

Ways of learning in a simple economic setting: a comparison.*

D. Colucci and V. Valori[†]

January 2005

Abstract

We analyse the capacity of a range of different learning rules to describe actual human behaviour in the experiment on expectation formation in a cobweb model conducted by Hommes et al. (2000). We find indication of a relative superiority in terms of descriptive capacity of forms of generalised adaptive expectations allowing for endogenous gain parameters.

1 Introduction

The economic literature on bounded rationality and learning has developed into a very rich field in the last decades and reference to behavioural and psychological issues of underlying human subjects acting in economic models is widespread nowadays as many aspects of the rational expectations/optimising behaviour paradigm are questioned from many standpoints. Quoting Rabin [21] (which contains a neat survey on psychological plausibility in economic modelling):

While still controversial, behavioral economics is on the verge of “going mainstream”, especially in top departments in the U.S. The number of recent hirings, tenurings, conferences, etc., based on behavioral-economic research reflects its growing acceptance.

*We are very grateful to Cars Hommes, Joep Sonnemans and Henk van de Velden for providing us with the experimental data used in this paper. We are also grateful to Jan Tuinstra. The usual caveats apply.

[†]DiMaD - Dipartimento di Matematica per le Decisioni - Università degli Studi di Firenze - Via C. Lombroso, 6/17 - 50134 Firenze - Italy. E-mail: domenico.colucci@dmd.unifi.it, vincenzo.valori@dmd.unifi.it

Within the wide current debate the mechanism of expectations formation retains an important and delicate role: first, because it is often a crucial choice in dynamic models and second, because there is no general agreement as to which assumptions ought to be imposed on model agents. Various papers have shown that, within specific contexts, rational expectations can be rejected as a good description of real agents behaviour: for example Schmalensee [22], Smith *et al.* [23] and Williams [28] conduct experiments with human subjects, whereas Figlewski and Wachtel [12] and Lovell [17] are econometric studies of survey data on expectations. On the other hand the bounded rationality literature has mostly focused on asymptotic properties of particular classes of learning rules, e.g. inquiring the learnability of various equilibria under indeterminacy as in Honkapohja and Mitra [16], investigating the stability properties of various attractors and the possibility of having roads to chaos as in Hommes [13] or checking certain long-run rationality requirements as in Marcet and Nicolini [19]. The perspective we take in this paper is partly different, in that we analyse the capacity of a range of different learning rules to describe actual human behaviour in the experiment on expectation formation in a cobweb model conducted by Hommes *et al.* [15]. This setting is especially interesting because it is a good test of how people deal with forecasting the future in situations of scarce knowledge of how the economic environment works but there is a strong feedback of their predictions and decisions on the observed state. The spirit of what we do in this paper is rather similar to the study carried out by J. D. Serman in the late 1980s whose results were published in a number of papers (see Serman [24, 25, 26, 27]): in particular, both in that study and in this paper the capacity of agents of detecting the expectations feedback plays a central role. We compare various learning algorithms, which are presented in Section 2, both with recursive and non-recursive formulation, and we try to understand which can best describe the subjects' behaviour. Given that the experiment by Hommes *et al.* [15] is a single person treatment repeated for 77 different subjects, we ranked various learning rules in terms of their capacity of describing each agent's actual behaviour. In particular we assigned each agent a forecast rule which best (within a given set of alternatives) resembles its behaviour in the experiment. This is described thoroughly in Section 3. Then we ran simulations to project on a longer run the behaviour of these rules and to evaluate the possibility of observing convergence to the unique rational expectation equilibrium asymptotically: this is the object of Section 4. Section 5 contains some concluding remarks. Section 6 gathers figures and tables.

2 Expectation formation: some alternative

We now briefly present the various expectations mechanisms/learning rules which will be used in the next Sections. We can distinguish them into four groups: *i*) linear predictors with ordinary least squares (or similar) estimation of parameters, *ii*) classic adaptive rules and *iii*) their generalisations, *iv*) predictors with a random component.

Least squares learning in the context of a linear perceived law of motion $y_t = a_t + b_t x_t$ is well known in the literature (see for instance Marcet and Sargent [18]). Ordinary least squares estimation of parameters a_t and b_t , setting $\phi_t = (a_t \ b_t)'$, $Y_t = (x_1, \dots, x_t)'$ and $Z_t = \begin{pmatrix} 1 & \dots & 1 \\ x_1 & \dots & x_t \end{pmatrix}'$, is given by

$$\phi_t = (Z_t' Z_t)^{-1} Z_t' Y_t \quad (1)$$

Also, (setting $z_t = (1 \ x_t)'$) the following recursive formulation can be given (see Evans and Honkapohja [10]):

$$\begin{cases} \phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} (y_t - \phi_{t-1}' z_{t-1}) \\ R_t = R_{t-1} + t^{-1} (z_{t-1} z_{t-1}' - R_{t-1}) \end{cases} \quad (2)$$

A simpler alternative can be written by substituting the matrix $t^{-1} R^{-1}$ in the above recursion with a deterministic gain sequence γ_t satisfying $\sum_{t=1}^{\infty} \gamma_t = \infty$ and $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$. This is known as **stochastic gradient learning** (see Evans and Honkapohja [9]) and assumes the recursive form:

$$\phi_t = \phi_{t-1} + \gamma_t z_{t-1} (y_t - \phi_{t-1}' z_{t-1}) \quad (3)$$

Among the classic adaptive rules consider first the usual version of **adaptive expectations**

$$x_{t+1}^e = x_t^e + \alpha (x_t - x_t^e) \quad (4)$$

where forecasts for the next period are recursively obtained through a convex combination (with a constant gain parameter) of the last observation and forecast. An extensive¹ form of adaptive expectations can be written for economies with infinite past (which is not our case), as a weighted mean of all available data. The counterpart of adaptive expectations in economies with finite past is called **fading memory learning** which has been studied for instance in Bischi-Gardini [2], Bischi-Naimzada [3], Barucci [1] and Chiarella *et al.* [7]. Expectations for the future are a weighted average of available

