$ (trend extrapolation coefficient), and 20 specifications differ only in the values of these coefficients. Importantly, these parameters are modelled as changing over time, as the agents repeatedly fine-tune the rule to adapt to the specific market conditions. For example, in an asset pricing market it may pay off to extrapolate the price trend and agents would try to find the optimal value of β, depending on current trend. This learning is embodied as a heuristic optimisation with the GA procedure, and introduces the individual heterogeneity to the model which is absent in the HSM or in any homogeneous expectations model.

Define $\mathbf{H}_{i,t}$ as the set of $H = 20$ heuristics of agent i at time t, where heuristic h is specified as a pair of parameters $(\alpha_{i,h}, \beta_{i,h}) \in \mathbf{H}_{i,t}$. Each pair is a 'chromosome' represented as a binary string of length 40 with 20 bits per coefficient. The bounds for the coefficients are chosen as follows. The price weight belongs to the unit interval $[0, 1]$. For the trend extrapolation coefficient we report two parametrisations, depending on the bounds. Under Parametrisation 1 (denoted as **GA-P1**), the restriction is symmetric, $\beta_{i,h} \in [-1.1, 1.1]$. Under Parametrisation 2 (denoted as **GA-P2**), the restriction is $\beta_{i,h} \in [0, 1.1]$, i.e., contrarian rules are not allowed.[16]

The heuristics are updated *independently* for each agent by GA evolutionary operators, see Table 1 for the specific parameter values. The updating is based on the relative forecasting performances of the heuristics. The experimental payoffs decreased with the mean squared error (MSE) of the prediction. Accordingly, at time t for every heuristic from $\mathbf{H}_{i,t}$ we compute the (hypothetical) mean squared error, $\text{MSE}_{i,h} = (p^e_{i,h,t} - p_t)^2$, and apply the logit transformation[17]

[15] In comparison with the estimated FOR (4), in rule (5) the coefficient in front of the fundamental price (which can be thought of as an anchor) is set to 0. We experimented with the full FOR with the anchor specified as either (i) the fundamental price p^f or (ii) the average realised price so far. Neither specification could closely match the experimental dynamics of the positive feedback treatment, where the anchor dampens the oscillations, see Appendix E.1. This is consistent with the fact that in the estimated rules of **HHST09** under positive feedback the anchor weight α_3 is typically insignificant. GA forecasting model in Hommes and Lux (2013) explains the **HSTV07** experiment with even simpler rule $p^e_{i,h,t} = \alpha_{i,h} + \beta_{i,h}(p_{t-1} - \alpha_{i,h})$. However, our simulations showed that this rule does not fit the positive feedback experiment well.

[16] Heuristics with negative extrapolation coefficient are often called the contrarian strategies. **HHST09** found only two subjects with such contrarian rules, but for the sake of completeness we report both parametrisations.

[17] We use the logit and not the power transformation as in Hommes and Lux (2013) to have a clear link with the HSM literature.

Parameter	Notation	Value
Number of agents	I	6
Number of heuristics per agent	H	20
Allowed α, price weight	$[\alpha_L, \alpha_H]$	$[0, 1]$
Allowed β, trend extrapolation coefficient		
Parametrisation 1 (**GA-P1**)	$[\beta_L, \beta_H]$	$[-1.1, 1.1]$
Parametrisation 2 (**GA-P2**)	$[\beta_L, \beta_H]$	$[0, 1.1]$
Number of bits per parameter	$\{L_1, L_2\}$	$\{20, 20\}$
Mutation rate	δ_m	0.01
Crossover rate	δ_c	0.9
Performance measure	$U(\cdot)$	$\exp(-\text{MSE}(\cdot))$

Table 1: Values of parameters used by the Genetic Algorithms agents.

to define the normalised performance ('fitness') of heuristic h that agent i uses in individual learning:

(6) $$\Pi_{i,h} = \frac{\exp(-\text{MSE}_{i,h})}{\sum_{k=1}^{H} \exp(-\text{MSE}_{i,k})} .$$

Before the market starts to operate, the set $\mathbf{H}_{i,1}$ of agents' heuristics is initialised at random. Every agent samples 800 initial bits (20 initial heuristics with 2 parameters, each encoded by 20 bits) independently as 0 or 1 with equal probability. Two other aspects of initialisation should be specified. First, in initial periods, with no past prices and predictions, the heuristics cannot be used. Here we sample random predictions from an exogenous distribution.[18] Second, in the first period when the heuristics can be used for prediction, their performances are still undefined. In this case, every agent picks one of own 20 heuristics with equal probabilities.

Once the agents have enough observations to use their heuristics and evaluate their performances, the timing at period t is as follows:

1. Agents forecast price, the market price p_t is realised according to (1), agents observe it;

2. Agents *independently* update their heuristics using one GA iteration. The criterion function is $\Pi_{i,h}$ computed in (6) from the *hypothetical* MSE's of different heuristics in predicting price p_t. To be specific, agent i uses four evolutionary operators:

   (a) *procreation*: agent samples H so-called 'child' heuristics from the pool of 'parent' heuristics, $\mathbf{H}_{i,t}$, with replacement using $\Pi_{i,h}$ as the corresponding probabilities;

[18] As it will be clear later, we use both experimental and random initial predictions. Specification of the distribution of the latter is important, since in the experiment the average initial prediction affected the group dynamics (cf., Anufriev, Hommes, and Philipse, 2013). Appendix B provides the details for every experiment.

(b) *mutation*: each bit of each child heuristic has probability $\delta_m = 0.01$ to switch its value;

(c) *crossover*: each pair of child heuristics has probability $\delta_c = 0.9$ to swap the last twenty bits (it corresponds to exchanging β's);

(d) *election*: each child heuristic (possibly modified after mutation and crossover) is compared in terms of MSE with a randomly chosen parent heuristic. The child joins $\mathbf{H}_{i,t+1}$ if it outperforms the parent. Otherwise, the parent is passed to $\mathbf{H}_{i,t+1}$.

3. Now, when the new sets $\mathbf{H}_{i,t+1}$ are formed, period $t+1$ starts.

4. With probabilities as in (6), but now based on the hypothetical MSE's of heuristics from the new pool, each agent i stochastically picks one heuristic from $\mathbf{H}_{i,t+1}$. Agent uses this heuristic to generate prediction $p^e_{i,h,t+1}$. The algorithm now returns to step 1.

For initialisation of the **HHST09** experiment, GA agents use the first-order rule (assuming a zero trend $\Delta p_1 = 0$ for each rule) already in the second period (with chromosome bits chosen randomly), and start to update their heuristics sets in the third period.

While the last step – the choice of heuristic – is the same as in the HSM, there are two important differences between HSM and our GA model. First, the set of heuristics evolves over time with $\mathbf{H}_{i,t} \neq \mathbf{H}_{i,t+1}$. As a result, the heuristics have time varying parameters adapted to the specific market dynamics. Second, this learning operates through a stochastic GA procedure and is independent between the agents. In practice thus the agents will learn different heuristics and remain heterogeneous with $\mathbf{H}_{i,t} \neq \mathbf{H}_{j,t}$ when $i \neq j$.

3.2 50-period ahead simulations

The first test for the fit of our GA model to the experimental data are 50-period ahead simulations for the **HHST09** experiment.[19] Thus we compare the *long-run* model dynamics with the experimental data.[20]

In the first Monte Carlo (MC) exercise, we begin by sampling the first predictions from an exogenous distribution calibrated from all experimental first period forecasts, see Appendix B for details.[21] Then the model is simulated for 50 periods with *no other information* from the

[19] All simulations were written in Ox matrix algebra language (Doornik, 2007) and are available upon request.

[20] We treat one of the groups in the positive feedback treatment as an outlier and omit this group from the analysis. In this group, in period 6 one of the subjects 'out of the blue' submitted the forecast which was ten times larger than the previous price and own forecasts. This destabilised the market for a number of periods. In total, we focus on six positive feedback and six negative feedback treatment groups.

[21] In the first period the subjects in the LtF experiments have limited, mostly qualitative information about the market, and have not yet interacted with each other. Their initial forecasts are necessarily more a matter of a guess than a reasoned out prediction. Thus, we treat these as coming from an exogenous distribution (see also Diks and Makarewicz, 2013, for a comprehensive discussion), and use the experimental one as we are interested in the model's dynamics fit to the experimental price and forecast paths.

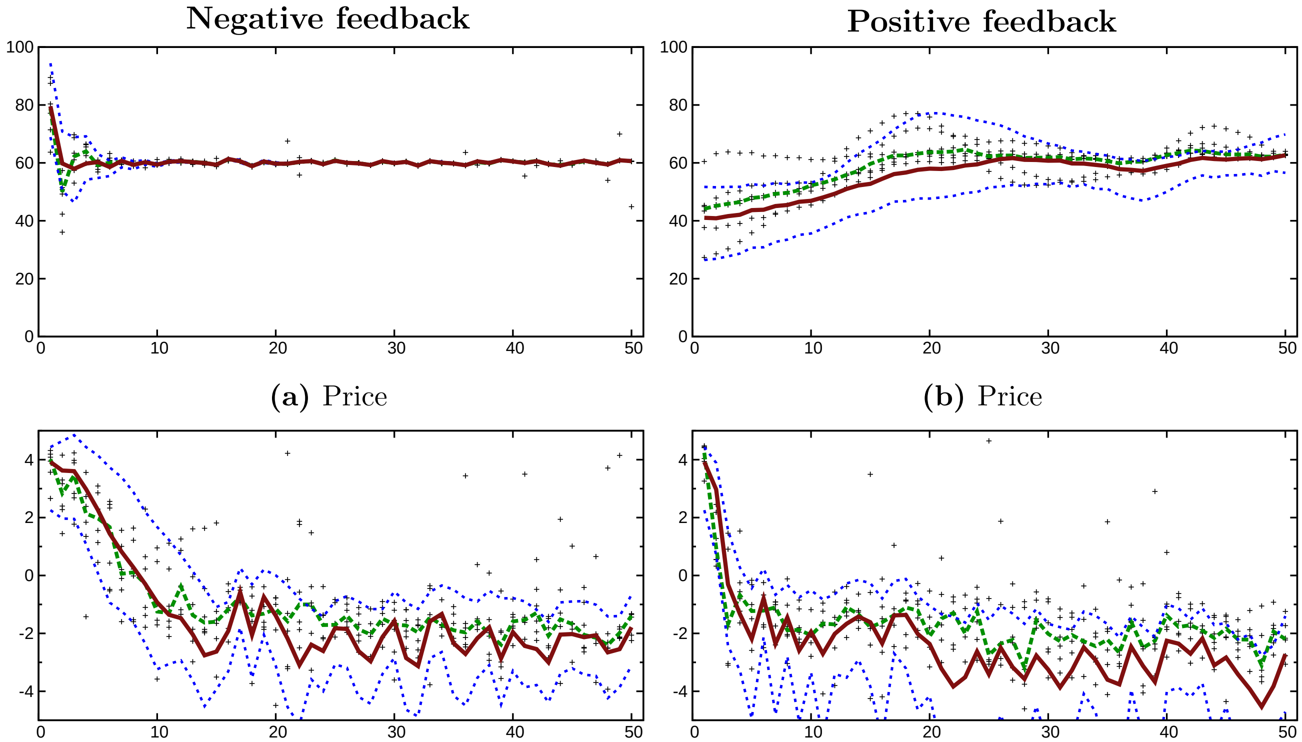

(a) Price **(b)** Price

(c) SD of individual predictions **(d)** SD of individual predictions

Figure 2: HHST09: 50-period ahead MC simulation (1000 markets) for **GA-P1** model compared with the experimental data. *Upper panels:* price. *Lower panels:* degree of coordination ($\log_2$ scale). Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median and blue dotted lines are the 95% confidence interval for the GA model.

experiment. To compute prices, Eq. (3) or Eq. (2) is applied for positive or negative feedback simulations, respectively. We resample the model 1000 times, including new initial predictions and realisations of the learning algorithm, to obtain a satisfactory MC distribution. The median of 1000 GA simulations with 95% confidence intervals (CI) for Parametrisation 1 (i.e., allowing contrarian rules) are shown in Fig. 2. For an example, we show two typical simulations of the model in the lower panels of Fig. 1. It is striking that these simulations are almost identical to the experimental data shown in the two upper panels.[22]

Fig. 2 shows the MC simulations of the realised prices (upper panels) and the degree of coordination computed as the standard deviation of six individual forecasts (lower panels that use the $\log_2$ scale). The model replicates the experimental outcomes well. Under negative feedback (left panels), prices are quickly pushed close to the fundamental, but individual heterogeneity of GA agents declines slowly and is visible until period 15, consistent with the experimental data. Under positive feedback, GA agents coordinate their forecasts in less than five periods, but the distribution of realised prices does not collapse into the fundamental even after 50 periods,

[22]Simulations presented in Fig. 1c and 1d were among the first that we run.

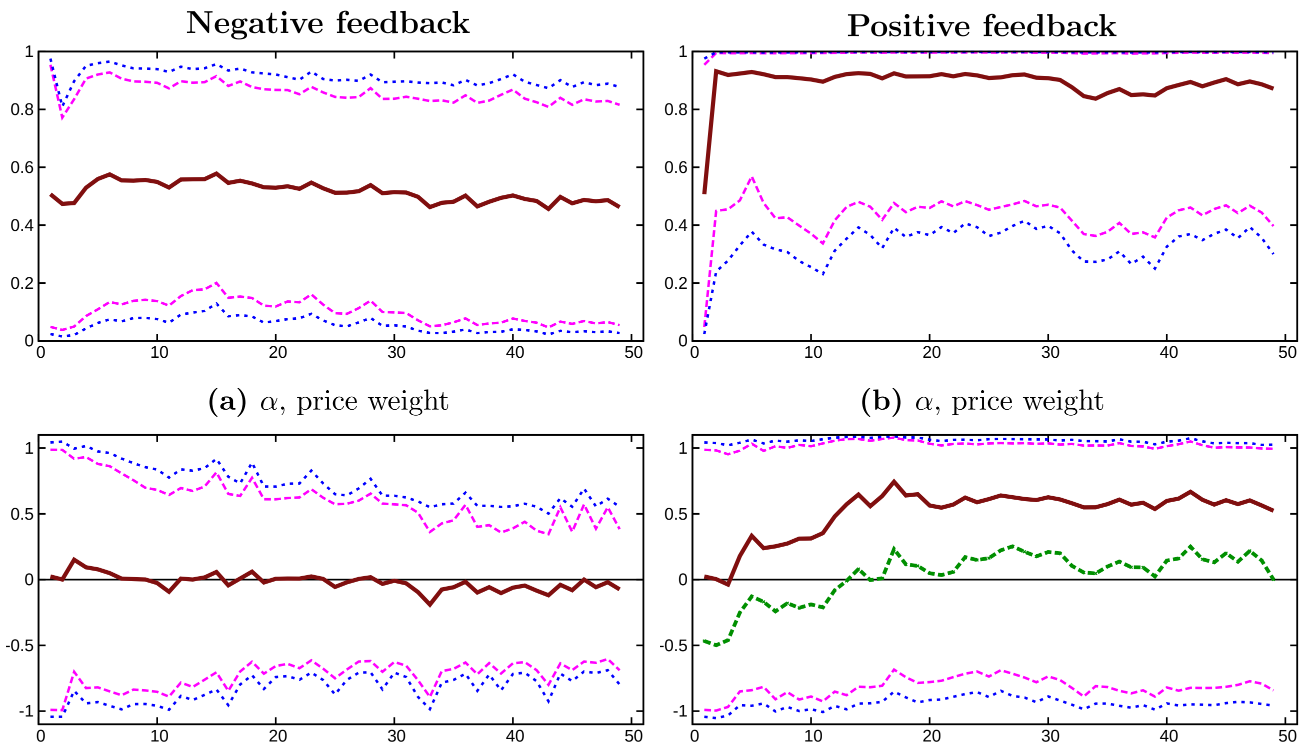

(a) α, price weight **(b)** α, price weight

(c) β, trend extrapolation coefficient **(d)** β, trend extrapolation coefficient

Figure 3: HHST09: Emerging heuristics in 50-period ahead MC simulation (1000 markets) for **GA-P1** model. The price weight α (*upper panels*) and the trend extrapolation coefficient β (*lower panels*) of the chosen heuristic are shown. Red thick line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model. The green star-line in panel (d) represents the 28% percentile of chosen β.

when the 95% CI of prices is as wide as $[55, 70]$. The median price resembles the experimental oscillations, including the typical amplitude and turning points. Overall, the 95% CI for our GA model captures 65% (resp. 81%) of the experimental prices and 81% (resp. 72%) of the degree of coordination for the negative (resp. positive) feedback treatment. In other words, we are able to replicate roughly 75% of the long-run (50-period ahead) behaviour of the experimental groups, both at the *aggregate* and *individual* levels.

