

Stock Market Volatility and Learning*

Klaus Adam Albert Marcet Juan Pablo Nicolini

August 14, 2007

Abstract

Introducing bounded rationality into a standard consumption based asset pricing model with a representative agent and time separable preferences strongly improves empirical performance. Learning causes momentum and mean reversion of returns and thereby excess volatility, persistence of price-dividend ratios, long-horizon return predictability and a risk premium, as in the habit model of Campbell and Cochrane (1999), but for lower risk aversion. This is obtained, although we restrict consideration to learning schemes that imply only small deviations from full rationality. The findings are robust to the particular learning rule used and the value chosen for the single free parameter introduced by learning, provided agents forecast future stock prices using past information on prices.

JEL Class. No.: G12, D84

*Thanks go to Luca Dedola and Jaume Ventura for interesting comments and suggestions. We particularly thank Phillip Weil for a very interesting discussion. Davide Debortoli has supported us with outstanding research assistance. Marcet acknowledges support from CIRIT (Generalitat de Catalunya), DGES (Ministry of Education and Science), CREI, the Barcelona Economics program of XREA and the Wim Duisenberg fellowship from the European Central Bank. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the European Central Bank. Author contacts: Klaus Adam (European Central Bank and CEPR) klaus.adam@ecb.int; Albert Marcet (Institut d'Anàlisi Econòmica CSIC, Universitat Pompeu Fabra) albert.marcet@upf.edu; Juan Pablo Nicolini (Universidad Torcuato di Tella) juanpa@utdt.edu.

"Investors, their confidence and expectations buoyed by past price increases, bid up speculative prices further, thereby enticing more investors to do the same, so that the cycle repeats again and again, .. "

Irrational Exuberance, Shiller (2005, p.56)

1 Introduction

The purpose of this paper is to show that a very simple asset pricing model is able to reproduce a variety of stylized facts if one allows for small departures from rationality. This result is somehow remarkable, since the literature in empirical finance had great difficulties in developing dynamic equilibrium rational expectations models accounting for all the facts we consider.

Our model is based on the representative agent time-separable utility endowment economy developed by Lucas (1978). It is well known that the implications of this model under rational expectations are at odds with basic asset pricing observations: the price dividend ratio is too volatile and persistent, stock returns are too volatile and should not be negatively related to the price dividend ratio in the long run, and the risk premium is too high. Our learning model introduces just one additional free parameter into Lucas' framework and quantitatively accounts for all these observations. Since the learning model reduces to the rational expectations model if the additional parameter is set to zero and since this parameter is close to zero throughout the paper, we consider the learning model to represent only a small departure from rationality. Nevertheless, the behavior of equilibrium prices differs considerably from the one obtained under rational expectations, implying that the asset pricing implications of the standard model are not robust to small departures from rationality. As we document, this non-robustness is empirically encouraging, i.e., the model matches the data much better if this small departure from rationality is allowed for.

A very large body of literature has documented that stock prices exhibit movements that are very hard to reproduce within the realm of rational expectations and Lucas' tree model has been extended in a variety of directions to improve its empirical performance. After several years of research, Campbell and Cochrane (1999) succeeded in reproducing all the facts, albeit at the cost of imposing complicated non-time-separabilities in preferences and high effective degrees of risk aversion. Our model retains simplicity and moderate curvature in utility, but instead deviates from full rationality.

The behavioral finance literature tried to understand the decision making process of individual investors by means of surveys, experiments and micro evidence, exploring the intersection between economics and psychology. One of the main themes of this literature was to test the rationality hypothesis in asset markets, see Shiller (2005) for a non-technical summary. We borrow some

of the economic intuition from this literature, but follow a different modeling approach: we aim for a model that is as close as possible to the original Lucas tree model, with agents who are quasi-rational and formulate forecasts using statistical models that imply only small departures from rationality.

In the baseline learning model, we assume agents form their expectations regarding future stock prices with the most standard scheme used in the literature: ordinary least squares learning (OLS).¹ This rule has the property that in the long run the equilibrium converges to rational expectations, but in the model this process takes a very long time, and the dynamics generated by learning along the transition cause prices to be very different from the rational expectations (RE) prices. This difference occurs for the following reasons: if expectations about stock price growth have increased, the actual growth rate of prices has a tendency to increase beyond the fundamental growth rate, thereby reinforcing the initial belief of higher stock price growth. Learning thus imparts ‘momentum’ on stock prices and beliefs and produces large and sustained deviations of the price dividend ratio from its mean, as can be observed in the data. The model thus displays something like the ‘naturally occurring Ponzi schemes’ described in Shiller’s opening quote above.

