ASSET PRICING WITH ADAPTIVE LEARNING*

Eva Carceles-Poveda[†] Chryssi Giannitsarou[‡]

February 27, 2007

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, since the algorithm converges relatively fast to rational expectations. On the other hand, constant gain learning may contribute towards explaining the stock price and return volatility as well as the predictability of excess returns in the endowment economy. In the production economy, however, 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 cannot fully 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 is often 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 whether and to what extent this assertion is true for asset pricing facts in a general equilibrium framework. We focus on three groups of asset pricing facts, namely first and second asset return moments, the predictability of future excess returns and the volatility of equity prices. Our work is both of qualitative and quantitative nature: we discuss how adaptive learning can help the relevant statistics move towards the right direction and whether it can generate statistics that are close to those observed in the data. At the same time we are interested in examining if and how much better adaptive learning can do relative to rational expectations.

Why would we expect adaptive learning to perform better than rational expectations in an asset pricing framework? Consider first the volatility of equity prices. Under rational expectations, this volatility depends in a direct way on the volatility of the underlying exogenous process that drives the uncertainty in the economy. On the other hand, adaptive learning may introduce an extra source of volatility due to the fact that certain parameters (that are known under rational expectations) are estimated via some statistical rule. Next consider the asset return moments, and in particular the equity premium and its volatility. If the equity price is

*We thank Seppo Honkapohja, two anonymous referees and seminar participants at the 2^{nd} Cambridge-Princeton Meeting, the Bank of England, University College London and University of Oxford for helpful comments. This project has been supported by the ESRC award RES-000-23-1152.

†Department of Economics, SUNY Stony Brook. *E-mail:* ecarcelespov@notes.cc.sunysb.edu.

‡Faculty of Economics, University of Cambridge and CEPR. *E-mail*: cg349@cam.ac.uk.

more volatile under adaptive learning, then the asset is perceived as being riskier than under rational expectations. This is because dividends are either exogenous or depend positively on equity prices. In turn, this results in a higher equity return and thus a higher equity premium. Finally, consider the predictability of future excess returns. If the equity price is below its long run average value, the future dividend yields and capital gains will be higher, leading to higher future returns. This mechanism generates a negative correlation between the current price-to-dividend ratio and future excess returns. If the equity price is more volatile under adaptive learning than under rational expectations, this negative correlation is magnified, therefore improving predictability.

We study the quantitative effects of adaptive learning on equity prices by incorporating 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. In particular, we consider log-linear versions of these general equilibrium models, with self-referential learning on the endogenous variables, under the assumption of a stationary dividend process. We deliberately restrict attention to standard modelling 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 departures from rational expectations can help explain stylized facts on equity prices and returns.

We start by presenting in more detail mechanisms through which recursive least squares or constant gain learning may do better than rational expectations in explaining the observed stylized facts. We also explain why the success of these mechanisms ultimately depends on the parameterization and the numerical specifications for adaptive learning.

Next, we evaluate the effects of the different learning algorithms by running numerical experiments based on standard calibrations for both models. First, we find that overall recursive least squares learning generates very little to almost no improvement of the statistics for neither of the two models. This is because recursive least squares is an algorithm that converges point wise to the rational expectations equilibrium and convergence is relatively fast; therefore its dynamics differ little from the rational expectations dynamics. Second, we find that constant gain learning may be able to drive certain asset statistics towards the correct direction. While the improvement of the statistics relative to those under rational expectations can be quite sizeable, the absolute magnitude of most of these improvements is generally too small to consider interesting. In more detail, for the benchmark parameterization, we find essentially no improvements with respect to the asset return statistics in the production economy. As in the fully rational model, this model performs very poorly under learning with respect to equity price behavior. On the other hand, we find a moderate effect of adaptive learning on the equity premium in the endowment economy. Moreover, constant gain learning can generate higher equity price volatility in both models, but the relative improvement in the production economy is much smaller. We attribute this difference between the two models to the fact that in the production economy there is an additional source of endogenous persistence (due to capital accumulation) that smooths out equity prices. Finally, constant gain learning can generate the predictability of future excess returns that we observe in the data in the context of the endowment economy. This is not surprising, since predictability is a relative feature that only requires a strong negative correlation between the price-to-dividend ratio and future excess returns rather than dependence on the absolute sizes of these.

Finally, we perform an extensive sensitivity analysis with respect to various parameters of the two models, as well as features that have to do with the specifications of the adaptive learning algorithms, such as initial conditions, length of simulations, etc. Regarding the latter, we find that our results can be quite sensitive to the initial conditions of the learning algorithm, given that the benchmark length of the simulations (corresponding to the length of the data time series) is relatively short. Furthermore, we find that the results are sensitive to the size of the gain when using constant gain learning. Specifically, we find that the improvements under learning relative to rational expectations become smaller the longer the memory of the constant gain algorithm is. This is because as the memory of the learning algorithm becomes longer, the equity price becomes less volatile, resulting in a smaller equity premium and weaker negative correlation between the price-to-dividend ratio and future excess returns.

Regarding the sensitivity of the results with respect to various parameters of the models, we first find (not surprisingly) that a higher coefficient of relative risk aversion improves the adaptive learning results on the volatility and the equity premium in absolute terms; however the relative improvements compared to the results under rational expectations are identical irrespective of the coefficient of relative risk aversion. As the coefficient of relative risk aversion increases, we also see that the results on predictability improve, since more volatile prices imply stronger negative correlation between the price-to-dividend ratio and future excess returns. Second, we find that the relative improvement in the equity price volatility under adaptive learning does not depend on the variance of the shocks in the two economies. Moreover, as the variance increases, the relative improvement in the equity premium is unchanged for the endowment economy but only increases slightly for the production economy. A higher variance also improves predictability for both models. Last, we perform sensitivity analysis with respect to the persistence of the exogenous shock and we find that as the persistence decreases, the system dynamics are less sensitive to the specifications of learning. This is because if there is an estimate that is very bad (e.g. very far from RE) then this will feed into the dynamics for many periods if the persistence is high, while it will disappear more quickly if the persistence is low.

In summary, we conclude that self-referential linear adaptive learning under the assumption of a stationary dividend process may provide some qualitative improvements relative to rational expectations. Overall, however, it does not seem to provide satisfactory explanations for the magnitude of various asset pricing statistics that we observe in the data. This is especially prevalent in models with capital accumulation.

Our findings are in contrast to the results in the well known work of Timmermann (1994, 1996). The three main differences with Timmermann's work are the following. First, his analysis is carried out in partial equilibrium while we study general equilibrium models. Note that his setting with constant rates of return can be interpreted as a general equilibrium framework only if utility is linear in consumption. Second, he assumes two different specifications for the dividend process, but both include a drift and a trend. We do not allow for any of the last two. Instead, we assume that the dividend process is stationary. Third, and most importantly, Timmermann considers two types of learning, which he calls *present value* learning and *self-referential* learning. The first is essentially standard OLS estimation written in recursive form. There is no self-referential element in this specification, since the estimation is on the (exogenous) dividend process. The second type of learning is self-referential, but it differs from ours, since it also allows for estimation of the exogenous dividend process. Instead, we assume that exogenous state variables are completely known. Moreover, self-referential learning may contain lags of the price in our production economy, while his estimates depend only on dividends. This lag induces endogenous persistence that reduces the volatility of equity prices and the predictability of the price-to-dividend ratio for future returns considerably.

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. Apart from the work of Timmermann (1994, 1996), this literature includes the papers of Brock and Hommes (1998), Cecchetti, Lam and Mark (2000), Brennan and Xia (2001), Bullard and Duffy (2001), Honkapohja and Mitra (2003), and more recently Adam, Marcet and Nicollini (2006) and Kim (2006).

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. 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. Cecchetti et al. (2000) consider a standard Lucas asset pricing model where agents are assumed to be boundedly rational and have misspecified beliefs. Adam, Marcet and Nicollini (2006) and Kim (2006) both analyze the effects of adaptive learning in the context of the Lucas Tree model. The former emphasize the relationship between adaptive learning and stock market crashes, while the latter work focuses on the combination of adaptive learning with structural shifts.

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 agents' expectations about prices are correctly specified, in the sense that all relevant variables are taken into account when forecasting, and that agents learn about deviations from a steady state. In particular, we do not allow for learning on the growth rate of dividends. 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 of these 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 stylized facts. Section 2 presents the model economies and section 3 discusses the calculation of the rational expectations and adaptive learning equilibria, as well as the mechanisms at work when studying the dynamics of adaptive learning. Section 4 presents the numerical results, section 5 presents the sensitivity analysis and section 6 concludes.

2. Stylized Facts

Table 1 presents the stylized asset pricing facts that we focus on and will use to compare the different models under rational expectations and adaptive learning. The numbers have been calculated using the data set in Campbell (2002).[1] The quarterly stock returns and the quarterly dividend series are obtained from the nominal CRSP NYSE/AMEX Value Weighted Indices. Following Campbell (2002), the price-to-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 corresponds to 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. Finally, 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.[2]

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-to-dividend ratio divided by its standard deviation. Thus, the slope coefficients reflect the effect of a one standard deviation change in the log price-to-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.[3] As reflected by the table, the predictive regressions exhibit the familiar pattern of an increasing R^2 and coefficient slope for longer horizons. The

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

[2] Several authors have argued that the US equity premium has declined considerably during the last three decades (see e.g. Jaganathan et all (2000)). However, generating a positive premium still poses a challenge for standard rational expectation models, particularly in the presence of a production sector (see for example Rouwenhorst (1995), Jermann (1998), Boldrin, Christiano and Fisher (2001) or Lettau (2003)).

[3] 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.

fact that the log price-to-dividend ratio may predict 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.[4]

Finally, since the price-to-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. The last panel reports the standard deviation of consumption and dividend growth.

3. The Environment

This section describes two standard general equilibrium asset pricing models. The first model, which we call the *production economy*, allows for capital accumulation, so that the model mimics the features of the neoclassical growth model. The second model, which we call the *endowment economy*, does not allow for capital accumulation or depreciation of capital and its dynamics can be viewed as a special case of the first by assuming constant capital over time.

For both economies, we will analyze the adaptive learning dynamics and compare them to rational expectation dynamics using log-linear approximations of the equilibrium conditions. This follows Jermann (1998), Lettau (2003) and Carceles-Poveda (2005) among others. Log-linear approximations may not always be very accurate, however they are known to perform reasonably well in general equilibrium models of the type studied here. Moreover, the log-linear framework provides a convenient platform for studying adaptive learning dynamics, since many more theoretical results have been developed for linear models than for non-linear ones. Besides, here we are mainly interested in relatively small deviations of variables from their stationary long-run averages, therefore a log-linear framework should be relatively accurate. Also, to avoid losing second order information when calculating the risk premium, we use the approach described in Jermann (1998) and Lettau (2003), which essentially corrects the log-linear asset pricing equations for Jensen terms.

3.1. 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

[4] Recent literature on the issue of predictability of future stock returns shows that the $t-$ statistics reported from such regressions might be misleading, due to the high autocorrelation of the price dividend ratio (Campbell and Yogo, 2005). However we report these to make the analysis comparable to existing literature.