Which heuristics were learned by our GA agents? Fig. 3 reports the median (with 95% and 90% CI) for the MC simulations of the price weight α and the trend extrapolation coefficient β. Large heterogeneity of individual rules persists, but there are clear differences between the two treatments. Under the positive feedback treatment, the median GA agent quickly converges towards an approximate rule

$$p^e_{i,t+1} \approx 0.9p_t + 0.1p^e_{i,t} + 0.6(p_t - p_{t-1}). \tag{7}$$

This median rule is close to a pure trend-following rule (i.e., with anchor p_t), but has a coefficient $\beta \approx 0.6$, smaller than $\beta = 1$ that Anufriev, Hommes, and Philipse (2013), AHP henceforth, used in the 2-type HSM. Furthermore, 72% of the GA agents never had negative β in the last 30 periods (see the green star-line in Fig. 3d for 28% percentile). For the distribution of β in period 50, see Fig. 10a. On the other hand, under negative feedback, the median GA agent learns a rule close to

$$p^e_{i,t+1} \approx 0.5p_t + 0.5p^e_{i,t} \tag{8}$$

with median trend coefficient β close to 0. Thus the median rule under negative feedback is adaptive expectations with price weight of 0.5; AHP used adaptive expectations with coefficient 0.75 on price in their 2-type HSM. Our learning dynamics therefore confirm the results by **HHST09** and yield empirical support for the 2-type HSM by AHP, albeit with slightly different parametrisation.

In the second MC study, we investigate how well our GA model can replicate long-run dynamics of a *specific* experimental group, focusing on both realised prices and individual forecasts. We fix experimental group X and initialise the 50-period ahead simulations of the GA model with the actual predictions submitted in the first period in this group. The rest of the simulation is performed in the same way as in the first MC study, i.e., we do not use any other information from group X. To quantify how close the simulations are to the actual predictions from group X, we define for every subject i and time t,

$$p^{e,\mathrm{GA}}_{i,t} = \sum_{h=1}^{H=20} \Pi_{i,h,t} p^e_{i,h,t}, \tag{9}$$

which is the average of the price forecasts given by the twenty different heuristics weighted by their fitnesses, $\Pi_{i,h,t}$, as defined in (6). The quantity $p^{e,\mathrm{GA}}_{i,t}$ in (9) is simply the *model prediction* of the subject i's price forecast at time t. Using this prediction and the price trajectory generated by the GA model, p^{GA}_t, we compute the mean squared error (MSE) in predicting the experimental data (both prices and individual price forecasts) for the last 47 periods, i.e., excluding the initialisation phase, as follows

$$\begin{aligned} \mathrm{MSE}^{\mathrm{prices}}_X &= \frac{1}{47}\sum_{t=4}^{50}\left(p^{\mathrm{Gr}X}_t - p^{\mathrm{GA}}_t\right)^2, \\ \mathrm{MSE}^{\mathrm{predictions}}_X &= \frac{1}{6\times 47}\sum_{i=1}^{6}\sum_{t=4}^{50}\left(p^{e,\mathrm{Gr}X}_{i,t} - p^{e,\mathrm{GA}}_{i,t}\right)^2, \end{aligned} \tag{10}$$

where $p^{\mathrm{Gr}X}_t$ and $p^{e,\mathrm{Gr}X}_{i,t}$ denote period t price and forecast of subject i in the experiment.

Table 2 reports these MSE averaged over 1024 sample GA model paths per experimental group and over the six groups for each treatment. We also include the results for a number of benchmark models, including several homogeneous expectation rules, RE, as well as the HSM from AHP.[23] The MSE for the best model is shown in bold and for the second best in italic. Two

[23] For the definition of the benchmark rules, please refer to Appendix A.

Model	Negative feedback		Positive feedback	
	Prices	Predictions	Prices	Predictions
Trend extrapolation	3421	1696	62.84	72.45
Adaptive	4.164	16.97	95.62	108.6
Contrarian	3.446	*16.18*	108.5	122.8
Naive	112.3	136.2	69.11	79.38
RE	**2.571**	**15.21**	46.835	54.811
HSM from AHP	19.64	34.02	55.15	63.98
GA-P1	*2.884*	20.03	*44.22*	*51.98*
GA-P2 (no contrarian)	9.392	29.51	**25.3**	**31.1**

Table 2: **HHST09**: 50-period ahead simulation. MSE of various models for experimental prices and subjects' predictions, averaged over six experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

simple models of adaptive and contrarian expectations as well as RE perform well under negative feedback, because they correctly predict convergence to the fundamental price. Our GA model performs only slightly worse. Under positive feedback, the contrarian and adaptive expectations perform badly, because they still predict convergence, in contrast to the experimental data. The HSM, trend extrapolation and naive expectations perform relatively well, but surprisingly they are not better than RE. The reason is that the price oscillations predicted by these three models at the long time horizon fall out of phase with the experimental oscillations. The best fit is achieved by our GA model, especially by **GA-P2** parametrisation, without contrarian rules. We conclude that most models are able to capture the long-run dynamics of possibly one feedback treatment, but not of both treatments at the same time. Only our GA model successfully predicts *long-run* behaviour in *both* treatments.

3.3 One-period ahead predictions

Another indicator of the model's fit is the precision of its one-period ahead predictions: how well the model predicts experimental outcomes in period $t+1$, conditional on the data available to the subjects of the experiment until period t.[24] For deterministic models such as HSM and the homogeneous expectations models, computing one period-ahead MSE is straightforward. For our GA model with its evolutionary operators, however, evaluating MSE is more complicated. Our model is both stochastic and highly non-linear: it evolves according to an analytically intractable

[24]Anufriev and Hommes (2012) and Anufriev, Hommes, and Philipse (2013) mostly focus on this measure to evaluate the fit of HSM. As Table 2 shows, the long-run simulations may generate oscillations that are qualitatively similar to the experimental data, but – being out of phase – have high MSE.

Model	Negative feedback		Positive feedback	
	Prices	**Predictions**	**Prices**	**Predictions**
Trend extrapolation	21.101	35.648	0.926	*4.196*
Adaptive	*2.3*	*14.912*	2.999	6.482
Contrarian	**2.249**	**14.856**	3.864	7.436
Naive	3.09	15.782	1.822	5.184
RE	2.571	15.21	46.835	54.811
HSM from AHP	2.999	17.106	0.889	**4.156**
GA-P1	4.95	25.017	*0.806*	4.235
GA-P2 (no contrarian)	4.496	25.012	**0.802**	4.198

Table 3: HHST09: one-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over six experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

period-to-period distribution. To address this issue, we compute the *expected* MSE using a simple Sequential Monte Carlo (SMC) approach designed as follows.

For each experimental group X, we run 1024 independent GA model simulations. In every simulation, we associate one GA agent with one subject, and in each period $t \geqslant 2$ every GA agent i (1) retains own heuristics from the previous period and (2) is given the *experimental* prices and the price forecasts of subject i until the previous period $t-1$. GA agents now use the experimental (not artificial) data to update their heuristics and forecast price in a usual way, which gives us the GA's price forecasts (9) and realised prices (1) for period t. We evaluate the fit of the model to the experimental group by computing the average MSE (10) over all 1024 GA simulations.

The results are similar to the 50-period ahead simulations, see Table 3. Under negative feedback many rules (RE, HSM, adaptive, contrarian, naive) capture the convergence of prices and forecasts to the fundamental price, slightly outperforming our GA model. Under positive feedback, these models (except for HSM) loose their predictive power and under-estimate the experimental oscillatory behaviour of individual forecasts. The GA model has the best fit for the positive feedback treatment and outperforms RE by a factor of 10.

4 Evidence from other experiments

Our GA model fits the **HHST09** experiment well. We will now move from the simple linear feedback to more complicated experimental settings. To be specific, we look at three other experiments that offer a hierarchy of challenges for the GA model of individual learning:

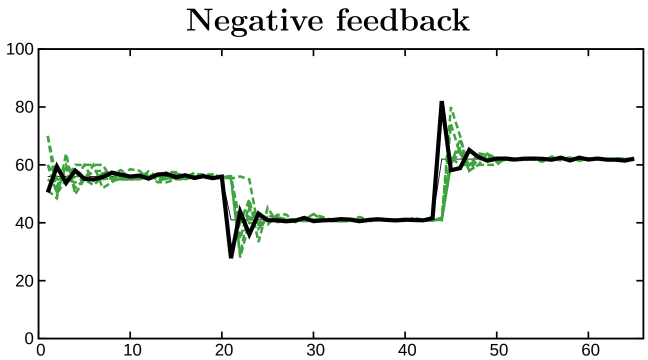

(a) Experimental group 8

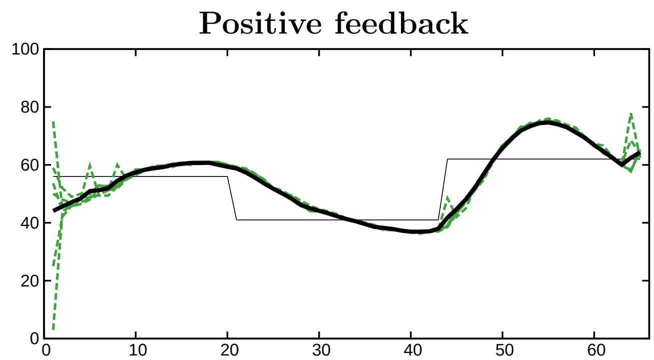

(b) Experimental group 8

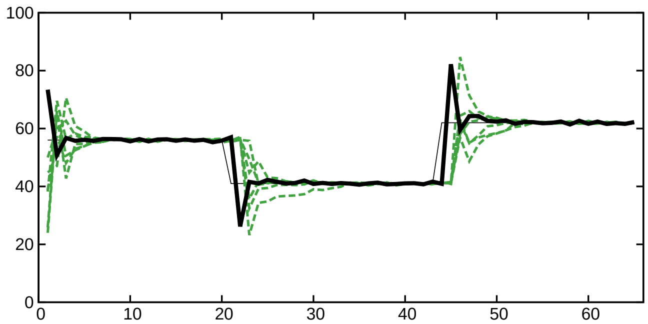

(c) Simulation of GA model

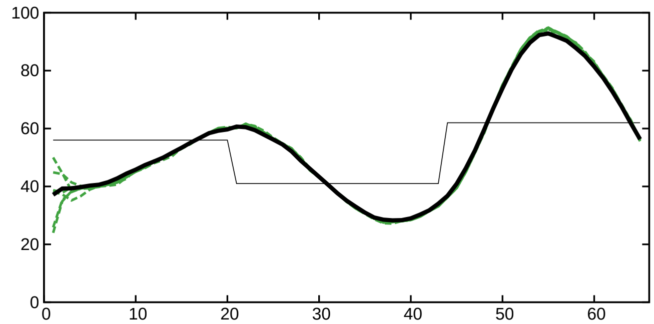

(d) Simulation of GA model

Figure 4: **BHST12**: experimental groups (*upper panels*) and sample 65-period ahead simulations of **GA-P1** model and random initial predictions (*lower panels*). Black thick line shows the price, green dashed lines show 6 individual predictions.

1. **BHST12**: linear feedback with large and unanticipated shocks to the fundamental price;
2. **V01**; **HSTV07**: non-linear (cobweb) negative feedback economy, investigated with a GA model by Hommes and Lux (2013);
3. **HSTV05**: non-linear positive feedback economy with the two-period ahead predictions;

4.1 Large shocks to the fundamental price

BHST12 report an LtF experiment with the same structure as **HHST09**: two treatments with positive and negative feedback, based on linear price equations (2) and (3) with the same dampening factor $\frac{20}{21}$. In this experiment, however, there are two large, permanent and unanticipated shocks to the fundamental price. First it changes from $p^f = 56$ to $p^f = 41$ in period $t = 21$ and then it changes again in period $t = 44$ and remains $p^f = 62$ until the last period $t = 65$.

The results of **BHST12** are similar to **HHST09** and typical time paths are shown in Fig. 4. Under negative feedback (Fig. 4a), a shock to the fundamental breaks the subjects' coordination and is followed by a quick convergence to the new fundamental price. Under positive feedback (Fig. 4b), shocks leave the coordination intact, and the predictions and prices move smoothly towards the new fundamental, eventually over- or undershooting it. The long-run dynamics of the GA model illustrated in Figs. 4c and 4d are very close to the experiment both under negative (*left panels*) and positive (*right panels*) feedback.

