ASSET PRICING WITH ADAPTIVE LEARNING[*]

EVA CARCELES-POVEDA[†]   CHRYSSI GIANNITSAROU[‡]

January 16, 2006

ABSTRACT.   We study the extent to which self-referential adaptive learning can explain stylized asset pricing facts in a general equilibrium framework. In particular, we analyze the effects of recursive least squares and constant gain algorithms in a production economy and a Lucas type endowment economy. We find that recursive least squares learning has almost no effects on asset price behavior, for either model, since the algorithm converges fast to rational expectations. At the other end, constant gain learning may sometimes contribute towards explaining the stock price volatility and the predictability of excess returns in the endowment economy. However, in the production economy the effects of constant gain learning are mitigated by the persistence induced by capital accumulation. We conclude that, contrary to popular belief, standard self-referential learning alone cannot resolve the asset pricing puzzles observed in the data.

KEYWORDS: Asset pricing, adaptive learning, excess returns, predictability.
JEL CLASSIFICATION: G12, D83, D84

1. INTRODUCTION

It has often been argued informally that adaptive learning should be able to generate statistics that can match stylized facts, in models where the traditional rational expectations paradigm fails. The aim of the present paper is to examine to what extent this assertion is true for asset pricing facts in a general equilibrium framework. To do this, we incorporate two popular adaptive learning algorithms, namely recursive least squares and constant gain, into two workhorse asset pricing models. The first is a production economy that mimics the behavior of the stochastic growth model. The second is an endowment economy of which the reduced form resembles the standard Lucas Tree model. We deliberately restrict attention to standard modeling frameworks and learning algorithms. In this way, we are able to isolate the pure effects of standard self-referential adaptive learning and examine whether such a departure from rational expectations can help explain stylized facts on asset returns.

The predictions of the two models with learning are compared to the data along several dimensions, including the first and second asset return moments, the predictability of future excess returns, the volatility of stock prices and the behavior of the price dividend ratio. Using

[†]SUNY Stony Brook. *E-mail:* ecarcelespov@notes.cc.sunysb.edu.
[‡]University of Cambridge and CEPR. *E-mail*: cg349@cam.ac.uk.

1

a standard parameterization for the production economy, we find that adaptive learning generates almost no improvements for the statistics and numbers we are interested in. As in the fully rational model, the model with learning performs very poorly with respect to asset price behavior. Moreover, both recursive least squares and constant gain learning have a relatively moderate effect on the first and second return moments in the Lucas type economy. However, constant gain learning can generate the predictability of future excess returns that we observe in the data. In addition, it generates considerably higher stock price volatility and approximately matches the behavior of the price dividend ratio.

To get some intuition for these results, consider first the endowment economy. In this case, adaptive learning can generate predictability of future excess returns through the following mechanism. Since the actual law of motion for the stock price is an increasing function of the stochastic dividend payment and of an estimated coefficient, a lower than average coefficient estimate leads to a higher price and a higher price to dividend ratio. In turn, this implies lower future expected returns, generating the negative correlation between the price to dividend ratio and the future returns that is observed in the data. A similar argument can be made for higher than average estimates. Moreover, these effects are reinforced with a higher shock variance and a higher constant gain for the constant gain algorithm. In contrast, the estimated price elasticity with respect to the lag that appears in the law of motion of the price in the production economy has the opposite effect on prices, and it turns out that both effects cancel out irrespective of the learning algorithm. In other words, adaptive learning does not generate any predictability in the presence of capital accumulation.

Regarding the price variability, the elasticity of the price with respect to the shock is constant in the fully rational models, but this may vary under adaptive learning. This implies that learning has the potential of generating additional volatility, an effect that turns out to be positive under constant gain learning and is almost negligible under recursive least squares learning in both models. Finally, we find that the effect of learning on the consumption and return elasticities that determine the equity premium is relatively small across models and algorithms, implying that learning has no potential of generating a sizeable equity premium. In summary, *self-referential adaptive learning is not enough to explain the stylized asset pricing facts that we observe in the data, particularly in models with capital accumulation.*

The literature addressing asset pricing facts is very large and a detailed review of it is beyond the scope of this paper. Kocherlakota (1996), Shiller (1981) and Campbell, Lo and MacKinlay (1997) provide extensive surveys on these topics. Our work is closely related to the part of the literature that attempts to explain asset pricing facts in the context of learning and bounded rationality. This literature includes the work of Timmermann (1994, 1996), Brock and Hommes (1998), Cecchetti, Lam and Mark (2000), Brennan and Xia (2001), Bullard and Duffy (2001) and Honkapohja and Mitra (2003).

The work of Brennan and Xia (2001) focuses on explaining the equity premium puzzle in a

general equilibrium pure exchange economy where non-observability of the exogenous dividend growth process induces extra volatility. Timmermann (1994, 1996) assumes that the exogenous dividend process is unknown and estimated by agents in the context of a present discounted value asset pricing model. As the estimated dividend process is more volatile than the true underlying process in the short run, this type of learning is able to account for some of the excess volatility that we see in the data. A similar mechanism to the one described earlier improves the predictability of stock returns. Brock and Hommes (1998) consider the same present discounted value asset pricing model with heterogenous beliefs and show how chaotic dynamics induce endogenous price fluctuations. Finally, Cecchetti et al. (2000) consider a standard Lucas asset pricing model where agents are assumed to be boundedly rational and have misspecified beliefs.

Our work differs from the previous papers in several important ways. First, we only consider self-referential learning, i.e. learning on the endogenous variable, so that agents' forecasts affect the realization of the variable. In addition, we assume that the steady state is known and that agents' expectations about prices are correctly specified, in the sense that all relevant variables are taken into account when forecasting. We do not allow for learning on the growth rate of dividends, a mechanism that has proven useful for generating stock price volatility and predictability in partial equilibrium models. Apart from the fact that we want to focus on self-referential learning, the reason is that this would involve introducing some type of structural learning in the production economy, where the dividends are endogenous. Given this, our findings can be considered as a lower bound of what adaptive learning can explain, since any additional features can only help to improve our results. In this sense, our work is closest to that of Bullard and Duffy (2001), who study the effects of self-referential recursive least squares learning in the context of a life cycle general equilibrium model. In contrast to this, we study standard asset pricing models with infinitely lived agents. Finally, our work is also closely related to the work of Honkapohja and Mitra (2003), who show that bounded memory adaptive learning can induce extra volatility in the economy. Here, however, we study constant gain learning, which is considered to be a variant of bounded memory adaptive learning, in the context of richer reduced form models.

The paper is organized as follows. Section 1 presents the model economies and section 2 discusses the calculation of the rational expectations and adaptive learning equilibria. Section 3 presents the numerical results and section 4 summarizes and concludes.

2. The Environment

We start by describing two standard general equilibrium asset pricing models. For the first model, which we call the *production economy*, we allow for capital accumulation, so that the model mimics the features of the neoclassical growth model. The second, which we call the *endowment economy*, does not allow for capital accumulation or depreciation of capital. The second model can be viewed as a special case of the first and its log-linear approximation

corresponds to the standard Lucas Tree model.

2.1. The Production Economy. The economy is populated by a large number of identical and infinitely lived households and firms. Each period, the representative household maximizes his expected lifetime utility subject to a sequential budget constraint

$$\max \ E_t \sum_{j=0}^{\infty} \beta^j u(C_{t+j}) \tag{1}$$

s.t.

$$C_t + P_t \Theta_t + P_t^b B_t = (P_t + D_t) \Theta_{t-1} + B_{t-1} + W_t N_t, \tag{2}$$

where

$$u(C) = \begin{cases} \frac{C^{1-\gamma}}{1-\gamma} & \text{if } \gamma > 1 \\ \ln C & \text{if } \gamma = 1 \end{cases}. \tag{3}$$

The parameters $\gamma \geq 1$ and $\beta \in (0,1)$ represent the household risk aversion and time discount factor respectively. The variables Θ_t and B_t are the holdings of equity shares and risk-free one period bonds, P_t and P_t^b represent the equity and bond prices, and D_t represents the equity dividends. The supply of equity is assumed to be constant and is normalized to one, whereas bonds are assumed to be in zero net supply.