¹In this paper with the term “*extensive*” we refer to the non-recursive form of a given predictor.

data, with geometrically decreasing weights. The general, extensive, form of a model with fading memory can thus be written as

$$\begin{cases} x_{t+1}^e = \frac{1}{W_t} \sum_{k=0}^t \rho^{t-k} x_k \\ W_t = \sum_{k=0}^t \rho^k, \quad 0 < \rho \leq 1 \end{cases} . \quad (5)$$

With the substitution $\alpha_t = \frac{1}{W_t}$, we get the recursive form

$$\begin{cases} x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e) \\ \alpha_{t+1} = \frac{\alpha_t}{\alpha_t + \rho} \end{cases} . \quad (6)$$

There are two extreme cases: when $\rho = 1$ all the past observations receive the same weight, so we are left with

$$x_t^e = \frac{1}{t} \sum_{k=0}^{t-1} x_k \quad (7)$$

that is, with a **simple average**² (see Bray [5]). The equivalent recursive formulation is

$$\begin{cases} x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e) \\ \alpha_{t+1} = \frac{\alpha_t}{\alpha_t + 1} \end{cases} . \quad (8)$$

Conversely when $\rho \rightarrow 0$ expectations reduce to the **myopic** case

$$x_t^e = x_{t-1} . \quad (9)$$

As suggested by several papers (e.g. the cited Schmalensee [22] and Figlewski and Wachtel [12]) adaptive rules often provide a good description of agents' expectations, particularly so when allowing for a parameter that is non-constant across agents and through time following the rate of uncertainty perceived by the agents. Fading memory and simple average can be considered first steps towards a general adaptive scheme of this type. Both cases are clearly very basic forms of recursive expectations updating. In particular, in (6), the internal variable α_t that determines the correction on the previous expectation in the direction of the last error, is itself only dependent on its own path and on a parameter. The way α is updated is therefore completely independent from external signals. Indeed, α_t converges to $1 - \rho$ regardless of the dynamics of the state variable x .

²Sometimes this way of producing expectations is labelled "least squares" in the sense that it is obviously the OLS estimation of the mean of x . We shall stick to the name "simple average" for this case and reserve the name of "least squares" for the genuine case in which at least two parameters are estimated.

There is a sort of natural way to generalise the fading memory rule to include past prediction performances in the determinants of the dynamics of α_t and that is by endogenising the parameter ρ : this can be called **generalised fading memory** (GFM) (as in Colucci and Valori [8]) and its recursive form is as follows

$$\begin{cases} x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e) \\ \alpha_t = \frac{\alpha_{t-1}}{\alpha_{t-1} + \rho_t} \\ \rho_t = H(x_t, x_t^e) \end{cases} \quad (10)$$

where the function H defines how agents react to forecast errors. The GFM also admits the following extensive form which highlights the fact that it can be seen as a weighted average of available observations:

$$\begin{cases} x_{t+1}^e = \frac{1}{W_t} (x_t + \rho_t x_{t-1} + \rho_t \rho_{t-1} x_{t-2} + \dots + \rho_t \rho_{t-1} \dots \rho_1 x_0) \\ W_t = 1 + \rho_t + \rho_t \rho_{t-2} + \dots + \rho_t \rho_{t-1} \dots \rho_1, \quad 0 < \rho_i \leq 1 \end{cases} \quad (11)$$

Obviously the basic fading memory case corresponds to the degenerate choice of a constant H between 0 and 1. Therefore, for appropriate H functions, the GFM expectations lie, at each step, between myopic expectations and simple average: in fact if we assume, for example, $H(0) = 1$ and $\lim_{y \rightarrow \infty} H(y) = 0$, the GFM is very close to a simple average when the forecast error is low, whereas with large errors we have α close to 1 (as with myopic expectations). This feature is rather similar to the spirit of the learning rule used in Marcet and Nicolini [19], which is as follows:

$$\begin{cases} x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e) \\ \alpha_{t+1} = \begin{cases} \frac{\alpha_t}{\alpha_t + 1} & \text{if } \left| \frac{x_t - x_t^e}{x_t^e} \right| < \nu \\ \bar{\alpha} & \text{otherwise} \end{cases} \end{cases} \quad (12)$$

This rule endogenously switches between simple average and adaptive expectations on the basis of the last forecast error.

Along the lines of GFM we shall also consider the following **generalised adaptive expectations**:

$$\begin{cases} x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e) \\ \alpha_t = H(x_t, x_t^e) \end{cases} \quad (13)$$

in which at each time period the gain parameter is tuned on the basis of the last forecast error.

Finally, we consider two predictors with a random component. The reason to do that is to have a neutral, minimal benchmark for the comparison we conduct in what follows.

In short, in the next section we will consider the following set of alternative predictors.

OLS₁: With a perceived law of motion equal to $x_{t+1} = a_t + b_t x_t$ we obtain a linear predictor $x_{t+1}^e = a_t + b_t x_t$ where the parameters a_t and b_t are estimated with recursive ordinary least squares of equation (2).

OLSE₁: OLS₁ in the extensive form of equation (1).

OLS₂: Now the perceived law of motion is $x_t = a_t + b_t x_t^e$, so the predictor is given by $x_t^e = \frac{a_t}{1-b}$ where a_t and b_t are estimated with recursive ordinary least squares.³

SG: In this case the parameters of the linear perceived law of motion $x_{t+1} = a_t + b_t x_t$, are estimated with stochastic gradient, a simplified version of recursive ordinary least squares, along the line of equation (3). The sequence γ_t is taken to be equal to $\frac{\alpha}{\beta+t}$.

AE: Adaptive expectations with constant gain, as in equation (4).

FM: Fading memory learning as in equation (6).

FME: Fading memory in extensive form as in equation (5).