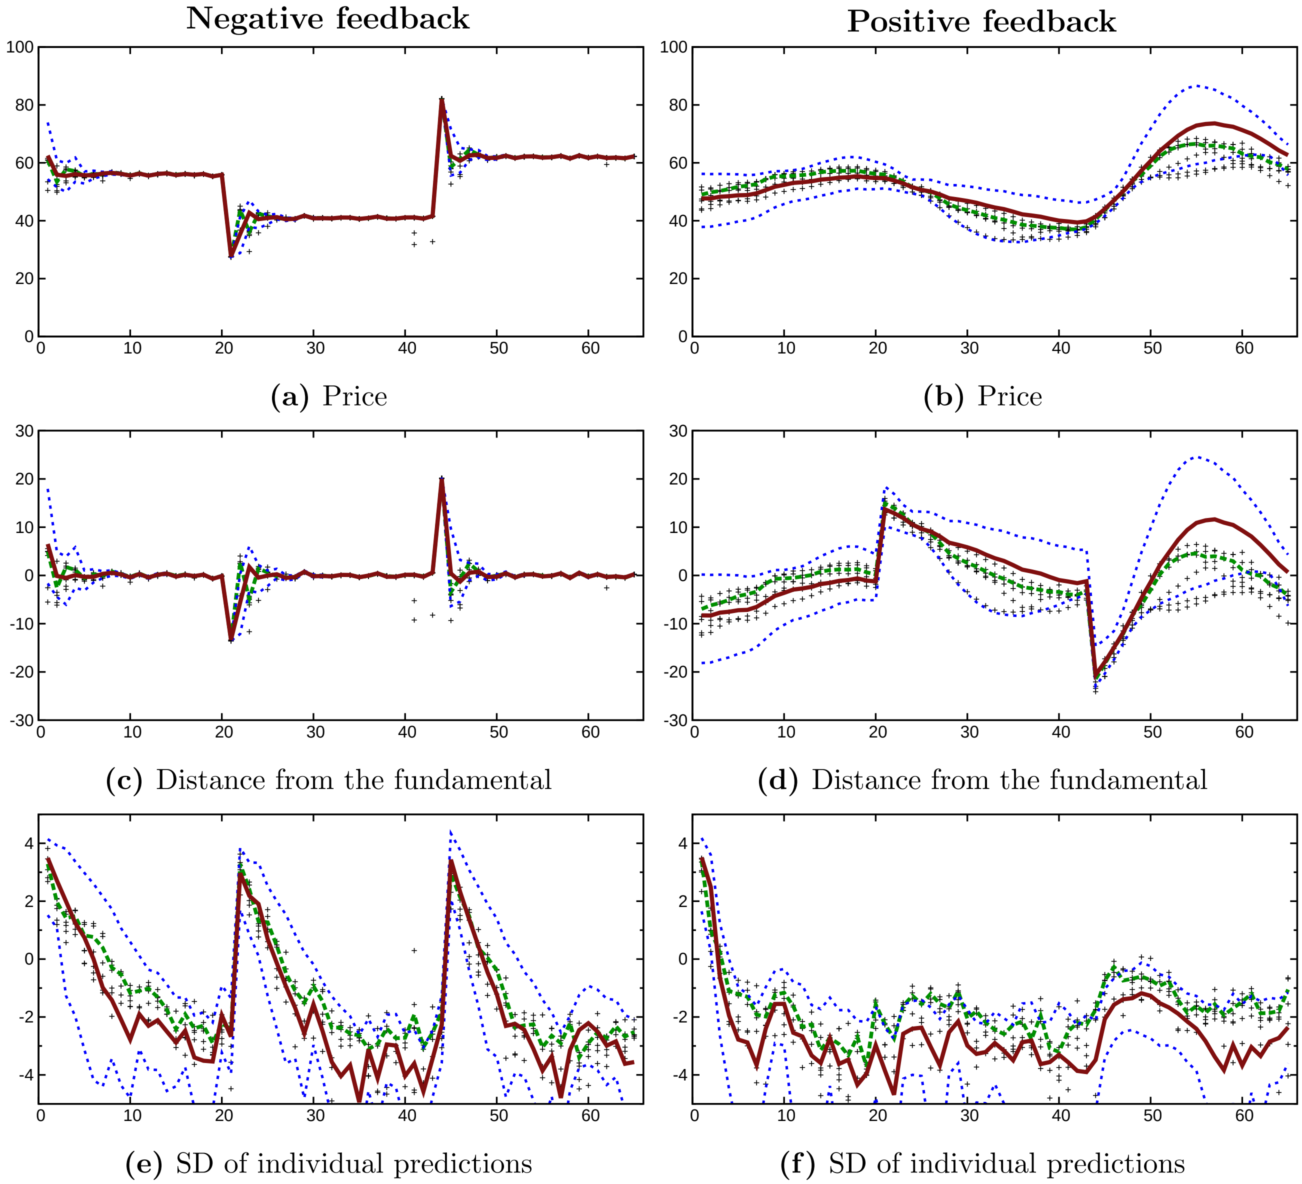

Figure 5: BHST12: 65-period ahead MC simulation (1000 markets) for **GA-P1** model compared with the experimental data. *Upper panels*: price. *Middle panels*: distance from the fundamental price. *Lower panels*: degree of coordination ($\log_2$ scale). Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median and blue dotted lines are the 95% confidence interval for the GA model.

This is further visible on Fig. 5, which illustrates 65-period ahead MC simulations of prices, shown both in levels (upper panels) and in deviations from the fundamental price (middle panels), and the degree of coordination (lower panels). The simulations closely follow the median experimental price paths for both treatments. They also replicate the difference in between-groups

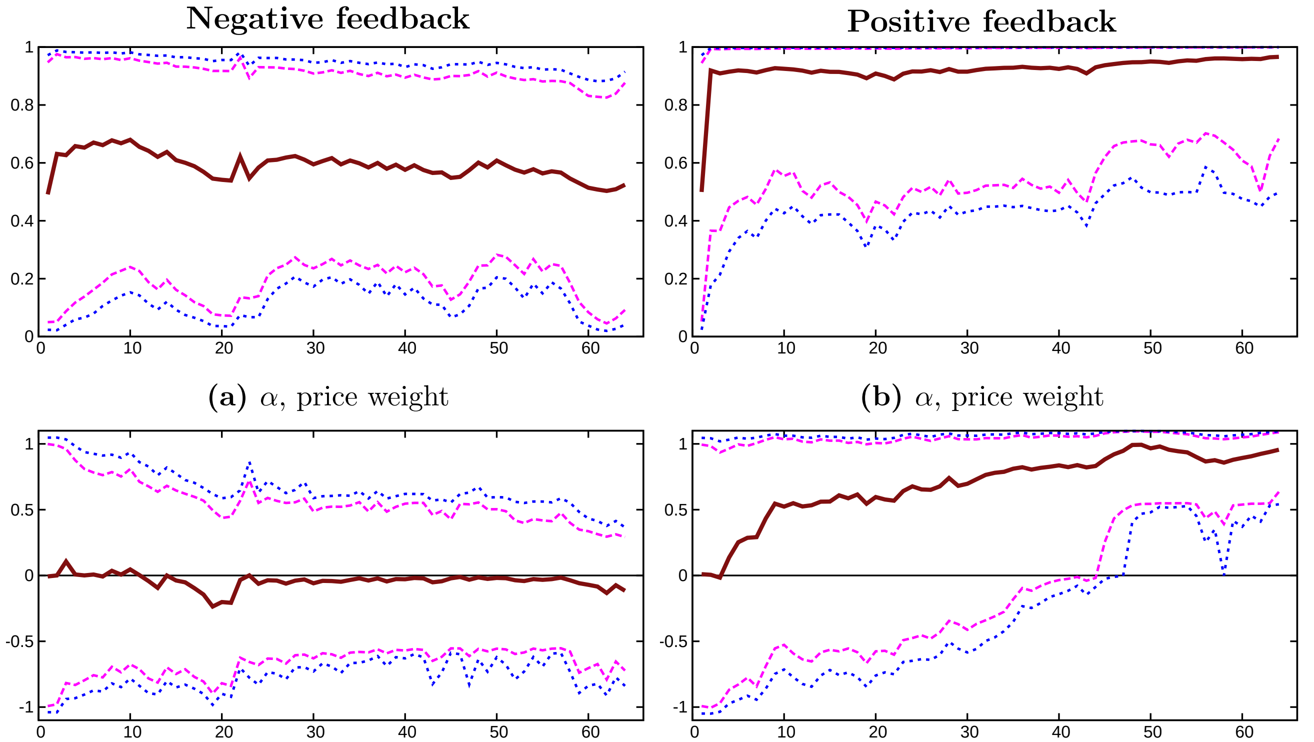

(a) α, price weight

(b) α, price weight

(c) β, trend extrapolation coefficient

(d) β, trend extrapolation coefficient

Figure 6: BHST12: Emerging heuristics in 65-period ahead MC simulation (1000 markets) for **GA-P1** model. The price weight α (*upper panels*) and the trend extrapolation coefficient β (*lower panels*) of the chosen heuristic are shown. Red thick line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model.

variability in the dynamics, which was observed only under positive feedback. The GA long-run simulations are also surprisingly good in evaluating the impact of the shocks on individual coordination under both treatments. The 95% CI of **GA-P1** model contains 66% (resp. 84%) of the experimental prices and 84% (resp. 67%) of the standard deviation of individual forecasts under negative (resp. positive) feedback. Overall, we can replicate around 75% of the experimental data from **BHST12** experiment with 65-period ahead simulations.

Fig. 6 illustrates the time evolution of the price weight α and trend extrapolation coefficient β, which were chosen by the GA agents in the 65-period ahead simulations. The median behaviour is similar to **HHST09** experiment discussed in the previous section. In fact, under negative feedback, the median GA agent learns the same adaptive expectations rule as before, $p^e_{i,t+1} \approx 0.5p_t + 0.5p^e_{i,t}$. Under positive feedback, the median GA agent converges to heuristic

$$p^e_{i,t+1} \approx 0.95p_t + 0.05p^e_{i,t} + 0.9(p_t - p_{t-1}), \tag{11}$$

which is a trend following rule with the trend extrapolation coefficient $\beta \approx 0.9$. This trend coefficient is significantly larger than the coefficient 0.6 in rule (7) used by the median GA agent under the positive feedback from **HHST09** experiment without large shocks. The 95%

Model	Negative feedback		Positive feedback	
	Prices	Predictions	Prices	Predictions
Trend extrapolation	2736	1289	101.3	113.3
Adaptive	**3.629**	**10.75**	55	62.14
Contrarian	6.984	*14.45*	58.46	65.95
Naive	94.44	110.9	46.62	52.9
RE	13.871	20.923	55.133	60.859
HSM from AHP	73.57	87.86	90.8	101.8
GA-P1	8.01	21.97	*43.49*	*49.44*
GA-P2 (no contrarian)	*6.333*	17.39	**43.49**	**49.64**

Table 4: BHST12: 65-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over eight experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

CI for the trend extrapolation coefficient β becomes significantly positive towards the end of the experiment (see also Fig. 10b for the histogram of β's chosen in period 65). Hence, *due to the large, unanticipated shocks in the positive feedback treatment, GA agents become strong trend followers.*

Table 4 reports the MSE for the 65-period ahead simulations initialised with the experimental initial predictions (1024 simulated markets per group for the GA models). We observe that the adaptive expectations have a good fit to the negative feedback treatment, while naive expectations perform well under positive feedback. Interestingly, RE are poor for both treatments: they cannot explain oscillations of the positive feedback and the short spells of volatility that follow the fundamental shocks under the negative feedback treatment. The HSM also performs below average. Our GA model performs, at a balance, very well: it is the second best for the negative feedback and the best for the positive feedback.

We also use the Sequential Monte Carlo (SMC) approach to compute the GA model's one-period ahead predicting power, reported in Table 5. The results are consistent with the 65-period ahead simulations. For both treatments, the GA model (especially without contrarian rules) ranks among the best of all reported models.

4.2 Cobweb economy

HSTV07 and **V01** conducted an LtF experiment in a setting of the cobweb economy. **HSTV07** investigate 18 markets with six subjects each, divided into three treatments of 6 groups: with stable, unstable (on the verge of stability) and strongly unstable parametrisation under the assumption of homogeneous naive expectations. **V01** report the latter treatment with 12 subjects.

Model	Negative feedback		Positive feedback	
	Prices	Predictions	Prices	Predictions
Trend extrapolation	114.061	121.329	1.183	2.165
Adaptive	**3.689**	**10.332**	3.776	4.618
Contrarian	5.92	*12.534*	4.737	5.559
Naive	9.979	16.81	2.411	3.286
RE	13.871	20.923	55.133	60.859
HSM from AHP	38.309	45.679	0.9996	**2.024**
GA-P1	10.247	21.464	*0.342*	2.059
GA-P2 (no contrarian)	**4.208**	15.267	**0.341**	*2.036*

Table 5: BHST12: one-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over eight experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

The experiment resulted in average prices very close to the RE fundamental price. However, the prices were excessively volatile, and – in contrast to the positive feedback experiments – also non-persistent (with weak autocorrelation structure). Hommes and Lux (2013), HL henceforth, study this experiment with a GA model in which agents learn parameters of a simple AR1 forecasting rule, $p_t^e = \alpha_{i,t} + \beta_{i,t} p_{t-1}$. It is now interesting to compare that specification, which we denote **GA-AR1**, and our GA model.

We test our model by conducting a MC exercise in the vein of HL. For each treatment, we compute *six* (as the number of groups per treatment) 50-period ahead simulations with different random seeds for sampling the initial predictions and the between period learning. Next we compute the mean and standard deviation of the realised prices and the individual price forecasts. We repeat this procedure 1000 times to obtain a *distribution* (including 95% CI) of the realised means and variances of prices and price predictions. We report the results in Table 6 for the two parametrisations of our GA model comparing them with the results of the GA model from HL.

Our 50-period ahead simulations explain well the experimental data and perform significantly better than RE. The 95% CI of **GA-P1** and **GA-P2** models replicate 12 and 11 out of 16 experimental statistics, respectively, see Table 6. Among 11 cases successful for **GA-P2** model, 9 statistics reported by HL are outside 95% CI of our model. It means that we can replicate around three quarters of experimental descriptive statistics, most of which with a significantly higher precision than the GA model in HL.[25]

[25]The GA simulations are also closer to the experimental data in terms of the autocorrelation of the prices. RE predicts zero autocorrelation, whereas benchmark models predict high autocorrelation up to the third lag. The experimental data exhibited weak autocorrelation, which is replicated by all three GA model specifications with comparable performance. See Table 12 in Appendix C for the results.