Apart from their asset income, households receive labor income, equal to the aggregate wage rate W_t times their labor supply N_t. Investors are endowed with one unit of productive time, which they can allocate to leisure or labor. Given that leisure does not enter the utility function, however, the entire time endowment is allocated to labor and N_t is therefore equal to one. The first order conditions for the household's problem give the usual Euler equations, which determine asset prices

$$P_t = E_t[M_{t,t+1}(P_{t+1} + D_{t+1})], \tag{4}$$

$$P_t^b = E_t[M_{t,t+1}], \tag{5}$$

where $M_{t,t+j} = \beta^j (C_{t+j}/C_t)^{-\gamma}$. Alternatively, we can rewrite the equations in terms of the gross asset returns as

$$1 = E_t[M_{t,t+1} R_{t+1}], \text{ where } R_{t+1} = \frac{D_{t+1} + P_{t+1}}{P_t}, \tag{6}$$

$$1 = E_t[R_{t+1}^f], \text{ where } R_{t+1}^f = \frac{1}{P_t^b}. \tag{7}$$

Each period, the representative firm combines the aggregate capital stock K_{t-1} with the labor input from the households to produce a single good Y_t according to the following constant

returns to scale technology

$$Y_t = Z_t K_{t-1}^{\alpha} N_t^{1-\alpha}, \tag{8}$$

where Z_t is a random productivity shock assumed to follow the stationary process

$$\log Z_t = \rho \log Z_{t-1} + \varepsilon_t, \ \varepsilon_t \sim iid(0, \sigma_\varepsilon^2). \tag{9}$$

Investment I_t is entirely financed by retained earnings or gross profits $X_t = Y_t - W_t N_t$ and the residual of gross profits and investment is paid out as dividends to the firm's owners. Thus, $D_t = X_t - I_t$. Furthermore, capital accumulates according to

$$K_t = I_t + (1 - \delta)K_{t-1}, \tag{10}$$

where $0 < \delta < 1$ is the capital depreciation rate. The representative firm maximizes the value of the firm to its owners, equal to the present discounted value of its nets cash flows or dividends $D_t = X_t - I_t$, subject to (8), (9) and (10)

$$\max \ E_t \sum_{j=0}^{\infty} M_{t,t+j} D_{t+j}. \tag{11}$$

The first-order conditions are

$$W_t = (1 - \alpha)Y_t, \tag{12}$$
$$1 = E_t \left\{ M_{t,t+1} \left[\alpha Z_{t+1} K_t^{\alpha-1} N_{t+1}^{1-\alpha} + (1 - \delta) \right] \right\}. \tag{13}$$

Finally, market clearing implies that

$$Y_t = C_t + K_t - (1 - \delta)K_{t-1}, \tag{14}$$
$$B_t = 0, \ \Theta_t = 1. \tag{15}$$

To derive the system of equations that describe the equilibrium, we substitute for $N_t = 1$, $B_t = 0$, $\Theta_t = 1$ and $W_t = (1 - \alpha)Y_t$. The budget constraint can be omitted, since it is redundant by Walras' Law. Moreover, it can be shown that, in equilibrium, $K_t = P_t$, and we can therefore omit the capital Euler equation. Finally, letting $x_t = \log(X_t / \bar{X})$ for any variable X_t, where $\bar{X}$ represents its steady state value, the original system of equations can be approximated by the

following system of linear equations:

$$z_{t+1} = \rho z_t + \varepsilon_t, \tag{16a}$$

$$y_t = z_t + \alpha k_{t-1}, \tag{16b}$$

$$c_t = \frac{1 - \beta(1-\delta)}{1 - \beta(1-\delta) - \alpha\beta\delta} y_t + \frac{(1-\delta)\alpha\beta}{1 - \beta(1-\delta) - \alpha\beta\delta} k_{t-1} - \frac{\alpha\beta}{1 - \beta(1-\delta) - \alpha\beta\delta} k_t, \tag{16c}$$

$$d_t = \frac{1 - \beta(1-\delta)}{1 - \beta} y_t + \frac{\beta(1-\delta)}{1 - \beta} k_{t-1} - \frac{\beta}{1 - \beta} k_t, \tag{16d}$$

$$p_t = E_t\left[-\gamma(c_{t+1} - c_t) + (1-\beta)d_{t+1} + \beta p_{t+1}\right], \tag{16e}$$

$$p_t^b = E_t\left[-\gamma(c_{t+1} - c_t)\right], \tag{16f}$$

$$k_t = p_t. \tag{16g}$$

2.2. The Endowment Economy. In the endowment economy, the household sector is the same as above. Capital is constant and does not depreciate over time. Therefore, the log-linear system of equilibrium equations can be obtained by setting $k_t = 0$ and $\delta = 0$ in the previous system of equations, resulting in the following log-linear model:

$$z_{t+1} = \rho z_t + \varepsilon_t \tag{17a}$$

$$c_t = d_t = y_t = z_t \tag{17b}$$

$$p_t = E_t\left[-\gamma(d_{t+1} - d_t) + (1-\beta)d_{t+1} + \beta p_{t+1}\right] \tag{17c}$$

$$p_t^b = E_t\left[-\gamma(d_{t+1} - d_t)\right] \tag{17d}$$

This economy can be viewed as an economy where a centralized technology or tree produces a single good Y_t using a constant amount of capital K and the labor supply from the households. Labor is paid its marginal product. Furthermore, households can decide how much labor to supply and how much to invest in the tree and in risk-free one period bonds, while the owners of the tree receive as dividend payments the total output net of labor payments. Moreover, the system of equations in (17a)-(17d) corresponds to the log-linear system of equations of a standard Lucas Tree model with equity and risk free one period bonds, where log-linearized consumption is equal to the log-linearized dividend payments of the tree, whereas the log-linearized dividends follow the same law of motion as the AR(1) process z_t. To see this, note that the equilibrium consumption of a standard Lucas Tree model is given by $C_t = D_t$, and the first-order conditions imply that the asset prices are equal to

$$P_t = \beta E_t \frac{D_{t+1}^{-\gamma}}{D_t^{-\gamma}}(D_{t+1} + P_{t+1}) \tag{18}$$

$$P_t^b = \beta E_t \frac{D_{t+1}^{-\gamma}}{D_t^{-\gamma}}. \tag{19}$$

Moreover, if we assume the AR(1) specification $\log D_t = \rho \log D_{t-1} + \varepsilon_t$, where $\varepsilon_t \sim iid\left(0, \sigma_\varepsilon^2\right)$, the log-linear system of the equations that describes the Lucas model is given by:

$$d_{t+1} = \rho d_t + \varepsilon_{t+1}, \tag{20a}$$

$$c_t = d_t, \tag{20b}$$

$$p_t = \beta E_t p_{t+1} + (1 - \beta - \gamma) E_t d_{t+1} + \gamma d_t \tag{20c}$$

$$p_t^b = E_t\left[-\gamma\left(d_{t+1} - d_t\right)\right]. \tag{20d}$$

3. RATIONAL EXPECTATIONS AND ADAPTIVE LEARNING

In order to calculate the rational expectations equilibria of the two models, we first rewrite the system (16a)-(16g) in reduced form by eliminating all variables but the state variables k_t and z_t in the Euler equation

$$p_t = a_1 E_t p_{t+1} + a_2 p_{t-1} + b z_t, \tag{21}$$

$$z_t = \rho z_{t-1} + \varepsilon_t, \tag{22}$$

where the coefficients a_1, a_2 and b are given by (39a)-(39c) in the appendix. Similarly, the reduced form for the endowment model is given by

$$p_t = a E_t p_{t+1} + b d_t, \tag{23}$$

$$d_t = \rho d_{t-1} + \varepsilon_t, \tag{24}$$

where $a = \beta$ and $b = (1 - \beta - \gamma)\rho + \gamma$.

