

INDIVIDUAL EXPECTATIONS AND  
AGGREGATE BEHAVIOR IN  
LEARNING TO FORECAST EXPERIMENTS \*

**CARS HOMMES**

*School of Economics CeNDEF  
University of Amsterdam E-Mail: C.H.Hommes@uva.nl*

**THOMAS LUX**

*Department of Economics Olshausen Str. 40, 24118 Kiel  
University of Kiel, Germany E-Mail: lux@bwl.uni-kiel.de  
and*

*Kiel Institute for the World Economy Düsternbrooker Weg 120, 24105 Kiel*

February 23, 2011

*Macroeconomic Dynamics*, forthcoming

**Abstract:** Models with heterogeneous interacting agents explain macro phenomena through interactions at the micro level. We propose genetic algorithms as a model for individual expectations to explain aggregate market phenomena. The model explains all stylized facts observed in aggregate price fluctuations and individual forecasting behaviour in recent learning to forecast laboratory experiments with human subjects (Hommes et al. 2007), simultaneously and across different treatments.

**Keywords:** learning, heterogeneous expectations, genetic algorithms, experimental economics.

**JEL-codes:** C91, C92, D83, D84, E3

---

\*Financial support from the EU, through the POLHIA grant no.225408 and Complex Markets grant no.516446 is gratefully acknowledged. We are also particularly grateful to Simone Alfarano, Jasmina Arifovic, Reiner Franke, Juliette Rouchier, Jan Tuinstra and Frank Westerhoff for helpful comments and discussions.

# 1 Introduction

An important feature of models with heterogeneous, interacting agents is that they can explain macro phenomena through simple interactions at the micro level (e.g. Kirman, 1993, 2006, Lux, 1995). Agent-based models have been particularly successful in explaining the main stylized facts of financial markets, such as fat tails and clustered volatility in asset returns (Arthur et al., 1997, Lux and Marchesi, 1999 and Hommes, 2002, among others). Duffy (2006) presents a recent overview of how agent-based models can explain individual behavior and aggregate phenomena in macroeconomics. The main purpose of our paper is to explain aggregate price behavior through interactions of individual learning. In particular, we provide a simple theory of individual learning through genetic algorithms explaining all stylized facts of aggregate price fluctuations in the recent learning to forecast laboratory experiments of Hommes et al. (2007).

Laboratory experiments with human subjects have become an important tool in economic analysis, complementing theoretical, computational and empirical work. A recurring observation from experiments is that individuals often do not behave fully rationally, but tend to use simple heuristics, possibly biased, in making decisions under uncertainty (Tversky and Kahneman, 1974). An extensive bounded rationality research program is developing (e.g. Sargent, 1993) and laboratory experiments are particularly suited to identify behavioral rules that individuals use in economic decision making out of an ocean of potential alternatives (e.g. Kahneman, 2003).

Individual expectations, their interaction and the aggregate outcome they create are at the heart of economics. Duffy (2008), for example, argues that laboratory experiments are important to study the adaptive process by which individuals learn and may or may not enforce convergence to a rational expectations outcome at the macro level. Recently, a number of *learning to forecast* experiments have been conducted to study individual expectation formation and aggregate outcomes, e.g. in Marimon and Sunder (1994), Gerber et al. (2002), Hommes et al. (2005), Sutan and Willinger (2005), Adam (2007) and Heemeijer et al. (2009); see Hommes (2011) for an overview. In these experiments, subjects must forecast the price of a good, which is determined by market clearing with feedback from

individual expectations. Aggregate demand and supply are computerized, e.g. derived from profit and utility maximization given subjects' individual forecasts. An advantage of this experimental setup is that it provides 'clean data' on expectations, *ceteris paribus*. These experimental data can therefore be used to test various theories of bounded rationality, individual expectations and learning at the micro level and test how their interaction matches aggregate behavior.

Hommes et al. (2007) conducted learning to forecast experiments in what is perhaps the simplest setting, the classical cobweb 'hog cycle' model describing a standard commodity market with a production lag. Historically, the cobweb model has served as a simple framework to develop and test various expectations hypotheses. Ezekiel (1938) started with naive expectations, Nerlove (1958) advocated adaptive expectations, Muth's seminal paper (Muth, 1961) used the cobweb framework to introduce rational expectations and, more recently, Brock and Hommes (1997), used it to introduce endogenous selection among heterogeneous expectations rules<sup>1</sup>.

In their learning to forecast experiments, Hommes et al. (2007) considered three different treatments, a stable, an unstable and a strongly unstable treatment. Stable here refers to the stability of the classical cobweb model under naive expectations<sup>2</sup>. They observed the following three *stylized facts*:

1. the sample mean of realized prices was very close to the RE benchmark *in all three treatments*;
2. the sample variance of realized prices, however, depended on the treatment:
  - a) it was close to the theoretical variance of the RE benchmark in the stable treatment, while
  - b) it was significantly higher than the RE benchmark (excess volatility) in the unstable and strongly unstable treatments;

---

<sup>1</sup>The cobweb model has been used in many different applications, ranging from markets for lawyers (Freeman, 1975), engineers (Freeman, 1976), public school teachers (Zarkin, 1985), oil (Krugman, 2001), cattle (Rosen et al., 1994) and beef (Chavas, 2000).

<sup>2</sup>The stability condition states that the ratio between marginal supply and marginal demand at steady state must be smaller than 1 in absolute value (Ezekiel, 1938).

3. in all treatments, realized market prices exhibited no significant linear autocorrelation.

These stylized facts were quite robust over a series of experiments, but they appeared hard to explain by standard learning mechanisms offered by the theoretical literature. While many adaptive learning rules lead to eventual convergence to rational expectations and some other learning rules may generate unstable dynamics and excess volatility, homogeneous expectations models are *unable* to explain the full set of stylized facts *simultaneously* (Hommes, 2009), suggesting that *heterogeneity* of forecasting rules plays a key role in explaining observed aggregate behavior. Hommes (2009) considers some simple examples with two different forecasting rules (e.g. fundamentalists versus naive expectations) and evolutionary competition between them. Although the results in these simple two type heterogeneous expectations examples improve compared to the homogeneous case, they are not capable of explaining *all* stylized facts simultaneously.

In this paper we propose a simple model for micro behavior in order to explain the observed experimental results at the macro level. We model individual learning through genetic algorithms (GAs)<sup>3</sup>. As it turns out, GA experiments with a small population of agents match *all* stylized facts simultaneously across different treatments within a market setting that *exactly* corresponds to that of the laboratory experiments. While it is certainly hard to imagine GAs as an accurate description of human learning in the literal sense, we argue that they may share key properties of the adapta-

---

<sup>3</sup>Arifovic (1994) used GA learning to explain the cobweb laboratory experiments of Wellford (1989), but there are a number of important differences with our approach. Most importantly, subjects assume the role of producers themselves whereas the subjects of Hommes et al. (2007) have to forecast next period's market price. Besides this major difference in the focus of the experiments, a number of additional differences exist. For example, Wellford used linear demand and supply curves in the experimental setup, implying that the market is either stable or explosive (except for a hairline case). Hommes et al. (2007) use a nonlinear supply curve, so that price dynamics remains bounded while price cycles become a generic possibility. Moreover, Arifovic (1994) only tests for differences of volatility between the stable and unstable treatments, while we match price volatility under GA learning directly to the experiments. Furthermore, Arifovic (1994), for example, did not look at the third stylized fact, the absence of linear predictability. Finally, we also study the average degree of heterogeneity in individual forecasting and how it varies over time.

tions of human subjects when exposed to a new situation that they cannot penetrate theoretically.

We also investigate the *degree of heterogeneity* in individual forecasting behaviour. Heterogeneity in forecasting future asset prices is supported by evidence from stock market survey data, e.g. Vissing-Jorgensen (2003) and Shiller (2000), and inflation expectations survey data, e.g. Branch (2004) and Mankiw et al. (2003). Moreover, in these survey data heterogeneity shows substantial variation through time. Consistent with the findings in the laboratory experiments in our GA learning simulations heterogeneity decreases over time, disappears in the stable treatment but heterogeneity persists in the (strongly) unstable treatment.

Using the GAs for individual learning, our paper makes another contribution that goes beyond the limitations of laboratory experiments. Laboratory experiments are costly, because subjects must be paid according to their performance, and typically experimental markets are small because of capacity limitations. After fitting our GA model to individual learning, we can easily investigate price behavior in alternative, more realistic market scenarios through numerical simulations. In particular, we investigate the occurrence of excess volatility when the number of subjects in the market becomes large and/or when the number of rules per individual becomes large. We also investigate how excess volatility depends on a continuum of parameters such as the ratio of marginal supply and demand.

The paper is organized as follows. Section 2 recalls the learning to forecast experiments, while Section 3 recalls some basic facts of GA-learning. Section 4 compares the stylized facts of the GA simulations to the laboratory experiments, while Section 5 presents simulations of GA learning in more realistic market scenarios. Finally, Section 6 concludes.

## 2 The forecasting experiment

Hommes et al. (2007) report on a set of cobweb experiments with  $K = 6$  participants per session. The participants were asked to predict next period's price under very limited information on the structural characteristics of the market. The realized price  $p_t$  in the experiments was determined by

the (unknown) market equilibrium between demand and supply:

$$D(p_t) = \frac{1}{K} \sum_{i=1}^K S(p_{i,t}^e) \quad (1)$$

with  $p_{i,t}^e$  the price forecast of participant  $i$  at time  $t$ . We normalize the supply side by dividing by the number of firms to facilitate comparison of settings with different  $K$ . Supply  $S(p_{i,t}^e)$  was determined by the nonlinear schedule

$$S(p_{i,t}^e) = \tanh(\lambda(p_{i,t}^e - 6)) + 1, \quad (2)$$

while demand was formalized via a simple linear schedule:

$$D(p_t) = a - dp_t + \eta_t \quad (3)$$

with  $\eta_t$  a small stochastic shock drawn from a Normal distribution. Both demand and supply can be derived from profit and utility maximization, and are thus consistent with rational behavior. The resulting equilibrium price is obtained as:

$$p_t = \frac{a - \frac{1}{K} \sum_{i=1}^K S(p_{i,t}^e)}{d} + \epsilon_t, \quad (4)$$

where  $\epsilon_t = \eta_t/d$ . Given the parameters  $a, d$  and  $\lambda$  the aggregate realized price  $p_t$  depends on individual price expectations as well as the realization of the stochastic shocks.

Participants were only informed about the basic principles of the cobweb-type market. They were advisors to producers, whose only job is to accurately forecast the price of the good for 50 subsequent periods. Pay-offs were defined as a quadratic function of squared forecasting errors, truncated at 0:<sup>4</sup>

$$E = \text{Max}\{1300 - 260(p_{i,t}^e - p_t^*)^2, 0\}. \quad (5)$$

Participants were informed that the price would be determined by market clearing and that it would have to be within the range  $[0, 10]$ . Furthermore,

---

<sup>4</sup>1300 points corresponded to 0.45 Euro, so that maximum earnings were 22.5 Euro's. Average earnings ranged from about 10 Euro in the strongly unstable to about 19 Euro in the stable treatment (in about 75 minutes; see Hommes et al. (2007, Table 2)).

they knew that there was (negative) feedback from individual price forecasts to realized market price in the sense that if their forecast would increase, the supply would increase and consequently the market price would decrease. Subjects however did not know how large these feedback effects would be, as they had no knowledge of underlying market equilibrium equations. One could say that subjects had *qualitative* information about the market, but no quantitative details.