	Prices		Predictions	
	Mean(p)	**Var(p)**	**Mean($\mathbf{p^e}$)**	**Var($\mathbf{p^e}$)**
		Stable		
Experiments	$5.64^{*\dagger}$	$0.36^{*\dagger}$	$5.56^{*\dagger}$	0.087^{*}
GA-AR1	$5.565^{¶}$	$0.326^{¶}$	$5.576^{¶}$	0.1
GA-P1	5.628	0.372	5.571	0.082
95% *CI*	[5.613, 5.643]	[0.359, 0.389]	[5.553, 5.59]	[0.065, 0.101]
GA-P2	5.649	0.353	5.548	0.0565
95% *CI*	[5.631, 5.667]	[0.341, 0.371]	[5.527, 5.57]	[0.043, 0.077]
		Unstable		
Experiments	$5.85^{\dagger}$	$0.63^{*\dagger}$	$5.67^{*\dagger}$	$0.101^{*\dagger}$
GA-AR1	5.817	0.647	$5.645^{¶}$	$0.16^{¶}$
GA-P1	5.792	0.598	5.705	0.103
95% *CI*	[5.744, 5.841]	[0.525, 0.746]	[5.667, 5.739]	[0.067, 0.171]
GA-P2	5.825	0.557	5.694	0.079
95% *CI*	[5.786, 5.863]	[0.487, 0.658]	[5.67, 5.719]	[0.052, 0.122]
		Strongly unstable		
Experiments	$5.93^{\dagger}$	2.62^{*}	5.73	0.429^{*}
GA-AR1	$6.2^{¶}$	2.161	5.434	0.769
GA-P1	5.809	2.172	5.832	0.345
95% *CI*	[5.693, 5.908]	[1.626, 2.875]	[5.735, 5.918]	[0.182, 0.598]
GA-P2	5.962	1.487	5.807	0.206
95% *CI*	[5.876, 6.045]	[1.188, 1.834]	[5.75, 5.858]	[0.113, 0.347]
		Strongly unstable, group size 12		
Experiments	$5.937^{\dagger}$	1.783^{*}	$5.781^{*\dagger}$	$0.204^{*\dagger}$
GA-AR1	$6.183^{¶}$	1.571	$5.515^{¶}$	$0.5^{¶}$
GA-P1	5.812	1.699	5.852	0.194
95% *CI*	[5.731, 5.892]	[1.368, 2.157]	[5.779, 5.918]	[0.122, 0.338]
GA-P2	5.972	1.316	5.804	0.173
95% *CI*	[5.918, 6.026]	[1.118, 1.553]	[5.768, 5.843]	[0.111, 0.253]

Table 6: HSTV07: 50-period ahead MC results for GA simulations for four treatments. Median statistics for average prices, predictions, and their variances are shown for the experiment and three GA models: with AR1 rule used in HL and our two parametrisations (also with 95% confidence intervals). $*$ and † denote experimental statistic which falls into 95% CI of **GA-P1** and **GA-P2**, respectively. ¶ denotes **GA-AR1** statistics which fall outside the 95% CI of **GA-P2** model containing the experimental statistics.

Treatments	Stable		Unstable		Strongly unstable	
Models	**Prices**	**Predictions**	**Prices**	**Predictions**	**Prices**	**Predictions**
Trend extrapolation	13.3	71.1	16.33	89.59	16.55	89.07
Adaptive	0.117	0.339	7.206	3.272	16.45	7.822
Contrarian	0.093	0.308	1.746	0.834	13.95	5.282
Naive	1.076	1.724	14.67	16.18	16.55	18.55
RE	*0.048*	**0.248**	0.364	**0.385**	**2.257**	**1.844**
HSM from AHP	0.178	0.422	7.446	3.431	16.46	7.885
GA-AR1	0.05742	0.3759	0.3552	0.6596	2.838	2.64
GA-P1	0.088	0.356	*0.346*	0.631	3.445	3.261
GA-P2	**0.043**	*0.275*	**0.223**	*0.449*	*2.376*	*2.114*

Table 7: HSTV07: 50-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over all six experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

Treatments	Stable		Unstable		Strongly unstable	
Models	**Prices**	**Predictions**	**Prices**	**Predictions**	**Prices**	**Predictions**
Trend extrapolation	1.176	1.997	2.122	3.719	5.856	14.39
Adaptive	0.108	0.328	0.434	0.549	2.784	2.863
Contrarian	0.102	0.318	0.414	*0.497*	*2.929*	2.729
Naive	0.196	0.448	0.577	0.788	3.095	3.731
RE	**0.048**	**0.248**	*0.364*	**0.385**	**2.257**	**1.844**
HSM from AHP	0.212	0.474	0.52	0.732	3.065	3.691
GA-AR1	0.054	0.36	0.51	0.674	5.36	3.432
GA-P1	0.13	0.393	0.866	0.795	5.547	3.25
GA-P2	*0.07*	*0.31*	**0.25**	0.531	3.079	*2.358*

Table 8: HSTV07: one-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over all six experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

We also check the 50-period ahead dynamics of the model conditional on the initial predictions submitted in a particular group of **HSTV07**, see Table 7. Homogeneous expectation models, as well as HSM for the two unstable treatments are outperformed by RE. The dynamics of this experiment (in contrast to the experiments with linear feedback) resemble a white noise around the fundamental price. As a result, predicting the mean (as RE do) of these dynamics is better than trying to capture them with structural models. However our GA model, in particular **GA-**

P2, keeps up with RE and performs better than **GA-AR1** used in HL.

The next MC exercise is the one-period ahead forecasting of the model with SMC approach for the 18 groups from **HSTV07**. Table 8 gives the summary of the results. It is apparent that the less stable the treatment is, the worse fit any model has. As for the 50-period ahead forecasts, the clear winners are RE and our GA model, which are able to explain the data well also for the strongly unstable treatment.[26] Our specification again prevails over **GA-AR1** model.

We conclude that in the cobweb experiments with unstable dynamics the simple homogeneous models, but also HSM, miss-identify any structure. As a result, their point forecasts are so poor that it is better to predict the mean price, as RE does missing, however, excess volatility. Our GA model (without contrarian behaviour) comes close to RE in terms of fitting the mean but also allows to explain the excess volatility observed in the experiments. Finally, it is clear that the use of experimental evidence for micro-behaviour has an advantageous effect: our GA model, with an empirically motivated anchor and adjustment rule (5), has a better fit to the data than the AR1 specification used by Hommes and Lux (2013).

4.3 Two-period ahead asset pricing

HSTV05 report an experiment based on a *2-period ahead* non-linear positive feedback market, an asset-pricing model. In this market the *current* price depends on the average of the subjects' expectations about the price in the *next* period, i.e., $p_t = F(\bar{p}^e_{t+1})$. There were two treatments with different fundamental price: in seven markets $p^f = 60$ and in three markets $p^f = 40$. Three different aggregate outcomes were observed: (i) monotonic convergence to the fundamental price (2 groups), (ii) dampened oscillations (3 groups) and (iii) volatile price oscillations (5 groups).[27]

Participants had to predict p_{t+1} without knowing p_t, and therefore their decisions were based on a different information set than in the previous one-period ahead experiments. The 2-period ahead version of our GA model is based on the following prediction heuristic:

$$p^e_{i,h,t+1} = \alpha_{i,h} p_{t-1} + (1 - \alpha_{i,h}) p^e_{i,t} + \beta_{i,h}(p_{t-1} - p_{t-2}) \,. \tag{12}$$

Once p_t is realised, the agents can evaluate their rules based on the hypothetical performance of predicting p_t *two periods ago*, i.e., their fitness is a normalised MSE $= (p^e_{i,h,t} - p_t)^2$, as before. This specification is the most straightforward adaptation of the baseline one-period ahead forecasting heuristic (5). Recall that in the two baseline parametrisations, **GA-P1** and **GA-P2**, we imposed the restrictions on the trend coefficients, $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$, respectively. **HSTV05**, however, found that many subjects used stronger trend extrapolation. Therefore, for the sake of

[26] Note that the scale of prices in this experiment is $[0, 10]$ in contrast with the two previous settings, where the prices belonged to $[0, 100]$ interval. The highest possible MSE in the linear experiments is 100 times higher than in the cobweb experiment.

[27] Bao, Hommes, and Makarewicz (2016) run a similar asset market experiment to compare learning to forecast versus learning to optimise designs. They report similar three types of price behaviour.

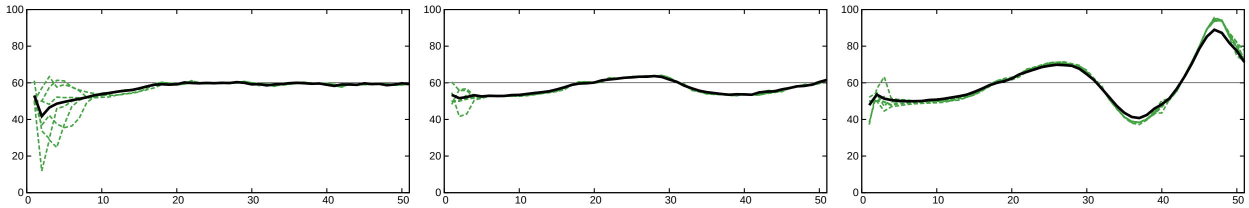

(a) Convergence **(b)** Unclear **(c)** Oscillations

Figure 7: HSTV05: three 50-period ahead simulations of **GA-P3** model for different seeds giving different initial predictions and learning. Green dotted lines are individual predictions and black thick line is the price.

completeness we will also report the results of our model with $\beta \in [-1.3, 1.3]$ (parametrisation **GA-P3**) and $\beta \in [0, 1.3]$ (parametrisation **GA-P4**).

Among seven groups with $p^f = 60$, **HSTV05** observe both the groups with the price converging to the fundamental price, and the non-converging groups with oscillations of different amplitude and frequency. Fig. 7 displays three typical simulated markets of **GA-P3** model for **HSTV05** economy with $p^f = 60$.[28] GA agents can either converge to the fundamental price (Fig. 7a) or coordinate on oscillations (Fig. 7c). Furthermore, sometimes an intermediate outcome occurs, which can be interpreted as transitory dynamics between the stable and unstable outcome. Fig. 7b shows a sample simulation, in which after the first 20 periods, the price seemingly stabilises at the fundamental value around 60. One may expect to see the same fundamental dynamics as in Fig. 7a. However, in the remaining periods, the price resume to oscillate mildly.

To further stress the volatile behaviour of this market structure, we report one long-run simulation for **GA-P3** model, see Fig. 8. The top panels display price dynamics with persistent oscillations of different amplitude, where large oscillations can reappear even after the market seemingly settled on the fundamental price. It suggests that the invariant distribution of our stochastic model may have several modes.[29] GA model generates economically relevant dynamics with *clustered volatility*, i.e., when phases of relatively stable price behaviour interchange with highly volatile price fluctuations. The bottom panel of Fig. 8 shows the average β chosen by the six GA agents. Despite continuing instability, a clear pattern is that the average β remains close to zero in the stable phase of the simulation, but stays close to the upper limit of 1.3 in volatile times. We interpret this pattern in the following way. If the price is stable and close to the fundamental

[28] The simulations are based on different initial predictions and learning realisation, though the supply shocks ε_t are the same.

[29] Due to the presence of mutation in the learning phase and the noise in the pricing equation, our GA model is an ergodic Markov process. Therefore, the invariant distribution exists, though it cannot be computed analytically due to the complexity of the model. As this paper is motivated by the experimental data, we simulated and compared in Fig. 10 the distributions of the trend extrapolation coefficient after the first 50 periods for all the positive feedback treatments discussed in the paper, leaving more systematic investigation of the asymptotic properties of GA dynamics to future research.

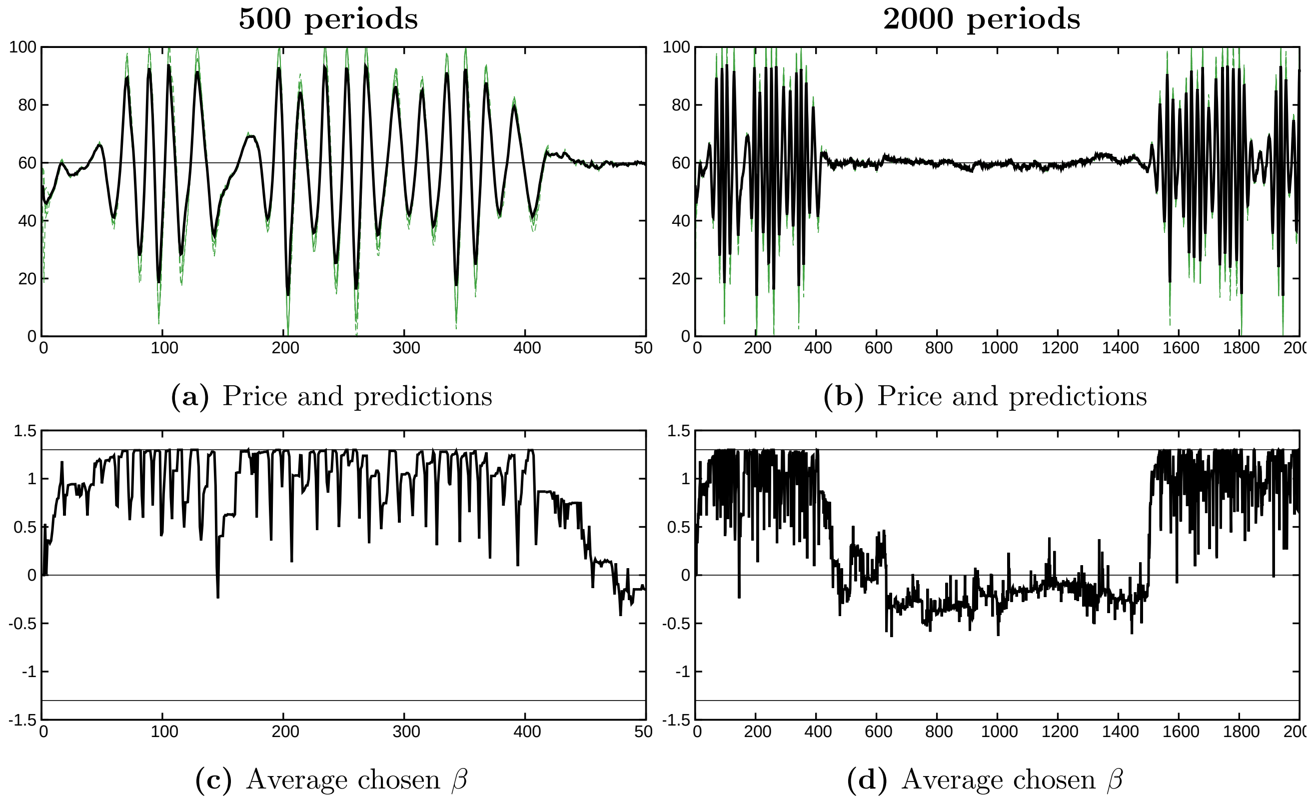

(a) Price and predictions **(b)** Price and predictions

(c) Average chosen β **(d)** Average chosen β

Figure 8: HSTV05: sample 2000-period ahead simulation (*right panels*) and its first 500 periods (*left panels*) of **GA-P3** model with $p^f = 60$. *Top panels:* individual predictions (green dashed lines) and price (black line). *Bottom panels:* average trend extrapolation coefficient β chosen by six GA agents.

value, the fittest heuristics give predictions that are close to the fundamental value. Due to averaging of the predictions of six GA agents and the artificial robotic fundamentalist, deviations from the fundamental price can be mitigated. This discourages GA agents from extrapolating an insignificant trend, reinforcing price stability. Nevertheless, the trend in prices may become sufficiently large, so that the predictions of GA agents become sufficiently coordinated to counterweight the stabilizing effect of the artificial fundamental robot traders. This leads to a drift of the extrapolating coefficients in the fittest heuristics towards the upper bound, and the price oscillations become self-reinforcing. Under the non-linear (due to the robotic trader) two-period ahead price feedback mechanism, the specific shape (i.e., phase, amplitude) of oscillations is diversified. As a result, there is still space for GA agents to experiment with the specific strength of trend following. In the two-period ahead feedback system, our GA model thus entails not only two 'attractors' (ie., two types of long run behaviour), fundamental price and large volatility (oscillations), but also generates endogenous switching between them. This corresponds well to the diversified dynamics observed in the experiment.

