



UvA-DARE (Digital Academic Repository)

Learning to forecast: Genetic algorithms and experiments

Makarewicz, T.A.

[Link to publication](#)

Citation for published version (APA):

Makarewicz, T. A. (2014). Learning to forecast: Genetic algorithms and experiments Amsterdam: Tinbergen Institute

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 2

Learning-to-Forecast with Genetic Algorithms

2.1 Introduction

Expectations are a cornerstone of many economic models, because economic agents often operate in a dynamic context. Consumers have to organize their life-time work and consumption paths, while companies decide on how to build up future production capabilities. In either case, the agents must forecast how the uncertain future may unfold. What makes modeling predictions difficult is that they typically form a feedback with the realizations through agents decisions. For instance, if producers expect an increased price of their consumption good, they will rise production. If the demand stays constant, this implies lower market clearing price in the future. It is therefore likely that agents would alter their predictions, leading to a new realized price.

Even if the agents know the structure of the economy, the price-expectation feedback can lead to non-trivial dynamics (Grandmont, 1985; Hommes, 2013; Tuinstra and Weddepohl, 1999). This picture becomes more complicated if the agents furthermore have to learn this structure (Bullard, 1994; Grandmont, 1998). Agents do want to form good price expectations, but how would they cope with this complexity?

The traditional literature (after Muth, 1961) emphasizes the Rational Expectations (RE) hypothesis, which states that in equilibrium the predictions have to be model consistent. Most economists would interpret RE as an ‘as-if’ approximation — real markets behave as if a representative agent was perfectly rational, because real people are rational enough to learn to avoid systematic, correlated errors.¹ However, this is

¹One interesting and straightforward explication of this approach can be found in the concluding section of Blundell and Stoker (2005).

not confirmed by the data. A recent example comes from the housing market in the US before the latest economic crisis, where people systematically misjudged the long-term value of their houses (Benítez-Silva et al., 2008; Case and Shiller, 2003; Goodman Jr. and Ittner, 1992). In a broader context of macroeconomics, inflation expectations formed by the ‘Jones’ are far from the RE predictions (Adam, 2007; Assenza et al., 2013; Charness et al., 2007; Pfajfar and Zakelj, 2011) and can be subject to cognitive biases (Malmendier and Nagel, 2009). Many firms similarly fail to use RE (see *e.g.* Nunes, 2010b, for a discussion about the Phillips Curve and survey expectations).

The failure of RE made many economists look for an alternative model with explicit learning. But alternatives face the so called ‘wilderness of bounded rationality’ problem: there is a myriad of possible learning mechanisms with varied restrictions on human memory and computational capabilities. These range from simple linear heuristic models (see Evans and Ramey, 2006, for a discussion of adaptive expectations), through adaptive or statistical learning (Evans and Honkapohja, 2001), through heuristic switching type of models (Brock and Hommes, 1997) to evolutionary learning mechanisms (Arifovic et al., 2012). Typically these mechanisms lead to different dynamics: for example Bullard (1994) and Tuinstra and Wagener (2007) show that for a standard OLG economy, where the agents use OLS learning for price forecasting, adaptive learning may lead both to stable and complicated, chaotic dynamics.

Learning-to-Forecast (LtF) experiments (Hommes, 2011; Hommes et al., 2005) offer a simple laboratory testing ground for learning mechanisms. These controlled experimental economies have a straightforward and unique fundamental (RE) equilibrium. As in real markets, subjects observe the realized 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. Many LtF laboratory experiments contradict the RE hypothesis. The subjects can coordinate on oscillating and serially correlated time series, and the exact dynamics depend greatly on the specific feedback structure of the experimental economy. Convergence to the fundamental equilibrium happens only under severe restrictions on the underlying law of motion (Hommes, 2011). Another important finding in the experiments is heterogeneity: within the same experimental group, subjects tend to use different forecasting rules, see Heemeijer et al. (2009) (henceforth HHST09).

The most successful attempt to explain the LtF experiments comes with the so-called Heuristic Switching Model (HSM; Brock and Hommes, 1997). The basic idea of the model is that the agents have a (small) set of simple forecasting heuristics (rules of thumb like adaptive or trend extrapolating expectations) and choose those that had a better relative past performance. The HSM has successfully been used

to explain different types of aggregate behavior — convergence and oscillations in various experimental settings (Anufriev and Hommes, 2012; Anufriev et al., 2013). A disadvantage of the HSM however is the small set of heuristics, which cannot fully account for the individual heterogeneity. Furthermore different experiments require the HSM to utilize differently specified sets of heuristics. It is unclear why the subjects would use only these particular forecasting rules and how they would learn them.

The purpose of this chapter is to explain LtF experiments with Genetic Algorithms (GA). The basic idea of the model is that the agents forecast prices using a possibly large set of heuristics from a simple but general class. The agents then *independently* use GA to update and select the heuristics based on their relative success. This results in an agent-based model with explicit individual learning and endogenous heterogeneity. We will argue that our model is able to capture the dynamics at both the *aggregate* and the *individual* level for different experimental settings. Furthermore, the model has a clear link to the existing behavioral literature. GA are a flexible optimization procedure, thus the GA-based model retains a basic economic interpretation. Agents, who use GA, have to rely on second-best forecasting rules. Nevertheless, they learn to use them efficiently. For example, if it is profitable to harvest speculative trade revenues, the agents will update their forecasting rule's parameters with GA, in the direction of stronger trend extrapolation. As a result, the GA model in spirit resembles the HSM, but is more flexible in the specification and evolution of forecasting heuristics.

Dawid (1996) provides a good overview of the first GA applications to economic problems. GA was initially applied in its social learning form to explore stylized facts from experimental data, outperforming the RE hypothesis (Arifovic, 1995), with the examples of the exchange rate volatility (Arifovic, 1996; Lux and Schornstein, 2005) or production level choices in a cobweb producers economy (Dawid and Kopel, 1998). More recently, Hommes and Lux (2013) investigate a model, in which agents use GA to optimize a forecasting heuristic (instead of directly optimizing a prediction) and, much like the actual subjects in LtF experiments, cannot observe each others behavior or strategies. The authors replicate the distribution of the predictions and prices (mean, variance and autocorrelations) of the cobweb experiments by Hommes et al. (2007) and van de Velden (2001) (henceforth HSTV07 and V01 respectively).

We follow Hommes and Lux (2013) in the basic design of our model, but the novelty of this chapter is fourfold. The first is that we will use a different heuristic space, based on the so called first order rule, which is a mixture of adaptive and trend extrapolating heuristics. This gives the model better micro-foundations, as HHST09 find this forecasting rule to well describe the individual expectations in their experiment. The GA agents then learn to optimize the parameters of this simple heuristic.

The second novelty is that our model allows for a simultaneous explanation of *different* LtF experiments, based on positive and negative feedback markets with or without breaks in the fundamental price, or with highly non-linear price-expectations feedback. Application to different market settings will demonstrate the generality of our model. In particular, we will look at four market settings: (1) the simple, linear feedback systems from HHST09; (2) the linear price-expectations feedback system with unexpected large shocks to the fundamental price (Bao et al., 2012) (henceforth BHST12); (3) a cobweb producers economy (HSTV07; V01), used also by Hommes and Lux (2013); and (4) a non-linear positive feedback asset pricing economy, where the subjects are asked for two-period ahead predictions (Hommes et al., 2005) (henceforth HSTV05).

The third novelty is that we explain the *individual* behavior observed in the LtF experiments. We will show with Monte Carlo studies that our model can replicate the long-run behavior of the data, both at the *aggregate* and *individual* level. Next we will 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 the *prices* and the *individual price forecasts* one period ahead. This is an important contribution to the literature on agent-based models, which usually focuses on a model's fit to the aggregate stylized facts.

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

The chapter is organized as follows. In Section 2, we present the setup and findings of the LtF experiments, and briefly discuss the HSM by Anufriev et al. (2013). In the third section, we introduce our GA model and fit it to the experimental setup by HHST09. In the fourth section we move to the remaining three experiments. Finally, a concluding section gives an overview of the results and suggestions for future research. A number of supplementary issues, including discussions on the initialization of the model and the forecasting rules based on a long-run anchor, are presented in the appendix.

2.2 Learning to Forecast and Heuristic Switching

Consider a market with a number of subjects $i \in \{1, \dots, I\}$, who are asked at each period t to forecast the price of a certain good. The subjects are explicitly informed that they act as forecasting consultants for firms and are rewarded only for the accuracy of the predictions.

The feedback relationship between the prices and predictions is summarized by a law of motion of the form

$$(2.1) \quad p_t = F(p_{1,t}^e, \dots, p_{I,t}^e),$$

where the realized price p_t is a function of all individual forecasts $p_{i,t}^e$. The mapping $F(\cdot)$ is generated from market clearing, with aggregate supply and demand derived from optimal choices of firms, consumers or investors, given the subjects' individual forecasts. Define the fundamental price p^f as the steady state RE outcome, the self-consistent prediction: $p^f = F(p^f, \dots, p^f)$. In all examples below the RE fundamental price is unique.

Unlike the RE agents, subjects in the experiment have limited information about the market. They are informed that their predictions affect the prices, but they are given only a qualitative story about this feedback. Moreover, they are not explicitly informed about the fundamental price.²

One important example investigated by HHST09 uses a linear version of (2.1):

$$(2.2) \quad p_t = A + B \left(\frac{\sum_{i=1}^I p_{i,t}^e}{I} - A \right) = A + B (\bar{p}_t^e - A),$$

where $\bar{p}_t^e = \frac{\sum_{i=1}^I p_{i,t}^e}{I}$ is the average prediction of all individuals at period t and $A = p^f$ is the fundamental price. There are two important cases: $B > 0$ (positive feedback) and $B < 0$ (negative 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 (although not perfectly if $B \neq 1$). Negative feedback arises *e.g.* in a supply driven market where producers face a lag in production. If they expect a high price in the future, they will increase production and so the market clearing price will go down.

²Usually it is possible to infer it from the experimental instructions. Anecdotal evidence suggests that even economics students, including graduate students, fail to realize it.

HHST09 used two simple linear treatments:

$$(2.3) \quad \text{Positive feedback: } p_t = 60 + \frac{20}{21}(\bar{p}_t^e - 60) + \varepsilon_t;$$

$$(2.4) \quad \text{Negative feedback: } p_t = 60 - \frac{20}{21}(\bar{p}_t^e - 60) + \varepsilon_t,$$

where $\varepsilon_t \sim NID(0, 0.25)$ is a small noise term. The experiment runs for 50 periods for each group of $I = 6$ subjects. The two linear treatments are symmetrically opposite. They have the same unique fundamental price $p^f = 60$ and the same absolute dampening factor $|B| = \frac{20}{21}$ (but with opposite signs). The dampening factors were chosen so that under naive expectations (*i.e.* $\bar{p}_t^e = p_{t-1}$), the fundamental price for both treatments is a (unique) stable steady state, but the system would still require some time to converge.

The two feedback treatments resulted in very different aggregate price behavior, illustrated in Figure 2.1a and 2.1b. Under the negative feedback after a short volatile phase of 7 – 8 periods, the price converges to the fundamental value $p^f = 60$, after which the subjects forecasts coordinate on the fundamental as well. In most of the positive feedback groups, persistent price oscillations arise (Figure 2.1b), where the price overshoots and undershoots p^f ; if the price converges to the fundamental at all, it does so only towards the end of the experiment (which happened for two out of seven cases). In spite of the price oscillations, the subjects' forecasts coordinate within 2 – 3 on a common value (different from the fundamental value) and remain so until the end of the experiment. In positive feedback markets, subjects' forecasts are thus strongly coordinated, but on a non-RE price.

To describe the subjects' forecasting behavior, HHST09 use the first-order rule (FOR):

$$(2.5) \quad p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + \alpha_3 p^f + \beta(p_{t-1} - p_{t-2}),$$

for $\alpha_1, \alpha_2, \alpha_3 \geq 0$, $\alpha_1 + \alpha_2 + \alpha_3 = 1$, $\beta \in [-1, 1]$. Rule (2.5) is an anchor and adjustment rule, extrapolating a price change from an anchor, which is given by a weighted average of past price, individual forecast and the fundamental price $p^f = 60$.³ HHST09 estimated this simple rule separately for each subject, based on their predictions from the last 40 periods. It described well the forecasting behavior of around 60% individuals.

HHST09 find that the individual forecasting rules varied between the subjects, even within the same treatment. The authors also report significant differences between

³Under RE, the FOR in (2.5) should be specified with $\alpha_1 = \alpha_2 = \beta = 0$, which implies that the subjects always predict the fundamental price, $p_{i,t}^e = p^f = 60$.

the two treatments. Under positive feedback, subjects focused on trend extrapolation ($\beta > 0$) and the estimated weight of the fundamental price α_3 coefficients were typically insignificant. Under the negative feedback, the reverse holds: trend extrapolation is barely used ($\beta \approx 0$), while the weight for the fundamental price ($\alpha_3 > 0$) is significant. This shows that a model with a homogeneous forecasting rule (RE, but also linear heuristics like trend extrapolation or naive expectations) may explain one of the two treatments, but not both at the same time. Moreover, a homogenous rule contradicts the significant differences between the subjects within each treatment.

This led Anufriev et al. (2013) to investigate the Heuristic Switching Model (HSM), in which the subjects are endowed with two prediction heuristics:

adaptive expectations: $p_{i,t}^e = \alpha p_{i,t-1} + (1 - \alpha)p_{i,t-1}^e$ with $\alpha \in [0, 1]$,

trend extrapolation: $p_{i,t}^e = p_{i,t-1} + \beta(p_{t-1} - p_{t-2})$ with $\beta \in [-1, 1]$,

where the authors used $\alpha = 0.75$ and $\beta = 1$. Notice that both heuristics are a special case of the first-order rule. 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 one with the higher relative past performance. Under positive feedback, agents easily coordinate their predictions, for example below the fundamental, close to the first observed price, but (by the construction of the positive feedback equation) the next realized price is then slightly higher than the average prediction. The trend extrapolation heuristic captures this gradual increase of initial prices and so becomes popular among the agents. This reinforces the trend and leads to persistent price oscillations. In contrast, under negative feedback there is no possibility of coordination of individual forecasts, unless the agents coordinate on the fundamental price. Under negative feedback the trend extrapolating rule performs poorly and agents switch to adaptive expectations, thus causing the price to converge to the fundamental.

HSM captures the essence of the aggregate forecasting behavior and successfully replicates the results of HHST09 in a stylized fashion. The drawback of the model is that the authors assume a limited number of only two heuristics, without explaining where these heuristics come from, that is, how the subjects are able to learn the two heuristics in the first place. Moreover, the HSM cannot account for heterogeneity of rules among subjects and hence does not explain the experiment at the individual level. To overcome these drawbacks, we will introduce a model with explicit individual learning through Genetic Algorithms.

2.3 The Genetic Algorithms model

2.3.1 Genetic Algorithms

Genetic Algorithms (GA) forms a class of numerical stochastic maximization procedures that mimic the evolutionary operations with which DNA of biological organisms adapts to the environment. GA were introduced to solve ‘hard’ optimization 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 (Haupt and Haupt, 2004). See *e.g.* 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), or chromosomes. They are retained into the next iteration with a probability that increases with their relative performance, which is defined directly in terms of a functional value (‘fitness’). This so called procreation operator means that with each iteration, the population of trial arguments is likely to have a higher functional value, *i.e.* be ‘fitter’. On top of the 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 (Arifovic, 1995).