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, and 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

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

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

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

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

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

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[5]

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

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

$$\log Z_t = \rho \log Z_{t-1} + \varepsilon_t, \quad (9)$$

where $\varepsilon_t \sim iid(0, \sigma_\varepsilon^2)$ and $\rho \in (0, 1)$. 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

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

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)

$$\max \; E_t \sum_{j=0}^{\infty} M_{t,t+j} D_{t+j}. \quad (11)$$

[5] The timing $t-1$ in the index of capital is conventional and does not affect the analysis that follows. Following a large amount of real business cycle literature, we use K_{t-1} instead of K_t in order to denote more clearly that capital is a state variable.

The first-order conditions are

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

Finally, market clearing implies that

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

To derive the system of equations that describe the equilibrium, we substitute for $N_t = 1$, $B_t = 0$, $\Theta_t = 1$ and $W_t = (1-\alpha)\,Y_t$. Moreover, we can omit the the resource constraint by Walras law, as well as the capital Euler equation (13), since $K_t = P_t$ in equilibrium. 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:

$$\begin{aligned} z_{t+1} &= \rho z_t + \varepsilon_{t+1}, & (16a)\\ y_t &= z_t + \alpha k_{t-1}, & (16b)\\ c_t &= \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)-\alpha\beta\delta}y_t + \frac{(1-\delta)\,\alpha\beta}{1-\beta(1-\delta)-\alpha\beta\delta}k_{t-1} - \frac{\alpha\beta}{1-\beta(1-\delta)-\alpha\beta\delta}k_t, & (16c)\\ d_t &= \frac{1-\beta(1-\delta)}{1-\beta}y_t + \frac{\beta(1-\delta)}{1-\beta}k_{t-1} - \frac{\beta}{1-\beta}k_t, & (16d)\\ p_t &= E_t\left[-\gamma\left(c_{t+1}-c_t\right)+(1-\beta)d_{t+1}+\beta p_{t+1}\right], & (16e)\\ p_t^b &= E_t\left[-\gamma\left(c_{t+1}-c_t\right)\right], & (16f)\\ k_t &= p_t. & (16g) \end{aligned}$$

This model is along the lines of well known general equilibrium asset pricing models with production (e.g. see Brock, 1982, Rouwenhorst, 1995 and Lettau, 2003).

3.2. The Endowment Economy. In the endowment economy 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 system of equations (16a) - (16g), resulting in the following log-linear model:

$$\begin{aligned} z_{t+1} &= \rho z_t + \varepsilon_{t+1} & (17a)\\ c_t &= d_t = y_t = z_t & (17b)\\ p_t &= E_t\left[-\gamma\left(d_{t+1}-d_t\right)+(1-\beta)d_{t+1}+\beta p_{t+1}\right] & (17c)\\ p_t^b &= E_t\left[-\gamma\left(d_{t+1}-d_t\right)\right] & (17d) \end{aligned}$$

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.

Note that 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, and 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

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

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

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

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

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

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

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

4. Rational Expectations and Adaptive Learning

In order to calculate the rational expectations equilibria of the production economy, 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

$$a_1 = \frac{-\gamma}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{23a}$$

$$a_2 = \frac{\gamma\left(\delta - 1 - \alpha\psi\right)}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{23b}$$

$$b = \frac{\psi\left(\gamma\left(\rho - 1\right) + \alpha\beta\left(\delta - \psi\right)\rho\right)}{\gamma\left(-2 + \delta - \alpha\psi\right) + \left(\delta - \psi\right)\left(1 + \beta\left(\delta - 1 - \alpha^2\psi\right)\right)}, \tag{23c}$$

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

Similarly, the reduced form for the endowment model is given by

$$p_t = a E_t p_{t+1} + b d_t, \tag{24}$$

$$d_t = \rho d_{t-1} + \varepsilon_t, \tag{25}$$

where

$$a = \beta, \tag{26}$$
$$b = (1 - \beta - \gamma)\,\rho + \gamma. \tag{27}$$

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

$$p_t = \bar{\phi}_p p_{t-1} + \bar{\phi}_z z_{t-1} + \eta_t, \tag{28}$$

where η_t is some white noise shock and[6]

$$\bar{\phi}_p = \frac{1}{2a_1}\left(1 - \sqrt{1 - 4a_1a_2}\right), \tag{29}$$
$$\bar{\phi}_z = \frac{b}{1 - a_1(\rho + \bar{\phi}_p)}\rho. \tag{30}$$

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

$$p_t = \bar{\phi} d_{t-1} + \eta_t, \tag{31}$$

where η_t is a white noise shock and

$$\bar{\phi} = \frac{(1 - \beta - \gamma)\,\rho + \gamma}{1 - \beta\rho}\rho. \tag{32}$$

Some points are worth noting. First, if we compare the models under rational expectations, the solution for the production economy (28) contains a lag of the price, while the solution of the endowment economy (31) 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 that is injected 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 (see the discussion in Rouwenhorst, 1995). Therefore, there seems to be a better 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. Third, the equilibrium consumption and dividend processes turn out to be equal in the endowment economy. Therefore when attempting to calibrate the model to match the data, we will only be able to match the

[6] 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).

behavior of one of these two variables at a time. Fourth, the equity price turns out to be equal to the capital stock in the production economy. This implies that we will not be able to increase the volatility of the equity price without compromising the volatility of capital, which is much lower than the volatility of the equity price in the data.

4.2. Adaptive Learning. Next, we make a small deviation from rational expectations by assuming that agents form expectations about future prices based on econometric forecasts. We should point out that under rational expectations, the only source of uncertainty in the two economies is the exogenous stochastic process. The rest of the parameters and laws of motions of variables are completely known. Thus, when households have to make consumption and savings decisions, they optimize conditional on the realizations of these exogenous shocks. In other words, under rational expectations, agents' forecasts are *on average* correct, since the only unknown element is the realization of the exogenous noise.

In contrast, adaptive learning it is implicitly assumed that the average forecasts of agents are not necessarily correct. This can be due to various reasons, but we will focus on the scenario where, although agents know the deep parameters of the model (e.g. preference parameters), they do not know in what way these parameters determine the evolution of prices variables over time. Moreover, the type of learning we analyze here is self-referential in the sense that agents' forecasts influence the laws of motion of the economic variables, which in turn then influence the subsequent future forecasts and so on. In this sense, adaptive learning introduces an additional source of uncertainty in the model that is eventually reflected in the dynamics of the economies: imprecise forecasts are used when agents and firms make decisions, leading to potentially non-optimal temporary equilibria.

Given this background and 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 are aware that they are estimating deviations from a 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. An interesting direction that is beyond of the scope of the present paper would be to relax A1 (i.e. to introduce model misspecification). Moreover, we conjecture that relaxing A2, i.e. allowing agents to estimate parameters ρ and σ^2, will not alter the results significantly. Such an extension would probably improve the results somewhat, since it would feed some extra volatility into the system. However, these parameters characterize an exogenous variable, implying that any econometric learning or estimation procedure in search of the true values would converge relatively fast without significantly affecting the evolution of the endogenous variables. Moreover, since we want to study types of learning that are the closest possible to rational expectations, we abstract from learning on the exogenous variable parameters.

Given these assumptions, agents' expectations for both models are formed according to

$$E_t^* p_{t+1} = x_t' \phi_t, \tag{33}$$

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[7]

$$\begin{cases} R_1 = S_0 + x_0 x_0' \\ \phi_1 = \phi_0 + R_1^{-1} x_0 (k_1 - x_0' \phi_0) \end{cases}, \tag{34a}$$

$$\begin{cases} R_t = R_{t-1} + g_t \left(x_{t-1} x_{t-1}' - R_{t-1} \right) \\ \phi_t = \phi_{t-1} + g_t R_t^{-1} x_{t-1} \left(p_t - x_{t-1}' \phi_{t-1} \right) \end{cases} \quad \text{for } t \in \{2, 3, ...\}, \tag{34b}$$

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, RLS learning 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, CG learning can be interpreted as the method that is used when the econometrician believes that recent realizations of the equity price are more important in forecasting next period's price.

Another difference between the two algorithms is their asymptotic behavior. First, convergence of the RLS algorithm is in the "almost surely" sense. It is global for the endowment economy and local of the production economy, whenever the E-stability conditions are satisfied (these are always satisfied for reasonable parameter ranges of the two models). Furthermore, to ensure local convergence for the production economy, a projection facility needs to be invoked (e.g. a restriction ensuring that the estimates ϕ_t imply a stationary endogenous state variable). This has interesting implications for the numerical results, as will become clearer later. Second, convergence of the CG algorithm is in the "distribution" sense, that is, CG learning converges to some distribution, for small positive gains.[8] In particular, 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 long-run average solution.

[7] See Carceles-Poveda and Giannitsarou (2007) for a derivation.

[8] More details on convergence issues and on the derivations of these conditions can be found in Evans and Honkapohja (2001), as well as in Carceles-Poveda and Giannitsarou (2007).

Note that initial conditions that are away from the REE are less important for the speed of convergence under CG learning than under RLS learning. This is because the CG algorithm is by definition much better at tracking large jumps of the estimates away from the long run average (such as structural shifts) than RLS: since more weight is assigned to recent observations, even if the initial condition is far from the REE, its effect will become less and less important over time and will eventually disappear much faster than if we used RLS.

Finally, we want to point out that we do not wish to provide a formal argument in favor of one algorithm over the other. Such an exercise would involve working out the optimal learning algorithm, in some appropriately defined sense of optimality. Instead, our aim is to compare the behavior of equity prices under various specifications of the two algorithms.

The rest of the section is devoted to describing mechanisms through which the adaptive learning algorithms we consider may or may not generate improved asset pricing statistics. We argue that the behavior of the statistics and facts that we are interested in depends crucially on the variance of the equity price under adaptive learning both in absolute terms and relative to the variance of equity prices under rational expectations. To see what this the case, note first that the volatility of equity prices under rational expectations depends in a direct way on the volatility of the underlying exogenous process that drives the uncertainty in the economy. At the other end, adaptive learning may introduce an extra source of variation in prices due to the fact that certain parameters are now estimated via some statistical rule. Moreover, the variance of the equity prices changes over time and may be higher or lower than the constant (rational expectations) variance.

To see how this would affect the other statistics of interest, consider first the asset return moments and in particular the equity premium and its variability. In general, if the equity price is more volatile under adaptive learning, then the asset is perceived as being riskier than under rational expectations. This is because dividends are either exogenous or they depend positively on the equity price. In turn, this will result in a higher equity return and thus a higher equity premium and premium variability. If on the other hand the equity price varies less under learning, then the asset is perceived as being safer than under rational expectations, resulting in a lower equity premium and premium variability.

Second, consider the predictability of future excess returns. If the equity price is below its long run average value, this will result in both a higher dividend yield and in higher future capital gains when the price adjusts upwards, leading to higher future returns. This mechanism generates a negative correlation between the current price-to-dividend ratio and future excess returns. Moreover, the correlation will be magnified if the equity price is more volatile under adaptive learning than under rational expectations, improving the predictability of the price to dividend ratio. The opposite will happen if the equity price is less volatile under adaptive learning. In conclusion, to understand how adaptive learning in our models can contribute towards explaining asset pricing statistics, it is very important to understand the learning dynamics of the variation of the equity price.