As we mentioned, OLS is the most standard assumption to model the evolution of expectations functions in the learning literature and its limiting properties have been used extensively as a stability criterion to justify or discard RE equilibria. Yet, models of learning are still not commonly used to explain data or for policy analysis.² It is still the standard view in the economics research literature that models of learning introduce too many degrees of freedom, so that it is easy to find a learning scheme that matches whatever observation one desires. One can deal with this important methodological issue in two ways: first, by using a learning scheme with as few free parameters as possible, and second, by imposing restrictions on the parameters of the learning scheme to only allow for small departures of rationality.³ These considerations prompted us to use an off-the-shelf learning scheme (OLS) that has only one free parameter. In addition, in the model at hand, OLS is the best estimator in the long run, and to make the departure from rationality during the transition small, we assume that initial beliefs are at the rational expectations equilibrium, and that agents have very strong - but less than complete - confidence in these initial beliefs, as we explain in detail in the main text.

Models of learning have been used before to explain some aspects of asset price behavior. Timmermann (1993, 1996), Brennan and Xia (2001) and Cogley and Sargent (2006) show that Bayesian learning can help explain various aspects of stock prices. These authors assume that agents learn about the dividend

¹We show that results are robust to using other standard learning rules.

²We will mention some exceptions along the paper.

³Marcet and Nicolini (2005) dealt with this issue by imposing bounds on the size of the mistakes agents can make in equilibrium. These bounds imposed discipline both on the type of learning rule and on the exact value of the parameters in the learning rule. For the present model we show that results are very robust to both the learning rule and the exact value of the single learning parameter.

process and use the Bayesian posterior on the dividend process to estimate the discounted sum of dividends that would determine the stock price under RE. Therefore, while the beliefs of agents influence the market outcomes, agents' beliefs remain unaffected by market outcomes because agents learn only about an exogenous dividend process. In the language of stochastic control, these models are not self-referential. In the language of Shiller, these models can not give rise to 'naturally occurring Ponzi schemes'. In contrast, we largely abstract from learning about the dividend process and consider learning on the future stock price using past observations of price, so that beliefs and prices are mutually determined. It is precisely the learning about future stock price growth and its self-referential nature that imparts the momentum to expectations and, therefore, is key in explaining the data.⁴

Other related papers by Bullard and Duffy (2001) and Brock and Hommes (1998) show that learning dynamics can converge to complicated attractors, if the RE equilibrium is unstable under learning dynamics.⁵ Branch and Evans (2006) study a model where agents' expectations switch depending on which of the available forecast models is performing best. By comparison, we look at learning about the stock price growth rate, we address more closely the data, and we do so in a model where the rational expectations equilibrium is stable under learning dynamics, so the departure from RE behavior occurs only along a transition related to the sample size of the observed data. Also related is Cárceles-Poveda and Giannitsarou (2006) who assume that agents know the mean stock price and learn only about the deviations from the mean; they find that the presence of learning does then not significantly alter the behavior of asset prices.⁶

The paper is organized as follows. Section 2 presents the stylized facts we focus on and the basic features of the underlying asset pricing model, showing that this model cannot explain the facts under the rational expectations hypothesis. In section 3 we take the simplest risk neutral model and assume instead that agents learn to forecast the growth rate of prices. We show that such a model can *qualitatively* deliver all the considered asset pricing facts and that learning converges to rational expectations. We also explain how the deviations from rational expectations can be made arbitrarily small. In Section 4 we present the baseline learning model with risk aversion and the baseline calibration procedure. We also explain why we choose to calibrate the model parameters in a slightly different way than in standard calibration exercises. Section 5 shows that the baseline model can *quantitatively* reproduce all the facts discussed in section 2. The robustness of our findings to various assumptions about the model, the learning rule, or the calibration procedure is illustrated in section 6.

⁴Timmerman (1996) analyzes in one setting also a self-referential model, but one in which agents use dividends to predict future price. He finds that this form of self-referential learning delivers lower volatility than settings with learning about the dividend process. It is thus crucial for our results that agents use information on past price behavior to predict future price.

⁵Stability under learning dynamics is defined in Marcet and Sargent (1989).

⁶Cecchetti, Lam, and Mark (2000) determine the misspecification in beliefs about future consumption growth required to match the equity premium and other moments of asset prices.

Readers interested in obtaining a glimpse of the quantitative performance of the baseline learning model may - after reading section 2 - directly jump to table 4 in section 5.