Mutation At each iteration, every entry 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 utilizes 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 bits corresponding to a 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 The election operator is meant to screen inefficient outcomes of the experi-

mentation phase. This operator transmits the new chromosomes (selected from the old generation and treated with mutation and crossover) into the new generation only if their functional value is greater than that of the original argument. This operator ensures that once the routine finds the global solution, it will not diverge from it due to unnecessary experimentation.

The procreation routine and the three evolutionary operators have a straightforward economic interpretation for a situation, in which the agents want to optimize their behavioral rules, *e.g.* price 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.

An important additional condition for a GA routine is that it requires a predefined interval for each parameter. For the above example of updating behavioral rules through GA, it means that we confine them to some predetermined, finite grid of heuristics. The specific formulation of our GA is given in Appendix A. For the technical discussion refer to Haupt and Haupt (2004).

2.3.2 Model specification

We consider a set of $I = 6$ GA agents in the price-expectation feedback economy (2.1). GA agents use a general forecasting rule which requires exact parameter specification, and each agent is endowed with $H = 20$ such specifications. In order to give our model robust *empirical micro-foundations*, we follow the estimations by HHST09, as well as the simple model discussed by Anufriev et al. (2013) and focus on the first order rule (FOR).⁴

To be specific, in order to predict price p_t agent $i \in \{1, \dots, I\}$ focuses on $H = 20$ linear prediction rules given by

$$(2.6) \quad p_{i,h,t}^e = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p_{i,t-1}^e + \beta_{i,h,t} (p_{t-1} - p_{t-2}),$$

where $p_{i,h,t}^e$ is the prediction of price p_t , formulated by agent i conditional on using the rule h at the beginning of period t , and $p_{i,t-1}^e$ is the prediction by agent i of the price p_{t-1} , which the agent submitted to the market in period $t-1$. Rule (2.6) is a simplified case of the general FOR (2.5) (with $\alpha_3 = 0$).⁵

⁴*Cf.* Hommes and Lux (2013), who use the anchor-and-adjustment rule $p_{i,t}^e = \alpha + \beta(p_{t-1} - \alpha)$.

⁵We experimented with the full FOR with the anchor ($\alpha_3 > 0$), namely with an anchor equal to (1) the fundamental price $p^f = 60$; and to (2) the average realized price so far. The two specifica-

Heuristic FOR (2.6) depends on two parameters only, namely on $\alpha_{i,h,t}$ (price weight) and $\beta_{i,h,t}$ (trend extrapolation coefficient). We emphasize that these parameters are time dependent, because the agents want to fine-tune the FOR (2.6) for their specific market. For example, in an asset pricing market it may pay off to focus on the trend of the asset price. The agents would like to find the optimal degree of trend following, by experimenting with different trend extrapolation coefficients $\beta_{i,h,t}$. This learning is embodied as a GA algorithm and constitutes the novel insight of our model, compared to HSM or any homogenous expectations model.

Define $\mathbf{H}_{i,t}$ as the set of heuristics of agent i at time t . Each agent has $H = 20$ heuristics which are specified as a pair of parameters $(\alpha_{i,h,t}, \beta_{i,h,t}) \in \mathbf{H}_{i,t}$. Each such pair is represented as a chromosome, a binary string of length 40, 20 bits per coefficient. This means that the coefficients have to be bound to a finite interval. Price weight simply spans a simplex $\alpha_{i,h,t} \in [0, 1]$. For the trend extrapolation coefficient $\beta_{i,h,t}$, we report two specifications, namely with $\beta \in [-1.1, 1.1]$ (contrarian rules allowed) and $\beta \in [0, 1.1]$ (contrarian rules not allowed).⁶

The chromosomes are updated *independently* for each agent by GA evolutionary operators. We focus on the same set of parameters as Hommes and Lux (2013), see Table 2.1. The updating of the heuristics is based on their relative forecasting performance, specifically on mean squared error (MSE). Let

$$(2.7) \quad MSE_{i,h,t} = (p_{h,i,t}^e - p_t)^2.$$

Define the normalized performance (or fitness) measure as:

$$(2.8) \quad \Pi_{i,h,t} = \frac{\exp(-MSE_{i,h,t})}{\sum_{j=1}^H \exp(-MSE_{i,j,t})},$$

which is a logit transformation of MSE. The normalized performance measure (2.8) can be directly interpreted as the probability attached to each heuristic h by agent i at time t .⁷

tions generated comparable results and neither could replicate the dynamics of the positive feedback. See Appendix C for a discussion of the latter specification. This is consistent with the fact that in the estimated rules of HHST09 under positive feedback, the anchor weight α_3 is typically insignificant. Finally, the FOR (2.6) is a combination of the two heuristics (adaptive expectations and trend extrapolation) used by Anufriev et al. (2013), who also did not use an anchor for their HSM.

⁶Experimentation has led us to take 1.1 as the upper bound of β , see Appendix C for a discussion of the role of trend extrapolation under positive feedback. As for the lower bound, experimental data suggests that subjects tend to avoid contrarian strategies (HHST09 report only two subjects with such rules), thus for the sake of completeness we report both specifications.

⁷Notice that (2.8) it is also independent between the agents, because they have different sets of heuristics. Measure (2.8) is different from the experimental payoff, which was used by Hommes and

Parameter	Notation	Value
Number of agents	I	6
Number of heuristics per agent	H	20
Number of parameters	N	2
Allowed α price weight	$[\alpha_L, \alpha_H]$	$[0, 1]$
Allowed β trend extrapolation		
Specification 1	$[\beta_L, \beta_H]$	$[-1.1, 1.1]$
Specification 2	$[\beta_L, \beta_H]$	$[0, 1.1]$
Number of bits per parameter	$\{L_1, L_2\}$	$\{20, 20\}$
Mutation rate	δ_m	0.01
Crossover rate	δ_c	0.9
Lower crossover cutoff point	C_L	20
Higher crossover cutoff point	C_H	-1 (none)
Performance measure	$U(\cdot)$	$\exp(-MSE(\cdot))$

Table 2.1: Parameter specification used by the Genetic Algorithms agents.

The timing of the model is as follows. Before the market starts to operate, the agents' heuristics are initialized at random from a 'uniform' distribution: agent i samples 800 initial bits (twenty initial heuristics with two parameters, each encoded by twenty bits) independently as 0 or 1 with equal probability. In some initial periods the agents cannot use their heuristics, as these require past prices and predictions. Here the agents sample random predictions from a predefined distribution which we take as exogenous (for instance this can mean the experimental initial predictions). Once the agents have enough observations to use their heuristics, the timing of the market at period t is as follows:

1. The market price p_t is realized according to (2.1) and agents observe it;
2. Agents *independently* update their heuristics using one GA iteration, where the GA criterion function is $\Pi_{i,h,t}$ (forecasting performance). To be specific, agent i uses four evolutionary operators:
 - (a) *procreation*: agent samples H child heuristics from $\mathbf{H}_{i,t}$ with $\Pi_{i,h,t}$ as the corresponding probabilities;
 - (b) *mutation*: each bit of each child heuristic has probability δ_m to switch its value;

Lux (2013) in their GA model. We decided to use logit transformation of MSE to have a clear link with HSM literature and to keep this model feature independent from the experimental design.

- (c) *crossover*: each pair of child heuristics has probability δ_c to swap last twenty bits, which corresponds to exchanging β 's.
 - (d) *election*: each child heuristic is compared with the parent heuristic in terms of $\Pi_{i,h,t}$: a child heuristic becomes part of the set $\mathbf{H}_{i,t+1}$ if it outperforms its parent, else its parent is passed to $\mathbf{H}_{i,t+1}$.
3. With the new $\mathbf{H}_{i,t+1}$, period $t + 1$ starts.
 4. Each agent i picks one particular heuristic $i, h, t + 1$, which is based on the hypothetical MSE of heuristics $\mathbf{H}_{i,t+1}$ in predicting the *last observed price* p_t (with probabilities $\Pi_{i,h,t}$). Agent i uses the chosen heuristic to generate her prediction $p_{i,t+1}^e$;
 5. New period $t + 1$ starts: the market price p_{t+1} is realized according to (2.1).

In the first period when the heuristics can be used, their hypothetical past performance is still undefined, and so the agents pick one with equal probabilities. For the HHST09 experiment, GA agents start to use the first-order rule in the second period (one at random) and start to update their heuristics in the third period (for a more detailed discussion of the initialization, see Appendix B).

The last step — the heuristic choice — is the same as in the HSM, but there are two important differences between HSM and our GA model. First, the heuristics evolve over time with $\mathbf{H}_{i,t} \neq \mathbf{H}_{i,t+1}$. As a result, we obtain a HSM in which 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 $\mathbf{H}_{i,t+1} \neq \mathbf{H}_{j,t+1}$, which gives us an agent-based counterpart for HSM.

2.3.3 50-period ahead simulations

The first test for the fit of our model to the experimental data are 50-period ahead simulations for the HHST09 experiment.⁸ We take the feedback equations (2.3) and (2.4) and simulate our model for 50 periods, without *any information* from the experiment after period 1, and hence compare the realized long-run model dynamics with the experimental data.⁹

⁸All simulations were written in Ox matrix algebra language (Doornik, 2007) and are available upon request.

⁹In one of the positive feedback treatment groups, one of the subjects ‘out of the blue’ predicted ten times higher price than both his previous forecast and the realized market price. This destabilized the whole market for a number of periods. In the following analysis, we follow Anufriev et al. (2013)

The model requires exogenous predictions for the first period. This is important, since in the experiment the average initial prediction affected the group dynamics under the positive feedback treatment (*cf.* Anufriev et al., 2013). In the first Monte Carlo (MC) exercise, we sample initial prediction from a distribution calibrated by Diks and Makarewicz (2013), in order to obtain a general picture of the model dynamics. We resample the model 1'000 times, including new initial predictions and realizations of the learning algorithm, to obtain a satisfactory MC distribution. The median of 1'000 GA simulations, with 95% confidence intervals (CI), for the model with contrarian rules $\beta \in [-1.1, 1.1]$ are shown in Figure 2.2. See also Figure 2.1 for sample experimental prices and predictions and realized 50 period ahead GA simulation of a representative group for each feedback treatment.

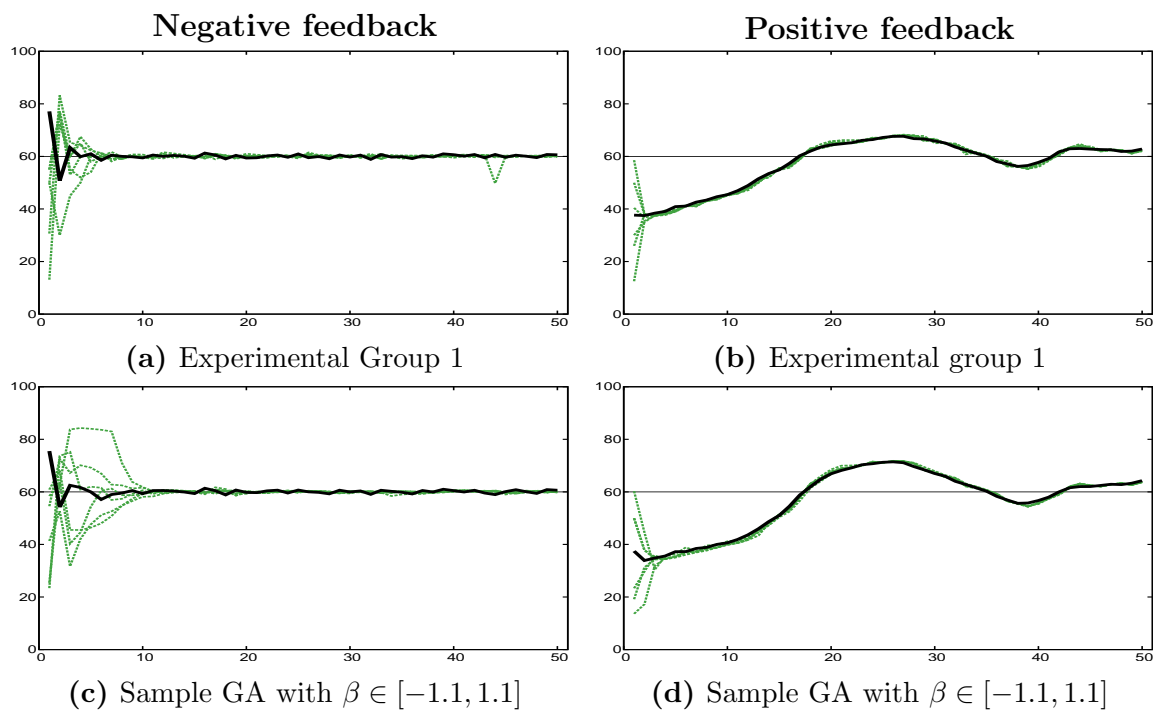


Figure 2.1: HHST09: experimental groups and sample 50-period ahead simulations of the GA model (with $\beta \in [-1.1, 1.1]$ and random initial predictions). Black line represents the price and green dashed lines are the individual predictions.

Figure 2.2 shows the MC simulations of the realized prices (top panel) and the degree of coordination, that is the standard deviation of six individual forecasts (bottom panel). 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 is visible until period 15, consistent with the experimental and omit this group and hence focus on six positive feedback and six negative feedback treatment groups.

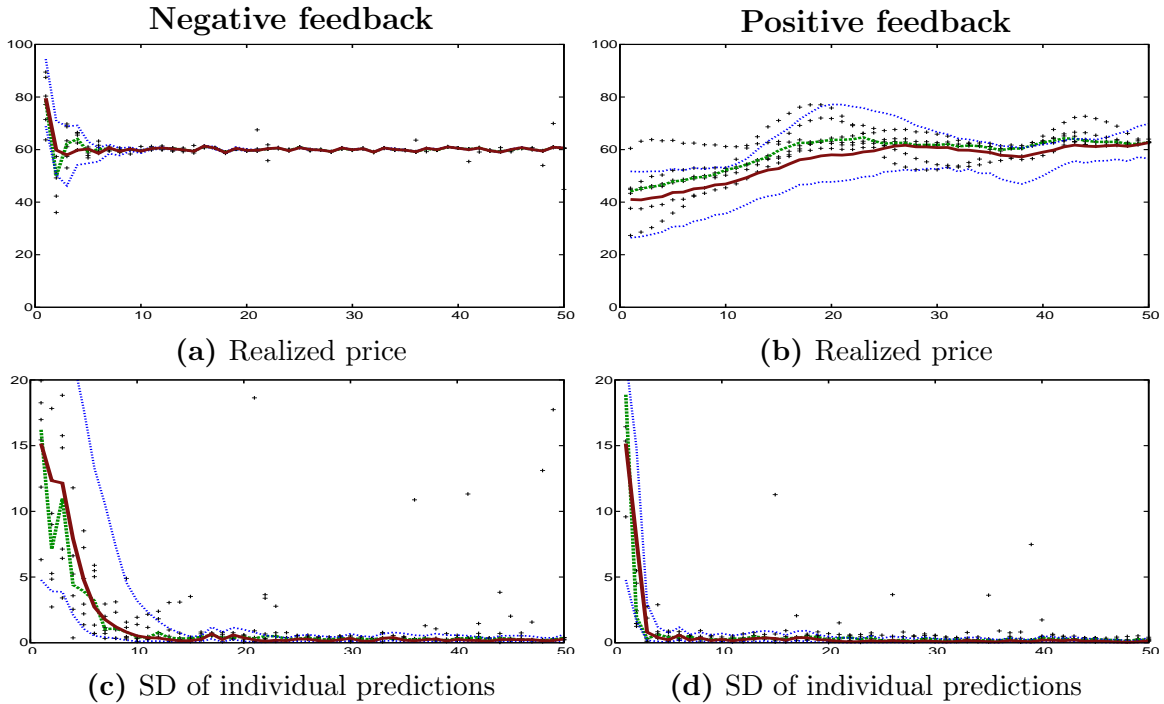


Figure 2.2: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.1, 1.1]$. Realized price (top) and degree of coordination (standard deviation of individual predictions; bottom) over time. Green dashed line and black pluses represent the experimental median and group observations; red line is the median and blue dotted lines are the 95% confidence interval for the GA model. Left panel displays the negative feedback, right panel the positive feedback.

data. Under positive feedback, GA agents coordinate their forecasts in less than five periods, but the distribution of realized prices does not collapse into the fundamental even after 50 periods, when the 95% CI of prices is as wide as $[55, 75]$. The median price resembles the experimental oscillations, including the typical amplitude and turning points. Overall, the 95% CI for our GA model captures 65% (81%) of the experimental prices and 81.33% (71.67%) of the degree of coordination for the negative (positive) feedback treatment. This means that we are able to evaluate roughly 75% of the long-run (50-period ahead) behavior of the experimental groups, both at the *aggregate* level (prices) and the *individual* level (coordination of individual predictions).