It should also be clear that the extent to which RLS can explain asset pricing facts within these two models depends on the initial values and the speed at which the algorithm converges to the rational expectations equilibrium. If for example the priors of the agents are close to the REE and the algorithm converges quickly, we should not expect to see any significant improvement in the results relative to rational expectations. On the other hand, since CG learning implies perpetual learning and does not converge point-wise to the REE, the initial values should not matter that much, and we may expect to see more interesting dynamics than under RLS.

With the preceding discussion in mind, we can now go deeper into the mechanisms that generate equity price volatility under adaptive learning**.** Due to the relative simplicity of the reduced form model for the endowment economy, we can go quite far analytically in this case. However, the reduced form of the production economy includes a lag of the endogenous state variable (i.e. the equity price), making the dynamics under learning too complicated to study analytically in a meaningful way. For the latter model, we will therefore be able to see clearer results through numerical experiments that are presented in the next section. For this reason, we focus on the endowment economy and we conjecture that one can apply loosely similar arguments for the production economy.

Consider first the dynamics of the equity price in the endowment economy. Under rational expectations, this is given by:

$$p_t^{RE} = h(\bar{\phi})d_t, \tag{35}$$

where $h(\phi) = a\phi + b = \beta\phi + (1 - \beta - \gamma)\rho + \gamma$, so that the variance of the equity price is

$$Var\left(p_t^{RE}\right) = h(\bar{\phi})^2\sigma_d^2. \tag{36}$$

The variance of the equity price under adaptive learning at a given period t is

$$Var\left(p_t^{AL}\right) = Var\left(h\left(\phi_{t-1}\right)d_t\right). \tag{37}$$

Furthermore, given our assumption of normal noise shocks and since $E\left(d_t\right) = 0$, the variance of this product can be expressed as follows (see Bacon, 1980)

$$Var\left(p_t^{AL}\right) = Var\left(h\left(\phi_{t-1}\right)d_t\right) = \left[m_{h,t}^2 + \left(1 + r_t^2\right)\sigma_{h,t}^2\right]\sigma_d^2, \tag{38}$$

where

$$\begin{aligned} m_{h,t} &= E\left(h\left(\phi_{t-1}\right)\right) = aE\left(\phi_{t-1}\right) + b, & (39)\\ \sigma_{h,t}^2 &= Var\left(h\left(\phi_{t-1}\right)\right) = a^2Var\left(\phi_{t-1}\right), & (40)\\ r_t &= Corr\left(h\left(\phi_{t-1}\right), d_t\right) = aCorr\left(\phi_{t-1}, d_t\right). & (41) \end{aligned}$$

Using (38), we can make the following observations. First, the variance of the equity price under adaptive learning depends positively on the variance of the exogenous shock (here the dividend). In other words, whenever the variance of the shock is higher, we should expect a more volatile equity price under learning. In addition, the variance at time t depends positively on the average estimate ϕ_{t-1}, the variance of the estimate and the correlation of the estimate dated $t-1$ with the exogenous shock at t.

To gain further insights, we define the *relative variance of the equity price at time* t as the variance of the equity price under adaptive learning over the variance of the equity price under rational expectations, i.e.:

$$\Sigma_t = \frac{Var\left(p_t^{AL}\right)}{Var\left(p_t^{RE}\right)} = \frac{m_{h,t}^2 + \left(1 + r_t^2\right)\sigma_{h,t}^2}{h(\bar{\phi})^2}. \tag{42}$$

If we use Σ_t as a measure of the change in variance of the equity price under adaptive learning relative to the variance under rational expectations, we now see that any *relative* improvement in the volatility of equity prices under learning *does not depend on the variance of the exogenous shock.* If Σ_t is larger than one, then the equity price will have a higher variance under adaptive learning than under rational expectations. In turn, since the dividends are exogenous and taken the same under both assumptions (there is no "learning" of the dividend process), a higher equity price variance under learning will imply that the equity is perceived as being riskier. Since the risk free rate is the same under both learning and rational expectations, this will in turn lead to a higher premium and to a higher premium volatility.

Regarding the predictability of excess returns, this will improve if $\sigma_{h,t}^2$ and r_t are high. The reason is the following. For a given dividend process, if ϕ_{t-1} is lower than average, then $h\left(\phi_{t-1}\right)$ will be lower than average and p_t^{AL} will be smaller than p_t^{RE}. In turn, this will imply that the current price-to-dividend ratio under learning will be lower than the one under rational expectations. Moreover, future returns will tend to be higher than under rational expectations due to the current high dividend yield and the future upward adjustment of the price to its long run average, generating capital gains. This mechanism will be reinforced if ϕ_{t-1} is more volatile, since this will lead to a more volatile $h\left(\phi_{t-1}\right)$ and to a more volatile price. In addition, if r_t is higher, these effects will be amplified even more, since a higher than average estimate combined with a higher shock will lead to an even higher price volatility.

What remains to be determined is how the variance of the equity price and the relative variance Σ_t behave for given parametrizations and learning specifications and, in particular, which of the three elements $m_{h,t}^2$, $\sigma_{h,t}^2$ and $\sigma_{h,t}^2 r_t^2$ is most important for determining the size of Σ_t. It is worth noting here that the last term $\sigma_{h,t}^2 r_t^2$ will be sizeable only if both r_t and $\sigma_{h,t}^2$ are quite high. In the next two sections we will explore these relations in more detail by performing various illuminating numerical experiments.

Finally, with regards to the production economy, it is not so straightforward to do a similar analysis. In this case, the equity price volatility under rational expectations is given by the following expression:[9]

$$Var(p_t^{RE}) = \frac{\gamma^2(1+\rho\bar{\phi}_p)}{(1-a_1\rho - a_1\bar{\phi}_p)^2(1-\rho\bar{\phi}_p)(1-\bar{\phi}_p^2)}. \tag{43}$$

This variance is a constant that depends on the parameters of the model. However, under adaptive learning, we have that

$$\begin{aligned} Var\left(p_t^{AL}\right) &= Var\left[\tau\left(\phi_{t-1}\right)p_{t-1} + h\left(\phi_{t-1}\right)d_t\right] \\ &= Var\left[\tau\left(\phi_{t-1}\right)p_{t-1}\right] + Var\left[h\left(\phi_{t-1}\right)d_t\right] + 2Cov\left[\tau\left(\phi_{t-1}\right)p_{t-1}, h\left(\phi_{t-1}\right)d_t\right], \end{aligned} \tag{44}$$

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

where

$$\tau(\phi) = \frac{a_2}{1 - a_1\phi_p} \text{ and } h(\phi) = \frac{a_1\phi_z + b}{1 - a_1\phi_p} \tag{45}$$

Unfortunately, this expression is too complicated to work with analytically and get meaningful conclusions. Nevertheless, we can use the breakdown in (44) when we do our numerical analysis to gain some insights about the importance of the different terms in determining the equity price volatility and its relationship with the other statistics.

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 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 (i.e. the average relative deviations under learning), as a proxy for the equity price volatility generated by adaptive learning relative to rational expectations.

We begin by describing the computing specifications and the calibrations. To implement the simulations we have used the adaptive learning toolbox for Matlab that accompanies Carceles-Poveda and Giannitsarou (2007). For each model, we run experiments with a number of $T = 211$ periods, corresponding to the number of quarters available in the data set. The statistics reported are the average statistics from replicating the experiments $N = 3000$ 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 (2007), 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 elasticities are ϕ_0 are 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 fifty observations (the larger the number of observations the closer the initial condition is to the REE). In the second, ϕ_0 is set at an ad-hoc value that is below or above the rational expectations value. These two values correspond to different initial priors of the households about the effects of the state variables on the current equity price. Note that one way to interpret initial conditions that are relatively far from the REE is that agents learn a new equilibrium after a structural change in the economy. Although we do not address structural shifts explicitly, such an interpretation is an interesting starting point for how such an assumption may, for example, explain the equity premium puzzle. A more thorough analysis of this assumption and its consequences for asset pricing statistics under learning is done by Kim (2006) in the context of the Lucas tree model. Finally, for each set of experiments, we simulate series under RLS learning and CG learning.

Turning to the parametrization of the gain, we use values of $g = 0.02$, $g = 0.2$ and $g = 0.4$. The size of the gain may be determined in various ways. For example, it may be estimated, so that it matches stylized facts, or it can be determined so that it gives the smallest possible mean squared forecasting error. Here, our choice of the gain values is based on the basic 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.[10] 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 quarters back is displayed in table 2.[11] The table also reports the half-life decay for these gains in quarters. For example, if the econometrician believes that only observations that date at most $i = 15$ 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 relatively high gain coefficient may be 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.

Note that the numbers for the gain turn out to be quite high due to the fact that we assume that the data are in quarterly frequency. It is true that forecasters in the financial sector use high frequency data (weekly, daily or even minute by minute), which would translate into a lower gain when considering 20 years of data; however, here we are working with quarterly data not only for comparability to existing work, but also because we care about the behavior of the aggregate macroeconomic variables. Finally, 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 >

The rest of the parameters are calibrated as follows. The risk aversion coefficient is set to $\gamma = 1$ in both models.[12] For the production economy, we have used the standard parametrization for US quarterly data, that is, the capital depreciation, the discount factor and the capital share are set to $\delta = 0.025$, $\beta = 0.99$ and $\alpha = 0.36$ respectively. Furthermore, the baseline parametrization for the productivity shock is $\sigma_\varepsilon = 0.00712$ and $\rho = 0.95$, as is usual in the real business cycle literature.

In the endowment economy, we again set $\gamma = 1$ and $\beta = 0.99$. 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 quarterly dividend series in the data on its first lag. In addition, we 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. It turns out that this last calibration approximately replicates the behavior of logged consumption growth in the data.

[10] The constant gain algorithm is some type of weighted least squares estimator. However, it does not necessarily follow the usual rule of assigning larger weight to observation points with smaller variance.

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

[12] As is well known, a high parameter γ improves the performance of rational expectations asset pricing models, such as consumption based models like our endowment economy. Although, it would help improve the results under learning as well, we prefer a low γ for our benchmark, since it has been documented empirically that values of γ larger than around 5 are implausible (e.g. see Hall, 1988).

We first present the results for the endowment economy and then discuss the results for the production economy. Moreover, we also present a sensitivity analysis for each economy with respect to the key model parameters.

5.1. The Endowment Economy. Tables 3A-3C contain the results for the calibration with the lower variance ($\sigma_\varepsilon = 0.00712$), whereas tables 3D-3F report the same results for the higher shock variance ($\sigma_\varepsilon = 0.06$). Tables 3A and 3D contain the first and second asset moments. Tables 3B and 3E contain (a) the standard deviation of the equity price under learning over the standard deviation of the equity price under rational expectations, (b) the average price-to-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 3000 replications of the experiment.

The first two rows of the tables display the numbers in the data and under rational expectations. Furthermore, the last six rows display the results under learning when the algorithms are initialized (a) from a distribution (DIS) as explained earlier, (b) below the REE, with elasticities 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).[13] 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.