2 Facts

This section describes stylized facts of U.S. stock price data and explains why it proved difficult to reproduce them using standard rational expectations models. The facts presented in this section have been extensively documented in the literature. We reproduce them here as a point of reference for our quantitative exercise in the latter part of the paper and using a single and updated data set.⁷

It is useful to start looking at the data through the lens of a simple dynamic stochastic endowment economy. Let D_t be the dividend of an inelastically supplied asset in period t , evolving according to

$$\frac{D_t}{D_{t-1}} = a\varepsilon_t \quad (1)$$

where $\log \varepsilon_t \sim N(-\frac{s^2}{2}, s^2)$ is *i.i.d.* and $a \geq 1$ the expected growth rate of dividends.⁸ Let the preferences of a representative consumer-investor be given by

$$E_0 \sum_{t=0}^{\infty} \delta^t U(C_t)$$

where C_t is consumption at time t , δ the discount factor and $U(\cdot)$ strictly increasing and concave. With S_t denoting the end-of-period t stock holdings, the budget constraint is

$$P_t S_t + C_t = (P_t + D_t) S_{t-1},$$

where P_t is the real price of the asset. In an equilibrium with rational expectations, the asset price must satisfy the consumer's first order condition evaluated at $C_t = D_t$

$$P_t = \delta E_t \left[\frac{U(D_{t+1})}{U(D_t)} (P_{t+1} + D_{t+1}) \right] \quad (2)$$

which defines a mapping from the exogenous dividend process to the stochastic process of prices.⁹ The nature of this mapping obviously depends on the way the intertemporal marginal rate of substitution moves with consumption. For instance, in the standard case of power preferences

$$U(C_t) = \frac{(C_t)^{1-\sigma} - 1}{1-\sigma}$$

⁷Details on the underlying data sources are provided in Appendix A.1.

⁸As documented in Mankiw, Romer and Shapiro (1985) and Campbell (2003), this is a reasonable first approximation to the empirical behavior of quarterly dividends in the U.S. It is also the standard assumption in the literature.

⁹In the data, consumption is much less volatile than dividends. This raises important issues that will be discussed later in the paper.

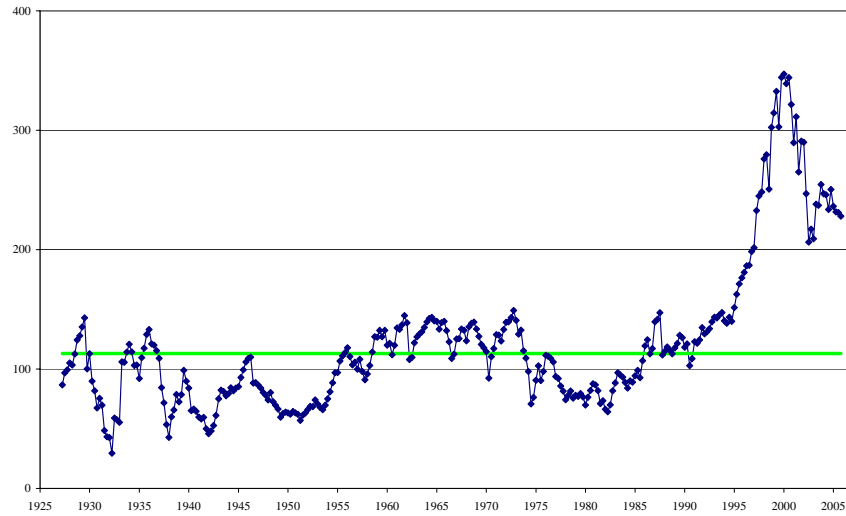


Figure 1: Quarterly U.S. price dividend ratio 1927:1-2005:4

equation (2) becomes

$$P_t = \delta E_t \left[\left(\frac{D_t}{D_{t+1}} \right)^\sigma (P_{t+1} + D_{t+1}) \right] \quad (3)$$

and the unique non-explosive price process is given by

$$P_t = \frac{\delta\beta}{1 - \delta\beta} D_t \quad (4)$$

where

$$\beta = a^{1-\sigma} e^{-\sigma(1-\sigma)\frac{\sigma^2}{2}}$$

The model then implies that the price dividend (PD) ratio is constant over time and states. Figure 1 confronts this prediction with the actual evolution of the quarterly price dividend ratio in the U.S.¹⁰ Compared to the simple model we just described, the PD ratio exhibits rather large fluctuations around its sample mean (the horizontal line in the graph). For example, the PD ratio takes on values below 30 in the year 1932 and values close to 350 in the year 2000. This large discrepancy between the prediction of the basic model and the data is also illustrated in table 1, which shows that the standard deviation of the PD ratio (σ_{PD}) is almost one half of its sample mean ($E(PD)$). We have the following asset pricing fact:

¹⁰Throughout the paper we follow Campbell (2003) and account for seasonalities in dividend payments by averaging actual payments over the last 4 quarters.