Participants thus solely had to rely on time series observations and their own behavior vis-à-vis their predictions. Due to the nonlinear aggregation of expectations, the superimposed noise and the ignorance of agents concerning the structural form and parameters, conscious coordination to some kind of rational expectations equilibrium would be extremely demanding if not impossible. This setting is close to the informational assumptions of various theoretical models in the literature on learning and bounded rationality (e.g. Sargent, 1993 and Evans and Honkapohja, 2001) so that the experimental subjects' behavior could be contrasted with various learning rules.

Following the classification of cobweb scenarios under naive expectations, Hommes et al. (2007) distinguished between three treatments. While the parameters of the demand function and the parameters of the noise component remained unchanged across all treatments at  $a = 2.3$ ,  $d = 0.25$  and  $\epsilon_t = \frac{\eta_t}{d} \sim N(0, 0.5)$ , the slope parameter of the supply function was varied over a relatively wide range. Treatment 1 had  $\lambda = 0.22$  which under naive expectations would guarantee convergence to the homogeneous rational expectations equilibrium (stable case), treatment 2 had a marginally unstable supply parameter  $\lambda = 0.5$  (unstable case), while treatment 3 had a strongly unstable supply parameter  $\lambda = 2$  (strongly unstable case). Both the unstable and the strongly unstable treatments lead to a 2-cycle under naive expectations. In all treatments in Hommes et al. (2007) the number of subjects was  $K = 6$ , but van de Velden (2001) also ran the same experiments in the strongly unstable treatment with  $K = 12$  subjects.

Under rational expectations, all individuals would predict  $p_t^e = p^*$ , that is, they would predict the price at which demand and supply intersect. Given that all individuals have rational expectations, realized prices will be given by

$$p_t = p^* + \epsilon_t. \tag{6}$$

Given the limited market information one can not expect that all individuals have rational expectations at the outset, but one can hope that in such a simple, stationary environment individuals would learn to have rational expectations. The *stylized facts* of these cobweb experiments have already been summarized in the introduction. We briefly recall them here: (1) the sample mean of realized prices is close to the RE benchmark  $p^*$  in all three treatments; (2) the sample variance of realized prices depends on the treatment: it is close to the RE benchmark in the stable treatment, but significantly higher in the unstable and strongly unstable treatments; (3) realized market prices do not exhibit significant linear autocorrelations. Item (3) indicates that even in the unstable and strongly unstable cases, agents did not leave any linear predictability unexploited. Apparently, the interaction of agents' individual forecasting rules washes out linear predictability in aggregate price behavior. While this points to a certain efficiency of their dispersed effects to predict market prices, market prices did fluctuate 'excessively' in the unstable and strongly unstable treatments. In these cases, price fluctuations exceeded those warranted by the exogenous noise component by more than one order of magnitude so that participants' attempts at learning about the market's behavior did apparently intensify price fluctuations. While these results were quite robust over a series of experiments, they appeared hard to explain by standard learning mechanisms offered by the theoretical literature.

Our goal here is to model individual learning via genetic algorithms (GAs), so that the interaction of these rules produces the stylized facts observed in the experiments simultaneously and across treatments.

### 3 Learning through Genetic Algorithms

Genetic algorithms have been introduced in the seventies as a stochastic learning algorithm (Holland, 1975). In order to solve complex optimization problems with multiple maxima or minima and possible non-continuities this approach mimics evolutionary operations in nature. One typically starts out with a randomly initiated population of candidate solutions. These initial blind trials are typically encoded as *chromosomes* (strings)



using a binary alphabet<sup>5</sup>. After evaluation of the fitness of the members of the initial population (in terms of the objective function), one applies the genetic operations of *reproduction*, *crossover*, and *mutation*. For certain reasons, economic applications have mostly added the *election* operator (Arifovic, 1996) as an additional step in the loop of genetic operations between successive generations of individuals. In the following we provide details of these operators and their implementation in the present setting:

1. *Reproduction*: in the transition from one generation to the next, the first step consists in sampling copies of strings from the old generation depending on their *fitness*. In conformity with the pay-off function used in the laboratory experiments fitness was defined as a negative quadratic function of forecast errors with truncation at zero:

$$f_i(t) = \text{Max}\{1300 - 260(p_{i,t}^e - p_t^*)^2, 0\}. \quad (7)$$

The most common reproduction operator is reproduction depending on relative fitness, i.e. copies are sampled from the old population with probabilities  $f_i / \sum_j f_j$  biasing the population of new agents towards strategies with higher fitness. Other algorithms in the literature are rank-dependent reproduction or tournament selection in which one draws repeatedly  $n_1$  individuals with replacement from the pool and accepts the  $n_2 < n_1$  with the highest fitness from the subsample until the new generation is complete. Throughout the present study we use fitness-based probabilities in order to mimic as closely as possible the incentives of the human subjects. To be concrete, we select new individuals with probabilities equal to the relative fitness of a strategy until the population is complete. Thus, a small difference in fitness would only lead to a similarly small evolutionary advantage of a strategy, while rank-based reproduction would be insensitive to the size of fitness differentials.<sup>6</sup>

---

<sup>5</sup>One could as well encode the population as real-coded chromosomes. However, this alternative encounters certain technical problems even if the problem at hand is properly defined for real values (Herrera et al., 1998).

<sup>6</sup>We have also experimented with a modification of the ‘election’ operation detailed below that replaces its built-in rank-based structure by a fitness-proportional choice among parents and off-spring. However, this modification had no remarkable effect

2. *Crossover*: when the pool of members of a new generation is complete, genetic material is exchanged between them in order to find new (possibly better) candidate solutions by recombination of the old ones. The simplest version is random selection of a pair of parent strings, determining a cut-off value within the string and sweeping part of the genetic material of the parents when creating their offspring. We follow this approach and take the genetic material of each of both offspring from the left (right) hand side of their ‘father’ and the right (left) hand side of their ‘mother’. This operation takes place with a probability  $p_{cross}$ , while with  $1 - p_{cross}$  the parent strings are transferred unchanged into the new generation. We note that both more involved crossover schemes as well as versions with more than two parents can be found in the literature as well.<sup>7</sup>
3. *Mutation*: means that any position (bit) within a chromosome might be flipped into another value (from ‘0’ to ‘1’ or vice versa in the binary alphabet). This happens with a probability  $p_{mut}$  once the reproduction and crossover operations are finished.
4. *Election*: in most economic applications, the usual range of genetic operators has been extended by the election algorithm. This compares new chromosomes that have emerged from crossover and mutation with their parents and only admits them to the population if their virtual fitness (measured with respect to the environment in which their parents had to compete) is at least as high as their parents’ fitness. This operator serves to prevent agents from adopting clearly inferior strategies. *Most* new strategies that emerge in a genetic process are far off the mark and conscious agents would not voluntarily adopt these new strategies if their trial performance ranks them way below the previous ones.

In many applications of genetic algorithms, the qualitative outcome is largely independent of the particular version of an operator that one adopts (cf. Lux and Schornstein, 2005, for a detailed comparison of various set-ups

---

on any of our results.

<sup>7</sup>Cross-over seems of fairly negligible influence as throughout even for extreme choices  $p_{cross} = 0$  or  $p_{cross} = 1$  results are virtually the same as with any intermediate choice. Similar findings are reported in Lux and Schornstein (2005).

within a learning context). One may even skip one or the other of the operators (e.g. crossover or election) without changing the overall qualitative results. In our simulations like in various previous economic applications, the results appear to be quite robust under variation of GA parameters and implementations of operators. Unfortunately, one has to rely exclusively on simulations since theoretical results for GAs within an interactive context seem to be essentially out of reach. In our setting with artificially intelligent agents, we tried to reproduce the design of the experiments as close as possible. This applies not only to the parametrization of demand and supply functions, and the choice of a fitness function identical to the payoff function in the laboratory experiments but also to the *number* of agents. Hence we report below experiments conducted with  $K = 6$  agents using genetic algorithms to evolve forecasting strategies.

Economic applications of GAs as a learning device have mostly applied it in the sense of ‘social learning’: the number of agents in these papers equals the number of chromosomes and each agent’s chromosome type determines her strategic behavior in the market place (e.g. Arifovic, 1996; Dawid, 1999; Arifovic and Gencay, 2000; Lux and Schornstein, 2005). When the genetic operations are applied to this pool of trader-chromosomes, information is effectively shared and incorporated into the entire new generation via the evolutionary dynamics. This design is certainly at odds with the set-up of the experimental market in which subjects are separated from each other and are not allowed to actively exchange information. We, therefore, assumed that each agent in our computer experiment had his own pool of strategies or forecasting rules which undergo their genetic evolution independently from the rules of other agents (cf. LeBaron et al., 1999, for a similar approach).<sup>8</sup> In our experiments reported below, we endowed each

---

<sup>8</sup>Vriend (2000) discusses differences between social learning and individual learning in agent-based models. He shows that in the context of a cobweb model, social learning converges to a competitive equilibrium since the Cournot strategy is not stable against deviations which then disseminate in the population. In contrast, with individual learning the firms in the artificial cobweb model converge to a Cournot equilibrium as they only compete against other rules of the same individual and do not interact with those of other firms. While we agree with Vriend that this difference might extend beyond the confines of the cobweb market structure, his considerations are not applicable to our setting. The reason is the difference in incentive structures: profits in Vriend’s case versus pay-offs depending on forecast accuracy in our cases.

agent with  $M = 10$  different chromosomes encoding pairs  $(\alpha_i, \beta_i)$  of the first order autoregressive forecasting rule detailed below. The active rule of each agent, i.e. the rule on which her actual forecast was based, was determined by random draws with probabilities equal to the relative fitness obtained in the last round (which is a monotonic function of the proximity of the forecast to the realized price, cf. eq. 7).

Genetic algorithms require a functional specification of the forecasting rule, whose fitness-maximizing parameter values would then be searched for via the evolutionary algorithm.<sup>9</sup> The simplest specification of a rule would be a constant price forecast. A slightly more complex version would use a constant together with a first order autoregressive component:

$$p_{i,t+1}^e = \alpha_i + \beta_i(p_t - \alpha_i). \quad (8)$$

This first order autoregressive (AR1) rule seems a natural forecasting scheme as agents could simply implement it via a linear autoregression using the sample average as their estimate of  $\alpha_i$  and the first order sample autocorrelation as the estimate of  $\beta_i$ . Moreover, the AR1 forecasting rule (8) has a simple *behavioral* interpretation, with  $\alpha_i$  representing an anchor or observed

---

Even though a spite effect would be present for firms, this effect does not carry over to the agents whose task is to predict aggregate market prices. It is, therefore, separation of the task of forecasting from the production decision that eliminates the tendency towards Cournot equilibria that was identified for individually learning firms in Vriend's (2000) paper. If we conduct our analyses reported below with social learning of forecasters, we still get convergence to the RE equilibrium in means, but higher volatility and more predictability (in the form of zig-zag patterns). The reasons are that i) social learning amounts to an overall reduction of the pool of available rules, and ii) social learning leads to strong correlations of the behaviour of agents. If a rule that was successful at time  $t$  is taken over by many agents at time  $t + 1$ , this easily generates a tendency towards the zig-zag patterns similar to the case of naive expectations in the textbook version of the cobweb model.

