Evolutionary and Adaptive Learning in Complex Markets: a brief summary

Cars H. Hommes[a]

[a]University of Amsterdam, School of Economics, CeNDEF, Roetersstraat 11, NL-1018 WB, Amsterdam, The Netherlands

ABSTRACT

We briefly review some work on expectations and learning in complex markets, using the familiar demand-supply cobweb model. We discuss and combine two different approaches on learning. According to the adaptive learning approach, agents behave as econometricians using time series observations to form expectations, and update the parameters as more observations become available. This approach has become popular in macro. The second approach has an evolutionary flavor and is sometimes referred to as reinforcement learning. Agents employ different forecasting strategies and evaluate these strategies based upon a fitness measure, e.g. past realized profits. In this framework, boundedly rational agents switch between different, but fixed behavioral rules. This approach has become popular in finance. We combine evolutionary and adaptive learning to model complex markets and discuss whether this theory can match empirical facts and forecasting behavior in laboratory experiments with human subjects.

Keywords: Complex systems, behavioral economics, evolutionary selection, nonlinear dynamics

1. INTRODUCTION

There are opposing views about expectation formation in economics and finance. According to the traditional, neoclassical view, agents form rational expectations (RE) without systematic forecasting errors.[30, 32] In the rational framework one usually assumes that agents have structural knowledge about the economy, and use all available information to compute a rational forecast. Moreover, typically it is assumed that *all* agents are fully rational, leading to the representative rational agent benchmark. In Ref. 17, Friedman gave an early argument supporting the representative rational agent framework: irrational agents would be driven out of the market, since rational agents earn higher profits or utility. Evolutionary selection thus prevents irrational behaviour and the economy may be described *as if* all agents are perfectly rational.

Simon[39] already criticized this view, arguing that deliberation and information gathering costs should be taken into account. Recent work on *bounded rationality* in the 1990s, has challenged the traditional view, emphasizing that the extreme assumptions concerning perfect knowledge of the economy and infinite computing capacities are highly unrealistic[11, 36] and in sharp contrast with observed behavior in laboratory experiments with human subjects.[41] In macroeconomics, the *adaptive learning* approach has become popular.[15] Agents do *not* know the underlying "law of motion" of the economy, but instead use time series observations to form expectations based upon their own "perceived law of motion", trying to learn the model parameters over time. Much of this literature has focussed on the stability of rational expectations equilibria (REE) and equilibrium selection, in an attempt to justify rationality by adaptive learning.

Stimulated by work at the Santa Fe Institute, the view that markets are *complex evolving systems* has gained popularity.[1, 2, 34] If the economy is a *complex system* with many interacting agents, it seems hard to justify perfect structural knowledge about the economy and fully rational expectations, since knowledge about the beliefs of all other agents would be required. A large population of boundedly rational heterogeneous agents, using different forecasting rules ranging from simple to sophisticated, seems much more natural and in line with human behavior. A problem of *bounded rationality* however is that there are *many degrees of freedom*, and which model of bounded rationality is an accurate description of learning behaviour at the individual level?

E-mail: C.H.Hommes@uva.nl

Here we briefly review some work on bounded rationality, expectation formation and learning in complex markets; see Ref. 22 for a much more detailed discussion. We will use the familiar demand-supply cobweb model, exactly the same framework employed in Ref. 32, in Muth's seminal paper introducing rational expectations. We emphasize two stories of bounded rationality: *adaptive learning* and *evolutionary selection*, and combine both stories. An important point of departure for both stories is that agents do not understand the world in its full complexity, but use relatively simple decision heuristics or forecasting strategies. According to the adaptive learning story agents are identical, and can be represented by an "average agent", who adapts his behavior trying to learn an optimal rule within a class of simple rules. An example is the consistent expectations equilibrium,[28] where agents try to learn the best linear rule in an unknown nonlinear economy. The optimal linear rule fits the observable sample mean and sample autocorrelation structure of the nonlinear economy. The second story is concerned with heterogeneous, interacting agents and evolutionary selection of different forecasting rules. Heterogeneous agent models are becoming increasingly popular in finance, where a distinction between fundamentalists and chartist trading strategies can be made; see e.g. Refs. 21, 29 for extensive surveys. Here, we consider the *adaptive belief systems*,[5,6] where agents can choose between a costly sophisticated forecasting strategy, such as rational expectations, and a freely available simple strategy, such as naive expectations. We will integrate both stories and consider an economy with evolutionary selection between a costly sophisticated adaptive learning rule and a cheap simple forecasting rule such as naive expectations.

Much theoretical work on expectations formation, learning and boundedly rational has been done, but surprisingly few laboratory experiments with human subjects on expectations and learning have been conducted done. A controlled laboratory environment is well suited to investigate how individuals form expectations and learn from experience, and how the market aggregates individual forecasting strategies. Recently, in Ref. 26 experiments on expectation formation within a cobweb framework have been performed. We confront theoretical work on expectation formation and learning with the observed "stylized facts" in these laboratory experiments.

Section 2 discusses cobweb dynamics under various expectations rules, such as naive, rational and adaptive expectations. Section 3 focuses on laboratory experiments with human subjects on expectation formation. In Section 4 we discuss adaptive learning, in particular the notion of consistent expectations equilibrium (CEE) and sample autocorrelation (SAC-)learning. Section 5 focuses on heterogeneity and evolutionary competition between different forecasting rules and ends with an example where adaptive learning and evolutionary selection are combined. Finally, Section 6 briefly discusses a future perspective.

2. THE COBWEB MODEL

The classical cobweb model describes commodity price fluctuations of a non-storable good, such as corn or hogs, that takes one time period to produce. It is one of the simplest benchmark models in economic dynamics. Producers form price expectations one period ahead and derive their optimal production decision from expected profit maximization. Given producers' price forecast p_t^e, optimal supply is given by

$$S(p_t^e) = \operatorname{argmax}_{q_t} \{p_t^e q_t - c(q_t)\} = (c')^{-1}(p_t^e). \tag{1}$$

The cost function $c(\cdot)$ is assumed to be strictly convex, so that supply is strictly increasing in expected price. The simplest case arises when the cost function is quadratic, $c(q) = q^2/(2s)$, yielding a linear supply curve

$$S(p^e) = sp^e, \qquad s > 0. \tag{2}$$

In general a strictly convex cost curve leads to a nonlinear, increasing, supply curve. As an example, we will consider an increasing, S-shaped supply curve

$$S(p^e) = b + \arctan(\lambda p^e), \qquad \lambda > 0, b > \pi/2, \tag{3}$$

where the parameter λ tunes the nonlinearity of the supply curve and $b > \pi/2$ is a parameter tuning the production level ensuring that production is always non-negative.

Consumer demand D depends upon the current market price p_t. We will simply work with a linearly decreasing demand curve

$$D(p_t) = a - dp_t + \epsilon_t, \qquad a, d > 0, \tag{4}$$

where $-d$ is the slope of the demand curve, a determines the demand level and ϵ_t is an IID stochastic series representing exogenous random demand shocks. If beliefs are *homogeneous*, i.e., all producers have identical price expectations p_t^e, market clearing implies

$$D(p_t) = S(p_t^e) \tag{5}$$

yielding the realized market price

$$p_t = D^{-1}(S(p_t^e)) = \frac{a + \epsilon_t - S(p_t^e)}{d}. \tag{6}$$

With an increasing supply curve and a decreasing demand curve, there can only be one price, denoted by $p = p^*$, where demand and supply intersect. The price dynamics in (6) thus depends upon the demand and supply curves, as well as on the assumed expectations hypothesis. How do producers form price expectations? We first consider the benchmarks of naive, rational and adaptive expectations.

