MANCHESTER 1824

The University of Manchester

Centre for Growth & Business Cycle Research

CGBCR

Discussion Paper Series

Fundamentalists vs. chartists: learning and predictor choice dynamics

By

Michele Berardi

Centre for Growth and Business Cycle Research, Economic Studies, University of Manchester, Manchester, M13 9PL, UK

September 2008
Number 104

Fundamentalists vs. chartists: learning and predictor choice dynamics

Michele Berardi*
The University of Manchester

September, 2008

Abstract

In a simple, forward looking univariate model of price determination we investigate the evolution of endogenous predictor choice dynamics in presence of two types of agents: fundamentalists and chartists. We find that heterogeneous equilibria in which fundamentalists and chartists coexist are possible, even when the fraction of agents is endogenized. We then combine evolutionary selection among heterogeneous classes of models with adaptive learning in the form of parameter updating within each class of rules and find that equilibria in which chartists constantly outperform fundamentalists seem never to be learnable. Simulations also show that, in general, interactions between learning and predictor choice dynamics do not prevent convergence of both processes towards their equilibrium values when the equilibrium is E-stable.

Key words: heterogeneity, expectations, predictor choice, learning.

JEL classification: C15, C62, D83, D84, E37.

1 Introduction

Expectations have long been recognized to play a key role in macroeconomic and finance models. After the rise of rational expectations (RE) in the 70s, in more recent years economists have acknowledged the importance that bounded rationality can have on (macro)economic dynamics and the role it should play in modelling real life situations (Sargent, 1993). In particular,

*I would like to thank Cars Hommes for useful comments on an earlier draft.

the RE assumption, at least in its strict sense, does not allow for heterogeneity of expectations, a feature that seems so commonly observed in real economic situations when agents have to make forecasts.

For some years now, a growing literature has been developing on learning in macroeconomics (Evans & Honkapohja, 2001). The main idea is to replace RE with adaptive learning schemes that can (potentially) converge in the limit to RE.

When making economic decisions, agents often need to forecast some future state of the economy: under adaptive learning, they do it by means of a model (perceived law of motion - PLM) that in their mind well represents the dynamics of the variable(s) they are trying to forecast. Moreover, the parameters of that model are fine tuned by recurrent estimation on the basis of the new information that becomes available over time.

We follow this approach and let agents choose between two possible models, one that corresponds to the MSV solution for the economy, and the other that is a simple AR(1) model. Agents who choose the first model are called fundamentalists, as they correctly recognize the exogenous driving force(s) for the economy and use them in their forecasts; agents who instead choose the AR(1) model are called chartists, as they only base their predictions on past data, trying to extract (and exploit) the autocorrelation structure present in the data.

In order to impose some rigor on what would otherwise be an arbitrary imposition of heterogeneity, we let agents choose endogenously among the two different predictors, i.e., we let the fraction of fundamentalists vs. chartists be determined endogenously in the economy.

We model the choice between predictors using a logit model, following the seminal work by Brock and Hommes (1997). This modeling strategy captures a sort of bounded rationality, in the sense that the best performing model (in any measure we care to use) is not chosen by all agents all the time. Moreover, parameters of the two models have to be fine tuned to the economy, in order to deliver the best possible forecasts, i.e., to minimize their mean squared errors. We start by assuming that these optimal values are known to agents, and then, in a later section, we drop this assumption and let agents learn those values through econometrics techniques such as recursive least squares (RLS). In this way, we are able to combine adaptive learning, in the form of parameter updating within a class of models, with evolutionary selection among heterogeneous classes of models.

The main question the paper tries to answer is whether fundamentalists and chartists can survive together in an economy, or if instead fundamentalists (chartists) will always run chartists (fundamentalists) out of the market in the long run. Moreover, we want to investigate the interaction between two types of learning, within a model and across models, to understand how they might affect each other.

In the literature, heterogeneous expectations with endogenous predictor choice have been analyzed in particular in the contest of a Cobweb model. In Brock and Hommes (1997) agents can have either rational or naive expectations, while in Lasselle et al. (2005) agents are split between rational expectations and adaptive expectations. In De Grauwe, Dewachter and Embrechts (1993) the authors consider models of exchange rates with two classes of agents, fundamentalists and chartists; while Chiarella (1992), Lux (1995) and Sethi (1996) analyze stock market models with fundamentalists and chartists. None of these works, though, couple the endogenous determination of the relative fraction of agents using each model with learning dynamics.

Guse (2005) and Berardi (2007) analyse heterogeneous learning, but a limitation of those works is that the proportion of agents using each model is exogenously determined. Guse (2007) tries to amend this problem in the same framework as Guse (2005) by introducing a game theoretic analysis of interactions between two type of agents endowed with different perceived laws of motion. Finally, Branch and Evans (2006) consider a Cobweb model with heterogeneous expectations, where agents can use different (but all misspecified) models, and show that in the limit intrinsic heterogeneity is possible when the correctly specified model is not available to agents.

Our work is somewhat in the spirit of Guse (2007), though it differs from it in many dimensions as it considers a different structural model, different types of expectations and a different endogenous predictor choice mechanism: in particular, we analyze a simple model of price determination where fundamentalists and chartists interact through learning and a discrete logit model of predictor choice. It is meant to give a simplified representation of how an asset price (or any other price that displays positive feedback from expectations of future prices on current values) is determined on the market. The main difference with respect to Branch and Evans (2006), a part from the type of expectational feedback that in their work is negative, is that we want to study whether heterogeneity can survive when the correct model of the economy is available to agents.