Which heuristics were learned by our GA agents? Figure 2.3 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

$$(2.9) \quad p_{i,t+1}^e \approx 0.9p_t + 0.1p_{i,t}^e + 0.6(p_t - p_{t-1}).$$

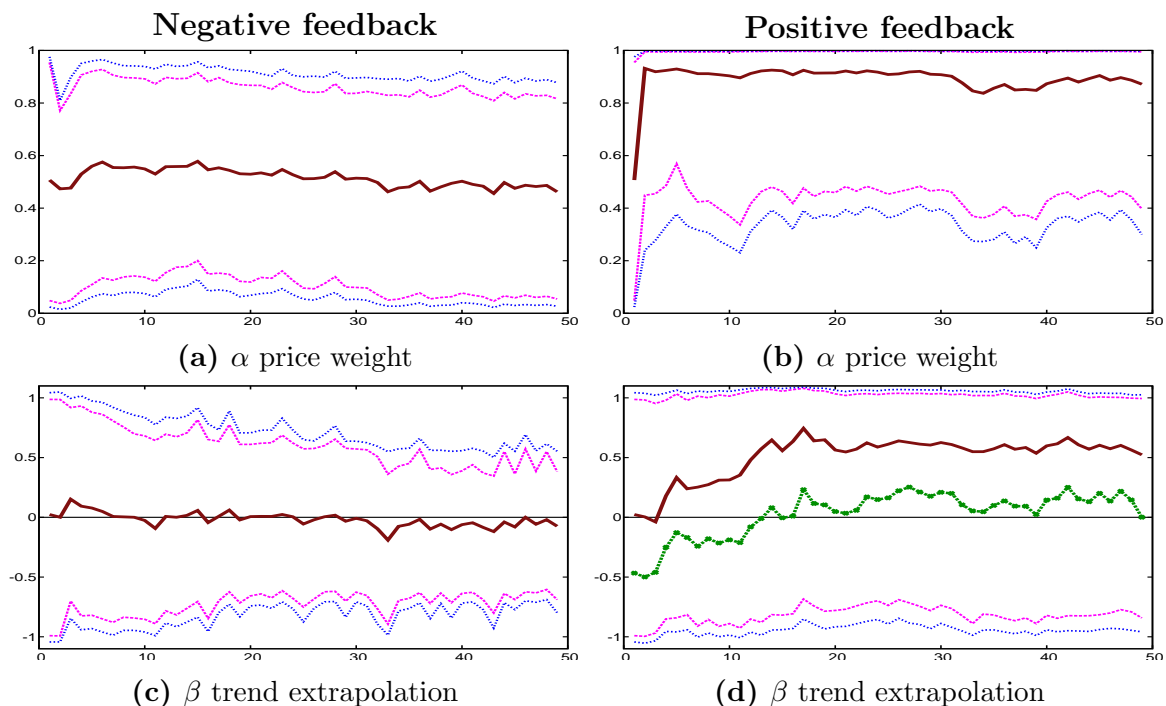


Figure 2.3: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.1, 1.1]$. The price weight α and the trend extrapolation β chosen by the agents over time. Red line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model. Left panel displays the negative feedback, right the positive feedback.

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 the coefficient $\beta = 1$ that Anufriev et al. (2013) are using in their 2-type HSM. Furthermore, 72.35% of the GA agents would never use a negative β in the last 30 periods (see the green star-line in Figure 2.3d for 27.65% percentile); and in the last period, the chosen β has a negatively skewed distribution (see Figure 2.10a).

On the other hand, under negative feedback the median GA agent learns a rule close to

$$(2.10) \quad p_{i,t+1}^e \approx 0.5p_t + 0.5p_{i,t}^e$$

(with median β trend coefficient close to 0). This median rule for negative feedback is adaptive expectations with equal coefficient 0.5; Anufriev et al. (2013) are using adaptive expectations with coefficient 0.75 on price in their 2-type HSM. Our learning dynamics therefore confirm the results by HHST09 and Anufriev et al. (2013), albeit with slightly different parametrization.

In the second MC study, we focus on how well our GA model can replicate long-

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	19.64	34.02	55.15	63.98
GA: $\beta \in [-1.1, 1.1]$	2.884	20.03	44.22	51.98
GA: $\beta \in [0, 1.1]$	9.392	29.51	25.3	31.1

Table 2.2: HHST09: 50-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over six negative feedback and six positive feedback groups.

run dynamics of a *specific* experimental group. We take initial predictions of each group and use them as initialization for 50-period simulations of the GA model. We investigate the realized prices and individual forecasts. Following Anufriev et al. (2013) we define the GA model expected individual price forecast as

$$(2.11) \quad \hat{p}_{i,t}^{e,GA} = \sum_{h=1}^{H=20} \Pi_{i,h,t}^t p_{i,h,t}^{e,t}.$$

For each sample model time path, we compute its mean squared error (MSE) in predicting the experimental data (both prices and individual price forecasts) for the last 47 periods (excluding the initialization phase)

$$(2.12) \quad \begin{aligned} MSE_X^{\text{prices}} &= \frac{1}{47} \sum_{t=4}^{50} (p_t^{Gr X} - p_t^{GA})^2, \\ MSE_X^{\text{price forecasts}} &= \frac{1}{6 \times 47} \sum_{i=1}^6 \sum_{t=4}^{50} (p_{i,t}^{e,Gr X} - \hat{p}_{i,t}^{e,GA})^2, \end{aligned}$$

where $p_t^{Gr X}$ and $p_{i,t}^{e,Gr X}$ denote the realized price and the price forecast of subject i at period t in an experimental group X and p_t^{GA} and $\hat{p}_{i,t}^{e,GA}$ are the price and the price forecast of agent i at period t predicted by the GA model for the group X .

Table 2.2 reports MSE averaged over the six groups for each treatment, with 1'024 sample GA model paths per experimental group. We also include results for a number

of benchmark models, including simple homogenous expectation rules, RE and HSM (two heuristics specification by Anufriev et al. (2013)).¹⁰ In terms of the long-run, 50 period ahead, predictions, RE and two simple models, adaptive and contrarian expectations, perform the best under negative feedback, as they correctly predict the agents to converge to the fundamental price. Our GA model performs only slightly worse. Under positive feedback, RE, contrarian and adaptive expectations still predict convergence, in contrast to the experimental oscillations. HSM, trend extrapolation and naive expectations perform comparatively well, but surprisingly they are not better than RE. The reason is that the price oscillations predicted by these three models at the longer time horizon fall out of phase with the experimental oscillations. The best fit is achieved by our GA model, especially the one without contrarian rules $\beta \in [0, 1.1]$. We conclude that all benchmark models are able to capture the long-run dynamics of possibly one feedback treatment, but not of two treatments at the same time. Only our GA model successfully evaluates both treatments.

2.3.4 One-period ahead predictions

A good indicator of the model's fit is the precision of its one-period ahead predictions (Anufriev et al., 2013): how well the model predicts experimental outcomes in period $t + 1$, conditional on the data until period t , in terms of MSE. For deterministic models such as HSM and the homogeneous expectations benchmark 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 period-to-period distribution. To address this issue, we compute the *expected* MSE using the Sequential Monte Carlo (SMC) approach.¹¹

Our SMC is designed in the following way. For each experimental group X , we run simultaneously 1'024 independent GA model simulations. We associate one GA agent with one subject, and in each period $t \geq 2$ every GA agent i (1) retains her 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 use the experimental data to update their heuristics and forecast the price p_t in the usual way, which gives us the GA's price forecasts (2.11) and realized prices (2.1) for period t . We evaluate the fit of

¹⁰For the definition of the benchmark rules, please refer to Appendix D.

¹¹We checked the robustness of our results by Sequential Importance Sampling (with Resampling) technique called Auxiliary Particle Filter (Doucet et al., 2000), see Appendix E, and obtained comparable results. One can also show that SMC is a restricted version of APF, and so can be used if the results are similar.

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	2.999	17.106	0.889	4.156
GA: $\beta \in [-1.1, 1.1]$	4.95	25.017	<i>0.806</i>	4.235
GA: $\beta \in [0, 1.1]$	4.496	25.012	0.802	4.198

Table 2.3: HHST09: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over six negative feedback and six positive feedback groups.

the model to the experimental group by computing the average MSE (2.12) over 1'024 GA simulations.

The results are similar to the 50-period ahead simulations, see Table 2.3. Under negative feedback, RE, HSM adaptive, contrarian and naive expectations all capture the convergence of prices and forecasts to the fundamental price, slightly outperforming our GA model. Under positive feedback, all these models (with the exception of HSM) loose their predictive power and under-estimate the experimental oscillatory behavior of individual forecasts. The GA model has the best fit for the positive feedback treatment and outperforms RE by a factor of 10.

These MC simulations show that our model captures both the aggregate and individual behavior in the LtF experiment reported by HHST09, both in terms of short and long-run dynamics. Moreover our model is the only one that captures the observed degree of heterogeneity of individual behavior between the experimental subjects, as measured by the coordination of the contemporary price forecasts (Figure 2.2c and 2.2c), together with the persistent heterogeneity of the forecasting heuristics.

2.4 Evidence from other experiments

Our GA model fits the HHST09 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:

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

2.4.1 Large shocks to the fundamental price

BHST12 report a LtF experiment with the same structure as HHST09: positive and negative feedback with the linear structure given by (2.2) and the same dampening factor $|B| = \frac{20}{21}$. In this experiment however there are two large and unanticipated shocks to the fundamental price A : it changes from $p^f = 56$ to $p^f = 41$ in period $t = 21$ and then to $p^f = 62$ in period $t = 44$ until the last period $t = 65$.

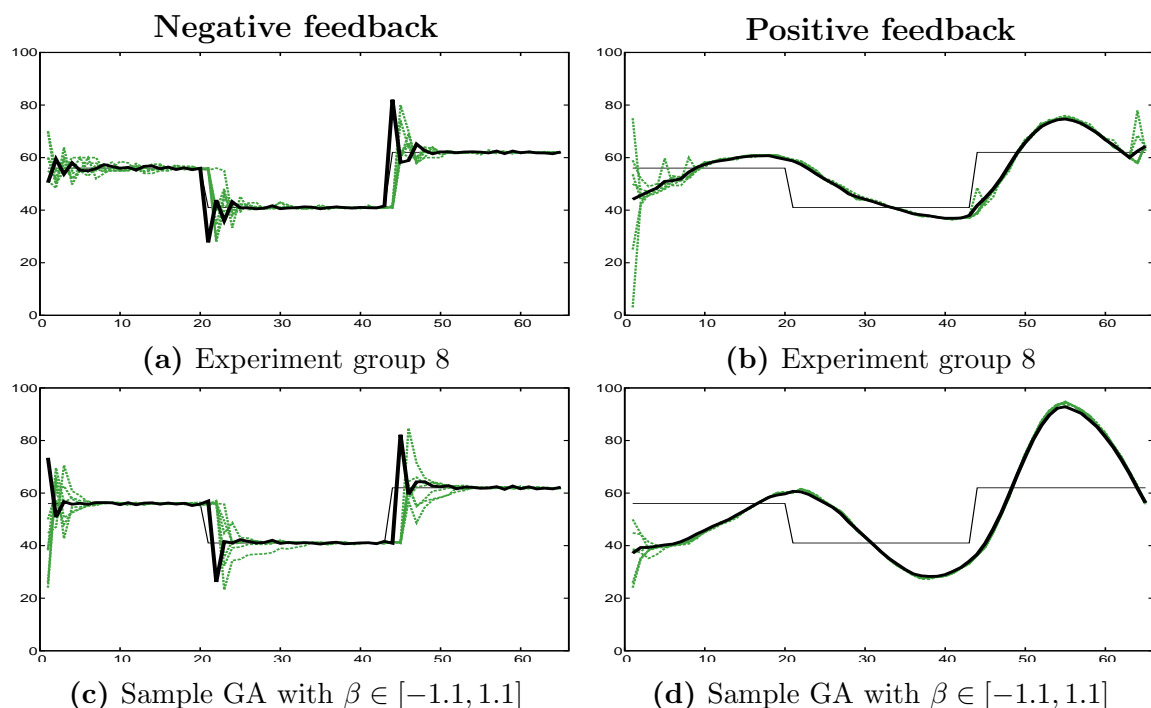


Figure 2.4: BHST12: experimental groups and sample 65-period ahead simulations of the GA model (with $\beta \in [-1.1, 1.1]$ and random initial predictions). Black line represents the price and green dashed lines are the individual predictions.

The results of this experiment are similar to HHST09 and typical time paths are shown in Figure 2.4. Under negative feedback (Figure 2.4a), a shock to the fundamental breaks the subjects' coordination and is followed by a quick convergence to the

new fundamental price. Under positive feedback (Figure 2.4b), shocks do not break the coordination, and the predictions and prices fluctuate smoothly towards the new fundamental, eventually over- or undershooting it.