2.1 Naive expectations

Before the rational expectations revolution it was common practice to use simple forecasting rules, such as naive expectations, where prediction equals the last observed price, $p_t^e = p_{t-1}$. Under naive expectations, the price dynamics (6) becomes

$$p_t = D^{-1}(S(p_{t-1})). \tag{7}$$

According to the well known cobweb theorem,[16] there are essentially two possibilities for the price dynamics under naive expectations: (i) if $-1 < S'(p^*)/D'(p^*) < 0$, then the steady state p^* is (locally) stable, and prices converge; if $S'(p^*)/D'(p^*) < -1$ the steady state p^* is (locally) unstable, and prices diverge. In the case of a nonlinear, bounded supply curve as in (3), if the steady state is unstable, prices will converge to a stable $2-$cycle, with regular up and down oscillations, as illustrated in Figure 2 in the next Section.

2.2 Rational expectations

Simple forecasting rules such as naive expectations, lead to *systematic forecasting errors*. This argument seems particularly strong when the model generates a 2-cycle. When producers expect a high (low) price, they will supply a high (low) quantity and consequently the realized market price will be low (high). Along a 'hog cycle' of up and down price oscillations, expectations are thus *systematically wrong*, and forecasting errors are strongly correlated. Rational agents would learn from their systematic errors and revise expectations accordingly. These considerations lead Muth in Ref.[32] to introduce *rational expectations*, where producers' subjective price expectations equal the objective conditional mathematical expectation of the market price, i.e. $p_t^e = E_t[p_t]$. The rational expectations forecast is given by

$$p_t^e = E_t[p_t] = p^*, \tag{8}$$

where p^* is the unique price corresponding to the intersection point of demand and supply. Given producers' rational price forecast $p_t^e = p^*$, the actual law of motion (6) becomes

$$p_t = p^* + \frac{\epsilon_t}{d}. \tag{9}$$

The cobweb model therefore has a unique REE, given by an IID process with mean p^*. Along a REE expectations are *self fulfilling* and producers make no systematic forecasting errors, since forecasting errors are uncorrelated. In order to form rational expectations however, perfect knowledge of underlying market equilibrium equations is required and, in particular, agents must be able to compute the intersection point p^*.

2.3 Adaptive expectations

Adaptive expectations is given by

$$p_t^e = (1 - w)p_{t-1}^e + wp_{t-1}, \qquad 0 \le w \le 1, \tag{10}$$

where w is the expectations weight factor. The expected price is a weighted average of yesterday's expected and realized prices, or equivalently, the expected price is adapted by a factor w in the direction of the most recent
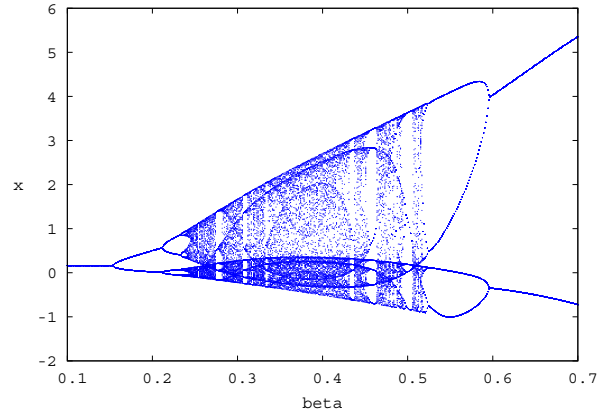
Figure 1. Bifurcation diagram with respect to the expectations weight factor w, $0.1 \leq w \leq 0.7$, with the other parameters fixed at $a = 0.7$, $d = 0.25$ and $\lambda = 4.8$ (x is the deviation from the inflection point of the nonlinear, S-shaped supply curve (3)). For large values of w prices converge to a regular 2-cycle with large amplitude. As w decreases the amplitude of price fluctuations decreases and a bifurcation route to chaos occurs. When w becomes very small, chaotic fluctuations are stabilized and prices converge to REE.

realization. Adaptive expectations may thus be seen as 'error learning' with a constant factor. Notice that for $w = 1$, adaptive expectations reduces to naive expectations. Under adaptive expectations and, given the linear demand curve (4), the dynamics of expected prices in the cobweb model becomes

$$p_t^e = (1 - w)p_{t-1}^e + w(\frac{a + \epsilon_t - S(p_{t-1}^e)}{d}). \tag{11}$$

Without any random shocks ϵ_t, for *nonlinear*, but monotonic, demand and/or supply curves, this nonlinear deterministic difference equation can easily generate chaotic price fluctuations.[9,20] Figure 1 shows a bifurcation diagram with respect to the expectations weight factor w, with the nonlinear, S-shaped supply curve (3). For high values of w, sufficiently close to $w = 1$ (i.e. close to naive expectations) prices converge to a stable 2-cycle, whereas for small values of w, sufficiently close to $w = 0$, prices converge to the RE steady state. For intermediate $w-$values however, chaotic price oscillations arise.

When prices fluctuate chaotically, the corresponding forecasting errors will be highly unpredictable and the question arises whether boundedly rational agents would be able to detect any structure in these chaotic forecasting errors and improve upon their simple adaptive forecasts. If patterns are indeed hard to discover, then adaptive expectations with chaotic price fluctuations might be a satisfactory (long run) boundedly rational equilibrium.

3. LABORATORY EXPERIMENTS

There is a lot of theoretical work on expectations formation and learning when agents are boundedly rational, but surprisingly few laboratory experiments with human subjects have been performed to study how individuals form expectations and learn from experience, and how the market aggregates individual forecasts.

Early experiments have been done in Ref. 37, using historical data on wheat prices and asking subjects to predict the mean wheat price for the next 5 periods. Ref. 33 studies price predictions in repeated double auction experimental asset markets, as in the famous bubble experiments of Ref. 40, and shows that forecasts tend to be biased and inconsistent with RE, but there is a tendency of forecasts to evolve in the direction of RE. In Refs. 31 expectation formation in laboratory experiments in inflationary overlapping generations economies is studied.