SA: Simple average as in equation (8).

SAE: Simple average in extensive form as in equation (7).

M: Myopic expectations as in equation (9).

GFM: Generalised fading memory as in equation (10). The function H is taken to be bell-shaped and symmetric: in particular we work with the functional form

$$H(x_t, x_t^e) = ke^{-\left[h\left(\frac{x_t - x_t^e}{x_t^e}\right)\right]^2} + d \quad (14)$$

with $h, d, k \geq 0$ and $k + d \leq 1$. A possible behavioural interpretation for this specification is as follows. Expectations are a weighted mean of the available data; as a consequence a large prediction error could be interpreted as the failure of the data to capture the present trend in the state variable, e.g. as in case of structural breaks or significant exogenous shocks. Therefore the most recent observation assumes a dominant role whereas a low weight is attributed to the bulk of older observations (by the choice of a small ρ_t). The opposite happens in the case of small prediction error. In this scheme, the variable ρ_t can be interpreted as a voluntary choice of how much it is worth recalling,

³In this case there is no need to distinguish an extensive form because it would be exactly the same as this.

the choice being the result of a form of “rational” assessment about the significance of the available information in terms of predicting the future. For more details about this specification see Colucci and Valori [8]. Experimental support for this kind of behaviour is provided by Marimon and Sunder [20].

GFME: Generalised fading memory in extensive form as in equation (11).

GAE: Generalised adaptive expectations as in equation (13). Again, we consider a bell-shaped and symmetric H function. We have chosen the following functional form:

$$H(x_t, x_t^e) = 1 - \left(ke^{-\left[h \left(\frac{x_t - x_t^e}{x_t^e} \right) \right]^2} + d \right) \quad (15)$$

where the parameters satisfy the same restrictions of those in equation (14). GAE is similar to GFM as, in both cases larger errors imply larger values for α ; it is simpler though in that forecast errors influence the gain parameter directly.

MN: Expectations *à la* “Marcet-Nicolini” (see [19]) as in equation (12). As for GFM and GAE the gain parameter is positively correlated with the absolute error (in percentage).

RE: Random expectations: x_t^e is drawn from the uniform distribution on a given support.

RAE: Adaptive expectations with a random gain $x_{t+1}^e = x_t^e + \alpha_t(x_t - x_t^e)$ where $\alpha_t \sim U(0, 1)$.

3 The experiment

Hommes *et al.* [15] have conducted one person experiments in which agents’ predictions are used to generate the price dynamics of a nonlinear cobweb model. We now introduce the model that underlies the experiment. Consider a single market for a perishable good that requires a time period to be produced. Demand for the good depends on its current price p_t . Due to the production lag, decisions on the supply side depend on the price expected by producers. Assuming a linearly decreasing demand

$$D(p_t) = a - bp_t \quad \text{with } a, b > 0$$

and an S-shaped supply curve, with a unique inflection point

$$S(p_t^e) = \tanh(\lambda(p_t^e - c)) + 1 \quad \text{with } \lambda, c > 0$$

the price is determined by market clearing:

$$p_t = \frac{a - \tanh \lambda(p_t^e - c) - 1}{b} \quad (16)$$

There is a unique fixed point and the dynamic of prices depends on agents' expectations. Under this specification of demand and supply and supposing that prices are forecasted using a simple adaptive expectations rule with constant gain the model displays very rich dynamics such as chaotic price fluctuations (see [13, 14] for more details).

A total of 77 subjects took part in the experiment of Hommes *et al.* [15].⁴ Participants were asked to predict a number, p_t^e , between 0 and 10 at each time period. Then the actual price was generated as follows

$$p_t = \frac{a - \tanh \lambda(p_t^e - c) - 1}{b} + \varepsilon_t$$

that is by perturbing equation (16) with a random additive shock. Throughout the experiment the participants could see on their screens their past predictions p_t^e , the true values p_t for all rounds up to the last, as well as last and total earnings. Each session lasted 50 time periods and rewards were proportional to the forecasting accuracy.⁵ There were some little differences in the experimental design across the agents. One group of agents was asked to predict a sequence of *values* without any information about how this sequence was generated; the other group of agents was asked to predict a sequence of *prices* and was informed that a certain market structure was underlying the model but no information was given either about the specific market or about the occurring shocks. The second difference is in the choice of parameters and additive shocks. For a first group the market equation was specified setting $a = 2.3$ and adding an i.i.d. shock distributed as a normal with zero mean and variance $\sigma^2 = 0.5$. For a second

⁴In each session computers were allocated an identification number: these go from 2 to 124 in the data set and obviously not all of them were used (due to the uncertainty about how many participants would actually show up). In the present paper we have used the identification numbers to distinguish the subjects in order to ensure the comparability with the paper by Hommes *et al.* [15] who in turn used them.

⁵Points were awarded at each time period on the basis of the payoff function $\max\{1300 - 260(x_t - x_t^e)^2, 0\}$. These points were then converted at a rate 1300 points = 1 dutch guilder.

group of agents the i.i.d. shock was uniformly distributed with zero mean and variance $\sigma^2 = 0.01\bar{3}$ while the parameter a was not constant, simulating successive demand shocks; in particular a was set to

periods	1 - 15	16 - 28	29 - 40	41 - 50
a	2	3	1.25	2.5

The other parameters, $b = 0.25$, $c = 6$, $\lambda = 2$, were constant across periods and agents.

3.1 Data analysis

In their paper Hommes *et al.* [15] have analysed the experimental data to understand whether any inference can be drawn on the subjects' behaviour. They show evidence of heterogeneity in the way people form expectations and find that there is a link between which rule agents used and how successful they were (how much they earned). In particular they distinguish three categories of agents on the basis of the particular learning rule they seem to have used: 1) agents using a rule that they call *markov* expectations (*i.e.* entailing a random switch between the last expected value and the last observed value), 2) agents using some kind of adaptive expectations, 3) agents who did not use any systematic rule. Agents belonging to the second category are those who, on average, earned more.