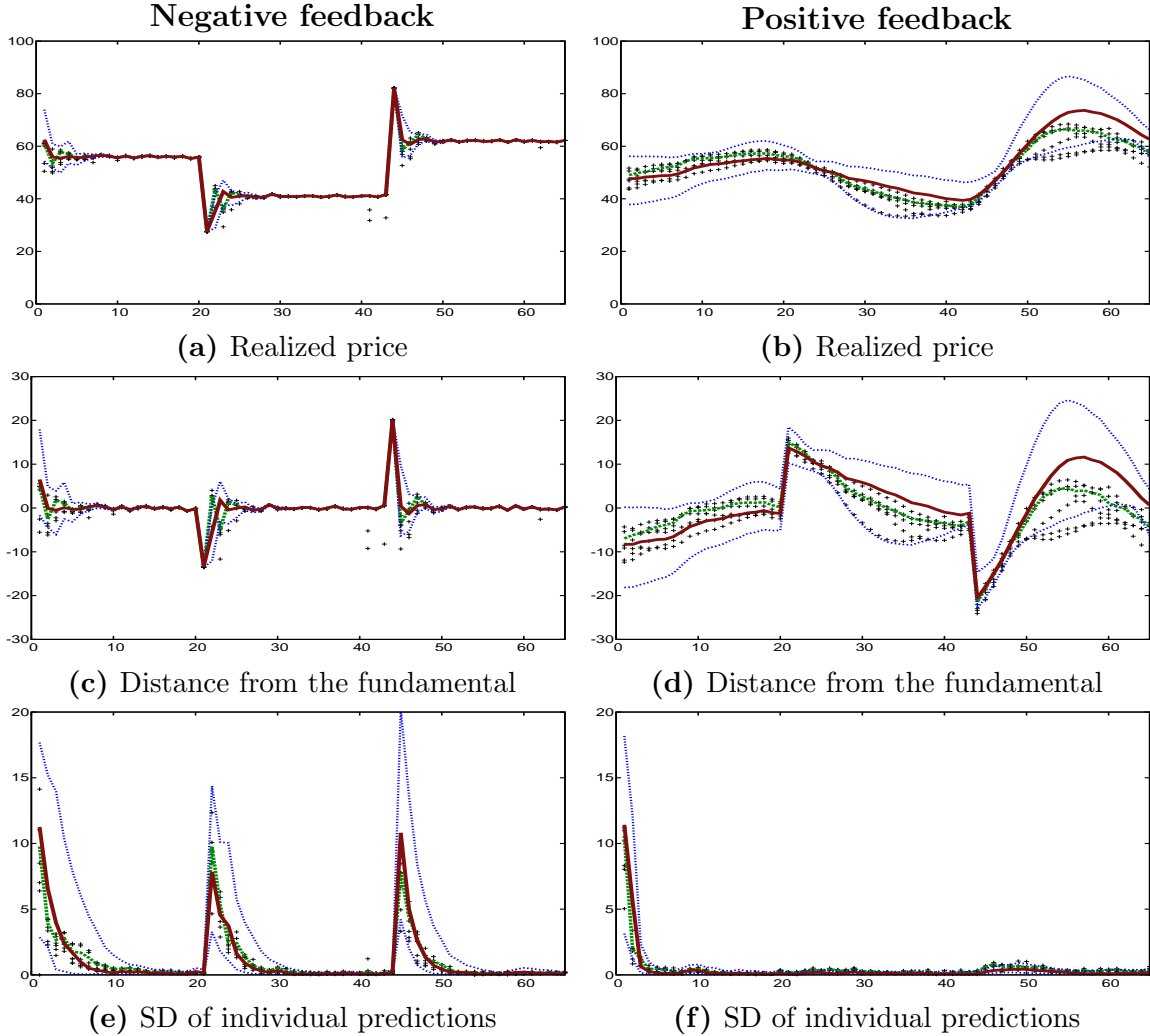


Figure 2.5: BHST12: 65-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.1, 1.1]$. Realized price (top) and degree of coordination (standard deviation of individual predictions; bottom) over time. Green dashed line and black pluses represent the experimental median and group observations; red line is the median and blue dotted lines are the 95% confidence interval for the GA model. Left panel displays the negative feedback, right panel the positive feedback.

Figure 2.5 shows 65-period ahead MC simulations of prices, individual price forecasts and the degree of coordination.¹² Our model replicates well the experimental price dynamics for both treatments, as well as the impact of the shocks to the fun-

¹²We estimate the distribution of the initial predictions as in Diks and Makarewicz (2013) and sample directly from it, see Appendix B.

damental price on individual coordination. For the $\beta \in [-1.1, 1.1]$ specification, the 95% CI of our GA model contain 65.58% (84.23%) of the experimental prices and 84.04% (66.73%) of the standard deviation of individual forecasts under negative (positive) feedback. Overall, we can replicate around 75% of the experimental data with 65-period ahead simulations.




Figure 2.6: BHST12: 65-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.1, 1.1]$. The price weight α and the trend extrapolation β chosen by the agents over time. Red line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model. Left panel displays the negative feedback, right the positive feedback.

Figure 2.6 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 behavior is similar to that from the experiment without large shocks HHST09. Under the negative feedback, the median GA agent learns the same adaptive expectations rule $p_{i,t+1}^e \approx 0.5p_t + 0.5p_{i,t}^e$. Under the positive feedback, the median GA agent converges to a heuristic

$$(2.13) \quad p_{i,t+1}^e \approx 0.95p_t + 0.05p_{i,t}^e + 0.9(p_t - p_{t-1}),$$

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 (2.9) used by the median GA agent under the positive feedback from the experiment without large

shocks HHST09. The 95% CI for the trend extrapolation β becomes significantly positive towards the end of the experiment (see also Figure 2.10b 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.

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	2736	1289	101.3	113.3
Adaptive	3.629	10.75	55	62.14
Contrarian	6.984	<i>14.45</i>	58.46	65.95
Naive	94.44	110.9	46.62	52.9
RE	13.871	20.923	55.133	60.859
HSM	73.57	87.86	90.8	101.8
GA: $\beta \in [-1.1, 1.1]$	8.01	21.97	<i>43.49</i>	<i>49.44</i>
GA: $\beta \in [0, 1.1]$	<i>6.333</i>	17.39	43.49	49.64

Table 2.4: BHST12: 65-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over eight negative feedback and eight positive feedback groups.

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	114.061	121.329	1.183	2.165
Adaptive	3.689	10.332	3.776	4.618
Contrarian	5.92	<i>12.534</i>	4.737	5.559
Naive	9.979	16.81	2.411	3.286
RE	13.871	20.923	55.133	60.859
HSM	38.309	45.679	0.9996	2.024
Genetic Algorithm model				
$\beta \in [-1.1, 1.1]$	10.247	21.464	<i>0.342</i>	2.059
$\beta \in [0, 1.1]$	4.208	15.267	0.341	<i>2.036</i>

Table 2.5: BHST12: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over eight negative feedback and eight positive feedback groups.

Table 2.4 reports the MSE for the 65-period ahead simulations initialized 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 for the 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 shocks to the fundamental under the negative feedback treatment. Also, HSM seems below average. In terms of long-run forecasting, our GA model is again second best for the negative feedback and the best for positive feedback.

We also use the SMC approach to compute the GA model's one-period ahead predicting power, reported in Table 2.5. The results are consistent with the 65-period ahead simulations. For both treatments, the GA model (especially without contrarian rules, $\beta \in [0, 1.1]$) is the best among all reported models.

2.4.2 Cobweb economy

V01 and HSTV07 report an LtF experiment in a Cobweb economy setting. 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 parametrization under the assumption of homogenous naive expectations. This is a follow-up on V01 who investigates the strongly unstable treatment with 12 subjects. The experiment resulted in average price equal to the RE fundamental price. However, the realized prices were excessively volatile, but — in contrast to positive feedback experiments — also non-persistent (with weak autocorrelation structure). Hommes and Lux (2013) study this experimental data set with the GA model based on an anchor-and-adjustment (AR1) forecasting rule. It therefore constitutes an important benchmark case for our GA model.

As a first test for our model, we conduct a MC exercise in the vein of Hommes and Lux (2013). For each treatment, we compute six 50-period ahead simulations with different random numbers.¹³ Next we compute the mean and standard deviation of the realized prices and the individual price forecasts. We repeat this procedure 1'000 times and thus obtain a *distribution* (including 95% CI) of the realized means and variances of prices and price forecasts. We report the results in Table 2.6 for the two GA model specifications.

Our 50-period ahead simulations explain well the experimental data and perform significantly better than RE. The 95% CI of our GA model with $\beta \in [-1.1, 1.1]$ and

¹³We estimate the distribution of the initial predictions as in Diks and Makarewicz (2013), see Appendix B.

	Mean(p)	Var(p)	Mean(p ^e)	Var(p ^e)
Stable				
Experiments	5.64*†	0.36*†	5.56*†	0.087*
GA: AR1	5.565¶	0.326¶	5.576¶	0.1
GA: FOR $\beta \in [-1.1, 1.1]$	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: FOR $\beta \in [0, 1.1]$	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†	0.63*†	5.67*†	0.101*†
GA: AR1	5.817	0.647	5.645¶	0.16¶
GA: FOR $\beta \in [-1.1, 1.1]$	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: FOR $\beta \in [0, 1.1]$	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†	2.62*	5.73	0.429*
GA: AR1	6.2¶	2.161	5.434	0.769
GA: FOR $\beta \in [-1.1, 1.1]$	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: FOR $\beta \in [0, 1.1]$	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, large group				
Experiments	5.937†	1.783*	5.781*†	0.204*†
GA: AR1	6.183¶	1.571	5.515¶	0.5¶
GA: FOR $\beta \in [-1.1, 1.1]$	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: FOR $\beta \in [0, 1.1]$	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 2.6: HSTV07: 50-period ahead MC results for GA simulations for four treatments, stable, unstable and strongly unstable with 6 or 12 subjects. Average price and prediction, and their variances. Mean experimental statistics; GA simulations with AR1 prediction rule for mutation rate $\delta_m = 0.01$ (mean statistics); GA simulations with FOR with or without contrarian rules (median statistics with 95% confidence intervals). * and † denote experimental statistic which falls into 95% CI of GA FOR with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$ respectively. ¶ denotes Hommes and Lux (2013) statistics which fall outside the 95% CI for GA model with $\beta \in [0, 1.1]$ when these CI contain the experimental statistics.

$\beta \in [0, 1.1]$ replicate 12 and 11 out of 16 experimental statistics respectively. Among the 11 cases successful for the GA model based on FOR rule (2.6) with $\beta \in [0, 1.1]$, 9 statistics reported by Hommes and Lux (2013) are outside 95% CI of our model. That means that we can replicate around three quarters of experimental descriptive statistics,

Treatments	Stable		Unstable		Strongly unstable	
	Prices	Forecasts	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	<i>0.048</i>	0.248	0.364	0.385	2.257	1.844
HSM	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: $\beta \in [-1.1, 1.1]$	0.088	0.356	<i>0.346</i>	0.631	3.445	3.261
GA: $\beta \in [0, 1.1]$	0.043	<i>0.275</i>	0.223	<i>0.449</i>	<i>2.376</i>	<i>2.114</i>

Table 2.7: HSTV07: 50-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (FOR with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over six groups for each treatment (stable, unstable, strongly unstable).

Treatments	Stable		Unstable		Strongly unstable	
	Prices	Forecasts	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	<i>0.497</i>	<i>2.929</i>	2.729
Naive	0.196	0.448	0.577	0.788	3.095	3.731
RE	0.048	0.248	<i>0.364</i>	0.385	2.257	1.844
HSM	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: $\beta \in [-1.1, 1.1]$	0.13	0.393	0.866	0.795	5.547	3.25
GA: $\beta \in [0, 1.1]$	<i>0.07</i>	<i>0.31</i>	0.25	0.531	3.079	<i>2.358</i>

Table 2.8: HSTV07: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over six groups for each treatment (stable, unstable, strongly unstable).

most of which with a significantly higher precision than the GA model specification in Hommes and Lux (2013).¹⁴

We also check the 50-period ahead dynamics of the model conditional on the initial predictions from particular groups from HSTV07, see Table 2.7. 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 linear experiments) resemble a white noise around the fundamental price. As a result, predicting the mean of these close-to-chaotic dynamics (as RE do) is better than trying to capture them with simplistic models. Only our GA model, in particular the one with $\beta \in [0, 1.1]$, keeps up with RE, and performs better than Hommes and Lux (2013) GA specification based on an AR1 rule.

The next exercise is the one-period ahead forecasting of the model with SMC approach for the 18 groups from HSTV07. Table 2.8 gives the summary results. It is apparent that the less stable the treatment, the worse fit has any model. 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.¹⁵ Our specification again prevails over the AR1 GA model of Hommes and Lux (2013).

We conclude that the cobweb experiments result in unstable, non-persistent prices, and simpler models like homogenous heuristics, but also HSM, miss-identify here any structure. As a result, their point predictions are so poor that it is better to predict the mean price, as in RE. Only our GA model (with $\beta \in [0, 1.1]$) comes close to RE in terms of this task. It furthermore allows to explain the volatility of the experimental markets, which RE cannot account for. Finally, it is clear that the use of experimental micro-foundations has an advantageous effect: our GA model has a better fit to the data than the AR1 specification used by Hommes and Lux (2013).

2.4.3 Two-period ahead asset pricing

HSTV05 report an experiment based on a non-linear positive feedback market; an asset-pricing model with a robotic fundamental trader, in which the current price depends on the subjects' expectations about the price in the *next* period: $p_t = F(p_{1,t+1}^e, \dots, p_{6,t+1}^e)$.

¹⁴The GA model simulations are also closer to the experimental data in terms of the autocorrelation of the prices. RE always 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 2.16 in Appendix 2.F for the results.

¹⁵Notice that the scale of the prices in this experiment is $[0, 10]$ in contrast with the two previous settings, where the prices belonged to $[0, 100]$ intervals. The highest possible MSE in the linear experiments is 100 times higher than in the cobweb experiment.

There were two treatments with the difference in the fundamental price: seven markets were based on $p^f = 60$ and three on $p^f = 40$. Subjects coordinated both on stable outcomes and diversified oscillations.

Notice that in this experiment the subject's decisions are based on a different information set than in the previous one-period ahead experiments. Upfront it is difficult to predict how this will influence subject behavior, specifically whether they will use more complicated strategies or extrapolate the trend to a different degree. After some experimentation, we decided that the two period ahead version of our GA model should be based on the following specification. Define the prediction of price from period $t + 1$ by the GA agent i based on her rule h as

$$(2.14) \quad p_{i,h,t+1}^e = \alpha_{i,h} p_{t-1} + (1 - \alpha_{i,h}) p_{i,t-1}^e + \beta_{i,h} (p_{t-1} - p_{t-2}).$$

Once p_t is realized, the agents can evaluate their rules based on the hypothetical performance of predicting p_t *two periods ago*. GA agents focus on $(p_t - p_{i,h,i,t}^e)^2$, where $p_{i,h,i,t}^e$ is function of $p_{i,t-2}^e$, p_{t-2} and p_{t-3} .

This specification is the most straightforward translation of the baseline one-period ahead forecasting heuristic (2.6). Again, there is no evidence that we need an anchor (see Appendix C). In the baseline simulations, we look at the allowed trend specified as before (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). HSTV05 report that many of their subjects use very strong trend extrapolation, thus for the sake of completeness we will also report the results of our model with $\beta \in [-1.3, 1.3]$ and $\beta \in [0, 1.3]$.




Figure 2.7: HSTV05: sample 50-period ahead simulations for GA model with $\beta \in [-1.3, 1.3]$ with different initial predictions and learning. The green lines are individual predictions, the black line is the realized price and the purple dashed line is the fundamental price.

In the seven treatment groups with the fundamental price $p^f = 60$, HSTV05 observe groups which have converged to this fundamental, as well as groups with oscillations of different amplitude and frequency. Figure 2.7 displays three typical simulated markets




Figure 2.8: HSTV05: sample 2'000-period ahead simulation (b) and its first 500 periods (a) of the GA model with $\beta \in [-1.3, 1.3]$ with fundamental price $p^f = 60$ and random initial predictions. The green lines are individual predictions, the black line is the realized price and the purple dashed line is the fundamental price.