Here we discuss some recent laboratory experiments of Ref. 26 on individual expectations and learning in the cobweb framework. See also Ref. 25 for similar experiments in an asset pricing framework. The participants

in the experiments were asked to predict next period's price of a certain, unspecified, good. The realized price p_t in the experiment was determined by the (unknown) cobweb market equilibrium equation

$$D(p_t) = \frac{1}{K} \sum_{i=1}^{K} S(p_{i,t}^e), \qquad (12)$$

where $D(p_t)$ is the demand for the good at price p_t, K is the size of the group, $p_{i,t}^e$ is the price forecast by participant i and $S(p_{i,t}^e)$ is the supply of producer i depending upon the forecast by participant i. Demand and supply curves D and S were fixed during all experiments (except for small random shocks to the demand curve) and unknown to the participants. We focus on the group experiments with $K = 6$, but one-person experiments (i.e. $K = 1$) have also been performed[27] and used to estimate various learning models.[10] Solving (12), with linear demand (4), the market equilibrium price is

$$p_t = \frac{a - \frac{1}{K} \sum_{i=1}^{K} S(p_{i,t}^e)}{d} + \epsilon_t, \qquad (13)$$

where ϵ_t are IID demand shocks, which are drawn from a normal distribution $N(0, 0.5)$. In the experiments the parameters were fixed at $a = 2.3$, $d = 0.25$ and $K = 6$, and we used the nonlinear, S-shaped supply curve (geometrically similar to the S-shaped supply curve (3)):

$$S(p_{i,t}^e) = \text{Tanh}(\lambda(p_{i,t}^e - 6)) + 1, \qquad (14)$$

Expectation formation of the producers is the only part of the model that is affected by the participants in the experiments. Participants did not know underlying market equilibrium equations, nor were they informed about the distribution of any exogenous shocks to demand and/or supply. The participants were told that they were advisors to producers of an unspecified good and that the price was determined by market clearing. Based upon this information the participants were asked to predict next period's price. The predicted price had to be between 0 and 10 and the realized price was also always between 0 and 10. Participants' earnings in each period were a quadratic function of their squared forecasting error. The better their forecast, the higher their earnings. After every period the participants were informed about the realized price in the experiment. Also a time series of the participant's own prediction and a time series of the realized price in the experiment was shown on their computer screen.

Participants in the experiments therefore had little information about the price generating process and had to rely mainly upon time series observations of past prices and predictions. The information in the experiment was thus similar to the information assumption underlying much of the bounded rationality literature, where agents form expectations based upon time series observations. Our setup enables us to test the expectations hypothesis in a controlled dynamic environment. The main question was whether agents can learn and coordinate on the unique REE, in a (relatively simple) world where underlying market equilibrium equations are unknown and agents only observe time series. Our choice for a nonlinear, S-shaped supply curve enables us to investigate whether agents can avoid systematic forecasting errors, as would e.g. occur along a 2-cycle under naive expectations, or can even learn a REE steady state.

In Ref. 26 a stable and an unstable treatment were considered, differing only in the parameter λ tuning the nonlinearity of the supply curve (14). In the stable treatment, if all subjects use naive expectations, prices converge to the RE steady state. In contrast, in the unstable treatment, if all subjects use naive expectations, prices diverge from the RE steady state and converge to the stable 2-cycle, with systematic forecasting errors, as illustrated in Figure 2 (top left panel). Figure 2 also illustrates what would happen in the unstable treatment of the experiment if all subjects would use one of the other well known benchmark expectations rules, namely adaptive expectations ($w = 0.2$), rational expectations (i.e. use the RE price p^* as forecast), learning by average, that is, use the sample average
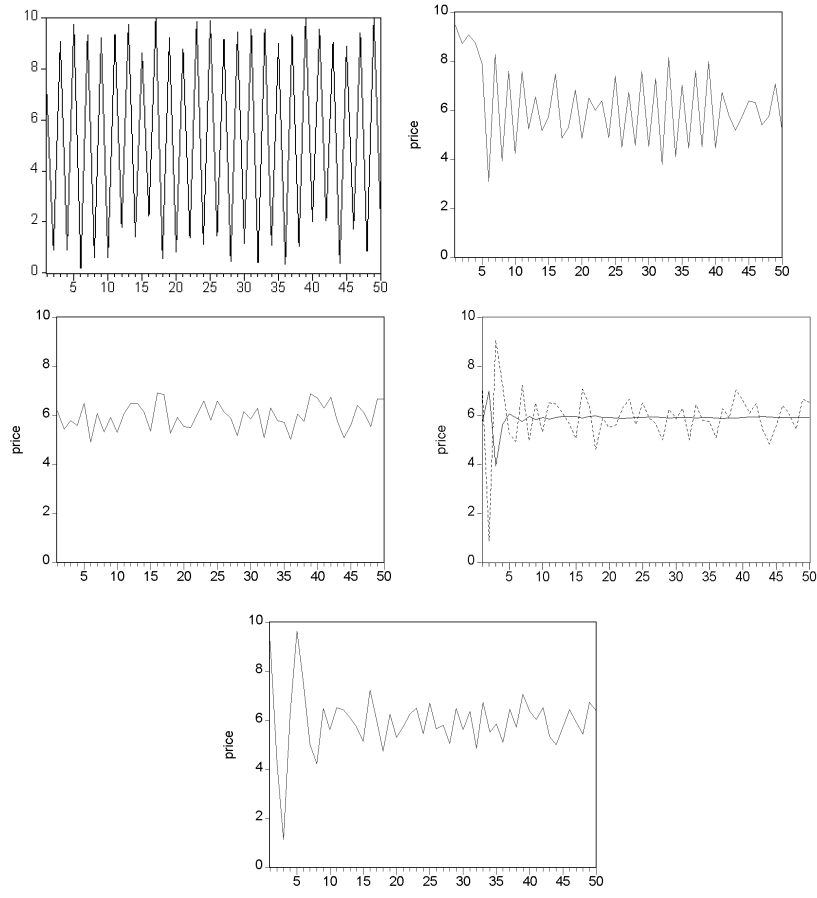
$$\bar{p}_t^e = (\sum_{j=0}^{t-1} p_j)/t, \qquad (15)$$

Figure 2. Price fluctuations in the cobweb model under naive expectations (top left), adaptive expectations (top right), rational expectations (middle left), average price forecast (middle right) and SAC-learning (bottom).
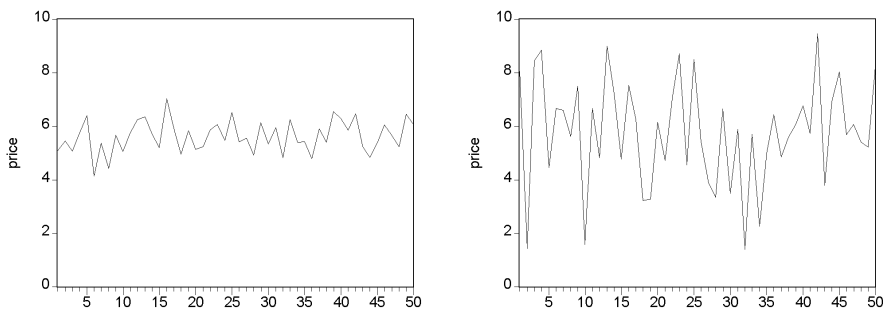

Figure 3. Realized market prices in two different cobweb group experiments. In the stable treatment (left panel; $\lambda = 0.22$) the price quickly converges to the RE price, whereas in the unstable treatment (right panel; $\lambda = 2$)) prices do not converge and exhibit excess volatility, with strongly fluctuating prices around the RE price.

and sample autocorrelation (SAC) learning (i.e. updating sample average and first order sample autocorrelation coefficient, as discussed in detail in Section 4).

Figure 3 shows time series of the realized prices in two typical group experiments, one stable and one unstable treatment. Hommes et al. (2007) summarize the *stylized facts* of realized market prices in the cobweb experiments as follows:

1. the sample mean of realized market prices is very close to the RE price;

2. the sample variance of realized market prices depends on the treatment

   (a) in the stable treatment the sample variance is very close to the RE benchmark;

   (b) in the unstable treatment the sample variance is significantly higher than the RE benchmark;

3. there is no linear autocorrelation left in realized market prices.