In what follows we want to take this analysis one step forward. In particular we try and answer the questions: Can we be more precise about what "types" seem to emerge? How do the rules discussed in Section 2 fare in describing the behaviour of each agent? Is there a difference between recursive and extensive formulations?

Our analysis is twofold.

First, we have tried to group the participants in categories homogeneous with respect to how (and if) they exploited the information conveyed by their own forecast errors. Supposing that the agents used an adaptive rule like

$$p_{t+1}^e = p_t^e + \alpha_t(p_t - p_t^e) \quad (17)$$

for each agent and each period we have calculated the value of α_t feeding the data from the experiment into equation (17). Then we have plotted, for each agent, the values of α_t against the last forecast error. Figure 1 shows the cumulative diagram, plotted for all the participants. At first glance no particular pattern seems to emerge. In fact it can be seen as the result of the superimposition of four typical cases. The participants to the experiment can be classified on the basis of the particular shape of their diagram in one of the

four types shown in Figure 2. To do so we have used the following heuristic rule. First we have analytically defined the three regions in evidence in Figure 2 (relative to types 1, 2 and 3)⁶. For each agent and each time period, we have checked if the couple $(error, \alpha)$ satisfies some of the inequalities. Finally we have assigned to each agent the type whose region is visited the maximum number of times subject to the fact that it has been visited at least 35 times. If no region has been visited at least 35 times over 50 we assign the residual type 4. So types are attributed when they characterise agents at least in 70% of the time periods: this threshold, while largely heuristic, is broad enough to permit actual classification of subjects in an experiment in which random components, high nonlinearity and strong feedbacks would otherwise make it impossible to detect any significant pattern in agents' behaviour.

Figure 3 shows 3 out of 20 participants who are classified as type 1: these agents typically react to their own forecast error $p_t - p_t^e$ either by setting $p_{t+1}^e = p_t^e$ or $p_{t+1}^e = p_t$ (see Figure 2), actually doing one thing or the other at random. This class of agents produces *markov* expectations in the terminology of Hommes *et al.* [15]. Agents using adaptive schemes correspond to type 2 (15 agents in total, see Figure 4 for three typical cases) and type 3 (5 agents, Figure 5): the latter are using simple adaptive expectations with constant gain while the former are using a variable gain parameter which implies that agents react to the forecast error (higher gain follows higher error) in a way that is captured, for example, by the GFM and GAE mechanisms described in Section 2. The residual class of type 4 is characterised by the lack of any clear criterion in the use of the forecast error.

The second building block of our analysis tries to associate each subject with one of the specific rules described in Section 2. To do so we assume that agents form expectations using a predictor chosen within a given set and, once chosen, they stick to it. We also suppose that agents are using the predictor with a certain degree of approximation, which can be the result of a lack of full attention to the data due to the short time or to their limited computational capacity. The same assumption can be found in Branch [4].

In detail, we suppose that agents form expectations according to a law of the form

$$p_{t+1}^e = G(p_t, \dots, p_0, p_t^e) + \varepsilon_t \quad (18)$$

where G is chosen by each agent, once and for all, in the set of predictors specified in Section 2 and ε_t represents a sort of "tremble". For each agent and each available predictor, we generate artificial time series of expectations

⁶The regions are defined as follows: (1) $(x \pm 7.5)^2 + [10(y + 0.5 \pm 0.45)]^2 < 9$; (2) $\frac{x^2}{100} - 0.2 < y < \frac{x^2}{40} - 0.5$; (3) $-0.1 < y < 0.4$.

p_t^a where

$$p_{t+1}^a = G(p_t, \dots, p_0, p_t^e)$$

and p, p^e are taken from the experimental data. Then we compute, for each predictor G , the mean square difference between artificial and experimental forecasts

$$\delta_G = \frac{1}{50} \sum_{t=1}^{50} (p_t^e - p_t^a)^2.$$

All the parameters in the prediction functions are calibrated for each agent. For the predictors that entail a parametric choice we have numerically selected the best parametrisation in terms of implied mean square difference; for the random predictors⁷ we have considered the average of the mean square difference on a 1000 times replication of the random choice. In the case of least squares and stochastic gradient a projection facility is applied to ensure that the predictions belong to the interval $[0, 10]$.

For the predictors that admit an *extensive* formulation which is equivalent to the recursive form, we have actually distinguished two cases in the generation of the artificial expectations. There is a rather subtle but important point here: indeed, under our assumptions (18) regarding the tremble ε_t , the two choices are not equivalent. If agents are genuinely using the extensive form and therefore are averaging the observations of p , each prediction differs from the value indicated by the predictor only because of ε_t . Conversely if agents follow the recursion, they use p_t^e (and therefore ε_{t-1}) to generate p_{t+1}^e so that the prediction will be different from what it would be if there were no trembles because of $\varepsilon_t, \varepsilon_{t-1}$ and recursively every ε_{t-k} with $k = 0, \dots, t-1$.

Table 1 contains the results of the above exercise. The Subject column identifies the person with a number running from 2 to 124 (see footnote 4). The second column, labelled \$, shows the total earning for the subject (subjects are in fact sorted according to earning starting with the one who earned the most down to those that earned nothing). The third column identifies the type according to our definitions above. The fourth column shows the rule that does best in describing the subject's actual behaviour in the experiment.