To support this story, we take a closer look at the trend extrapolation coefficient β chosen by the GA agents during the first 50 periods. Fig. 9 shows the results for MC 50-period ahead

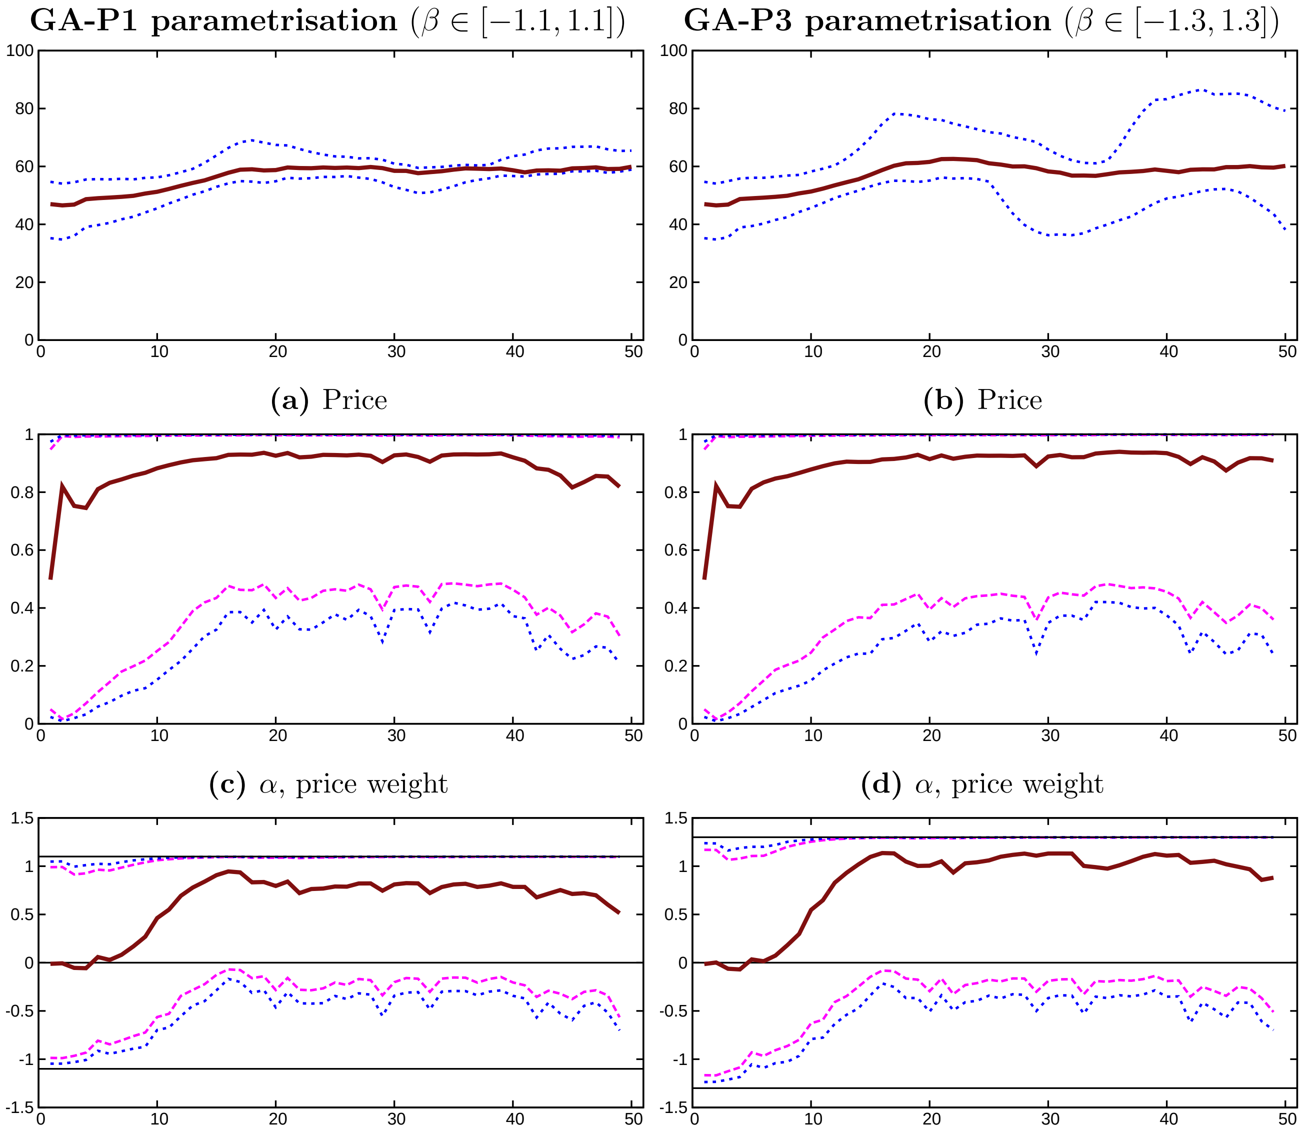

(a) Price **(b)** Price

(c) α, price weight **(d)** α, price weight

(e) β, trend extrapolation coefficient **(f)** β, trend extrapolation coefficient

Figure 9: HSTV05: Emerging heuristics in 50-period ahead MC simulation (1000 markets) for **GA-P1** (*left panel*) and **GA-P3** (*right panel*) parametrisations. The price (*upper panels*), the price weight α (*middle panels*) and the trend extrapolation coefficient β (*lower panels*) of the chosen heuristic are shown. Red thick line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model.

simulations for two GA model parametrisations, **GA-P1** and **GA-P3**. Under the latter setting, the agents are allowed to experiment with higher β. The median price has a very similar oscillatory shape in both cases, but the difference is seen in the 95% CI. Both parametrisations are likely to generate two price bubbles within 50 periods, but **GA-P3** model with higher β's has larger potential oscillations (Fig. 9b), and the second bubble can be even bigger than the first (unlike in the linear positive feedback). In both parametrisations, the median GA agent converges to a strong trend extrapolation rule, close to $p^e_{i,t+1} = p_{t-1} + (p_{t-1} - p_{t-2})$, which is consistent with

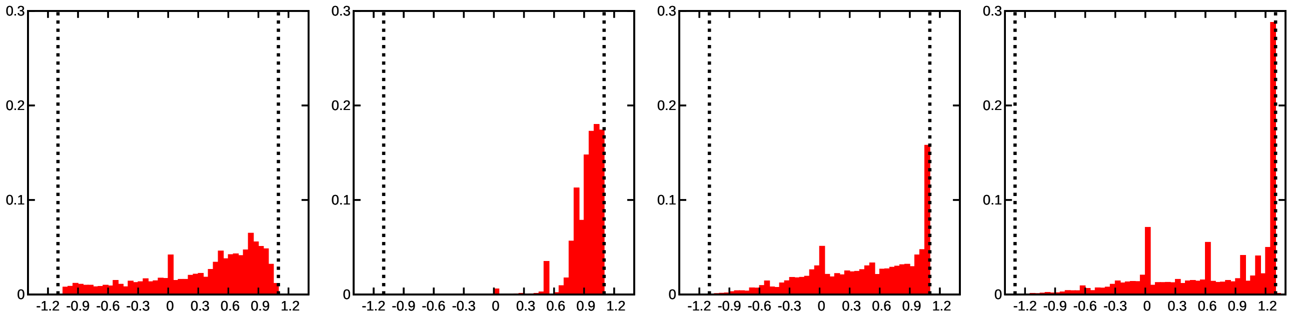

(a) HHST09, GA-P1 **(b) BHST12, GA-P1** **(c) HSTV05, GA-P1** **(d) HSTV05, GA-P3**

Figure 10: Positive feedback treatments: **HHST09**, **BHST12** and **HSTV05** with $p^f = 60$. Distribution of trend extrapolation coefficient β chosen by the agents in period $t = 50$ across the whole MC sample for **GA-P1** and **GA-P3** (last panel) parametrisations.

Models	**Prices**	**Predictions**
Trend extrapolation	178.2	174.9
Adaptive	*96.12*	*145.9*
Contrarian	157	146.8
Naive	**95.29**	**144.6**
RE	96.0328	145.998
GA-P1	103.9	155.8
GA-P2	114.9	169.1
GA-P3	139.4	201.5
GA-P4	226.5	318.5

Table 9: **HSTV05**: 50-period ahead predictions. MSE of various models for experimental prices and subjects' predictions, averaged over all experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

the behaviour of our model in the previous experiments. Nevertheless, the 95% CI of the chosen trend coefficient remain wide and the distribution of this variable in period 50 is close to bi-modal (see Figs. 10c and 10d), with a relatively large mass centred around zero, i.e., weak or no trend extrapolation, and a peak around the maximum possible trend coefficient.

Even though our GA model leaves space for improvement,[30] it is the only model which is comparatively good in predicting the experimental results of **HSTV05** both in the long- and

[30]For instance, the GA model seemingly does not generate the dynamics with relatively frequent (with period of 8-9 periods) oscillations of constant amplitude around the fundamental price, observed in some sessions of **HSTV05**. HSM with four heuristics (adaptive, two different trend extrapolation and anchor and adjustment) did actually capture such dynamics and also had a good one-period ahead fit to these faster price oscillations (Anufriev and Hommes, 2012). In order to improve the GA model's fit to the observed subjects learning in this set-up, one could experiment with higher order rules, but we leave this for future investigations.

Models	**Prices**	**Predictions**
Trend extrapolation	17.4527	55.0898
Adaptive	44.125	*25.3157*
Contrarian	59.3905	30.8646
Naive	31.6864	**20.8416**
RE	96.0328	145.998
HSM (4 heuristics)	*6.798*	—
GA-P1	42.224	74.95
GA-P2	**5.934**	30.341
GA-P3	21.192	53.238
GA-P4	16.29	42.125

Table 10: **HSTV05**: one-period ahead predictions. MSE of various models, including 4-type Heuristic Switching Model (source: Anufriev and Hommes, 2012), for experimental prices and subjects' predictions, averaged over all experimental groups for the corresponding treatment. MSE of the best model is in bold, of the second best is in italic.

the short-run. Table 9 reports the MSE of 50-period ahead simulations initialised with the experimental initial predictions. The long-run predictive power is relatively poor for all models. The best three models are naive, adaptive and RE, though our model (with 1.1 as the upper bound for trend extrapolation) yields similar results. Table 10 shows the MSE of one-period ahead predictions for our GA model and other benchmark models. The GA model is now among the best, especially in terms of predicting the experimental prices. Surprisingly, the models that did well in 50-period ahead predictions are poor now, while trend extrapolation is comparable with our model. Anufriev and Hommes (2012) investigated the **HSTV05** experiment with a *four-heuristics* HSM, which is a richer model than the two-heuristic HSM we used as a benchmark for the previous experiments. Interestingly, only our GA model (specifically with β restricted to $[0, 1.1]$) is able to compete with this richer HSM in terms of predicting the experimental prices.

5 Conclusions

In the model of this paper agents independently use Genetic Algorithms (GA) to optimise a simple forecasting heuristic. The model dynamics was compared with the outcomes of Learning-to-Forecast experiments, where the realised market price depends on individual forecasts. These experiments are used to study how human subjects adapt to the price-predictions feedback in a controlled environment. We showed that GAs capture individual forecasting behaviour in the experiments quite well and also reproduce the aggregate outcomes. GA agents use a linear first-order heuristic (a mixture of adaptive and trend extrapolating expectations) to forecast prices.

They independently optimise the two parameters of their forecasting rule with GAs, learning to fine-tune them to the specific market conditions. Agents thus learn simple forecasting heuristics that make them smart.

Experimental data can be used to test various theories. Our goal was to compare the prediction accuracy of the GA model with other models: rational expectations, a number of homogeneous expectation models, and the Heuristic Switching Model of Anufriev and Hommes (2012). We focused on the out-of-sample one-period and $50-$periods ahead predictions and showed that in comparison with other models, the GA model is able to account for both the *aggregate* outcomes and the *individual* behaviour across four *different* experiments.

The strength of the model lies in its parsimony, flexibility, and generality. Under GA, the coefficients in the active heuristic are *time-varying.* When agents face a negative feedback type of economy, a median GA agent will increasingly rely on adaptive expectations, enforcing convergence of the market to the fundamental equilibrium. In contrast, positive feedback induces the agents to follow the observed price trend and median forecasting behaviour converges to a trend extrapolation rule, which amplifies price oscillations. Also, the more 'complex' the positive feedback is (in terms of shocks to the fundamental solution, or a non-linear law of motion of the price), the stronger trend extrapolation chosen by the median agent is, and the more volatile the price fluctuations will be.

The GA model is general and can be used to investigate settings with more complicated interactions between individual agents. This can include economies with heterogeneous preferences, unequal market power, information networks, decentralised price setting, etc. In any of these cases, heterogeneous price expectations may have important consequences for market efficiency and price dynamics. Our Genetic Algorithms model gives a realistic explanation of how such heterogeneity between the agents emerges from their individual learning and, for each environment, which heuristics make them smart.

References

ADAM, K. (2007): "Experimental Evidence on the Persistence of Output and Inflation," *The Economic Journal*, 117(520), 603–636.

ALLEN, H., AND M. P. TAYLOR (1990): "Charts, Noise and Fundamentals in the London Foreign Exchange Market," *Economic Journal*, 100(400), 49–59.

ANUFRIEV, M., T. ASSENZA, C. HOMMES, AND D. MASSARO (2013): "Interest rate rules and macroeconomic stability under heterogeneous expectations," *Macroeconomic Dynamics*, 17(08), 1574–1604.

ANUFRIEV, M., AND C. HOMMES (2012): "Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments," *American Economic Journal: Microeconomics*, 4(4), 35–64.

ANUFRIEV, M., C. H. HOMMES, AND R. H. PHILIPSE (2013): "Evolutionary selection of expectations in positive and negative feedback markets," *Journal of Evolutionary Economics*, 23(3), 663–688.

ANUFRIEV, M., AND V. PANCHENKO (2009): "Asset prices, traders behavior and market design," *Journal of Economic Dynamics and Control*, 33(5), 1073–1090.

ARIFOVIC, J. (1994): "Genetic algorithm learning and the cobweb model," *Journal of Economic Dynamics and Control*, 18(1), 3–28.

——— (1995): "Genetic algorithms and inflationary economies," *Journal of Monetary Economics*, 36(1), 219–243.

——— (1996): "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy*, 104(3), 510–541.