<sup>9</sup>Genetic programs, in contrast, would allow for the evolution of arbitrary functional specifications using a set of basic functional elements (Chen and Wang, 2002). Since we are able to replicate the experimental stylized facts already with the simpler concept of GAs, we abstain from using the more intricate evolutionary dynamics of genetic programs.

average price level around which the market price fluctuates, and  $\beta_i$  representing the observed persistence or anti-persistence of price fluctuations<sup>10</sup>.

As discussed by Hommes (2009) a representative agent model where all agents employ the same fixed rule, e.g. the rule (8), or where all agents adopt the same adaptive learning scheme, e.g. sample autocorrelation or least squares learning, as a uniform learning mechanism for the whole population, can *not* explain all stylized facts of the experiments simultaneously. A homogenous adaptive learning rule either always enforces convergence to RE (i.e. does *not* explain the second observed stylized fact, the excess volatility in the strongly unstable case) or, in cases where the adaptive learning rule leads to excess volatility, it generates anti-persistent price behavior with significantly negative first order autocorrelation, violating the third stylized fact in the laboratory experiments. In our GA-model, we apply the same functional scheme in a *heterogenous agent framework* with genetically evolved sets of parameters  $\alpha_i, \beta_i$  that could differ across individuals. The key question then is whether the *interaction* of individual forecasting rules (8) can explain all stylized facts observed in aggregate price behaviour simultaneously.

In our simulations the two parameters  $\alpha_i$  and  $\beta_i$  are encoded in one string of length  $l = 40$ , the first (last) 20 bits representing the parameters  $\alpha_i$  and  $\beta_i$ , respectively (the number of bits is quite arbitrary and only needs to be large enough for a sufficiently fine-grained structure of the resulting real-valued strategies).  $\alpha_i$  is restricted to the interval  $[0,10]$  just as in the instructions to participants in the laboratory experiments. The interval for  $\beta_i$  is more arbitrary and had been set symmetrically around zero,  $\beta_i \in [-1, 1]$ , allowing for quite strong serial correlation or anti-correlation up to the point where individual rules would lead to an unstable dynamics,  $|\beta_i| = 1$ . We have also tried wider intervals such as  $[-2, 2]$  for the autoregressive coefficients. Results were qualitatively similar, but not quantitatively as close to the laboratory experiment as the ones reported below. The transition from the binary coded evolutionary process to the real-valued forecasts requires to compute:

---

<sup>10</sup>In similar cobweb type laboratory experiments Heemeijer et al. (2009) recently estimated individual forecasting rules, and many individuals actually used forecasting rules of the simple form (8).

$$\alpha_{i,t} = 10 \sum_{j=1}^{20} a_{i,t}^j \frac{2^{j-1}}{2^{20} - 1}, \quad \beta_{i,t} = -1 + 2 \sum_{j=21}^{40} a_{i,t}^j \frac{2^{j-21}}{2^{20} - 1} \quad (9)$$

with  $a_{i,t}^j \in \{0, 1\}$  the bits at position  $j$  ( $j = 1, \dots, 40$ ) of chromosome  $i$  at time  $t$ .<sup>11</sup>

It is well known that short run GA-simulations may be sensitive to the initialization of the GA's. Therefore, in our (short run) simulations we have chosen an initialization in line with individual forecasts in the first and the second period of the experiment, as discussed in more detail below.

## 4 Experiments with Genetic Algorithms

In this section we report the results from GA simulations. Unless reported otherwise, each of the six agents will be endowed with  $M = 10$  chromosomes

---

<sup>11</sup>We have also checked the robustness of our results by running the experiments below with real-coded GAs. Although it might appear more plausible to encode an initially real-coded problem (like the present one) with real rather than binary 'chromosomes', the latter approach is more popular in the literature. For real coded GAs the formulation of a sensible cross-over operation is particularly cumbersome (cf. Herrera et al., 1998). With two real-valued parents, say  $(\alpha_{1,t}, \beta_{1,t})$  and  $(\alpha_{2,t}, \beta_{2,t})$ , it seems straightforward to define cross-over as the random selection from the interval spanned by the two parents. However, this approach would generate a strong mechanical tendency towards the mean of the support and would often lead to a degenerate population. In order to compensate for this tendency, one would have to extend the interval by a certain factor to the right and left of the parents' strategic values in order to give the off-spring a chance to escape from this range. Mutation operators for real-valued GAs typically use random draws from a Normal distribution with small variance to perturb the values inherited from the parents. While this leads to close proximity of the strategies of parents and offspring, it also implies that the influence exercised by the mutation operator is typically much smaller in real-valued than in binary GAs (where the mutation could be arbitrarily large depending on what bit is modified). With these features in mind, we have replicated our experiments with a real-coded GA with extended range for cross-over. Again, results are pretty similar to those reported in main part of this paper. For initialisation, we have used random draws from a Normal distribution with the mean and standard deviation computed from the initial forecasts of the experiments.

and the crossover probability  $p_{cross} = 0.6$  (but different values yield similar results). Genetic learning would converge to the rational expectations equilibrium if uniformly across the population all  $\beta_{i,t}$  would tend to zero and the  $\alpha_{i,t}$  would converge to the RE price  $p^*$ . Since the experiments run over a limited number of rounds, an appropriate alignment of our GA simulation with the lab settings is required. Note that in order to start the evaluation of the fitness of agents' strategies, we need two realizations of the market price: the first one serves as the anchor value for the AC strategies in eq. (8) and the subsequent realization serves to evaluate the quality of the AC forecast using eq. (9). As a consequence, evolutionary strategies could be evaluated for the first time at  $t = 3$ . In order to align the GAs to the lab experiments, we therefore, choose for periods 1 and 2 forecasts and prices from the experiments while the GA population of each agent is initialized randomly.<sup>12</sup> In this way, the 'initial conditions' of the GAs are set equal to those of the lab experiments and our artificial agents initially are subject to the same incentives like the human subjects. As it turns out, this alignment typically guarantees greater similarity than, say, a randomized choice of forecasts at  $t = 1$  and  $t = 2$  (while still qualitative results are pretty much the same under different initialization schemes). Obtaining closer proximity by accurate alignment should be seen as an encouraging finding: It is worth emphasizing that this is not a fine-tuning of our algorithm but rather an attempt at matching as closely as possible the experimental scenario.

In our simulations, consistent with the laboratory experiments, we find that the market price fluctuates around the RE benchmark with a sample mean very close to the RE benchmark, but that the level of volatility depends

---

<sup>12</sup>We have also conducted experiments with asymmetric initialization. As it turned out, results were completely insensitive to the initial distribution up to extreme degrees of asymmetry. E.g., results from a binary GA remained practically unchanged even if we fixed the first 16 (out of 20) digits for  $\alpha_i$  and  $\beta_i$  to zero or one. For the first 16 digits equal to zero that means that we drew  $\alpha_i$  and  $\beta_i$  from the reduced ranges  $[0, 0.0024]$  and  $[-1, -0.4495]$ , respectively. Apparently, the evolutionary forces exerted by the first two realizations (that are taken from the experiments) are strong enough to overcome completely the influence of such a very biased distribution of initial forecast parameters. A bit of an influence of the initial distribution can be seen for even higher numbers of fixed bits. However, it seems perfectly plausible that evolution would need longer to modify an almost completely degenerate set of initial conditions.

strongly on the treatment. In the stable case, the sample variance is close to its rational expectations benchmark, but it increases significantly beyond the RE benchmark if we proceed to the ‘unstable’ and ‘strongly unstable’ scenarios. Fig. 1 shows snapshots from longer simulation runs and Table 1 summarizes some key statistics, for all three treatments, averaged over 1,000 simulations of 50 periods each. The sample mean of individual forecasts ( $\text{Mean}(p^e)$ ) has been obtained by averaging the individual forecasts over all subjects ( $K=6$ ) and all experiments ( $J = 6$ ) for each treatment. The sample variance of individual forecasts ( $\text{Var}(p^e)$ ) has been computed as follows. Let  $\hat{p}_{t,k}^j$  be the price forecast for time period  $t$ , by subject  $k$ , in experiment  $j$ , then the mean forecast for period  $t$  in experiment  $j$  is  $\mu_t^j = \frac{1}{K} \sum_k \hat{p}_{t,k}^j$ , with  $K = 6$  in the experiments. The sample variance of this mean forecast over all rounds ( $T = 50$ ) of experiment  $j$  is given by

$$\text{Var}^j(p^e) = \frac{1}{T-1} \sum_t (\mu_t^j - \frac{1}{T} \sum \mu_t^j)^2. \quad (10)$$

The sample average of individual forecasts can then be obtained by averaging over all experiments ( $J = 6$ ) or over all simulations ( $J = 1,000$ ) respectively, within each treatment:<sup>13</sup>

$$\text{Var}(p^e) = \frac{1}{J} \sum_j \text{Var}^j(p^e). \quad (11)$$

In a similar vein, the variance of realized prices,  $\text{Var}(p)$  has been computed according to (11), averaging over  $J = 1,000$  simulations. Table 1 shows the statistics for the GA simulations, the laboratory experiments and the RE benchmark<sup>14</sup>. The table shows that the GA-simulations are surprisingly

---

<sup>13</sup>There are other ways of defining the sample mean and sample variance of individual forecasts. For example, one can pool the forecasts for each period  $t$  over all individuals  $i$  and all experiments  $j$ , compute the variance of the pooled forecasts and then average over all rounds. Alternatively, one can start out with the variance over the individual forecasts for each experiment  $j$ , for each period  $t$  and then average over all experiments and all rounds. The results for these alternative ways of averaging of individual forecasts are very similar to those reported below. In particular, independently of the details of the averaging method, the GA-learning simulations match the laboratory experiments quite nicely.

<sup>14</sup>For the RE benchmark,  $\text{Var}(p^e)$  has been computed as the variance of the average of individual forecasts drawn randomly from the RE stochastic process  $p^* + \epsilon_t$  in (6), yielding  $\text{Var}(p^e) = \sigma_{\epsilon_t}^2 / K \approx 0.042$  for  $K = 6$ .