The evidence of a connection between the total earning gained by each subject and the rule that best describes the subject is certainly not overwhelming, with the possible exception of the OLS₂ rule which seems to have been used by subjects in the upper earnings ranking. Nonetheless a link between earnings and type can be clearly inferred; higher earnings are clearly

⁷We have specified RE as $x_t^e \sim U(0, 10)$.

associated with types 2 and 3. Indeed, consider the following table:

Type	δ_G	Average earnings
1	15.16242958	2530.45
2	2.745266594	35806.3
3	0.977983936	46368.8
4	6.106350931	11814.7

Column 2 shows the average, on the total of agents belonging to each class, of the mean square difference between expectations from the experiment and those artificially generated. That is, δ_G is a measure of the goodness of fit of our model (18) for the four types of agents. Column 3 has the average earnings. Agents of type 2 and 3 not only are those who earn more: they are also the ones whose behaviour can be well approximated using a predictor chosen in the set described in Section 2. This is not the case for agents of types 1 and 4. Furthermore, as expected, agents of type 2 and 3 are best represented by one of the generalisations of adaptive expectations (GFM, GA, MN) in 13 out of 20 cases; by fading memory in 2 cases (again an adaptive law with non-constant gain) while in 5 cases the best fitting predictor is in the OLS family. Three cases of type 2 participants are shown in Figure 4; their best fitting predictor are, GA, GFM and MN, from left to right respectively. As required by this family of predictors the scatter of points is, vaguely, bell-shaped. Figure 5 shows type 3 participants; they are associated, from left to right, to FM, MN and OLS₂. Participant 110 is typical of those well represented by a predictor of the OLS family. Most points are concentrated around the origin showing that the dynamics rapidly converges to the steady state. Observe that predictors of the generalized adaptive family do better than the others also in representing behaviour of type 1 and 4 participants. Obviously this does not mean that those agents are mostly adaptive. In fact, none of the rules we considered does well in reproducing the behaviour of such agents.

A few more facts seem to emerge. The first is that there is no single rule that can effectively describe the behaviour of all agents: therefore a significant amount of heterogeneity can clearly be detected. The behaviour of most of the experimental subjects is best described by the rules that generalise the adaptive baseline in the sense of paying more attention to the most recent available observation after a relatively high forecast error and vice-versa. Indeed, as shown in Table 2, almost 70% of the subjects can be described as using one among ‘‘Marcet-Nicolini’’ (MN), Generalized adaptive (GA) and Generalized fading memory (GFM).

To have a measure of how well each predictor works we calculate the average over all the participants of the mean square differences. The results are

shown in Table 3. The four algorithms that are more successful in describing the agents' behaviour are in recursive form: this agrees with the idea that when decisions have to be taken quickly people rely on "short memory" information and simple rule of thumbs. Remark that the rule labeled OLS_2 fares a good deal better than its counterparts OLS_1 and OLS_1 extensive. This probably reflects the fact that OLS_1 is misspecified with respect to the true data generating process, in the sense that it implies that agents using it would have missed the feedback from their expectations to the data: which most agents actually detected in some form or other. OLS_2 instead correctly posits an expectations feedback (even though in an underparametrised form). Further note that a rule like the simple average (in both its possible forms), which is simple and is known to be very efficient in a context of strong expectations' feedback, is definitely not a good description of the way the subjects tackled the experiment. The same applies to stochastic gradient learning. In fact these two rules perform similarly to the two rules with random components in terms of descriptive capacity. Indeed the random adaptive rule fares, on average, better than several other (non-adaptive) predictors: this can be taken as an indirect signal of the good performance of the adaptive family.

Data about the worst and best performance of each predictor over all the participants, recorded in Tables 4 and 5, confirm the above observations.

4 Simulations

On the basis of our above analysis of the experimental data of Hommes *et al.* [15], each agent has been assigned a forecast rule (and a set of parameters where needed) which best describes its behaviour during the experiment. It is a natural question to ask which kind of asymptotic behaviour these rules would have implied were they used on a long time horizon. In particular, due to the particular economic model underlying the experiment, which has a single rational expectations equilibrium, it is interesting to see whether or not there is evidence of convergence towards the RE equilibrium if we observe the behaviour of the rules assigned to the agents in a 500 time periods simulated exercise.

In detail: for each agent we have taken the forecast rule best describing its behaviour and simulated the model used in the experiment (in fact we have only used the scenario with a constant $a_t = 2.3$ and the normally distributed shock $N(0, 0.5)$ for simplicity) for 500 periods (with the same shocks used for all agents). We have collected data on the variance for the simulated price expectations in the following way. Let $Var(T)$ be the variance calculated on

p_T^e, \dots, p_{T-49}^e that is on a window of 50 time periods ending at time T . We have recorded in particular $Var(50), Var(150), Var(450)$. This has been repeated 1000 times for each agent and averages (over the replications) for the variances at the various T have been computed. Remark that the distinction between recursive and extensive form of a predictor ceases to be important in this context because the two versions are identical.

A summary of the results is presented in Table 6. The first column contains the specifications used for the rules that involve parameters.⁸ The last row in each box displays averages over the different parametric specifications: it can be taken as a rough measure of how much each type of rule is inclined to imply convergence toward the rational expectations equilibrium (a lower value means convergence is more likely). Notice that MN and the non-parametric rules (OLS1, OLS2 and Simple Average) have the lowest values: this does not mean that convergence is warranted in all cases. In fact, a number of types can be detected in terms of qualitative behaviour which we can describe as follows.

1. Fast convergence to the RE equilibrium: this means that within the initial 10 time periods expectations enter a very small neighbourhood of the equilibrium and stay there
2. Eventual convergence to the RE equilibrium: expectations fluctuate around the equilibrium in the initial time periods and then eventually they converge
3. Erratic behaviour: expectations continue to bounce erratically around the equilibrium
4. Persistent small fluctuations: expectations fluctuate but stay near the equilibrium

We present various typical examples of the four types above for the different rules in Table 7. Notice that:

- With adaptive expectations (under the two parametrisations used here) the outcome is always erratic behaviour
- Fading memory and Generalised adaptive expectations are the only rules under which persistent small fluctuations show up

⁸The various sets of parameters have been (numerically) determined to best describe the agents' behaviour, as explained in Section 3.1.

