UvA-DARE (Digital Academic Repository)

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

Anufriev, M.; Hommes, C.; Makarewicz, T.

DOI

[10.1093/jeea/jvy028](https://doi.org/10.1093/jeea/jvy028)

Publication date

2019

Document Version

Final published version

Published in

Journal of the European Economic Association

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Anufriev, M., Hommes, C., & Makarewicz, T. (2019). Simple Forecasting Heuristics that Make us Smart: Evidence from Different Market Experiments. *Journal of the European Economic Association*, 17(5), 1538-1584. <https://doi.org/10.1093/jeea/jvy028>

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: <https://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.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

SIMPLE FORECASTING HEURISTICS THAT MAKE US SMART: EVIDENCE FROM DIFFERENT MARKET EXPERIMENTS

Mikhail Anufriev

University of Technology Sydney,
Business School

Cars Hommes

CeNDEF, University of Amsterdam
and Tinbergen Institute

Tomasz Makarewicz

University of Bamberg

Abstract

In this paper we address the question of how individuals form expectations and invent, reinforce, and update their forecasting rules in a complex world. We do so by fitting a novel, parsimonious, and empirically validated genetic algorithm learning model with explicit heterogeneity in expectations to a set of laboratory experiments. Agents use simple linear first order price forecasting rules, adapting them to the complex evolving market environment with a Genetic Algorithm optimization procedure. The novelties are: (1) a parsimonious experimental foundation of individual forecasting behavior; (2) explanation of individual and aggregate behavior in three different experimental settings, (3) improved one- and 50-period ahead forecasting of experiments, and (4) characterization of the mean, median, and empirical distribution of forecasting heuristics. The median of the distribution of GA forecasting heuristics can be used in designing or validating simple Heuristic Switching Models. (JEL: C53, C63, C91, D03, D83, D84)

The editor in charge of this paper was Dirk Krueger.

Acknowledgments: The authors would like to thank the Editor and five referees for their detailed reports and feedback that significantly improved our results. We are thankful for stimulating discussions to Blake LeBaron and Murat Yildizoglu, as well as to the participants of the Summer School of the Econometric Society, Seoul, 2013, Australian Conference of Economists, Brisbane, 2015, the Workshop on the Economic Science with Heterogeneous Interacting Agents (WEHIA), Sophia Antipolis, 2015, the “Computation in Economics and Finance” (CEF) conference, Taipei, 2015, the MACFINROBODS conference, Frankfurt, 2017, the Macro, Banking & Finance Conference, Milan, 2017, and to the seminar participants at the University of Amsterdam, Higher School of Economics, St. Petersburg, Surrey and the DG-R seminar at the ECB, Frankfurt, 2018. We gratefully acknowledge the financial support from NWO (Dutch Science Foundation) Project No. 40611142 “Learning to Forecast with Evolutionary Models.” The research leading to these results has also received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement Integrated Macro-Financial Modeling for Robust Policy Design (MACFINROBODS, grant no. 612796). Mikhail Anufriev acknowledges financial support from the Australian Research Council through Discovery Project (DP140103501).

E-mail: Mikhail.Anufriev@uts.edu.au (Anufriev); C.H.Hommes@uva.nl (Hommes); tomasz.makarewicz@uni-bamberg.de (Makarewicz)

1. Introduction

Expectations are a cornerstone of many dynamic economic models. In this paper we address the question of how individuals form expectations and invent, reinforce and update their forecasting rules in a complex world. We do so by fitting a novel, parsimonious, and empirically validated genetic algorithm learning model with explicit heterogeneity in expectations to a set of laboratory experiments. The traditional literature after Muth (1961), Lucas (1972), and others emphasizes the Rational Expectations (RE) hypothesis, which states that the expectations of all agents have to be model consistent. While the RE hypothesis is a natural benchmark theory of expectation formation, as a realistic description of real world behavior it faces challenges both theoretically and empirically. A substantial literature addresses many of the shortcomings of RE by an alternative boundedly rational approach attributing to individuals an adaptive learning rule. These learning models are typically motivated by restricting forecasting rules to nest the RE of interest and to specify a reasonable learning algorithm that adjusts the forecasting model in light of forecast errors, thereby, preserving the cross-equation restrictions that are the hallmark of RE models. Much of this literature is motivated by a cognitive consistency principle, formulated by Sargent (1993), Evans and Honkapohja (2001), and others, that holds that individuals should forecast like a good econometrician who specifies a forecasting model and revises the model in light of data.¹

An influential set of papers following Brock and Hommes (1997) developed a learning model where individuals entertain a set of possible forecasting rules and select the rule that performs best. In this model expectations can be heterogeneous and time-varying in macroeconomic and asset-pricing settings with a strong feedback mechanism from expectations onto equilibrium outcomes. Applications of this learning model include asset prices (Brock and Hommes 1998; Anufriev and Panchenko 2009; Branch and Evans 2010) and business cycle dynamics and monetary policy (Branch and McGough 2010; Branch and McGough 2011; Anufriev et al. 2013a). Moreover, empirical evidence for this predictor-selection learning model can be seen in survey data² (Branch 2004), estimated financial models (Boswijk et al. 2007), estimated DSGE models (Cornea-Madeira et al. 2017), housing market models (Bolt et al. 2014),

1. Nonlearning streams of macroeconomic literature on bounded rationality include the rational inattention approach, see Sims (2010) for a comprehensive review, the rational or “near-rational” beliefs approach, see Woodford (2010) and Kurz and Motolese (2011), and the educative approach of Guesnerie (2005). In the rational inattention literature agents do not react on all relevant information quickly but instead process information at some finite rate. Similarly to the adaptive learning models it induces sluggish behavior that then can be translated into sluggishness of economic variables. “Near-Rational” expectations allow distortions of expectations with respect to the RE case within certain bounds. Educative learning means that agents’ expectations are consistent with the actual law of motion and *some* common knowledge assumption about expectations of others. All these threads of literature produce a richer set of equilibrium dynamics than RE, possibly leading to complex dynamics (Bullard 1994). Woodford (2013) gives a recent survey.

2. Perhaps, the most prominent recent example on failure of REs comes from the housing market in the United States, which in the last decade exhibited first a boom and then a collapse. Case et al. (2012)

and, most importantly for this paper, in a series of “learning-to-forecast” experiments reviewed in Hommes (2011) and Wagener (2014).³ Heterogeneity in expectations has been a general feature in these empirical works, survey data analyses, and laboratory experiments.

Learning-to-Forecast (LtF) experiments offer a simple laboratory testing ground for adaptive learning mechanisms (Lucas 1986). These controlled experimental economies have a straightforward and unique fundamental equilibrium consistent with RE. As in real markets, subjects observe the 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. The outcomes of many LtF laboratory experiments contradict the RE hypothesis, see the review in Hommes (2011). The experiments in Hommes et al. (2005), henceforth **HSTV05**, showed that subjects can coordinate on oscillating and serially correlated time series, and that convergence to the fundamental equilibrium happens only under severe restrictions on the underlying law of motion. Further experiments in Heemeijer et al. (2009), henceforth **HHST09**, and Bao et al. (2012), henceforth **BHST12**, demonstrated that the expectations feedback structure plays a crucial role. Negative feedback systems (i.e., where more optimistic forecasts lead to lower market prices, as in supply driven commodity markets) tend to generate convergence to the fundamental equilibrium rather easily, whereas positive feedback systems (i.e., where more optimistic forecasts lead to higher market prices, as in speculative asset markets) typically generate behavior with the price oscillating around the fundamental equilibrium dynamics.⁴

conduct a survey of households’ expectations about changes in their home value over the next years and reject the RE hypothesis. They conclude that people’s expectations are consistent with trend-extrapolation and that people systematically misjudge the long-term value of their houses. Similar effects were observed with expectations before the previous housing bubble in the late 1980s, see Goodman and Ittner (1992). In fact, economic history knows many similar examples of prolonged asset misvaluation, see, for example, Reinhart and Rogoff (2009) and Kindleberger and Aliber (2011). Many studies use surveys of inflation expectations. For example, Malmendier and Nagel (2009) study the responses in the Reuters/Michigan Survey of Consumers and find support for the backward looking, learning from experience model. Branch (2004) shows that the responses are consistent with a mixed model where nonrational expectations (such as naive or adaptive) have a high weight. A similar conclusion is reached in Nunes (2010) who uses, instead, the Survey of Professional Forecasters.

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

4. In this paper the one-variable LtF experiments are used as a test bed for the Genetic Algorithms (GA) model. Their relatively long duration of 50 periods and more is well suited to evaluate both short and long run performance of the GA model and compare it with other models. Recent LtF experiments in macroeconomics (Adam 2007; Assenza et al. 2014; Pfajfar and Žakelj 2014) investigate forecasting of *two* variables simultaneously (Duffy 2016). Experiments with repeated play of the “beauty contest”

Anufriev and Hommes (2012) and Anufriev et al. (2013b) show that an extension of the predictor-selection learning model of Brock and Hommes (1997)—called the *heuristics switching model (HSM)*—fits these different experimental settings quite well. Individuals choose from a small set of simple linear forecast rules (heuristics) that depend on a weighted average of past prices and a trend following component based upon their relative performance. However, a drawback to the HSM is that the modeler needs to specify the set of rules available for individuals to choose from. Moreover, the set of rules providing the best fit to the experimental data (or survey data, or DSGE model) may depend on the specifics of the economic model. So, although the HSM addresses some of the “wilderness of bounded rationality” criticism by endogenizing the distribution of agents across learning models, there is still the issue of how people come to select from this particular set of models, and why some models are used in some settings (like negative feedback experiments) and other models in other settings (like positive feedback experiments). Finally, a general model is needed to account for within-group heterogeneity in experimental data.

The current paper breaks new ground by further disciplining the wilderness of bounded rationality and proposing an evolutionary model, based on the genetic algorithm, that provides a general framework for expectation formation and the formulation of simple sets of forecasting models or heuristics. The GA learning model that we propose is *empirically validated and parsimonious*. Heterogeneous agents use linear first order price forecasting rule *with only two parameters*, that nests most of the existing ways the literature would model the formulation of expectations. The GA allows the agents to fine tune that set of forecasting rules by comparing past performance, experimenting, and refining the set to include rules likely to perform well in a mean square error sense. We will use three different experimental frameworks, initialize the genetic algorithm using the (estimated distribution of) initial outcomes from the experiment, and then run Monte Carlo simulations. A comparison is then made between the model-implied data and the data from the learning-to-forecast experiments. The results are striking. The GA model provides a very good fit to the experimental data and the same type of behavior—convergence in negative feedback models, oscillations in positive feedback—emerges. The model can also account for much of the cross-sectional heterogeneity. The ability of the GA model to generate a good fit to the results across a *variety* of lab experiments provides strong evidence in favor of the evolutionary predictor-selection model.

GAs are a prominent tool in the economic literature to model individual learning (see, e.g., Sargent 1993 and Dawid 1996). From the very first economic application in Arifovic (1994), GA were used to model both the social and individual learning and to explain the results of experiments with human subjects. Areas of GA applications include the overlapping generation monetary economies (Arifovic 1995), exchange

game (Nagel 1995) are also closely related. Despite somewhat different design and shorter duration, their outcomes are similar to the LtF experiments (Ho et al. 1998; Sutan and Willinger 2009; Sonnemans and Tuinstra 2010). Application of our GA model to these experiments is beyond the scope of this paper and left for future research.

rate volatility (Arifovic 1996; Lux and Schornstein 2005), production level choices in a cobweb producers economy (Dawid and Kopel 1998), financial markets (LeBaron et al. 1999; Duffy and Ünver 2006) and monetary policy (Arifovic et al. 2013). Recently in a related paper Hommes and Lux (2013) investigate a model in which agents use GA to optimize an AR1 forecasting heuristic (instead of directly optimizing a prediction) and, much like the actual subjects in the LtF experiments, cannot observe each other's behavior or strategies. The authors replicate the distribution (mean, variance and first order autocorrelation) of the predictions and prices of the cobweb experiments by Hommes et al. (2007).

We compare our GA model to 11 alternative learning models: seven homogeneous models: RE, naive, adaptive, trend-following and contrarian expectations, Least Square learning and Constant Gain learning, and four heterogeneous expectation models: a HSM with two rules, a HSM with four rules, and two different specifications of the GA, an action-based GA, and a GA based on the AR1 rule. This set of models includes the benchmark RE model, a standard set of homogeneous models that were popular in macroeconomics before the RE, and that constitute the building blocks for a recent successful HSM model, the two standard adaptive learning models as used in the literature discussed before, and different variations of GA models. Our GA model almost always outperforms these alternative benchmark models both in the short and in the long run. In particular, it substantially outperforms the action-based GA, where the action (i.e., the forecast) is optimized, and the AR1 based GA, where the two parameters of the AR1 rule are optimized. An important and novel insight from our analysis is that the use of an appropriate forecasting heuristic, that takes the trend-extrapolation in positive feedback systems into account, makes our heuristic-based GA “smart,” in the sense that it fits well with the observed behavior of human subjects in the experiments, cf. Gigerenzer and Todd (1999).

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

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

The paper is organized as follows. In Section 2 we present the set-up and findings of the LtF experiments. Section 3 presents some benchmark models of expectation formation. In Section 4 we introduce our GA model. Section 5 investigates how our model fits three different experimental settings. Finally, the concluding Section 6 gives an overview of the results and suggestions for future research. The appendices contain GA simulation details and various robustness checks.