- **Fact 1: The PD ratio is very volatile.**

It follows from equation (2) that matching the observed volatility of the PD ratio under rational expectation requires alternative preference specifications. Indeed, maintaining the assumptions of *i.i.d.* dividend growth and of a representative agent, the behavior of the marginal rate of substitution is the only degree of freedom left to the theorist. This explains the development of a large and interesting literature exploring non-time-separability in consumption or consumption habits. Introducing habit amounts to consider consumers whose preferences are given by

$$E_0 \sum_{t=0}^{\infty} \delta^t \frac{(\mathbb{C}_t)^{1-\sigma} - 1}{1-\sigma},$$

where $\mathbb{C}_t = H(C_t, C_{t-1}, C_{t-2}, \dots)$ is a function of current and past consumption.¹¹ A simple habit model has been studied by Abel (1990) who assumes

$$\mathbb{C}_t = \frac{C_t}{C_{t-1}^\kappa}$$

with $\kappa \in (0, 1)$.¹² In this case, the asset price under rational expectations is

$$\frac{P_t}{D_t} = A (a\varepsilon_t)^{\kappa(\sigma-1)} \quad (5)$$

for some constant A , which shows that this model can give rise to a volatile PD ratio. Yet, with ε_t being *i.i.d.* the PD ratio will display no autocorrelation, which is in stark contrast to the empirical evidence. As figure 1 illustrates, the PD ratio displays rather persistent deviations from its sample mean. Indeed, as table 1 shows, the quarterly autocorrelation of the PD ratio ($\rho_{PD_t, PD_{t-1}}$) is very high. Therefore, this is the second fact we focus on:

- **Fact 2. The PD ratio is persistent.**

The previous observations suggest that matching the volatility and persistence of the PD ratio under rational expectations would require preferences that give rise to a volatile and persistent marginal rate of substitution. This is the avenue pursued in Campbell and Cochrane (1999) who engineer preferences that can match the behavior of the PD ratio we observe in Figure 1. Their specification also helps in replicating the asset pricing facts mentioned later in this section, as well as other facts not mentioned here.¹³ Their solution requires, however, imposing a very high degree of relative risk aversion and relies on a rather complicated structure for the habit function $H(\cdot)$.¹⁴

¹¹We keep power utility for expositional purposes only.

¹²Importantly, the main purpose of Abel's model was to generate an 'equity premium' - a fact we discuss below - not to reproduce the behavior of the price dividend ratio.

¹³They also match the pro-cyclical variation of stock prices and the counter-cyclical variation of stock market volatility. We have not explored conditional moments in our learning model, see also the discussion at the end of this section.

¹⁴The coefficient of relative risk aversion is 35 in steady state and higher still in states with 'low surplus consumption ratios'.

In our model we maintain the assumption of standard time-separable consumption preferences with moderate degrees of risk aversion. Instead, we relax the rational expectations assumption by replacing the mathematical expectation in equation (2) by the most standard learning algorithm used in the literature. Persistence and volatility of the price dividend ratio will then be the result of adjustments in beliefs that are induced by the learning process.

Before getting into the details of our model, we want to mention three additional asset pricing facts about stock returns. These facts have received considerable attention in the literature and are qualitatively related to the behavior of the PD ratio, as we discuss below.

- **Fact 3. Stock returns are ‘excessively’ volatile.**

Starting with the work of Shiller (1981) and LeRoy and Porter (1981) it has become well-known that the volatility of the PD ratio is largely the result of large changes in stock prices. Related to this is the observation that the volatility of stock returns (σ_{r^s}) in the data is much higher than the volatility of dividend growth ($\sigma_{\Delta D/D}$), see table 1.¹⁵ The observed amount of return volatility has been called ‘excessive’ mainly because the rational expectations model with time separable preferences predicts approximately identical volatilities. To see this, let r_t^s denote the stock return

$$r_t^s = \frac{P_t + D_t - P_{t-1}}{P_{t-1}} = \left[\frac{\frac{P_t}{D_t} + 1}{\frac{P_{t-1}}{D_{t-1}}} \right] \frac{D_t}{D_{t-1}} - 1 \quad (6)$$

and note that with time-separable preferences and *i.i.d.* dividend growth, the PD ratio is constant and the term in the square brackets above approximately equal to one.

From equation (6) follows that excessive return volatility is *qualitatively* related to Fact 1 discussed above, as return volatility depends partly on the volatility of the PD ratio.¹⁶ Yet, *quantitatively* return volatility also depends on the volatility of dividend growth and - up to a linear approximation - on the first two moments of the cross-correlogram between the PD ratio and the rate of growth of dividends. Since the main contribution of the paper is to show the ability of the learning model to account for the *quantitative* properties of the data, we treat the volatility of returns as a separate asset pricing fact.

¹⁵This is not due to accounting for seasonalities in dividends by taking averages: stock returns are also about three times as volatile as dividend growth at yearly frequency.