< TABLES 3A - 3F HERE>

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 0.4 percent with the higher benchmark shock variance. Furthermore, the standard deviation and the autocorrelation of the price-to-dividend ratio are far from the data, and this variable generates absolutely no predictability for the excess stock returns. This 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 the different initializations do not seem to have an overall very significant effect on the outcomes. However, we do observe interesting 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. As discussed earlier, 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

[13]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.

dynamics under RLS and rational expectations disappear quickly. This observation is also clear from tables 3B and 3E, where we can see that the relative variability of the equity price is close to one with both parametrizations of the variance σ_{ε}^2. Furthermore, Tables 3A and 3D illustrate that the asset return moments are almost identical to the rational expectation values.

We also note that the relative price variability is less than one and the equity premium under learning is below the one under RE when the initial condition is below the REE value, and vice versa when the initial condition is above the REE. When the initial condition is drawn from a distribution (i.e. close to the REE), both the relative variability is close to one and the equity premium is close to the one of RE. Finally, Tables 3C and 3F illustrate that the model under RLS also performs poorly regarding predictability. The coefficients have the right sign and are higher in absolute value than the ones under rational expectations, but they are still very far from the data under both variance calibrations. In addition, the percentage of significant estimates is relative small.

With this discussion in mind, it should now 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.

Regarding the average asset returns, Table 3D illustrates that CG learning with the benchmark parameterization and the higher gain can generate an average stock return that is 20% higher than its rational expectations value, leading to a premium that is twice as high as its value under rational expectations. However, since the average risk free rate generated is too high compared to its value in the data, the average equity premium is still relatively small.[14]

On the other hand, we see a considerable improvement regarding the volatility and predictability. Tables 3B and 3E reflect that the equity price under CG learning can be 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 and of the stock return and equity premium, whereas the standard deviation of the price-dividend ratio is just about half of the one observed in the data. In addition, 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 very 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 ranges from approximately 40% to 70%, a large improvement compared to the results under RLS and rational expectations.

Finally, we note that these improvements are smaller when the model is calibrated to con-

[14] To lower the risk free rate, we can either recalibrate the discount factor or choose a different initial value for the price of the bond. However, this also lowers the risk free rate under rational expectations, generating the same differences with respect to the rational expectations statistics. Since our aim is to compare the performance of learning with respect to rational expectations, we opt for using the current calibration.

sumption behavior ($\sigma = 0.00712$). In this case, the model does not generate the right behavior for the dividend growth or the price-to-dividend ratio. Furthermore, although this calibration leads to a higher equity return and equity price volatility and generates a much higher predictability under learning than under rational expectations, the slope coefficients with a gain of 0.4 do not provide a satisfactory match with the data.

In what follows, we provide some intuitive comments to help understand our findings. We first consider the volatility of the equity price. As explained earlier, whether equity prices are more volatile under learning than under rational expectations depends on the behavior of three terms: $m_{h,t}^2 = \left[aE\left(\phi_{t-1}\right) + b\right]^2$, $\sigma_{h,t}^2 = a^2 Var\left(\phi_{t-1}\right)$ and $r_t^2\sigma_{h,t}^2 = aCorr\left(\phi_{t-1}, d_t\right)\sigma_{h,t}^2$. Under RLS learning, our numerical results show that the estimated coefficient ϕ_{t-1} is very close to the REE value, while the variance and correlation terms are relatively small. As a result, the relative price variability Σ_t is very close to one, and the equity return is not perceived as being riskier under learning. This explains why the equity premia are very close to those under REE. As noted before, we see an effect of the initial conditions on the relative price variability Σ_t, which is below (above) one when we initialize the coefficient below (above) the REE value. This is due to the fact that $m_{h,t}^2$ is one average below (above) the REE value when we initialize below (above), generating premia that follow the same pattern. However, since the RLS algorithm converges relatively fast, the differences are very small.

Consider now the case of CG learning. First, we see that the initialization does not seem to matter in this case. This is due to the fact that $m_{h,t}^2$ is much less important than $\sigma_{h,t}^2$ and $r_t^2\sigma_{h,t}^2$ for determining the relative price variability Σ_t. In particular, our numerical results show that $\sigma_{h,t}^2$ is approximately 20 times higher than under RLS with CG learning and a gain of 0.4, while $r_t^2\sigma_{h,t}^2$ is also considerably bigger. In turn, this leads to a higher price variability Σ_t than under RLS learning, generating a higher premium. Moreover, the higher price variability, combined with a higher correlation of the estimate with the exogenous shock process, improves on the predictability of excess returns through the mechanisms explained earlier.

To illustrate how the three terms $m_{h,t}^2$, $\sigma_{h,t}^2$ and $r_t^2\sigma_{h,t}^2$ contribute to the relative variability of the equity price, we present as an example how these evolve on average for the cases of RLS and CG learning with $g = 0.4$, using the initial condition that is set below the REE. These are shown in Figures 1 and 2. The panels in the Figures show these three terms, as well as Σ_t. Note that the term $r_t^2\sigma_{h,t}^2$ is essentially insignificant for both cases. In the case of CG learning where $r_t^2\sigma_{h,t}^2$ is larger than under RLS learning, this term is of order of magnitude of 10^{-3}. This is because, even in this case with high variance and gain, the maximum correlation between beliefs ϕ_{t-1} and the exogenous state d_t is very small, at around 5%.

< FIGURES 1 - 2 HERE >

Confirming our earlier derivations, the relative variability of the equity price is independent of the variance of the exogenous shock process. Thus, while the absolute value of the equity premium is higher as the shock variance increases, the relative improvement of the equity premium under learning is the same as with the low variance. In contrast, the model does perform much better regarding the predictability of excess returns with the higher shock variance. This is because predictability only requires a strong negative correlation between the price-to-dividend ratio and the future excess returns. While the relative improvement in equity price volatility is the same for both variances, both the absolute volatility of the equity price and the correlation of the estimates with the dividends increase as the variance increases, generating a stronger negative correlation between the price-to-dividend ratio and future excess returns.

To summarize, RLS learning generates results that are very close to their rational expectations counterparts and the initialization seems to matter somewhat. CG learning can generate some extra volatility of the equity price with a high gain, leading to a higher equity premium and premium variability. The initial conditions matter less than under RLS, as expected, and they matter even less as the gain increases (i.e. as the learning memory reduces).[15] However, the values are still too small compared to the data. We therefore conclude that standard adaptive learning cannot fully explain the mean equity premium in the endowment economy model. On the other hand, our numerical results confirm that adaptive learning is able to generate the excess return predictability that we see in the data.

5.2. The Production Economy. Tables 4A-4C report the results for the benchmark parameterization of the production economy, organized in the same way as the results for the endowment economy. Table 4A contains asset moments, table 4B contains various statistics and table 4C reports the results for predictability.

$<$ TABLES 4A - 4C HERE $>$

As with the endowment economy, the tables indicate that the rational production economy performs very poorly in explaining the first and second asset moments (see for example Rouwenhorst, 1995, or Lettau, 2003). The implied equity premium is approximately 0.002 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-to-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 in general it has a relatively small effect on the different asset moments regardless of the algorithm. The equity premium only increases from 0.0027 up to 0.0154 percent and its variability only increases from 0.007 to 0.023 percent. In addition, although learning improves the behavior of the price-to-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, much more than for the endowment economy. In particular, starting above the rational expectations value generates a premium that is ten times higher than if we start below rational expectations under RLS learning. In addition, Table 4B reflects that the relative variability of the equity price can be considerably lower than under rational expectations if the initial coefficient is set below its rational expectations value, while it can be considerably

[15] This is not so easy to see from the two sets of numbers we report here based on $g = 0.2$ and 0.4, but the pattern is pretty clear once one starts experimenting with a wider range for the gain.

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. Finally, while the differences are still bigger than in the endowment economy with a gain of 0.2, the results exhibit a non-monotonic behavior when the gain becomes larger.

Next, we provide some further insights that will help us understand these results. As already noted in the previous section, the relative equity price variability Σ_t cannot be decomposed nicely as in the endowment economy. However, we can decompose it into three terms: the variance of the term $h\left(\phi_{t-1}\right) z_t$, the variance of the term $\tau\left(\phi_{t-1}\right) p_{t-1}$, and the covariance of these two terms. Our numerical results (not reported here) illustrate that the term $Var\left[\tau\left(\phi_{t-1}\right) p_{t-1}\right]$ is the one that mostly drives the behavior of the equity price variance. This is partly due to the fact that the elasticity of the price with respect to the exogenous shock z_t is much smaller than that in the endowment economy, both under REE and under learning. Moreover, the term $\tau\left(\phi_{t-1}\right)$ only depends on the estimate $\phi_{p,t-1}$, and this therefore the estimated coefficient that is key for determining the behavior of the statistics of interest.

We first consider how learning affects the equity price volatility. Using similar arguments to those in the endowment economy, it is easy to see that the term $Var\left[\tau\left(\phi_{t-1}\right) p_{t-1}\right]$ will increase if the estimate $\phi_{p,t-1}$ is more variable and if it remains well above the rational expectations value $\bar{\phi}_p$ for a large number of periods, while the opposite will happen if the estimated coefficient is less variable and if it remains well below the rational expectations value $\bar{\phi}_p$ 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. In fact, Table 4B reflects that the relative price variability is only slightly above one. Furthermore, a gain of 0.2 makes the coefficients more volatile, so that they are more often above but also below the REE. As a result, the difference between the relative price volatilities is even smaller. Note that increasing the gain coefficient above 0.2 also implies that the coefficients violate the stationarity condition $\left|\phi_{p,t-1}\right| < 1$ very often, since the long-run average value $\bar{\phi}_p$ is very close to 1. 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. Finally, when we initialize the algorithm above or below the REE, the estimated coefficients will remain on average above or below the REE respectively, for many periods. This happens because of endogenous persistence, which is due to the presence of a lag in the law of motion of the equity price. This explains why the different initializations generate different results in the production economy, while they matter much less in the endowment economy.

As in the endowment economy, the results reported here reflect a strong relationship between the relative price variability and the equity premium. In the present model, a higher price volatility under AL will generate a higher premium for two reasons. For a given dividend process, higher equity price volatility implies more risk. Moreover, the higher price volatility will feed back into the dividend process, since the dividend is a function of the equity price as well, inducing more volatility and thus more risk. In fact, Tables 4A and 4B reflect that the model can generate a premium of up to five times larger than the one under REE with the initialization that starts above the REE. Nevertheless, the premium is still very far from the data in all the cases considered.

Regarding the predictability of excess returns, we first note that the arguments that we have used in the endowment economy to explain the negative correlation between the price to dividend ratio and the future excess returns do not necessarily apply here. First, the elasticity of the price with respect to the shock is much smaller than that in the endowment economy and the term that generated the predictability results before does not play a role, at least with a low shock variance. Suppose now that the coefficient ϕ_p is higher than average, leading to a higher than average price. This will feed back into the current and future dividend payments and it will not necessarily lead to a low future dividend yield. In addition, the adjustment of the price to its long run average is likely to be slower than in the endowment economy. Given this, the correlation between the future excess returns and the current price-to-dividend ratio will be weak and not necessarily negative. As we see, this translates into a behavior that is very similar to the one under rational expectations in terms of the predictability of excess returns.