2. Learning to Forecast Experiments

2.1. Overview

Learning to Forecast (LtF) Experiments are experimental markets where participants repeatedly forecast the market price and are rewarded for their *forecasting accuracy*. The underlying law of motion for realized prices, as a function of individual forecasts, is given by one of the following:

$$p_t = F(\bar{p}_t^e) + \varepsilon_t, \quad \text{or} \quad (1a)$$

$$p_t = F(\bar{p}_{t+1}^e) + \varepsilon_t, \quad (1b)$$

where p_t is the price at period t , $\bar{p}_t^e$ (or $\bar{p}_{t+1}^e$) denotes subjects' average forecast for the price, and ε_t are small IID shocks. In all experiments discussed in this paper there were 6 participants per group, so that $\bar{p}_t^e = (p_{1,t}^e + \dots + p_{6,t}^e)/6$. The price law of motion $F(\cdot)$ is obtained from a market clearing condition with aggregate supply and demand derived from optimal (i.e., profit/utility maximizing) choices of firms, consumers or investors, given the individual forecasts.⁵ In all examples in what follows there exists a unique price p^f such that $p^f = F(p^f)$, the fixed point of the expectations feedback system F . This price is called the *fundamental price* and when all subjects forecast it, the Rational Expectation steady state emerges.⁶ An important question for LtF experiments is whether the price will converge to its fundamental value.

The subjects in the LtF experiments are only informed about qualitative aspects of the market, but not about the exact law of motion. They know that their forecasts affect realized prices and whether the law of motion exhibits *negative feedback* (i.e., a higher average forecast results in a lower price) or *positive feedback* (i.e., a higher average forecast results in a higher price). Subjects do not know the exact number of other

5. The LtF experiments focus only on the forecasting behavior and abstract from other considerations (e.g., trading) by assuming that the subjects' actions are rational conditional on the submitted forecast. See Hommes (2011) for an in-depth discussion on methodology of the LtF experiments. See also Bao et al. (2013) and Bao et al. (2017) for experimental studies on how trading/production decision making is related to price forecasting.

6. In LtF experiments the subjects typically face an upper constraint on the price forecasts, which excludes exploding "rational bubbles" in asset pricing treatments.

participants and their forecasts, and are not explicitly informed about the fundamental price.⁷

The forecasts are submitted repeatedly and the experimental screen of a participant shows past prices, own past forecasts and earnings of the participant. The earnings per period decrease with the squared forecasting error, $(p_t - p_{i,t}^e)^2$. In the experiments with the law of motion (1a), subjects make *one-period ahead* forecasts, but in the experiments with the law of motion (1b), subjects forecast *two-periods ahead*, since when forecasting p_{t+1} the last observed price is p_{t-1} .

In this paper we consider three LtF experiments that differ by their law of motion.

HHST09 in Heemeijer et al. (2009): a linear law of motion (1a) with negative or positive feedback;

BHST12 in Bao et al. (2012): law of motion as in **HHST09**, with two large and unanticipated shocks to the fundamental price;

HSTV05 in Hommes et al. (2005): nonlinear positive feedback asset market with the two-period ahead law of motion (1b).

We explain the details of these experiments in what follows.

2.2. Experimental Data and “Stylized Facts”

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

$$\text{Negative feedback: } p_t = p^f - \frac{20}{21} (\bar{p}_t^e - p^f) + \varepsilon_t, \quad (2)$$

$$\text{Positive feedback: } p_t = p^f + \frac{20}{21} (\bar{p}_t^e - p^f) + \varepsilon_t, \quad (3)$$

where $\bar{p}_t^e = \sum_{i=1}^I p_{i,t}^e / I$ is the average prediction of all individuals at period t and p^f is the unique constant RE solution of the price dynamics, which we will refer to as the *fundamental price*. The experiment ran for 50 periods with 13 groups of $I = 6$ subjects with the same realization of the shocks ε_t drawn independently from a normal distribution $N(0, 0.25)$. The two treatments are symmetrically opposite, with the same

7. The fundamental price can sometimes be inferred from the experimental instructions. For example, in the asset pricing experiment **HSTV05**, the fundamental price is equal to the present value of future dividends, which is the ratio of the average dividend to the interest rate. Both variables were provided to the subjects, but most of the individual (first period) forecasts were not at the fundamental level.

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

fundamental price $p^f = 60$, and dampening factors of the same absolute value, but with opposite signs.⁸ Under homogeneous naive expectations (i.e., $\bar{p}_t^e = p_{t-1}$) the fundamental price for both treatments is a stable steady state of the price dynamics.

The aggregate price dynamics in the two feedback treatments of the experiment were very different, see Figures 1(a) and (b) for typical examples (the GA simulations in the two lower panels are explained in Section 5). Under negative feedback after a short volatile phase of 6–7 periods, the price converged to the fundamental value $p^f = 60$, after which the subjects' forecasts coordinated on the fundamental price as well. In most of the positive feedback groups⁹ persistent price oscillations arose where the price over- and undershoots p^f . In spite of the price oscillations the subjects' forecasts became very close to each other after only 2–3 periods and remained so until the end of the experiment. In positive feedback markets subjects' forecasts are thus strongly coordinated, but on a nonfundamental price. The almost self-fulfilling character of the

8. In an asset pricing market, the near unit root coefficient 20/21 arises from a realistic discount factor $1/(1+r)$ with interest rate $r = 5\%$. To have symmetric treatments, the factor in the negative feedback was set to $-20/21$.

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

FIGURE 2. **BHST12**: experimental groups (*upper panels*) and 65-period ahead simulations of GA-P1 model with random initial predictions (*lower panels*). Black thick line shows the price, green dashed lines show 6 individual predictions, thin line shows the fundamental price.

near-unit root positive feedback is a key feature of the system that allows subjects to coordinate on trend following behavior, which results in price oscillations (Hommes 2013).

BHST12 report an LtF experiment with the same structure as **HHST09**, with a positive and a negative feedback treatment, based on linear price equations (2) and (3) with dampening factors $\pm 20/21$. There are $I = 6$ participants in every group, and the experiment runs for 65 periods. The key difference in this experiment is that there are two large, permanent and unanticipated shocks to the fundamental price: the fundamental price 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 $t = 65$. Typical time paths of **BHST12** are shown in the upper panels of Figure 2 (two lower panels are explained in Section 5). Under negative feedback (Figure 2a), a shock to the fundamental breaks the subjects' coordination and is followed by quick convergence to the new fundamental price. In contrast, under positive feedback (Figure 2b), shocks leave the coordination intact, and the predictions and prices move smoothly toward the new fundamental, eventually over- or undershooting it.

HSTV05 report an experiment based on a *two-period ahead* nonlinear positive feedback market asset-pricing model with law of motion (1b) specified by

$$p_t = F(\bar{p}_{i,t+1}^e) = \frac{1}{1+r} \left(n_t p^f + (1 - n_t) \bar{p}_{i,t+1}^e + y \right) + \varepsilon_t, \quad (4)$$

FIGURE 3. **HSTV05**: experimental groups (upper panels) and 50-period ahead simulations of GA-P3 model for different seeds giving different initial predictions and learning (lower panels). Black thick line shows the price, green dashed lines show 6 individual predictions.

where $r = 5\%$ is the interest rate, y is the mean dividend, and n_t is the share of computerized, robotic forecasters predicting the fundamental price p^f . This fraction changes endogenously as¹⁰

$$n_t = 1 - \exp\left(-\left|p_{t-1} - p^f\right|/200\right). \quad (5)$$

The most important feature of this market is that, differently from **HHST09** and **BHST12**, the *current* price depends on the average of the subjects' expectations about the price in the *next* period. Thus participants in this experiment had to predict p_{t+1} without knowing p_t . Moreover, the presence of robotic forecasters makes the feedback between forecasts and the price, $F(\cdot)$, highly nonlinear. There were two treatments with different values of mean dividend, and thus different fundamental price: in seven markets the mean dividend was set to $y = 3$ leading to $p^f = 60$, and in three markets the mean dividend was set to $y = 2$, leading to $p^f = 40$.

The two upper panels of Figure 3 illustrate the results of **HSTV05**. In fact, in this experiment three different aggregate outcomes were observed: (i) monotonic convergence to the fundamental price (in 2 groups; see Figure 3a for an example),

10. In this market, with constant interest rate r and mean dividend y , the unique constant RE solution, that is, fundamental price, coincides with the discounted value of future dividends, $p^f = y/r$. The participants were not explicitly informed about the fundamental price, though they knew values of y and r . In the field setting, so-called *fundamental* traders have a better understanding of the underlying process and often stabilize markets, especially when the price deviation from the fundamental level is large. Robotic forecasters were introduced in this experiment to mimic those fundamentalists from real financial markets.

(ii) dampened oscillations (in 3 groups), and (iii) persistent price oscillations (in 5 groups; see Figure 3b for an example).¹¹

To summarize, the LtF experiments revealed the following characteristic “stylized fact”:

- S1. Price dynamics under negative feedback is more stable than under positive feedback;
- S2. There are generally different types of price dynamics under positive feedback: monotonic convergence, dampened oscillations, and permanent oscillations of prices were observed;
- S3. The expectations of 6 different participants in the same group were typically coordinated even though participants could not observe the forecasts of others.

HHST09 described the *subjects’ forecasting behavior* in the experiment with the first-order rule FOR:

$$p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + \alpha_3 p^f + \beta(p_{t-1} - p_{t-2}), \quad (6)$$

for $\alpha_1, \alpha_2, \alpha_3 \geq 0$, $\alpha_1 + \alpha_2 + \alpha_3 = 1$, $\beta \in [-1, 1]$. Rule (6) is an *anchor and adjustment rule* extrapolating a price change from an anchor given by a weighted average of the previous price, the previous forecast and the fundamental price p^f .¹²

HHST09 estimated the FOR separately for each subject, fitting well the forecasting behavior of around 60% of all individuals. In Section 4 we will use these findings to motivate the specification of our GA model.

HHST09 found significant variability in terms of individual forecasting, within the same treatment, but even more so between the two treatments. The main difference lies in the trend extrapolation, which is popular under positive feedback (i.e., $\beta > 0$), but disregarded under negative feedback (i.e., $\beta \approx 0$). **BHST12** and **HSTV05** report qualitatively similar results.

3. Models of Expectation Formation

Before specifying our Genetic Algorithm model, we discuss a number of alternative benchmark models of expectation formation. In particular, we consider Rational Expectations, simple adaptive learning (LS learning), a number of simple linear heuristics (naive, adaptive, trend following, and contrarian expectations) and a Heuristic Switching Model. For technical details on these models, we refer to Appendix A.

The *Rational Expectations* model (RE) for the LtF experiments implies that the subjects should predict the fundamental price, $p_{i,t}^e = p^f$, the fixed point of the

11. Bao et al. (2017) run a similar asset market experiment to compare “Learning to Forecast” versus “Learning to Optimize” designs. They report three similar types of price behavior.

12. RE is nested as a special case of the FOR in (6), with $\alpha_1 = \alpha_2 = \beta = 0$ and $\alpha_3 = 1$, so that the forecast reduces to the fundamental price, $p_{i,t}^e = p^f$.

experimental feedback map $F(\cdot)$. Naturally, one cannot expect the subjects' forecasts to immediately converge to the fundamental value. Nevertheless, under negative feedback we typically observe that the subjects learn the RE equilibrium within a few periods (see Figures 1a and 2a). On the other hand, RE contradicts subject heterogeneity and price oscillations under positive feedback treatments, as observed in all three LtF experiments discussed previously (see Figures 1b, 2b, and 3b).

Adaptive learning offers a less restrictive alternative to RE. The broad idea of adaptive learning is that agents learn to forecast prices by means of some statistical inference, for example, by econometrically fitting some forecasting rule. As a benchmark, we will focus on *LS learning*, where agents use a perceived law of motion (PLM)

$$p_t^e = \alpha_t^{LS} p_{t-1} + (1 - \alpha_t^{LS}) p_{i,t-1}^e + \beta_t^{LS} (p_{t-1} - p_{t-2}), \quad (7)$$

where α_t^{LS} and β_t^{LS} are found by simple OLS inference on the regression equation

$$(p_s - p_{s-1}^e) = (\alpha \quad \beta) \begin{pmatrix} p_{s-1} - p_{s-1}^e \\ p_{s-1} - p_{s-2} \end{pmatrix} + u_s \quad (8)$$

in every period t , based on the full information set in that period (i.e., data from all periods $s < t$). We focus on rule (7), because it is the same first-order forecasting heuristic (10), which will form the core of our GA model (see Section 4). Table A.1 in Appendix A provides initialization details on LS learning model. LS learning gives equal weight to all past observations and, as time goes to infinity, the weight given to the last observation tends to 0. It may therefore be more realistic to give more weight to recent observations. Constant Gain (CG) learning is a modification of LS learning that takes this into account and gives more weight to more recent observations. Appendix A describes the LS and CG learning algorithms in more detail.

The basic class of nonrational models of expectation formation is a collection of simple homogeneous linear forecasting heuristics. In this paper, we focus on:

$$\begin{aligned} \text{naïve expectations:} \quad & p_{i,t}^e = p_{i,t-1}, \\ \text{adaptive expectations:} \quad & p_{i,t}^e = \alpha p_{i,t-1} + (1 - \alpha) p_{i,t-1}^e \\ & \text{with } \alpha = 0.75, \\ \text{trend following rule:} \quad & p_{i,t}^e = p_{i,t-1} + \beta (p_{t-1} - p_{t-2}) \\ & \text{with } \beta = 1, \\ \text{contrarian expectations:} \quad & p_{i,t}^e = p_{i,t-1} + \beta (p_{t-1} - p_{t-2}) \\ & \text{with } \beta = -0.5. \end{aligned} \quad (9)$$

Notice that all these heuristics are special cases of the first-order rule (6). The main drawback of any of these heuristics is that each disregards subject heterogeneity and cannot explain differences between experimental treatments (we will show this formally in later analysis).