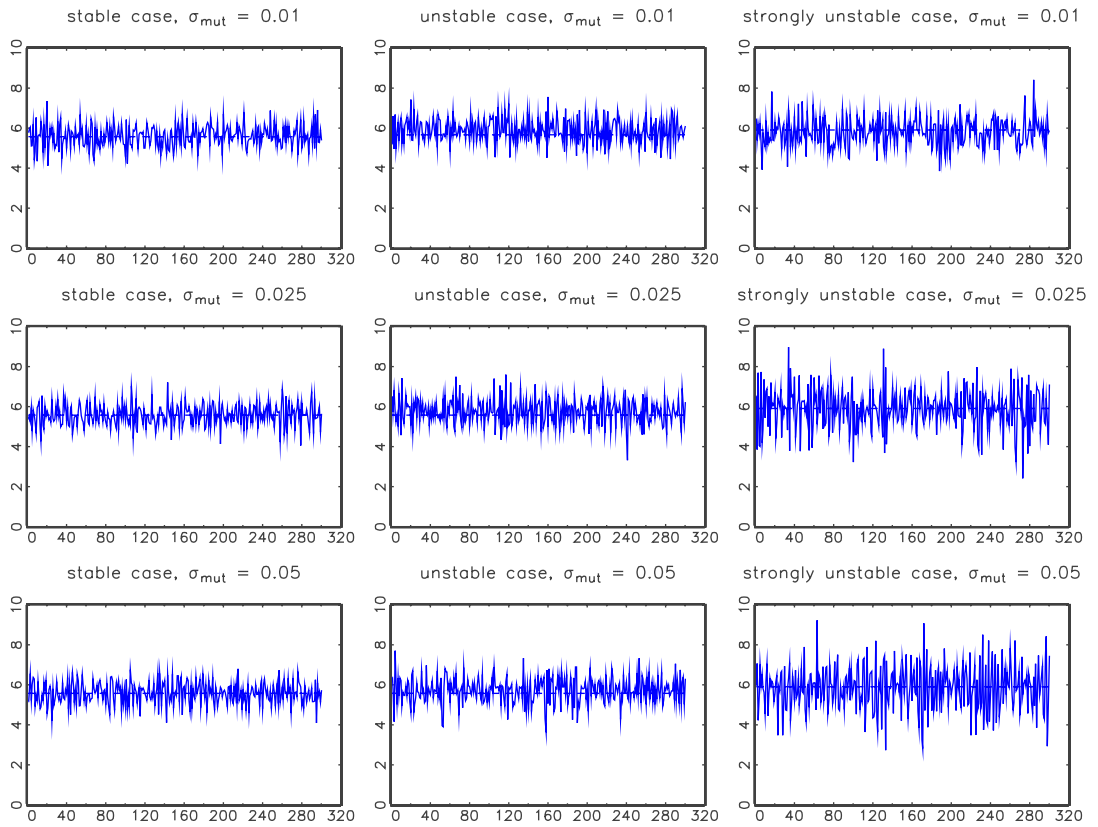


Figure 1: Snapshots from simulations of GA-learning: realized prices (solid) and RE benchmarks (broken lines) for all three treatments (stable, unstable and strongly unstable) and three different values of the mutation probability  $\sigma_{mut}$ .

Table 1: Sample means and sample variances of individual expectations and realized market prices

	Mean( $p^e$ )	Mean(p)	Var( $p^e$ )	Var(p)
<b>Stable case (K=6)</b>				
RE	5.57	5.57	0.042	0.25
Experiments	5.56	5.64	0.087	0.36
$p_{mut} = 0.01$	5.576	5.565	0.100	0.326
$p_{mut} = 0.025$	5.572	5.569	0.095	0.320
$p_{mut} = 0.05$	5.571	5.568	0.096	0.319
<b>Unstable case (K=6)</b>				
RE	5.73	5.73	0.042	0.25
Experiments	5.67	5.85	0.101	0.63
$p_{mut} = 0.01$	5.645	5.817	0.169	0.647
$p_{mut} = 0.025$	5.662	5.791	0.155	0.618
$p_{mut} = 0.05$	5.675	5.773	0.151	0.611
<b>Strongly unstable case (K=6)</b>				
RE	5.91	5.91	0.042	0.25
Experiments	5.73	5.93	0.429	2.62
$p_{mut} = 0.01$	5.434	6.200	0.769	2.161
$p_{mut} = 0.025$	5.587	6.027	0.625	2.241
$p_{mut} = 0.05$	5.653	5.914	0.614	2.472
<b>Strongly unstable case (K=12)</b>				
RE	5.91	5.91	0.021	0.25
Experiments	5.781	5.937	0.204	1.783
$p_{mut} = 0.01$	5.515	6.183	0.500	1.571
$p_{mut} = 0.025$	5.637	6.031	0.401	1.576
$p_{mut} = 0.05$	5.695	5.931	0.368	1.643

*Notes:* all parameters have been chosen exactly as in Hommes et al. (2007), i.e., there are  $K = 6$  GA agents whose task is to forecast the next period's price. Crossover probability is 0.6 and the election operator is applied. The first and second moments are computed from 1,000 runs of 50 periods each (i.e., using 50,000 observations). The last five rows of the table correspond to the experiments with  $K = 12$  in van de Velden (2001) and simulations with  $K = 12$  GA agents.

close to the laboratory experiments across all treatments. Besides the three treatments of the laboratory experiments, we also distinguish between different settings for the mutation probability,  $p_{mut} = 0.01, 0.025$  and  $0.05$ , as this appears to be the more interesting aspect of the GA design. As can be seen, price fluctuations also increase *ceteris paribus* with higher mutation probability due to the higher rate of new forecasting rules entering the population. Like in other applications of GAs (cf. Lux and Schornstein, 2005) varying other parameters as well as choosing different specifications of the operators appears to cause no major changes in the overall outcome.

Fig. 2 shows the autocorrelations of prices for the nine scenarios under investigation. All autocorrelations are small, with the first one or two lags exhibiting small negative values, in nice agreement with the laboratory findings (cf. the autocorrelation plots of realized prices in the experiments, Hommes et al., 2007, Fig. 5, pp.21-22). The slight increase of the autocorrelation at the first few lags for a higher mutation probability is simply due to the mechanics of a population with high rate of change as the random elements invoke a mean reverting tendency towards the average of the population.

Fig. 3 shows that even under the ‘strongly unstable’ scenario, the population mean values of the AR1 parameters  $\alpha_i$  and  $\beta_i$  are close to their benchmark values under rational expectations:  $\alpha_i$  fluctuates in the vicinity of the RE equilibrium price ( $p^* = 5.91$ ) while the average  $\beta_i$  is close to zero. The same applies in the other cases. However, fluctuations around the RE benchmark are stronger for the unstable and strongly unstable cases which leads to stronger fluctuations and excess volatility of market prices consistent with the laboratory experiments.

Overall, the GA experiments with the same parameter setting and incentive structure as in the laboratory experiments appear to closely mimic the set of stylized facts above. We note that the qualitative outcome was quite robust under various modifications of the GA learning mechanism. For example, we get similar results when dispensing with the election operator. The main difference in this setting is that fluctuations in the ‘unstable’ and ‘strongly unstable’ treatment become more pronounced and that we see somewhat more significant negative autocorrelation in the first few lags: as mentioned above, this feature can be easily explained by the mean-reverting nature of

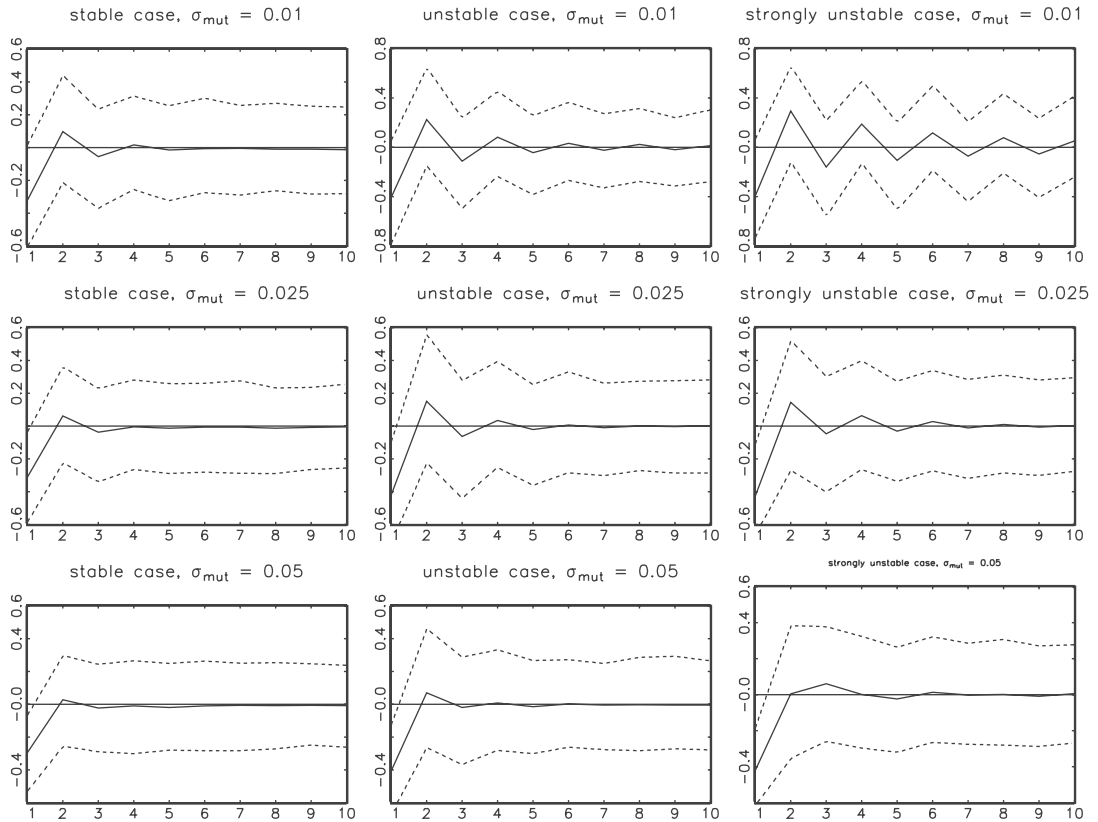


Figure 2: Autocorrelations of realized prices under GA-learning: the process only shows slight negative correlations over one or two lags (consistent with the laboratory experiments). The plots show the means and 95% confidence intervals from 1,000 GA-simulations each extending over 50 periods.

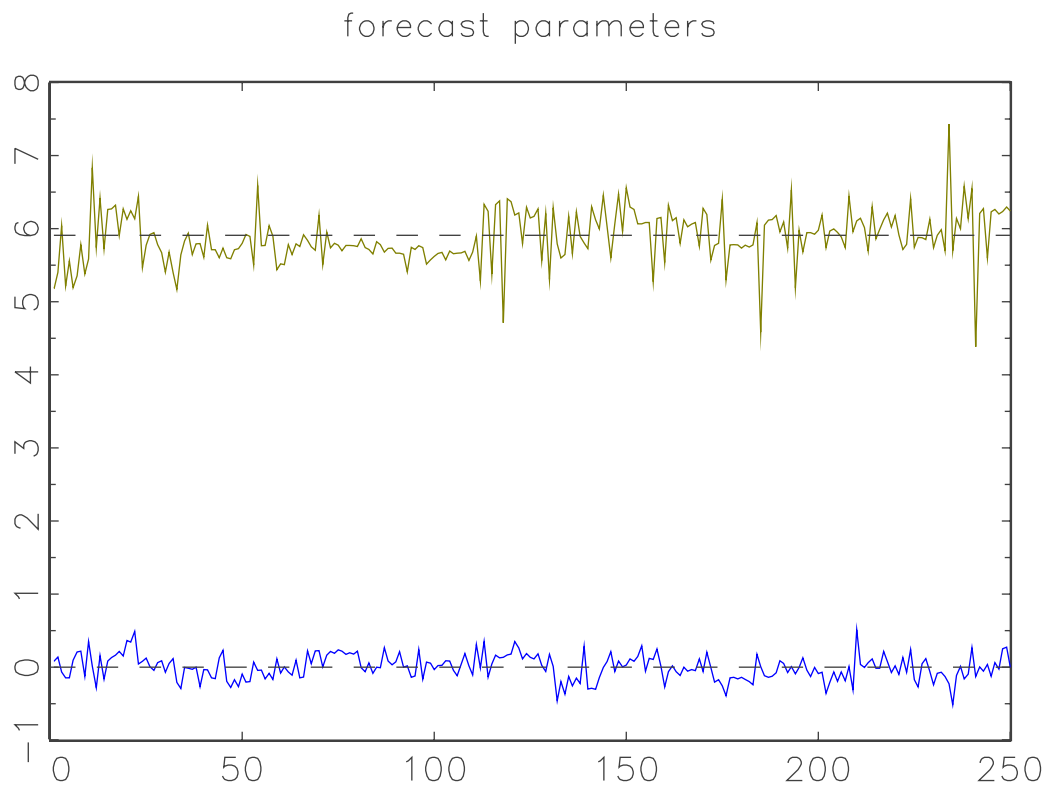


Figure 3: Snapshots of the development of average AR1 parameters  $\alpha_i$  and  $\beta_i$ , together with their RE benchmarks  $\alpha_i = p^*, \beta_i = 0$  (broken lines), under GA-learning in the ‘strongly unstable’ treatment.

the dynamics with more random mutations admitted to the population. On economic grounds we might, however, argue that agents should not allow obviously unsuccessful strategies to enter their set of forecast functions (which is why the election operator had been introduced in economic settings) so that we would not place too much weight on these results.