MSE	Prices	Forecasts
Trend extr.	178.2	174.9
Adaptive	96.12	145.9
Contrarian	157	146.8
Naive	95.29	144.6
RE	96.0328	145.998
GA: $\beta \in [-1.1, 1.1]$	103.9	155.8
GA: $\beta \in [0, 1.1]$	114.9	169.1
GA: $\beta \in [-1.3, 1.3]$	139.4	201.5
GA: $\beta \in [0, 1.3]$	226.5	318.5

Table 2.9: HSTV05: 50-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over all experimental groups.

of the GA model (with $\beta \in [-1.3, 1.3]$) for the HSTV05 economy (with $p^f = 60$).¹⁶ The GA agents can both converge to the fundamental price (Figure 2.7a) as well as coordinate on unruly oscillations (Figure 2.7c). Furthermore, sometimes both outcomes are present at the same time. Figure 2.7b shows a sample simulation, in which the price seemingly stabilizes at the fundamental value between periods 18 and 20, but then resumes to oscillate mildly.

To further stress the volatile behavior of this market structure, we report one long run simulation for the GA model with $\beta \in [-1.3, 1.3]$. Figure 2.8 displays its first 500

¹⁶See Appendix B for initialization.




Figure 2.9: HSTV05: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.1, 1.1]$ (left panel) and $\beta \in [-1.3, 1.3]$ (right panel). The price weight α and the trend extrapolation β chosen by the agents over time. Red line is the median, blue dotted lines are 95% CI, purple dashed are 90% CI for the GA model.

(Figure 2.8a) and 2'000 (Figure 2.8b) periods. Oscillations of different amplitude are persistent and can reappear even if the market settles on the fundamental price for some time, as seen in Figure 2.8b around period 800 or after period 1200. This means that in the system the fundamental price is not a unique attractor.

To explain this outcome, we take a closer look at the trend extrapolation chosen by the GA agents. Figure 2.9 shows results for MC 50-period ahead simulations for two GA model specifications, with $\beta \in [-1.1, 1.1]$ and $\beta \in [-1.3, 1.3]$. If the agents are allowed to experiment with higher β , the median price has a very similar oscillatory shape. The difference is seen in the 95% CI: both specifications are likely to generate two price

MSE	Prices	Forecasts
Trend extr.	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: $\beta \in [-1.1, 1.1]$	42.224	74.95
GA: $\beta \in [0, 1.1]$	5.934	30.341
GA: $\beta \in [-1.3, 1.3]$	21.192	53.238
GA: $\beta \in [0, 1.3]$	16.29	42.125

Table 2.10: HSTV05: one-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, 4-type Heuristic Switching Model (source: Anufriev and Hommes, 2012) and GA models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over all experimental groups.

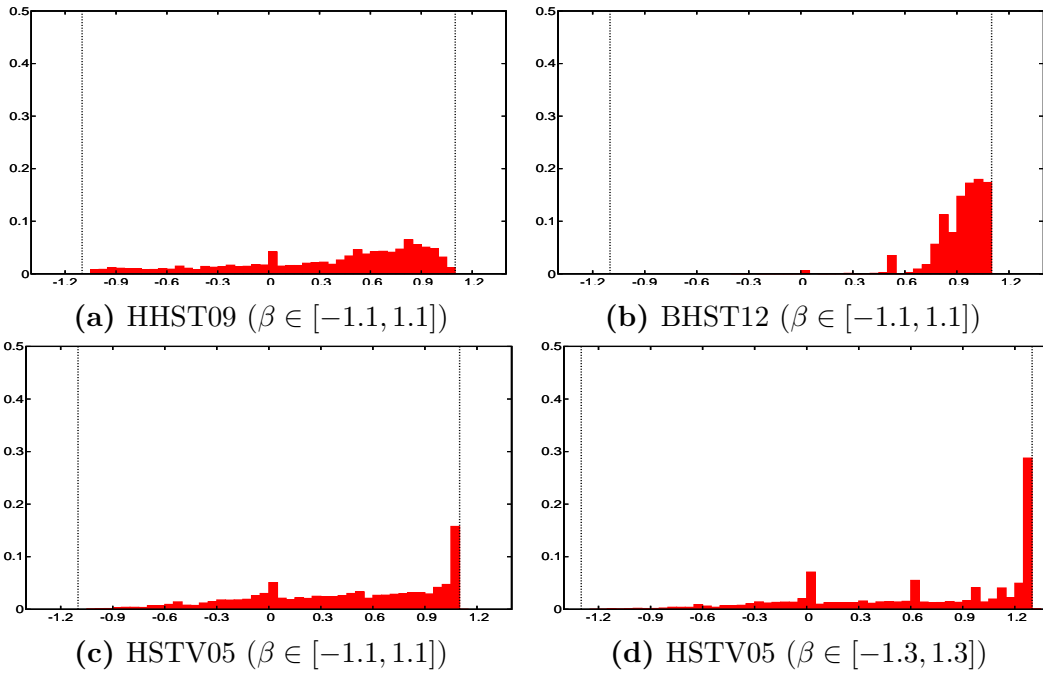


Figure 2.10: Positive feedback treatments: HHST09, BHST12 and HSTV05 with $p^f = 60$: 50-period ahead predictions. Distribution of trend extrapolation coefficient β chosen by the agents in the last period $t = 50$ across the whole MC sample for each treatment, and two β specifications for HSTV05.

bubbles within 50 periods, but the model with $\beta \in [-1.3, 1.3]$ has larger potential oscillations (Figure 2.9b), and the second bubble can be even bigger than the first one (unlike in the linear positive feedback). Regardless, the median GA agent converges

to a strong trend extrapolation rule, close to $p_{i,t+1}^e = p_{t-1} + (p_{t-1} - p_{t-2})$, which is consistent with the behavior 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 (Figures 2.10c and 2.10d) is bimodal, with a relatively large mass centered around zero (*i.e.* weak or no trend extrapolation).

We interpret this finding in the following way. If the price is sufficiently stable and close to the fundamental value, the robotic fundamental trader is powerful enough to mitigate additional price deviations. This discourages GA agents to extrapolate the insignificant trend, and so the price stability becomes self-reinforcing. However, if the trend in prices is sufficiently large, the stabilizing effect of the robotic trader can be counter-weighted by the GA agents coordinating on trend extrapolation. The non-linear and the two-period ahead price feedback amplifies the realized price oscillations (which become self-reinforcing), but also allows their specific shape to be diversified. For this reason we speculate that the two period ahead feedback entails two types of attractors in our model, which corresponds well to the diversified dynamics observed in the experiment.

We note that the model does not predict fast price oscillations, and more than two bubbles within 50 periods are rarely observed (in contrast to the experiment). This is independent from the allowed trend extrapolation and cannot be explained by adding an anchor to the forecasting rule (2.14) (see Appendix C). We speculate that one should experiment with higher order rules to replicate all the oscillations from HSTV05, but we leave this for future investigations.¹⁷

Even though our GA model leaves space for improvement, it is the only one which is comparatively good in predicting the experimental results of HSTV05 both in the long- and the short-run. Table 2.9 reports the MSE of 50-period ahead simulations initialized with the experimental initial predictions. These are comparatively 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 2.10 shows the MSE of one-period ahead predictions for our GA model and 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 $\beta \in [0, 1.1.]$) is able

¹⁷HSM does explain these faster oscillations with higher order AR2 rule (Anufriev and Hommes, 2012).

to compete with this richer HSM in terms of predicting experimental prices.

This outcome resembles the results for the cobweb economy Hommes et al. (2007). A natural interpretation is given by our GA model, which predicts that the dynamics of HSTV05 economy are unruly price oscillations. It follows that one can successfully predict the subjects forecasts and prices in the short run by a trend extrapolation heuristics. However, the oscillations can arbitrarily change shape, which together with the co-existence of two attractors renders any long-run forecasting virtually impossible. It also means that, despite potentially oscillatory dynamics, no single trend extrapolation model can replicate the long-run dynamics of this experiment. We leave it open for other research whether any structural model can cope with such an environment.

2.5 Conclusions

In this chapter we discuss a model in which agents independently use Genetic Algorithms to optimize a simple forecasting heuristic. We argue that our model is able to replicate many findings from different Learning-to-Forecast experiments, both at the aggregate and individual level.

In Learning-to-Forecast experiments, subjects are asked to forecast prices, while the realized price depends on their predictions. This mimics many well studied economic environments, such as asset pricing markets or cobweb economies. These experiments can be used as a controlled setting to study how the human subjects try to adapt to the price-predictions feedback. Their major insight is that the market converges to the rational expectations equilibrium only if the relationship between the average price expectation and the realized price is linear and negative (Heemeijer et al., 2009). In the case of markets in which this relationship is positive, subjects may coordinate on extrapolating observed price trends, which reinforces price oscillations (Hommes et al., 2005).

The most successful attempt to replicate these dynamics comes from the Heuristic Switching Models (Anufriev and Hommes, 2012). The main intuition of this approach is that among different prediction heuristics, the agents focus on those that have a relatively good hypothetical past performance. On the other hand, Heuristic Switching Models cannot explain the full degree of observed individual heterogeneity, nor does it explain how the agents could learn their heuristics.

Hommes and Lux (2013) explicitly model such individual learning with Genetic Algorithms. We enhance the GA model of Hommes and Lux (2013) with the empirical micro-foundations identified by Heemeijer et al. (2009). Our GA agents use a first-

order heuristic (a mixture of adaptive and trend extrapolating expectations) to forecast the prices. They independently optimize the parametrization of their heuristics with Genetic Algorithms, thus learning to fine-tune their forecasting rules of thumb to the specific market conditions. This gives an agent-based model of explicit learning-to-forecast with strong empirical motivation for the particular forecasting behavior.

We use our Genetic Algorithms model to investigate four Learning-to-Forecast experiments. The simple linear setting of the experiment reported by Heemeijer et al. (2009) enables us to set up the model. The experiment reported by Bao et al. (2012) adds large and unanticipated shocks to the basic linear structure of Heemeijer et al. (2009). The third experiment, reported by van de Velden (2001) and Hommes et al. (2007) focuses on a non-linear cobweb economy, and is an important benchmark already investigated by Hommes and Lux (2013). Finally, the asset pricing experiment reported by Hommes et al. (2005) introduces two-periods ahead feedback between the predictions and the realized prices.

We evaluate the out of sample one-period ahead and 50 period ahead prediction accuracy of our model in comparison with benchmark models: rational expectations, a number of simple homogenous expectations models (including adaptive and naive expectations) and the Heuristic Switching Model. To our best knowledge, this chapter is the first to present an explicit econometric evaluation of how a full fledged agent-based model performs in explaining experimental data, including the individual level. For the difficult task of one-period ahead model predictions, we develop a Sequential Monte Carlo technique, a special case of the Auxiliary Particle. This is a novelty in the literature, which would rather focus on explaining the aggregate experimental outcomes.

Across the four discussed experiments, we observe a clear pattern in how different models can predict subject behavior. Rational expectations tend to explain comparatively well the negative types of price-predictions feedback. Also, for every experimental economy one can find a simple homogenous expectations model that fits this particular experiment well, but only in the short-run. Typically, contrarian or adaptive expectations have a good one-period ahead fitness to the data under negative feedback, while trend extrapolation or naive expectations outperform other models for positive feedback type of economies. However, there is no single homogenous expectations rule, including rational expectations, that can explain all the experimental economies at the same time. On the other hand, this is where the strength of our Genetic Algorithms model lies, which is able to account for both the *aggregate* outcomes and the *individual* behavior across different experiments.

Homogenous expectations models, as well as Heuristic Switching Model, take agents

as using static heuristics, whereas in reality people try to adjust their behavioral rules to the particular circumstances. This means that the simple homogenous models can replicate subject behavior only for very simple experimental economies and typically only in short out-of-sample studies. Across the first three discussed experiments, only our model remains realistically close to the experimental aggregate and individual behavior in the short-run as well as after 50 periods. For the Hommes et al. (2005) data, our model is relatively the best one and directly explains why no model performs well here: this experimental design incurs two type of attractors.

In addition, we conduct a Monte Carlo study of 50-period ahead simulations based on random initialization. With these, we can study the evolution of learning and price dynamics in the four experimental economies, by evaluating the median and 95% confidence intervals of the heuristic coefficients, which were chosen by the agents. Simulations of our model replicate the stylized results by Anufriev et al. (2013) and identify a clear pattern of individual learning. When agents face a negative feedback type of economy, a median agent will rely on adaptive expectations. If the feedback is sufficiently simple, this implies convergence of the market to the fundamental equilibrium. However, strong non-linearities will rather result in near-to-chaotic dynamics around the perfectly rational solution.

On the other hand, positive feedback induces the agents to follow a price trend. Median agent converges to a trend extrapolation rule with little emphasis on her own past predictions, which typically causes price oscillations. The more ‘difficult’ the feedback is (in terms of shocks to the fundamental solution, or non-linear law of motion of the price), the stronger trend extrapolation will be chosen by the median agent. However, positive feedback markets can also settle on the fundamental solution. We emphasize that this is not a sign for support to rational expectations framework. In fact, our model shows that the agents will switch between following price oscillations and settling on the fundamental solution, if they have to act in a complex economic environment, such as the two-period ahead type of expectation-price feedback.

The strength of our model lies in its generality, seen in the good fit to different Learning-to-Forecast experiments. Furthermore its agent-based structure allows for replicating the *individual* behavior observed in these experiments, by a realistic account of heterogeneity and learning. We therefore argue that it can be used to investigate settings with a more complicated interactions between individual agents. This can include economies with heterogeneous preferences, unequal market power, information networks or decentralized price setting. In any of these cases, heterogeneous price expectations may have important consequences for market efficiency or dynamics. Our Genetic Algorithms model gives a realistic explanation of how such heterogene-

ity between the agents emerges from their individual learning, and what can be the consequences for the aggregate market outcomes.

This contrasts the dominating framework of the perfectly rational expectations. Traditionally, economists assumed that people use sophisticated concepts such as a fundamental price or a long run equilibrium. Economists for a long time disregarded the fact that in the market practice the agents face constraints on their rationality and thus may be forced to use second-best prediction rules. As a result, rational expectation fail to describe economic phenomena, unless these are extremely simple. In the context of price expectations, we propose a model where the agents use simple behavioral rules, but adapt them to the current environment with a smart optimization procedure. This allows for a realistic description of human behavior, which explains the experimental data, both at the individual and the aggregate level, to the degree that was unattainable for the traditional literature.

Appendix 2.A 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).

2.A.1 Optimization procedures: traditional and Genetic Algorithms

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

Traditional maximization 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 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 modeled 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 a_n, b_n .

2.A.2 Binary strings

A Genetic Algorithm (GA) uses H chromosomes $g_{h,t} \in \mathbb{H}$ which are binary strings divided into N genes $g_{h,t}^n$, each encoding one candidate parameter $\theta_{h,t}^n$ for the argument

θ^n . A chromosome $h \in \{1, \dots, H\}$ at time $t \in \{1, \dots, T\}$ has predetermined length L and is specified as

$$(2.15) \quad g_{h,t} = \{g_{h,t}^1, \dots, g_{h,t}^N\},$$

such that each gene $n \in \{1, \dots, N\}$ 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)

$$(2.16) \quad g_{h,t}^n = \{g_{h,t}^{n,1}, \dots, g_{h,t}^{n,L_n}\}, \quad g_{h,t}^{n,l} \in \{0, 1\} \text{ for each } j \in \{1, \dots, L_n\}.$$

The relation between the genes and the arguments is straightforward. An integer θ^n is simply encoded by (2.16) with its binary notation. Consider now an argument θ^n which is a probability. Notice that $\sum_{l=0}^{L_n-1} 2^l = 2^{L_n} - 1$. It follows that a particular gene $g_{h,t}^n$ can be decoded as a normalized sum

$$(2.17) \quad \theta_{h,t}^n = \sum_{l=1}^{L_n} \frac{g_{h,t}^{n,l} 2^{l-1}}{2^{L_n} - 1}.$$

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 a linear transformation of a probability:

$$(2.18) \quad \theta_{h,t}^n = a_n + (b_n - a_n) \sum_{l=1}^{L_n} \frac{g_{h,t}^{n,l} 2^{l-1}}{2^{L_n} - 1}$$

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.