Another issue with any homogeneous forecasting rule is that it underestimates the sophistication of the experimental subjects. As explained before, **HHST09** report that subjects are much more likely to forecast price trends under positive feedback, where trends occurred frequently. This indicates that subjects are *smart* and try to adapt their behavior to the particular experimental economy. These findings led Anufriev et al. (2013b) to investigate the Heuristic Switching Model (HSM), in which the subjects are endowed with two prediction heuristics (henceforth 2-type HSM), namely *adaptive expectations* and *trend extrapolation* as specified in (9). The idea of the HSM model is that the subjects can at any time use any of the two heuristics, but tend to focus on the rule with a higher relative past performance (see Anufriev et al. 2013b, for a technical description). The dynamics of the HSM are similar to the experimental outcome and have an intuitive behavioral interpretation. Under positive feedback agents quickly coordinate on the trend extrapolation heuristic, leading to persistent price oscillations and thus self-confirming trend chasing forecasting. In contrast, under negative feedback the trend extrapolation rule performs poorly and agents switch to adaptive expectations, which eventually causes the price to converge to the fundamental price.

An extended HSM with 4 heuristics (henceforth 4-type HSM) was used by Anufriev and Hommes (2012) to explain the experimental data **HSTV05**. These four heuristics contain (i) a so-called *anchoring and adjustment* rule, given by

$$p_{i,t+1}^e = 0.5(p_{t-1} + p_{t-1}^{av}) + (p_{t-1} - p_{t-2}) \text{ where } p_{t-1}^{av} = (1/(t-1)) \sum_{s=1}^{t-1} p_s$$

is the average price so far; and three rules similar to the 2-type HSM specification, namely (ii) adaptive expectations (with $\alpha = 0.65$), (iii) weak and (iv) strong versions of the trend following rule (with $\beta = 0.4$ and $\beta = 1.3$ respectively). Anufriev and Hommes (2012) demonstrate the impact of the trend coefficient in the trend following heuristic on the amplitude of price oscillations.

The HSM captures the essence of the aggregate forecasting behavior in the LtF experiment by successfully replicating the results of **HSTV05** and of the two treatments of **HHST09**. It is a stylized model, however, and leaves open two important questions about the origins of the forecasting heuristics: (1) where do these particular rules (with these coefficients) come from? and (2) which rules (and how many of them) should be used in a more general setting? Moreover, the HSM cannot fully account for the within-treatment heterogeneity of predictions and hence does not fully explain the experiment at the individual level. To overcome these drawbacks, we introduce a model with explicit individual heuristic-learning through Genetic Algorithms.

4. The Genetic Algorithm Model

Genetic Algorithms (GA) form a class of numerical stochastic maximization procedures that mimic the evolutionary operations with which DNA of biological organisms adapts to the environment. GA were invented by John Holland in the 1960s

to study the mechanism of adaptation and have since then been used to solve “hard” optimization problems, which may involve noncontinuities or high dimensionality with complicated interrelations between the arguments. GA are flexible and efficient and so are often used in computer sciences and engineering. See Holland (1975) for the original introduction of GA, Haupt and Haupt (2004) for technical discussion, De Jong (2006) and Mitchell (1996) for an overview of their use and Dawid (1996) for economic applications. We note that the ability of GA-models to replicate human behavior in laboratory experiments does not prove the ability of reproducing their deliberate process. The LtF experiments abstract from all economic dimensions of the problems that the agents have to solve in general, since they reduce the economy to an aggregate forecasting feed-back environment. We rather view GA-models as a convenient statistical way to represent an evolutionary selection process that, as we will see, fits the LtF experimental data rather well.

4.1. Core Algorithm

A GA routine starts with a population of random trial solutions to the problem. Individual trial arguments are encoded as binary strings (strings of ones and zeros), so-called “chromosomes”.¹³ They are retained into the next iteration with a probability that increases with their relative performance (or “fitness”). This so-called *reproduction* operator means that with each iteration, the trial arguments are likely to have a higher performance, that is, be “fitter”. On top of reproduction, GAs use three operators that allow for an efficient search through the problem space: mutation, crossover and election, where the last operator was introduced in the economic literature in Arifovic (1995).

Mutation. At each iteration, every bit in each chromosome has a small probability to mutate, in which case it changes its value from zero to one or 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 (small) probability a mutation of a bit at the beginning of the chromosome can occur, which changes the argument drastically. Mutation is thus a form of parameter experimentation, enabling the GA to search through the whole parameter space and have a good chance of shifting from a local maximum toward 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 often set to exchange subsets of the arguments. For example, if the objective function has two

13. We use a binary representation, as in a majority of economic GA applications, for the sake of parsimony. The real number variant of the GA requires additional parametrization, such as a distribution of the mutation changes. An attractive feature of the binary representation is that agents’ experimentation will include both local and global changes.

arguments, crossover would swap the first argument between pairs of trial arguments. Crossover is thus a form of experimentation with the set of heuristics.

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

These four operators have a straightforward economic interpretation for a situation in which agents optimize their behavioral rules such as forecasting heuristics. The reproduction 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 people disregard unsuccessful experimentation.¹⁴

4.2. Model Specification

GAs can be used to model both individual and social learning, see Arifovic (1994) and Vriend (2000) for comparison. Under social learning different agents would be able to imitate the actions of others. However, in the LtF experiments the subjects did not have access to the predictions and performances of other subjects. They could learn only individually by experimenting and evaluating own ideas, and for this reason, we use a GA implementation of individual learning.

We populate the price-expectation feedback economies from Section 2 by $I = 6$ GA agents. For now we focus on the one-period ahead LtF economy (1a), and will extend our framework to the two-period ahead LtF design (1b) in Section 5.

At the beginning of each period t an agent i submits a forecast $p_{i,t}^e$, and learns to do so with a GA. The GA can be used to directly optimize the forecast (action-based GA) or alternatively, to optimize the parametrization of some forecasting heuristic (heuristic-based GA). In the latter case, one has to further specify the heuristic of the GA, which can have important implications on the model dynamics.

In our model every agent i uses one of *her own* $H = 20$ parametrizations¹⁵ of a linear first-order forecasting heuristic. The forecast $p_{i,h,t}^e$ of agent i with heuristic h at time t is given by

$$p_{i,h,t}^e = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p_{i,t-1}^e + \beta_{i,h,t} (p_{t-1} - p_{t-2}), \quad (10)$$

where $p_{i,t-1}^e$ denotes the prediction of the price p_{t-1} submitted by agent i in period $t - 1$. The forecasting heuristic (10) is a simplified version of the general FOR (6) as

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

15. Simulations with $H = 10$ and $H = 100$ yield similar results, see Appendix B.

estimated in **HHST09** on individual data.¹⁶ In Section 4.3 we will introduce alternative specifications of the GA, which were suggested in the literature (such as the action-based GA). We will demonstrate that our heuristic-based GA works particularly well. *It is the heuristic (10) that makes our GA agents as smart as human subjects in the experiment.*

Heuristic (10) depends on two parameters, $\alpha_{i,h,t}$ (price weight) and $\beta_{i,h,t}$ (trend extrapolation coefficient), and the H specifications only differ in these coefficients. Importantly, these parameters are time varying, as the agents repeatedly fine-tune the rule to adapt to the specific market conditions. For example, in an asset pricing market it may pay off to extrapolate the price trend and agents would try to find the optimal value of β , depending on the most recent trend. This learning is embodied as a heuristic optimization with the GA procedure, and introduces individual heterogeneity into the model that is absent in any homogeneous expectation model, and in the HSM where each individual is confined to the same fixed set of a few heuristics.

Define $\mathbf{H}_{i,t}$ as the set of H heuristics of agent i at time t , where heuristic h is specified as a pair of parameters $(\alpha_{i,h,t}, \beta_{i,h,t}) \in \mathbf{H}_{i,t}$. Each pair is a “chromosome” represented as a binary string of length 40 with 20 bits per coefficient. The bounds for the coefficients are chosen as follows. From a behavioral viewpoint, the price weight $\alpha_{i,h,t}$ must belong to the unit interval $[0, 1]$. For the trend extrapolation coefficient we report two parametrizations, depending on the bounds. Under Parametrization 1 (denoted as *GA-P1*), the restriction is symmetric, $\beta_{i,h,t} \in [-1.1, 1.1]$. Under Parametrization 2 (denoted as *GA-P2*), the restriction is $\beta_{i,h,t} \in [0, 1.1]$, that is, contrarian rules are not allowed.¹⁷

The heuristics are updated *independently* for each agent by one iteration of the GA operators, see Table 1 for the specific parameter values.¹⁸ The updating is based on the relative forecasting performances of the heuristics. The experimental payoffs decrease with the squared error (SE) of the forecast. Accordingly, at time t for every heuristic from $\mathbf{H}_{i,t}$ we compute the (hypothetical) squared error, $SE_{i,h,t} = (p_{i,h,t}^e - p_t)^2$, and apply the logit transformation¹⁹ to define the normalized performance (“fitness”) of

16. In comparison with the estimated FOR (6), in rule (10) the coefficient in front of the fundamental price (which can be thought of as an anchor) is set to 0. We experimented with the general FOR (6) with the anchor specified as either (i) the fundamental price p^f or (ii) the average realized price so far. Neither specification could closely match the experimental dynamics of the positive feedback treatment, where the anchor dampens the oscillations, see Online Appendix F. This is consistent with the fact that in the estimated rules of the experimental data in **HHST09** under positive feedback, the anchor weight α_3 in (6) is typically insignificant.

17. Heuristics with negative extrapolation coefficient are often called contrarian strategies. **HHST09** found only two subjects with such contrarian rules, but for the sake of completeness we report both parametrizations. The estimated positive trend coefficients in **HHST09** range from 0.3 to 1.

18. Extensive simulation runs show that the model is robust to small changes in the parametrization. Only when the number of heuristics falls significantly below $H = 10$, the GA model loses empirical fit.

19. We use the logit and not the power transformation as in Hommes and Lux (2013) to have a clear link with the literature on HSMs.

TABLE 1. Values of parameters used by the Genetic Algorithms agents.

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

heuristic h that agent i uses in individual learning:

$$\Pi_{i,h,t} = \frac{\exp(-SE_{i,h,t})}{\sum_{k=1}^H \exp(-SE_{i,k,t})}. \quad (11)$$

Before the market starts to operate, the GA model and the set $H_{i,1}$ of agents' heuristics have to be initialized. In particular three aspects of the initialization are relevant:

1. *Initialization of heuristics:* At period $t = 0$, every agent samples 20 random heuristics.²⁰
2. *Uninformed initial forecast:* At period $t = 1$, with no past prices and forecasts, the heuristics' forecasts are yet undefined. To highlight different aspects of our GA model, we will either draw forecasts $p_{i,1}^e$ from an *exogenous* distribution (Section 5.1.1), or use the forecasts submitted in the experiment (Sections 5.1.2 and 5.2). Price p_1 is then realized according to (1a).
3. *Initial use of heuristics:* At period $t = 2$ (when the agents observe the first realized price p_1 and their first forecasts $p_{i,1}^e$)²¹ the heuristics can already be used for forecasting, but their performances are still undefined. In this period, every agent randomly picks one of her own heuristics with equal probabilities.

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

20. In particular, every agent samples 800 bits (20 heuristics with 2 parameters, each encoded with 20 bits) as 0 or 1 with equal probability. Note that the range of parameters will then affect the forecasts in the initial periods.

21. The agents still do not observe a price trend $\Delta p_1 = p_1 - p_0$, since p_0 is undefined. We assume no initial trend with $\Delta p_1 = 0$, as if agents were forced to use $\beta_{i,h,1} = 0$ within this period.

1. Agents forecast price; the market price p_t is realized according to (1a); agents observe it;
2. Agents *independently* update their heuristics using one GA iteration. The criterion function $\Pi_{i,h,t}$ in (11) is computed from the *hypothetical* SE's of all different heuristics in predicting the price p_t . To be specific, agent i uses four operators:
 - a *reproduction*: agent samples H so-called 'child' heuristics from the pool of 'parent' heuristics $H_{i,t}$, with replacement, using $\Pi_{i,h,t}$ as the corresponding probabilities;
 - b *mutation*: each bit of each child heuristic has probability $\delta_m = 0.01$ to switch its value;
 - c *crossover*: each pair of child heuristics has probability $\delta_c = 0.9$ to swap the last twenty bits (it corresponds to exchanging β 's);
 - d *election*: each child heuristic (possibly modified after mutation and crossover) is compared in terms of SE with a randomly chosen parent heuristic. The child joins $H_{i,t+1}$ if it strictly outperforms the parent. Otherwise, the parent is passed to $H_{i,t+1}$.
3. Now, when the new sets $H_{i,t+1}$ are formed, period $t + 1$ starts.²²
4. With probabilities as in (11), but now based on the hypothetical SE's of heuristics from the new pool, each agent i randomly picks one heuristic from $H_{i,t+1}$. Each agent uses this heuristic to generate her forecast $p_{i,t+1}^e$. The algorithm now returns to step 1.

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

4.3. Alternative GA Specifications

The success of the GA model relies to a large extent on a smart choice of the forecasting heuristic of the GA agents. To emphasize this point, we compare our GA model with two alternative specifications from the literature: an action-based GA and a GA with an AR1 heuristic. These are simple GA benchmarks, but unlike the GA-P1 and GA-P2 specifications, they have no direct empirical motivation based on laboratory data. In the next section it will become apparent that these two alternative GA models result in

22. All the simulations of our GA model are based on different initial predictions and learning realizations, but the supply shocks ε_t are the same within each treatment.

a worse fit to the data.²³ Note that the timing of the two alternative GA models is the same as for our GA-P1 and GA-P2.