One may say that the stable treatment converges to RE*, whereas the unstable treatment exhibits excess volatility, with prices fluctuating irregularly (no autocorrelations) and with high amplitude around the RE benchmark.

It is useful to compare these experimental results to the theoretical benchmarks illustrated in Figure 2. These are *representative agent* benchmarks, where all agents use the same forecasting rule, and demand and supply are exactly the same as in the unstable treatment in the experiment. Naive expectations is clearly very different from the experiments, since it leads to high amplitude price fluctuation with regular, predictable up and down (noisy) period 2 oscillations. Adaptive expectations is also inconsistent with the experiments. Although the amplitude is smaller, the price fluctuations are too regular, with frequent up and down oscillations. In contrast to the experiments, the price series under adaptive expectations, for example, exhibits strong negative first order autocorrelation. The time series under rational expectations is very similar to the time series in the stable treatment (the exogenous shocks in the experiments are the same as for the RE benchmark simulation), but very different from the unstable treatment, which has a much larger amplitude. RE is therefore a good description in the stable treatment, but not in the unstable treatment. Finally, learning by average or by sample autocorrelation always leads to (quick) convergence to RE, which is inconsistent with the observed excess volatility in the unstable treatment of the experiments. None of these representative agent learning models thus can explain the cobweb experiments, suggesting that *heterogeneous expectations* play a key role in expectation formation of boundedly rational agents. Before turning to heterogeneous expectations models in Section 5, we discuss adaptive learning by an "average agent" in Section 4.

4. ADAPTIVE LEARNING

Adaptive learning refers to the situation where agents use some parameterized rule, and update the parameters over time as additional observations become available. Agents thus try to learn the parameters of their rule, for example behaving as a time series econometrician using a recursive ordinary least squares (OLS) updating rule. Ref. 15 contains an extensive and excellent overview of adaptive learning in macroeconomics. Within the cobweb framework adaptive learning has been applied by Bray and Savin (1986).

Adaptive learning may provide a *learning story* how agents may learn a REE, without structural knowledge of market equilibrium equations. In fact, we have seen an example already, since the average price forecast rule (15) can be obtained from OLS regression of prices on a constant. The average price forecast rule enforces convergence to the unique REE in the cobweb model. In cases when there are multiple REE, adaptive learning may be used as an equilibrium selection device, providing a justification of RE equilibria that are stable under learning.

However, adaptive learning need not always converge to REE. In particular, when the perceived law of motion (i.e. the law of motion agents believe in) is *misspecified* (i.e. different from the true law of motion), the learning process need not converge to a REE steady state, but may lead to some *boundedly rational learning equilibrium*, leading to expectations driven periodic or even chaotic fluctuations.[7, 38]

This section discusses a simple adaptive learning scheme, sample autocorrelation (SAC-)learning.[28] Agents are trying to learn the best *linear* rule (according to forecasting performance) in an unknown, *nonlinear* economy. In a consistent expectations equilibrium the linear rule has the same autocorrelation structure as the unknown nonlinear system.

---

*For different market settings, these results may off course change. The cobweb model has negative expectations feedback. Heemeijer et al. (2007) show in fact that the results are quite different in markets with positive feedback, such as demand driven speculative asset markets. Positive feedback may lead to persistent deviations from the fundamental benchmark, with the sample mean of realized prices e.g. much higher than the RE fundamental benchmark.

4.1 Consistent Expectations Equilibrium (CEE)

There are simple chaotic processes whose behavior looks random to an observer, and e.g. exhibits no autocorrelations. More generally, the well known asymmetric tent maps (a class of piecewise linear interval map) can generate exactly the same autocorrelation structure as a stochastic AR(1) process.[8, 35] Boundedly rational agents observing time series generated by such an unknown nonlinear process and using linear statistical techniques, might wrongly believe that the time series is generated by a stochastic AR(1) process. This example motivated the concept of consistent expectations equilibrium,[28] building on the concept of a *self-fulfilling mistake*.[18]

If agents believe that prices are generated by a stochastic AR(1), the predictor for p_t minimizing the mean squared prediction error is

$$p_t^e = \alpha + \beta(p_{t-1} - \alpha), \tag{16}$$

where the parameters α and β, $\beta \in [-1, 1]$, represent the long run average and the first order autocorrelation coefficient. Given the linear predictor (16), the *implied actual law of motion* for the cobweb model becomes

$$p_t = F_{\alpha,\beta}(p_{t-1}) := D^{-1}S(\alpha + \beta(p_{t-1} - \alpha)). \tag{17}$$

The *sample average* of a time series $(p_t)_{t=0}^{\infty}$ is defined as

$$\bar{p} = \lim_{T \to \infty} \frac{1}{T+1} \sum_{t=0}^{T} p_t \tag{18}$$

and the *sample autocorrelation coefficients* are given by

$$\rho_j = \lim_{T \to \infty} \frac{c_{j,T}}{c_{0,T}}, \qquad j \geq 1, \qquad \text{where} \qquad c_{j,T} = \frac{1}{T+1} \sum_{t=0}^{T-j} (p_t - \bar{p})(p_{t+j} - \bar{p}), \qquad j \geq 0. \tag{19}$$

A CEE is now defined as

DEFINITION 4.1. A triple $\{(p_t)_{t=0}^{\infty}; \alpha, \beta\}$, where $(p_t)_{t=0}^{\infty}$ is a sequence of prices and α and β are real numbers, $\beta \in [-1, 1]$, is called a *consistent expectations equilibrium* (CEE) if

1. the sequence $(p_t)_{t=0}^{\infty}$ satisfies the implied actual law of motion (17),

2. the sample average $\bar{p}$ in (18) exists and is equal to α, and

3. the sample autocorrelation coefficients ρ_j, $j \geq 1$, in (19) exist and the following is true:
   a. if $(p_t)_{t=0}^{\infty}$ is a convergent sequence, then $\text{sgn}(\rho_j) = \text{sgn}(\beta^j)$, $j \geq 1$;
   b. if $(p_t)_{t=0}^{\infty}$ is not convergent, then $\rho_j = \beta^j$, $j \geq 1$.

Stated differently, a CEE is a price sequence together with an AR(1) belief such that expectations are self-fulfilling in terms of the *observable* sample average and sample autocorrelations. Along a CEE expectations are thus correct in a linear statistical sense.

4.2 Sample Autocorrelation Learning

Now consider the more flexible situation of *adaptive learning* with agents updating their AR(1) belief parameters α_t and β_t over time, as additional observations become available. A natural learning scheme fitting the framework of CEE is based upon sample average and sample autocorrelation coefficients. For any finite set of past observations $\{p_0, p_1, \ldots, p_t\}$ the sample average is

$$\alpha_t = \frac{1}{t+1} \sum_{i=0}^{t} p_i, \qquad t \geq 1 \tag{20}$$

and the first order sample autocorrelation coefficient is

$$\beta_t = \frac{\sum_{i=0}^{t-1}(p_i - \alpha_t)(p_{i+1} - \alpha_t)}{\sum_{i=0}^{t}(p_i - \alpha_t)^2}, \qquad t \geq 1. \tag{21}$$

When the belief parameters are updated according to (20) and (21) the (temporary) law of motion (17) becomes

$$p_{t+1} = F_{\alpha_t, \beta_t}(p_t) = D^{-1}S(\alpha_t + \beta_t(p_t - \alpha_t)), \qquad t \geq 0. \tag{22}$$

The dynamical system (20)-(22) is called the actual dynamics with *sample autocorrelation learning* (SAC-learning).