To summarize, with the bechmark parameterization, the production economy is not able to generate any excess return predictability. Regarding the equity premium, the initial conditions matter more than in the endowment economy, generating a higher premium when we initialize the coefficients above their RE value. On the other hand, the magnitude of the premium and its variability are still very far from the data.

6. Sensitivity Analysis

This final section presents some sensitivity analysis of our numerical results with respect to three key parameters of the two models, namely the coefficient of relative risk aversion γ, the variance σ_ε^2 and the persistence ρ of the exogenous shocks. The sensitivity results for the two models are displayed in tables 5A-5B and 6A-6C respectively. The numbers are based on experiments that have the same specifications as our baseline analysis, with 3000 replications. The Tables display the mean and standard deviation of the equity premium and the price-to-dividend ratio, as well as the coefficients of the regressions of the price-to-dividend ratio on the future excess returns. We only report the results for the ad-hoc initializations starting below and above the REE value under RLS and CG learning, with $g = 0.2$. For comparison, the tables contain also the results with our benchmark parametrizations. Finally, we have not included a separate table for sensitivity with respect to σ_ε^2 for the endowment economy, since this case is covered extensively in tables 3A-3F.

6.1. Endowment Economy. Table 5A displays the results for levels of risk aversion of $\gamma = 1$ and $\gamma = 3$. The first important observation is that the relative price variability is the same both risk aversion values. In a way, this is not surprising: altering the deep preference parameters of the models (such as the relative risk aversion) does not alter the learning dynamics *relative* to rational expectations dynamics, because such parameters are not directly relevant for the determination of the learning dynamics. As in the baseline case, we see that the relative price variability is higher under CG learning, which generates an equity premium that is very close to the one we see in the data. Note, however, that the relative improvement with respect to the REE value is approximately the same as with a risk aversion of 1. In addition, the premium volatility is twice as high as in the data, indicating that the model is not able to match the first and second moments simultaneously. In contrast, a higher risk aversion does lead to a considerable improvement with respect to the predictability of excess returns. The regression coefficients generated by the CG algorithm now match the coefficients in the data with a gain of 0.2 and a risk aversion of $\gamma = 3$. The reason why this happens is the same as when we increase the variance of the shock: while the relative price variability with a higher risk aversion is the same, the variability of the equity price and the correlation of the estimated coefficient with the dividends is higher in *absolute* terms, generating a stronger negative correlation between the price-to-dividend ratio and the future excess returns.

< TABLES 5A - 5B HERE >

Table 5B displays the results with levels of the shock persistence of $\rho = 0.1$, $\rho = 0.5$ and $\rho = 0.95$. This table reveals a clear pattern: as the persistence of the shock decreases, the relative variability and consequently all the rest of the statistics of interest are much less sensitive to the variations in the specifications of learning, e.g. the initial conditions or the learning algorithm. The intuition of why this happens is the following. As the persistence becomes smaller, the effects of learning cannot propagate or remain in the dynamics of the system for too long. If for example there is an estimate (under learning) that is very bad, i.e. very far from REE, then this will feed into the dynamics for many periods if the persistence is high, while its effects will disappear more quickly if the persistence is low.

In sum, increasing the risk aversion does not improve the results relative to rational expectations, while a lower persistence tends to make the dynamics more robust to changes in the learning specifications.

6.2. Production Economy. Table 6A displays the results for shock standard deviations of $\sigma_\varepsilon = 0.00712$, $\sigma_\varepsilon = 0.02$ and $\sigma_\varepsilon = 0.04$. As in the endowment economy, we find that the relative price variability does not depend on the variance of the exogenous shock process. On the other hand, we do see that the equity premium and its variability improve more than under REE with a higher shock variance. In particular, with the AH-A initialization, when the shock variance increases from 0.00712 to 0.04, the premium increases from 0.002 to 0.0907 under REE, whereas it increases from 0.0113 to 1.8277 under CG and from 0.154 to 0.7252 under RLS.

The more than proportional improvement of the premium when we increase the shock variance is due to the fact that in the production economy there are several effects that reinforce each other. On one hand, if the exogenous shock is more volatile, the price will be more volatile due to a more volatile term $\tau(\phi_{t-1})p_{t-1}$ and to a higher feedback of the estimate of $\phi_{z,t-1}$ into the price through a more volatile second term $h(\phi_{t-1})z_t$. In addition, this will lead to more volatile dividends that will in turn feed into the price. These effects can be substantial, and the model generates an equity premium that approximately matches the data with CG learning,

$\sigma_\varepsilon^2 = 0.04$ and the AH-A initialization. With high variance we also observe that the production economy can generate the negative correlation between current price-to-dividend ratio and future excess returns that is needed to produce some predictability results.

< TABLES 6A - 6C >

Table 6B displays the results for levels of risk aversion of $\gamma = 1$ and $\gamma = 3$. As in the endowment economy, increasing the risk aversion does lead to a higher premium in absolute value. However, the improvement of the results relative to rational expectations is again very similar to the one with $\gamma = 1$, while there is no improvement regarding the predictability of excess returns. Finally, Table 6C displays the results for persistence levels of $\rho = 0.95$, $\rho = 0.5$ and $\rho = 0.1$. As in the endowment economy, we find that the results are much less sensitive to the learning specifications with a lower shock persistence and this can be explained using similar arguments to the ones in the endowment economy.

In sum, the performance of the production economy regarding the asset return statistics can only be improved by increasing the variance of the exogenous shock. The model can generate statistics that are much closer to the data when the variance is high. Of course, any improvements come at the expense of unrealistically high values for the moments of the price-to-dividend ratio and of the real aggregate macroeconomic variables. Given this, we conclude that learning in the presence of capital accumulation does not provide a satisfactory explanation for most asset pricing puzzles.

7. 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 equity prices. The main conclusions from our results are the following. For reasonable parametrizations, (a) constant gain adaptive learning does better in generating a higher equity return, a higher equity 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.

In general, standard adaptive learning has less potential for explaining the mean excess returns in the data than for generating volatility and predictability. This is due to the fact that the average estimated coefficients from the law of motion of the equity price typically fluctuates around the rational expectations equilibrium, which is known to fail in generating a sizeable premium for reasonable parametrizations. As to the equity 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 equity 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 a higher equity return and considerably more volatility and predictability in the endowment economy, especially when it is calibrated to match the dividend behavior in the data.

We also investigate the sensitivity of the results with respect to key parameter values. We find that changing the risk aversion or the shock of the variance does not alter the relative improvements in equity price volatility and equity premium when switching from rational expectations to adaptive learning (except when we increase the variance in the production economy). However, we find that lower persistence of the exogenous shock implies less sensitivity of the results to the various characteristics of adaptive learning, such as the algorithm, initial conditions, etc.

In general, our findings suggest that tracking algorithms such as CG are more likely to explain asset pricing facts than RLS in models where there is no inherent persistence in the equity price, such as our 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 only has a chance of generating asset statistics that are closer to the data at the expense of an unreasonable behavior for the macroeconomic variables. We can moreover generally claim that more model persistence (either in the form of endogenous persistence as in the production economy or in the form of the exogenous shock persistence ρ) makes the dynamics much more sensitive to how we set up adaptive learning, i.e. learning algorithms, initial conditions, gain, etc. than when there is less persistence (e.g. endowment economy and/or low ρ).

Overall, we conclude that self-referential adaptive learning, while moving the results towards the right direction, is not sufficient to generate most of the basic asset pricing stylized facts. We conjecture however that adaptive learning may be part of a story that explains equity price dynamics, if combined with more elaborate model ingredients, such as learning on the exogenous shock process, learning on the growth of dividend or consumption, model misspecification, etc.

As a final comment, our paper also provides a contribution to the literature of general use of adaptive learning for quantitative analysis in a macroeconomic framework. In Carceles-Poveda and Giannitsarou (2007), we show how the initial conditions for a learning algorithm may or may not matter (in terms of dynamics and speed of convergence) in the context of reduced form linear or linearized models like the ones we used here. In this paper, we provide a more thorough quantitative exploration of this issue; we illustrate via our extensive sensitivity analysis that one really needs to turn to the quantitative model before concluding whether initial conditions matter a lot or not for the learning dynamics.

References

[1] Adam, K., A. Marcet and J. P. Nicolini. Learning and Stock Market Volatility. Mimeograph.

[2] Bacon, R. W. A Note on the Properties of Products of Random Variables with Reference to Economic Applications. Oxford Bulletin of Economics and Statistics 42, 4, 337-44.

[3] Boldrin, M., L. J. Christiano and J. D. M. Fisher, 2001. Habit Persistence, Asset Returns, and the Business Cycle. The American Economic Review 91, 149-166.

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

[5] Brock, W. A., 1982. Asset Prices in a Production Economy. In the "Economics of Information and Uncertainty", ed. by J. J. McCall. The University of Chicago Press.

[6] 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.

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

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

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

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

[11] Campbell, J. and M. Yogo, 2005. Efficient Tests of Stock Return Predictability. Forthcoming in Journal of Financial Economics.

[12] Carceles-Poveda, E., 2005, Idiosyncratic Shocks and Asset Returns in the RBC Model, an Approximate Analytical Approach, Macroeconomic Dynamics 9, 295-320.

[13] Carceles-Poveda, E. and C. Giannitsarou, 2007. Adaptive Learning in Practice. Journal of Economic Dynamics and Control, forthcoming.

[14] 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.

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

[16] Evans, G., S. Honkapohja and N. Williams, 2005. Generalized Stochastic Gradient. Mimeograph.

[17] Fama E. and K. French, 1988, Dividend Yields and Expected Stock Returns. Journal of Financial Economics 22, 3-27.

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

[19] Hall, R., 1988. Intertemporal Elasticity in Consumption. Journal of Political Economy 98, 339-357.

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

[21] Jaganathan, R., E. McGrattan and A. Sherbina, 2000. The Declining US Equity Premium. Federal Reserve Bank of Minneapolis Quarterly Review 24, 4, 3-19.

[22] Jermann, U. J., 1998. Asset pricing in production economies, Journal of Monetary Economics 41, 2, 257-275.

[23] Kim, Y. S., 2006. Learning, Structural Breaks, and Asset-Return Dynamics. Mimeograph.

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

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

[26] 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.