ARIFOVIC, J., J. BULLARD, AND O. KOSTYSHYNA (2013): "Social learning and monetary policy rules," *The Economic Journal*, 123(567), 38–76.

ASSENZA, T., P. HEEMEIJER, C. HOMMES, AND D. MASSARO (2014): "Managing Self-organization of Expectations through Monetary Policy: a Macro Experiment," CeNDEF Working paper 14-07, University of Amsterdam.

BAO, T., C. HOMMES, AND T. MAKAREWICZ (2016): "Bubble Formation and (In)efficient Markets in Learning-to-Forecast and -Optimize Experiments," *The Economic Journal*, in press.

BAO, T., C. HOMMES, J. SONNEMANS, AND J. TUINSTRA (2012): "Individual expectations, limited rationality and aggregate outcomes," *Journal of Economic Dynamics and Control*, 36(8), 1101–1120.

BOLT, W., M. DEMERTZIS, C. G. DIKS, C. H. HOMMES, AND M. VAN DER LEIJ (2014): "Identifying booms and busts in house prices under heterogeneous expectations," Discussion Paper 14-157/II, Tinbergen Institute.

Boswijk, H. P., C. H. Hommes, and S. Manzan (2007): "Behavioral heterogeneity in stock prices," *Journal of Economic Dynamics and Control*, 31(6), 1938–1970.

Branch, W. A. (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *The Economic Journal*, 114(497), 592–621.

Branch, W. A., and G. W. Evans (2010): "Asset return dynamics and learning," *Review of Financial Studies*, 23(4), 1651–1680.

Branch, W. A., and B. McGough (2010): "Dynamic predictor selection in a new Keynesian model with heterogeneous expectations," *Journal of Economic Dynamics and Control*, 34(8), 1492–1508.

——— (2011): "Business cycle amplification with heterogeneous expectations," *Economic Theory*, 47(2-3), 395–421.

Brock, W., and C. Hommes (1998): "Heterogeneous beliefs and routes to chaos in a simple asset pricing model," *Journal of Economic Dynamics & Control*, 22(8-9), 1235–1274.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65(5), 1059–1095.

Bullard, J. (1994): "Learning Equilibria," *Journal of Economic Theory*, 64(2), 468 – 485.

Camerer, C., and H. T. Ho (1999): "Experience-weighted Attraction Learning in Normal Form Games," *Econometrica*, 67(4), 827–874.

Case, K. E., R. J. Shiller, and A. K. Thompson (2012): "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets," *Brookings Papers on Economic Activity*, 2012(Fall), 265–298.

Cornea, A., C. Hommes, and D. Massaro (2013): "Behavioral heterogeneity in U.S. inflation dynamics," Discussion Paper 13-015/II, Tinbergen Institute.

Dawid, H. (1996): *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*. Springer-Verlag New York, Inc.

Dawid, H., and M. Kopel (1998): "On economic applications of the genetic algorithm: a model of the cobweb type," *Journal of Evolutionary Economics*, 8, 297–315.

Diks, C., and T. Makarewicz (2013): "Initial Predictions in Learning-to-Forecast Experiment," in *Managing Market Complexity*, ed. by A. Teglio, S. Alfarano, E. Camacho-Cuena, and M. Gins-Vilar, vol. 662, pp. 223–235. Springer Berlin Heidelberg.

Doornik, J. (2007): *Object-oriented matrix programming using Ox*. Timberlake Consultants Press, London, 3rd edn.

Duffy, J. (2014): "Macroeconomics: A Survey of Laboratory Research," Discussion paper, University of Pittsburgh, Department of Economics, Available online at http://www.socsci.uci.edu/~duffy/papers/DuffyHEEJune14.pdf.

Erev, I., and A. E. Roth (1998): "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88(4), 848–81.

Evans, G., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

Evans, G., and G. Ramey (2006): "Adaptive expectations, underparameterization and the Lucas critique," *Journal of Monetary Economics*, 53(2), 249–264.

Feltovich, N. (2000): "Reinforcement-based vs. Belief-based Learning Models in Experimental Asymmetric-information Games," *Econometrica*, 68(3), 605–641.

Gigerenzer, G., and P. M. Todd (1999): *Simple Heuristics That Make Us Smart*. Oxford University Press.

Goodman, J. L., and J. B. Ittner (1992): "The accuracy of home owners' estimates of house value," *Journal of Housing Economics*, 2(4), 339 – 357.

Guesnerie, R. (2005): *Assessing Rational Expectations 2: "Eductive" Stability in Economics*. MIT Press.

Haupt, R., and S. Haupt (2004): *Practical Genetic Algorithms*. John Wiley & Sons, Inc., New Jersey, 2nd edn.

Heemeijer, P., C. Hommes, J. Sonnemans, and J. Tuinstra (2009): "Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation," *Journal of Economic Dynamics and Control*, 33(5), 1052–1072.

Hommes, C. (2011): "The heterogeneous expectations hypothesis: Some evidence from the lab," *Journal of Economic Dynamics and Control*, 35(1), 1–24.

——— (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press.

Hommes, C., and T. Lux (2013): "Individual expectations and aggregate behavior in learning to forecast experiments," *Macroeconomic Dynamics*, 17(2), 373–401.

Hommes, C., J. Sonnemans, J. Tuinstra, and H. van de Velden (2005): "Coordination of Expectations in Asset Pricing Experiments," *The Review of Financial Studies*, 18(3), 955–980.

——— (2007): "Learning in Cobweb Experiments," *Macroeconomic Dynamics*, 11(S1), 8–33.

Kahneman, D. (2011): *Thinking, Fast and Slow*. Macmillan.

Kindleberger, C. P., and R. Z. Aliber (2011): *Manias, Panics and Crashes: A History of Financial Crises*. Palgrave Macmillan.

Kurz, M., and M. Motolese (2011): "Diverse Beliefs and Time Variability of Risk Premia," *Economic Theory*, 47(2-3), 293–335.

Lucas, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(2), 103–124.

——— (1986): "Adaptive Behavior and Economic Theory," *The Journal of Business*, 59(4), S401–S426.

Lux, T., and S. Schornstein (2005): "Genetic learning as an explanation of stylized facts of foreign exchange markets," *Journal of Mathematical Economics*, 41(12), 169–196.

Malmendier, U., and S. Nagel (2009): "Learning from Inflation Experiences," Discussion paper, Available online at `http://faculty-gsb.stanford.edu/nagel/documents/InflExp.pdf`.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29(3), 315–335.

Nunes, R. (2010): "Inflation Dynamics: The Role of Expectations," *Journal of Money, Credit and Banking*, 42(6), 1161–1172.

Pesaran, M. H. (1987): *The Limits to Rational Expectations*. Blackwell Oxford.

Pfajfar, D., and B. Žakelj (2014): "Experimental evidence on inflation expectation formation," *Journal of Economic Dynamics and Control*, 44, 147–168.

Reinhart, C. M., and K. Rogoff (2009): *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.

Sims, C. A. (2010): "Rational Inattention and Monetary Economics," in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3, pp. 155–181. Elsevier.

Tversky, A., and D. Kahneman (1974): "Judgment under Uncertainty: Heuristics and Biases," *Science*, 185(4157), 1124–1131.

van de Velden, H. (2001): "An experimental approach to expectation formation in dynamic economic systems," Ph.D. thesis, Tinbergen Institute and Universiteit van Amsterdam.

Wagener, F. (2014): "Expectations in Experiments," *Annual Review of Economics*, 6(1), 421–443.

Woodford, M. (2010): "Robustly Optimal Monetary Policy with Near-Rational Expectations," *American Economic Review*, 100(1), 274–303.

——— (2013): "Macroeconomic Analysis Without the Rational Expectations Hypothesis," *Annual Review of Economics*, 5(1), 303–346.

Appendices

A Definition of forecasting rules

Table 11 provides the exact specification for all prediction rules used in the paper. For the full specification of the HSM with two heuristics, see Anufriev, Hommes, and Philipse (2013). For the full specification of the HSM with four heuristics, see Anufriev and Hommes (2012).

Rule	**Prediction**
Homogeneous rules	
Trend extrapolation	$p_{t-1} + \gamma(p_{t-1} - p_{t-2})$ with $\gamma = 1$
Adaptive	$wp_{t-1} + (1-w)p^e_{\text{previous}}$, with $w = 0.75$
Contrarian	$p_{t-1} - 0.5(p_{t-1} - p_{t-2})$
Naive	p_{t-1}
RE	p^f
Heterogeneous rules	
HSM from AHP	switching between 2 heuristics: trend extrapolation and adaptive expectations, as specified above; learning parameters are $\beta = 1.5$, $\eta = 0.1$, $\gamma = 0.1$
HSM (4 heuristics)	switching between 4 heuristics: adaptive with $w = 0.65$, two trend extrapolation (with $\gamma = 0.4$ and $\gamma = 1.3$), and the anchor-and-adjustment rule; learning parameters are $\beta = 0.4$, $\eta = 0.7$, $\gamma = 0.9$
GA model	$\alpha_{i,t}p_{t-1} + (1-\alpha_{i,t})p^e_{i,\text{previous}} + \beta_{i,t}(p_{t-1} - p_{t-2})$
GA-P1	with restrictions $\alpha_{i,t} \in [0,1]$ and $\beta_{i,t} \in [-1.1, 1.1]$
GA-P2	with restrictions $\alpha_{i,t} \in [0,1]$ and $\beta_{i,t} \in [0, 1.1]$
GA-P3	with restrictions $\alpha_{i,t} \in [0,1]$ and $\beta_{i,t} \in [-1.3, 1.3]$
GA-P4	with restrictions $\alpha_{i,t} \in [0,1]$ and $\beta_{i,t} \in [0, 1.3]$

Table 11: Specification of the forecasting rules used in the paper. For the one-period ahead environments (**HHST09**; **BHST12**; **HSTV07**; **V01**), the rules generate prediction p^e_t, the adaptive rule includes $p^e_{\text{previous}} = p^e_{t-1}$, whereas GA model includes $p^e_{i,\text{previous}} = p^e_{i,t-1}$. For the two-period ahead environment **HSTV05**, the rules generate prediction p^e_{t+1}, and the adaptive rule includes $p^e_{\text{previous}} = p^e_t$, whereas GA model includes $p^e_{i,\text{previous}} = p^e_{i,t}$.

B Initialisation of the model

In this Appendix we discuss one aspect of initialisation of the GA model for the 50-period Monte Carlo simulations, namely the choice of distribution for the initial predictions. Recall that our task is to demonstrate that GA model can replicate experimental stylised facts. Two examples can be given for **HHST09** experiment to show that the initialisation of the model can be crucial in achieving this task.

First, under negative feedback, the individual price forecasts coordinated only after the price itself has already converged. To replicate this feature in our simulations, one has to start with a similar degree of initial heterogeneity in the agents predictions and then to show that due to the learning of GA agents it can disappear as it happened in the experiment.

Second, under positive feedback, as Anufriev, Hommes, and Philipse (2013) suggest, price oscillations emerged in the groups where the average of the first predictions was relatively far from the fundamental price. Therefore, in this set-up the initial individual predictions influenced later outcomes, such as appearance and characteristics of oscillations, or dynamics of coordination. One would like to have a model that can mimic this *path-dependence*. But without a realistic initialisation, the path-dependent model will not fit the data well.

How did subjects make predictions in the very first period of the experiment, when the information set of past prices and predictions is empty? Diks and Makarewicz (2013) investigate this issue in a systematic fashion for the case of the **HHST09** experiment. They argue that the initial subject predictions can be regarded as a sample from a common distribution, which they estimate. We use their methodology and estimate a distribution of initial predictions for all other experiments. In those MC simulations, where the initial predictions are sampled from the distribution, this distribution is the one estimated from the respective experiment.

HHST09 For this experiment we use the estimated Winged Focal Point (WFP) reported by Diks and Makarewicz (2013), which is given by

$$p^e_{i,1} = \begin{cases} \varepsilon^1_i \sim U(9.546, 50) & \text{with probability } 0.45739\,, \\ 50 & \text{with probability } 0.30379\,, \\ \varepsilon^2_i \sim U(50, 62.793) & \text{with probability } 0.23882\,, \end{cases} \tag{13}$$

where $U(a,b)$ is the uniform distribution on interval $[a,b]$. Around 1/3 would predict 50, a mid-point of the suggested interval for the initial price forecast $[0,100]$. Others were spread around this focal point with more people predicting low price and almost nobody predicting price higher than 60. Hence the distribution is a composite of a unit mass at 50 and two 'wings', uniform distributions spreading from the focal point. See Fig. 11 for visualisation.

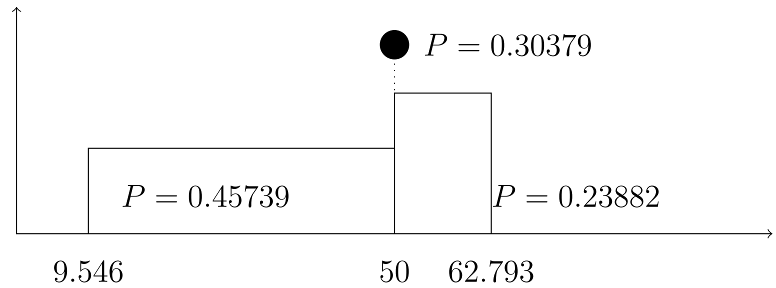

Figure 11: Estimated density of winged focal point distribution for **HHST09** as given in Eq. 13. The sizes of the wings around the mass point $p = 50$ are scaled to their masses and lengths.