An *action-based GA* does not require a heuristic such as the FOR rule (10). Instead, in this specification each agent directly optimizes the forecast through GAs. In particular, agent i has $H = 20$ chromosomes, where each encodes exactly one argument from the forecasting range $[0, 100]$, and which is then used directly as the (constant) forecast. Such an *action-based GA* does *not* take the experimental trend-following behavior into account, and as we will see in what follows, is among the worst performing models.

An *ARI GA* was used by Hommes and Lux (2013). This model is similar to our GA model, with the exception that here agents optimize not the FOR heuristic (10), but an AR1 rule

$$p_{i,t}^e = \mu + \rho (p_{t-1} - \mu), \quad (12)$$

where $\mu \in [0, 100]$ and $\rho \in [-1, 1]$. This heuristic, like FOR, has two parameters: the long-run mean μ and the mean-reversion coefficient ρ . Like the action-based GA, the AR1 GA does not take trend-following into account, but rather learns the mean-reversion parameter ρ .

5. Empirical Validity of the GA Model

In order to test the goodness of fit of our GA model to the experimental data, we focus on two types of evidence: long run out-of-sample 50-period ahead predictions in Section 5.1, and short run out-of-sample one-period ahead predictions in Section 5.2. We consider the fitness of the model to the price dynamics (aggregate behavior), but also to the *subject heterogeneity and individual forecasts*. We first present jointly the results for **HHST09** and **BHST12**, since they share a similar interpretation. Section 5.3 will focus on the experiment in **HSTV05**, in which our GA model will yield more complex dynamics, consistent with the experimental data.

5.1. 50-Period Ahead Simulations for Linear LtF Economies

The first empirical test for the fit of our GA model to the experimental data are 50- and 65-period ahead simulations for the **HHST09** and **BHST12** experiments respectively.²⁴ In these simulations we only use experimental data from the first period and we run two

23. As mentioned earlier, a third alternative GA model is one with the general FOR rule (6) with an additional coefficient for the anchor. This extended GA did not perform well in any of the empirical tests. We also checked a GA model based on a hybrid rule that nested the FOR heuristic (10) and an additional “action” term $\gamma_{i,h,t} \in [0, 100]$ (as in the action-based GA). Again this specification had a poor fit to the data.

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

types of Monte Carlo (MC) simulations that differ in how these data are used for model initialization. The GA model is then simulated for the duration of the experimental sessions with *no other information* from the experiment. In this way, we test if our GA model is able to capture the *long-run* dynamics observed in the experiment, that is, to reproduce the “stylized facts”.

5.1.1. Long-Run Average Group Behavior. In the first MC exercise we look at “unconditional” dynamics by starting each new simulation run with *new sampling of the initial predictions* from an exogenous distribution. We calibrated this distribution separately for the two experiments, from the set of all forecasts submitted in the first period of the experiment, see Appendix C for details.²⁵ To compute prices, the laws of motion (2) and (3) are applied for negative and positive feedback simulations, respectively, from either experiment. Recall that the difference between the two experiments lies only in the level of the fundamental price, which is constant in **HHST09** and changes twice by large unanticipated shocks in **BHST12**. For each treatment, we run 1000 simulations that differ by the sampled initial predictions and by the random realizations in all steps of the GA learning algorithm.

Figure 4 (upper panels) show the median price of 1000 GA simulations (solid line) and 95% confidence intervals (CI, dotted lines) for the two treatments from **HHST09**, for the GA-P1 parametrization of the model. To compare these simulations with the data, we superimpose the experimental data on this figure. The pluses “+” represent prices in 6 different groups in the treatment²⁶ and the green dashed line shows the median of those data points. See also Figure D.1 in Appendix D for a comparison of the experimental and GA model (GA-P1 and GA-P2) forecasts in period 49, for both types of feedback. The lower panels of Figure 4 illustrate the degree of coordination among agents by showing the dynamics of the standard deviation of six individual forecasts (in $\log_2$ scale). For examples of model dynamics in the **HHST09** environment, we also show two typical simulations of the GA model in the lower panels of Figure 1. It is striking that these simulations are almost identical to the experimental data shown in the two upper panels.²⁷

The same set of figures of 65-period ahead unconditional simulations of the GA-P1 model for **BHST12** can be found in Figure 5 (with two additional middle panels for the distance of the experimental and the GA model prices from the fundamental price). Two typical sample simulations for this experiment are exhibited in Figures 2(c) and (d).

25. In the first period the subjects in the LtF experiments have limited, mostly qualitative information about the market and do not see any price history. Their initial forecasts are necessarily more a matter of an initial guess than a reasoned forecast. Thus, we treat initial forecasts as coming from an exogenous distribution, as discussed in the model initialization in Section 4.2 (see also Diks and Makarewicz 2013 for a comprehensive discussion).

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

27. Simulations presented in Figures 1(c) and (d) were among the first that we ran.

FIGURE 4. **HHST09**: 50-period ahead “unconditional” simulation (1000 runs) of GA-P1 model compared with the experimental data. *Upper panels*: dynamics of price. *Lower panels*: dynamics of standard deviation of 6 forecasts (in $\log_2$ scale). Black pluses denote the group observations in the experiment, green dashed line shows the median. Red thick line shows the median simulation; blue dotted lines show the 95% confidence interval for the GA model.

Figures 4 and 5 show that the GA model replicates well the outcomes of the linear LtF experiments. Under negative feedback of the **HHST09** experiment (left panels of Figure 4), prices quickly converge close to the fundamental, but individual heterogeneity of GA agents declines only slowly and is visible until period 15, consistent with the experimental data. Under positive feedback, GA agents coordinate their forecasts in less than five periods, but the distribution of realized prices does not collapse into the fundamental even after 50 periods, when the 95% CI of prices is as wide as [55,70]. The median price resembles the experimental oscillations, including the typical amplitude and turning points. Large shocks to the fundamental (**BHST12** experiment in Figure 5) exaggerate these dynamics. Under negative feedback, both GA agents and the experimental subjects quickly converge to the new fundamental prices, whereas under positive feedback fundamental shocks cause smooth, high amplitude oscillations both in the experiment and in the GA model (as shown by the oscillating 95% CI).

Overall, the 95% CI for the GA-P1 model captures 65% (resp. 81%) of the experimental prices and 81% (resp. 72%) of the degree of coordination for the negative (resp. positive) feedback treatment in the **HHST09** experiment. For simulations for the **BHST12** experiment these numbers are 66% (resp. 84%) of the prices and 84% (resp. 67%) of the individual coordination under negative (resp. positive) feedback.

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

In other words, we are able to replicate roughly 75% of the *long-run* behavior of the experimental groups in the two linear LtF experiments, both at the *aggregate* and *individual* levels.

Which heuristics were learned by our GA agents in the **HHST09** environment? Figure 6 reports the median and the mean (with 95% and 90% CI) for the MC simulations of the price weight α and the trend extrapolation coefficient β , which were selected by the GA agents (6 pairs of α and β in each period in one simulation, with the same set of 1000 simulations as in Figure 4). A first observation is that large heterogeneity of individual rules persists, consistent with the estimated rules in **HHST09**. Secondly, there are clear differences between the two treatments. Under

FIGURE 6. **HHST09**: Emerging heuristics in 50-period ahead MC simulation (1000 runs) for GA-P1 model. The price weight α (*upper panels*) and the trend extrapolation coefficient β (*lower panels*) of the chosen heuristic are shown. Red thick line is the median, black dot-dashed line is the mean, blue dotted and purple dashed lines show the 95% and 90% confidence intervals, respectively, for the GA model. The green thick dashed line in panel (d) represents the 28% percentile of chosen β .

positive feedback the median GA agent²⁸ quickly converges toward an approximate rule

$$p_{i,t+1}^e \approx 0.9p_t + 0.1p_{i,t}^e + 0.6(p_t - p_{t-1}). \quad (13)$$

This median rule is close to a pure trend-following rule (i.e., with anchor p_t), but has a coefficient $\beta \approx 0.6$, smaller than $\beta = 1$ as used in the 2-type HSM of Anufriev et al. (2013b). Furthermore, 72% of the GA agents never used negative β in the last 30 periods (see the green thick dashed line in Figure 6d for 28% percentile). For the distribution of β in period 50, see Figure 9(a). On the other hand, under negative feedback, the median GA agent learns a rule close to

$$p_{i,t+1}^e \approx 0.5p_t + 0.5p_{i,t}^e \quad (14)$$

with median trend coefficient β close to 0. Thus the median rule under negative feedback is adaptive expectations with price weight of 0.5; Anufriev et al. (2013b) used adaptive expectations with coefficient 0.75 on the price in their 2-type HSM.

28. Median rules of the GA agents in all treatments are similar to average rules, see dashed black lines in Figures 6 and 7 for **HHST09** and **BHST12**, respectively.

FIGURE 7. **BHST12**: Emerging heuristics in 65-period ahead MC simulation (1000 runs) for GA-P1 model. The price weight α (*upper panels*) and the trend extrapolation coefficient β (*lower panels*) of the chosen heuristic are shown. Red thick line is the median, black dot-dashed line is the mean, blue dotted and purple dashed lines show the 95% and 90% confidence intervals, respectively, for the GA model.

Our learning dynamics therefore confirm the results by **HHST09** and yield empirical support for the 2-type HSM by Anufriev et al. (2013b), albeit with slightly different parametrizations.

Similar results are found in the **BHST12** environment, as seen in Figure 7. In fact, under negative feedback, the median GA agent learns the same adaptive expectations rule as before, $p_{i,t+1}^e \approx 0.5p_t + 0.5p_{i,t}^e$. Under positive feedback, the median GA agent converges to the heuristic

$$p_{i,t+1}^e \approx 0.95p_t + 0.05p_{i,t}^e + 0.9(p_t - p_{t-1}), \quad (15)$$

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 (13) used by the median GA agent under the positive feedback from the **HHST09** experiment. The 95% CI for the trend extrapolation coefficient β becomes significantly positive toward the end of the experiment (see also Figure 9b 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 stronger trend followers.

Figures 6 and 7 suggest that convergence to the mean and median parameters is relatively fast, within the first 10–15 periods. To check that learning in the GA model occurs very quickly, we performed two simulation experiments. In the first experiment

we switched off the GA evolution at time $t = 15$. Specifically, after that time we skipped step 2 of the GA algorithm described in Section 4.2, so that the set of rules for each GA agent did not evolve after period 15. The goodness of fit of this model turned out to be comparable to the benchmark GA simulations, showing that learning of the set of rules indeed has taken place within first 15 periods.²⁹ In the second experiment we switched off the GA model at the outset, fixing the set of rules by random draws of the coefficients around the median. This *noGA*-model performs worse. These two simulation experiments show that GA learning is important in the first 15 periods, but not thereafter for the **HHST09** experiment.³⁰

We also checked what happens when replacing the heterogeneous rules with the median forecasting type. The dynamics with such a median representative agent is close to the median trajectory of the GA simulations. However, the median forecasting model does not match the heterogeneity across different groups, which is well explained by the GA model, as we will see in Section 5.1.2.

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

$$p_{i,t}^{e,GA} = \sum_{h=1}^{H=20} \Pi_{i,h,t} p_{i,h,t}^e, \quad (16)$$

which is the average of the price forecasts given by the twenty different heuristics weighted by their fitness, $\Pi_{i,h,t}$, as defined in (11). The quantity $p_{i,t}^{e,GA}$ in (16) is simply the *model prediction* of the subject i 's price forecast at time t .³¹ Using this

29. In the next sections we introduce the goodness-of-fit measures. Tables **B.1** and **B.2** report the results for these two simulation experiments under labels *GA: 15Learn* and *GA: noGA* (with two specifications), respectively.

30. The same holds for the **BHST12** data. However, these linear one-period ahead feedback environments are particularly simple. The results for the **HSTV05** data indicate that in a more complex nonlinear and two-period ahead feedback environment, GA agents do not have to remain locked in one specific set of rules after the initial learning phase. Instead, they can switch between some “local equilibria” (e.g., following price oscillations versus staying close to the fundamental steady state), if given enough time, and each shift between “local equilibria” is accompanied by another set of rules; see Section 5.5 and Figure 10 for more details.

31. Notice that within the simulation, the realized price itself *does not* depend on the set of mean forecasts (16) of the GA agents. Instead, as explained in Section 4.2, every GA agents samples one from her $H = 20$ heuristics to generate her forecast, which is then used in the price law of motion.

model prediction and the price trajectory generated by the GA model, p_t^{GA} , we compute the mean squared error (MSE) in predicting the experimental data (both prices and individual price forecasts), excluding the initialization phase of the first three periods, as follows:

$$\text{MSE}_X^{\text{prices}} = \frac{1}{T-3} \sum_{t=4}^T \left(p_t^{\text{GrX}} - p_t^{\text{GA}} \right)^2, \quad (17)$$

$$\text{MSE}_X^{\text{forecasts}} = \frac{1}{6 \times (T-3)} \sum_{i=1}^6 \sum_{t=4}^T \left(p_{i,t}^{e,\text{GrX}} - p_{i,t}^{e,\text{GA}} \right)^2, \quad (18)$$

where p_t^{GrX} and $p_{i,t}^{e,\text{GrX}}$ denote period t price and forecast of subject i in the experiment, and $T = 50$ ($T = 65$) for groups from the **HHST09** (**BHST12**) experiment.