The structure of the paper is as follows: Section 2 introduces the model; Section 3 solves it under heterogeneous expectations and analyses its E-stability; Section 4 introduces endogenous predictor choice dynamics, analyses E-stability of the endogenous equilibrium and simulate the ensuing system; Section 5 introduces real-time adaptive learning and analyzes its interplay with predictor choice dynamics; Section 6 concludes.

2 The model

The model representing the system is a simple univariate forward looking model

$$y_t = AE_t y_{t+1} + Bw_t \tag{1}$$

where y is the price level and w the exogenous driving force represented by a stationary AR(1) process:

$$w_t = \rho w_{t-1} + \varepsilon_t \tag{2}$$

with $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ and $0 \leq \rho < 1$.

The model can be thought of as a model of financial markets, where net demand (and therefore current price) of the asset depends positively on its future expected price and on a fundamental driving process (here w_t) that could for example represent firm's profits. A is strictly positive, meaning that there is a positive feedback from expectations of future values on current values, as it is common in financial markets. The semi-reduced form (1) could also represent a model of exchange rates, with the exogenous process representing, for example, interest rates or inflation differentials.

There are two types of traders, fundamentalists and chartists. Fundamentalists use the right model, in the sense that it is consistent with the MSV solution in which the endogenous variable depends only on the exogenous driving process. Chartists use instead an AR(1) model in the attempt to discover statistical regularities in the time-series for prices.

With homogeneous expectations, for $A < 1$ the model is determinate and for $A < 1/\rho$ the MSV equilibrium is E-stable. The MSV REE is

$$y_t = \frac{B}{1 - \rho A} w_t. \tag{3}$$

For $A > 1$ the model is indeterminate and there exist an infinite number of AR(1) solutions, none of which is E-stable. This last result is derived assuming that agents use a PLM of the type $y_t = aw_t + by_{t-1}$. Note that while this solution is never learnable for any parameterizations, the restricted perception equilibrium (RPE) that would emerge if agents were to use the model $y_t = by_{t-1}$ will turn out to be learnable for a wide range of parameterizations. This RPE is in fact an example of a consistent expectations equilibrium (CEE) in the sense of Hommes and Sorger (1998): here the true law of motion is a linear stochastic process driven by an exogenous linear random process, while agents believe in a simple AR(1) process driven by i.i.d. noise.

In the heterogeneous setting, there is a unitary mass of agents split between α fundamentalists (denoted by superscript f) who correctly recognize the exogenous driving process and use information about it to form their expectations

$$PLM^f : y_t = aw_t \Rightarrow E_t^f y_{t+1} = a\rho w_t; \tag{4}$$

and $(1-\alpha)$ chartists (denoted by superscript c), who use past observations to predict future ones, using a simple AR(1) model

$$PLM^c : y_t = by_{t-1} \Rightarrow E_t^c y_{t+1} = b^2 y_{t-1}. \tag{5}$$

Parameters a and b are chosen by agents as optimal linear projections, so as to minimize the mean square errors (MSEs) incurred in making forecasts:

$$\bar{a} \quad : \quad E[w_t(y_t - aw_t)] = 0 \tag{6}$$

$$\bar{b} \quad : \quad E[y_{t-1}(y_t - by_{t-1})] = 0. \tag{7}$$

Note that when there are chartists in the economy, the model used by fundamentalists is also misspecified (underparameterized), as it misses the lag component that is introduced into the ALM by chartists.

3 Heterogeneous expectations equilibrium

The equilibrium values $\bar{a}$ and $\bar{b}$ depend on the parameters A, B, ρ and α. While the first three are structural parameters, α should be regarded as endogenous, coming from the choice of agents about their forecasting model. As a first step, we start here by treating α as a structural and exogenous parameter and compute the heterogeneous expectations equilibrium (HEE)[1] for different values of it. In a later section, we will then relax this assumption and model α as the outcome of endogenous predictor choice dynamics.

Substituting agents' expectations from (4) and (5) into (1) we can find the ALM for y_t

$$y_t = (A\alpha a\rho + B)\, w_t + A(1-\alpha)b^2 y_{t-1}$$

[1] For a definition of the concept of HEE used here, see Berardi (2007). From a learning perspective, in a HEE parameter estimates have converged to the optimal value, the one that minimizes the MSE, for all agents invoved.

which can then be used to solve (6) and (7)[2] for the equilibrium values of a and b:

$$\bar{a} = \frac{B}{1 - A\alpha\rho - A(1-\alpha)\rho\bar{b}^2} \tag{8}$$

and

$$\bar{b} : b^3 A\rho(1-\alpha) - b^2 A(1-\alpha) + b - \rho = 0. \tag{9}$$

As it can be seen, $\bar{a}$ depends on $\bar{b}$ and the solution for the latter comes from a 3^{th}-degree polynomial, which we solve numerically. There are of course 3 solutions, of which at least one real. Using Descartes' Sign Rule we can see that, for A and ρ positive, there are a maximum of 3 positive real roots and a maximum of 0 negative real roots, i.e., all the real roots will be positive.[3] Numerical computations show that, in the region of the parameter space we will consider, there is almost always only one real solution; 3 real solutions emerge only for low values of ρ, and in a very narrow region of the space. From (8), we can see that as $\alpha \to 1$, we obtain again the homogeneous MSV RE solution (3).[4]

Note that for $\rho = 0$, there are two solutions: $\{\bar{a} = B, \bar{b} = 0\}$ and $\{\bar{a} = B, \bar{b} = 1/A(1-\alpha)\}$. If we substitute the ensuing expectations into the structural model, in the first case (with $\bar{b} = 0$) we obtain the MSV RE solution $y_t = Bw_t$ (and fundamentalists turn out to be correct), while for $\bar{b} = 1/A(1-\alpha)$ we have $y_t = 1/A(1-\alpha)y_{t-1} + Bw_t$, which is a linear combination of the two models used by the two groups. Ex post, thus, both fundamentalists and chartists are making mistakes, the first because they ignore the effect on prices introduced by the second. This last equilibrium, though, is never E-stable. Nonetheless, imposing $\rho = 0$ will allow us to derive some analytical results in a later section, where we will use this case to help understand the more complicated case with $\rho > 0$.