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

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

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

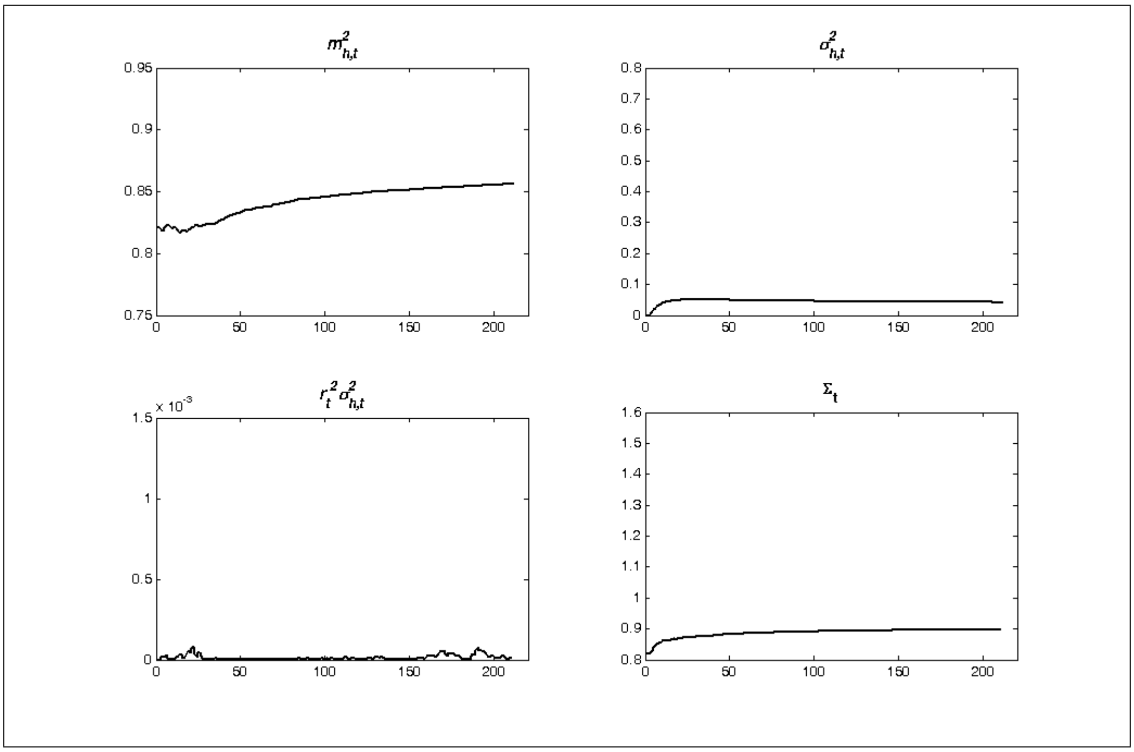

Figure 1: **Decomposition of Relative Variance** Σ_t **with RLS**. Endowment Economy, initial conditions AH-B**.** Results based on experiments of 3000 simulations of 211 periods, with σ_ε = 0.06.

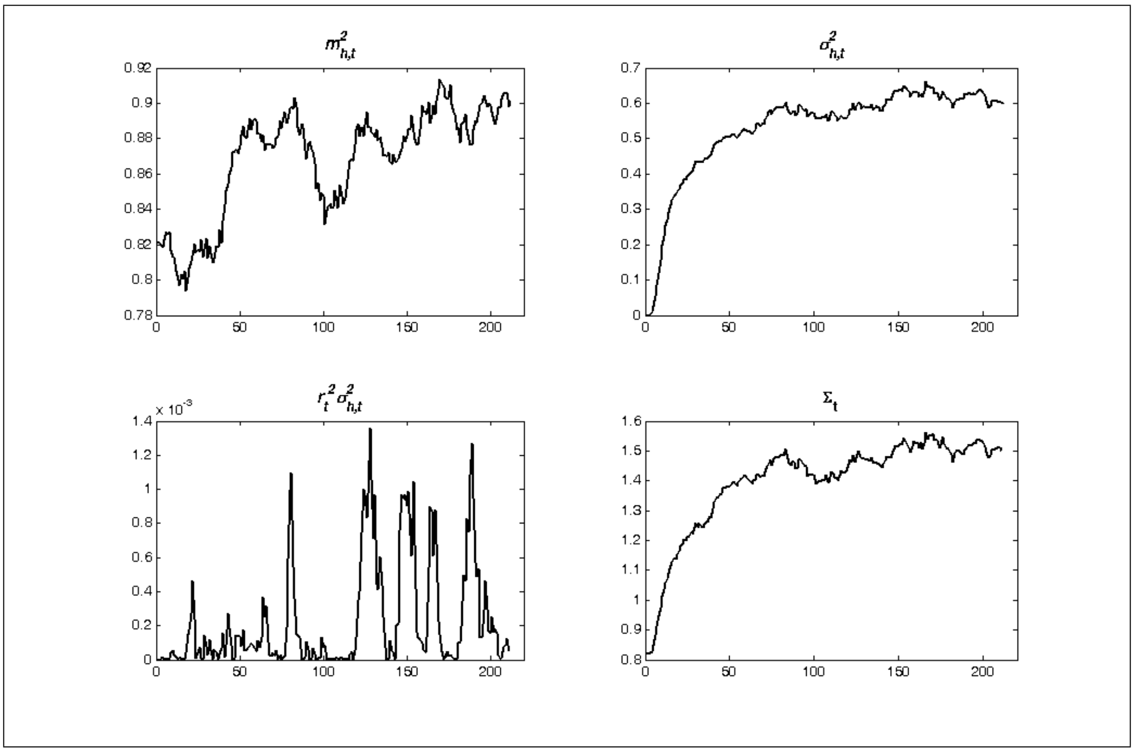

Figure 2: **Decomposition of Relative Variance** Σ_t **with CG**. Endowment Economy, initial conditions AH-B**.** Results based on experiments of 3000 simulations of 211 periods, with σ_ε = 0.06.

Asset Moments			
	Mean	*Std.*	
r^e	2.3447	7.7287	
r^f	0.2268	0.8719	
$r^e - r^f$	2.0335	7.6214	
Predictability			
Horizon	*Slope*	R^2	$t - statistic$
1	-0.0534	0.0928	-2.3317
2	-0.1076	0.2052	-2.5135
4	-0.1858	0.3687	-3.2161
Moments for P/D			
	Mean	*Std.*	*Autocor.*
	28.3117	9.0578	0.9656
Moments for Δc and Δd			
	$Std(\Delta d)$	$Std(\Delta c)$	
	6.0300	0.5362	

Table 1: Asset pricing facts 1947.2-1998.4. Standard deviations, asset returns and the risk premium are in percentage terms.

Gain	0.02	0.04	0.09	0.17	0.31	0.37	0.46	0.60	0.85
Full Decay (Quarters)	1600	800	400	200	100	80	60	40	20
Half Life Decay (Quarters)	34.30	17.00	7.34	3.72	1.86	1.50	1.12	0.75	0.36

Table 2: Gains for the CG algorithm.

	σ = 0.00712	Equity Return		Risk Free Rate		Equity Premium	
		Mean	*St. Dev.*	*Mean*	*St. Dev.*	*Mean*	*St. Dev.*
Data		**2.3447**	**7.7287**	**0.2268**	**0.8719**	**2.0335**	**7.6214**
RE		1.0128	0.7285	1.0103	0.1015	0.0026	0.7542
DIS	RLS	1.0129	0.7241	1.0103	0.1015	0.0027	0.7506
	CG, g = 0.2	1.0134	0.7639	1.0103	0.1015	0.0032	0.7921
	CG, g = 0.4	1.0149	0.8729	1.0103	0.1015	0.0047	0.9008
AH-B	RLS	1.0125	0.6705	1.0103	0.1015	0.0023	0.6977
	CG, g = 0.2	1.0131	0.7361	1.0103	0.1015	0.0029	0.7647
	CG, g = 0.4	1.0147	0.8546	1.0103	0.1015	0.0045	0.8826
AH-A	RLS	1.0130	0.7437	1.0103	0.1015	0.0029	0.7699
	CG, g = 0.2	1.0135	0.7741	1.0103	0.1015	0.0033	0.8022
	CG, g = 0.4	1.0150	0.8797	1.0103	0.1015	0.0048	0.9075

Table 3A: Endowment Economy, Statistics for Returns. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. The standard deviations, returns and the equity premia are in percentage terms.

	σ = 0.00712	sd(p_{AL})/sd(p_{RE})	Mean(P/D)	STD(P/D)	Corr(P/D)	STD(Δc)	STD(Δd)
Data			**28.3117**	**9.0578**	**0.9656**	**0.5362**	**6.0300**
RE			24.7502	0.1632	0.5378	0.7212	0.7212
DIS	RLS	**1.0002**	24.7503	0.1922	0.6431	0.7212	0.7212
	CG, g = 0.2	**1.1846**	24.7505	0.2806	0.7910	0.7212	0.7212
	CG, g = 0.4	**1.4391**	24.7576	0.4330	0.8664	0.7212	0.7212
AH-B	RLS	**0.9262**	24.7505	0.1809	0.6575	0.7212	0.7212
	CG, g = 0.2	**1.1438**	24.7506	0.2659	0.7887	0.7212	0.7212
	CG, g = 0.4	**1.4104**	24.7571	0.4198	0.8654	0.7212	0.7212
AH-A	RLS	**1.0272**	24.7507	0.1966	0.6384	0.7212	0.7212
	CG, g = 0.2	**1.1995**	24.7506	0.2861	0.7919	0.7212	0.7212
	CG, g = 0.4	**1.4495**	24.7576	0.4379	0.8666	0.7212	0.7212

Table 3B: Endowment Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. The third column in boldface letters shows the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The standard deviations are in percentage terms.

		BETAS						R-SQUARE		
	σ = 0.00712	*1 year*		*2 years*		*4 years*		*1 year*	*2 years*	*4 years*
Data		**-0.0534**		**-0.1076**		**-0.1858**		**0.0928**	**0.2052**	**0.3687**
		Aver.	**% Sig.**	**Aver.**	**% Sig.**	**Aver.**	**% Sig.**			
RE		-0.0001	7.2	-0.0001	9.5	-0.0002	13.8	0.0048	0.0036	0.0055
DIS	RLS	-0.0006	9.6	-0.0011	14.4	-0.0019	20.8	0.0077	0.0133	0.0222
	CG, g = 0.2	-0.0024	20.2	-0.0046	34.6	-0.0080	44.9	0.0218	0.0437	0.0724
	CG, g = 0.4	-0.0049	43.4	-0.0091	58.6	-0.0146	64.1	0.0562	0.0986	0.1414
AH-B	RLS	-0.0003	7.2	-0.0006	10.2	-0.0009	15.3	0.0070	0.0117	0.0186
	CG, g = 0.2	-0.0021	18.7	-0.0042	31.3	-0.0073	41.5	0.0200	0.0403	0.0667
	CG, g = 0.4	-0.0047	42.8	-0.0088	57.3	-0.0141	62.5	0.0546	0.0957	0.1367
AH-A	RLS	-0.0007	10.3	-0.0013	15.5	-0.0022	22.9	0.0080	0.0140	0.0236
	CG, g = 0.2	-0.0024	20.6	-0.0048	35.3	-0.0082	46.0	0.0225	0.0450	0.0744
	CG, g = 0.4	-0.0049	43.5	-0.0092	59.2	-0.0148	64.6	0.0568	0.0997	0.1431

Table 3C: Endowment Economy, Predictability of Excess Returns. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. Columns 3-8 show average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation, as well as the percentage of these regressions for which the estimated slope is significant.