Which type of CEE exist in the cobweb model, and to which of them will the SAC-learning dynamics converge? In Ref. 28 it is shown that in the most relevant case, when demand is decreasing and supply is increasing, the *only* CEE is the RE steady state price p^*. This means that, even when underlying market equilibrium equations are *not* known, agents should be able to learn and coordinate on the REE price simply by looking at sample averages and sample autocorrelations. Although other simple forecasting rules, such as adaptive expectations, might lead to chaotic price fluctuations, these forecasting rules are *inconsistent* in terms of sample autocorrelations. Hence, in a nonlinear cobweb economy with monotonic demand and supply, boundedly rational agents should, at least in theory, be able to learn the unique REE from time series observations.

In general however, given an AR(1) belief, there are at least three possible types of CEE:

- a *steady state CEE* in which the price sequence $(p_t)_{t=0}^{\infty}$ converges to a steady state p^*, with $\alpha = p^*$ and $\beta = 0$;

- a *2-cycle CEE* in which the price sequence $(p_t)_{t=0}^{\infty}$ converges to a period two cycle $\{p_1^*, p_2^*\}$, $p_1^* \neq p_2^*$, with $\alpha = (p_1^* + p_2^*)/2$ and $\beta = -1$;

- a *chaotic CEE* in which the price sequence $(p_t)_{t=0}^{\infty}$ is chaotic, with sample average α and autocorrelations β^j.

Which of these cases occurs in the cobweb model depends on the composite mapping $D^{-1}S$ in (17), determined by demand and supply curves. The chaotic case can only arise when demand and/or supply curves are *non-monotonic*; see Ref. 24 for an example. With monotonic demand and supply curve as in the laboratory experiments SAC-learning will always enforce convergence to the REE price.

5. HETEROGENEOUS BELIEFS AND EVOLUTIONARY SELECTION

So far we have focused on a representative agent cobweb model, where all producers have identical expectations. But why would all agents have the same expectations? Laboratory experiments have shown that, even when individuals face the same information, they may disagree and take different decisions. In a complex market it seems more appropriate to model agents as boundedly rational and heterogeneous, using different types of forecasting rules. But this raises an immediate problem: which rules will boundedly rational agents choose from an ocean of infinitely many possible rules?

Models with heterogeneous agents are becoming increasingly popular. In particular, in finance models with fundamentalists and chartists have received much attention; see e.g Refs. 21, 29 for extensive recent reviews of this rapidly expanding literature. In this section we discuss a model with heterogeneous expectations, as proposed by Brock and Hommes in Ref. 5 (henceforth BH) based on three underlying assumptions: (i) agents choose from a class of rules varying from very simple to very sophisticated; (ii) more sophisticated rules require more effort and are therefore more costly than simple rules, and (iii) agents tend to switch to rules that have performed better in the recent past. Evolutionary selection thus disciplines the forecasting rules to be used. In the cobweb framework, producers can choose between different forecasting rules H_j. The fractions $n_{j,t}$ of producers using predictor H_j at date t, will be updated over time based upon a publically available evolutionary fitness measure, given by realized net profits, associated to each predictor.

BH focus on a simple two type case with rational expectations, which can be obtained at costs $C \geq 0$ per time period, versus naive expectations, which is freely available. This case may be viewed as an extreme case, with rational expectations representing the most sophisticated forecasting rule, and naive expectations representing the simplest forecasting rule. BH show the occurrence of a *rational route to randomness*, i.e. a bifurcation route to strange attractors and chaos as traders become more rational in the sense that they become more sensitive to differences in past performance and switch more quickly to a better predictor.

Rational agents have perfect knowledge about market equilibrium equations and are aware of the fact that the market equilibrium price is affected by the presence of naive traders. Hence, in a heterogeneous world rational agents have perfect knowledge about prices and quantities, but also about *beliefs* of *all* other traders. Although this case is theoretically appealing, it seems highly unrealistic in real markets that some agents have (perfect) information about beliefs of other agents. Therefore we will focus here on some perhaps more realistic cases, where agents only use information extracted from observable quantities, such as prices. As a starting point of the discussion, we consider the case of two simple linear predictors.

5.1 Linear forecasting rules

Consider the two linear AR(1) prediction rules

$$H_j(p_{t-1}) = \alpha_j + \beta_j p_{t-1}, \qquad j = 1, 2, \tag{23}$$

with fixed parameters α_j and β_j. The supply curve is *linear* as in (2), with corresponding cost function $c(q) = q^2/(2s)$. The market clearing price in the cobweb model with linear demand and supply and two trader types, with linear predictors as in (23), is determined by[†]

$$a - dp_t = n_{1t}s(\alpha_1 + \beta_1 p_{t-1}) + n_{2t}s(\alpha_2 + \beta_2 p_{t-1}), \tag{24}$$

where n_{1t} and n_{2t} denote the fractions of agents using respectively H_1 and H_2, at the beginning of period t. These fractions will be updated according to an evolutionary fitness measure based on past realized profits. Realized net profit in period t for traders using predictor H_j is given by

$$\pi_{j,t} = sp_t H_j(p_{t-1}) - \frac{s}{2}(H_j(p_{t-1}))^2 - C_j, \tag{25}$$

where C_j represents the average costs per time period for obtaining predictor H_j. For a simple habitual rule of thumb predictor, such as naive or adaptive expectations, these costs C_j will be zero, whereas for more sophisticated predictors such as fundamentalists beliefs based on fundamental analysis, information gathering costs C_j may be positive. The fitness measure underlying evolutionary selection is given by

$$U_{jt} = wU_{j,t-1} + (1-w)\pi_{j,t}, \tag{26}$$

where $0 \leq w \leq 1$ is a memory parameter. A smaller w puts more memory on recent observations and in the case $w = 0$ fitness is given by most recent observed realized net profits.

BH considered this model with *synchronous* updating of strategies, that is, in each period *all* agents update their strategies. Here we consider the more general case of *asynchronous* updating.[13, 23] Per time unit only a fraction $1 - \delta$ of agents, distributed randomly among agents of both types and independently across time, is assumed to reconsider their strategy on the basis of the most recent information available. The remaining fraction δ sticks to their current strategy. The corresponding dynamics of the fractions is given by a modified version of the discrete choice, logit probabilities:

$$n_{jt} = (1-\delta)e^{\gamma U_{j,t-1}}/Z_{t-1} + \delta n_{j,t-1}, \tag{27}$$

where $Z_{t-1} = \sum_h e^{\gamma U_{h,t-1}}$ is a normalization factor so that fractions add up to 1. For $\delta = 0$, we are back in the case of synchronous updating. A key feature of this evolutionary predictor selection is that agents are *boundedly*

[†]In our simulations we will work in deviations $x_t = p_t - p^*$ from the fundamental RE steady state price p^*. This is equivalent to setting the parameter $a = 0$, so that the RE steady state $p^* = a/(d+s) = 0$.

rational, in the sense that predictors with higher evolutionary fitness attract more followers. The parameter γ is called the *intensity of choice*, measuring how fast producers switch between different prediction strategies. For $\gamma = 0$ the fractions always converge to equal shares $1/H$, whereas for the other extreme $\gamma = \infty$, in each period *all* producers who update in that period (i.e., a fraction $1 - \delta$) switch to the optimal predictor. Hence, the higher the intensity of choice, the more rational agents are in the sense that they switch more quickly to the best strategy in terms of past performance. According to (24) and (27) market equilibrium prices and fractions of different trading strategies co-evolve over time.