3.1. Rational Expectations Equilibrium. With the equilibrium conditions in place, we next solve for the rational expectations equilibria of the two models using the method of undetermined coefficients. For the production economy, the (unique stationary) rational expectations equilibrium is given by

$$p_t = \bar{\phi}_p p_{t-1} + \bar{\phi}_z z_{t-1} + \eta_t, \tag{25}$$

where η_t is some white noise shock and[1]

$$\bar{\phi}_p = \frac{1}{2a_1}\left(1 - \sqrt{1 - 4a_1 a_2}\right), \tag{26}$$

$$\bar{\phi}_z = \frac{b}{1 - a_1(\rho + \bar{\phi}_p)}\rho. \tag{27}$$

---

[1] The log-linear system for the production economy has two solutions, corresponding to the so-called minimum state variable (MSV) solutions. Moreover, it is known that this reduced form model is regular, i.e. it has a unique stationary solution, if and only if $|a_1 + a_2| < 1$. In the present model, and given the parameter restrictions, it can be verified that a_1, $a_2 \in (0, 1)$ and that $b > 0$. It can further be shown that $|a_1 + a_2| < 1$. Therefore, the solution with the minus is the unique stationary solution (see Evans and Honkapohja, 2001).

For the endowment economy, the rational expectations equilibrium is given by

$$p_t = \bar{\phi} d_{t-1} + \eta_t, \tag{28}$$

where η_t is a white noise shock and

$$\overline{\phi} = \frac{(1 - \beta - \gamma)\,\rho + \gamma}{1 - \beta\rho}\rho. \tag{29}$$

Note that this solution exists and is finite only under the assumption that $\beta\rho < 1$.

If we compare the two models under rational expectations, the first difference is that the solution for the production economy (25) contains a lag of the price, while the solution of the endowment economy (28) does not. This means that, for an identical parametrization of the exogenous shock, the price series in the production economy has an additional source of persistence due to the lag. Second, it can easily be shown that the elasticity with respect to the shock $\bar{\phi}_z$ in the production economy is smaller than the one in the endowment model for the same parametrization. These observations imply that, under rational expectations, the amount of exogenous volatility fed into the price series of the production economy can be considerably smaller than that in the endowment economy. This is a well-known result which is attributed to the fact that a production economy induces additional consumption smoothing via capital accumulation. Therefore, there is a higher chance of matching the stylized facts of asset prices under rational expectations in the endowment economy. These observations will prove to be useful later on.

3.2. Adaptive Learning. Next, we make a small deviation from rational expectations, by assuming that agents form expectations about future prices based on econometric forecasts. In particular, since we want to keep the economies as close as possible to the rational expectations framework, we make the following assumptions:

A1. Agents know the correct specifications of the models; in other words, they know the steady state and they know which variables are relevant for forecasting prices (no omission or inclusion of extra variables).

A2. Agents know the true parameters that characterize the exogenous shock, i.e. they know ρ and σ_ε^2.

By making these assumptions, we aim in isolating the effects of self-referential learning on the asset pricing statistics and examining if this type of learning alone can provide a better match for the stylized facts. Given these assumptions, agents expectations for both models are formed according to

$$E_t^* p_{t+1} = x_t' \phi_t,$$

where x_t is the vector of state variables, i.e. $x_t = (p_t, z_t)'$ for the production economy and $x_t = d_t$ for the endowment economy. The vector ϕ_t is now an estimate of the true coefficients which is obtained by the recursive algorithm[2]

$$\begin{cases} R_1 = S_0 + x_0 x_0' \\ \phi_1 = \phi_0 + R_1^{-1} x_0 (k_1 - x_0' \phi_0) \end{cases} , \tag{30a}$$

$$\begin{cases} R_t = R_{t-1} + g_t \left(x_{t-1} x_{t-1}' - R_{t-1} \right) \\ \phi_t = \phi_{t-1} + g_t R_t^{-1} x_{t-1} \left(p_t - x_{t-1}' \phi_{t-1} \right) \end{cases} \quad \text{for } t \in \{2, 3, ...\} , \tag{30b}$$

S_0 and ϕ_0 given.

The sequence $\{g_t\}$ is known as the *gain* and represents the weight of the forecasting errors when updating the estimates. We consider two standard and broadly used specifications for the gain, namely $g_t = 1/t$ and $g_t = g$, $0 < g < 1$. The former is a recursive least squares (RLS) algorithm, whereas the latter is known as a tracking or constant gain (CG) algorithm.

A first difference between the two algorithms is that, when written in a non-recursive way, RLS assigns equal weights to all past forecasting errors, while CG assigns weights that decrease geometrically. As a consequence, the RLS can be interpreted as the forecasting method that is used when the econometrician believes that all past information is equally important for forecasting future prices. On the other hand, the CG can be interpreted as the method that is used when the econometrician believes that recent realizations of the stock price are more important in forecasting next period's price.

Another difference between the two algorithms is related to their asymptotic behavior. Since $1/t \to 0$ as $t \to \infty$, the contribution of the forecasting error in the estimate of ϕ under RLS disappears in the limit and the forecasting algorithm eventually converges to the rational expectations equilibrium $\bar{\phi}$. In contrast, the CG algorithm implies that there is always some non-zero correction of the estimate (perpetual learning) which prevents the algorithm from converging to a constant. Instead, the estimate from the CG algorithm converges to some stationary distribution that fluctuates around the rational expectations solution.[3]

4. Stylized Facts

Table 1 presents the stylized asset pricing facts that we will use to compare the different models under rational expectations and adaptive learning. The data set is the one used in Campbell (2002).[4] The quarterly stock market data set is obtained from the nominal CRSP NYSE/AMEX Value Weighted Indices. The aggregate dividend series is extracted from these indices to con-

---

[2]See Carceles-Poveda and Giannitsarou (2005).

[3]Convergence to the rational expectations solution under RLS, for both models, is achieved under certain conditions, which we omit here since these are always satisfied for all reasonable parametrizations. Convergence under CG is achieved for small gains. The details for the derivations of these conditions can be found in Evans and Honkapohja (2001), as well as in Carceles-Poveda and Giannitsarou (2005).

[4]The dataset is available at the author's website.

struct the quarterly dividend and stock return series. Following Campbell (2002), the price dividend ratio is constructed as the stock price index associated with returns excluding dividends, divided by the total dividends paid during the last four quarters. The nominal risk-free rate is the three-month quarterly T-Bill rate. The nominal stock return is deflated using current inflation and the nominal risk-free rate is deflated using the inflation next period. The consumption series corresponds to real per capita consumption of non-durables and services.

< TABLE 1 HERE >

The first part of table 1 reports our estimates for the quarterly mean and standard deviation of stock returns, the risk-free rate and the equity premium in percentage terms. The stock return has been around 2.3 % per quarter against a risk-free rate of 0.2 %, leading to a quarterly premium of around 2 % during the postwar period. We also see a much higher volatility for the equity return and equity premium of around 7.6 %, in contrast to the volatility of around 1 % for the risk-free rate. Replicating the first and second asset moments still represents a challenge for standard rational expectations models.

The second panel of table 1 reports results from regressions of the $k = 1, 2, 4$ year ahead equity premium on the current log price dividend ratio divided by its standard deviation. Thus, the slope coefficients reflect the effect of a one standard deviation change in the log price dividend ratio on the cumulative excess returns in natural units. The table reports the regression slopes, the adjusted R^2 and the t-statistic, adjusted for heteroskedasticity and serial correlation with the Newey-West method.[5] As reflected by the table, the predictive regressions exhibit the familiar pattern of an increasing R^2 and coefficient slope for longer horizons. The fact that the log price dividend ratio predicts future excess returns was first documented by Fama and French (1988) and Campbell and Shiller (1988) and it still poses a puzzle for standard rational expectations models.