	σ = 0.06	Equity Return		Risk Free Rate		Equity Premium	
		Mean	***St. Dev.***	***Mean***	***St. Dev.***	***Mean***	***St. Dev.***
Data		**2.3447**	**7.7287**	**0.2268**	**0.8719**	**2.0335**	**7.6214**
RE		1.1969	6.1561	1.0155	0.8560	0.1829	6.3721
DIS	RLS	1.2055	6.1216	1.0155	0.8560	0.1914	6.3445
	CG, *g* = 0.2	1.2421	6.4746	1.0155	0.8560	0.2277	6.7120
	CG, *g* = 0.4	1.3634	7.5067	1.0155	0.8560	0.3489	7.7356
AH-B	RLS	1.1770	5.6654	1.0155	0.8560	0.1629	5.8946
	CG, *g* = 0.2	1.2245	6.2368	1.0155	0.8560	0.2102	6.4771
	CG, *g* = 0.4	1.3453	7.3325	1.0155	0.8560	0.3308	7.5671
AH-A	RLS	1.2159	6.2881	1.0155	0.8560	0.2019	6.5086
	CG, *g* = 0.2	1.2486	6.5623	1.0155	0.8560	0.2343	6.7986
	CG, *g* = 0.4	1.3694	7.5635	1.0155	0.8560	0.3550	7.7957

Table 3D: Endowment Economy, Statistics for Returns. Results based on experiments of 3000 simulations of 211 periods, with $\sigma = 0.06$. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and *g* is the corresponding gain. The standard deviations, returns and the equity premia are in percentage terms.

	σ = 0.06	sd(p_{AL})/sd(p_{RE})	Mean(*P/D*)	STD(*P/D*)	Corr(*P/D*)	STD(Δc)	STD(Δd)
Data			**28.3117**	**9.0578**	**0.9656**	**0.5362**	**6.0300**
RE			24.7610	1.3751	0.5373	6.0782	6.0782
DIS	RLS	**1.0002**	24.7839	1.6237	0.6427	6.0782	6.0782
	CG, *g* = 0.2	**1.1846**	24.8861	2.4284	0.7899	6.0782	6.0782
	CG, *g* = 0.4	**1.4391**	25.3400	4.4524	0.8624	6.0782	6.0782
AH-B	RLS	**0.9262**	24.7767	1.5263	0.6572	6.0782	6.0782
	CG, *g* = 0.2	**1.1438**	24.8679	2.2931	0.7877	6.0782	6.0782
	CG, *g* = 0.4	**1.4104**	25.2917	4.2492	0.8618	6.0782	6.0782
AH-A	RLS	**1.0272**	24.7898	1.6617	0.6379	6.0782	6.0782
	CG, *g* = 0.2	**1.1995**	24.8942	2.4705	0.7907	6.0782	6.0782
	CG, *g* = 0.4	**1.4495**	25.3557	4.5201	0.8626	6.0782	6.0782

Table 3E: Endowment Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth. Results based on experiments of 3000 simulations of 211 periods, with $\sigma = 0.06$. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and *g* is the corresponding gain. The third column in boldface letters shows the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The standard deviations are in percentage terms.

		BETAS						R-SQUARE		
	σ = 0.06	*1 year*		*2 years*		*4 years*		*1 year*	*2 years*	*4 years*
Data		**-0.0534**		**-0.1076**		**-0.1858**		**0.0928**	**0.2052**	**0.3687**
		Average	**% Sig.**	**Average**	**% Sig.**	**Average**	**% Sig.**			
RE		-0.0009	7.3	-0.0010	9.6	-0.0018	13.7	0.0048	0.0056	0.0055
DIS	RLS	-0.0052	9.7	-0.0095	14.3	-0.0160	20.9	0.0077	0.0133	0.0222
	CG, *g* = 0.2	-0.0202	20.6	-0.0396	34.6	-0.0677	45.2	0.0219	0.0439	0.0726
	CG, *g* = 0.4	-0.0422	43.8	-0.0784	59.1	-0.1256	64.7	0.0572	0.1000	0.1432
AH-B	RLS	-0.0029	7.1	-0.0050	10.2	-0.0076	15.1	0.0070	0.0117	0.0186
	CG, *g* = 0.2	-0.0184	18.6	-0.0362	31.2	-0.0617	41.7	0.0201	0.0404	0.0669
	CG, *g* = 0.4	-0.0407	43.3	-0.0757	57.5	-0.1211	62.8	0.0555	0.0970	0.1383
AH-A	RLS	-0.0061	10.3	-0.0112	15.5	-0.0192	23.1	0.0080	0.0140	0.0236
	CG, *g* = 0.2	-0.0209	20.6	-0.0409	35.2	-0.0700	46.2	0.0226	0.0452	0.0747
	CG, *g* = 0.4	-0.0427	44.2	-0.0794	59.6	-0.1271	65.1	0.0578	0.1011	0.1449

Table 3F: Endowment Economy, Predictability of Excess Returns. Results based on experiments of 3000 simulations of 211 periods, with $\sigma = 0.06$. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and *g* is the corresponding gain. Columns 3-8 show average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation, as well as the percentage of these regressions for which the estimated slope is significant.

		Equity Return		Risk Free Rate		Equity Premium	
	σ = 0.00712	***Mean***	***St. Dev.***	***Mean***	***St. Dev.***	***Mean***	***St. Dev.***
Data	**Data**	**2.3447**	**7.7287**	**0.2268**	**0.8719**	**2.0335**	**7.6214**
RE	**RE**	1.0131	0.0618	1.0103	0.0564	0.0027	0.0070
DIS	RLS	1.0159	0.0686	1.0112	0.0632	0.0048	0.0089
	CG, g = 0.2	1.0182	0.0694	1.0103	0.0621	0.0078	0.0167
	CG, g = 0.4	1.0157	0.0660	1.0100	0.0570	0.0057	0.0178
AH-B	RLS	1.0114	0.0583	1.0101	0.0533	0.0013	0.0051
	CG, g = 0.2	1.0143	0.0676	1.0101	0.0617	0.0043	0.0112
	CG, g = 0.4	1.0143	0.0652	1.0099	0.0576	0.0043	0.0143
AH-A	RLS	1.0275	0.0929	1.0121	0.0863	0.0154	0.0198
	CG, g = 0.2	1.0218	0.0727	1.0105	0.0634	0.0113	0.0229
	CG, g = 0.4	1.0182	0.0686	1.0103	0.0566	0.0079	0.0239

Table 4A: Production Economy, Statistics for Returns. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. The standard deviations, returns and the equity premia are in percentage terms.

	σ = 0.00712	**sd(p_{AL})/sd(p_{RE})**	**Mean(*P/D*)**	**STD(*P/D*)**	**Corr(*P/D*)**	**STD(Δc)**	**STD(Δd)**
Data			**28.311**	**9.0578**	**0.9656**	**0.5362**	**6.0300**
RE	RE		24.810	1.7209	0.9762	0.2246	2.9272
DIS	RLS	**1.0546**	24.830	1.9020	0.9769	0.2851	2.9126
	CG, g = 0.2	**1.0317**	24.912	2.5494	0.9693	0.5460	3.7855
	CG, g = 0.4	**0.7139**	24.864	2.1178	0.9499	0.7837	4.3781
AH-B	RLS	**0.3877**	24.779	1.0914	0.9795	0.4309	1.6975
	CG, g = 0.2	**0.8467**	24.859	1.8321	0.9698	0.5550	2.7111
	CG, g = 0.4	**0.6515**	24.836	1.8223	0.9528	0.7540	3.7712
AH-A	RLS	**1.6990**	25.056	3.4067	0.9824	0.1832	4.4266
	CG, g = 0.2	**1.1575**	24.986	3.1258	0.9691	0.5700	4.6448
	CG, g = 0.4	**0.7804**	24.887	2.5231	0.9488	0.8368	5.2190

Table 4B: Production Economy, Statistics for Price Dividend Ratio, Consumption and Dividend Growth. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. The third column in boldface letters shows the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The standard deviations are in percentage terms.

σ = 0.00712				BETAS				R-SQUARE		
		1 year		*2 years*		*4 years*		*1 year*	*2 years*	*4 years*
Data	**Data**	**-0.0534**		**-0.1076**		**-0.1858**		**0.0928**	**0.2052**	**0.3687**
		Aver.	**% Sig.**	**Aver.**	**% Sig.**	**Aver.**	**% Sig.**			
RE	RE	0.0001	0.0	0.0003	0.0	0.0004	0.0	0.4768	0.3898	0.2618
DIS	RLS	0.0001	4.4	0.0002	6.3	0.0003	5.7	0.3794	0.3072	0.2086
	CG, g = 0.2	0.0000	6.3	0.0001	6.3	0.0002	8	0.1824	0.1537	0.1196
	CG, g = 0.4	0.0000	6.6	0.0000	7.3	0.0001	9.6	0.0984	0.0880	0.0780
AH-B	RLS	0.0000	11.1	0.0000	6.4	0.0000	4.6	0.0946	0.0927	0.0945
	CG, g = 0.2	0.0000	6.3	0.0001	6.4	0.0002	7.9	0.1835	0.1524	0.1192
	CG, g = 0.4	0.0000	9	0.0000	9.2	0.0001	9.6	0.1005	0.0871	0.0752
AH-A	RLS	0.0001	16.8	0.0002	16.4	0.0002	18.4	0.3053	0.2639	0.2024
	CG, g = 0.2	0.0000	6.9	0.0001	7.4	0.0001	9.4	0.1664	0.1426	0.1144
	CG, g = 0.4	0.0000	2.9	0.0000	4.6	0.0000	6.8	0.0962	0.0888	0.0817

Table 4C: Production Economy, Predictability of Excess Returns. Results based on experiments of 3000 simulations of 211 periods, with σ = 0.00712. RE stands for rational expectations. DIS is for simulations with an initial condition for learning that is drawn from an appropriate distribution around the REE. AH-B is for simulations with an initial condition for learning that is below the REE (ϕ_0 = 0.9*REE) and AH-A is for simulations with an initial condition for learning that is above the REE (ϕ_0 = 1.035*REE). RLS is for recursive least squares, CG stands for constant gain and g is the corresponding gain. Columns 3-8 show average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation, as well as the percentage of these regressions for which the estimated slope is significant.

		$sd(p_{AL})/sd(p_{RE})$	**m(ep)**	**std(ep)**	**m(P/D)**	**std(P/D)**	β **(1 year)**	β **(2 years)**	β **(4 years)**
	DATA		**2.0335**	**7.6214**	**28.3117**	**9.0578**	**-0.0534**	**-0.1076**	**-0.1858**
γ **= 1**	REE		0.1829	6.3721	24.7610	1.3751	-0.0009	-0.0010	-0.0018
	AH-B, RLS	0.9262	0.1629	5.8946	24.7767	1.5263	-0.0029	-0.0050	-0.0076
	AH-B, CG	1.1438	0.2102	6.4771	24.8679	2.2931	-0.0184	-0.0362	-0.0617
	AH-A, RLS	1.0272	0.2019	6.5086	24.7898	1.6617	-0.0061	-0.0112	-0.0192
	AH-A, CG	1.1995	0.2343	6.7986	24.8942	2.4705	-0.0209	-0.0409	-0.0700
γ **= 3**	REE		1.3531	17.3536	26.1659	8.1821	-0.0306	-0.0584	-0.1088
	AH-B, RLS	0.9262	1.2050	16.0521	26.1598	7.6641	-0.0233	-0.0466	-0.0915
	AH-B, CG	1.1438	1.6284	17.9733	28.6879	15.3138	-0.0553	-0.1131	-0.2078
	AH-A, RLS	1.0272	1.5065	17.7966	26.6942	9.2123	-0.0365	-0.0702	-0.1298
	AH-A, CG	1.1995	1.8241	18.9458	29.3614	17.2583	-0.0639	-0.1283	-0.2319

Table 5A. Endowment economy, sensitivity analysis with respect to risk aversion. Results based on experiments of 3000 simulations of 211 periods. The abbreviations in the second column are the same as in tables 3 and 4. For CG learning we used $g = 0.2$. The third column gives the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The fourth and fifth column show mean and st. deviation of the equity premium. The sixth and seventh columns give the mean and standard deviation of the price dividend ratio. The last three columns give estimated average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation. Standard deviations, returns and the equity premia are in percentage terms.