Note that for $\rho = 0$, there is no intrinsic autocorrelation in the driving process, and all the autocorrelation in prices comes from chartists' behavior: in this sense, it is a sort of "self-confirming" autcorrelation; for $\rho > 0$,

[2] The main difficulty in solving for the optimal values for a and b is that variances and covariances involving the endogenous variable depend, in equilibrium, on the value of the parameters that are to be found. Variances and covariances needed in order to find the solutions are available in the Appendix.

[3] This result relies on the restrictions imposed above, namely A, $\rho > 0$. We are not considering here models with negative feedback from expectations or with negative autocorrelation in the exogenous driving process.

[4] Note that for $\alpha = 1$, the model implies $b = \rho$: as there are only fundamentalists in the economy, the autocorrelation in the price level is the same as that in the exogenous driving force. Of course, there are no chartists actually using the equilibrium value for b.

instead, there is some persistence in the driving process that is passed on to prices, and this can be picked up (and increased) by chartists. This could explain why in the first case the HEE is not E-stable, while in the second case it can be.

3.1 E-stability of HEE

We are interested now in checking whether or not the equilibrium just found could be the outcome of adaptive learning dynamics. Suppose agents did not know the equilibrium values for a and b, could they learn these values through adaptive learning schemes such as recursive least squares? To answer this question, we will make use of the concept of E-stability.[5]

In order to assess E-stablity, we need to derive the ODEs governing the limiting behavior of the parameters being learned by agents. Given the misspecifications in the PLMs, we can not map directly the PLMs into the ALM, but we need to project the ALM onto the restricted space of the PLMs. Using the stochastic approximation approach, we can thus derive the differential equations representing the T-maps from PLMs to the ALM:

$$\dot{a} = \frac{\Gamma}{1-\rho\Omega} - a \tag{10}$$

$$\dot{b} = \Gamma\left(\sigma_y^2\right)^{-1}\sigma_{w,y-1}^2 + \Omega - b = \frac{\left(1-\Omega^2\right)\rho}{(1+\rho\Omega)} + \Omega - b \tag{11}$$

where Γ, Ω, σ_y^2 and $\sigma_{w,y-1}^2$ are as defined in Appendix A and all depend on the value of a and b, but for simplicity of notation we do not write out this dependence explicitly. By evaluating the Jacobian J (see Appendix B) at the points $(\bar{a},\bar{b})$ that solve (6), (7),[6] we can evaluate E-stability of the equilibrium: this requires the two eigenvalues of J to have negative real part.

We consider the parameter space delimited by $.01 \leq A \leq 3$,[7] $.01 \leq \alpha \leq .99$ and $\rho = \{.3, .6, .9\}$ (with $B = .1$) and analyze E-stability of the real solutions: it turns out that only one solution can be E-stable and the contour map in Figure 1 shows the region of E-stability (light green) for this solution.

The region of E-stability for the HEE seems to be related with the region of E-stability for the homogeneous REE ($A < 1/\rho$), except that for low values of α the region here tends to shrink down and E-stability requires

[5] For a detailed explanation of the concept of E-stability, see Evans and Honkapohja (2001).

[6] These $(\bar{a},\bar{b})$ are of course also the solutions to (10), (11).

[7] This includes a region of determinacy ($A < 1$) and one of indeterminacy ($A > 1$).

a lower A. Note that for high values of A and low values of α, there are combinations of these two parameters that lead to a violation of the restriction $\Omega\rho < 1$ necessary for having a finite covariance structure in the system (see Appendix): all these combinations, though, reside in a subregion of the E-instability area.[8]

By setting $\alpha = 0$, we are also able to analyse E-stability for the restricted perceptions equilibrium (RPE) that emerges when the economy is populated only by chartists. As we have mentioned before, and is shown in Figue 2, the RPE turns out to be E-stable in a wide region of the parameter space.[9]

4 Endogenous predictor choice dynamics

In order to endogenize the relative size of the two groups (α), we look at the relative performance of the two forecasting models for different values of A, ρ and α. As a measure of performance, we use the difference in the mean squared errors ($MSE^c - MSE^f$) and show it in Figure 3 as contour map, where the green color represents regions of the parameter space where $MSE^c - MSE^f > 0$, while the blue color shows regions where chartists outperform fundamentalists. We show 3 cases, for $\rho = .4$, .6 and .9: as it can be seen, in all cases there is a stripe cutting across the graph where chartists are able to outperform fundamentalists.[10] For values of ρ lower than .4, though, fundamentalists always have smaller forecast errors than chartists: this means that there must be enough persistence in the exogenous data generating process for an AR(1) model to be able to replicate successfully the dynamics of the endogenous variable.

4.1 The general case

We can now use this measure of relative performance to endogenize the fraction of fundamentalists/chartists in the economy. In the previous section we have derived equilibrium values for the endogenous parameters $(\bar{a}, \bar{b})$ as function of the parameters (A, B, ρ, α). Now we endogenize α and look

[8] If we decrease ρ even further below .3, the region of E-stability increases rapidly and comprises an area where the restriction on $\Omega\rho$ is not satisfied: here, as in the rest of the paper, we do not consider such cases, since when the restriction on $\Omega\rho$ is violated the model breaks down and can not deliver meaningful results.