Finally, since the price dividend ratio is a crucial variable for addressing the predictability puzzle, the third panel of the table displays its mean, standard deviation and first order autocorrelation in levels. Furthermore, the last panel reports the standard deviation of consumption and dividend growth. It is important to note that these two variables are the same in the endowment economy, but they have a very different behavior in the data. Given this, we will only be able to match one of them with a single calibration.

5. Numerical Results

This section presents the numerical results for the two models under rational expectations and adaptive learning. For each of the two models, we calculate the same statistics as the ones

---

[5]For the truncation lag, we follow Campbell, Lo and MacKinlay (1997), who use $q = 2\,(k-1)$. The results are very similar if we use $q = k - 1$ or the default value of $q = floor\left(4(T/100)^{2/9}\right)$ suggested to Eviews by Newey and West. Similar qualitative results can be obtained by regressing the k-period ahead stock returns on the current log price dividend ratio.

reported in table 1. Additionally, we report the ratio of the standard deviation of the price under learning over the standard deviation under rational expectations, as a proxy for the stock price volatility generated by adaptive learning.

We begin by describing the computing specifications. To implement the simulations we have used the adaptive learning toolbox for *Matlab* that accompanies Carceles-Poveda and Giannitsarou (2005). For each model, we run experiments with a number of $T = 211$ periods, corresponding to the number of quarters available from our data set. Furthermore, the statistics reported are the average statistics from replicating the experiments $N = 1000$ times. To make all results comparable, shocks are generated from normal distributions with the same state value for the *Matlab* pseudorandom number generator, which was set to 98.

As shown in Carceles-Poveda and Giannitsarou (2005), the initialization of adaptive learning algorithms can have important effects on the model dynamics. We therefore use two different initializations. In the first, the initial elasticity ϕ_0 is drawn from a distribution around the rational expectations equilibrium $\bar{\phi}$, with a variance which approximates the variance of an OLS estimator of ϕ based on five observations. In the second, ϕ_0 is set at a value that is below, above and exactly at the rational expectations value. These three values correspond to different initial priors of the households about the effects of the state variables on the current stock price.

Finally, for each set of experiments, we simulate series under RLS learning and CG learning using gain coefficients of $g = 0.02$, $g = 0.2$ and $g = 0.4$. Our choice of these gain values is based on the interpretation of CG learning. As explained earlier, the CG algorithm assigns geometrically decreasing weights to observations across time, so that recent observations matter a lot for the current estimate, even in the limit. In this sense, we can interpret the constant gain algorithm as the tool of an econometrician that believes that recent observations are more relevant for forecasting than observations that date very far back. Specifically, an observation that dates i periods back is assigned a weight equal to $(1 - g)^{i-1}$.

The size of the gain g corresponding to a weight of approximately zero for observations that date more than i years back is displayed in table 2.[6] For example, if the econometrician believes that only observations that date $i = 15$ years back are important for the forecast, the corresponding gain is $g = 0.46$, or if $i = 20$ years, then $g = 0.37$. Since professional forecasters typically use rather short and recent data series from the stock markets, we believe that a constant gain learning (with a relatively high gain coefficient) is a more appropriate modeling framework for asset pricing forecasting. Given this, we have calculated our results with gain values of 0.2 and 0.4, corresponding approximately to using data from the last 20 to 50 years. Furthermore, to get a sense of how our results depend on the size of the gain, we have also calculated the results with a gain of $g = 0.02$, corresponding to approximately using data from the last 400 years to make the forecasts.

< TABLE 2 HERE >

---

[6] To calculate the gains, we have used the default tolerance level of *Matlab*, as an approximation of zero.

We first present the results for the endowment economy and then discuss the results for the production economy.

5.1. The Endowment Economy. In the endowment economy, the risk aversion value is set to $\gamma = 1$. Further, since use a quarterly time period, we set $\beta = 0.99$. As for the dividend process, the benchmark calibration assumes that $\rho = 0.95$ and $\sigma_\varepsilon = 0.06$, corresponding to the estimated slope coefficient and error standard deviation of regressing the log of the seasonally adjusted real dividend series in the data on its first lag. We also repeat the experiments with $\rho = 0.95$ and $\sigma_\varepsilon = 0.00712$ in order to make the findings comparable to those from the production economy. Moreover, this last calibration approximately replicates the behavior of logged consumption growth in the data.

Tables 3A-3C contain the results for the calibration with the lower variance, whereas tables 3D-3F report the same results for the higher shock variance. Tables 3A and 3D contain the first and second asset moments. Tables 3B and 3E contain (a) the standard deviation of the stock price under learning over the standard deviation of the stock price under rational expectations, (b) the average price dividend ratio, its standard deviation and its first autocorrelation, (c) the standard deviation of consumption growth and (d) the standard deviation of dividend growth. Finally, tables 3C and 3F report the results for predictability. To obtain these, we run the same regressions as with the true data. The table reports the average estimated slope coefficients, the average adjusted R^2 and the percentage of estimated coefficients that are negative and significant out of 1000 replications of the experiment.

The first two rows of the tables display the numbers in the data and under rational expectations (RE). Furthermore, the last six rows display the results under learning when the algorithms are initialized (a) from a distribution (DIS), (b) below the REE, with an elasticity set to $0.9 \times \bar{\phi}$ (AH-B), or (c) above the REE, with an elasticity set to $1.035 \times \bar{\phi}$ (AH-A).[7] For each initialization, we report the results for the recursive least squares (RLS) and constant gain (CG) algorithms with gains of $g = 0.2$ and $g = 0.4$. The case with $g = 0.02$ is omitted, since the results are almost identical to the ones under RLS.

<center>< TABLES 3A - 3F HERE></center>

Starting with the results under RE, we see that the model performs very poorly in all dimensions. With the lower shock variance, the premium is only around 0.002 percent, while it only increases to approximately 0.2 percent with the benchmark parameterization. Furthermore, whereas the standard deviation and the autocorrelation of the price dividend ratio are far from the data, this variable generates absolutely no predictability for the excess stock returns. This

---

[7]The percentage 1.035 above the REE has been chosen for both models, so that the stationarity condition $|\phi_0| < 1$ is satisfied for the production economy. Although such a restriction is not necessary for the endowment economy, we use the same number to keep the results comparable. Furthermore, the case where the initial elasticity starts at the rational expectations value $\bar{\phi}$ has been omitted, since it generates the same results on average as when we initialize from a distribution.

is not surprising, since it is well documented in the literature that the Lucas tree model with a low risk aversion parameter value is unsuccessful in reproducing the asset pricing moments in the data.

Turning to the results under adaptive learning, the first important observation is that they are almost identical across the different initializations. In other words, different initial priors for the elasticity of the stock price with respect to the stochastic dividend process do not alter the results in the endowment economy. On the other hand, we do observe important differences across the different learning algorithms and parametrizations. We discuss the results with each algorithm in turn.

Starting with RLS learning, we see that the asset return moments are very close to those generated by rational expectations. In general, the reason why RLS cannot generate any significant improvements in the predictions of the model is that the algorithm converges relatively fast to the rational expectations equilibrium. Therefore, any differences between the dynamics under RLS and rational expectations disappear quickly. This is also reflected in tables 3B and 3E, illustrating that the standard deviation of the price under RLS learning over the one under rational expectations is approximately one with both parametrizations. Finally, Tables 3C and 3F illustrate that the model under RLS also performs very poorly regarding predictability. Whereas the coefficients have the right sign and are higher in absolute value than the ones under rational expectations, they are still far from the data under both calibrations. In addition, the percentage of significant simulations is relative small.

It should be pointed out that if we had also assumed that agents estimate the exogenous dividend process, some additional variation between RLS and rational expectations would be present. However, we have chosen not to allow learning on the exogenous dividend process for several reasons. First, since the dividend process is endogenous in the stochastic growth model, this would involve introducing some form of structural learning, while we want to focus on self-referential learning. Moreover, since the least squares estimates for the dividend process would be reached quite fast, we conjecture that this addition would not provide any significant variations, unless the assumed time horizon was very short.