5.2 Fundamentalists versus naive expectations

The linear predictors (23) specialize to the case with fundamentalists versus naive expectations when $\alpha_1 = p^* = a/(d + s)$ (the steady state price), $\beta_1 = 0$, $\alpha_2 = 0$ and $\beta_2 = 1$:

$$H_1(p_{t-1}) = p^* = \frac{a}{d + s} \tag{28}$$

$$H_2(p_{t-1}) = p_{t-1}. \tag{29}$$

Figure 4 shows attractors for different values of the intensity of choice γ and some corresponding time series. A *rational route to randomness*, i.e. a bifurcation route from simple to complicated, chaotic dynamics and strange attractors occurs as the intensity of choice increases. The market switches between periods of low volatility, with prices close to the fundamental price, and high volatility, with irregularly switching prices. Prices diverge slowly from the fundamental steady state price, as long as most agents use the simple, freely available naive forecast. When forecasting errors increase, it becomes worthwhile to buy the sophisticated fundamental forecast, and more agents start switching to the fundamental forecast, thus stabilizing price fluctuations, etc. Due to the *asynchronous* updating of strategies, agents switch more gradually between strategies, and the time series of fractions of fundamentalists shows much more persistence than in the case with synchronous updating. Figure 4 (bottom panel) also illustrates the sample average and first order sample autocorrelation (SAC) of the price series. Sample average quickly settles down to a value close to $0^\ddagger$, whereas the first order SAC is clearly negative, converging to approx. -0.85.

5.3 Adaptive learning versus naive expectations

In the case of fundamentalists versus naive, price series exhibit strong first order negative autocorrelations, even when the dynamics is chaotic. This has been illustrated in Figure 4 showing that, for $\gamma = 3$, the sample autocorrelations of prices converges to a negative value around -0.85. An agent who behaves as a time series econometrician would easily detect this strong negative autocorrelation and adapt her forecasts. Even without the use of any statistical software, a smart agent might detect negative autocorrelation, simply by observing that positive (negative) deviations from the average price are always followed by negative (positive) deviations. What would happen if agents recognize this structure from observing realized market prices?

The next example combines evolutionary strategy selection and adaptive learning. Consider a group of agents using *SAC-learning*, to exploit the negative first order autocorrelation in observed prices. That is, replace the fundamental forecast by a SAC-learning forecasting rule

$$H_1(p_{t-1}) = \alpha_{t-1} + \beta_{t-1}(p_{t-1} - \alpha_{t-1}), \tag{30}$$

where α_t and β_t are determined through SAC-learning as in (20) and (21) respectively. This approach widens the range of forecasting rules to all linear AR(1) rules. The sophisticated agent type tries to learn the optimal linear rule through adaptive learning, within a heterogeneous agent environment. Recently, in different contexts similar heterogeneous agent models with adaptive learning have been introduced.[3,12,14]

Figure 5 illustrates the dynamics in the case of SAC-learning versus naive expectations. Agents *learn to be contrarians*, as $\beta_t \to -0.62$, consistent with the SAC in realized prices. In this example there is still fairly strong

---

‡Recall that the simulations are in deviations $x_t = p_t - p^*$ from the fundamental, so that the sample average of prices converges to fundamental value.
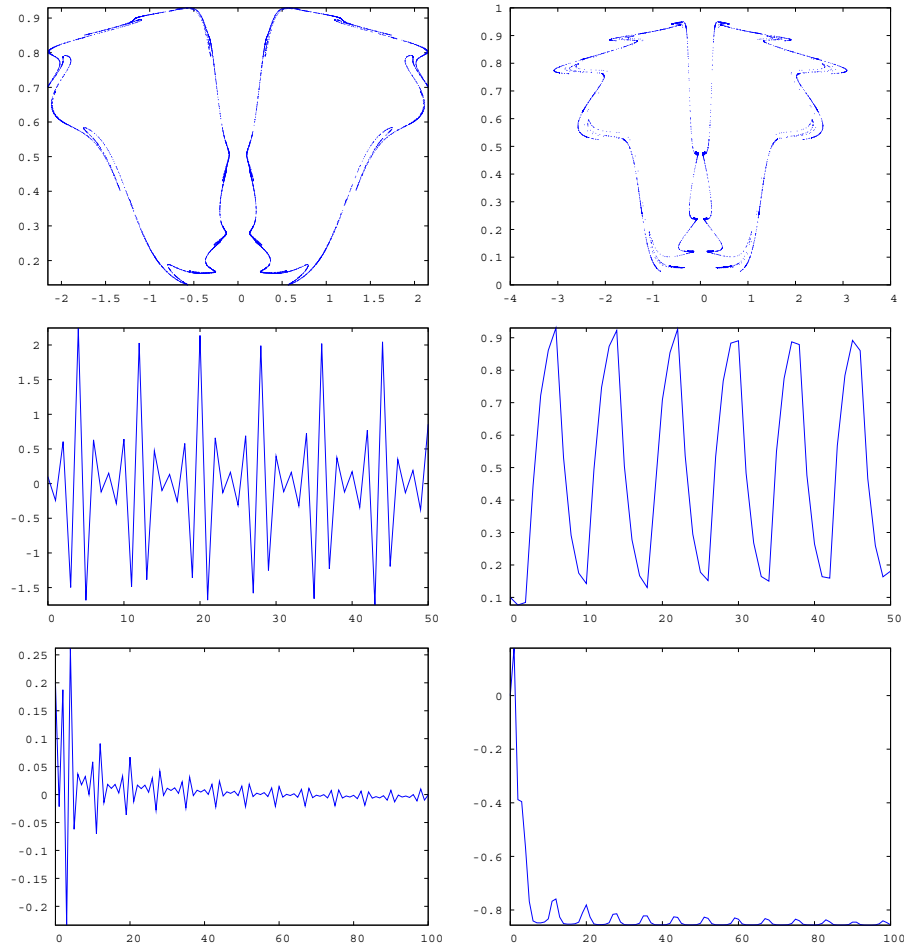
Figure 4. Fundamentalists versus naive. Strange attractors and time series for different γ-values, with other parameters fixed at $a = 0$, $d = 0.5$, $s = 1.35$, $\delta = 0.5$, $\alpha_1 = 0$, $\beta_1 = 0$, $C_1 = 1$, $\alpha_2 = 0$, $\beta_2 = 1$ and $C_2 = 0$. Although price dynamics is chaotic, there is still clear linear autocorrelation structure. Sample average of prices converges (close) to fundamental value, while sample autocorrelations converge (close) to -0.85, indicating significantly negative first order autocorrelation.
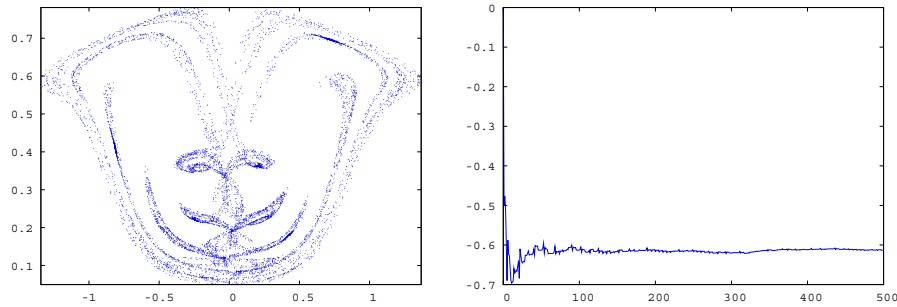