[9] It turns out that, also in this case, only one solution to (7) can be real and E-stable in the region where the restriction on $\Omega\rho$ is satisfied. In particular now, also for very low values of ρ this condition remains satisfied, and the RPE is E-stable.

[10] All these regions are in the area where the restriction $\Omega\rho < 1$ is satisfied. This condition is violated only in a subregion of the area below the blue stripes.

for an equilibrium triple $(\bar{a}, \bar{b}, \bar{\alpha})$ dependent on the parameters (A, B, ρ): α will be now determined endogenously by the relative performance of the two models. This equilibrium concept, where heterogeneity is endogenized, is similar in spirit to the misspecifications equilibrium (ME) introduced in Branch and Evans (2006), except that in their case all agents use a misspecified model while here one group of agents is in fact using the correct model for the economy.[11] This is an important difference, because it will allow us to see whether a correctly specified model always outperforms a misspecified one or if instead intrinsic heterogeneity can persist also when the correct model is available to agents (at no extra cost).

The measure we use to describe the performance of each predictor is the mean squared error (MSE), which comes from assuming a quadratic loss function for the forecast errors:

$$MSE^f = E(y_t - aw_t)^2 \tag{12}$$
$$MSE^c = E(y_t - by_{t-1})^2. \tag{13}$$

Agents choose their predictor (i.e., their identity as fundamentalists or chartists) on the basis of relative performance, according to a discrete logit model. Denoting with α the fraction of fundamentalists, we have:

$$\alpha = \frac{\exp\{-\beta MSE^f\}}{\exp\{-\beta MSE^f\} + \exp\{-\beta MSE^c\}} \tag{14}$$

or equivalently

$$\alpha = \frac{1}{2}\left(\tanh\left[\frac{\beta}{2}(MSE^c - MSE^f)\right] + 1\right), \tag{15}$$

where β is the "intensity of choice" parameter (Brock and Hommes, 1997), a measure of how fast agents switch predictors. For $\beta \to \infty$, we have a deterministic choice model where agents are fully rational and all choose the predictor with the smaller MSE, while for $\beta = 0$ agents split equally between the two predictors independently of the relative performance. The parameter β also acts as a re-scaling factor: as the absolute magnitude of the MSEs (but not their relative one) depends ultimately on the variance of the noise ε_t, different combinations of β and σ_ε^2 give similar dynamics for α. Throughout our reported simulations we will keep $\beta = 10$ and $\sigma_\varepsilon^2 = .1$.

Denoting $\Delta = MSE^c - MSE^f$, and noting how this measure depends on the structural parameters (A, B, ρ) and on α (see Appendix A), we can

[11] Here correct has to be interpreted as consistent with the MSV REE.

write

$$\begin{aligned} \alpha &= f(\Delta;\beta) \\ \Delta &= g(\alpha; A, B, \rho, \sigma_\varepsilon^2) \end{aligned}$$

and therefore, with some economy of notation

$$\alpha = f\left[g(\alpha)\right]. \tag{16}$$

As this function can not be analyzed analytically, we consider the system in real time, transform the function into a map

$$\alpha_{t+1} = f\left[g(\alpha_t)\right] \tag{17}$$

and simulate the system in order to find the long-run behavior of α. Note that as α itself is a function of the structural parameters, the long-run behavior of the system will be ultimately determined by $(A, B, \rho, \beta, \sigma_\varepsilon^2)$.

We will analyse the system for different values of A and ρ, while keeping $B = .1, \beta = 10, \sigma_\varepsilon^2 = .1$. For low values of A, fundamentalists always prevail, and α converges to a fixed point $\bar{\alpha} \in (.5, 1]$. Things become more interesting when there is a strong feedback from expectations to current values. We will thus focus on three scenarios, respectively defined by $(A, \rho) = (1.5, .5), (1.5, .62), (1.5, .99)$.[12] In Figure 4 we report results for the dynamics of α in each of the three different cases. As it can be seen, for low values of ρ, fundamentalists still prevail even when there is high feedback from expectations to current variables. For high values of ρ, instead, chartists are able to outperform fundamentalists, as shown in the third quadrant. For intermediate values (ρ about .62), the system keeps cycling, as depicted in the second quadrant.[13] These cycles increase their frequency as ρ increases to about .63, and then disappear as the system quickly enters into a region where the restriction on $\Omega\rho$ is violated. The same type of behavior happens for different combination of the parameter values: in particular, if A is increased, then a lower ρ is sufficient to trigger cycles and to make chartists outperform fundamentalists.

These simulations show different types of intrinsic heterogeneity that can emerge in this model: the relative size of the two groups can stabilize on a fixed value, with one group or the other prevailing in size, or it can keep changing in a cyclic behavior, as the two groups alternate in taking the lead.

[12] Other combinations of A and ρ can produce the same three scenarios: in general, lower values of A require higher values of ρ in order to produce the same scenario.

[13] In all results reported, we have checked that the condition on $\Omega\rho$ is satisfied.

Note that as α cycles, the equilibrium values of the two forecasting models $(\bar{a}, \bar{b})$ cycle together: this means that agents, in their learning activity, need to learn parameters that are drifting around.

A regularity that has emerged from our simulations is that there must be enough persistence in the exogenous driving process for chartists' AR(1) model to be able to compete with the fundamentalists' one. When there is a strong (positive) feedback effect from expectations to current values, and a high autocorrelation coefficient in the exogenous driving precess, chartists can in fact outperform fundamentalists.