With this discussion in mind, it should be clear that any improvements in the predictions of the model can only come from some type of learning algorithm that does not converge to the rational expectations equilibrium. Constant gain learning is such an algorithm, since its dynamics fluctuate perpetually around the rational expectations equilibrium and the size of the fluctuations depends positively on the size of the gain function. Indeed, turning to the results generated by CG learning, the results appear to be quite different from those under RLS and rational expectations.

Whereas the improvements regarding the asset return moments are again relatively modest, we see a considerable improvement regarding the stock price volatility and predictability. Tables 3B and 3E reflect that the asset price under CG learning can be significantly more volatile than

under rational expectations. In addition, with the benchmark parameterization and the higher gain, the model matches the standard deviation of the logged dividend growth, whereas the standard deviation of the price-dividend ratio is just about half of the one observed in the data. Finally, tables 3C and 3F reflect that the model performs much better than under rational expectations regarding predictability. As we see, the average slope coefficients, the percentage of significant and negative estimates and the R^2 display the increasing pattern with a longer horizon that we see in the data. Furthermore, the slope coefficients are surprisingly close to the ones in the data when the model is calibrated to dividend behavior ($\sigma = 0.06$) and the gain is equal to 0.4. In this case, the number of significant estimates in this case ranges from approximately 40% to 70%, a large improvement compared to the results under RLS and rational expectations.

These improvements are smaller when the model is calibrated to consumption behavior ($\sigma = 0.00712$). In this case, the model does not generate the right behavior for the dividend growth or the price dividend ratio. Further, although this calibration leads to a higher stock price volatility and generates a much higher predictability under learning than the rational expectations version, the slope coefficients with a gain of 0.4 do not provide a satisfactory match with the data.

Summarizing, RLS learning generates results that are very close to their rational expectations counterpart. On the other hand, CG generates some improvements when the model is calibrated to dividend behavior, particularly with respect to the stock price volatility and the predictability of excess returns. In what follows, we provide some intuitive comments to help understand our findings.

Consider first the volatility of the stock price. It is easy to see that the variance of the price under rational expectations is only affected by the variance of the shock, since it is equal to

$$Var(p_t^{re}) = V(\overline{\phi})^2 Var(z_t) = Var(z_t). \tag{31}$$

On the other hand, under adaptive learning, the behavior of the estimated coefficient ϕ also affects the price variance, which is given by

$$Var(p_t^{al}) = a^2 Var\left(\phi_{t-1}z_t\right) + b^2 Var\left(z_t\right) + 2ab\, Cov(\phi_{t-1}z_t). \tag{32}$$

The variance in (32) will be higher than variance of the shock in (31) if $Var\left(\phi_{t-1}z_t\right) > Var\left(\overline{\phi}z_t\right)$, whereas the opposite will happen if $Var\left(\phi_{t-1}z_t\right) < Var\left(\overline{\phi}z_t\right)$. Our numerical results show that $Var\left(\phi_{t-1}z_t\right)$ is on average very similar to $Var\left(\overline{\phi}z_t\right)$ under RLS, explaining why the two prices have almost the same variability. For the case of CG learning, however, the coefficients deviate more from the REE solution, and $Var\left(\phi_{t-1}z_t\right)$ is on average higher than $Var\left(\overline{\phi}z_t\right)$, explaining the higher price variability. Finally, since $Var\left(\phi_{t-1}z_t\right)$ is increasing with the variability of ϕ_{t-1}, which increases with a higher gain, this would also explain the higher

variance of the price with higher gain coefficients.

While learning may generate some extra volatility with a high gain, note that the mean equity premium is very close to its value under rational expectations. To see why this is the case, note that the expected equity premium under rational expectations is given by[8]

$$r_{t+1}^{rp} = \gamma \eta_{cz} \eta_{r_z^e} \, Var\left(z_t\right), \tag{33}$$

where η_{cz} and $\eta_{r_z^e}$ represent the elasticities of the unexpected consumption growth and the unexpected equity return with respect to the shock. Moreover, these elasticities depend on the deep parameters of the model and on the underlying elasticity $\bar{\phi}$. First, similar arguments to the ones for volatility can be used to explain why more volatile estimated coefficients can generate some extra excess return volatility through their effects on the two elasticities η_{cz} and $\eta_{r_z^e}$. However, since the estimated coefficients are on average close to their REE values, the consumption and equity return elasticities will not be very far from the REE solution. We therefore conclude that, unless some form of misspecification is introduced, standard adaptive learning does not have a lot of potential for explaining mean returns in the present model.

Finally, to explain why self-referential learning can generate excess return predictability, note that the actual law of motion for the price under learning is given by

$$p_t = T\left(\phi_{t-1}\right) z_{t-1} + V\left(\phi_{t-1}\right) \varepsilon_t = V\left(\phi_{t-1}\right) z_t, \tag{34}$$

where

$$
\begin{aligned}
T\left(\phi_{t-1}\right) &= \left(a\phi_{t-1} + b\right)\rho, \\
V\left(\phi_{t-1}\right) &= \left(a\phi_{t-1} + b\right).
\end{aligned} \tag{35}
$$

Consider an estimate of ϕ that is below the "true" or rational expectations value. Since $V\left(\phi_{t-1}\right)$ is increasing in ϕ_{t-1}, the stock price will be low and this will generate a low price to dividend ratio (or a high dividend to price ratio) which, in turn, will generate higher future returns. This is because, first, a higher dividend to price ratio implies higher payoffs and, second, because when the price re-adjusts upwards, it will lead to higher capital gains. A similar argument would indicate that an estimate of ϕ that is above the true value will lead to a higher than average price to dividend ratio and to lower future returns. This indicates that learning has the potential of generating a negative correlation between the current price to dividend ratio and the future excess returns. Finally, it is worth noting that ϕ and $V\left(\phi\right)$ are constant under rational expectations. This implies that the effects we have just described are not present in the fully rational setting, providing a possible explanation for its poor performance concerning predictability.

---

[8]See Lettau (2003) or Carceles-Poveda (2005) for details.

5.2. The Production Economy. For the production economy, we have used the standard parametrization for US quarterly data. The risk aversion coefficient is again set to $\gamma = 1$. Furthermore, the capital depreciation, the discount factor and the capital share are set to $\delta = 0.025$, $\beta = 0.99$ and $\alpha = 0.36$ respectively. Finally, regarding the productivity shock, our baseline parametrization is $\sigma_\varepsilon = 0.00712$ and $\rho = 0.95$, as is usual in the real business cycle literature.

Tables 4A-4C report the results for the production economy, organized in the same way as the results for the endowment economy. In particular, table 4A contains asset moments, table 4B contains various statistics and table 4C reports the results for predictability.

<center>< TABLES 4A - 3C HERE></center>

As with the endowment economy, the tables indicate that the rational production economy performs very poorly in explaining the first and second asset moments.[9] The implied equity premium is approximately 0.005 percent, whereas the asset variabilities are very similar across the two assets and very far from their counterparts in the data. Furthermore, the standard deviation of the price dividend ratio is much lower than the one in the data, and it does not have any predictive power for the excess stock returns.

Turning to the results under learning, we see that it has a very small effect on the different asset moments. In particular, the equity premium only increases from 0.005 to 0.01 percent and its variability only increases from 0.01 to 0.025 percent. In addition, constant gain learning does induce stock price volatility, but this seems to be non-monotonic with respect to the gain coefficient. Finally, although learning improves the behavior of the price dividend ratio, the average regression coefficients are practically zero and rarely significant for all the horizons considered, as can be seen from table 4C. These findings suggest that, with a standard calibration, adaptive learning does not seem to provide an explanation for the behavior of asset returns in the production economy.

An interesting observation is that the results depend on the different initializations of the learning algorithms, unlike in the endowment economy. In particular, starting above the rational expectations value generates a premium that is ten times higher than the premium if we start below rational expectations. This is also reflected in the relative variability of the stock price. As can be seen in Table 4B, the variability of the stock price can be considerably lower than under rational expectations if the initial coefficient is set below its rational expectations value, while it is always higher if we start above. Thus, contrary to the common view that learning can only generate higher volatility, we find that the size of the volatility actually depends on the initialization of the algorithm in the stochastic growth model and may very well be below the one generated by rational expectations.[10]

---

[9]See for example Rouwenhorst (1995), or Lettau (2003).