¹⁶Cochrane (2005) provides a detailed derivation of the qualitative relationship between facts 3 and 1 for *i.i.d.* dividend growth.

U.S. asset pricing facts, 1927:2-2000:4 (quarterly real values, growth rates & returns in percentage terms)			
Fact 1	Volatility of PD ratio	$E(PD)$ σ_{PD}	113.20 52.98
Fact 2	Persistence of PD ratio	$\rho_{PD_t, PD_{t-1}}$	0.92
Fact 3	Excessive return volatility	σ_{r^s} $\sigma_{\frac{\Delta D}{D}}$	11.65 2.98
Fact 4	Excess return predictability	c_2^5 R_5^2	-0.0048 0.1986
Fact 5	Equity premium	$E[r^s]$ $E[r^b]$	2.41 0.18

Table 1: Stylized asset pricing facts

- **Fact 4. Excess stock returns are predictable over the long-run.**

While stock returns are difficult to predict in general, the PD ratio is negatively related to future excess stock returns in the long run. This is illustrated in Table 2, which shows the results of regressing future cumulated excess returns over different horizons on today's price dividend ratio.¹⁷ The absolute value of the parameter estimate and the R^2 both increase with the horizon. As argued in Cochrane (2005, chapter 20), the presence of return predictability and the increase in the R^2 with the prediction horizon are *qualitatively* a joint consequence of Fact 2 and *i.i.d.* dividend growth. Nevertheless, we keep excess return predictability as an independent result, since we are again interested in the *quantitative* model implications. Yet, Cochrane also shows that the absolute value of the regression parameter increases approximately linearly with the prediction horizon, which is a quantitative result. For this reason, we summarize the return predictability evidence using the regression outcome for a single prediction horizon. We choose to include the 5 year horizon in table 1.¹⁸

¹⁷More precisely, the table reports results from OLS estimation of

$$X_{t,t+s} = c_1^s + c_2^s PD_t + u_t^s$$

for $s = 1, 3, 5, 10$, where $X_{t,t+s}$ is the observed real excess return of stocks over bonds between t and $t + s$. The second column of Table 2 reports estimates of c_2^s .

¹⁸We also used longer and shorter horizons. This did not affect our findings.

Years	Coefficient on PD, c_1^s	R^2
1	-0.0008	0.0438
3	-0.0023	0.1196
5	-0.0048	0.1986
10	-0.0219	0.3285

Table 2: Excess stock return predictability

- **Fact 5. The equity premium puzzle.**

Finally, and even though the emphasis of our paper is on moments of the PD ratio and stock returns, it is interesting to note that learning also improves the ability of the standard model to match the equity premium puzzle, i.e., the observation that stock returns - averaged over long time spans and measured in real terms - tend to be high relative to short-term real bond returns. The equity premium puzzle is illustrated in table 1, which shows the average quarterly real return on bonds ($E(r_t^b)$) being much lower than the corresponding return on stocks ($E(r_t^s)$).

Unlike Campbell and Cochrane (1999) we do not include in our list of facts any correlation between stock market data and real variables like consumption or investment. In this sense, we follow more closely the literature in finance. In our model, it is the learning scheme that delivers the movement in stock prices, even in a model with risk neutrality in which the marginal rate of substitution is constant. This contrasts with the habit literature where the movement of stock prices is obtained by modeling the way the observed stochastic process for consumption generates movements in the marginal rate of substitution. The latter explains why the habit literature focuses on the relationship between particularly low values of consumption and low stock prices. Since this mechanism does not play a significant role in our model, we abstract from these asset pricing facts.

3 The risk neutral case

In this section we analyze the simplest asset pricing model assuming risk neutrality and time separable preferences ($\sigma = 0$ and $\mathbb{C}_t = C_t$). The goal of this section is to derive *qualitative* results and to show how the introduction of learning improves the performance compared to a setting with rational expectations. Sections 4 and 5 present a formal quantitative model evaluation, extending the analysis to risk-averse investors.

With risk neutrality and rational expectations the model misses almost all of the asset pricing facts described in the previous section:¹⁹ the PD ratio is

¹⁹Since the RE model implies a constant PD ratio, it trivially has a ‘persistent’ PD ratio.

constant, stock returns are unforecastable (*i.i.d.*) and approximately as volatile as dividend growth, i.e., do not display excess volatility. In addition, there is no equity premium. For these reasons, the risk-neutral model is particularly suited to illustrate how the introduction of learning qualitatively improves model performance.

The consumer has beliefs about future variables, these beliefs are summarized in expectations denoted \tilde{E} that we now allow to be less than fully rational. Under the assumptions of this section, equation (3) becomes²⁰

$$P_t = \delta \tilde{E}_t (P_{t+1} + D_{t+1}) \quad (7)$$

This asset pricing equation will be the focus of our analysis in this section.