Figure 5 shows an enlargement of the cycle arising for $\rho = .62$. We have initialized $\alpha = .5$. As it can be seen, α decreases slowly for a while, then drops abruptly before jumping up to 1; at this point we have only fundamentalists in the economy, but this situation doesn't last long, as immediately α starts dropping, the fraction of chartists is restored and the cycle repeats itslef.

It is interesting to verify, at this point, the E-stability of the equilibria in each of the three scenario under investigation. Remember that for each quintuple $(A, B, \rho, \beta, \sigma_\varepsilon^2)$ there is an endogenous HEE, defined as a triple $(\bar{a}, \bar{b}, \bar{\alpha})$ that can represent either a fixed point or a cycle. If the equilibrium is a fixed point, we can then evaluate its E-stability directly. If the equilibrium is instead a cycle, it may be that some of the states of the cycle are E-stable and some are not (as indeed it turns out to be the case in our simulations). This in practice means that the whole equilibrium is not learnable, as when the system is at a point in the cycle that is not E-stable, the real-time parameter estimates would diverge and the cycle break down. Figure 6 shows the E-stability results for the equilibria in the three scenarios shown in Figure 4: as it can be seen, in the first case ($\rho = .5$), where fundamentalists outperform chartists, the equilibrium is E-stable,[14] while in the last case ($\rho = .99$), where chartists outperform fundamentalists, the equilibrium is not E-stable. In the intermediate case ($\rho = .62$), where the equilibrium value of α keeps fluctuating, E-stability results are mixed: the system is E-stable in the majority of its states, but becomes E-unstable when α drops to its lowest three values. Thus, according to our criterion, the equilibrium as a whole must be considered E-unstable. These results show that the system can never converge, under endogenous predictor choice and learning dynamics, to an equilibrium where chartists prevail over fundamentalists, as these outcomes are not E-stable.

[14] Here E-stability corresponds to a value of 1 on the y-axes, while E-instability corresponds to 0.

With regards to the E-stability analysis, we must point out that, for this analysis to be meaningful, it must be that agents in our economy update the parameters (a, b) in their model with a higher frequency than that with which they revise the choice of the model (the frequency of the "update" of α in the economy). This is an assumption that seems to be quite reasonable: agents will try to fine-tune continuously the model they are using, but will consider the option of switching to another model only from time to time, as this last option is more costly in terms of information requirements. The minimum relative frequency required would depend on the speed of convergence of the learning process, an issue that we do not address in this paper and leave for future work. We will relax this assumption later on when we simulate real time learning dynamics and allow agents to choose, in each period, both the model and the parameter values in that model.

4.2 The case with $\rho = 0$

We consider now the case with $\rho = 0$, even though we have seen that in this case the HEE (the one that emerges for $\bar{b} = 1/(A(1-\alpha))$) is not E-stable. The reason why we think this case is worth consideration anyway is that in this case we are able to compute analytically the solution and the MSE for the two groups of agents.

It turns out that when $\rho = 0$, $\Delta = MSE^c - MSE^f > 0$ iff $-1 < A(1-\alpha) < 1$: since A and α are positive, the condition requires $A < \frac{1}{1-\alpha}$. In this case, chartists are outperformed by fundamentalists. If instead $A > \frac{1}{1-\alpha}$, then chartists perform better than fundamentalists. This means that in order for chartists to outperform fundamentalists, the positive feedback effect from expectations to current price must be high enough, a result that we had already inferred before from our simulations. Only in this case there is enough persistence in the data that can be picked up and exploited by a forecasting model based purely on past data series.

If we endogenize α, now, we have that

$$\begin{aligned} \Delta &= MSE^c - MSE^f = -\frac{\left[B^2 \left(A\left(1-\alpha\right)\right)^2 + 2B\right]}{\left(A\left(1-\alpha\right)\right)^2 - 1}\sigma_\varepsilon^2 \\ \alpha &= \frac{1}{2}\left(\tanh\left[\frac{\beta}{2}\Delta\right] + 1\right) \end{aligned}$$

and by combining the two expressions above

$$\alpha = \frac{1}{2}\left(\tanh\left[-\frac{\beta}{2}\frac{\left[B^2\left(A\left(1-\alpha\right)\right)^2+2B\right]}{\left(A\left(1-\alpha\right)\right)^2-1}\sigma_\varepsilon^2\right]+1\right).$$

We can thus compute the equilibrium value for α, depending on the fundamental parameters A, B, β and σ_ε^2. In particular, we fix B, β and σ_ε^2 as before and focus our attention on the parameter A.[15] In Figure 7 we show the results: as it can be seen, for $A < 2$ we have that fundamentalists outnumber chartists and $\alpha > .5$, while for $A > 2$ the opposite happens and $\alpha < .5$. Note that when exactly $A = 2$, a 2-perid cycle emerges, in which A oscillates between 0.17 and 0.48.

5 Learning dynamics

We turn now to analyze real time learning dynamics and their interactions with predictor choice dynamics: each period, agents have now to choose which model to use and select the best value for the parameters in that model.[16] In order to choose the model (i.e., to compute MSE^f and MSE^c) they use the most recent estimates of the parameters, both for their model and for the alternative. We assume here that once learned, the value of all parameters becomes common knowledge. This is a common assumption in the literature (see, e.g., Guse (2007)), and it is innocuous here as there are no costs involved in the learning process.

Note that the learning speed for the two PLMs could affect the results, so that even though one model might be better with optimal parameter values, the other could be easier to learn and thus provide more accurate forecasts over the learning path. The time needed to learn a process (i.e., the speed of convergence), and its consequences on the dynamics of the system, is something that is often neglected in the learning literature, with some notable exceptions (e.g., Marcet and Sargent (1995), Beneviste et al (1990) and for an application to monetary policy: Ferrero 2007).