[10]An extensive discussion of the effects of different initializations of learning can be found in Carceles-Poveda and Giannitsarou (2005).

In what follows, we discuss intuitively why adaptive learnings fails to improve the predictions of the model in the production economy. First, regarding the volatility of asset prices, it can be shown that it is equal to the following expression under rational expectations:[11]

$$Var(p_t^{ree}) = \frac{\gamma^2(1 + \rho\bar{\phi}_p)}{(1 - a_1\rho - a_1\bar{\phi}_p)^2(1 - \rho\bar{\phi}_p)(1 - \bar{\phi}_p^2)} \tag{36}$$

Moreover, for our calibration, it is possible to show that the previous expression is increasing in $\bar{\phi}_p$. Using this result, we can then heuristically argue that the variance of the endogenous state under learning will be lower than the variance under rational expectations if the estimated coefficients ϕ_p remain well below $\bar{\phi}_p$ for a large number of periods.

If we initialize the algorithm using a distribution, RLS implies that these coefficients will be relatively close to the REE, leading to a very similar variance between the two cases. On the other hand, a gain of 0.2 makes them more volatile, so that they are more often above the REE. Finally, increasing the gain coefficient above 0.2 also implies that the coefficients violate the stationarity condition requiring that $\left|\phi_{p,t-1}\right| < 1$, since the REE value $\bar{\phi}_p$ is very close to one. Since it would not be sensible to allow the elasticity to be larger than one, the learning algorithm is augmented with a projection facility, which simply resets ϕ to its last value when this condition is violated. In this case, the projection facility generates a downward bias that reduces the price volatility when increasing the gain from 0.2 to 0.4, explaining the non-monotonicity in the results that is seen in the tables.[12]

On the other hand, if we initialize the algorithm above or below the REE, the estimated coefficients will remain on average above or below the REE respectively. This happens because of persistence, which due to the presence of a lag in the law of motion of the stock price. This explains why the different initializations generate different results in the production economy, while they do not matter in the endowment economy.

Finally, turning to the predictability results, note that the actual law of motion for the price in the production economy is given by

$$p_t = T(\phi_{t-1})p_{t-1} + V(\phi_{t-1})z_t, \tag{37}$$

where

$$
\begin{aligned}
T(\phi) &= \frac{a_2}{1 - a_1\phi_p}, \\
V(\phi) &= \frac{a_1\phi_z + b}{1 - a_1\phi_p}.
\end{aligned}
\tag{38}
$$

---

[11]See Giannitsarou (2005) for a derivation.

[12]For details on the projection facility and how it affects the behavior of the estimates see Carceles-Poveda and Giannitsarou (2005). Note that we do not need to impose it in the endowment economy, where the equilibrium is globally stable.

Comparing to the case of the endowment economy, a lower estimate of $\phi_{z,t-1}$ does not directly translate into a lower than average price level anymore, since the T and V maps are also affected by $\phi_{p,t-1}$. In particular, our numerical results suggest that these two coefficients move on average in opposite directions, whereas their variances are similar. Thus, while a lower estimate of ϕ_z leads to a lower $V(\phi_{t-1})$ and therefore to a lower price, a higher estimate of ϕ_p has the opposite effect, since it increases $V(\phi_{t-1})$ and $T(\phi_{t-1})$. Moreover, as noted earlier, the elasticity of the price with respect to the shock is much smaller than that in the endowment economy. Based on these arguments, these findings suggest that any effects learning might induce may cancel out for all algorithms when the variance of the shock is relatively small.

Yet, a higher shock variance could generate some better results in this respect, since it will lead to a higher feedback of the estimate of $\phi_{z,t-1}$ on the price through a higher second term $V(\phi_{t-1})z_t$. To get a sense of how big such an effect would be, we have also simulated the production economy with a higher shock variance.[13] In this case, we find that a higher shock variance generates some predictability, especially with the initialization above the REE. In particular, the estimated slope coefficients in this case can be around one half of their value in the data. The number of significant simulations, however, is still considerably smaller than in the endowment economy. In addition, we find that the model can generate a higher premium and return volatility than under rational expectations, specially with the initialization above the rational expectations equilibrium. As before, this can be explained by the persistence induced by the lag in the law of motion of the price, which implies that the elasticities that affect the equity premium are further away from their rational expectations value than in the endowment economy. Of course, such improvements come at the expense of unrealistically high values for the moments of the price dividend ratio and of the real macroeconomic variables. Given this, we conclude that learning does not provide a satisfactory explanation for the asset pricing puzzles in the presence of capital accumulation.

6. Conclusion

We studied the effects of self-referential adaptive learning on asset returns in the framework of standard general equilibrium asset pricing models. In particular, we have considered recursive least squares and constant gain learning, with a variety of specifications, in a production economy and a Lucas type exchange economy. Both models were evaluated with respect to the first and second equity premium moments, the predictability of excess returns and the volatility of stock prices. The main conclusions from our results are that (a) constant gain adaptive learning has a chance of generating stock price volatility and predictability in the endowment economy, when the gain coefficient is relatively high, (b) constant gain learning does not generate any interesting improvements in the production economy framework and (c) recursive least squares learning does not generate any improvements for any of the two models.

---

[13] To avoid the proliferation of tables, we do not report the results of these simulations but they can be provided by the authors upon request.

In general, standard adaptive learning has little potential of explaining the mean excess returns in the data, since the average estimated coefficients from the law of motion of the stock price fluctuate around the rational expectations equilibrium, which is known to fail in generating a sizeable premium for reasonable parametrizations. As to the stock price volatility and predictability of excess returns, we find important differences across models and across learning algorithms. In particular, recursive least squares learning has relatively small effect on the stock price volatility and it generates no predictability in the production economy and almost no predictability in the endowment economy. The effects of constant gain learning with a relatively small gain are very similar. Nevertheless, a higher gain coefficient, reflecting the fact that forecasters give more importance to recent observations, generates considerably more volatility and predictability in the endowment economy, especially when it is calibrated to match the dividend behavior in the data.

In general, our findings suggest that tracking algorithms such as CG have more potential than RLS to explain asset pricing facts in models where there is no inherent persistence in the stock price, such as the Lucas Tree endowment economy. On the other hand, in the presence of capital accumulation, where the endogenous variables exhibit more persistence and where consumption smoothing plays an important role, adaptive learning is not sufficient to generate any of the stylized facts in the data.

A.   Reduced Form Coefficients

The coefficients of the reduced form for the production economy are

$$a_1 = \frac{-\gamma}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{39a}$$

$$a_2 = \frac{\gamma\left(\delta - 1 - \alpha\psi\right)}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{39b}$$

$$b = \frac{\psi\left(\gamma\left(\rho - 1\right) + \alpha\beta\left(\delta - \psi\right)\rho\right)}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{39c}$$

where $\psi = \left(1 - \beta + \delta\beta\right)/\left(\alpha\beta\right)$.

References

[1] Brennan, M. J. and Y. Xia, 2001. Stock Price Volatility and Equity Premium. *Journal of Monetary Economics*, 47, 249-283.

[2] Brock, W. A. and C. H. Hommes, 1998. Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics and Control*, 22, 1235-1274.

[3] Bullard, J. and J. Duffy, 2001. Learning and Excess Volatility. *Macroeconomic Dynamics*, April 2001, 5, 272-302.

[4] Campbell, J. Y., 2003. Consumption-Based Asset Pricing. In the Handbook of the Economics of Finance, *Elsevier North-Holland.*

[5] Campbell, J. Y., A. W. Lo and A. G. MacKinlay, 1997. The Econometrics of Financial Markets. *Princeton University Press.*

[6] Campbell, J. and R. Shiller, 1988, Stock Prices, Earnings and Expected Dividends, *Journal of Finance* 43(3), 661-676.