Interestingly, adopting a simpler concept of learning, that dispenses with the AC parameter  $\beta_i$  and restricts forecast rules to a constant  $\alpha_i$  also leads to results that share some of the stylized facts. While this scenario leads to similar outcomes for volatility in the three treatments, it is, however, also characterized by anti-persistent price behavior and more significant zigzag patterns of autocorrelations in the strongly unstable case. The simplest GA with only constant rules, therefore, seems to inherit at least part of the oscillating dynamics of the benchmark case of homogenous naive or adaptive expectations. A more intelligent type of forecast rules such as our AR1 rules, taking into account both the average price level and first order autocorrelation, is required to remove linear forecastability. Stated differently, individual learning of the mean alone is not consistent with the laboratory experiments, but some more sophisticated form of individual learning taking into account whether prices are persistent or anti-persistent is needed to remove autocorrelations in aggregate prices. The interaction of individual rules, learning both the price level and the first order autocorrelation, leads to the correct aggregate price level and washes out all autocorrelations in aggregate price fluctuations.

Why do the GA experiments reproduce the experimental results so well? Our conjecture is that GAs and human subjects share the tendency of “learning by experience” and of shifting their strategies towards a specification that would have performed well in the recent past. This is actually the consequence of the built-in genetic operators of GAs. While this leads to an optimal solution for static problems (at least, if the evolutionary process is allowed to run long enough), with the changing objective functions of an interactive environment it could also lead to repetitive patterns (cf. Lux and Schornstein, 2005). In the absence of structural knowledge about the underlying mechanisms of a decision problem, humans can also still determine what actions or decisions would have performed well in the past. Quite clearly, the laboratory experiments with their unknown forms of the

underlying market functions and added stochasticity could not have been fully penetrated by the experimental subjects. However, they could easily focus on the past success and failure of their forecasts and learn to maintain successful strategies. Exploiting the mean together with short-run autocorrelations seems to be one of the simplest strategies that could be pursued in a rule-of-thumb manner without computing the sample autocorrelation exactly (which would normally not be possible given the time pressure of most experimental settings). These rough computations lead to stochastic fluctuations that are similar to the fluctuations caused by the evolutionary dynamics of the GA. The latter feature distinguishes our heterogeneous learning scenario from homogenous learning models (Hommes, 2009) which seem unable to explain the full set of stylized facts. Heterogeneity in individual forecasting thus seems to be a key element in explaining all stylized facts at the aggregate level simultaneously.

In summary, our conjecture is that the orientation at successful performance in the past within a reasonable class of forecasting heuristics together with the heterogeneity of the GA design explains its proximity to human behaviour in the lab. We may note that such conformity has also been found in a number of other cases, e.g. in an experimental foreign exchange market (Arifovic, 1996) or public good experiments (Casari, 2004). It is also related to the work of Erev and Roth (1999) on reinforcement learning to explain experiments with repeated games<sup>15</sup>. Anufriev and Hommes (2009,

---

<sup>15</sup>We also investigated a reinforcement learning (RL) algorithm with similar strategy sets as in our GA setting. Individuals were endowed with AC strategies with parameter space  $\alpha_i \in [0, 10]$  and  $\beta_i \in [-1, 1]$ . Initial parameters were drawn randomly from a uniform distribution and updated with probabilities computed via relative fitness (as in the GA experiments) with fitness defined by the payoff function (5). We distinguished between a myopic RL algorithm, only using the last payoff as fitness function, and a full memory RL scheme that computed fitness as the arithmetic average of all previous payoffs. Qualitative results of both settings were not too different, however. The basic outcome of these RL experiments was that (i) agents never converged to the RE benchmark. In particular, the variance was always above the RE benchmark, even in the ‘stable’ scenario. Mean values were slightly further away from RE prices than under GA learning, (ii) in all cases, realized market prices showed pronounced cyclical patterns indicating that RL agents were not able to exploit all linear structure, (iii) results seemed to be entirely insensitive to the number of strategies (we allowed the strategy set of agents to vary from  $M = 50$  over 500 up to 5000). Detailed results are available upon request.

2012) have recently used another form of reinforcement learning to explain learning to forecast experiments in a different, asset pricing framework.

## The degree of heterogeneity

In the stable treatment of the laboratory experiments it seems that prices converge to RE and agents learn to coordinate on the RE forecast. In contrast, in the (strongly) unstable treatment prices do not converge, but exhibit excess volatility. Does forecast heterogeneity disappear in the stable treatment and does persistent forecast heterogeneity explain the observed excess volatility in the (strongly) unstable treatment? Whether or not heterogeneity persists is important for economic theory. If beliefs converge to a common rule, long run aggregate price behavior can be described by a representative agent model. If on the other hand beliefs do not converge, heterogeneous expectations models become relevant as a description of short run as well as long run aggregate market behavior. See Hommes (2006) for an extensive discussion of heterogeneous expectations models.

Figure 4 (top panel) shows some typical examples of time series of the six individual forecasts in one group for each of the three treatments. These examples already suggest that heterogeneity quickly disappears in the stable treatment, while heterogeneity is highly persistent in the strongly unstable treatment. Figure 4 (bottom panel) also shows the average degree of heterogeneity, that is, the time development of the standard deviations of individual forecasts ( $K = 6$  individuals per group) averaged over all groups in each treatment. More precisely, let  $\hat{p}_{t,k}^j$  be the price forecast for time period  $t$ , by subject  $k$ , in experiment  $j$ , then the mean forecast for period  $t$  in experiment  $j$  is  $\mu_t^j = \frac{1}{K} \sum_k \hat{p}_{t,k}^j$ . The standard deviation of the mean forecast at date  $t$ , over  $K = 6$  subjects of experiment  $j$  is

$$\sigma_t^j = \sqrt{\frac{1}{K-1} \sum_k (\hat{p}_{t,k}^j - \mu_t^j)^2}. \quad (12)$$

The *average degree of heterogeneity* at date  $t$  over all experiments ( $J = 6$ ) within a treatment is then defined as the average standard deviation

$$\sigma_t(p^e) = \frac{1}{J} \sum_j \sigma_t^j. \quad (13)$$



The time development of the average degree of heterogeneity in Figure 4 (bottom panel) exhibits two important features: (1) for all treatments forecast heterogeneity decreases over time, and (2) forecast heterogeneity is persistent in the unstable treatment and highly persistent in the strongly unstable treatment.

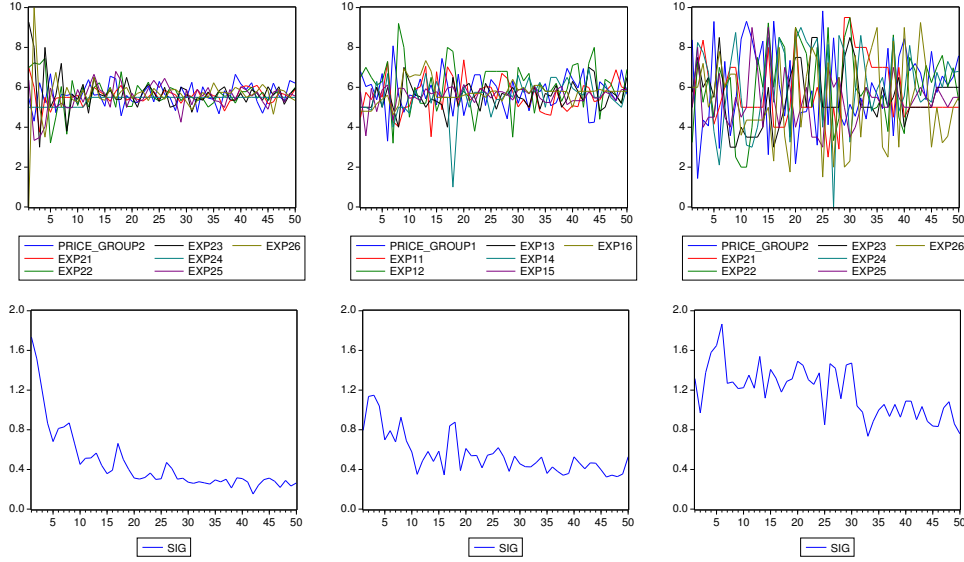


Figure 4: Top panel: Six individual forecasts in one group of the laboratory experiments in the stable (left), unstable (middle) and strongly unstable (right) treatments. Bottom panel: Time development of the average degree of heterogeneity, i.e. the standard deviations of individual forecasts (six individuals) averaged over all (six) groups in the stable treatment (left), the unstable (middle) and the strongly unstable treatment (right).