The third statistic of interest measures the degree to which our GA model predicts subject heterogeneity, that is the dispersion of individual forecasts in a given period in a given experimental group. Denote the average forecast of a certain group (group of experimental subjects or group of GA model agents) in period t as $\bar{p}_t^e = (1/6) \sum_{i=1}^6 p_{i,t}^e$. Then, the period coordination measure at t is

$$\text{Coord}_t = \frac{1}{6} \sum_{i=1}^6 \left(p_{i,t}^e - \bar{p}_t^e \right)^2. \quad (19)$$

Notice that Coord_t often changes over the session's course in the experimental groups. For example, under negative feedback in **HHST09** it was typically large in the initial periods, but would collapse close to zero later on (see Figure 4c). In order to measure how well a model traces the *evolution* of the subjects' coordination in a particular group, we use the following coordination proximity statistic:

$$\text{CoordProx}_{\text{Model}} = \frac{1}{T-3} \sum_{t=4}^T \left(\ln \left(\frac{\text{Coord}_t^{\text{Model}}}{\text{Coord}_t^{\text{Data}}} \right) \right)^2 \geq 0, \quad (20)$$

where T denotes the duration of the experimental session (note that the square guarantees this measure to be nonnegative). If the model perfectly predicts the evolution of the dispersion of the subjects forecasts, the coordination proximity measure (20) is equal to zero. The measure increases for both under- and overestimation of the individual dispersion, and explodes to infinity for homogeneous expectation models.

Tables 2 and 3 show the results of the Monte Carlo simulations for **HHST09** and **BHST12**, respectively. We report the three statistics (17)–(20) averaged separately over six (eight) groups for each treatment from **HHST09** (**BHST12**), for all the benchmark models and our GA model (averaged over 1024 sample GA simulated paths per group). The statistics for the best model is shown in bold and for the second best in italic.

Under negative feedback, two simple models of adaptive and contrarian expectations as well as RE perform well in terms of long-run predicting prices and forecasts, because they correctly predict convergence to the fundamental price. Our GA

TABLE 2. **HHST09.**

Model	Negative feedback			Positive feedback		
	Prices	Forec.	Coord.	Prices	Forec.	Coord.
Trend extrapolation	3421	1696	NA	62.84	72.45	NA
Adaptive	4.164	16.97	NA	95.62	108.6	NA
Contrarian	3.446	<i>16.18</i>	NA	108.5	122.8	NA
Naive	112.3	136.2	NA	69.11	79.38	NA
RE	2.571	15.21	NA	46.835	54.811	NA
LS learning	229.9	250.7	NA	334.2	342.7	NA
LS const. gain learning	230	250.8	NA	328.6	337.2	NA
HSM 4-type AH (2012)	499.78	634.53	23.84	45.96	53.84	4.88
HSM 2-type AHP (2013)	19.64	34.02	48.05	55.15	63.98	38.06
GA-P1	2.884	20.03	6.446	44.22	51.98	8.059
GA-P2 (no contrarians)	9.392	29.51	<i>7.687</i>	25.3	31.1	<i>8.216</i>
GA: Action-based	32.86	71.05	NA	82.78	120.89	NA
GA: AR1	10.6	40.31	15.17	106.8	122.1	8.727

Notes: 50-period ahead simulation, fitness of various models to the experimental data, experimental prices (MSE of predicted prices (17), columns Prices); subjects' forecasts (MSE of predicted forecasts (18), columns Forec.) and subject heterogeneity (predicted dispersion of forecasts (20), columns Coord.). Statistics are averaged over six experimental groups for the corresponding treatment, and the best model is in bold, the second best in italic. NA denotes infinite statistics.

TABLE 3. **BHST12.**

Model	Negative feedback			Positive feedback		
	Prices	Forec.	Coord.	Prices	Forec.	Coord.
Trend extrapolation	2736	1289	NA	101.26	113.3	NA
Adaptive	3.629	10.75	NA	55	62.14	NA
Contrarian	6.984	<i>14.45</i>	NA	58.46	65.95	NA
Naive	94.44	110.9	NA	46.62	52.9	NA
RE	13.871	20.923	NA	55.133	60.859	NA
LS learning	262.1	230.7	NA	228.8	235.8	NA
LS const. gain learning	384.2	367.7	NA	333.1	338.6	NA
HSM—4 type	236.08	267.592	31.402	32.18	37.01	4.57
HSM—2 type	73.57	87.86	52.73	90.8	101.8	65.62
GA-P1	8.01	21.97	<i>14.66</i>	<i>43.49</i>	49.44	<i>4.83</i>
GA-P2 (no contrarians)	<i>6.333</i>	17.39	13.34	<i>43.49</i>	<i>49.64</i>	5.112
GA: Action-based	29.73	90.7	40.79	179.8	200.2	8.985
GA: AR1	21.28	59.71	29.83	88.02	98.71	5.216

Notes: 65-period ahead predictions, fitness of various models to the experimental data, experimental prices (MSE of predicted prices (17), columns Prices); subjects' forecasts (MSE of predicted forecasts (18), columns Forec.) and subject heterogeneity (predicted dispersion of forecasts (20), columns Coord.). Statistics are averaged over six experimental groups for the corresponding treatment, and the best model is in bold, the second best in italic. NA denotes infinite statistics.

model performs only slightly worse. Furthermore, RE has problems with predicting the short spell of instability that follows fundamental shocks in **BHST12**, as seen in worsening of the statistics in Table 3.

Under positive feedback, the contrarian and adaptive expectations lose goodness of fit, because they still predict convergence, in contrast to the experimental data. The 2-type HSM, trend extrapolation and naive expectations perform relatively well, but surprisingly they are not better than RE, especially for **BHST12**. The reason is that the price oscillations predicted by these three models at the long run horizon fall out of phase with the experimental oscillations. LS learning performs poorly in long run forecasting for both positive and negative feedback. Our GA model is consistently either the best or the second best, especially the GA-P2 parametrization.³² Finally, under positive feedback, the model that comes closest to our GA model in terms of long horizon predictions is the 4-type HSM (in particular for **BHST12**), but at the expense of a poor fit under negative feedback.

Finally, our GA model is always among the best three models, if not the best, in terms of predicting the subject forecasts' heterogeneity, regardless of the experiment or the treatment. Notice that the two other GA specifications (Action-based and AR1) do not perform well and are typically among the worst models in any statistics. We conclude that (i) some of the alternative models are able to capture the price and forecast long-run dynamics of some treatments from either experiment, but (ii) only our heuristic-based GA model with FOR heuristic (10) consistently and successfully predicts *long-run* behavior of prices and forecasts in *all four* treatments; and (iii) is the only one to properly pick up the long-run evolution of subject heterogeneity in each treatment from these two experiments. This shows that the GA model is the only model to properly replicate the two experiments also *at the individual level*.

5.2. One-Period Ahead Simulations for Linear LtF Experiments

Another important indicator of the model's fit is the precision of its one-period ahead predictions: how well the model predicts experimental outcomes in period $t + 1$, conditional on the data available to the subjects of the experiment until period t .³³ For deterministic models such as the HSM and the homogeneous expectations models, computing the one period-ahead MSE is straightforward. For our GA model with its

32. As mentioned in Section 4.2 (see footnote 20) the range of β affects the initial set of heuristics and hence their forecasts in the initial periods. Comparison of performances of GA-P1 and GA-P2 in Table 2 across negative and positive feedback shows that initialization matters. For instance, GA-P1 will generate initially heuristics with both positive and negative β , and it performs better than GA-P2 under negative feedback. Instead, GA-P2 will generate initially heuristics with only positive β s, and it performs better under positive feedback. Notice that subjects in the lab experiment had general information about the market and could in fact infer that they were in a positive feedback environment and might therefore ignore negative coefficients under the positive feedback treatments.

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

TABLE 4. **HHST09**.

Model	Negative feedback			Positive feedback		
	Prices	Forec.	Coord.	Prices	Forec.	Coord.
Trend extrapolation	21.101	35.648	NA	0.926	4.196	NA
Adaptive	2.3	<i>14.912</i>	NA	2.999	6.482	NA
Contrarian	2.249	14.856	NA	3.864	7.436	NA
Naive	3.09	15.782	NA	1.822	5.184	NA
RE	2.571	15.21	NA	46.835	54.811	NA
LS learning	2.999	15.682	NA	0.889	<i>4.155</i>	NA
LS const. gain learning	4.323	17.142	NA	4.099	7.695	NA
HSM 4-type AH (2012)	9.05	22.35	4.80	0.82	4.07	6.16
HSM 2-type AHP (2013)	2.999	15.106	19.868	0.889	4.156	19.144
GA-P1	4.95	25.017	3.4291	<i>0.806</i>	4.235	4.9
GA-P2 (no contrarians)	4.496	25.012	<i>3.446</i>	0.802	4.198	5.558
GA: Action-based	32.855	71.047	NA	82.783	120.848	NA
GA: AR1	2.949	21.904	4.246	1.635	6.24	4.464

Notes: One-period ahead predictions, fitness of various models to the experimental data, experimental prices (MSE of predicted prices (17), columns Prices); subjects' forecasts (MSE of predicted forecasts (18), columns Forec.) and subject heterogeneity (predicted dispersion of forecasts (20), columns Coord.). Statistics are averaged over six experimental groups for the corresponding treatment, and the one the best model is in bold, of the second best is in italic. NA denotes infinite statistics.

probabilistic reproduction and three operators, however, evaluating the MSE is more complicated. Our model is both stochastic and highly nonlinear: it evolves according to an analytically intractable period-to-period distribution. To address this issue, we compute the *expected* MSE using a simple Sequential Monte Carlo (SMC) approach designed as follows.

For each experimental group X from the three experiments that we consider, and for each GA model specification, we run 1024 independent GA model simulations. In every simulation, we associate one GA agent with one subject, and in each period $t \geq 2$ every GA agent i (1) retains her own heuristics from the previous period and (2) is given the *experimental* prices and the price forecasts of “her” subject i until the previous period $t - 1$. Every GA agent uses the experimental (not artificial) data to update her heuristics and picks one heuristic to forecast the next price, which gives us the GA's price forecasts (16) and realized prices (according to the experimental price equation) for period t . We evaluate the fit of the model to the experimental group by computing the same three statistics (17)–(20) as for 50-period ahead simulations, averaged over all 1024 GA simulations.

The results for the one-period ahead simulations are summarized in Tables 4 and 5 for **HHST09** and **BHST12**, respectively. Under negative feedback many rules (RE, HSM, adaptive, contrarian, naive) capture the convergence of prices and forecasts to the fundamental price, putting our GA model at some disadvantage. Under positive feedback, these models (except for HSM) lose their predictive power and underestimate the experimental oscillatory behavior of individual forecasts. LS learning

TABLE 5. BHST12.

Model	Negative feedback			Positive feedback		
	Prices	Forec.	Coord.	Prices	Forec.	Coord.
Trend extrapolation	103.93	121.329	NA	0.548	2.165	NA
Adaptive	3.251	10.332	NA	2.797	4.618	NA
Contrarian	5.248	<i>12.534</i>	NA	3.653	5.559	NA
Naive	9.127	16.81	NA	1.583	3.286	NA
RE	12.856	20.923	NA	53.757	60.857	NA
LS learning	21.109	30.023	NA	3.283	5.156	NA
LS const. gain learning	23.485	32.643	NA	2.91	4.744	NA
HSM—4 type	6.28	13.68	15.06	0.42	2.01	7.77
HSM—2 type	15.45	23.78	36.022	0.42	<i>2.02</i>	16.23
GA-P1	10.247	21.464	<i>11.213</i>	<i>0.342</i>	2.059	5.892
GA-P2 (no contrarians)	<i>4.208</i>	15.267	11.127	0.341	2.036	7.051
GA: Action-based	34.865	67.753	32.176	40.333	66.451	23.347
GA: AR1	7.939	24.022	15.7	1.111	3.555	4.8

Notes. One-period ahead predictions, fitness of various models to the experimental data, experimental prices (MSE of predicted prices (17), columns Prices); subjects' forecasts (MSE of predicted forecasts (18), columns Forec.) and subject heterogeneity (predicted dispersion of forecasts (20), columns Coord.). Statistics are averaged over six experimental groups for the corresponding treatment, and the best model is in bold, the second best is in italic. NA denotes infinite statistics.

does surprisingly well in the one-period ahead forecasting of **HHST09**, as it is able to adapt to the changing trends in a stationary environment. In the **BHST12** experiments, however, LS performs slightly worse as it responds too slowly to the large structural breaks. The GA model has the best, or comparable to the best, fit for the positive feedback treatment and outperforms RE by a factor of 10. Furthermore, under both treatments, the GA model is particularly good at tracking the evolution of the heterogeneity of the individual forecasts, as seen in the Coordination statistic. As for the case of the long-run simulations, the GA model is the only model that consistently explains different experiments in the short-run, in terms of the prices and in terms of the distribution of the individual forecasts. This depends crucially on the specification of the GA model, since AR1 GA performs well only in some treatments, whereas Action-based GA is among the worst models.