[7] Carceles-Poveda, E., 2005, Idiosyncratic Shocks and Asset Returns in the RBC Model, an Approximate Analytical Approach, *Macroeconomic Dynamics*, forthcoming.

[8] Carceles-Poveda, E. and C. Giannitsarou, 2005. Adaptive Learning in Practice. *Mimeograph.*

[9] Cecchetti, S., P. S. Lam and N. C. Mark, 2000. Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to be True?. *American Economic Review*, 90, 787-805.

[10] Evans, G. and S. Honkapohja, 2001. Learning and Expectations in Macroeconomics. *Princeton University Press.*

[11] Fama E. and K. French, 1988, Dividend Yields and Expected Stock Returns, *Journal of Financial Economics* 22(1), 3-27.

[12] Giannitsarou, C., 2005, E-stability Does Not Imply Learnability, *Macroeconomic Dynamics* 9, 276-287.

[13] Honkapohja, S. and K. Mitra, 2003. Learning with Bounded Memory in Stochastic Models. *Journal of Economic Dynamics and Control*, 27, 1437-1457.

[14] Kocherlakota N., 1996, The Equity Premium: It's still a Puzzle, *Journal of Economic Literature*, Vol. 34, No. 1, pp. 42-71.

[15] Lettau M., 2003, Inspecting the Mechanism: The Determination of Asset Prices in the Real Business Cycle Model, *The Economic Journal*, 113, 550–575.

[16] Rouwenhorst K. G., 1995, Asset Prices Implications of Equilibrium Business Cycle Models. In New Frontiers of Modern Business Cycle Research, ed. by T. Cooley, *Princeton University Press.*

[17] Shiller, R. J., 1989. Market Volatility. *MIT Press.*

[18] Timmermann, A., 1994. How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices. *Quarterly Journal of Economics*, 108, 1135-1145.

[19] Timmermann, A., 1996. Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning. *Review of Economic Studies*, 63, 523-557.

Asset Moments		
	Mean	*Std.*
r^e	2.3454	7.7378
r^f	0.2268	0.8719
$r^e - r^f$	2.0341	7.6291
Predictability		
Horizon	*Slope*	R^2	$t-statistic$
1	-0.0533	0.0926	-2.3317
2	-0.1073	0.2044	-2.5135
4	-0.1856	0.3683	-3.2161
Moments for P/D		
	Mean	*Std.*	*Autocor.*
	28.3065	9.0611	0.9654
Moments for Δc and Δd		
	$Std(\Delta d)$	$Std(\Delta c)$
	12.0599	1.0725

Table 1: Asset pricing facts 1947.2-1998.4

$Years$	400	200	100	50	25	20	15	10	5
$Gain$	0.02	0.04	0.09	0.17	0.31	0.37	0.46	0.60	0.85

Table 2: Gains for the CG algorithm

	$\sigma = 0.00712$	Equity Return		Risk Free Rate		Equity Premium	
		Mean	*St. Dev.*	*Mean*	*St. Dev.*	*Mean*	*St. Dev.*
Data		**2.3454**	**7.7378**	**0.2268**	**0.8719**	**2.0341**	**7.6291**
RE		1.0123	0.7303	1.0106	0.1020	0.0017	0.7562
DIS	RLS	1.0122	0.7156	1.0106	0.1020	0.0017	0.7424
	CG, $g = 0.2$	1.0128	0.7629	1.0106	0.1020	0.0024	0.7911
	CG, $g = 0.4$	1.0146	0.8788	1.0106	0.1020	0.0043	0.9067
AH-B	RLS	1.0122	0.7143	1.0106	0.1020	0.0017	0.7459
	CG, $g = 0.2$	1.0128	0.7646	1.0106	0.1020	0.0024	0.7928
	CG, $g = 0.4$	1.0146	0.8810	1.0106	0.1020	0.0043	0.9089
AH-A	RLS	1.0122	0.7143	1.0106	0.1020	0.0017	0.7459
	CG, $g = 0.2$	1.0128	0.7646	1.0106	0.1020	0.0024	0.7928
	CG, $g = 0.4$	1.0146	0.8810	1.0106	0.1020	0.0043	0.9089

Table 3A: Endowment Economy, Statistics for Returns

	$\sigma = 0.00712$	P_{AL} / P_{RE}	Mean(*P/D*)	STD(*P/D*)	Corr(*P/D*)	STD(Δc)	STD(Δd)
Data			**28.3065**	**9.0611**	**0.9654**	**1.0725**	**12.0599**
RE			24.7521	0.1639	0.5395	1.4432	1.4432
DIS	RLS	0.9895	24.7531	0.1944	0.6554	1.4432	1.4432
	CG, $g = 0.2$	1.1516	24.7542	0.2806	0.7909	1.4432	1.4432
	CG, $g = 0.4$	1.4138	24.7631	0.4398	0.8642	1.4432	1.4432
AH-B	RLS	0.9949	24.7542	0.1903	0.6437	1.4432	1.4432
	CG, $g = 0.2$	1.1544	24.7548	0.2805	0.7899	1.4432	1.4432
	CG, $g = 0.4$	1.4180	24.7633	0.4415	0.8642	1.4432	1.4432
AH-A	RLS	0.9949	24.7542	0.1903	0.6437	1.4432	1.4432
	CG, $g = 0.2$	1.1544	24.7548	0.2805	0.7899	1.4432	1.4432
	CG, $g = 0.4$	1.4180	24.7633	0.4415	0.8642	1.4432	1.4432

Table 3B: Endowment Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth

	$\sigma = 0.00712$	BETAS						R-SQUARE		
		1 year		*2 years*		*4 years*		*1 year*	*2 years*	*4 years*
Data		**-0.0533**		**-0.1073**		**-0.1856**		**0.0926**	**0.2044**	**0.3683**
		Aver.	% Sig.	Aver.	% Sig.	Aver.	% Sig.			
RE		-0.0000	7.6	-0.0001	10	-0.0002	14.4	0.0046	0.0056	0.0057
DIS	RLS	-0.0005	7	-0.0009	11.6	-0.0017	17.8	0.0066	0.0119	0.0205
	CG, $g = 0.2$	-0.0234	16.4	-0.0477	31.8	-0.0083	45.2	0.0207	0.0432	0.0722
	CG, $g = 0.4$	-0.0049	44.4	-0.0093	59.2	-0.0152	67.8	0.0561	0.0983	0.1424
AH-B	RLS	-0.0004	6.8	-0.0008	12	-0.0015	16.4	0.0064	0.0106	0.0179
	CG, $g = 0.2$	-0.0233	17.8	-0.0475	31.8	-0.0082	44.8	0.0205	0.0424	0.0713
	CG, $g = 0.4$	-0.0049	43.8	-0.0093	60	-0.0153	68	0.0561	0.0982	0.1424
AH-A	RLS	-0.0004	6.8	-0.0008	12	-0.0015	16.4	0.0064	0.0106	0.0179
	CG, $g = 0.2$	-0.0233	17.8	-0.0475	31.8	-0.0082	44.8	0.0205	0.0424	0.0713
	CG, $g = 0.4$	-0.0049	43.8	-0.0093	60	-0.0153	68	0.0561	0.0982	0.1424

Table 3C: Endowment Economy, Predictability of Excess Returns

	$\sigma = 0.06$	Equity Return		Risk Free Rate		Equity Premium	
		Mean	*St. Dev.*	*Mean*	*St. Dev.*	*Mean*	*St. Dev.*
Data		**2.3454**	**7.7378**	**0.2268**	**0.8719**	**2.0341**	**7.6291**
RE		1.1940	6.1721	1.0185	0.8600	0.1769	6.3894
DIS	RLS	1.1984	6.0497	1.0185	0.8600	0.1813	6.2752
	CG, $g = 0.2$	1.2389	6.4672	1.0185	0.8600	0.2226	6.7052
	CG, $g = 0.4$	1.4244	7.7093	1.0185	0.8600	0.4088	7.8943
AH-B	RLS	1.1974	6.0808	1.0185	0.8600	0.1805	6.3049
	CG, $g = 0.2$	1.2386	6.4812	1.0185	0.8600	0.2224	6.7188
	CG, $g = 0.4$	1.4115	7.6763	1.0185	0.8600	0.3959	7.9095
AH-A	RLS	1.1974	6.0808	1.0185	0.8600	0.1805	6.3049
	CG, $g = 0.2$	1.2386	6.4812	1.0185	0.8600	0.2224	6.7188
	CG, $g = 0.4$	1.4115	7.6763	1.0185	0.8600	0.3959	7.9095