[15] Simulations show that results do not change qualitatively for β ranging in a wide are between 1 and 1000: relatively high values of β only tend to extremize the equilibrium values for α.

[16] We could alternatively choose an asynchronous timing: every period agents reestimate the parameters in their model, but only from time to time they reconsider the choice of the model specification. This, though, would effectively split the analysis in two parts: one of learning dynamics with a fixed α (see section 3.1); the other of the dynamics of α with fixed $\bar{\mathrm{a}}$ and $\bar{\mathrm{b}}$ (see section 4).

In order to analyze real time learning dynamics, we substitute the two differential equations (10) (11) with the following stochastic difference equations, representing recursive least squares learning in real time:

$$\begin{aligned} a_{t+1} &= a_t + t^{-1} R_t^{-1} w_t (y_t - a_t w_t) \qquad (18) \\ R_t &= R_{t-1} + t^{-1}(w_t^2 - R_{t-1}) \\ b_{t+1} &= b_t + t^{-1} S_t^{-1} y_{t-1}(y_t - b_t y_{t-1}) \qquad (19) \\ S_t &= S_{t-1} + t^{-1}(y_{t-1}^2 - S_{t-1}). \end{aligned}$$

The E-stability principle assures that the limiting behavior of these stochastic difference equations is well approximated by the aforementioned differential equations, though only in the limit the two systems are equivalent. We are here interested in investigating how things change over the learning path, when adaptive learning about the parameters interacts with the evolutionary learning about the model.

At each time t, agents now form their expectations according to

$$E_t^f y_{t+1} = a_t \rho w_t \qquad (20)$$

$$E_t^c y_{t+1} = b_t^2 y_{t-1} \qquad (21)$$

and given these expectations and the current value of α_t, y_t can be obtained. Then, on the basis of the new information available, the values for a, b and α are updated according to (18), (19) and (15).

Simulations show that, in the parameter region where the equilibrium is E-stable and α converges to a fixed point (e.g., $A = 1.5$; $B = .1$; $\beta = 10$; $\rho = .5$), the interactions between real time heterogeneous learning and predictor choice dynamics do not prevent convergence of all processes to their equilibrium values. Moreover, convergence of the learning algorithms is quite fast for both models, with no clear advantage of one over the other. Where the equilibrium is E-unstable, instead, parameter estimates (and MSEs) of both models quickly diverge and the system collapses. These simulations, thus, confirm our results above that when endogenous predictor choice dynamics are coupled with learning dynamics, only situations in which fundamentalists constantly outperform chartists can represent long-run equilibria for the system, as they are the only E-stable equilibria.

6 Conclusions

In this paper we have analyzed a simple univariate forward-looking model populated by agents that are heterogeneous in terms of their forecasting models. Following a growing literature in economics and finance, we have allowed for two types of agents, fundamentalists and chartists, whose relative group size is determined endogenously on the basis of the predictive power of their model. While fundamentalists form their expectations using a model that correctly recognizes the role of a fundamental exogenous driving process, chartists employ a model that only uses past prices to predict future ones.

Parameters in each model are chosen by agents so as to minimize the expected mean squared errors of the ensuing predictions, and if these values are not known, they are learned using econometrics techniques such as RLS. A main finding of this paper is that fundamentalists do not necessarily outperform chartists, and are not always able to drive them out of the market permanently. This means that intrinsic heterogeneity can persist even when the correct model is available for selection, and without the need to introduce ad-hoc costs of information gathering or processing on it. Though this result has been derived in a simple setting, it could help explaining why we continue to observe chartist behavior in financial markets. Interestingly, though, it seems that equilibria where chartists constantly outperform fundamentalists are not learnable by adaptive learners.

Another interesting finding is that, when the equilibrium is E-stable, coupled dynamics of adaptive learning and evolutionary predictor choice can simultaneously converge to their equilibrium values and are not displaced by their interactions. Agents are therefore able to learn simultaneously in two dimensions: "within" a model, and "across" models.

7 Appendix

7.1 Appendix A

Denoting

$$\begin{aligned}\Gamma(a) &= A\alpha a\rho + B \\ \Omega(b) &= A(1-\alpha)b^2\end{aligned}$$

we can write recursively

$$\begin{aligned}
\sigma_w^2 &= Ew_t^2 = \frac{1}{1-\rho^2}\sigma_\varepsilon^2 \\
\sigma_{w,w_i}^2 &= Ew_t w_{t-i} = \rho^i \sigma_w^2 \\
\sigma_{w,y}^2 &= Ew_t y_t = \frac{\Gamma}{1-\Omega\rho}\sigma_w^2 \\
\sigma_{w,y_1}^2 &= Ew_t y_{t-1} = \Gamma\rho\sigma_w^2 \sum_{i=0}^{\infty}(\Omega\rho)^i = \frac{\Gamma\rho}{1-\Omega\rho}\sigma_w^2 \\
\sigma_y^2 &= Ey_t^2 = \left(\frac{\Gamma^2(1+\Omega\rho)}{(1-\Omega\rho)\left(1-\Omega^2\right)}\right)\sigma_w^2 \\
\sigma_{y,y_1}^2 &= Ey_t y_{t-1} = \Gamma\sigma_{w,y_1}^2 + \Omega\sigma_y^2.
\end{aligned}$$

The expression for σ_{w,y_1}^2 of course holds only for $\Omega\rho < 1$, which is a restriction we impose on the model in order to have finite covariance structures in our economy.