5.3. Two-Period Ahead Asset Pricing LtF Experiment

In **HSTV05**, with the underlying law of motion (4)–(5), participants had to forecast p_{t+1} without knowing p_t , and therefore their two-period ahead price forecasts were based on a different information set than in the previous one-period ahead experiments. The two-period ahead version of our GA model is based on the following forecasting heuristic:

$$p_{i,h,t+1}^e = \alpha_{i,h,t} p_{t-1} + (1 - \alpha_{i,h,t}) p_{i,t}^e + \beta_{i,h,t} (p_{t-1} - p_{t-2}). \quad (21)$$

Once p_t is realized, the agents can evaluate their rules based on the hypothetical performance of predicting p_t *two periods ago*, that is, their fitness is a normalized SE = $(p_{i,h,t}^e - p_t)^2$, as before. This specification is the most straightforward adaptation of the baseline one-period ahead forecasting heuristic (10). Recall that in the two baseline parametrizations, GA-P1 and GA-P2, we imposed the restrictions on the trend coefficients, $\beta \in [-1.1, 1.1]$ and $\beta \in [0, 1.1]$, respectively. **HSTV05**, however, found that many subjects used stronger trend extrapolation. Therefore, for the sake of completeness we will also report the results of our GA model with $\beta \in [-1.3, 1.3]$ (parametrization GA-P3) and $\beta \in [0, 1.3]$ (parametrization GA-P4).³⁴

The lower panel of Figure 3 displays two typical simulated markets of the GA-P3 model for the **HSTV05** experiment. GA agents can either converge to the fundamental price (Figure 3c) or coordinate on large oscillations (Figure 3d). Figure 8 shows the results for MC 50-period ahead simulations for two GA model parametrizations, GA-P1 and GA-P3. Under the latter setting, the agents are allowed to experiment with higher trend coefficients β . The median price has a very similar oscillatory shape in both cases, but the difference is seen in the 95% CI. Both parametrizations are likely to generate two price bubbles within 50 periods, but the GA-P3 model with higher β s has larger potential oscillations (Figure 8b), and the second bubble can be even bigger than the first (unlike in the linear one-period ahead positive feedback case). In both parametrizations, the median GA agent converges to a strong trend extrapolation rule, close to $p_{i,t+1}^e \approx p_{t-1} + (p_{t-1} - p_{t-2})$, 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 is close to bi-modal (see Figure 9c and d), with a relatively large mass centered around zero, that is, weak or no trend extrapolation, and a peak around the maximum possible trend coefficient.

Even though our GA model leaves space for improvement,³⁵ it is the only model that is on par with Anufriev and Hommes (2012) in terms of predicting the experimental results of **HSTV05** both in the long- and the short-run. Table 6 reports the MSE of both 50-period and 1-period ahead simulations, initialized with the experimental initial predictions. The long-run predictive power is relatively poor for all models, as typical MSEs are approximately 10–25 times larger than for the two linear one-period ahead forecasting experiments. The best four models for 50-period ahead predictions are naive, adaptive, RE and the 4-type HSM, though our model (in particular GA-P1 and GA-P2 specifications) yields similar results. Surprisingly, the models that did well in terms of the long-run predictions become poor for one-period ahead predictions, with the exception of the 4-type HSM, whereas our GA model takes over as the best one.

34. Notice that subjects in this experiment extrapolate the trend two periods ahead, hence a trend coefficient 1.3 should be interpreted as an approximately trend extrapolation coefficient of 1.14 per period.

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

FIGURE 8. **HSTV05**: Emerging heuristics in 50-period ahead MC simulation (1000 markets) for GA-P1 (left panel) and GA-P3 (right panel) parametrizations. The price (upper panels), the price weight α (middle panels) and the trend extrapolation coefficient β (lower panels) of the chosen heuristic are shown. Red thick line is the median, blue dotted and purple dashed lines show the 95% and 90% confidence intervals, respectively, for the GA model.

FIGURE 9. Positive feedback treatments: **HHST09**, **BHST12**, and **HSTV05** with $p^f = 60$. Distribution of trend extrapolation coefficient β chosen by the agents in period $t = 50$ across the whole MC sample for GA-P1 and GA-P3 (last panel) parametrizations.

TABLE 6. **HSTV05**.

Predictions	50-period ahead			1-period ahead		
Models	Prices	Forec.	Coord.	Prices	Forec.	Coord.
Trend extrapolation	178.2	174.9	NA	17.45	55.09	NA
Adaptive	96.12	<i>145.9</i>	NA	44.13	25.32	NA
Contrarian	157	146.8	NA	59.39	30.86	NA
Naive	95.29	144.6	NA	31.67	20.84	NA
RE	<i>96.03</i>	146	NA	96.03	146.00	NA
LS learning	183.1	210.9	NA	138.84	109.41	NA
LS const. gain learning	170.1	176	NA	143.72	113.37	NA
HSM 4-type AH (2012)	96.38	144.42	11.42	7.88	53.25	2.89
HSM 2-type AHP (2013)	101.37	151.59	35.41	30.54	33.36	31.76
GA-P1	103.9	155.8	28.32	42.22	74.95	2.77
GA-P2	114.9	169.1	31.46	5.93	30.34	4.28
GA-P3	139.4	201.5	22.12	21.19	53.24	4.47
GA-P4	226.5	318.5	19.67	16.29	42.13	2.56
GA: AR1	121.6	182.2	<i>12.97</i>	45.55	136.62	3.24
GA: Action	154	221.8	NA	54.49	153.95	6.83

Notes. 50-period ahead and one-period ahead predictions, fitness of various models to the experimental data, experimental prices (MSE of predicted prices (17), columns Prices); subjects' forecasts (MSE of predicted forecasts (18), columns Forec.) and subject heterogeneity (predicted dispersion of forecasts (20), columns Coord.). Statistics are averaged over six experimental groups for the corresponding treatment, and the one the best model is in bold, of the second best is in italic. NA denotes infinite statistics.

The disparity between short- and long-run predicting power of different models has a natural interpretation. Our simulations show that the 2-period ahead economy is intrinsically unstable, also in terms of the specific amplitude, phase and frequency of price oscillations (see also the discussion on the long-run behavior of the model in Section 5.5). As a result any model that can actually generate proper price cycles and is initialized with the experimental data, is likely to “disconnect” itself from the data in the medium to long run, that is, generate oscillations out of phase with the experimental data. The most successful long-run prediction of the data is then just to predict its average, which is close to the fundamental value, as RE and adaptive expectations tend to do. Another successful prediction rule in such an unstable environment is naive expectations, taking account of the highly persistent data oscillations. In contrast, in the short run an average or RE fundamental predictions perform poorly, whereas the HSM and our GA model, both based on trend following heuristics, perform much better.

5.4. Trend Following Heuristics Across Different Experiments

It is clear from our analysis that the GA model can successfully replicate stylized facts and predict experimental data from the three LtF experiments reported by **HHST09**, **BHST12**, and **HSTV05**. The virtue of the GA model is that it allows agents to learn

quickly and adapt their behavior to diversified environments. The strength of our GA model lies in the optimization of the first order heuristic (10), because subjects exhibited this behavior in the lab.

As discussed before, positive feedback environments (in particular asset markets) reinforce trend following behavior, and the particular strength of trend chasing depends on the environment. Figure 9 displays the histogram (over 1000 runs) of the trend coefficients of the first-order heuristic (10), which were chosen in the last period by our GA agents in the “unconditional” 50-period ahead simulations across different positive feedback economies. Going from the simple linear feedback of **HHST09**, through the same linear feedback with additional large shocks to the fundamental value of **BHST12**, to the highly nonlinear two-period ahead feedback of **HSTV05**, we observe that the distribution of the trend following coefficient becomes more extreme. In comparison with **HHST09**, the GA agents in **BHST12** use stronger trend chasing rules, whereas the GA agents in **HSTV05** exhibit much more polarized behavior with more modes (either no trend whatsoever, or extreme trend following). We conclude that our GA model shows that the more complicated the asset market is, the less likely we will observe convergent type of behavior, whereas coordination on trend-following behavior may arise more easily.

5.5. Long-Run Dynamics

In comparison to the linear, one-period ahead experiments **HHST09** and **BHST12**, the two-period ahead **HSTV05** experiment with nonlinear feedback (due to robotic forecasters, see equation (5)) report more diversified dynamics. Among the 10 groups only 3 converged to the vicinity of the fundamental price, 3 generated dampening oscillations, and the remaining 4 exhibited persistent oscillations. This leads to a question about the *long-run* behavior of this two-period ahead asset market: does our GA model predict eventual convergence,³⁶ persistent instability, or maybe switching between these two outcomes? To stress the volatile behavior of this market structure, we report one long-run simulation for GA-P3 model in Figure 10. The top panels display price dynamics with persistent oscillations of different amplitude, where large oscillations can reappear even after the market seemingly settled on the fundamental price. This simulation suggests that the invariant distribution of our stochastic model may have several modes.³⁷ The GA model generates economically relevant dynamics with *clustered volatility*, that is, phases of relatively stable price behavior interchange with highly volatile price fluctuations. The bottom panel of Figure 10 shows the average

36. Simulations of the GA model show that for the **HHST09** and **BHST12** environments the long run dynamics converges to the fundamental price.

37. Due to the presence of mutation in the learning phase and the noise in the pricing equation, our GA model is an ergodic Markov process. Therefore, the invariant distribution exists, though it cannot be computed analytically due to the complexity of the model. As this paper is motivated by the experimental data, we simulated and compared in Figure 9 the distributions of the trend extrapolation coefficient after the first 50 periods for all the positive feedback treatments discussed in the paper, leaving more systematic investigation of the asymptotic properties of GA dynamics to future research.

FIGURE 10. **HSTV05**: sample 2000-period ahead simulation (*right panels*) and its first 500 periods (*left panels*) of GA-P3 model with $p^f = 60$. *Top panels*: individual predictions (green dashed lines) and price (black line). *Bottom panels*: average trend extrapolation coefficient β chosen by six GA agents.

β chosen by the six GA agents. Despite continuing instability, a clear pattern is that the average β remains close to zero in the stable phase of the simulation, but stays close to the upper limit of 1.3 in volatile times.

We interpret this pattern in the following way. If the price is stable and close to the fundamental value, the fittest heuristics give predictions that are close to the fundamental value. Due to averaging of the predictions of six GA agents and the robotic forecasters (who always predict the fundamental value) in the pricing equation (4), random deviations from the fundamental price will typically be mitigated. At the same time, although the price dynamics are near the fundamental price, the coefficients of the GA heuristic (10) are not identified (Youssefmir and Huberman 1997), allowing agents to experiment with strong trend extrapolation. If now a sufficiently large shock triggers an initial trend in prices, and if the GA agents coordinate sufficiently well on following that trend, they are able to counter-weight the stabilizing effect of robotic forecasters. This adds momentum to the initial price trend, and hence leads to a drift of the extrapolating coefficients in the fittest heuristics toward the upper bound. As a result, the price oscillations become self-reinforcing. A reverse scenario can occur as well, when a sufficiently large shock mitigates price trend, which stabilizes the forecasts of the GA agents and pushes the market toward the stable regime.

These switches in learning dynamics are possible, because under nonlinear (due to presence of robotic forecasters) two-period ahead price feedback mechanism, the specific shape (i.e., frequency, amplitude) of oscillations is diversified. As a result, there is still space for GA agents to experiment with the specific strength of trend following.

Therefore, in this economy our GA model not only entails two “attractors” (i.e., two types of long run behavior: stable fundamental price and large oscillations), but also generates endogenous switching between them. Note that the [HSTV05](#) experiment exhibited different dynamic patterns (e.g., convergence to the fundamental price and persistent oscillations around it) under the same experimental conditions. Whereas switching between these patterns has not been observed in the lab, simulations of our model suggest that if the experiment could last long enough, endogenous switching between converging and oscillating price behavior may be observed.

6. Conclusions

This paper presented a model where in a complex environment agents learn to use smart heuristics. Agents independently use a GA to optimize a simple anchor and adjustment forecasting heuristic. The GA model dynamics was compared with the stylized facts of Learning-to-Forecast experiments, where the realized market price depends on individual forecasts. These experiments are used to study how human subjects adapt to the price-predictions feedback in a controlled environment. We showed that GAs capture individual forecasting behavior in the experiments quite well and also reproduce the aggregate outcomes. GA agents use a parsimonious linear first-order heuristic (a mixture of adaptive and trend extrapolating expectations) to forecast prices. They independently optimize the two parameters of their forecasting rule with GAs, learning to fine-tune them to the specific market conditions.

Experimental data can be used to test various theories. Our goal was to compare the prediction accuracy of the GA model with other benchmark models: Rational Expectations, simple LS learning, constant gain learning, a number of homogeneous expectation models, and the Heuristic Switching Model of Anufriev and Hommes (2012) and Anufriev et al. (2013b). We focused on the out-of-sample one-period and 50-period ahead predictions and showed that in comparison with other models, the GA model is able to account for both the *aggregate* outcomes and the *individual* behavior across three *different* experiments. The strength of the model lies in its parsimony, flexibility, and generality. Under GA, the coefficients in the active heuristic are *time-varying*. When agents face a negative feedback type of economy, a median GA agent will increasingly rely on adaptive expectations, enforcing convergence of the market to the fundamental equilibrium. In contrast, positive feedback induces the agents to follow the observed price trend and median forecasting behavior converges to a trend extrapolation rule, which amplifies price oscillations. The more “complex” the positive feedback is (in terms of shocks to the fundamental solution or a two-period ahead nonlinear forecasting environment), the stronger trend extrapolation chosen by the median agent is, and the more volatile the price fluctuations will be.

Heterogeneity is a key feature of our GA-model that is consistent with individual and aggregate experimental data. Individuals are heterogeneous and different market environments generate heterogeneous aggregate outcomes. The evolutionary selection among different forecasting heuristics, a key feature of our GA-model, describes well the adaptive human behavior in such different market environments.

In the experiments that we used in our paper participants had to forecast only one variable, such as the price of a financial asset, for at least 50 periods. There are several other related experiments where we expect our heuristic-based GA model will perform equally well. Learning-to-Forecasts experiments in macroeconomic setting investigate the dynamics when participants forecast two variables, for example, inflation and output gap. The results of these experiments are roughly consistent with stylized facts **S1–S3**. Interestingly, the set of facts may be enriched allowing for more interactions between variables, as now, for example, one variable may have positive feedback, whereas the other variable may have negative feedback. As a natural next step, it would be interesting to generalize our GA model to such two-dimensional environment and compare its performance and predictions with existing learning theories. There is also a growing literature on a so-called “beauty contest” game, see Nagel (1995). Repeated version of this game resembles the Learning-to-Forecast experiments, and it is found that its dynamic properties depend on the strength and type of the feedback. It would be interesting to investigate how well our GA model captures behavior of participants in these experiments.