Figure 5 shows the time development of the average degree of heterogeneity in GA learning simulations in the first 50 periods, averaged over 1000 runs, in the stable, the unstable and the strongly unstable treatments. This figure shows that GA-learning simulations reproduce the patterns of the average degree of heterogeneity in the laboratory experiments quite nicely: a quick decrease of forecasting heterogeneity in the stable treatment and a much slower decrease in the (strongly) unstable treatment. In fact, both in the experiments and the GA learning simulations the unstable and strongly unstable treatments exhibit a non-monotonic development of forecasting heterogeneity with an increase in heterogeneity in the early stage of the

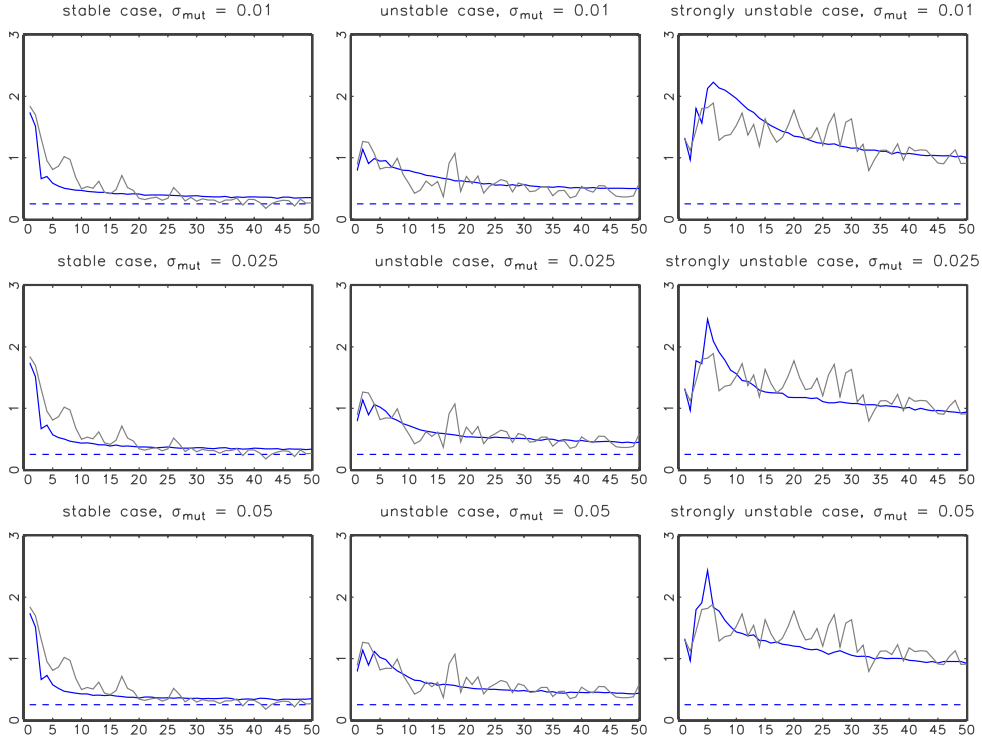


Figure 5: Time development of the average degree of heterogeneity, i.e. the standard deviations of the six individual forecasts in GA learning simulations (dark line) averaged over 1000 runs, in the stable (left), the unstable (middle) and the strongly unstable treatment (right) for different values of the mutation probability  $p_{mut}$ . The time series of the average degree of heterogeneity in the corresponding experiments (light line) as well as the RE benchmark (dotted line) are also shown.

experiment/simulations due to overshooting and a decrease in heterogeneity after periods 5 – 7 due to learning.

## 5 Beyond the Laboratory Setting

In contrast to experiments with human subjects, additional experiments with genetic algorithms can be conducted at essentially zero cost. In this section we expand our previous experiments into various directions not covered by the laboratory experiments. Among others, we investigate long run price behaviour, how price behaviour depends upon parameter values, in particular the parameter tuning the nonlinearity of the supply curve, and we investigate the consequences of an increase of the number of agents and forecasting strategies in the GA populations.

### 5.1 Long run behaviour

Table 2 summarizes the long run statistics for all three treatments and three different mutation probabilities,  $p_{mut} = 0.01, 0.025$  and  $0.05$ . As can be seen, in the stable case the long run average degree of heterogeneity,  $\text{Var}(p^e)$ , is small and price volatility is quite close to the RE benchmark 0.25. In the unstable treatment price volatility is slightly above the RE variance, while in the strongly unstable treatment the long run price variance is significantly above the RE benchmark due to a larger average degree of heterogeneity. The strongly unstable treatment thus exhibits persistent heterogeneity and long run excess price volatility. Moreover, both the average degree of heterogeneity and excess price volatility increases with the mutation probability. This seems intuitively clear, since a higher mutation probability leads to a higher rate of new forecasting rules entering the population.

### 5.2 Parameter sensitivity

In the next set of simulations we explore the transition between the ‘nice’ price behaviour of the ‘stable’ treatment and the excessive price volatility

Table 2: Long-run simulations

	Mean( $p^e$ )	Mean(p)	Var( $p^e$ )	Var(p)
<b>Stable case (K=6)</b>				
RE	-	5.57	-	0.25
Experiments	5.56	5.64	0.087	0.36
$p_{mut} = 0.01$	5.580	5.567	0.012	0.257
$p_{mut} = 0.025$	5.585	5.564	0.018	0.262
$p_{mut} = 0.05$	5.589	5.557	0.030	0.274
<b>Unstable case (K=6)</b>				
RE	-	5.73	-	0.25
Experiments	5.67	5.85	0.101	0.63
$p_{mut} = 0.01$	5.728	5.731	0.009	0.283
$p_{mut} = 0.025$	5.734	5.724	0.016	0.306
$p_{mut} = 0.05$	5.737	5.708	0.028	0.343
<b>Strongly unstable case (K=6)</b>				
RE	-	5.91	-	0.25
Experiments	5.73	5.93	0.429	2.62
$p_{mut} = 0.01$	5.889	5.869	0.012	0.475
$p_{mut} = 0.025$	5.881	5.826	0.050	0.824
$p_{mut} = 0.05$	5.848	5.768	0.173	1.516

*Notes:* Long run simulations with  $K = 6$  GA agents (chromosomes) and different mutation probabilities  $p_{mut}$ . The first and second moments for market prices,  $\text{Mean}(p^e)$  and  $\text{Var}(p^e)$ , are computed from simulations over 50,000 time steps after discarding the first 10,000 observations as transient sample.  $\text{Mean}(p^e)$  is the mean over the whole simulation of the average forecast across the 6 ‘agents’ in each period. The average degree of heterogeneity,  $\text{Var}(p^e)$  has been computed according to (10), averaged over  $T = 50,000$  periods after a transient of 10,000 periods.

of the unstable treatments. Recall that the difference between these treatments is the parameter  $\lambda$  tuning the nonlinearity of the supply curve. We ran the same type of GA experiments with 800 different values of the slope parameters  $\lambda$  ranging from 0.005 to 4 (with increments of 0.005). Fig. 6 reports the mean values and variances of realized prices over 50,000 rounds together with their RE benchmark. It turns out that the variance of realized prices is close to its RE benchmark of 0.25 only for very small values of  $\lambda$  with an almost perfectly linear increase with  $\lambda$  thereafter.<sup>16</sup> In contrast, the average price stays close to its RE benchmark over the whole range of our experiments. While there appears to be a slightly increasing wedge between average price and the RE solution for increasing  $\lambda$  the deviation is always very small compared to the difference between the realization of the second moment and its RE benchmark. We conjecture that this increasing wedge might be more an artifact of our simulation design than a ‘true’ indication of (small) deviations of the mean price from the rational expectations price. Since  $p_{RE}$  is slightly above the center of the admissible range  $[0, 10]$  larger fluctuations would generate some asymmetries in realized prices with a slight dominance of lower rather than higher prices. The slight deviation from RE in the first moment (which remains smaller than 2% percent in all scenarios) would then be a numerical consequence of the large deviation in the second moment from its RE benchmark.

### 5.3 Simulations with many agents

Another set of Monte Carlo runs investigates what happens if we increase the pool of participants in our forecasting experiments. While laboratory settings are typically restricted to small numbers of agents due to technical restrictions, the availability of subjects and the costs of running large experiments, we can easily extend our previous GA setting to much larger numbers of artificial agents. Since we have normalized supply by dividing through the number of firms in eq. (1), the RE benchmarks for first and second moments remain the same for all population sizes,  $K$ .

---

<sup>16</sup>A regression of the variance of realized prices on the parameter  $\lambda$  over the second half of our experiments produced a slope parameter 0.12 and a constant 0.25 with  $R^2$  of the regression equal to 0.94. Note that the constant is equal to the variance in RE equilibrium.

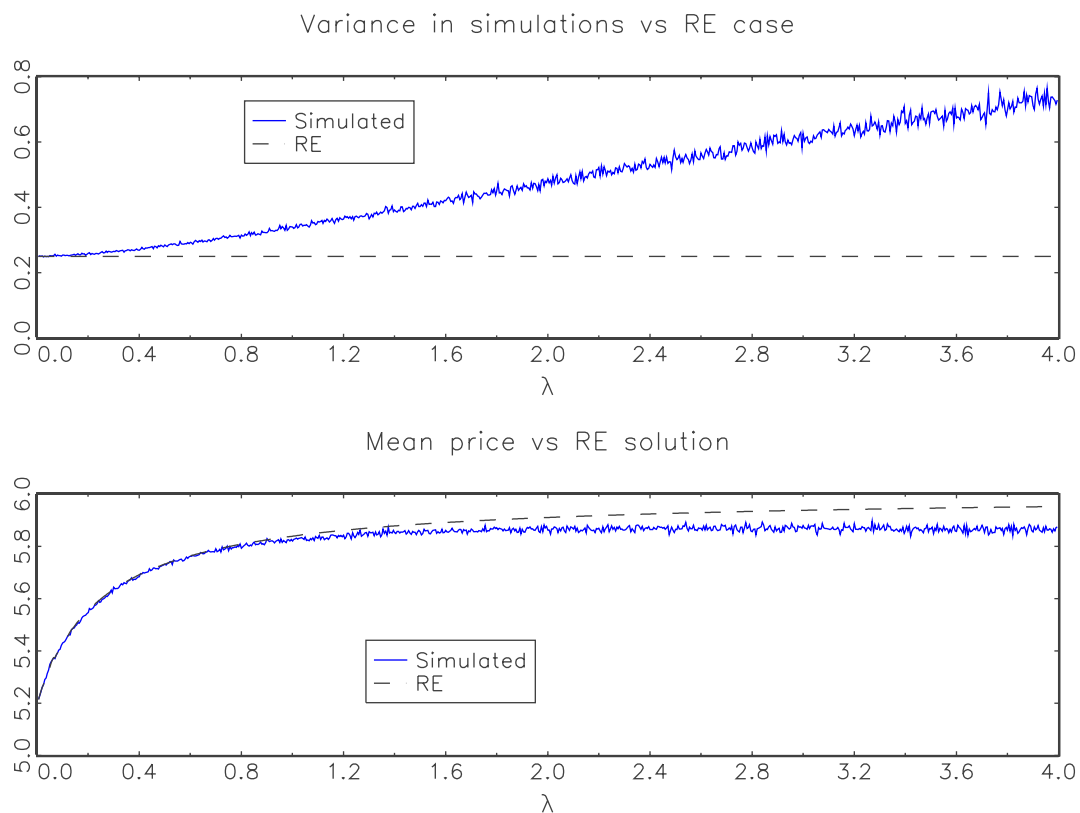


Figure 6: Mean prices (bottom panel) and variances (top panel) from GA simulations of markets with  $\lambda$  ranging from 0.005 to 4. Except for the variation of  $\lambda$ , all parameters are the same as before. Results are sample moments over simulations of 50,000 rounds (after discarding a transient period of 10,000 rounds).

Table 3 compares the results for population numbers  $K \in \{6, 12, 30, 100, 150, 600\}$ . Initialization of the GA simulation is done based on the first and second period individual forecasts in the experiments with  $K = 6$  subjects in Hommes et al. (2007) as well as the experiments with  $K = 12$  subjects in the strongly unstable treatment in van de Velden (2001). For the first period, all experiments (with 6 or 12 subject, with stable, unstable and strongly unstable treatments) have been pooled and the resulting distribution has been fitted with a normal  $N(5.271, 1.393)$ . Second period forecasts in the experiments differ between treatments, but are very similar for the 6 and 12 subjects cases in the strongly unstable treatment. We, therefore, pooled the forecasts of period 2 over all experiments of each treatment and fitted Normal distributions  $N(5.279, 1.698)$  for the stable treatment (only cases with 6 subjects),  $N(5.952, 1.266)$  for the unstable treatment (only cases with 6 subjects), and  $N(6.885, 1.225)$  for the strongly unstable treatment (pooled over all experiments with 6 or 12 subjects). The above experiments with  $K = 6$  to  $K = 600$  agents have been initialized by random draws from the pertinent Normal distribution, i.e., the same over all settings in period 1 but the treatment-dependent ones for period 2.