- The three non-parametric rules imply fast convergence (this feature is well known from the theory)
- The simulated dynamics of MN agents show convergence in all but one case. This is at odds with the fact that many agents in the bottom of the earning ranking in the experiment are associated with an MN rule: this contradiction stems from the already mentioned difficulty in representing the behaviour of agents of types 1 and 4 within our set of predictors.

5 Conclusions

As concluding remarks we can say that our analysis confirms previous speculations (and evidence) that agents in experimental settings show a preference for simple rules. In the experiment we have analysed, many subjects seem to have tested several of these simple rules: some eventually settled on one choice; others, particularly in the lower part of the earning ranking, went on trying (mostly without success). An implication of this is that adaptive schemes with variable gain parameter have a potential to describe rather well what real agents actually do. It is well known that these rules do not possess general properties of convergence to rational expectations: this is reflected in the fact that a good number of the subjects earned very little or no money at all in the experiment. The precise way to model this type of adaptive behaviour in fruitful and analytically tractable ways is the subject of an interesting and lively area of current research. For instance Evans and Ramey [11] analyse this issue with an eye to the scope of applicability of the Lucas Critique. On the empirical side, the particular model underlying the experiment of Hommes *et al.* [15] features a strong expectations feedback which could have had a good deal of influence on the outcome of the analysis: in particular the behaviour of type 1 agents could have been the result of the switch between two different rules. These aspects could be focused on, analysing experiments on expectations formation without feedback and/or assuming some kind of discrete choice dynamics in the line of Brock and Hommes [6]: an example of this (with survey data) can be found in Branch [4].

6 Figures and tables




Figure 1: Gain parameter vs forecast errors: all agents.




Figure 2: Four scatter types




Figure 3: Type 1: Markov expectations




Figure 4: Type 2: Adaptive expectations with non-constant gain




Figure 5: Type 3: Adaptive expectations with constant gain




Figure 6: Type 4: No pattern

Subject	\$	Type	Best fitting rule
103	55929	3	Marcet-Nicolini
110	52449	3	OLS2
31	52007	3	Fading memory
7	51625	2	Generalised adaptive
122	45372	2	OLS2
124	45360	2	Generalised adaptive
3	43216	2	Marcet-Nicolini
52	40790	2	OLS2
106	40591	2	OLS1 - ext.
5	38972	2	Generalised adaptive
39	36812	2	Marcet-Nicolini
8	36551	3	Generalised adaptive
34	35720	2	Marcet-Nicolini
79	34908	3	Marcet-Nicolini
15	31402	2	Generalised fading memory
18	29231	2	Generalised fading memory
46	28482	2	Marcet-Nicolini
19	26210	2	Fading memory - ext.
120	25294	4	Marcet-Nicolini
65	24743	2	Generalised fading memory - ext.
43	21216	4	Generalised fading memory
2	19978	4	Marcet-Nicolini
113	18835	4	Marcet-Nicolini
104	18569	2	OLS2
51	17797	4	Fading memory - ext.
119	16821	4	Generalised adaptive
41	16700	4	Generalised adaptive
75	16325	4	Adaptive
69	16307	4	Adaptive
73	15657	4	Generalised adaptive
121	15619	4	Generalised fading memory
109	15546	4	Generalised adaptive
115	15314	4	Generalised fading memory - ext.
81	15292	1	Generalised fading memory
45	15083	4	OLS1 - ext.
84	14528	4	Fading memory
107	14337	4	Generalised fading memory - ext.
53	13828	4	Generalised adaptive
123	13249	4	OLS1

Subject	\$	Type	Best fitting rule
47	12775	4	Marcet-Nicolini
35	12340	4	Generalised fading memory - ext.
49	11633	4	Generalised adaptive
117	11304	4	Generalised adaptive
40	10888	4	Marcet-Nicolini
114	10771	4	OLS2
77	10635	1	Marcet-Nicolini
36	9582	4	Fading memory - ext.
118	9035	4	Generalised adaptive
17	8889	4	Simple Average - ext.
101	8819	1	Generalised fading memory
9	7761	1	Fading memory - ext.
48	7009	4	Simple Average - ext.
4	6074	4	Marcet-Nicolini
13	5649	4	Generalised fading memory
105	5158	4	Marcet-Nicolini
42	4809	1	OLS2
33	4248	4	Generalised fading memory
63	2814	4	Marcet-Nicolini
44	2639	4	Generalised fading memory
64	2109	1	Marcet-Nicolini
66	1469	4	Generalised adaptive
108	1336	4	Generalised adaptive
12	1184	1	Marcet-Nicolini
37	1097	4	Generalised adaptive
6	0	1	Marcet-Nicolini
10	0	1	Marcet-Nicolini
14	0	1	Fading memory
16	0	1	OLS1 - ext.
20	0	1	Marcet-Nicolini
30	0	1	Marcet-Nicolini
32	0	1	Marcet-Nicolini
38	0	1	Marcet-Nicolini
50	0	1	Marcet-Nicolini
62	0	1	Marcet-Nicolini
68	0	1	Fading memory
78	0	1	Fading memory
102	0	1	Marcet-Nicolini

Table 1: Detailed result.

Learning rule	Nr
Marcet-Nicolini	26
Generalized adaptive	15
Generalized fading memory	9
OLS2	6
Fading memory	5
Fading memory - ext.	4
Generalized fading memory - ext.	4
OLS1 - ext.	3
Adaptive	2
Simple average - ext.	2
OLS1	1
Random adaptive	0
Simple average	0
Stochastic gradient	0
Myopic	0
Random	0

Table 2: Nr of agents best described by each rule

Learning rule	Avr. δG
Generalized adaptive	6.525994509
Marcet-Nicolini	6.560275815
Generalized fading memory	6.589226707
Adaptive	6.797988622
Fading memory	6.857585565
OLS2	7.585184099
Generalized fading memory - ext.	7.687932099
Fading memory - ext.	7.811571059
Simple average - ext.	8.100513318
Random adaptive	10.04380802
Simple average	11.39200503
OLS1 - ext.	12.17410884
Stochastic gradient	15.18161102
Random	16.00872408
Myopic	17.79218868
OLS1	21.88277349