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

Appendix A: Definition of Forecasting Rules and Models

Table A.1 provides the specification for all forecasting rules and models used in the paper. For the full specification of the 2-type and 4-type HSM, see Anufriev et al. (2013b) and Anufriev and Hommes (2012), respectively.

A.1. Recursive Least Squares and Constant Gain Learning

Consider an agent in a one-period ahead LtF setting, who tries to minimize forecasting error $p_t^e - p_t$. That is, this agent is looking for coefficients α_t and β_t that optimize the fit of

$$p_t = \alpha_t p_{t-1} + (1 - \alpha_t) p_{t-1}^e + \beta_t (p_{t-1} - p_{t-2}),$$

with $\alpha_t \in [0, 1]$ and $\beta \in [-1.1, 1.1]$, the same constraints as in GA-P1 model specification.³⁸ The agent uses standard LS to estimate the optimal projection of

38. In the adaptive learning literature these bounded parameter intervals are referred to as projection facilities.

TABLE A.1. Specification of the forecasting rules used in the paper. For the one-period ahead environments (BHST09; BHST12), the rules generate prediction p_t^e , the adaptive rule includes $p_{\text{previous}}^e = p_{t-1}^e$, whereas the GA model includes $p_{i,\text{previous}}^e = p_{i,t-1}^e$. For the two-period ahead environment HSTV05, the rules generate prediction p_{t+1}^e , and the adaptive rule includes $p_{\text{previous}}^e = p_t^e$, whereas the GA model includes $p_{i,\text{previous}}^e = p_{i,t}^e$.

Rule/Model	Forecast p_t^e
<i>Homogeneous rules</i>	
Trend extrapolation	$p_{t-1} + \gamma(p_{t-1} - p_{t-2})$ with $\gamma = 1$
Adaptive	$w p_{t-1} + (1 - w) p_{\text{previous}}^e$, with $w = 0.75$
Contrarian	$p_{t-1} - 0.5(p_{t-1} - p_{t-2})$
Naive	p_{t-1}
RE	p_t^f
LS learning	$w_t = 1/(t - 2)$, detailed description below this table.
LS const. gain learning	$w_t = 0.05$, detailed description below this table.
<i>Heterogeneous rules</i>	
HSM 2-type AHP	Switching between 2 heuristics: trend extrapolation and adaptive expectations, as specified previously; learning parameters are $\beta = 1.5$, $\eta = 0.1$, $\gamma = 0.1$.
HSM 4-type AH	Switching between 4 heuristics: adaptive with $w = 0.65$, two trend extrapolation (with $\gamma = 0.4$ and $\gamma = 1.3$), and the anchor-and-adjustment rule; learning parameters are $\beta = 0.4$, $\eta = 0.7$, $\gamma = 0.9$.
GA model	$\alpha_{i,t} p_{t-1} + (1 - \alpha_{i,t}) p_{i,\text{previous}}^e + \beta_{i,t} (p_{t-1} - p_{t-2})$
GA-P1	With restrictions $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [-1.1, 1.1]$
GA-P2	With restrictions $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [0, 1.1]$
GA-P3	With restrictions $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [-1.3, 1.3]$
GA-P4	With restrictions $\alpha_{i,t} \in [0, 1]$ and $\beta_{i,t} \in [0, 1.3]$

$\hat{p}_t^e - p_{t-1}^e = p_t - p_{t-1}^e$ on the variables $p_{t-1} - p_{t-1}^e$ and $p_{t-1} - p_{t-2}$. Define the information set of the agent at the beginning of period t as

$$Y_t \equiv \begin{pmatrix} p_2 - p_1^e \\ \vdots \\ p_{t-1} - p_{t-2}^e \end{pmatrix},$$

$$X_t \equiv \begin{pmatrix} p_1 - p_1^e & p_1 - p_0 \\ \vdots & \vdots \\ p_{t-2} - p_{t-2}^e & p_{t-2} - p_{t-3} \end{pmatrix} \text{ with } p_1 - p_0 \equiv 0.$$

Notice that the initial trend $p_1 - p_0 = 0$ is set to zero as in the GA model and that both X_t and Y_t contain $t - 2$ rows. Second, define x_t and y_t as the $t - 2$ (the last) row of X_t and Y_t respectively. Then the OLS estimator $(\hat{\alpha}_t, \hat{\beta}_t)' = (X_t' X_t)^{-1} X_t' Y_t$ can be

represented as a recursion of the form

$$\begin{aligned} V_t &= V_{t-1} + w_t (x_t' x_t - V_{t-1}), \\ (\hat{\alpha}_t, \hat{\beta}_t)' &= (\hat{\alpha}_{t-1}, \hat{\beta}_{t-1})' + w_t V_t^{-1} x_t' (y_t - x_t \gamma_{t-1}), \end{aligned}$$

with weight $w_t = 1/(t - 2)$. If the recursion yields a coefficient outside its constraint, we set it to the relevant bound of the allowed interval.

For each experimental group, we initialize this recursion in the following way. The first three prices p_1 , p_2 , and p_3 are set equal to the experimental prices. The first three forecasts p_1^e , p_2^e , and p_3^e are set to the average forecast of the subjects from the respective periods of the group session. In period $t = 4$, the learning can start, and the initial estimation in to the baseline OLS is

$$\begin{aligned} V_4 &= X_4' X_4, \\ (\hat{\alpha}_4, \hat{\beta}_4)' &= V_4^{-1} X_4' Y_4. \end{aligned}$$

Under normal LS learning, the weight $w_t = 1/(t - 2)$ is time decreasing and converges to 0. We also consider constant gain LS learning with time invariant $w_t \equiv w = 0.05$ as is common in the literature. Robustness checks suggest that the qualitative results of this algorithm applied to our experimental data do not depend much on the specific value of the weight parameter w . Finally, the LS model has a straightforward translation to the two-period ahead setting, where the LS algorithm can start operating in period $t = 5$.

Appendix B: Robustness Checks

To illustrate the robustness of our GA model, we run additional 50- and 1-period ahead simulations (as discussed in Sections 5.1.2 and 5.2) for the GA-P1 model, but with alternative parametrization and heuristics.

Tables B.1 and B.2 summarize the results for **HHST09** experiments, for the following versions of GA-P1 model:

- HK: model with K heuristics ($H = K$), where we consider $K \in \{2, 4, 6, 8, 10, 100\}$,
- dM05: mutation rate of 5% ($\delta_M = 0.05$),
- dM10: mutation rate of 10% ($\delta_M = 0.1$),
- dM75: mutation rate of 75% ($\delta_M = 0.75$),
- dC50: crossover rate of 50% ($\delta_C = 0.5$),
- dC00: crossover rate of 0%, that is, no crossover at all ($\delta_C = 0.0$),
- AltCross: alternate crossover position—crossover swaps the first and the last ten bits,
- NoElec: election operated turned off (no election),

TABLE B.1. **HHST09**: 50-period ahead simulation.

Model	Negative feedback			Positive feedback		
	Prices	Preds.	Coord.	Prices	Preds.	Coord.
Trend extr.	3421	1696	NA	62.84	72.45	NA
Adaptive	4.164	16.97	NA	95.62	108.6	NA
Contrarian	3.446	16.18	NA	108.5	122.8	NA
Naive	112.3	136.2	NA	69.11	79.38	NA
RE	2.571	15.21	NA	46.835	54.811	NA
LS learning	229.9	250.7	NA	334.2	342.7	NA
LS const. gain learning	230	250.8	NA	328.6	337.2	NA
HSM—4 type	499.78	634.53	23.84	45.96	53.84	4.88
HSM—2 type	19.64	34.02	48.05	55.15	63.98	38.06
GA-P1	2.884	20.03	6.446	44.22	51.98	8.059
GA-P2	9.392	29.51	7.687	25.3	31.1	8.216
GA: H2	7.68	31.73	NA	123.2	139.8	5.298
GA: H4	3.513	21.79	5.837	81.04	92.74	5.555
GA: H10	3.496	21.7	5.973	56.98	66.09	6.933
GA: H100	3.145	19.35	8.062	46.54	54.5	11.69
GA: dM05	2.982	18.77	6.472	45.6	53.47	7.721
GA: dM10	3.331	19.11	6.221	47.36	55.42	7.502
GA: dM75	3.309	18.27	6.167	49.87	58.19	5.235
GA: dC50	2.929	20.2	6.416	44.04	51.79	8.181
GA: dC00	2.885	20.04	6.452	44.63	52.44	7.986
GA: AltCross	2.885	20.02	6.458	44.44	52.22	8.055
GA: NoElec	2.956	19.99	5.7	42.64	50.24	6.837
GA: 15Learn	2.967	19.91	5.553	46.65	54.68	6.9
GA: NoGANeg	4.265	20.19	5.589	60.59	70.02	4.232
GA: NoGAPos	2945	7239	NaN	14.43	19.08	21.13
GA: Actions	7.35	46.47	18.88	213.3	239.7	8.634
GA: AR1 rule	10.6	40.31	15.17	106.8	122.1	8.727
GA: FOR+Action	602	3981	66.16	460200	510900	NA
GA: FOR+Anchor	2.533	16.08	8.405	39.97	47.45	5.819

Notes: To complement Table 2 of the paper. MSE of various models for experimental prices and subjects' predictions, and coordination measure, averaged over six experimental groups for the corresponding treatment. Robustness check for model parameters (compare with GA-P1). NaN denotes statistics, which are larger than the standard floating number.

- 15Learn: only 15 periods of GA procreation-learning and a fixed set of rules for each agent thereafter (see Section 5.1.1 in the paper),
- noGANeg: the model is initialized as in GA-P1, but the heuristics are never updated with the GA. Hence, every agent has 20 fixed heuristics, which on average are equal to the median rule (14) learned by the GA-P1 agents under negative feedback,
- noGAPos: similar to noGANeg, but with heuristics initialized under constraints $\alpha \in [0.8, 1]$ and $\beta \in [0.1, 1.1]$, which results in heuristics that on average are equal to the median rule (13) of the GA agents under positive feedback,
- Action: GA directly optimize forecast (see Section 4.3 in the paper),

TABLE B.2. **HHST09**: one-period ahead predictions.

Model	Negative feedback			Positive feedback		
	Prices	Preds.	Coord.	Prices	Preds.	Coord.
Trend extr.	21.101	35.648	NA	0.926	4.196	NA
Adaptive	2.3	14.912	NA	2.999	6.482	NA
Contrarian	2.249	14.856	NA	3.864	7.436	NA
Naive	3.09	15.782	NA	1.822	5.184	NA
RE	2.571	15.21	NA	46.835	54.811	NA
LS learning	2.999	15.682	NA	0.889	4.155	NA
LS const. gain learning	4.323	17.142	NA	4.099	7.695	NA
HSM—4 type	9.05	22.35	4.80	0.82	4.07	6.16
HSM—2 type	2.999	15.106	19.868	0.889	4.156	19.144
GA-P1	4.95	25.017	3.4291	0.806	4.235	4.9
GA-P2	4.496	25.012	3.446	0.802	4.198	5.558
GA: H2	4.074	26.417	3.363	1.468	5.827	3.373
GA: H4	4.51	26.62	3.22	1.11	5.07	3.35
GA: H10	4.739	25.797	3.252	0.849	4.434	3.922
GA: H100	6.175	24.331	4.091	0.824	4.119	7.835
GA: dM05	6.441	24.434	4.129	0.835	4.263	5.816
GA: dM10	7.138	25.197	4.447	0.883	4.334	5.496
GA: dM75	6.04	23.481	4.589	1.006	4.83	4.223
GA: dC50	4.975	25.089	3.445	0.807	4.229	5.007
GA: dC00	4.946	25.062	3.419	0.806	4.244	4.863
GA: AltCross	4.95	25.024	3.423	0.81	4.242	4.887
GA: NoElec	5.223	25.262	3.539	0.824	4.248	4.283
GA: 15Learn	4.024	24.873	3.293	0.823	4.371	4.444
GA: NoGANeg	6.582	24.319	4.255	1.063	4.923	3.486
GA: NoGAPos	8.544	23.833	15.468	0.924	4.001	13.948
GA: Actions	32.855	71.047	NA	82.783	120.85	NA
GA: AR1 rule	2.949	21.904	4.246	1.635	6.24	4.464
GA: FOR+Action	25.166	75.302	7.579	3.658	11.625	7.827
GA: FOR+Anch.	4.455	21.754	4.342	1.134	5.033	3.219

Notes: To complement Table 4 of the paper. MSE of various models for experimental prices and subjects' predictions, and coordination measure, averaged over six experimental groups for the corresponding treatment. Robustness check for model parameters (compare with GA-P1).

- AR1: GA optimize AR1 rule (see Section 4.3 in the paper).
- FOR+Action: GA optimizes hybrid heuristic with adaptive and trend expectations (FOR rule) together with an “action” term (as in Action GA), namely $p_{i,t}^e = \alpha p_{t-1} + (1 - \alpha)p_{i,t-1}^e + \beta \Delta p_{t-1} + \gamma$, where $\alpha \in [0, 1]$ and $\beta \in [-1.1, 1.1]$ are the price weight and trend coefficient as in GA-P1 and $\gamma \in [0, 100]$ corresponds to the level from Action GA.
- FOR+Anchor: GA optimizes First Order Heuristic with anchor, that is $p_{i,t}^e = \alpha_1 p_{t-1} + \alpha_2 p_{i,t-1}^e + (1 - \alpha_1 - \alpha_2)p_t^{av} + \beta \Delta p_{t-1}$, where the anchor term $p_t^{av} = (1/(t - 1)) \sum_{s=1}^{t-1} p_s$ is defined as the average price so far.