Moreover:

$$\begin{aligned}
MSE^f &= E(y_t - aw_t)^2 = \sigma_y^2 + a^2\sigma_w^2 - 2a\sigma_{w,y}^2 \\
MSE^c &= E(y_t - by_{t-1})^2 = \sigma_y^2 + b^2\sigma_y^2 - 2b\sigma_{y,y_1}^2
\end{aligned}$$

7.2 Appendix B

The Jacobian for analyzing E-stability of equilibria is

$$J = \begin{bmatrix} \frac{\delta\dot{a}}{a} & \frac{\delta\dot{a}}{b} \\ \frac{\delta\dot{b}}{a} & \frac{\delta\dot{b}}{b} \end{bmatrix}$$

with

$$\begin{aligned}
\frac{\delta\dot{a}}{a} &= \frac{A\alpha\rho}{1-A(1-\alpha)\rho b^2} - 1 \\
\frac{\delta\dot{a}}{b} &= \frac{2A(1-\alpha)\rho b\left(A\alpha\rho a + B\right)}{\left(1-A(1-\alpha\rho b^2)\right)^2} \\
\frac{\delta\dot{b}}{a} &= 0 \\
\frac{\delta\dot{b}}{b} &= 2bA(1-\alpha) + \frac{-4\rho\left[A(1-\alpha)\right]^2 b^3 + 6\rho^2\left[A(1-\alpha)\right]^3 b^5 - 2\rho^2 A(1-\alpha)b}{\left(1+\rho A(1-\alpha)b^2\right)^2} - 1
\end{aligned}$$

8 Figures

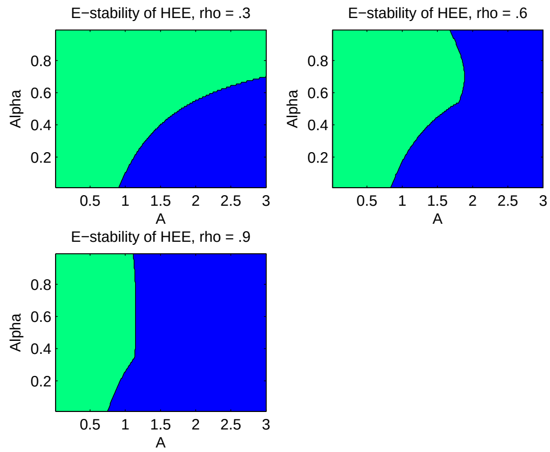

Figure 1: Countour map for the E-stability region (green region).

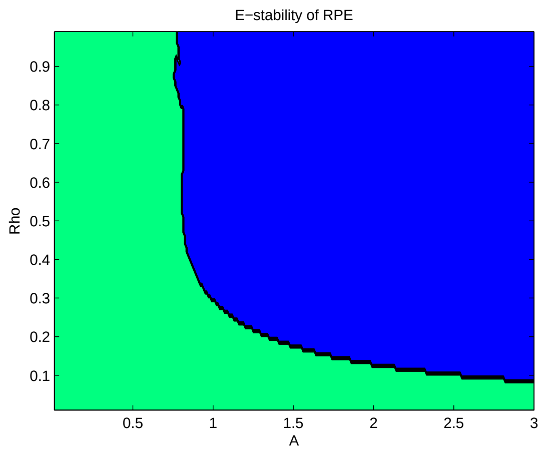

Figure 2: Contour map for the E-stability region for the RPE emerging when $\alpha = 0$.

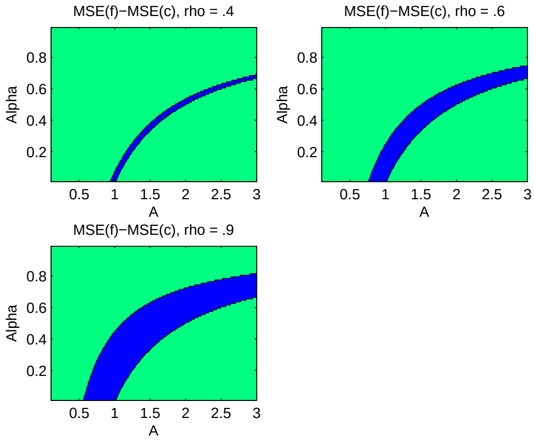

Figure 3: Relative performance of chartists and fundamentalists. Blue regions indicate chartists outperforming fundamentalists.

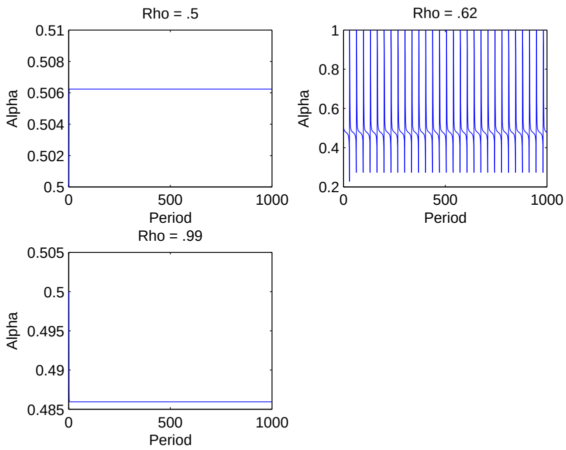

Figure 4: Endogenous dynamics for α.

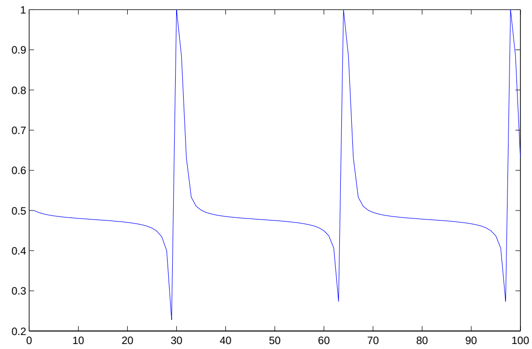

Figure 5: Cyclic behaviour of α emerging for $A = 1.5, \rho = .62$.