Some papers in the learning literature²¹ have studied stock prices when agents formulate expectations about the discounted sum of all future dividends and set

$$P_t = \tilde{E}_t \sum_{j=1}^{\infty} \delta^j D_{t+j} \quad (8)$$

and the evaluation of the expectation is based on the Bayesian posterior distribution of the parameters in the dividend process. It is well known that under RE and some limiting condition on price growth the one-period ahead formulation of (7) is equivalent to the discounted sum expression for prices.²² However, under the standard learning rules used in the literature, this is not the case.

If agents learn on the price according to (8), the posterior is about parameters of an exogenous variable, i.e., the dividend process. Market prices do then not influence expectations. As a result, learning in these papers is not self-referential and Bayesian beliefs are straightforward to formulate. Yet, this lack of feedback from market prices to expectations also limits the ability of the model to generate interesting ‘data-like’ behavior. Here instead, we use the formulation in equation (7), where agents are assumed to have a forecast model of next period’s price and dividend. They try to estimate the parameters of this forecast model and will have to use data on stock prices to do so. Our point will be that it is precisely when agents formulate expectations on prices to satisfy (7) that there is a large effect of learning and that many moments of the data are better matched. It is in fact this self-referential nature of our model that makes it attractive in explaining the data.

Focusing on equation (7) instead of (8) can be justified by a number of other arguments based on principles. Informally, one can say that most participants in the stock market care much more about the selling price of the stock than

²⁰This equation is similarly implied by many other models, for example, it can be interpreted as a no-arbitrage condition in a model with risk-neutral investors or can be derived from an overlapping generations model. All that is required is that investors formulate expectations about the future payoff $P_{t+1} + D_{t+1}$ and for investors’ choice to be in equilibrium, today’s price has to equal next period’s discounted expected payoff.

²¹For example, Timmermann (1993, 1996), Brennan and Xia (2001), Cogley and Sargent (2006).

²²More precisely, equivalence is obtained if $\tilde{E}_t[\cdot] = E_t[\cdot]$ and if the no-rational-bubble requirement $\lim_{j \rightarrow \infty} \delta^j E P_{t+j} = 0$ must hold.

about the discounted dividends. More formally, this would arise in a model with only short-run investors.²³ Also, the discounted sum formula implicitly assumes that agents know perfectly the process for the *market* interest rate, therefore it either assumes a lot of knowledge about interest rates on the part of the agents or it ignores issues of learning about the interest rate.²⁴ Because of these and a number of other reasons^{25,26} we conclude that our one-period formulation in terms of prices is an interesting avenue to explore.

3.1 Analytical results

In this section we show that the introduction of learning changes qualitatively the behavior of stock prices in the direction of matching the stylized facts described above. At this point we consider a wide class of learning schemes that includes the standard linear rules used in the literature; later on we will restrict attention to learning schemes that forecast well within the model.

We first trivially rewrite the expectation of the agent by splitting the sum in the expectation:

$$P_t = \delta \tilde{E}_t(P_{t+1}) + \delta \tilde{E}_t(D_{t+1}) \quad (9)$$

We assume that agents know how to formulate the conditional expectation of the dividend $\tilde{E}_t(D_{t+1}) = aD_t$, which amounts to assuming that agents have rational expectations about the dividend process. This simplifies the discussion and highlights the fact that it is learning about future prices that allows the model to better match the data. Appendix A.4 shows that the pricing implications are very similar if agents also learn how to forecast dividends.²⁷

Agents are assumed to use a learning scheme in order to form a forecast $\tilde{E}_t(P_{t+1})$ based on past information. Equation (4) shows that under rational

²³Allen, Morris and Shin (2006) also make this argument informally. It is possible to justify it formally in an overlapping generations model. We do not pursue this further in this paper.

²⁴This point can be formalized in a model of heterogeneous agents where the market interest rate is not equal to the discount factor of a single agent. The agent's knowledge about his/her own discount factor does then not imply knowledge of the market interest rate. A Bayesian agent would have to formulate expectations about the interest rate also and the learning formula would have to be altered significantly.

²⁵If the model is slightly misspecified and agents use robust forecasting rules the 'law of iterated expectations', which is required to obtain the discounted sum, may not hold. Learning about stock price is then likely to be a more robust formulation of expectations. Adam (2007) provides experimental evidence of the breakdown of the law of iterated expectations in an experiment where agents become gradually aware that they use a possibly misspecified forecasting model.