Table 3: Average mean square difference

Learning rule	Min. δG
OLS2	0.007211171
Marcet-Nicolini	0.007842847
Generalized adaptive	0.015394033
Generalized fading memory	0.035534587
Adaptive	0.047396939
Fading memory	0.050138809
Generalized fading memory - ext.	0.157404916
Simple average	0.169711112
Fading memory - ext.	0.19646911
Simple average - ext.	0.254927161
OLS1 - ext.	0.360572596
OLS1	0.396592126
Random adaptive	0.410735704
Myopic	1.361879592
Stochastic gradient	1.376401445
Random	9.169400848

Table 4: Minimum mean square difference

Learning rule	Max. δG
Marcet-Nicolini	16.31400247
Adaptive	16.77935286
Generalized adaptive	16.77935286
Generalized fading memory	16.96644259
Fading memory	16.97041393
OLS2	17.78178814
Generalized fading memory - ext.	18.97558235
Fading memory - ext.	19.11940716
Simple average - ext.	19.11940716
Random adaptive	22.37439778
Random	26.03696498
Simple average	28.79897912
OLS1 - ext.	33.052582
Stochastic gradient	40.5530226
OLS1	41.23737483
Myopic	42.99791429

Table 5: Maximum mean square difference

Adaptive			
parameters	Var(50)	Var(150)	Var(450)
alpha=0.53	1.939206	1.947048	1.954276
alpha=0.63	3.378096	3.400377	3.407029
Average	2.658651	2.673712	2.680652

GFM			
parameters	Var(50)	Var(150)	Var(450)
k=0.65, h=1, d=0.34	0.465866	0.000623	0.000539
k=0.58, h=0.9, d=0.35	0.510573	0.003817	0.003874
k=0.82, h=0.4, d=0.17	1.047905	0.001336	0.000642
k=0.64, h=0.2, d=0.03	8.047219	7.956341	7.958007
k=0.41, h=1, d=0.49	0.401164	0.006447	0.006566
k=0.22, h=1, d=0.5	0.920233	0.553788	0.553382
k=0.16, h=1, d=0.42	2.053869	1.755098	1.740896
k=0.6, h=0.5, d=0.07	6.975363	6.806976	6.828818
k=0.46, h=1, d=0.44	0.444791	0.006679	0.006799
k=0.16, h=0.2, d=0.83	0.304392	0.000386	0.000274
k=0.31, h=0.6, d=0.65	0.329682	0.001553	0.001557
k=0.7, h=0.1, d=0.22	1.177673	0.013268	0.005255
k=0.74, h=0.1, d=0.01	7.858344	7.684601	7.701003
Average	2.349006	1.906993	1.908278

GADAP			
parameters	Var(50)	Var(150)	Var(450)
k=0.87, h=0.6, d=0.01	0.410712	0.017791	0.017459
k=0.75, h=0.2, d=0.05	0.573391	0.225698	0.226369
k=0.14, h=0.2, d=0.41	1.788105	1.566604	1.562495
k=0.25, h=1, d=0.43	0.846437	0.582759	0.586046
k=0.81, h=0.5, d=0.02	0.466721	0.098349	0.097376
k=0.21, h=0.1, d=0.23	3.956567	3.745693	3.749983
k=0.41, h=0.1, d=0.15	2.897057	2.698823	2.671274
k=0.02, h=0.1, d=0.3	4.757518	4.584501	4.58719
k=0.66, h=0.7, d=0.23	0.335844	0.009488	0.009398
k=0.41, h=0.1, d=0.29	0.929043	0.670795	0.674945
k=0.51, h=0.1, d=0.2	1.039608	0.768191	0.769532
k=0.79, h=0.1, d=0	0.717605	0.366272	0.36675
k=0.71, h=0.1, d=0	1.640462	1.370344	1.381399
k=0.69, h=0.3, d=0	1.925656	1.656495	1.676736
Average	1.591766	1.311557	1.312639

OLS1		
Var(50)	Var(150)	Var(450)
1.475198	0.000251	6.78E-05

OLS2		
Var(50)	Var(150)	Var(450)
1.164862	4.92E-05	5.4E-06

Simple Average		
Var(50)	Var(150)	Var(450)
0.495319	0.000151	1.92E-05

FM			
parameters	Var(50)	Var(150)	Var(450)
rho=0.69	0.626739	0.282944	0.286039
rho=0.8	0.363911	0.032439	0.031913
rho=0.75	0.436926	0.096076	0.095354
rho=0.66	0.746773	0.42246	0.421826
rho=0.78	0.391682	0.049884	0.049199
rho=0.77	0.407651	0.062182	0.061421
rho=0.96	0.314041	0.001306	0.001269
rho=0.87	0.336399	0.008676	0.008502
Average	0.453015	0.119496	0.11944