Figure 5. SAC-learning versus naive. Agents learn to be contrarians, as the first order autocorrelation coefficient converges, $\beta_t \rightarrow -0.62$. Parameters: $\gamma = 3$, $a = 0$, $d = 0.5$, $s = 1.35$, $\delta = 0.5$, $w = 0$, $C_1 = 1$, $\alpha_2 = 0$, $\beta_2 = 1$ and $C_2 = 0$.
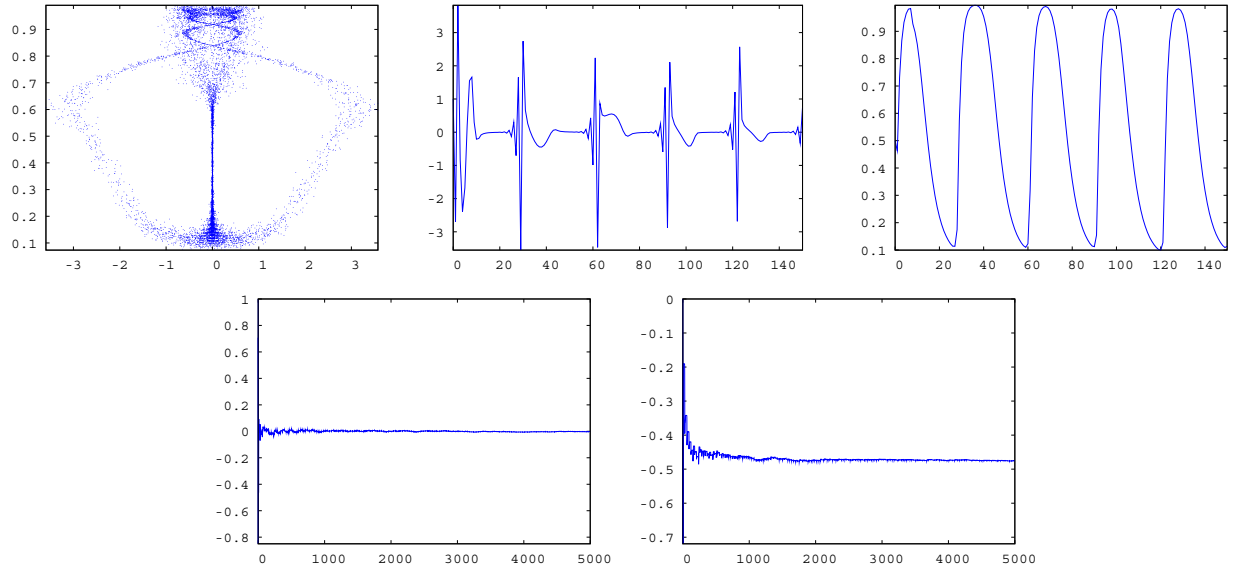
Figure 6. SAC-learning versus naive expectations with memory in the fitness measure. Attractor (a) and time series of prices p_t (deviations from fundamental), fraction n_{1t} of SAC-learners, sample average α_t and sample autocorrelation coefficient β_t. With more memory in the fitness meausure, the remaining autocorrelation in prices is weaker ($\beta_t \to -0.48$). Parameters: $\gamma = 3$, $a = 0$, $d = 0.5$, $b = 1.35$, $\delta = 0.5$, $w = 0.9$, $C_1 = 1$, $\alpha_2 = 0$, $\beta_2 = 1$ and $C_2 = 0$.

negative first order autocorrelation in prices, although it is less than in the case of fundamentalists versus naive, and it is consistent with the behavior of the sophisticated type, who have learned the first order autocorrelation coefficient consistent with realized market prices. Figure 6 illustrates another example with memory in the fitness measure, where the (first order) autocorrelation in prices becomes even weaker ($\beta_t \to -0.48$).

6. CONCLUDING REMARKS

We have summarized bounded rationality and learning in the familiar cobweb, hog-cycle framework. Two stories of bounded rationality have been emphasized. The story of adaptive learning assumes a representative "average" agent trying to optimize a simple, (linear) misspecified rule in an unknown complex (nonlinear) economy. The other story assumes heterogeneous forecasting strategies and endogenous, evolutionary selection based upon past performance. We have also presented an example where both stories are integrated, with evolutionary selection between an adaptive learning rule and a simple, fixed rule.

In a cobweb economy with nonlinear, but monotonic demand and supply curves, many adaptive learning processes enforce convergence to the unique REE steady state price. For example, the steady state price forecast is the only (linear) forecast, where sample averages and sample autocorrelations of realized market prices are consistent with beliefs. Simply by looking at sample averages and sample autocorrelations, in particular trying to learn the negative first order autocorrelation so typical for the 'hog cycle', boundedly rational agents should be able to learn the unique REE.

Laboratory experiments with human subjects show however that this is not as easy as theory suggests. Only in the stable treatment of the experiment (i.e. when the market is stable under naive expectations) do market prices converge to REE. In the unstable treatment of the experiments, realized market prices are characterized by three stylized facts: (i) the sample mean is close to the RE price; (ii) there is excess volatility, i.e., the sample variance is much higher than the RE variance, and (iii) there is no linear predictability (no autocorrelations) in realized prices. The observed excess volatility is inconsistent with convergence of adaptive learning of a representative agent. For other simple expectations rules, such as adaptive expectations, irregular price fluctuations around the RE benchmark arise, but these fluctuations, even when chaotic, typically still exhibit negative first order

autocorrelations, inconsistent with stylized fact (iii) in the experiments. Some form of heterogeneity is therefore needed to explain the laboratory experiments.

We have also reviewed some results on heterogeneous agent models with endogenous, evolutionary strategy selection, including several two-type cases with a costly sophisticated forecasting rule (fundamentalists or SAC-learning) versus a free, simple forecasting rule (naive expectations). These two type models will converge to the RE price in the stable treatment of the experiment, and at the same time may generate instability and excess volatility in the unstable treatment, when agents switch fast enough between strategies, similar to the stylized facts in the experiments. However, it is not clear whether a two type model can simultaneously explain stylized fact (iii), i.e. linear unpredictability. A two type model with fundamentalists versus naive expectations generates strongly negative first order autocorrelation in prices, even when the system is chaotic. The typical up and down 'hog cycle' oscillations are still present, and would be observable to a careful agent. When fundamentalists are replaced by SAC-learning, who try to learn and exploit the negative first order autocorrelation in prices, the first order autocorrelation gets weaker, but does not disappear completely. In the cobweb framework, adaptive agents *learn to become contrarians* and "arbitrage away" part of the linear predictability, but do not completely wash out the autocorrelations in market prices. These results suggest that, in order to match all stylized facts of the experiments, either the simple strategy (naive expectations) in these 2-type models needs to be replaced by a somewhat more complicated strategy (perhaps adaptive expectations or a 2-period average forecast), or more heterogeneity, i.e. more types of forecasting rules, are needed to fully explain the laboratory experiments. Matching the stylized facts of laboratory experiments on expectations formation remains an important challenge for theories of bounded rationality and learning, in the simple cobweb framework as well as for other, more realistic expectations feedback settings.