2.A.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

$V(\mathcal{F}(\theta_{h,t})) \equiv V(h_{k,t}) \rightarrow \mathbb{R}^+ \cap \{0\}$. For example, a non-negative function can be used directly as the fitness. If the problem is to minimize 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).

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 z th element (where $z \in \{1, \dots, \kappa\}$) is usually defined by the power function:

$$(2.19) \quad \text{Prob}(g_z = g_{h,t}) = \frac{V(g_{h,t})}{\sum_{j \in \mathbb{X}} V(g_{j,t})}.$$

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, $\chi \ll H$. For instance, GA could divide the chromosomes into pairs and sample two offspring from each pair.

Procreation is modeled 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.

Mutation

For each generation $t \in \{1, \dots, T\}$, 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.

Crossover

Let $0 \leq C_L, C_H \leq \sum_{n=1}^N L_n = L$ 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 $0 \leq \nu_L \leq \nu_H \leq N$ 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.

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.

Appendix 2.B Initialization of the model

In this appendix we discuss the initialization of the GA model for the 50-period Monte Carlo simulations, which we use to show that our model replicates experimental stylized facts. Initialization is crucial, since in the experiments the initial individual predictions influenced later outcomes, such as appearance and characteristics of oscillations, or dynamics of coordination. Two examples can be given for HHST09. Under negative feedback, the individual price forecasts coordinated only after the price itself has already converged; in our simulations we want to start with a similar degree of non-coordination between the agents, to show that disappears in the same way as happened in the experiment. Anufriev et al. (2013) suggest that under positive feedback, price oscillations require the subjects to start relatively far from the fundamental price, as was also the case for their HSM. Therefore, proper distribution of initial predictions of the experimental subjects is a crucial aspect of model calibration; without a realistic initialization, the model will not fit the data well.

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 next estimate. We use their methodology to calibrate the initial period of our model to all the other experiments, that we investigate for our GA model. In each MC simulation, we sample the initial predictions from the distribution calibrated to the respective experimental data.

HHST09

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

$$(2.20) \quad p_{i,1}^e = \begin{cases} \varepsilon_i^1 \sim U(9.546, 50) & \text{with probability } 0.45739, \\ 50 & \text{with probability } 0.30379, \\ \varepsilon_i^2 \sim U(50, 62.793) & \text{with probability } 0.23882. \end{cases}$$

With WFP we replicate the observed behavior of the subjects in the first period. Around 1/3 would predict 50, a mid-point of the suggested interval for the initial price forecast $[0, 100]$. Others were evenly spread around this focal point, with more people choosing < 50 and almost nobody choosing > 60 . Hence the distribution is a composite of a unit mass at 50 and two ‘wings’, uniform distributions preading from the focal point. ee Figure 2.11 for a visualization of the density function for this distribution.

BHST12




Figure 2.11: Density function of winged focal point distribution for HHST09. Initial prediction will be equal to $p_{i,1}^e = 50$ with probability $P = 0.30379$ (mass point); with probability $P = 0.45739$ it will fall into the left wing, where its value is drawn from $Uniform(9.546, 50)$; with probability $P = 0.23882$ it will fall into the right wing, where its value is drawn from $Uniform(50, 62.793)$. The size of the wings is scaled to their masses and lengths.

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

$$(2.21) \quad p_{i,1}^e = \begin{cases} \varepsilon_i^1 \sim U(16.406, 50) & \text{with probability } 0.32296, \\ 50 & \text{with probability } 0.35159, \\ \varepsilon_i^2 \sim U(50, 70.312) & \text{with probability } 0.32296. \end{cases}$$

HSTV07; V01

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

$$(2.22) \quad p_{i,1}^e = \begin{cases} \varepsilon_i^1 \sim U(1.875, 5) & \text{with probability } 0.17983, \\ 5 & \text{with probability } 0.36344, \\ \varepsilon_i^2 \sim U(5, 7.5) & \text{with probability } 0.45673. \end{cases}$$

HSTV05

In this experiment, the predictions are two-period ahead, hence the subjects would have to give *two* initial predictions, $p_{i,1}^e$ and $p_{i,2}^e$. 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 $p_{i,1}^e = 50$ 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:

$$(2.23) \quad p_{i,1}^e = \begin{cases} \varepsilon_i^1 \sim U(4.712, 50) & \text{with probability } 0.31306, \\ 5 & \text{with probability } 0.45536, \\ \varepsilon_i^2 \sim U(50, 64.062) & \text{with probability } 0.23158. \end{cases}$$

Define the auxiliary draw

$$(2.24) \quad p_{i,2}^{aux} = \begin{cases} \varepsilon_i^1 \sim U(3.125, 50) & \text{with probability } 0.44958, \\ 5 & \text{with probability } 0.018761, \\ \varepsilon_i^2 \sim U(50, 67.227) & \text{with probability } 0.53166. \end{cases}$$

Thus, the second period predictions are given by

$$(2.25) \quad p_{i,2}^e = \begin{cases} p_{i,2}^{aux} & \text{always if } p_{i,1}^e \neq 50, \\ p_{i,2}^{aux} & \text{with probability } 1/3 \text{ if } p_{i,1}^e = 50, \\ 50 & \text{with probability } 2/3 \text{ if } p_{i,1}^e = 50. \end{cases}$$

Appendix 2.C Parametrization of the forecasting heuristic

In this appendix, we will address two issues. First, following HHST09 we will look on the importance of the anchor, for the said experiment and the two-period ahead Hommes et al. (2005) setting. Second, we study the proper degree of allowed trend extrapolation, based on the linear feedback from HHST09.

2.C.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:

$$(2.26) \quad p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + \alpha_3 60 + \beta(p_{t-1} - p_{t-2})$$

where the fundamental price 60 serves as an anchor,¹⁸ the three α_i span a simplex and β is the trend extrapolation coefficient. Our rule (2.6) is a special case of (2.26) with the restriction that $\alpha_3 = 0$, which implies that fixed anchor is not used by the agents.

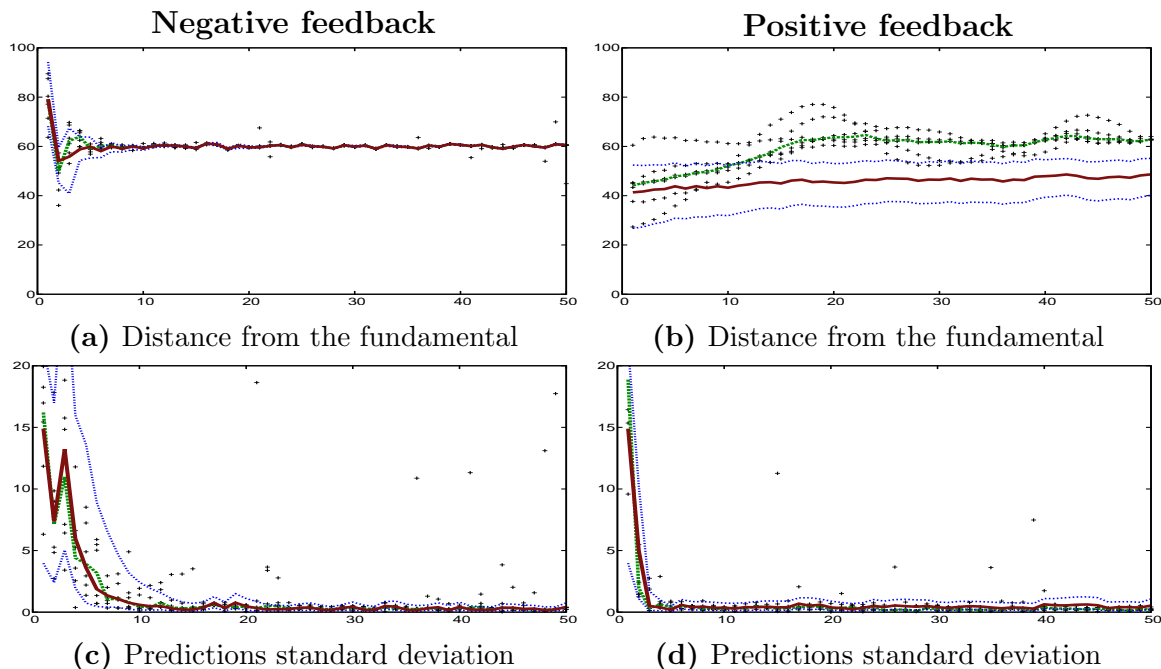


Figure 2.12: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with anchored-FOR and $\beta \in [-1.1, 1.1]$. Realized price and coordination over time. Green dashed line and black pluses represent the experimental median and group observations; red line is the median and blue dotted lines are the 95% confidence interval for the GA model. Left panel displays the negative feedback, right the positive feedback.

Experimental literature suggests that in general anchors and focal points are important in explaining human behavior. 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, the authors specify the anchor as the fundamental price 60, which however was not directly observed by the subjects. It is more plausible that they used the average price so far as an anchor, $p_t^a = p^t \equiv \sum_{s=1}^t p_s$. We will use

¹⁸Notice that what is the anchor, can be a matter of interpretation. One may think of the (2.6) rule as an anchor-based rule as well, since it can be rewritten as a rule that adjusts the previous price forecast with the latest observed price and trend.

thus anchored-FOR specified as

$$(2.27) \quad p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + \alpha_3 \left(\sum_{s=1}^t p_s \right) + \beta (p_{t-1} - p_{t-2}).$$

We consider the Monte Carlo (MC) simulations exactly as in the first part of Section 3.3, but for the GA model based on (2.27) with $\beta \in [-1.1, 1.1]$. The results are presented on Figure 2.12. We observe for the positive feedback that, in contrast to our restricted model without an anchor, the GA model based on FOR as in (2.27) does not predict oscillations at all, but rather a sluggish convergence towards the fundamental. This is seen in the stable median price, bounded by relatively narrow 95% CI. This specification misses most of the dynamics observed in half of the experiment. 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.




Figure 2.13: HSTV05 with $p^f = 60$: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with anchored-FOR and $\beta \in [-1.1, 1.1]$ and $\beta \in [-1.3, 1.3]$. Realized price over time: red line is the median and blue dotted lines are the 95% confidence interval for the GA model.

2.C.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 unraveled around the fundamental price, which again suggests that the subjects tried to extrapolate the trend around it. To address this issue, we run the 50-period ahead MC simulation like in Section 4.3, but with FOR (2.14) replaced by the anchored-FOR rule (2.27) adapted for the two-period ahead setting, assuming that

the fundamental price is $p^f = 60$.

Results for two specifications (with allowed trend extrapolation $\beta \in [-1.1, 1.1]$ and $\beta \in [-1.3, 1.3]$) are presented on Figure 2.13. 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 the GA model would decrease its fitness to the experimental data.

2.C.3 Degree of trend extrapolation

Recall that the GA requires a predefined finite interval for the optimized parameters. In the case of our GA model based on (2.6), 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 (2.6) rule to use 1.1 as the upper bound to the trend.

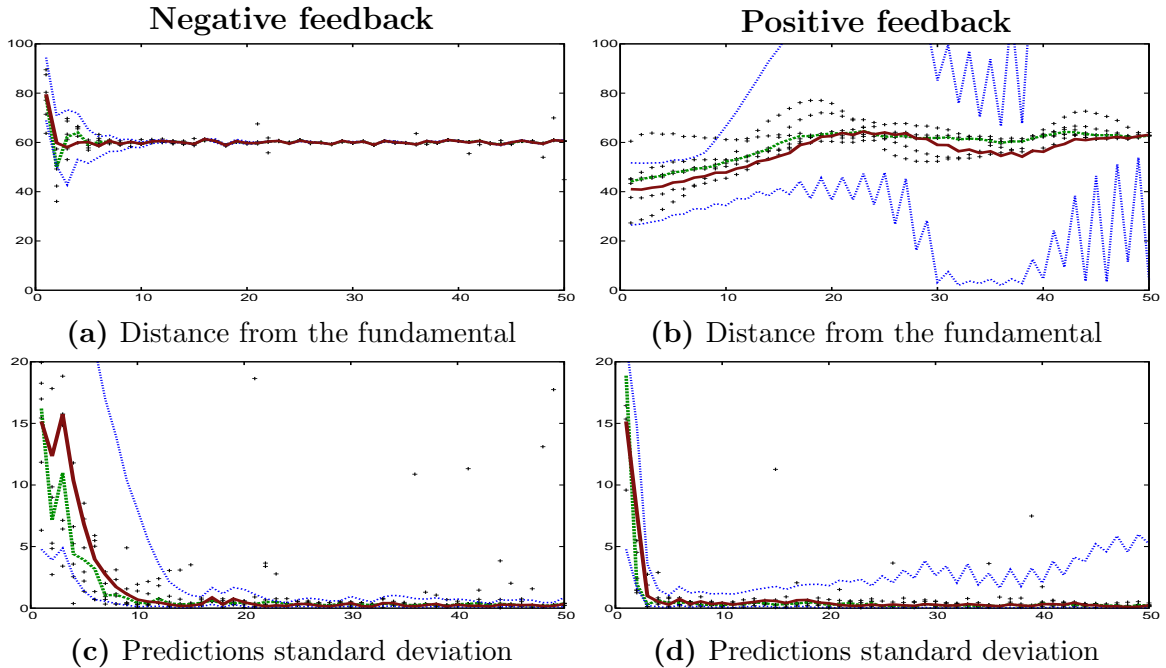


Figure 2.14: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-1.5, 1.5]$. Realized price and coordination over time. Green dashed line and black pluses represent the experimental median and group observations; red line is the median and blue dotted lines are the 95% confidence interval for the GA model. Left panel displays the negative feedback, right the positive feedback.

It turns out (not surprisingly) that the allowed trend extrapolation interval has little effect on the behavior of our GA model under the negative feedback. However, the larger the interval $\beta \in [\beta_L, \beta_H]$ is, the bigger the amplitude of the price fluctuations generated under the positive feedback. Thus we experimented with different β 's, trying to calibrate the model to the the experimental oscillations. We used the same Monte Carlo experiments as in the first part of Section 3.3.

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 Figure 2.14. On the other hand, specification with $\beta \in [-0.5, 0.5]$ has narrow CI, but predicts small oscillations, see Figure 2.15. We found the model with $\beta \in [-1.1, 1.1]$ to be the best trade-off between fit and explanatory power of the experiment.