MN			
parameters	Var(50)	Var(150)	Var(450)
alpha=0.29, nu=0.76	0.337208	0.000151	2.08E-05
alpha=0.49, nu=1	0.313904	0.000151	2.07E-05
alpha=0.39, nu=0.78	0.323923	0.000154	2.08E-05
alpha=0.31, nu=0.78	0.317174	0.000151	2.08E-05
alpha=0.46, nu=0.96	0.314836	0.000151	2.07E-05
alpha=0.35, nu=0.75	0.328854	0.000153	2.08E-05
alpha=0.4, nu=0.9	0.314063	0.00015	2.07E-05
alpha=0.58, nu=0.85	0.464584	0.00016	2.11E-05
alpha=0.45, nu=0.84	0.339098	0.000154	2.09E-05
alpha=0.44, nu=0.88	0.320249	0.000152	2.08E-05
alpha=0.69, nu=0.97	0.478401	0.000156	2.1E-05
alpha=0.32, nu=0.96	0.301261	0.000149	2.07E-05
alpha=0.39, nu=0.7	0.337438	0.000156	2.09E-05
alpha=0.41, nu=0.87	0.313554	0.000151	2.07E-05
alpha=0.64, nu=0.99	0.404797	0.000156	2.09E-05
alpha=0.54, nu=0.9	0.372316	0.000155	2.1E-05
alpha=0.35, nu=0.95	0.307026	0.000149	2.07E-05
alpha=0.71, nu=0.93	0.541562	0.00016	2.11E-05
alpha=0.31, nu=0.8	0.337008	0.00015	2.07E-05
alpha=0.57, nu=0.61	1.628573	0.248979	2.92E-05
alpha=0.49, nu=1	0.312749	0.000151	2.08E-05
alpha=0.55, nu=0.96	0.360369	0.000153	2.09E-05
alpha=0.35, nu=0.73	0.330127	0.000153	2.09E-05
alpha=0.48, nu=0.81	0.376919	0.000157	2.11E-05
alpha=0.37, nu=0.15	0.82518	0.558192	0.557236
alpha=0.37, nu=0.57	0.384997	0.000166	2.14E-05
Average	0.422545	0.031187	0.021453

Table 6: Mean variances over 1000 replications




Table 7a: Asymptotic behaviour of simulated dynamics: some typical cases.











Table 7b: Asymptotic behaviour of simulated dynamics: some typical cases.

References

- [1] E. Barucci: Exponentially fading memory learning in forward looking economic models. *Journal of Economic Dynamics and Control*, **24**(5-7), 1027-46, (2000).
- [2] G. I. Bischi, L. Gardini: Basin fractalization due to focal points in a class of triangular maps. *International Journal of Bifurcation and Chaos*, **7**(7), 1555-1577, (1997).
- [3] G. I. Bischi, A. Naimzada: Global analysis of a nonlinear model with learning. *Economic Notes*, **26**(3), 445-476, (1997).
- [4] W. Branch: The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations. *The Economic Journal*, **114**, 592-621, (2004).
- [5] M. Bray: Convergence to rational expectations equilibrium. In R. Friedman and E.S. Phelps (eds.), *Individual forecasting and aggregate outcomes*, Cambridge University Press, (1983).
- [6] W.A. Brock, C.H. Hommes: A rational route to randomness. *Econometrica*, **65**(5), 1059-1095, (1997).
- [7] C. Chiarella, X. He, P. Zhu: Fading memory learning in the Cobweb model with risk averse heterogeneous producers. *Research Paper Series*, **108**, Quantitative Finance Research Centre, University of Technology, Sydney, (2003).
- [8] D. Colucci, V. Valori: Error learning behaviour and stability revisited. *Journal of Economic Dynamics and Control*, **29**(3), 371-388, (2005).
- [9] G.W. Evans, S. Honkapohja: Stochastic gradient learning in the cobweb model. *Economics Letters*, **61**(3), 333-337, (1998).
- [10] G.W. Evans, S. Honkapohja: *Learning and expectations in macroeconomics*. Princeton University Press, (2001).
- [11] G.W. Evans, G. Ramey: *Adaptive expectations, underparametrization and the Lucas Critique*. University of Oregon Economics Department Working Papers, **8**, (2001).
- [12] S. Figlewski, P. Wachtel: The formation of inflationary expectations. *Review of Economics and Statistics*, **63**(1), 1-10, (1981).

- [13] C. H. Hommes: Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. *Journal of Economic Behavior and Organization* **24**, 315-335, (1994).
- [14] C. H. Hommes: On the consistency of backward-looking expectations: The case of the cobweb. *Journal of Economic Behavior and Organization* **33**, 333-362, (1998).
- [15] C. H. Hommes, J. Sonnemans, H. van de Velden: Expectations formation in a Cobweb Economy: Some One Person Experiments. In D. Delli Gatti, M. Gallegati and A. Kirman (eds.), *Interaction and market structure Essays on heterogeneity in economics*. Lecture Notes in Economics and Mathematical Systems - Springer Verlag, 253-266, (2000).
- [16] Honkapohja, S. and Mitra, K.: Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics* **51**(8), 1743-1770, (2004)
- [17] M. Lovell: Tests of the Rational Expectations Hypothesis. *American Economic Review*, **76**(1), 110-24, (1986).
- [18] A. Marcet, T. Sargent: Convergence of least square learning in self-referential stochastic models. *Journal of Economic Theory*, **48**, 337-368, (1989).
- [19] A. Marcet, J. Nicolini: Recurrent hyperinflations and learning. *American Economic Review*, **93**(5), 1476-1498, (2003).
- [20] R. Marimon, S. Sunder: Expectations and learning under alternative monetary regimes: an experimental approach. *Economic Theory*, **4**(1), 62-131, (1994).
- [21] M. Rabin: A Perspective on Psychology and Economics. *European Economic Review*, **46**(4-5), 657-685, (2002).
- [22] R. Schmalensee: An experimental study of expectation formation. *Econometrica*, **44**, 17-41, (1976).
- [23] V. Smith, G. Suchanek, A. Williams: Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, **56**, 1119-1151, (1988).
- [24] J. D. Sterman: Deterministic chaos in models of human behavior: methodological issues and experimental results. *Systems Dynamics Review*, **4**, 148-178, (1988).

- [25] J. D. Sterman: Misperception of feedback in dynamic decision making, *Organizational Behavior and Human Decision Processes*, **43**, 301-355, (1989).
- [26] J. D. Sterman: Deterministic chaos in an experimental economic system. *Journal of Economic Behavior and Organization*, **12**, 1-28, (1989).
- [27] J. D. Sterman: Modeling managerial behavior: misperceptions of feedback structure in a dynamic decision making experiment. *Management Science*, **35**, 321-339, (1989).
- [28] A. Williams: The Formation of Price Forecasts in Experimental Markets. *Journal of Money, Credit and Banking*, **19**(1), 1-18, (1987).