ACKNOWLEDGMENTS

REFERENCES

1. Anderson, P.W, Arrow, K.J. and Pines, D. (eds.), (1988), *The Economy as an Evolving Complex System II* Addison-Wesley, Reading, MA.
2. Arthur, W.B., Durlauf, S.N and Lane, D.A. (eds.), (1997), *The economy as an evolving complex system II* Addison-Wesley, Reading, MA.
3. Branch, W.A. and Evans, G.W. (2006), Intrinsic heterogeneity in expectation formation, *Journal of Economic Theory* 127, 264-295.
4. Bray, M.M., and Savin, N.E., (1986) Rational expectations equilibria, learning and model specification, *Econometrica* 54, 1129-1160.
5. Brock, W.A., and Hommes, C.H., (1997) A rational route to randomness, *Econometrica*, 65, 1059-1095.
6. Brock, W.A., and Hommes, C.H., (1998), Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* 22, 1235-1274.
7. Bullard, J., (1994) Learning equilibria, *Journal of Economic Theory* 64, 468-485.
8. Bunow, Barry, and Weiss, George H. How Chaotic is Chaos? Chaotic and Other 'Noisy' Dynamics in the Frequency Domain. *Mathematical Biosciences* 47 (1979): 221-237.
9. Chiarella, C., (1988) The Cobweb Model: Its Instability and the Onset of Chaos, *Economic Modeling*, 5, 377-384.
10. Colucci, D. and Valori, V. (1986), Ways of learning in a simple economic setting: a comparison, *Chaos, Solitons and Fractals* 29, 653-670.
11. Conlisk, J. (1996), Why Bounded Rationality?, *Journal of Economic Literature* 34, 669 - 700.
12. DeGrauwe, P. and Markiewicz, A. (2006), Learning to Forecast the Exchange Rate: Two Competing Approaches, *CESifo Working Paper* no. 1717, May 2006
13. Diks, C.G.H. and Weide, R. van der (2005), Herding, a-synchronous updating and heterogeneity in memory in a CBS, *Journal of Economic Dynamics and Control* 29, 741-763.
14. Diks, C.G.H. and Dindo, P.D.E., (2006), Informational differences and learning in an asset market with boundedly rational agents, CeNDEF Working paper 06-11 University of Amsterdam.
15. Evans, G.W. and Honkapohja, S, (2001), Learning and expectations in macroeconomics, Princeton University Press, Princeton.
16. Ezekiel, M., (1938) The cobweb theorem, *Quarterly Journal of Economics* 52, 255-280.
17. Friedman, M., (1953) The case of flexible exchange rates, In: Essays in positive economics, Univ. Chicago Press.

18. Grandmont, J.-M., (1998) Expectations Formation and Stability in Large Socio-Economic Systems, *Econometrica*, 66, 741-781.
19. Heemeijer, P., Hommes, C.H., Sonnemans, J. and Tuinstra, J. (2007), Price Stability and Volatility in Markets with Positive and Negative Expectations Feedback: An Experimental Investigation, *CeNDEF Working Paper* 07-01, University of Amsterdam.
20. Hommes, C.H., (1994) Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand, *Journal of Economic Behavior & Organization* 24, 315-335.
21. Hommes, C.H. (2006), Heterogeneous Agent Models In Economics and Finance, In: *Hand-book of Computational Economics, Volume 2: Agent-Based Computational Economics*, Edited by L. Tesfatsion and K.L. Judd, Elsevier Science, 1109-1186.
22. Hommes, C.H. (2007), Bounded Rationality and Learning in Complex Markets, In: J. Barkley Rosser, Jr. (Ed.), Cheltenham: Edward Elgar, to appear.
23. Hommes, C.H., Huang, H. and Wang, D. (2005), A robust rational route to randomness in a simple asset pricing model, *Journal of Economic Dynamics and Control* 29, 1043-1072.
24. Hommes, C.H., and Rosser, J.B., (2001) Consistent Expectations Equilibria and Complex Dynamics in Renewable Resource Markets, *Macroeconomic Dynamics*.
25. 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.
26. Hommes, C.H., Sonnemans, J., Tuinstra, J. and van de Velden, H., (2007), Learning in cobweb experiments, *Macro-economic Dynamics*, in press.
27. Hommes, C.H., Sonnemans, J., and van de Velden, H., (2000) Expectation formation in a cobweb economy: some one person experiments, In: Delli Gatti, M. et al. (eds.), Interaction and Market Structure. Essays on Heterogeneity in Economics, *Lect. Notes in Economics and Math. Systems* 484, Springer Verlag, Berlin/Heidelberg, pp.253–266.
28. Hommes, C. H., and Sorger, G., (1998) Consistent Expectations Equilibria, *Macroeconomic Dynamics* 2, 287-321.
29. LeBaron, B., (2006), Agent-based Computational Finance, In: *Hand-book of Computational Economics, Volume 2: Agent-Based Computational Economics*, Edited by L. Tesfatsion and K.L. Judd, Elsevier Science.
30. Lucas, R.E., (1972) Econometric testing of the natural rate hypothesis, In: O. Eckstein (ed.) The econometrics of price determination Conference. Board of Governors of the Federal Reserve System and Social Science Research Council.
31. Marimon, R. and S. Sunder, (1993), Indeterminacy of equilibria in a hyperinflationary world: experimental evidence, *Econometrica* 61, 1073-1107.
32. Muth, J.F., (1961) Rational expectations and the theory of price movements, Econometrica 29, 315-335.
33. Peterson, S.P. (1993) Forecasting dynamics and convergence to market fundamentals. Evidence from experimental asset markets, *Journal of Economic Behaviour and Organization* 269-284.
34. Rosser, J.B., (2004), *Complexity in Economics: The International Library of Critical Writings in Economics* 174, (3 volumes), Aldergate: Edward Elgar, 2004.
35. Sakai, H., and Tokumaru, H., (1980), Autocorrelations of a certain chaos, *IEEE Transactions on Acoustics, Speech and Signal Processing* 28, 588-90.
36. Sargent, T.J., (1993) Bounded rationality in macroeconomics, Clarendon Press, Oxford.
37. Schmalensee, R., 1976, An Experimental Study of Expectation Formation, *Econometrica* 44, 17–41.
38. Schönhofer, M. (1999), Chaotic Learning Equilibria, *Journal of Economic Theory* 89, 1-20.
39. Simon, H.A. (1957), Models of man, Wiley, New York, NY.
40. Smith, V., Suchanek, G.L. and Williams, A.W., (1988) Bubbles, crashes and endogenous expectations in experimental spot asset markets, Econometrica 56, 1119-51.
41. Tversky, A. and Kahneman, D. (1974), Judgment under uncertainty: heuristics and biases, *Science* 185, 1124-1131.