Apparently, higher numbers of agents have a tendency to dampen fluctuations. While there is not much difference in the experiments with stable slope parameter  $\lambda = 0.22$ , the effect is more pronounced in the unstable and strongly unstable scenarios. The stable case stays close to the RE benchmark for all sizes of the population with the variance of price fluctuations close to the variance of the random term. In the other cases, the excess fluctuations are clearly reduced when increasing the size of the population. In the ‘unstable’ case ( $\lambda = 0.5$ ) price volatility decreases from  $\text{Var}(p) = 0.629$  for  $K = 6$ , to  $\text{Var}(p) = 0.373$  for  $K = 600$ . In the strongly unstable case ( $\lambda = 2$ ) price volatility is reduced by about 40% for  $K = 600$ , but is still significantly higher than the RE benchmark ( $\text{Var}(p) = 1.235$ ). In the strongly unstable treatment in the short run, i.e. for the first 50 periods, excess volatility thus persists when increasing the number of agents to  $K = 600$ .

Another striking feature of these GA-simulations is that an increase of the number of agents beyond  $N = 30$  has little effect upon aggregate behavior. In all treatments, price volatility and the average degree of heterogeneity

Table 3: Short run simulations for increasing number of subjects

	Mean( $p^e$ )	Mean( $p$ )	Var( $p^e$ )	Var( $p$ )
<b>Stable case</b>				
$K = 6$	5.558	5.581	0.100	0.325
$K = 12$	5.557	5.584	0.075	0.308
$K = 30$	5.559	5.584	0.061	0.302
$K = 100$	5.558	5.583	0.054	0.292
$K = 150$	5.557	5.580	0.053	0.293
$K = 600$	5.558	5.583	0.051	0.293
<b>Unstable case</b>				
$K = 6$	5.665	5.784	0.159	0.629
$K = 12$	5.679	5.774	0.085	0.474
$K = 30$	5.685	5.766	0.059	0.416
$K = 100$	5.691	5.764	0.047	0.386
$K = 150$	5.694	5.763	0.044	0.376
$K = 600$	5.694	5.764	0.043	0.373
<b>Strongly unstable case</b>				
$K = 6$	5.605	6.011	0.625	2.297
$K = 12$	5.623	6.046	0.445	1.716
$K = 30$	5.643	6.055	0.361	1.429
$K = 100$	5.660	6.046	0.314	1.264
$K = 150$	5.659	6.047	0.313	1.260
$K = 600$	5.663	6.044	0.301	1.235

*Notes:* Effects of the variation of the number of agents,  $K$ . Other parameters are  $M = 10$  and  $p_{mut} = 0.025$ . Increasing  $K$  leads to convergence towards the RE benchmark in the stable and unstable treatment, but excess volatility persists in the ‘strongly unstable’ treatment. The first and second moments are computed from 1,000 runs with 50 periods each (i.e., using 50,000 observation).



Table 4: Long run simulations for increasing number of subjects

	Mean( $p^e$ )	Mean( $p$ )	Var( $p^e$ )	Var( $p$ )
<b>Stable case</b>				
$K = 6$	5.587	5.564	0.019	0.265
$K = 12$	5.582	5.566	0.014	0.262
$K = 30$	5.582	5.566	0.011	0.257
$K = 100$	5.581	5.562	0.009	0.255
$K = 150$	5.582	5.564	0.010	0.256
$K = 600$	5.583	5.565	0.010	0.256
<b>Unstable case</b>				
$K = 6$	5.733	5.721	0.016	0.307
$K = 12$	5.732	5.722	0.010	0.285
$K = 30$	5.734	5.722	0.007	0.277
$K = 100$	5.732	5.723	0.006	0.273
$K = 150$	5.733	5.720	0.006	0.270
$K = 600$	5.732	5.721	0.006	0.271
<b>Strongly unstable case</b>				
$K = 6$	5.882	5.834	0.047	0.797
$K = 12$	5.892	5.868	0.014	0.513
$K = 30$	5.895	5.876	0.006	0.399
$K = 100$	5.896	5.883	0.004	0.354
$K = 150$	5.896	5.882	0.004	0.353
$K = 600$	5.896	5.882	0.003	0.346

*Notes:* Long run statistics when the number of agents,  $K$ , increases. Other parameters are  $M = 10$  and  $p_{mut} = 0.025$ . Increasing  $K$  leads to convergence towards the RE benchmark in all three treatments. Even in the strongly unstable treatment long run market volatility decreases with  $K$ , to a value fairly close to the RE benchmark of 0.25. The first and second moments are computed using 50,000 observation, after disregarding the first 10,000 observations as transient.

drop significantly when the number of agents is increased from  $N = 6$  to  $N = 30$ , but hardly drop when the number of agents is further increased from  $N = 30$  to  $N = 600$ . This suggests that it may be possible to study macro phenomena in relatively small laboratory experiments with about 30 subjects, a size that is manageable in most experimental laboratories. See also the discussion about the relevance of laboratory experiments in macro in Duffy (2008).

Table 4 gives an overview of the same statistics in long run simulations, based on 50,000 periods, after a transient of 10,000 periods. Both the stable and the unstable treatments converge closely to the RE benchmark, with price volatility of  $\text{Var}(p) = 0.256$  and  $\text{Var}(p) = 0.271$  respectively for  $K = 600$ . Also the strongly unstable treatment approaches the RE benchmark relatively closely in the long run, although price volatility  $\text{Var}(p) = 0.399$  for  $K = 30$  and  $\text{Var}(p) = 0.346$  for  $K = 600$  respectively are still more than 35% above the RE-benchmark.

The laboratory experiments provide a simple and stylized framework that is stationary for 50 periods. In real markets with fluctuating prices, one would perhaps expect larger exogenous shocks to occur occasionally. From this perspective, what we have called the “short run”, i.e. the first 50 periods, could be more relevant to real markets than the “long run” where the market is stationary for a very long time.

The decrease of volatility with increasing population is probably easy to explain: adding more agents evokes a law of large numbers. Since our GA agents are effectively independent stochastic processes, their individual fluctuations should be averaged out when aggregating over more and more individuals.<sup>17</sup> This is what seems to happen in our experiments. Note, in particular the strong decrease of the variance of average price *expectations* in all settings. Of course, if price expectations would converge to the RE benchmark, realized prices would only fluctuate due to the exogenous random noise component.

---

<sup>17</sup>While there should always be a quasi-deterministic limit for the dynamics of GA populations, this need not necessarily lead to convergence towards some kind of steady state. Lux and Schornstein (2005) provide an example for how adaptive GA agents could converge to perfectly oscillatory dynamics due to their interactions in a large population.

## 5.4 Simulations with many rules

Finally, Table 5 reports results of GA-learning simulations where the number of agents has again been fixed to  $K = 6$ , while the number of rules,  $M$ , available to each agent increases from  $M = 4$  to  $M = 60$ .

Note that in a GA setting, the set of rules of an individual need not necessarily be different. In fact, convergence of the GA would imply that the population of rules of an agent becomes fully homogeneous. Increasing  $M$  thus does not necessarily mean that an individual has more different rules in each period, but it only increases the potential *sophistication* of the set of rules. We have again performed a comprehensive series of experiments varying the rules available to each individual from a low of  $M = 4$  to a maximum of  $M = 60$ . Over the short run, results are completely insensitive to the number of rules. Monte Carlo results along the lines of Table 1 yield practically no differences in first and second moments over this wide range of variation of  $M$  (detailed results are available upon request). However, notable differences show up if we look at the long run results, replicating our forecasting exercise with the GA agents over 50000 periods, cf. Table 5. As it turns out, at least in the ‘strongly unstable’ treatment, this higher sophistication leads to an increase of the volatility of realized (as well as expected) prices. With  $M = 60$  rules per agent, price volatility has almost doubled from  $\text{Var}(p) = 0.867$  to  $\text{Var}(p) = 1.433$ . In contrast, variation of the number of rules  $M$  seems to leave the results of the stable and unstable treatments almost unchanged. We conjecture that the higher number of chromosomes allows agents more easily to react on price fluctuations around the RE benchmark. With a high  $\lambda$ , under naive expectations a small deviation from  $p_{RE}$  would lead to a step-wise increase or decrease of the price for some time. Autocorrelation detection by some agents could reinforce this tendency as they would already forestall the direction of the subsequent price changes. With a large number of chromosomes, chances are increasing to evolve such momentarily advantageous rules. If such rules are admitted to the population, they would enhance fluctuations. It might, therefore, be a mixture of ‘naive’ adaptation of some agents (modifications of  $\alpha_i$ ) and trend chasing of others (adapting  $\alpha_i$  and  $\beta_i$ ) that generates the higher volatility in this case. Unfortunately, a systematic analysis of the interplay between the number of agents and their behavior in experimental

Table 5: Long-run simulations for increasing number of rules per agent

	Mean( $p^e$ )	Mean( $p$ )	Var( $p^e$ )	Var( $p$ )
<b>Stable case (K=6)</b>				
$M = 4$	5.576	5.569	0.034	0.277
$M = 6$	5.585	5.563	0.025	0.269
$M = 10$	5.579	5.567	0.020	0.266
$M = 20$	5.586	5.564	0.016	0.261
$M = 30$	5.582	5.568	0.015	0.261
$M = 40$	5.579	5.568	0.015	0.264
$M = 50$	5.576	5.574	0.015	0.260
$M = 60$	5.575	5.571	0.015	0.266
<b>Unstable case (K=6)</b>				
$M = 4$	5.731	5.714	0.024	0.327
$M = 6$	5.735	5.713	0.018	0.312
$M = 10$	5.731	5.725	0.016	0.305
$M = 20$	5.732	5.724	0.014	0.299
$M = 30$	5.732	5.724	0.013	0.295
$M = 40$	5.735	5.720	0.014	0.302
$M = 50$	5.736	5.719	0.013	0.296
$M = 60$	5.737	5.718	0.013	0.293
<b>Strongly unstable case (K=6)</b>				
$M = 4$	5.850	5.815	0.063	0.867
$M = 6$	5.870	5.821	0.046	0.783
$M = 10$	5.881	5.833	0.047	0.798
$M = 20$	5.874	5.834	0.083	0.962
$M = 30$	5.871	5.831	0.110	1.070
$M = 40$	5.862	5.827	0.144	1.194
$M = 50$	5.863	5.820	0.172	1.299
$M = 60$	5.850	5.817	0.213	1.433

*Notes:* Effects of the variation of the number of rules,  $M$ , per GA-agent. Other parameters are  $K = 6$  and  $p_{mut} = 0.025$ . Increasing  $M$  seems to leave the results practically unchanged in the ‘stable’ and ‘unstable’ cases, but increases both the volatility of predicted prices and realized market prices in the ‘strongly unstable’ scenario. The moments are extracted from simulations over 50,000 periods after a transient of 10,000 steps.

settings is beyond the limit of available laboratory resources. Given the autonomous adaptation of human subjects to different environments, it is not clear whether their learning behavior would remain unchanged in groups of different sizes. Our simulations suggest that, at least in the strongly unstable treatment, an increase of the number of learning rules per agent may be a potentially destabilizing force counterbalancing the stabilizing force of an increase in the number of agents in the market. It therefore seems possible, that changes of behavior might compensate for the law-of-large-number tendency in larger groups.