²⁶It is also worth noting that most applications of the discounted dividend formula do not make a strict use of the principles of Bayesian learning. For example, Timmermann assumes that agents form a posterior for the serial correlation of dividends ρ . Then prices satisfy $P_t = \sum_{j=1}^{\infty} \delta^j E_t^{Bay}(\rho^j) D_t$ where E_t^{Bay} is the expectation with the Bayesian posterior. Therefore, a strictly Bayesian agent would use the posterior distribution to form a different expectation $E_t^{Bay}(\rho^j)$ for each j . Instead, he uses $P_t = \sum_{j=1}^{\infty} \delta^j [E_t^{Bay}(\rho)]^j D_t$, where $E_t^{Bay}(\rho)$ is the OLS estimate. This is a valuable simplification but, of course, it is not a fully rational model.

²⁷In section 6 we verify this also quantitatively in a model with learning about dividends and prices.

expectations $E_t [P_{t+1}] = aP_t$. As we restrict our analysis to learning rules that behave close enough to rational expectations, we specify expectations under learning as

$$\tilde{E}_t [P_{t+1}] = \beta_t P_t \quad (10)$$

where $\beta_t > 0$ denotes agents' time t estimate of gross stock price growth. For $\beta_t = a$, learning agents beliefs are rational. Also, if over time agents' beliefs converge to the RE equilibrium ($\lim_{t \rightarrow \infty} \beta_t = a$), investors will realize in the long-run that they were correct in using the functional form (10). Yet, during the transition beliefs may deviate from rational ones.

It remains to specify how agents update their beliefs β_t . We consider the following general learning mechanism

$$\Delta\beta_t = f_t \left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) \quad (11)$$

for some exogenously chosen functions $f_t : R \rightarrow R$ with the property

$$\begin{aligned} f_t(0) &= 0 \\ f_t' &> 0 \end{aligned}$$

We thus assume beliefs to adjust in the direction of the last prediction error, i.e., agents revise beliefs upwards (downwards), if they underpredicted (over-predicted) stock price growth in the past.²⁸ Arguably, a learning scheme that violates these conditions would appear quite unnatural. For technical reasons, we also need to assume that the functions f_t are such that

$$0 < \beta_t < \delta^{-1} \quad (12)$$

at all times. This rules out beliefs $\beta_t > \delta^{-1}$ which would imply that expected stock returns exceed the inverse of the discount factor, prompting the representative agent to have an infinite demand for stocks at any stock price.

Deriving quantitative and convergence results on learning will require specifying the learning scheme more explicitly. At this point, we show that the key features of the model emerge within this more general specification.

3.1.1 Stock prices under learning

Given the perceptions β_t , the expectation function (10), and the assumption on perceived dividends, equation (9) implies that prices under learning satisfy

$$P_t = \frac{\delta a D_t}{1 - \delta \beta_t}. \quad (13)$$

²⁸Note that β_t is determined from observations up to period $t - 1$ only. The assumption that the current price does not enter in the formulation of the expectations is common in the learning literature and is entertained for simplicity. Difficulties emerging with simultaneous information sets are discussed in Adam (2003).

Since β_t and ε_t are independent, the previous equation implies that

$$\text{Var} \left(\ln \frac{P_t}{P_{t-1}} \right) = \text{Var} \left(\ln \frac{1 - \delta\beta_{t-1}}{1 - \delta\beta_t} \right) + \text{Var} \left(\ln \frac{D_t}{D_{t-1}} \right), \quad (14)$$

and shows that price growth under learning is more volatile than dividend growth. Clearly, this occurs because the volatility of beliefs adds to the volatility generated by fundamentals. While this intuition is present in previous models of learning, e.g., Timmermann (1993), it will be particular to our case that under more specific learning schemes $\text{Var} \left(\ln \frac{1 - \delta\beta_t}{1 - \delta\beta_{t+1}} \right)$ is very high and remains high for a long time.

Equation (13) shows that the beliefs β_t are monotonically related to the PD ratio. One can thus understand the qualitative dynamics of the PD ratio by studying the belief dynamics. We now derive these dynamics.

$$\frac{P_t}{P_{t-1}} = T(\beta_t, \Delta\beta_t) \varepsilon_t \quad (15)$$

where

$$T(\beta, \Delta\beta) \equiv a + \frac{a\delta \Delta\beta}{1 - \delta\beta} \quad (16)$$

Substituting (15) in the law of motion for beliefs (11) and using also (13) shows that the dynamics of β_t ($t \geq 1$) can be described as a function of the shocks ε_t ($t \geq 1$), the initial conditions (D_0, P_{-1}), and the initial belief β_0 . Alternatively, the evolution of β_t is described by a second order stochastic non-linear difference equation

$$\Delta\beta_{t+1} = f_{t+1} (T(\beta_t, \Delta\beta_t)\varepsilon_t - \beta_t). \quad (17)$$

This equation can not be solved analytically, but it is still possible to gain qualitative insights into the belief dynamics of the model. We do this in the next section.