BHST12 We reestimate WFP model for the data reported by **BHST12** using the same methodology as reported by Diks and Makarewicz (2013). This leads to WFP specified as

$$p^e_{i,1} = \begin{cases} \varepsilon^1_i \sim U(16.406, 50) & \text{with probability } 0.32296, \\ 50 & \text{with probability } 0.35159, \\ \varepsilon^2_i \sim U(50, 70.312) & \text{with probability } 0.32296. \end{cases}$$

V01; HSTV07 In the case of the cobweb economy experiment, the subjects were asked to predict prices in the $[0, 10]$ interval. Interestingly, the initial predictions still have the WFP form, with a large proportion equal to the midpoint 5 and the rest (not necessarily rounded to a full integer) distributed around this new focal point. To account for that, we reestimate the WFP and obtain

$$p^e_{i,1} = \begin{cases} \varepsilon^1_i \sim U(1.875, 5) & \text{with probability } 0.17983, \\ 5 & \text{with probability } 0.36344, \\ \varepsilon^2_i \sim U(5, 7.5) & \text{with probability } 0.45673. \end{cases}$$

HSTV05 In this experiment, the predictions are two-period ahead, hence the subjects would have to give *two* initial predictions, $p^e_{i,1}$ and $p^e_{i,2}$. First period forecasts are similar to those from the other experiments. As for the second period, one can notice that 2/3 of the subjects, who would predict $p^e_{i,1} = 50$ the focal point in the first period, would do the same in the second period; otherwise they would again draw predictions resembling WFP, but with a substantially small weight on the focal point 50. Hence we follow Diks and Makarewicz (2013) and get the following estimations for the first period:

$$p^e_{i,1} = \begin{cases} \varepsilon^1_i \sim U(4.712, 50) & \text{with probability } 0.31306, \\ 50 & \text{with probability } 0.45536, \\ \varepsilon^2_i \sim U(50, 64.062) & \text{with probability } 0.23158. \end{cases}$$

To generate the second period predictions, we define the auxiliary draw

$$p_{i,2}^{aux} = \begin{cases} \varepsilon_i^1 \sim U(3.125, 50) & \text{with probability } 0.44958, \\ 50 & \text{with probability } 0.018761, \\ \varepsilon_i^2 \sim U(50, 67.227) & \text{with probability } 0.53166. \end{cases} \tag{14}$$

With the realisation from this draw, the second period predictions are defined as

$$p_{i,2}^{e} = \begin{cases} p_{i,2}^{aux} & \text{always if } p_{i,1}^{e} \neq 50, \\ p_{i,2}^{aux} & \text{with probability } 1/3 \text{ if } p_{i,1}^{e} = 50, \\ 50 & \text{with probability } 2/3 \text{ if } p_{i,1}^{e} = 50. \end{cases} \tag{15}$$

C Price autocorrelation in the cobweb experiment

Table 12 gives the first three autocorrelations of the experimental groups in **HSTV07** and the 50-period ahead simulations of the GA and benchmark models.

Treatments	**Stable**			**Unstable**			**Strongly unstable**		
Autocorrelation	ρ_1	ρ_2	ρ_3	ρ_1	ρ_2	ρ_3	ρ_1	ρ_2	ρ_3
Experiment	−0.1878	0.06323	−0.12	−0.2948	0.01363	−0.01114	−0.1973	0.211	0.02144
Trend extrapolation	−0.9661	0.9423	−0.9209	−0.9655	0.9404	−0.9159	−0.9639	0.9403	−0.918
Adaptive	−0.5996	0.3446	−0.3078	−0.9628	0.9235	−0.8927	−0.964	0.94	−0.9176
Contrarian	−0.257	−0.3006	0.1604	−0.4556	−0.4895	0.8202	−0.4756	−0.4704	0.8974
Naive	−0.9043	0.837	−0.7911	−0.967	0.9394	−0.9143	−0.9639	0.9403	−0.918
RE	0	0	0	0	0	0	0	0	0
HSM from AHP	−0.6528	0.4224	−0.3438	−0.9561	0.9153	−0.8816	−0.9639	0.9399	−0.9175
GA-AR1	−0.1161	0.008603	−0.1253	−0.1686	−0.005697	−0.1028	−0.2346	−0.09282	−0.02373
GA-P1	−0.1102	−0.3232	0.002674	−0.2201	−0.2013	0.0362	−0.2478	−0.3148	0.2432
GA-P2	−0.2955	0.1059	−0.171	−0.3882	0.1405	−0.1848	−0.6206	0.4428	−0.356

Table 12: **HSTV07**: 50-period ahead predictions. First three autocorrelations in prices for various models compared with the experimental data. Autocorrelations are averaged over six groups for each treatment.

D Formal definition of Genetic Algorithms

In this Appendix we present a formal definition of the Genetic Algorithms (GA) version, which served as the cornerstone of our model. It closely follows the standard specification suggested by Haupt and Haupt (2004) and used by Hommes and Lux (2013).

D.1 Optimisation procedures: traditional and Genetic Algorithms

Consider a maximisation problem where the target function $\mathcal{F}$ of N arguments $\theta = (\theta^1, \ldots, \theta^N)$ is such that a straightforward analytical solution is unavailable. Instead, one needs to use a numerical optimisation procedure.

Traditional maximisation algorithms, like the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, iterate a candidate argument for the optimum of the target function $\mathcal{F}$ by (1) estimating the curvature around the candidate and (2) using this curvature to find the optimal direction and length of the change to the candidate solution. This so called 'hill-climbing' algorithm is very efficient in its use of the shape of the target function. On the other hand, it will fail if the target function is 'ill-behaved': non-continuous or almost flat around the optima, has kinks or breaks. Here the curvature cannot be reliably estimated. Another problem of a computational nature is that the BFGS may perform poorly for a problem of large dimensionality.

The Genetic Algorithms are based on a fundamentally different approach and therefore can be used for a wider class of problems. The basic idea is that we have a population of arguments which compete *only* in terms of their respective function value. This competition is modelled in an evolutionary fashion: mutation operators allow for a blind-search experimentation, but the probability that a particular candidate will survive over time is relative to its functional value. As a result, the target function may be as general as necessary, while the arguments can be of any kind, including real numbers, integers, probabilities or binary variables. The only constraint is that each argument must fall into a predefined dense interval (or set) $[a_n, b_n]$.

D.2 Binary strings

A Genetic Algorithm (GA) uses H chromosomes $g_{h,t} \in \mathbb{H}$ which are binary strings divided into N genes $g^n_{h,t}$, each encoding one candidate parameter $\theta^n_{h,t}$ for the argument θ^n. A chromosome $h \in \{1, \ldots, H\}$ at time $t \in \{1, \ldots, T\}$ has predetermined length L and is specified as

$$g_{h,t} = (g^1_{h,t}, \ldots, g^N_{h,t}), \tag{16}$$

such that each gene $n \in \{1, \ldots, N\}$ has its length equal to an integer L_n (with $\sum_{n=1}^{N} L_n = L$) and is a string of binary entries (bits)

$$g^n_{h,t} = (g^{n,1}_{h,t}, \ldots, g^{n,L_n}_{h,t}), \quad g^{n,l}_{h,t} \in \{0, 1\} \text{ for each } j \in \{1, \ldots, L_n\}. \tag{17}$$

The relation between the genes and the arguments is straightforward. An integer θ^n is simply encoded by (17) with its binary notation. Consider now an argument θ^n which is a probability. Notice that $\sum_{l=0}^{L_n-1} 2^j = 2^{L_n} - 1$. It follows that a particular gene $g^n_{h,t}$ can be decoded as a normalised sum

$$\theta^n_{h,t} = \sum_{l=1}^{L_n} \frac{g^{n,l}_{h,t} 2^{l-1}}{2^{L_n} - 1}. \tag{18}$$

A gene of zeros only is therefore associated with $\theta_n = 0$, a gene of ones only – with $\theta_n = 1$, while other possible binary strings cover the $[0, 1]$ interval with an $\frac{1}{2^{L_n}-1}$ increment. Any desired precision can be achieved with this representation. Since $2^{-10} \approx 10^{-3}$, the precision close to one over trillion (10^{-12}) is obtained by a mere of 40 bits.

A real variable θ^n from an $[a_n, b_n]$ interval can be encoded in a similar fashion, by an affine transformation of a probability:

$$\theta^n_{h,t} = a_n + (b_n - a_n) \sum_{l=1}^{L_n} \frac{g^{n,l}_{h,t} 2^{l-1}}{2^{L_n} - 1} \tag{19}$$

where the precision of this representation is given by $\frac{b_n - a_n}{2^{L_n}-1}$. Notice that one can approximate an unbounded real number by reasonably large a_n or b_n, since the loss of precision is easily undone by a longer string.

D.3 Evolutionary operators

The core of GA are evolutionary operators. GA iterates the population of chromosomes for T periods, where T is either large and predefined, or depends on some convergence criterion. First, at each period $t \in \{1, \dots, T\}$ each chromosome has its fitness equal to a monotone transformation of the function value $\mathcal{F}$. This transformation is defined as $V(\mathcal{F}(\theta_{h,t})) \equiv V(h_{k,t}) \to \mathbb{R}^+ \cap \{0\}$. For example, a non-negative function can be used directly as the fitness. If the problem is to minimise a function, a popular choice is the exponential transformation of the function values, similar to the one used in the logit specification of the Heuristic Switching Model (Brock and Hommes, 1997).

Chromosomes at each period can undergo the following evolutionary operators: procreation, mutation, crossover and election. These operators first generate an offspring population of chromosomes from the parent population t and therefore transform both populations into a new generation of chromosomes $t + 1$ (notice the division of the process).

D.3.1 Procreation

For the population at time t, GA picks subset $\mathbb{X} \subseteq \mathbb{H}$ of χ chromosomes and picks $\kappa < \chi$ of them into a set $\mathbb{K}$. The probability that the chromosome $h \in \mathbb{X}$ will be picked into $\mathbb{K}$ as its zth element (where $z \in \{1, \ldots, \kappa\}$) is usually defined by the power function:

$$Prob(g_z = g_{h,t}) = \frac{V(g_{h,t})}{\sum_{j \in \mathbb{X}} V(g_{j,t})}. \tag{20}$$

This procedure is repeated with differently chosen $\mathbb{X}$'s until the number of chromosomes in all such sets $\mathbb{K}$'s is equal to H. For instance, the *roulette* is procreation with $\chi = H$ and $\kappa = 1$: GA picks randomly one chromosome from the whole population, where each chromosome has probability of being picked equal to its function value relative to the function value of all other chromosomes. This is repeated exactly H times.

So called *tournaments* are often used for the sake of computational efficiency. Here, $\chi << H$. For instance, GA could divide the chromosomes into pairs and sample two offspring from each pair.

Procreation is modelled as the basic natural selection mechanism. We consider subsets of the original population (or maybe the whole population at once). Out of each such a subset, we pick a small number of chromosomes, giving advantage to these which perform better. We repeat this procedure until the offspring generation is as large as the old one. Thus the new generation is likely to be 'better' than the old one.

D.3.2 Mutation

For each generation $t \in \{1, \ldots, T\}$, after the procreation has taken place, each binary entry in each new chromosome has a predefined δ_m probability to mutate: ones turned into zeros and vice versa. In this way the chromosomes represent different numbers and may therefore attain better fit.

The mutation operator is where the binary representation becomes most useful. If the bits, which are close to the beginning of the gene, mutate, the new argument will be substantially different from the original one. On the other hand, small changes can be obtained by mutating bits from the end of the gene. Both changes are equally likely! In this way, GA can easily evaluate arguments which are both far away from and close to what the chromosomes are currently encoding. As a result, GA efficiently converges to the maximum, but are also likely *not* to get stuck on a local maximum. This is clearly independent of the initial conditions, which gives GA additional advantage over hill-climbing algorithms (like BFGS), where a good choice of the initial argument can be crucial to obtain the global maximum.

D.3.3 Crossover

Let $0 \leqslant C_L, C_H \leqslant \sum_{n=1}^{N} L_n = L$ be two predefined integers. The crossover operator divides the population of chromosomes into pairs. If $C_L < L - C_H$, it exchanges the first C_L and the last C_H bits between chromosomes in each pair with a predefined probability δ_c. Otherwise, the crossover operator exchanges $max\{C_L, C_H\}$ bits in each pair of chromosomes with this predefined probability δ_c. This operator facilitates experimentation in a different way than the mutation operator. Typically, it is set to exchange whole arguments, that is there are $0 \leqslant \nu_L \leqslant \nu_H \leqslant N$ such that $C_L = \sum_{n=1}^{\nu_L} L_n$ and $C_H = \sum_{n=\nu_H}^{N} L_n$. This allows the chromosomes to experiment with different compositions of the individual arguments, which on their own are already successful.

D.3.4 Election

The experimentation done by the mutation and crossover operators does not need to lead to efficient binary sequences. For instance, a chromosome which actually decodes the optimal argument should not mutate at all. To counter this effect, it is customary to divide the creation of a new generation into two stages. First, the chromosomes procreate and undergo mutation and crossover in some predefined order. Next, the resulting set of chromosomes is compared in terms of fitness with the parent population. Thus, offspring will be passed to the new generation only if it *strictly* outperforms the parent chromosome. In this way each generation will be at least as good as the previous one, what in many cases facilitates convergence.

E Parametrisation of the forecasting heuristic

In this Appendix, we will address two issues. First, we will investigate the importance of the anchor in the forecasting heuristic both for the one-period ahead **HHST09** and for the two-period ahead **HSTV05** settings. Second, we study the proper degree of allowed trend extrapolation, based on the linear feedback from **HHST09**.

E.1 Is the anchor important for HHST09?

HHST09 show that most of their subjects (around 60%) use First-Order prediction rule with heterogeneous parameter specification:

$$p^e_{i,t} = \alpha_1 p_{t-1} + \alpha_2 p^e_{i,t-1} + \alpha_3 60 + \beta(p_{t-1} - p_{t-2}) \tag{21}$$

where the fundamental price 60 serves as an anchor,[31] α' span a simplex ($\alpha_1 + \alpha_2 + \alpha_3 = 1$) and β is the trend extrapolation coefficient. Our rule (5) is a special case of (21) with the restriction that $\alpha_3 = 0$, which implies that fixed anchor is not used by the agents.

Experimental literature suggests that, in general, anchors and focal points are important in describing human behaviour. However, **HHST09** report that the anchor weight α_3 is typically significant for the subjects under negative feedback treatment, while most of the subjects under positive feedback treatment would not use it. Furthermore, under negative feedback prices and predictions converge to the vicinity of 60, which in practice makes the coefficients α sample-unidentifiable; and could also make redundant the anchor itself. When designing our GA model, we therefore investigated whether the anchor has any additional explanatory power.

To simplify econometric issues, in the previous literature the anchor was set at the fundamental level, which however was not directly given to the subjects. It is more plausible that the subjects used the average of all previous prices as an anchor. We will use thus anchored-FOR specified as

$$p^e_{i,t} = \alpha_1 p_{t-1} + \alpha_2 p^e_{i,t-1} + \alpha_3 \left(\frac{1}{t-1} \sum_{s=1}^{t-1} p_s \right) + \beta(p_{t-1} - p_{t-2}). \tag{22}$$

We run the Monte Carlo (MC) simulations exactly as in the first part of Section 3.2, but for the GA model based on (22) with the restriction for $\beta \in [-1.1, 1.1]$. The results are presented on Fig. 12. We observe for the positive feedback that, in contrast to our restricted model without an anchor, the GA model based on FOR as in (22) does not predict oscillations at all. Instead a sluggish convergence towards the fundamental is generated, as can be seen in the stable median price, bounded by relatively narrow 95% CI. In other words, this specification misses most of the dynamics observed in half of the experimental groups. We conclude that there is no evidence for a need of an anchor, specified as a long-run average of the observed prices, in our GA model.

[31] Notice that what is the anchor, can be a matter of interpretation. One may think of rule (5) as an anchor-based rule as well, since it can be rewritten as a rule that adjusts the anchor given by the previous price forecast with the latest observed price and trend.