Table 3D: Endowment Economy, Statistics for Returns

	$\sigma = 0.06$	P_{AL} / P_{RE}	Mean(P/D)	STD(P/D)	Corr(P/D)	STD(Δc)	STD(Δd)
Data			**28.3065**	**9.0611**	**0.9654**	**1.0725**	**12.0599**
RE			24.7621	1.3805	0.5390	12.187	12.187
DIS	RLS	0.9895	24.7941	1.6589	0.6566	12.187	12.187
	CG, $g = 0.2$	1.1516	24.8985	2.4223	0.7898	12.187	12.187
	CG, $g = 0.4$	1.4138	24.4047	4.6046	0.8599	12.187	12.187
AH-B	RLS	0.9949	24.7995	1.6084	0.6432	12.187	12.187
	CG, $g = 0.2$	1.1544	24.9028	2.4219	0.7887	12.187	12.187
	CG, $g = 0.4$	1.4180	24.4083	4.6187	0.8528	12.187	12.187
AH-A	RLS	0.9949	24.7995	1.6084	0.6432	12.187	12.187
	CG, $g = 0.2$	1.1544	24.9028	2.4219	0.7887	12.187	12.187
	CG, $g = 0.4$	1.4180	24.4083	4.6187	0.8528	12.187	12.187

Table 3E: Endowment Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth

	$\sigma = 0.06$	BETAS						R-SQUARE		
		1 year		2 years		4 years		1 year	2 years	4 years
Data		-0.0533		-0.1073		-0.1856		0.0926	0.2044	0.3683
		Average	% Sig.	Average	% Sig.	Average	% Sig.			
RE		-0.0001	7.8	-0.0008	10.2	-0.0017	14	0.0045	0.0056	0.0057
DIS	RLS	-0.0041	7.2	-0.0083	11.8	-0.0144	18	0.0066	0.0119	0.0205
	CG, $g = 0.2$	-0.0198	16.6	-0.0403	32.4	-0.0701	45.9	0.0217	0.0430	0.0725
	CG, $g = 0.4$	-0.0443	43.6	-0.0834	60.4	-0.1362	68.2	0.0578	0.1016	0.1449
AH-B	RLS	-0.0035	7	-0.0071	12	-0.0127	17	0.0064	0.0106	0.0179
	CG, $g = 0.2$	-0.0197	18.2	-0.0401	32.4	-0.0700	54.8	0.0215	0.0426	0.0716
	CG, $g = 0.4$	-0.0442	43.4	-0.0831	64.2	-0.1357	68.6	0.0579	0.1006	0.1457
AH-A	RLS	-0.0035	7	-0.0071	12	-0.0127	17	0.0064	0.0106	0.0179
	CG, $g = 0.2$	-0.0197	18.2	-0.0401	32.4	-0.0700	54.8	0.0215	0.0426	0.0716
	CG, $g = 0.4$	-0.0442	43.4	-0.0831	64.2	-0.1357	68.6	0.0579	0.1006	0.1457

Table 3F: Endowment Economy, Predictability of Excess Returns

	$\sigma = 0.00712$	Equity Return		Risk Free Rate		Equity Premium	
		Mean	*St. Dev.*	*Mean*	*St. Dev.*	*Mean*	*St. Dev.*
Data	**Data**	**2.3454**	**7.7378**	**0.2268**	**0.8719**	**2.0341**	**7.6291**
RE	**RE**	1.0133	0.0622	1.0106	0.0568	0.0047	0.0088
DIS	RLS	1.0192	0.0703	1.0112	0.0625	0.0079	0.0174
	CG, $g = 0.2$	1.0163	0.0672	1.0106	0.0568	0.0057	0.0072
	CG, $g = 0.4$	1.0163	0.0675	1.0102	0.0574	0.0063	0.0185
AH-B	RLS	1.0112	0.0586	1.0099	0.0536	0.0012	0.0051
	CG, $g = 0.2$	1.0144	0.0678	1.0101	0.0617	0.0043	0.0114
	CG, $g = 0.4$	1.0146	0.0663	1.0100	0.0579	0.0046	0.0155
AH-A	RLS	1.0263	0.0915	1.0122	0.0852	0.0141	0.0186
	CG, $g = 0.2$	1.0229	0.0750	1.0113	0.0640	0.0116	0.0249
	CG, $g = 0.4$	1.0189	0.0697	1.0109	0.0568	0.0080	0.0246

Table 4A: Production Economy, Statistics for Returns

	$\sigma = 0.00712$	P_{AL}/P_{RE}	**Mean(***P/D***)**	**STD(***P/D***)**	**Corr(***P/D***)**	**STD(**Δc**)**	**STD(**Δd**)**
Data			28.306	9.0611	0.9654	1.0725	12.059
RE	RE		24.809	1.7315	0.9762	0.4508	5.8688
DIS	RLS	1.0073	24.809	1.8862	0.9772	0.5931	5.7070
	CG, $g = 0.2$	1.2457	25.914	2.5706	0.9690	1.1014	7.5479
	CG, $g = 0.4$	1.0854	24.980	2.4780	0.9499	1.5785	8.7195
AH-B	RLS	0.6677	24.784	1.7315	0.9799	0.8716	3.3341
	CG, $g = 0.2$	0.9849	24.870	1.8758	0.9696	1.1101	5.4975
	CG, $g = 0.4$	0.9855	24.847	1.8692	0.9521	1.5292	7.6803
AH-A	RLS	1.3405	25.031	3.3826	0.9823	0.3638	8.7862
	CG, $g = 0.2$	1.4361	25.303	3.8062	0.9681	1.1481	9.2412
	CG, $g = 0.4$	1.2350	24.877	2.5207	0.9475	1.6857	10.443

Table 4B: Production Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth

$\sigma = 0.00712$		BETAS						R-SQUARE		
		1 year		*2 years*		*4 years*		*1 year*	*2 years*	*4 years*
Data	**Data**	-0.0533		-0.1073		-0.1856		0.0926	0.2044	0.3683
		Aver.	% Sig.	Aver.	% Sig.	Aver.	% Sig.			
RE	RE	0.0001	0.0	0.0003	0.0	0.0004	0.0	0.4741	0.3859	0.2583
DIS	RLS	0.0001	4.6	0.0002	4.8	0.0003	5.8	0.3683	0.2930	0.2025
	CG, $g = 0.2$	-0.0001	6.8	0.0001	6.6	0.0002	9	0.1854	0.1553	0.1194
	CG, $g = 0.4$	0.0000	7.4	0.0000	7.6	0.0000	9.4	0.0927	0.0835	0.0698
AH-B	RLS	0.0000	10.8	0.0000	3.2	0.0000	3.2	0.0883	0.0835	0.0935
	CG, $g = 0.2$	0.0001	5.6	0.0001	6	0.0002	6	0.1895	0.1541	0.1150
	CG, $g = 0.4$	-0.0000	8.8	0.0000	10	0.0001	10	0.0960	0.0832	0.0695
AH-A	RLS	0.0001	16.2	0.0002	17	0.0002	17	0.2945	0.2547	0.1967
	CG, $g = 0.2$	0.0000	6.4	0.0000	9.6	0.0001	9.6	0.1695	0.1560	0.1358
	CG, $g = 0.4$	0.0001	7.4	0.0000	13.2	0.0000	13.2	0.0957	0.0892	0.0799

Table 4C: Production Economy, Predictability of Excess Returns