Appendix C: Distribution of Random Initial Forecasts

In this appendix we discuss one aspect of initialization of the GA model for the “unconditional” 50-period Monte Carlo simulations discussed in Section 5.1, namely the choice of the distribution for the initial predictions. Recall that our task is to demonstrate that GA model can replicate experimental stylized facts. Two examples for the **HHST09** experiment illustrate that the initialization of the model can be crucial in achieving this task.

First, under negative feedback, the coordination in individual price forecasts was observed only after the price has converged to the fundamental. To replicate this feature in our simulations, one has to start with a similar degree of initial heterogeneity in the agents forecasts and then show that due to the learning of GA agents, coordination arises as in the experiment.

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

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

HHST09. For this experiment we use the estimated Winged Focal Point (WFP; see Figure C.1 for visualization) reported by Diks and Makarewicz (2013), which is given by

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

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

FIGURE C.1. Estimated density of winged focal point distribution for **HHST09** as given in equation (C.1). The sizes of the wings around the mass point $p = 50$ are scaled to their masses and lengths.

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

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

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 smaller weight on the focal point 50. Hence we follow Diks and Makarewicz (2013) and get the following estimations for the first period:

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

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

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

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

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

Appendix D: Distributions of the Last Period Predictions

Figure D.1 compares the distributions of the 49th period individual forecasts simulated by our GA model (specifications GA-P1 and GA-P2) in the “unconditional” MC simulations (Section 5.1.1), with the last period forecasts distribution in the **HHST09** experiment under negative and positive feedback. Every panel is divided into two parts by a horizontal black line. The upper part shows (in red) the distributions from the GA model. The lower part shows the distribution of forecasts in the experiment (which is reflected for the ease of comparison).

FIGURE D.1. **HHST09**: 49th period of 50-period ahead MC simulation (1000 runs) for GA-P1 model (*upper panels*) and GA-P2 model (*lower panels*) compared with the experimental data from negative (*left panels*) and positive feedback (*right panel*). Every panel contains two histograms: the histogram of the individual price forecasts of the GA agents (*upper part of panels*, red bars) and the reflected histogram of the experimental subjects' forecasts (*lower part of panels*, blue bars). Vertical dotted line represents the fundamental price.

References

- Adam, Klaus (2007). "Experimental Evidence on the Persistence of Output and Inflation." *The Economic Journal*, 117, 603–636.
- Anufriev, Mikhail and Cars Hommes (2012). "Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments." *American Economic Journal: Microeconomics*, 4, 35–64.
- Anufriev, Mikhail and Valentyn Panchenko (2009). "Asset Prices, Traders' Behavior and Market Design." *Journal of Economic Dynamics and Control*, 33, 1073–1090.
- Anufriev, Mikhail, Tiziana Assenza, Cars Hommes, and Domenico Massaro (2013a). "Interest Rate Rules and Macroeconomic Stability under Heterogeneous Expectations." *Macroeconomic Dynamics*, 17, 1574–1604.
- Anufriev, Mikhail, Cars H. Hommes, and Raoul H. S. Philipse (2013b). "Evolutionary Selection of Expectations in Positive and Negative Feedback Markets." *Journal of Evolutionary Economics*, 23, 663–688.
- Arifovic, Jasmina (1994). "Genetic Algorithm Learning and the Cobweb Model." *Journal of Economic Dynamics and Control*, 18, 3–28.
- Arifovic, Jasmina (1995). "Genetic Algorithms and Inflationary Economies." *Journal of Monetary Economics*, 36, 219–243.
- Arifovic, Jasmina (1996). "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies." *Journal of Political Economy*, 104, 510–541.
- Arifovic, Jasmina, James Bullard, and Olena Kostyshyna (2013). "Social Learning and Monetary Policy Rules." *The Economic Journal*, 123, 38–76.
- Assenza, Tiziana, Peter Heemeijer, Cars Hommes, and Domenico Massaro (2014). "Managing Self-Organization of Expectations through Monetary Policy: A Macro Experiment." CeNDEF Working Paper 14-07, University of Amsterdam.
- Bao, T., C. Hommes, J. Sonnemans, and J. Tuinstra (2012). "Individual Expectations, Limited Rationality and Aggregate Outcomes." *Journal of Economic Dynamics and Control*, 36, 1101–1120.
- Bao, Te, John Duffy, and Cars Hommes (2013). "Learning, Forecasting and Optimizing: An Experimental Study." *European Economic Review*, 61, 186–204.
- Bao, Te, Cars Hommes, and Tomasz Makarewicz (2017). "Bubble Formation and (In)efficient Markets in Learning-to-Forecast and -Optimize Experiments." *The Economic Journal*, 127, F581–F609.
- Bolt, Wilko, Maria Demertzis, Cees G. H. Diks, Cars H. Hommes, and Marco Van der Leij (2014). "Identifying Booms and Busts in House Prices under Heterogeneous Expectations." Discussion Paper 14-157/II, Tinbergen Institute.
- Boswijk, H. Peter, Cars H. Hommes, and Sebastiano Manzan (2007). "Behavioral Heterogeneity in Stock Prices." *Journal of Economic Dynamics and Control*, 31, 1938–1970.
- Branch, William A. (2004). "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations." *The Economic Journal*, 114, 592–621.
- Branch, William A. and George W. Evans (2010). "Asset Return Dynamics and Learning." *Review of Financial Studies*, 23, 1651–1680.
- Branch, William A. and Bruce McGough (2010). "Dynamic Predictor Selection in a New Keynesian Model with Heterogeneous Expectations." *Journal of Economic Dynamics and Control*, 34, 1492–1508.
- Branch, William A. and Bruce McGough (2011). "Business Cycle Amplification with Heterogeneous Expectations." *Economic Theory*, 47, 395–421.
- Brock, W. A. and C. Hommes (1998). "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model." *Journal of Economic Dynamics and Control*, 22, 1235–1274.
- Brock, William A. and Cars H. Hommes (1997). "A Rational Route to Randomness." *Econometrica*, 65, 1059–1095.
- Bullard, James (1994). "Learning Equilibria." *Journal of Economic Theory*, 64, 468–485.

- Camerer, Colin and Hua Teck Ho (1999). "Experience-weighted Attraction Learning in Normal Form Games." *Econometrica*, 67, 827–874.
- Case, Karl E., Robert J. Shiller, and Anne K. Thompson (2012). "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets." *Brookings Papers on Economic Activity*, 2012, 265–298.
- Cornea-Madeira, Adriana, Cars Hommes, and Domenico Massaro (2017). "Behavioral Heterogeneity in US Inflation Dynamics." *Journal of Business & Economic Statistics*, 1–13, doi:10.1080/07350015.
- Dawid, Herbert (1996). *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*. Springer-Verlag Inc., New York.
- Dawid, Herbert and Michael Kopel (1998). "On Economic Applications of the Genetic Algorithm: A Model of the Cobweb Type." *Journal of Evolutionary Economics*, 8, 297–315.
- De Jong, Kenneth A. (2006). *Evolutionary Computation: A Unified approach*. MIT Press, Cambridge, MA.
- Diks, Cees and Tomasz Makarewicz (2013). "Initial Predictions in Learning-to-Forecast Experiment." In *Managing Market Complexity*, Vol. 662, edited by A. Teglio, S. Alfarano, E. Camacho-Cuena, and M. Ginés-Vilar, pp. 223–235. Springer, Berlin/Heidelberg.
- Doornik, Jan (2007). *Object-Oriented Matrix Programming Using Ox*, 3rd ed. Timberlake Consultants Press, London.
- Duffy, John (2016). "Macroeconomics: A Survey of Laboratory Research." In *Handbook of Experimental Economics*, Vol. 2, edited by J. H. Kagel and A. E. Roth. Princeton University Press, Princeton, pp. 1–90.
- Duffy, John and M. Utku Ünver (2006). "Asset Price Bubbles and Crashes with Near-Zero-Intelligence Traders." *Economic Theory*, 27, 537–563.
- Erev, Ido and Alvin E. Roth (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria." *American Economic Review*, 88(4), 848–881.
- Evans, G. W. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- Feltovich, Nick (2000). "Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games." *Econometrica*, 68, 605–641.
- Gigerenzer, Gerd and Peter M. Todd (1999). *Simple Heuristics That Make Us Smart*. Oxford University Press, New York.
- Goodman, John L. and John B. Iitner (1992). "The Accuracy of Home Owners' Estimates of House Value." *Journal of Housing Economics*, 2, 339–357.
- Guesnerie, Roger (2005). *Assessing Rational Expectations 2: "Eductive" Stability in Economics*. MIT Press.
- Haupt, Randy and Sue Haupt (2004). *Practical Genetic Algorithms*, 2nd ed. John Wiley & Sons, Inc., Hoboken, NJ.
- Heemeijer, Peter, Cars Hommes, Joep Sonnemans, and Jan Tuinstra (2009). "Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback: An Experimental Investigation." *Journal of Economic Dynamics and Control*, 33, 1052–1072.
- Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998). "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'." *American Economic Review*, 88(4), 947–969.
- Holland, John H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Oxford, England.
- Hommes, Cars (2011). "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab." *Journal of Economic Dynamics and Control*, 35, 1–24.
- Hommes, Cars (2013). *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, Cambridge.
- Hommes, Cars and Thomas Lux (2013). "Individual Expectations and Aggregate Behavior in Learning to Forecast Experiments." *Macroeconomic Dynamics*, 17, 373–401.
- Hommes, Cars, Joep Sonnemans, Jan Tuinstra, and de Velden van Henk (2005). "Coordination of Expectations in Asset Pricing Experiments." *The Review of Financial Studies*, 18, 955–980.

- Hommel, Cars, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden (2007). "Learning in Cobweb Experiments." *Macroeconomic Dynamics*, 11, 8–33.
- Kahneman, Daniel (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, NY.
- Kindleberger, Charles P. and Robert Z. Aliber (2011). *Manias, Panics and Crashes: A History of Financial Crises*. Palgrave Macmillan, Basingstoke.
- Kurz, Mordecai and Maurizio Motolese (2011). "Diverse Beliefs and Time Variability of Risk Premia." *Economic Theory*, 47, 293–335.
- LeBaron, Blake, W. Brian Arthur, and Richard Palmer (1999). "Time Series Properties of an Artificial Stock Market." *Journal of Economic Dynamics and Control*, 23, 1487–1516.
- Lucas, Robert E. (1972). "Expectations and the Neutrality of Money." *Journal of Economic Theory*, 4, 103–124.
- Lucas, Robert E. (1986). "Adaptive Behavior and Economic Theory." *The Journal of Business*, 59, S401–S426.
- Lux, Thomas and Sascha Schornstein (2005). "Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets." *Journal of Mathematical Economics*, 41, 169–196.
- Malmendier, Ulrike and Stefan Nagel (2009). "Learning from Inflation Experiences." <http://faculty-gsb.stanford.edu/nagel/documents/InflExp.pdf>.
- Mitchell, Melanie (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
- Muth, John F. (1961). "Rational Expectations and the Theory of Price Movements." *Econometrica*, 29, 315–335.
- Nagel, Rosemarie (1995). "Unraveling in Guessing Games: An Experimental Study." *American Economic Review*, 85(5), 1313–1326.
- Nunes, Ricardo (2010). "Inflation Dynamics: The Role of Expectations." *Journal of Money, Credit and Banking*, 42, 1161–1172.
- Pfajfar, Damjan and Blaž Žakelj (2014). "Experimental Evidence on Inflation Expectation Formation." *Journal of Economic Dynamics and Control*, 44, 147–168.
- Reinhart, Carmen M. and Kenneth Rogoff (2009). *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press, Princeton.
- Sargent, Thomas J. (1993). *Bounded Rationality in Macroeconomics*. Oxford University Press, New York, NY.
- Sims, Christopher A. (2010). "Rational Inattention and Monetary Economics." In *Handbook of Monetary Economics*, Vol. 3, edited by Benjamin M. Friedman and Michael Woodford. Elsevier, pp. 155–181.
- Sonnemans, Joep and Jan Tuinstra (2010). "Positive Expectations Feedback Experiments and Number Guessing Games as Models of Financial Markets." *Journal of Economic Psychology*, 31, 964–984.
- Sutan, Angela and Marc Willinger (2009). "Guessing with Negative Feedback: An Experiment." *Journal of Economic Dynamics and Control*, 33, 1123–1133.
- Tversky, Amos and Daniel Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases." *Science*, 185, 1124–1131.
- Vriend, Nicolaas J. (2000). "An Illustration of the Essential Difference between Individual and Social Learning, and its Consequences for Computational Analyses." *Journal of Economic Dynamics and Control*, 24, 1–19.
- Wagener, Florian (2014). "Expectations in Experiments." *Annual Review of Economics*, 6, 421–443.
- Woodford, Michael (2010). "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review*, 100(1), 274–303.
- Woodford, Michael (2013). "Macroeconomic Analysis Without the Rational Expectations Hypothesis." *Annual Review of Economics*, 5, 303–346.
- Youssefmir, Michael and Bernardo A. Huberman (1997). "Clustered Volatility in Multiagent Dynamics." *Journal of Economic Behavior & Organization*, 32, 101–118.

Supplementary Data

Supplementary data are available at [JEEA](https://www.jeeaonline.org) online.