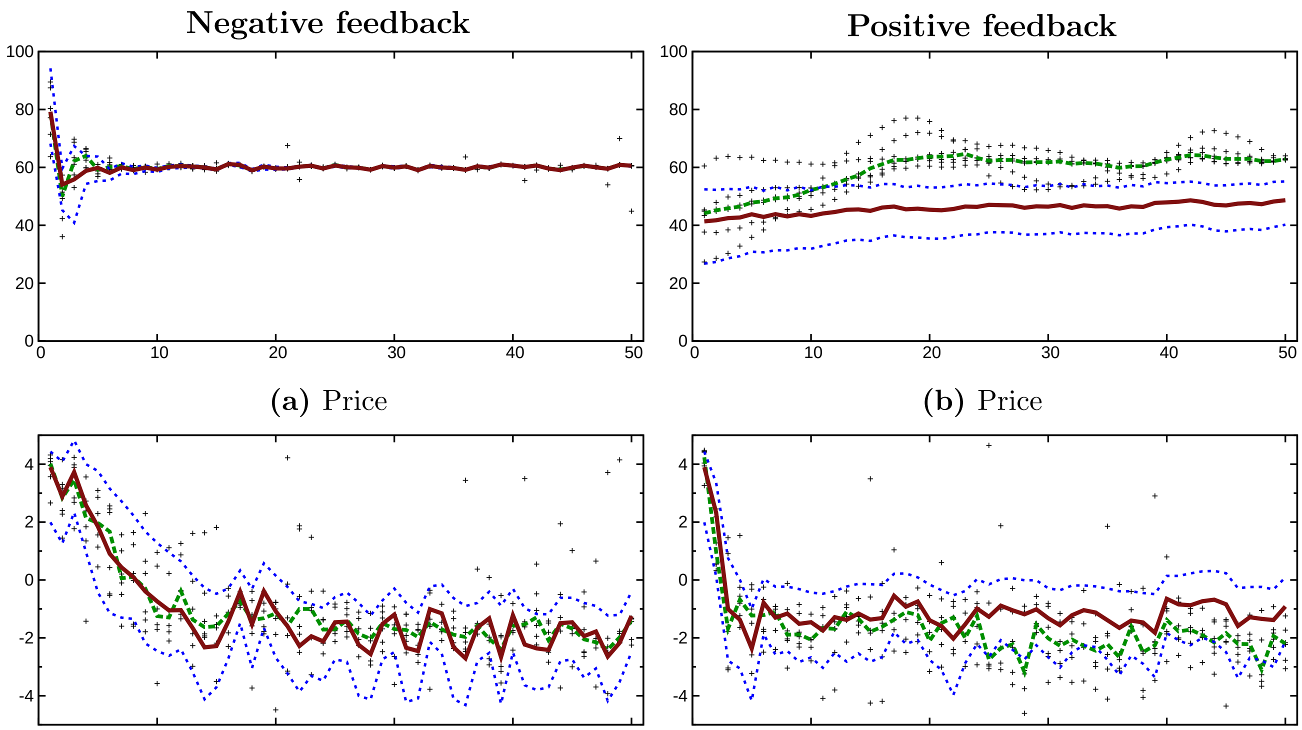

(a) Price

(b) Price

(c) SD of individual predictions

(d) SD of individual predictions

Figure 12: **HHST09**: 50-period ahead Monte Carlo simulation (1000 markets) for the **GA-P1** model with the anchored-FOR compared with the experimental data. *Upper panels:* price. *Lower panels:* degree of coordination ($\log_2$ scale). Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median and blue dotted lines are the 95% confidence interval for the GA model.

E.2 Anchor and HSTV05

The **HSTV05** non-linear, two-period ahead LtF asset pricing market resulted in much more unruly oscillations than those observed in the simple linear experiment **HHST09** under positive feedback. One could therefore think that some kind of a long-run anchor might have been important for the subjects, even though they would not use it in one-period ahead forecasting setting. Furthermore, in the experiment the oscillations typically unravelled around the fundamental price, which again suggests that the subjects tried to anchor the price changes to it. To address this issue, we run the 50-period ahead MC simulation like in Section 4.3, but where the rule (12) is replaced by the anchored-FOR rule (22) appropriately adapted for the two-period ahead setting, and where the anchor was given by the fundamental price $p^f = 60$.

Results for two parametrisations (with allowed trend extrapolation $\beta \in [-1.1, 1.1]$ and $\beta \in [-1.3, 1.3]$) are presented on Fig. 13. Just as in the case of **HHST09**, we find that the GA model with anchored-FOR rule generates sluggish convergence towards the fundamental price from below. Indeed, in contrast to **HHST09**, the 95% CI of the GA model's prices do not include the fundamental $p^f = 60$ even after 50 periods. This indicated that adding an anchor to

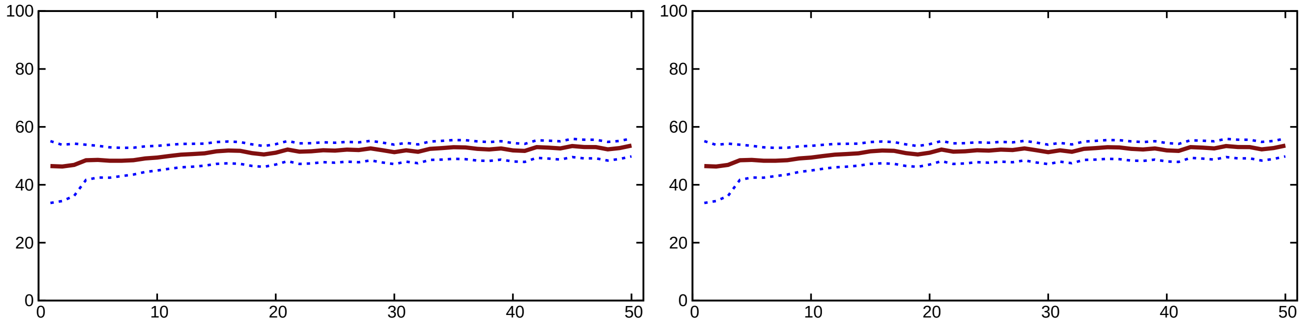

(a) GA-P1, $\beta \in [-1.1, 1.1]$ **(b) GA-P3**, $\beta \in [-1.3, 1.3]$

Figure 13: HSTV05 with $p^f = 60$: 50-period ahead Monte Carlo simulation (1000 markets) for the **GA-P1** (*left panel*) and **GA-P3** (*right panel*) models with anchored-FOR. Price evolution is shown. Red line is the median and blue dotted lines are the 95% CI.

the GA model would decrease its fitness to the experimental data.[32]

E.3 Degree of trend extrapolation

Recall that the GA requires a predefined finite interval for the optimised parameters. In the case of our GA model based on (5), the price weight is confound to $\alpha \in [0, 1]$, but *prima facie* there is no 'natural' bound for the trend extrapolation $\beta \in [\beta_L, \beta_H]$, since *a priori* we do not know the degree of trend extrapolation that people consider while forecasting prices. As mentioned in Section 3, we argue that the model performs well if we specify the (5) rule to use 1.1 as the upper bound to the trend (as in **GA-P1** and **GA-P2** models).

It turns out (not surprisingly) that the allowed trend extrapolation interval has little effect on the behaviour of our GA model under negative feedback. However, the effect exists for the model under positive feedback: the larger the interval $\beta \in [\beta_L, \beta_H]$ is, the bigger the amplitude of the price fluctuations is. We experimented with different bounds, trying to calibrate the model to the experimental oscillations. We used the same Monte Carlo experiments as in the first part of Section 3.2.

Allowing for a high trend extrapolation $\beta \in [-1.5, 1.5]$ results in a model with huge possible oscillations and little predictive power, see Fig. 14. On the other hand, parametrisation with $\beta \in [-0.5, 0.5]$ has narrow CI, but predicts small oscillations, see Fig. 15. We found the model with $\beta \in [-1.1, 1.1]$ is the best trade-off between in-sample fit and out-sample predictive power of the model.

This result reflects the experimental findings. **HHST09** find that under positive feedback, four out of twenty estimated rules had $\beta > 0.9$ and further five rules had $\beta > 0.75$. Nevertheless, **HHST09** in their estimations impose a restriction that $\beta \in [-1, 1]$. Our GA model suggests that such a restriction is inconsistent with the degree of experimental price oscillations.

[32]We found similar results when the anchor was specified as the average price so far $\frac{1}{t-1}\sum_{s=1}^{t-1} p_s$.

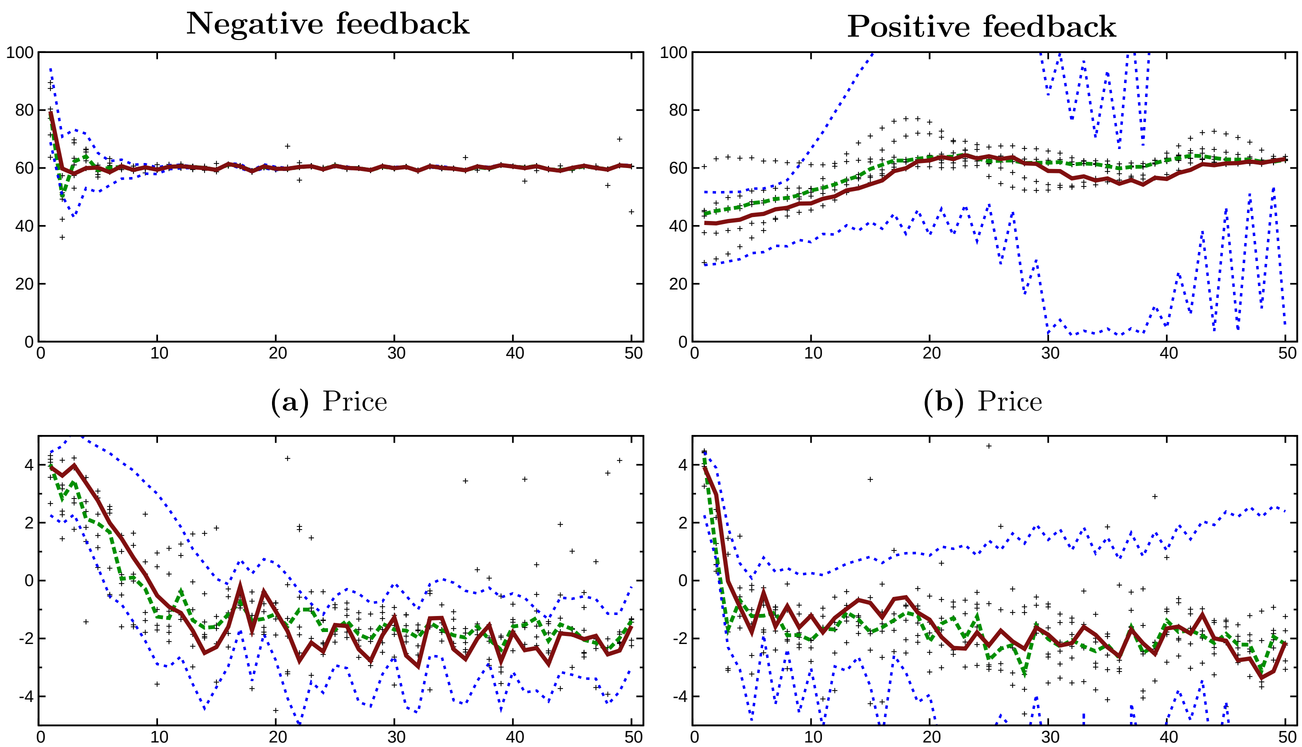

(a) Price

(b) Price

(c) SD of individual predictions

(d) SD of individual predictions

Figure 14: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the model with restriction $\beta \in [-1.5, 1.5]$ compared with the experimental data. *Upper panels:* price. *Lower panels:* degree of coordination. Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median and blue dotted lines are the 95% confidence interval for the GA model.

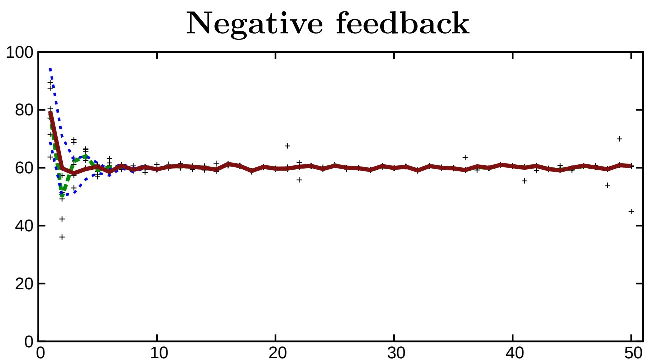

(a) Price

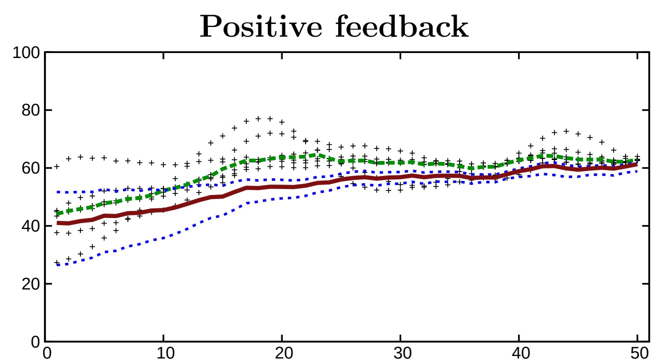

(b) Price

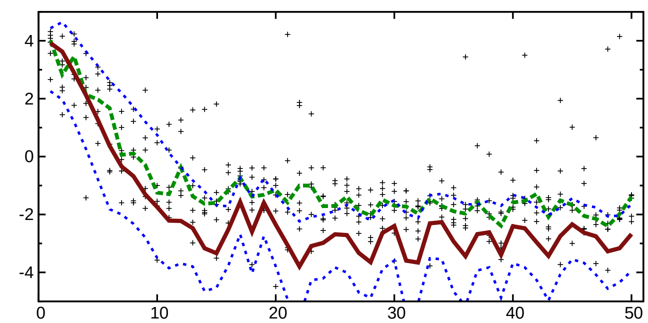

(c) SD of individual predictions

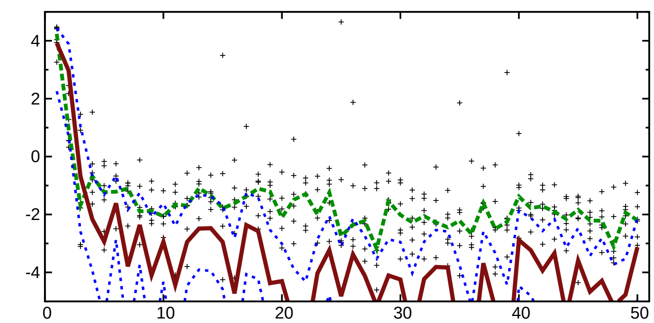

(d) SD of individual predictions

Figure 15: **HHST09**: 50-period ahead Monte Carlo simulation (1000 markets) for the model with restriction $\beta \in [-0.5, 0.5]$) compared with the experimental data. *Upper panels:* price. *Lower panels:* degree of coordination. Green dashed line and black pluses show the experimental median and group observations, respectively. Red thick line is the median and blue dotted lines are the 95% confidence interval for the GA model.