		$sd(p_{AL})/sd(p_{RE})$	**m(ep)**	**std(ep)**	**m(P/D)**	**std(P/D)**	β **(1 year)**	β **(2 years)**	β **(4 years)**
	DATA		**2.0335**	**7.6214**	**28.3117**	**9.0578**	**-0.0534**	**-0.1076**	**-0.1858**
ρ **= 0.95**	REE		0.1829	6.3721	24.7610	1.3751	-0.0009	-0.0010	-0.0018
	AH-B, RLS	0.9262	0.1629	5.8946	24.7767	1.5263	-0.0029	-0.0050	-0.0076
	AH-B, CG	1.1438	0.2102	6.4771	24.8679	2.2931	-0.0184	-0.0362	-0.0617
	AH-A, RLS	1.0272	0.2019	6.5086	24.7898	1.6617	-0.0061	-0.0112	-0.0192
	AH-A, CG	1.1995	0.2343	6.7986	24.8942	2.4705	-0.0209	-0.0409	-0.0700
ρ **= 0.5**	REE		0.1835	9.2813	24.7556	1.3027	-0.0230	-0.0241	-0.0243
	AH-B, RLS	0.9835	0.1844	9.1968	24.7564	1.2996	-0.0230	-0.0242	-0.0241
	AH-B, CG	1.0520	0.1933	9.3679	24.7581	1.3449	-0.0286	-0.0319	-0.0329
	AH-A, RLS	0.9944	0.1899	9.2645	24.7573	1.3126	-0.0237	-0.0250	-0.0249
	AH-A, CG	1.0560	0.1948	9.3876	24.7584	1.3489	-0.0289	-0.0321	-0.0332
ρ **= 0.1**	REE		0.1835	12.5717	24.7510	1.2792	-0.0461	-0.0458	-0.0458
	AH-B, RLS	1.0000	0.1876	12.5858	24.7515	1.2847	-0.0464	-0.0462	-0.0461
	AH-B, CG	1.0350	0.2107	12.7812	24.7544	1.3428	-0.0492	-0.0491	-0.0492
	AH-A, RLS	1.0012	0.1879	12.5897	24.7516	1.2853	-0.0464	-0.0462	-0.0462
	AH-A, CG	1.0355	0.2109	12.7833	24.7544	1.3432	-0.0492	-0.0491	-0.0492

Table 5B. Endowment economy, sensitivity analysis with respect to exogenous shock persistence. Results based on experiments of 3000 simulations of 211 periods. The abbreviations in the second column are the same as in tables 3 and 4. For CG learning we used $g = 0.2$. The third column gives the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The fourth and fifth column show mean and st. deviation of the equity premium. The sixth and seventh columns give the mean and standard deviation of the price dividend ratio. The last three columns give estimated average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation. Standard deviations, returns and the equity premia are in percentage terms.

		$sd(p_{AL})/sd(p_{RE})$	m(ep)	std(ep)	m(P/D)	std(P/D)	β(1 year)	β(2 years)	β(4 years)
	DATA		**2.0335**	**7.6214**	**28.3117**	**9.0578**	**-0.0534**	**-0.1076**	**-0.1858**
σ=0.007	REE		0.0027	0.0069	24.810	1.7209	0.0001	0.0030	0.0004
	AH-B, RLS	0.3877	0.0013	0.0051	24.779	1.0914	0.0000	0.0000	0.0000
	AH-B, CG	0.8467	0.0043	0.0112	24.859	1.8321	0.0000	0.0000	0.0002
	AH-A, RLS	1.6990	0.0154	0.0198	25.056	3.4067	0.0001	0.0001	0.0001
	AH-A, CG	1.1575	0.0113	0.0229	24.989	0.1738	0.0000	0.0001	0.0001
σ = 0.02	REE		0.0219	0.0320	25.218	4.9448	0.0004	0.0006	0.0010
	AH-B, RLS	0.3877	0.0104	0.0196	24.973	3.0994	0.0000	0.0000	0.0001
	AH-B, CG	0.8467	0.0374	0.0888	25.580	5.7938	-0.0004	-0.0006	-0.0005
	AH-A, RLS	1.6989	0.1328	0.1710	27.575	10.974	-0.0013	-0.0025	-0.0042
	AH-A, CG	1.1575	0.1050	0.2546	26.841	10.998	-0.0019	-0.0028	-0.0032
σ = 0.04	REE		0.0907	0.1239	26.670	10.670	0.0000	0.0001	0.0004
	AH-B, RLS	0.3877	0.0425	0.0644	25.648	6.4406	-0.0002	-0.0002	-0.0002
	AH-B, CG	0.8467	0.2383	0.9168	29.209	20.430	-0.0070	-0.0108	-0.0123
	AH-A, RLS	1.6989	0.7252	1.1820	39.624	35.961	-0.0177	-0.0318	-0.0512
	AH-A, CG	1.1575	1.8277	11.160	40.259	71.014	-0.0235	-0.03505	-0.0405

Table 6A. Production economy, sensitivity analysis with respect to innovation variance. Results based on experiments of 3000 simulations of 211 periods. The abbreviations in the second column are the same as in tables 3 and 4. For CG learning we used g = 0.2. The third column gives the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The fourth and fifth column show mean and st. deviation of the equity premium. The sixth and seventh columns give the mean and standard deviation of the price dividend ratio. The last three columns give estimated average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation. Standard deviations, returns and the equity premia are in percentage terms.

		$sd(p_{AL})/sd(p_{RE})$	m(ep)	std(ep)	m(P/D)	std(P/D)	β(1 year)	β(2 years)	β(4 years)
	DATA		**2.0335**	**7.6214**	**28.3117**	**9.0578**	**-0.0534**	**-0.1076**	**-0.1858**
γ **= 1**	REE		0.0027	0.0070	24.810	1.7209	0.0001	0.0003	0.0004
	AH-B, RLS	0.3877	0.0013	0.0051	24.779	1.0914	0.0000	0.0000	0.0000
	AH-B, CG	0.8467	0.0043	0.0112	24.859	1.8321	0.0000	0.0001	0.0002
	AH-A, RLS	1.6990	0.0154	0.0198	25.056	3.4067	0.0001	0.0002	0.0002
	AH-A, CG	1.1575	0.0113	0.0229	24.986	3.1258	0.0000	0.0001	0.0001
γ **= 3**	REE		0.0042	0.0079	24.8424	1.9218	0.0001	0.0002	0.0004
	AH-B, RLS	0.3124	0.0011	0.0051	24.7780	1.0389	0.0000	0.0000	0.0000
	AH-B, CG	0.6236	0.0032	0.0090	24.8378	1.5022	0.0000	0.0001	0.0001
	AH-A, RLS	1.4979	0.0198	0.0221	25.1751	3.2769	0.0000	0.0000	0.0000
	AH-A, CG	0.9463	0.0107	0.0208	24.9759	2.8686	0.0000	0.0001	0.0001

Table 6B. Production economy, sensitivity analysis with respect to risk aversion. Results based on experiments of 3000 simulations of 211 periods. The abbreviations in the second column are the same as in tables 3 and 4. For CG learning we used g = 0.2. The third column gives the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The fourth and fifth column show mean and st. deviation of the equity premium. The sixth and seventh columns give the mean and standard deviation of the price dividend ratio. The last three columns give estimated average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation. Standard deviations, returns and the equity premia are in percentage terms.

		sd(p_{AL})/sd(p_{RE})	m(ep)	std(ep)	m(P/D)	std(P/D)	β(1 year)	β(2 years)	β(4 years)
	DATA		**2.0335**	**7.6214**	**28.3117**	**9.0578**	**-0.0534**	**-0.1076**	**-0.1858**
ρ **= 0.95**	REE		0.0027	0.0070	24.810	1.7209	0.0001	0.0003	0.0004
	AH-B, RLS	0.3877	0.0013	0.0051	24.779	1.0914	0.0000	0.0000	0.0000
	AH-B, CG	0.8467	0.0043	0.0112	24.859	1.8321	0.0000	0.0001	0.0002
	AH-A, RLS	1.6990	0.0154	0.0198	25.056	3.4067	0.0001	0.0002	0.0002
	AH-A, CG	1.1575	0.0113	0.0229	24.986	3.1258	0.0000	0.0001	0.0001
ρ **= 0.5**	REE		0.0010	0.0159	24.7505	0.8164	0.0000	0.0000	0.0000
	AH-B, RLS	0.6508	0.0007	0.0157	24.7465	0.6435	0.0000	0.0000	0.0000
	AH-B, CG	0.6289	0.0007	0.0156	24.7495	0.6618	0.0000	0.0000	0.0000
	AH-A, RLS	0.9087	0.0011	0.0159	24.7476	0.8156	0.0000	0.0000	0.0000
	AH-A, CG	0.7036	0.0008	0.0157	24.7499	0.7237	0.0000	0.0000	0.0000
ρ **= 0.1**	REE		0.0008	0.0256	24.7412	0.5541	0.0000	0.0000	0.0000
	AH-B, RLS	0.6351	0.0007	0.0253	24.7401	0.4627	0.0000	0.0000	0.0000
	AH-B, CG	0.5112	0.0005	0.0237	24.7428	0.4095	0.0000	0.0000	0.0000
	AH-A, RLS	0.7556	0.0007	0.0239	24.7401	0.5030	0.0000	0.0000	0.0000
	AH-A, CG	0.5601	0.0006	0.0237	24.7424	0.4260	0.0000	0.0000	0.0000

Table 6C. Production economy, sensitivity analysis with respect to exogenous shock persistence. Results based on experiments of 3000 simulations of 211 periods. The abbreviations in the second column are the same as in tables 3 and 4. For CG learning we used g = 0.2. The third column gives the ratio of the st. deviation of the equity price under learning over the st. deviation of the price under RE. The fourth and fifth column show mean and st. deviation of the equity premium. The sixth and seventh columns give the mean and standard deviation of the price dividend ratio. The last three columns give estimated average slopes from regressions of 1, 2, or 4 year ahead excess returns on the current log(P/D), divided by its standard deviation. Standard deviations, returns and the equity premia are in percentage terms.