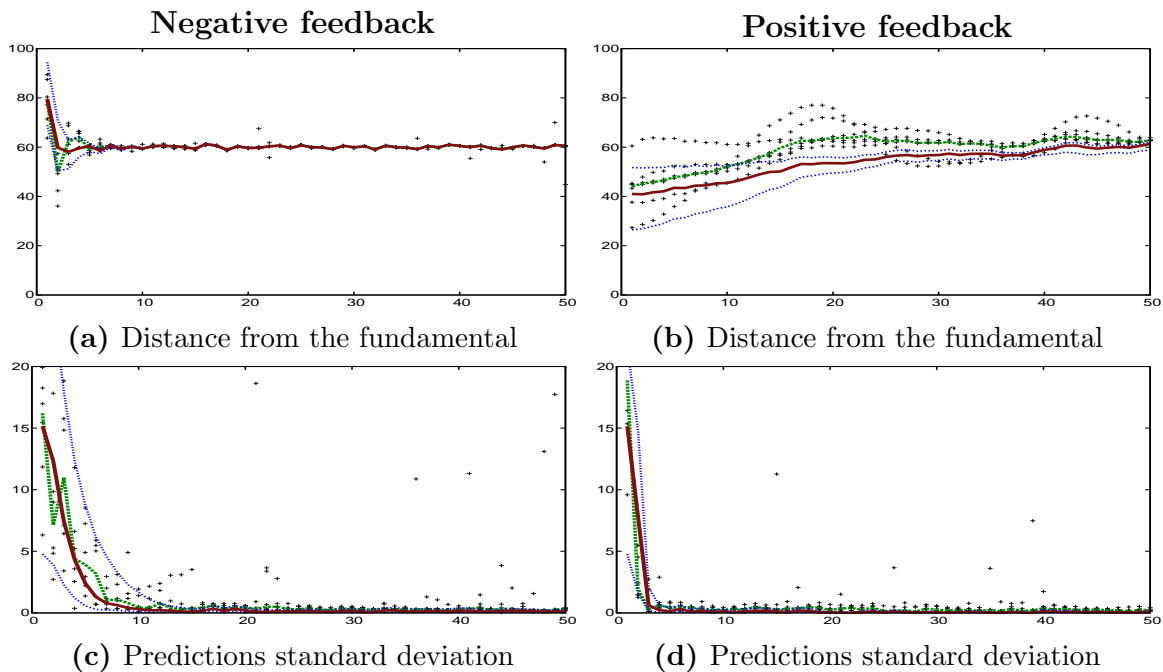


Figure 2.15: HHST09: 50-period ahead Monte Carlo simulation (1000 markets) for the GA model with $\beta \in [-0.5, 0.5]$. Realized price and coordination over time. Green dashed line and black pluses represent the experimental median and group observations; red line is the median and blue dotted lines are the 95% confidence interval for the GA model. Left panel displays the negative feedback, right the positive feedback.

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.

Appendix 2.D Definition of forecasting rules

Table 2.11 gives the exact specification for all the prediction rules used in the one-period ahead forecasting exercises for the four experiments. The full HSM specification can be found in Anufriev et al. (2013).

Rule	Prediction
<i>Homogeneous rules</i>	
Trend extr.	$p_t^e = p_{t-1} + (p_{t-1} - p_{t-2})$
Adaptive	$p_t^e = 0.75p_{t-1} + 0.25p_{t-1}^e$
Contrarian	$p_t^e = p_{t-1} - 0.5(p_{t-1} - p_{t-2})$
Naive	$p_t^e = p_{t-1}$
RE	$p_t^e = p^f$
<i>Heterogeneous rules</i>	
HSM	two heuristic model (trend extrapolation vs. adaptive expectations)
GA model	$p_{i,t}^e = \alpha_{i,t}p_{t-1} + (1 - \alpha_{i,t})p_{i,t-1}^e + \beta_{i,t}(p_{t-1} - p_{t-2})$
	$\beta \in [-1.1, 1.1]$ $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [-1.1, 1.1]$
	$\beta \in [0, 1.1]$ $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [0, 1.1]$

Table 2.11: Specification of the forecasting rules p_t^e for one-period ahead forecasting environment.

Appendix 2.E APF for the GA model

In this appendix we discuss how Auxiliary Particle Filter can be used to estimate the mean squared error of one-period ahead predictions of our GA model. This is complicated task, since the model operates on unobservable evolution of heuristics. Thus what it predicts for period t is in fact a distribution, which is a function of all periods until and including $t - 1$. The following discussion explains formally how it can be represented as a Sequential Importance Sampling problem. Next we show how exactly one can estimate the period-to-period distribution of the heuristic evolution through a simple Monte Carlo integral. The final part of this appendix presents the estimated MSE for the four discussed experiments, along with exemplary price and price forecasts from the experiments and their one-period ahead predictions of our GA model.

2.E.1 General specification

We introduce the following notation. Let a_t denote the state of the model at the beginning of period t . Specifically, we mean the set of the six sets of chromosomes H_{t-1} , which correspond to the six lists of twenty heuristics $H_{i,t-1}$ of each agent. Be the beginning of the period t we mean that the price p_{t-1} is already observed, but the heuristics are not yet updated (hence the subscript of H_{t-1}). Consider an experimental group X , for which we can observe its underlying chromosomes only indirectly, through the realized prices and predictions picked by the subjects (observational variables). Both the state and observed variables are evolving according to a distribution $q(\cdot)$ over time $t \in \{1, \dots, 50\}$. Denote also $p_t^{e,GRX} = \{p_{1,t}^{e,GRX}, \dots, p_{6,t}^{e,GRX}\}$ as the set of six individual predictions from period t in the experimental group X . Henceforth t in the superscript denotes history of the variable, so $p^{t-1,GRX} = \{p_1^{GRX}, \dots, p_{t-1}^{GRX}\}$ and $p^{e,GRX,t-1} = \{p_1^{e,GRX}, \dots, p_{t-1}^{e,GRX}\}$.

Our problem is to define the baseline distribution $q(p_t^{e,GRX}, a_{t-1})$, that is, to evaluate the distribution of the real predictions $p_{i,t}^{e,GRX}$ given the predictions from the period $t-1$ and what they signal could have been the chromosomes H_{t-1} from the period $t-1$. This is a typical state-space model problem. Essentially, $q(\cdot)$ can be decomposed as

$$\begin{aligned}
 q(p_t^{e,GRX}, a_t) &= q(p_t^{e,GRX} | a_t) \times q(a_t) \\
 &= q(a_t | p_t^{e,GRX}) \times q(p_t^{e,GRX}), \\
 q(p_t^{e,GRX} | a_{t-1}) &= q(p_t^{e,GRX} | a_t) \times q(a_t | a_{t-1}), \\
 q(a_t | p_{t-1}^{e,GRX}) &= q(a_t | a_{t-1}) \times q(a_{t-1} | p_{t-1}^{e,GRX}), \\
 q(a_t | a^{t-1}) &= q(a_t | a_{t-1}), \\
 (2.28) \quad q(p_t^{e,GRX} | p^{e,GRX,t-1}) &= q(p_t^{e,GRX} | p_{t-1}^{e,GRX}).
 \end{aligned}$$

The first three equalities are a simple consequence of the structure of our model, whereas the two last ones are implied by it being a Markovian process.

Given the information structure of the experiment we can assume that in each period, *conditional on the history until $t-1$ and the chromosome set H_t* , the individual predictions are independent between the agents, and their joint density is a simple

product of the marginal densities of individual forecasts:

$$\begin{aligned}
 p_t^{e,GRX} &\sim q\left(p_t^{e,GRX}|a_t, p^{GRX,t-1}, p^{e,GRX,t-1}\right) \\
 (2.29) \qquad &= \prod_{i=1}^6 q\left(p_t^{e,GRX}|a_t, p^{GRX,t-1}, p^{e,GRX,t-1}\right).
 \end{aligned}$$

Unfortunately, $q(a_t|a_{t-1})$ is not that simple to work with. As explained earlier, it is not feasible to represent this problem analytically or to linearize it. Nevertheless, it is fairly simple to simulate a_t conditional on a_{t-1} . Therefore, we focus on SISR technique known as Auxiliary Particle Filter (APF) (Johansen and Doucet, 2008).¹⁹ APF works on two distributions (baseline $q(\cdot)$ and importance $g(\cdot)$), which are used to generate MC weights and updates for the particles, and these we will approximate with simple MC integrals.

The general APF algorithm is discussed by Johansen and Doucet (2008) and below we provide its interpretation to our problem. The starting point of the APF is the question: what is $q(a^t|p^{e,GRX,t})$ the distribution of possible heuristics throughout the experiment, conditional on the observed individual predictions from group X , until some period $t = 1, \dots, T$. We need this to estimate $MSE_{X,t}^{\text{price forecasts}}$, MSE of the model predictions *specifically for period t*, which for brevity we will denote in this appendix simply as $MSE_{X,t}^e$. To be specific,

$$\begin{aligned}
 MSE_{X,t}^e &= \mathbb{E}_q \{MSE_X^e(a^t) | p^{e,GRX,t}\} \\
 (2.30) \qquad &= \int MSE_X^e(a^t) q(a^t|p^{e,GRX,t}) da^t,
 \end{aligned}$$

the expected MSE of the GA model conditional on the observed price forecasts in the group X . The subscript q in \mathbb{E}_q signals that the expectation is taken in respect to the $q(\cdot)$ density. It turns out that in practice this estimator has better properties if we use importance sampling approach (the specific reason is that we obtain a better approximation of the tails of the distribution, see Johansen and Doucet, 2008). Denote the importance distribution as $g(\cdot)$ and assume that $g(a_t|a_{t-1}) = q(a_t|a_{t-1})$ (which is a standard assumption for APF). It follows that $g(p_t^{e,GRX}|a_{t-1})$ can be decomposed in the same manner as the baseline distribution in equation (2.28). Hence we focus

¹⁹It is extremely difficult, also in conceptual terms, to define a reverse distribution of the model at period t conditional on period $t + 1$, given the complexity of GA operators. As a result, we leave open the question whether econometrically more efficient filtering-smoothing techniques can be used for the case of our model.

instead

$$\begin{aligned}
 MSE_{X,t}^e &= \int MSE_X^e(a^t) \frac{q(a^t|p^{e,GRX,t})}{g(a^t|p^{e,GRX,t})} g(a^t|p^{e,GRX,t}) da^T \\
 &= \mathbb{E}_g \left\{ MSE_X^e(a^t) \frac{q(a^t|p^{e,GRX,t})}{g(a^t|p^{e,GRX,t})} \right\} \\
 (2.31) \quad &\equiv \mathbb{E}_g \{ MSE_X^e(a^t) w_t \},
 \end{aligned}$$

where $g(\cdot)$ is the importance density and w_t denote the importance weights.

This expected value can be approximated with a MC integral, based on samples from $g(\cdot)$. It still not easy to work with. With (2.31) we estimate the MSE of the model's one period-ahead predictions for a *specific* period t . Since we do it *conditional on the whole history* until this period, we would need to sample *all* t periods conditional on t periods of the data. This means that MSE is a function of the contemporary heuristics, but the MC weights are computed as probabilities for the whole history of the heuristics' evolution!

The next step is to represent this integral in a recursive way. To be specific, we can transform density $q(\cdot)$ from period t as

$$\begin{aligned}
 q(a^t|p^{e,GRX,t}) &= \frac{q(a^t, p^{e,GRX,t})}{q(p^{e,GRX,t})} \\
 &= \frac{q(a_t|p^{e,GRX,t}) q(a^{t-1}, p^{e,GRX,t})}{q(p^{e,GRX,t})} \\
 (2.32) \quad &= q(a_t|p^{e,GRX,t-1}) q(a^{t-1}|p^{e,GRX,t-1}),
 \end{aligned}$$

where the last equality holds because the formation of heuristics at period t is independent from further forecasting at that period, that is $q(a^{t-1}|p^{e,GRX,t}) = q(a^{t-1}|p^{e,GRX,t-1})$ (the relation goes the other way around). In the same way we can decompose

$$(2.33) \quad g(a^t|p^{e,GRX,t}) = g(a_t|p^{e,GRX,t-1}) g(a^{t-1}|p^{e,GRX,t-1}).$$

It follows that

$$\begin{aligned}
 w_t &= \frac{q(a^t|p^{e,GRX,t})}{g(a^t|p^{e,GRX,t})} \\
 &= \frac{q(a_t|p^{e,GRX,t-1}) q(a^{t-1}|p^{e,GRX,t-1})}{g(a_t|p^{e,GRX,t-1}) g(a^{t-1}|p^{e,GRX,t-1})} \\
 (2.34) \quad &= \frac{q(a_t|p^{e,GRX,t-1})}{g(a_t|p^{e,GRX,t-1})} w_{t-1}.
 \end{aligned}$$

Next denote $\tilde{w}_t \equiv q(p^{e,GRX,t}) w_t$. The Markovian nature of both $q(\cdot)$ and $g(\cdot)$, together with the decomposition (2.28) and (2.34) gives us

$$\begin{aligned}
 \tilde{w}_t &= \frac{q(a^t, p^{e,GRX,t})}{g(a^t | p^{e,GRX,t})} \\
 &= \tilde{w}_{t-1} \frac{q(a_t | a_{t-1}) q(p_t^{e,GRX} | a_t)}{g(a_t | a_{t-1}) g(p_t^{e,GRX} | a_{t-1})} \\
 (2.35) \quad &= \tilde{w}_{t-1} \frac{q(p_t^{e,GRX} | a_t)}{g(p_t^{e,GRX} | a_{t-1})},
 \end{aligned}$$

where the last equality follows from the APf assumption that $q(a_t | a_{t-1}) = g(a_t | a_{t-1})$.

These lengthy derivations can now be put to a good use. If the MC integral for period $t - 1$ is based on a good sample of potential market outcomes, we can simulate them for one extra period and hence use (2.35) recursion to reevaluate the fit of each such market for period t . This is the basic idea of APF: every counter-factual market is a particle that follows the data. If our model is well specified, the particles should have little problems doing so. This is measured through two things. MSE is the direct signal of the particle about the reliance of the model. Second, APF weights \tilde{w}_t show a relative reliance of the particle itself. The former depends only on the current period. Strictly speaking, the latter depends on the whole history, but the recursion (2.35) allows to express it as a straightforward recursive process.

In order to keep the pool of particles representative, it is customary in the literature to resample them. This is done based on the APF weights, since we want to have more of the ‘reliable’ particles. Thus, the whole algorithm becomes an iterative procedure over $b \in \{1, \dots, B\}$ particles:

1. Compute resample weights $g(p_t^{e,GRX} | a_{t-1}^b) w_{t-1}^b$, normalize them, and use them as probabilities to resample (with replacement) B particles a_{t-1}^{b*} from the existing a_{t-1}^b .
2. Draw a_t^b from a_{t-1}^{b*} , compute \tilde{w}_t^b from (2.35), normalize them and use them to compute MSE for period t .

Specifically for our model, we follow Anufriev et al. (2013) and we assume that the baseline distribution of the six realized predictions is given by standard normal

distribution.²⁰ To be specific, we represent it as

$$(2.36) \quad p_t^{e,GRX} \sim q(p_t^{e,GRX} | a_t, p^{GRX,t-1}, p^{e,GRX,t-1}) = \prod_{i=1}^6 N(\hat{p}_{i,t}^{e,GA}, 1).$$

We simplify the notation by suppressing $p_t^{e,GRX} \sim q(p_t^{e,GRX} | a_t)$. For a general case, (2.36) density of the experimental predictions $p_t^{e,GRX}$ is just a product of normal standard densities centered around the forecasts predicted by a model.

For $g(p_t^{e,GRX} | a_t)$ we use (a product of six) Student-t with one degree of freedom. This density is again analytically straightforward and compares $p_t^{e,GRX}$ with $\hat{p}_t^{e,GA}$ price forecasts predicted by the GA model as in equation (2.11). We define $\hat{p}_{i,t}^{e,GA}$ predicted individual price forecasts and predicted realized price \hat{p}_t^{GA} as in the previous subsection.