## 6 Concluding Remarks

Genetic algorithm learning of simple forecasting strategies provides an accurate description of *individual expectations* at the micro level and, at the same time, the *interaction* of these individual rules matches observed aggregate price behavior at the macro level surprisingly well. In the simple framework of the classical cobweb model, the interaction of individual GA-learning rules is able to reproduce all stylized facts in aggregate prices –correct sample mean, excess volatility depending on demand/supply characteristics and no linear predictability – observed in recent learning to forecast laboratory experiments with human subjects. In contrast to homogeneous learning rules, the interaction of heterogeneous GA-learning rules explains all stylized facts simultaneously and across various treatments. It should be emphasized that these results are robust and not sensitive to the GA-specification or the two GA-parameters (the mutation probability  $p_{mut}$  and the crossover probability  $p_{cross}$ ). The GA-algorithms attempt to learn two parameters –the sample mean and the first order autocorrelation coefficient– in a simple AR(1) forecasting rule. Evolutionary selection within a simple class of individual forecasting heuristics, that take into account both the observed sample mean and the first order sample autocorrelation, thus explains aggregate price behavior surprisingly well.

We have also looked at the average degree of heterogeneity of individual forecasting behaviour. In all treatments, heterogeneity decreases over time. In the stable treatment heterogeneity quickly disappears and the price set-

tles down to its RE benchmark. In the (strongly) unstable treatment heterogeneity decreases somewhat due to learning, but heterogeneity persists, even in the long run. These results suggest that economic theory needs to go beyond representative agent models with homogeneous expectations. The matching of our GA-simulation results with laboratory experiments are consistent with a theory of endogenous selection of heterogeneous expectations, for example, as in Evans and Ramey (1992), Brock and Hommes (1997) and, more recently, in Reis (2006).

Fitting a GA-learning model to the laboratory experiments allows one to go beyond experiments and simulate alternative and more realistic market environments. Through GA-simulations, we have seen that adding more agents to the market has a stabilizing effect, that is, price volatility decreases as the number of agents increases. However, in the strongly unstable treatment, excess price volatility persists when the number of agents becomes large. On the other hand, increasing the potential sophistication by allowing more strategies per individual has a destabilizing effect and makes price behavior more volatile. Additional laboratory experiments could reveal more information about the number of strategies subjects are using, in order to explore which of these two forces will dominate.

We have also seen that an increase in the number of agents beyond 30 has relatively little impact on aggregate price behavior. This suggests that laboratory experiments with 30 interacting subjects may reveal useful information about macro phenomena. Such larger macro experiments as well as applying GA's to other laboratory experiments, in particular other learning to forecast experiments, is a challenge for future work and may shed more light on individual expectation formation, their interaction and aggregate outcomes in alternative market settings.

## References

- [1] Adam, K., (2007), Experimental evidence on the persistence of output and inflation, *The Economic Journal* 117 (520) , 603–636.
- [2] Anufriev, M. and Hommes, C.H. (2009), Evolutionary selection of individual expectations and aggregate outcomes, *Cendef Working Paper* 09-09, University of Amsterdam.
- [3] Anufriev, M. and Hommes, C.H. (2012), Evolution of market heuristics, *Knowledge Engineering Review*, forthcoming.
- [4] Arifovic, J., (1994) Genetic algorithm learning and the cobweb model, *Journal of Economic Dynamics and Control* 18, 3-28.
- [5] Arifovic, J. (1996), The behavior of the exchange rate in the genetic algorithm and experimental economies, *Journal of Political Economy* 104, 510–541.
- [6] Arifovic, J. and R. Gencay (2000), Statistical properties of genetic learning in a model of exchange rate, *Journal of Economic Dynamics and Control* 24, 981–1005.
- [7] Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Tayler, P., (1997) Asset pricing under endogenous expectations in an artificial stock market, in Arthur, W., Lane, D. and Durlauf, S., (eds.) *The economy as an evolving complex system II*, Addison-Wesley, 1997, pp.15-44.
- [8] Branch, W.A. (2004), The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations, *Economic Journal* 114, 592-621.
- [9] Brock, W.A. and Hommes, C.H. (1997), A rational route to randomness, *Econometrica* 65, 1059-1095.
- [10] Casari, M. (2004), Can genetic algorithms explain experimental anomalies?, *Computational Economics* 24, 257–275.
- [11] Chavas, J.-P. (2000), On information and market dynamics: the case of the U.S. beef market, *Journal of Economic Dynamics and Control* 24, 833-853.

- [12] Chen, S.-H. and P. Wang (2002), *Computational Intelligence in Economics and Finance*. Berlin, Springer.
- [13] Dawid, H. (1999), *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*, 2nd ed., Berlin: Springer.
- [14] Duffy, J. (2006), Agent-based models and human-subject experiments, in Leigh Tesfatsion and Kenneth L. Judd (ed.), *Handbook of Computational Economics*, edition 1, volume 2, chapter 19, pages 949–1011.
- [15] Duffy, J. (2008), Experimental Macroeconomics, Entry in: S. Durlauf and L. Blume, (Eds.), *The New Palgrave Dictionary of Economics*, 2nd ed., New York: Palgrave Macmillan.
- [16] Erev, I. and Roth, A.E., (1999), Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria, *American Economic Review* 88, 848-881.
- [17] Evans, G.W. and Honkapohja, S, (2001), *Learning and expectations in macroeconomics*, Princeton University Press, Princeton.
- [18] Evans, G.W. and Ramey, G. (1992), Expectation, calculation and macroeconomics dynamics, *American Economic Review* 82, 207-224.
- [19] Ezekiel, M., (1938) The cobweb theorem, *Quarterly Journal of Economics* 52, 255-280.
- [20] Freeman, R.B. (1975), Legal "cobwebs: a recursive model of the market for new lawyers, *Review of Economics and Statistics* 57, 171-179.
- [21] Freeman, R.B. (1976), A cobweb model of the supply and starting salary of new engineers, *Industrial and Labor Relations Review* 29, 236-248.
- [22] Heemeijer, P., Hommes, C.H., Sonnemans, J. and Tuinstra, J. (2009), Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback: An Experimental Investigation, *Journal of Economic Dynamics and Control* 33, 1052-72.
- [23] Herrera, F. et al. (1998), Tackling real-coded genetic algorithms: oper-



- ators and tools for behavioural analysis, *Artificial Intelligence Review* 12, 265–319.
- [24] Holland, J. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.
- [25] Hommes, C.H., (2002), Modeling the stylized facts in finance through simple nonlinear adaptive systems, *Proceedings of the National Academy of Sciences* 99, 7221-7228.
- [26] Hommes, C.H., 2006. Heterogeneous agent models in economics and finance, In: Tesfatsion, L. and Judd, K.J. (Eds.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, Elsevier, pp.1109-1186.
- [27] Hommes, C. (2009), Bounded rationality and learning in complex markets, In: *Handbook of Research on Complexity*, Edited by J. Barkley Rosser, Jr., Cheltenham: Edward Elgar, pp.87-123.
- [28] Hommes, C. (2011), The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab, *Journal of Economic Dynamics and Control* 35, 1-24.
- [29] Hommes, C.H., Sonnemans, J., Tuinstra, J., and van de Velden, H., (2005), Coordination of expectations in asset pricing experiments, *Review of Financial Studies* 18, 955-980.
- [30] Hommes, C., J. Sonnemans, J. Tuinstra and H. van de Velden (2007) Learning in cobweb experiments, *Macroeconomic Dynamics* 11 (S1), 8-33.
- [31] Kahneman, D. (2003), Maps of bounded rationality: Psychology for behavioral economics, *American Economic Review* 93, 1449-1475.
- [32] Kirman, A. (1993), Ants, rationality and recruitment, *Quarterly Journal of Economics* 108, 137–156.
- [33] Kirman, A. (2006), Heterogeneity in economics, *Journal of Economic Interaction and Coordination* 1, 89-117.
- [34] Krugman, P.R. (2001), The oil-hog cycle, *The New York Times*, November 18, p.14.

- [35] LeBaron, B., B. Arthur, and R. Palmer (1999), The time series properties of an artificial stock market, *Journal of Economic Dynamics and Control*, 23, 1487–1516.
- [36] Lux, T., (1995) Herd Behavior, Bubbles and Crashes, *The Economic Journal* 105, 881–896.
- [37] Lux, T. and Marchesi, M. (1999) Scaling and criticality in a stochastic multi-agent model of a financial market, *Nature* Vol. 397, February 1999, 498–500.
- [38] Lux, T. and Marchesi, M. (2000) Volatility clustering in financial markets: a micro-simulation of interacting agents, *International Journal of Theoretical and Applied Finance* 3, 675-702.
- [39] Lux, T. and S. Schornstein (2005), Genetic learning as an explanation of stylized facts of foreign exchange markets, *Journal of Mathematical Economics* 41, 169–196.
- [40] Mankiw, N.G., Reis, R.A.M.R. and Wolfers, J. (2003), Disagreement about inflation expectations, *NBER Working Paper* No. W9796.
- [41] Marimon, R. and S. Sunder (1994), Expectations and learning under alternative monetary regimes: an experimental approach,” *Economic Theory*, 4, 131-162.
- [42] Muth, J.F., (1961) Rational expectations and the theory of price movements, *Econometrica* 29, 315-335.
- [43] Nerlove, M., (1958), Adaptive expectations and cobweb phenomena, *Quarterly Journal of Economics* 72, 227-40.
- [44] Reis, R. (2006), Inattentive producers, *Review of Economic Studies* 73, 793-821.
- [45] Rosen, S., Murphy, K. and Scheinkman, J., (1994), Cattle cycles, *Journal of Political Economy* 102, 468-92.
- [46] Sargent, T.J., (1993) Bounded rationality in macroeconomics, Clarendon Press, Oxford.
- [47] Shiller, R.J., (2000), Measuring bubble expectations and investor confidence, *Journal of Psychology and Financial Markets* 1, 49-60.

- [48] Sutan, A. and Willinger, M., (2005), Why do we guess better in negative feedback situations? An experiment of beauty contest games with negative feedback and interior equilibria, Working paper University of Montpellier.
- [49] Tversky, A. and Kahneman, D. (1974), Judgment under uncertainty: heuristics and biases, *Science* 185, 1124-1131.
- [50] van de Velden, H. (2001), An Experimental Approach to Expectation Formation in Dynamic Economic Systems, Ph-D Dissertation, Tinbergen Institute Research Series 268.
- [51] Vissing-Jorgensen, A. (2003), Perspective on behavioral finance: does ‘irrationality’ disappear with wealth? Evidence from expectations and actions, In: Gertler, M., Rogoff, K. (Eds.) *NBER Macroeconomics Annual*, MIT Press, Cambridge.
- [52] Vriend, N.(2000), An illustration of the essential difference between individual and social learning and its consequences for computational analyses, *Journal of Economic Dynamics & Control*, 24, 1-14.
- [53] Wellford, C.P. (1989), A laboratory analysis of price dynamics and expectations in the cobweb model, *Discussion Paper* No. 89-15, Department of Economics, University of Arizona.
- [54] Zarkin, G.A. (1985), Occupational choice: an application to the market for public school teachers, *Quarterly Journal of Economics* 100, 409-446.