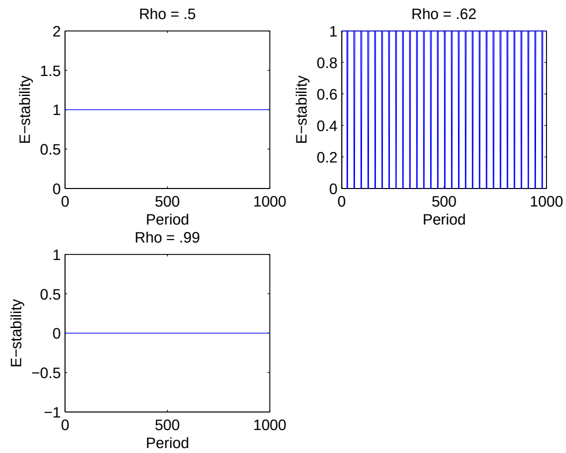

Figure 6: E-stability of equilibria shown in Figure 4.

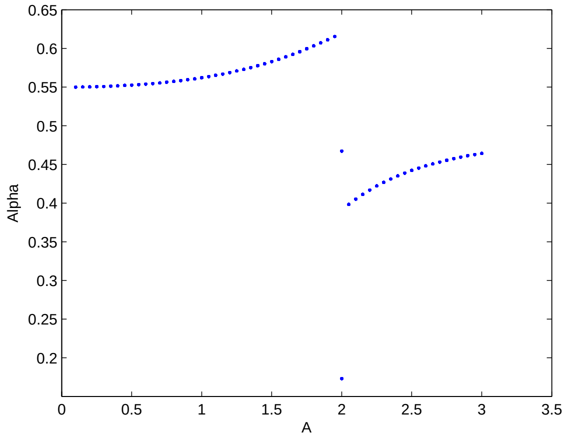

Figure 7: Endogenous α with $\rho = 0$.

References

[1] Benveniste, A., Metivier, M., Priouret, P., 1990. Adaptive Algorithms and Stochastic Approximations. Springer, Berlin.

[2] Berardi, M., 2007. Heterogeneity and misspecifications in learning. Journal of Economic Dynamics and Control 31, 3203-3227.

[3] Branch, W.A., Evans, G.W., 2007. Model uncertainty and endogenous volatility. Review of Economic Dynamics 10, 207-237.

[4] Branch, W.A., Evans, G.W., 2006. Intrinsic heterogeneity in expectation formation. Journal of Economic Theory 127, 264–295.

[5] Brock W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica 65, 1059-1096.

[6] Chiarella, C., 1992. The dynamics of speculative behavior. Annals of Operations Research, 37, 101-123.

[7] Evans, G.W., Honkapohja S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, New Jersey.

[8] Evans, G.W., Honkapohja S., 1996. Least square learning with heterogeneous expectations. Economics Letters 53, 197-201.

[9] Ferrero, G., 2007. Monetary policy, learning and the speed of convergence. Journal of Economic Dynamics and Control 31, 3006-30041.

[10] De Grauwe, P., Dewachter, H., Embrechts, M., 1993. Exchange rate theory. Chaotic models of foreign exchange markets. Blackwell, Oxford.

[11] Goeree, J.K., Hommes, C.H., 2000. Heterogeneous beliefs and the non-linear cobweb model. Journal of Economic Dynamics and Control 24, 761-798.

[12] Guse, E.A., 2007. Heterogeneous expectations, adaptive learning and evolutionary dynamics. University of Cambridge Working Paper 0547, October 2005, revised version 2007.

[13] Guse, E.A., 2005. Stability properties for learning with heterogeneous expectations and multiple equilibria. Journal of Economic Dynamics and Control 29, 1623-1642.

[14] Hommes, C.H., 2007. Bounded rationality and learning in complex markets. CeNDEF working paper 07-01.

[15] Hommes, C.H., 2000. Cobweb dynamics under bounded rationality. In: Dockner E.J., Hartl, R.F., Luptacik, M., Sorger, G. (Editors), Optimization, Dynamics, and Economic Analysis: Essays in Honor of Gustav Feichtinger. Physica-Verlag, 134-150.

[16] Hommes, C.H., 1998. On the consistency of backward-looking expectations: The case of the cobweb. Journal of Economic Behavior and Organization 33, 333-362.

[17] Hommes, C.H., 1994. Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. Journal of Economic Behavior and Organization 24, 315-335.

[18] Hommes, C.H., Sorger, G., 1998. Consistent expectations equilibria. Macroeconomic Dynamics 2, 287-321.

[19] Lasselle, L., Svizzero, S., Tisdell, C., 2005. Stability and cycles in a cobweb model with heterogeneous expectations. Macroeconomic Dynamics 9, 630–650.

[20] Lux, T., 1995. Herd behavior, bubbles and crashes. The Economic Journal, 105, 881-896.

[21] Marcet, A., Sargent, T.J., 1995. Speed of convergence of recursive least squares: learning with autoregressive moving-average perceptions. In: Kirman, A., Salmon, M. (Eds.), Learning and Rationality in Economics. Basil Blackwell, Oxford, pp. 179–215.

[22] Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Clarendon Press, Oxford.

[23] Strogatz, S.H., 1994. Nonlinear dynamics and chaos. Perseus Books, Cambridge.

[24] Sethi, R., 1996. Endogenous regime switching in speculative markets. Structural Change and Economic Dynamics, 99, 99-118.