We use 128 particles $a_t^{}$ for $b \in \{1, \dots, 128\}$ (128 sets of six heuristic sets for six agents), with full resampling.²¹ The core problem of our investigation lies with the $q(a_t | a_{t-1})$ and $g(a_t | a_{t-1})$ distributions, which cannot be tracked analytically. We approximate them with a MC integral in the following fashion. Consider first the baseline distribution. At the beginning of each iteration $t > 1$ of AFP, for each particle $a_{t-1}^{}$ (that is, for each set of six agents and their chromosomes), we simulate 256 counter-factual iterations, with each market iteration $s \in \{1, \dots, 256\}$ of the following structure (see section 3 for details of the model):

1. Heuristics evolve with the GA operators based on the experimental group data, which is observable until period $t - 1$.
2. Agents report their forecasts $p_{i,t}^{e,<b,s>}$, which hence also generates the counter-factual price $p_t^{<b,s>}$.

The counter-factual price forecasts $\hat{p}_{i,t}^{e,<b,s>}$ (again, we look at the expected price forecasts, see formula (2.11)) and prices are evaluated against the actual experimental prices and forecasts, which gives the baseline density

$$(2.37) \quad q(p_t^{e,GRX} | a_{t-1}^{}) = \frac{1}{256} \sum_{s=1}^{256} \prod_{i=1}^6 N\left(\hat{p}_{i,t}^{e,<b,s>} - p_{i,t}^{e,GRX}, 1\right).$$

The importance distribution is generated in a symmetric way. For each period t , we simulate 256 counter-factual two-period market iterations for each particle b , that is for each counter-factual market $s \in \{1, \dots, \}$ we repeat the above iteration twice. As

²⁰For APF, the choice of variance of the distributions is not important. We use standard normal for the sake of computational efficiency.

²¹Larger particle samples are computationally involved and do not seem to increase the efficiency of the estimates.

one can expect, for period t the importance density is based on the data until period $t - 2$, and the simulated counter-factual prices and forecasts from period $t - 1$ are used by the GA agents to act in period t . We use these to compute price $\hat{p}_{t,importance}^{<b,s>}$ and $\hat{p}_{i,t,importance}^{e,<b,s>}$ individual price forecasts predicted by the model. We therefore define

$$(2.38) \quad g(p_t^{e,GRX} | a_{t-1}^{}) = \frac{1}{256} \sum_{s=1}^{256} \prod_{i=1}^6 T_1 \left(\hat{p}_{i,t,importance}^{e,<b,s>} - p_{i,t}^{e,GRX} \right),$$

where T_1 denotes Student-t distribution with one degree of freedom.²²

We use the baseline (2.36) and the importance (2.38) distributions for the APF updating as described above.²³ We run a separate APF for each of the twelve investigated experimental groups. Like in the 50-period ahead simulations, the chromosomes (or the particles) are initialized at random from the ‘uniform’ distribution defined above. Notice, however, that in this case we do not have the problem of the initial predictions or prices, since the APF works on the period-to-period basis, independently for each experimental group.

For deterministic models like RE or HSM or homogeneous heuristic models, the APF effectively reduces to the procedure reported by Anufriev et al. (2013), since all the particles would be the same and had the same weights. Moreover, APF procedure reduces to the simple SMC described in section 3, if for the importance distribution we use standard normal instead of Student-t.

For each experimental group, we focus on seven variables in total, which we obtain by using the APF weighting of the particles. For each of the last 47 periods in each group, we look at the (mean) one-period ahead prediction of the price, as well as at the (mean) one-period ahead predictions of individual price forecasts. Notice that we compute the expected APF MSE’s, instead of MSE’s for the expected prices or predictions. We compute these variables for each group, and average them separately over the two treatments to obtain the average MSE of the model’s prediction of the prices and individual forecasts.

The drawback of this method is the computational time. For a dual-core Pentium

²²For the sake of computational efficiency, we perform the resampling $a_t \sim a_{t-1}$ of particle b in the following way. We pick at random one counter-factual s market from the importance MC density for period $t + 1$ and use the chromosomes as they are at the end of the first iteration for this counter-factual market — that is, one of the 256 chromosome sets updated based on the data until $t - 1$ in the first step of the importance distribution estimation.

²³Both the baseline and importance densities are a product of six independent densities, which can take very low values in the first periods for some experimental groups. To ensure numerical stability, we multiply both joint densities by 10^{60} (or each of the six marginal distributions by 10^{10}) and truncate them at 10^{-100} .

with 2.7GHz clock and 3.21 GB RAM, a shot of APF for *one experimental group* from HHST09 takes approximately 50 minutes for a relatively small number of 32 particles.

2.E.2 Results for the four experiments

Results for the four experiments are presented in this subsection: Table 2.12 and Figure 2.16 for HHST09; Table 2.12 and Figure 2.17 for BHST12; Table 2.14 and Figure 2.18 for HSTV07; and Table 2.15 and Figure 2.19 for HSTV05.

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	2.999	17.106	0.889	4.156
Genetic Algorithm model				
<i>Baseline Sequential Monte Carlo</i>				
$\beta \in [-1.1, 1.1]$	4.95	25.017	0.806	4.235
$\beta \in [0, 1.1]$	4.496	25.012	0.802	4.198
<i>Auxiliary Particle Filter</i>				
$\beta \in [-1.1, 1.1]$	3.889	21.12	0.716	3.918
$\beta \in [0, 1.1]$	3.558	23.03	0.696	3.943

Table 2.12: HHST09: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over 6 negative feedback and 6 positive feedback groups.




Figure 2.16: HHST09: APF one-period ahead predictions of the GA model with $\beta \in [-1.1, 1.1]$ for prices and price forecasts of subject 1 from sample groups from each treatment. Black line denotes the experimental variable and red boxes display the model predictions.

MSE	Negative feedback		Positive feedback	
	Prices	Forecasts	Prices	Forecasts
Trend extr.	114.061	121.329	1.183	2.165
Adaptive	<i>3.689</i>	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	38.309	45.679	0.9996	2.024
Genetic Algorithm model				
<i>Baseline Sequential Monte Carlo</i>				
$\beta \in [-1.1, 1.1]$	10.247	21.464	0.342	2.059
$\beta \in [0, 1.1]$	4.208	15.267	0.341	<i>2.036</i>
<i>Auxiliary Particle Filter</i>				
$\beta \in [-1.1, 1.1]$	7.902	16.534	<i>0.335</i>	2.067
$\beta \in [0, 1.1]$	3.348	<i>11.667</i>	0.329	1.923

Table 2.13: BHST12: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over 8 negative feedback and 8 positive feedback groups.




Figure 2.17: BHST12: APF one-period ahead predictions of the GA model with $\beta \in [-1.1, 1.1]$ for prices and price forecasts of subject 1 from sample groups from each treatment. Black line denotes the experimental variable and red boxes display the model predictions.

Treatments	Stable		Unstable		Strongly unstable	
	Prices	Forecasts	Prices	Forecasts	Prices	Forecasts
Trend extr.	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	<i>0.364</i>	0.385	2.257	1.844
HSM	0.212	0.474	0.52	0.732	3.065	3.691
Genetic Algorithm model						
<i>Baseline Sequential Monte Carlo</i>						
$\beta \in [-1.1, 1.1]$	0.13	0.393	0.866	0.795	5.547	3.25
$\beta \in [0, 1.1]$	0.07	0.31	0.25	0.531	3.079	2.358
<i>Auxiliary Particle Filter</i>						
AR1	0.042	0.267	0.321	0.476	3.187	2.28
$\beta \in [-1.1, 1.1]$	0.085	0.294	0.639	0.625	3.173	2.196
$\beta \in [0, 1.1]$	<i>0.0545</i>	<i>0.266</i>	0.229	<i>0.448</i>	<i>2.465</i>	<i>1.861</i>

Table 2.14: HSTV07: one-period ahead predictions. MSE of the experimental prices and forecasts, for the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and Genetic Algorithms models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over 6 groups for each treatment (stable, unstable, strongly unstable).

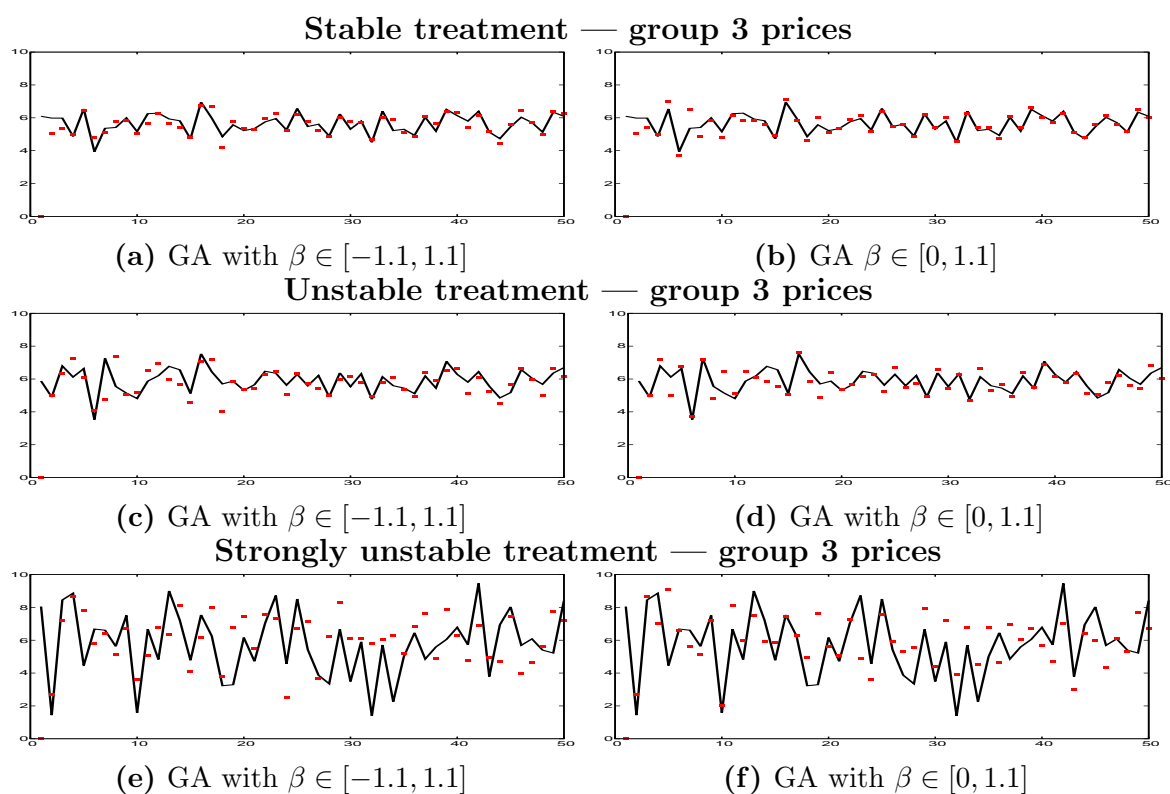


Figure 2.18: HSTV07: APF one-period ahead predictions for GA model with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$, for prices from the third group of each treatments (stable, unstable, strongly unstable). Black line denotes the experimental variable and red boxes display the model predictions.

MSE	Prices	Forecasts
Trend extr.	17.4527	55.0898
Adaptive	44.125	25.3157
Contrarian	59.3905	30.8646
Naive	31.6864	20.8416
RE	96.0328	145.998
Genetic Algorithm model		
<i>Baseline Sequential Monte Carlo</i>		
$\beta \in [-1.1, 1.1]$	42.224	74.95
$\beta \in [0, 1.1]$	5.934	30.341
<i>Auxiliary Particle Filter</i>		
$\beta \in [-1.1, 1.1]$	19.794	39.226
$\beta \in [-0, 1.1]$	17.899	39.256
$\beta \in [-1.3, 1.3]$	20.55	45.274
$\beta \in [0, 1.3]$	17.104	39.291

Table 2.15: HSTV05: one-period ahead predictions. MSE of the experimental prices and forecasts, for Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). MSE averaged over all experimental groups.

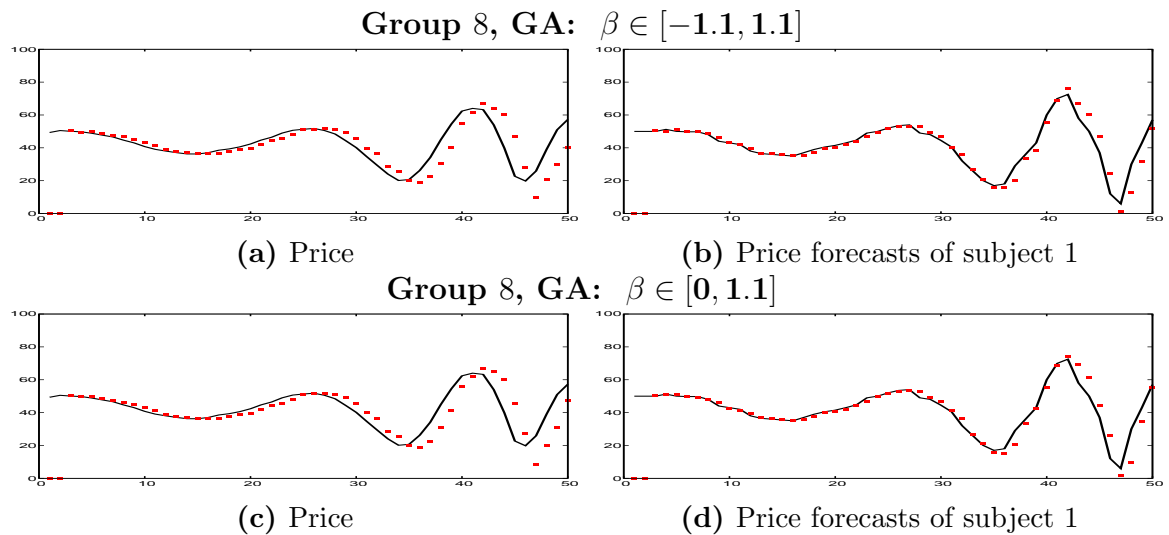


Figure 2.19: HSTV05: APF one-period ahead predictions of the GA model with $\beta \in [-1.1, 1.1]$, for prices and price forecasts of subject 1 from group 8. Black line denotes the experimental variable and red boxes display the APF one-period ahead predictions.

Appendix 2.F Price autocorrelation in the cobweb experiment

Table 2.16 gives the first three autocorrelations of the experimental groups in HSTV07 and the 50-period ahead simulations of the GA and benchmrak models.

Treatments	Stable			Unstable			Strongly unstable		
	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 extr.	-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	-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: $\beta \in [-1.1, 1.1]$	-0.1102	-0.3232	0.002674	-0.2201	-0.2013	0.0362	-0.2478	-0.3148	0.2432
GA: $\beta \in [0, 1.1]$	-0.2955	0.1059	-0.171	-0.3882	0.1405	-0.1848	-0.6206	0.4428	-0.356

Table 2.16: HSTV07: 50-period ahead predictions. First three autocorrelations in the experimental groups, and in the 50-period ahead simulations of the Trend Extrapolation, Adaptive, Contrarian, Naive and Rational Expectations, Heuristic Switching Model and GA models (FOR with $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$). Autocorrelations averaged over six groups for each treatments (stable, unstable and strongly unstable).