3.1.2 Deterministic dynamics

To discuss the dynamics of beliefs β_t under learning, we simplify matters by considering the deterministic case in which $\varepsilon_t \equiv 1$. Equation (17) then simplifies to

$$\Delta\beta_{t+1} = f_{t+1} (T(\beta_t, \Delta\beta_t) - \beta_t). \quad (18)$$

We thus restrict attention to the endogenous stock price dynamics generated by the learning mechanism rather than the dynamics induced by exogenous disturbances. Equation (18) shows that beliefs are increasing whenever $T(\beta_t, \Delta\beta_t) > \beta_t$, i.e., whenever actual stock price growth exceeds expected stock price growth. Understanding the evolution of beliefs thus requires studying the T -mapping.

We start by noting that the actual stock price growth implied by T depends not only on the level of price growth expectations β_t but also on the change $\Delta\beta_t$, and that this imparts momentum on the stock prices, leading to a feedback between increases in expectations and increases in actual stock price growth:

Momentum: For all $\beta_t \in (0, \delta^{-1})$

$$\begin{aligned} T(\beta_t, \Delta\beta_t) &> a && \text{if } \Delta\beta_t > 0 \\ T(\beta_t, \Delta\beta_t) &< a && \text{if } \Delta\beta_t < 0 \end{aligned}$$

Therefore, if agents arrived at the rational expectations belief $\beta_t = a$ from below ($\Delta\beta_t > 0$), the price growth generated by the learning model exceeds the fundamental growth rate a . Therefore,

$$\Delta\beta_{t+1} > 0 \quad \text{if } \beta_t = a \text{ and } \Delta\beta_t > 0$$

Just because agents' expectations have become more optimistic (in what a journalist would perhaps call a 'bullish' market), the price growth in the market has a tendency to be larger than fundamental growth. Since agents will use this higher-than-fundamental stock price growth to update their beliefs in the next period, β_{t+1} will tend to overshoot a , which will further reinforce the upward tendency. Since beliefs are monotonically related to the PD ratio, see equation (13), there will be momentum the asset price, which could be interpreted as a 'naturally arising Ponzi process'. Conversely, if $\beta_t = a$ in a bearish market ($\Delta\beta_t < 0$), beliefs and prices display downward momentum, i.e., a tendency to undershoot the RE value.

Stock prices and beliefs, however, can not stay at levels unjustified by fundamentals forever. Appendix A.2 proves that under some additional technical assumptions we have

Mean reversion: For any $\eta > 0$ and t such that $\beta_t > a + \eta$ ($\beta_t < a - \eta$), there is a finite time $\bar{t} > t$ such that $\beta_{\bar{t}} < a + \eta$ ($\beta_{\bar{t}} > a - \eta$).

Since η can be chosen arbitrarily small, the previous statement shows that beliefs will eventually move back arbitrarily close to fundamentals or even move to the 'other side' of fundamentals. This occurs even if agents' beliefs are currently far away from fundamentals. The monotone relationship between beliefs and the PD ratio then implies mean reverting behavior of the PD ratio.

Figure 2 illustrates the momentum and mean reversion by depicting the phase diagram for the dynamics of the beliefs (β_t, β_{t-1}) implied by equation (18).²⁹ The arrows in the figure thereby indicate the direction in which the difference equation is going to move for any possible state (β_t, β_{t-1}) and the solid lines indicate the boundaries of these areas.³⁰ Since we have a difference rather than a differential equation, we cannot plot the evolution of beliefs exactly. Nevertheless, the arrows suggest that the beliefs are likely to move in ellipses around the rational expectations equilibrium $(\beta_t, \beta_{t-1}) = (a, a)$. Consider, for example, point A in the diagram. Although at this point β_t is already below

²⁹Appendix A.3 explains the construction of the phase diagram.

³⁰The vertical solid line close to δ^{-1} is meant to illustrate the restriction $\beta < \delta^{-1}$ imposed on the learning rule.




Figure 2: Phase diagram illustrating momentum and mean-reversion

its fundamental value, the phase diagram indicates that beliefs will fall further. This is the result of the momentum induced by the fact that $\beta_t < \beta_{t-1}$ at point A. If beliefs go to point B, they will start to revert direction and move on to points such as C and D: beliefs thus display mean reversion. The elliptic movements imply that beliefs (and thus the PD ratio) are likely to go up and down in sustained and persistent swings.

Momentum together with the mean reversion allows the model to match the volatility and serial correlation of the PD ratio (facts 1 and 2). Also, according to our discussion around equation (14), momentum imparts variability to the ratio $\frac{1-\delta\beta_{t-1}}{1-\delta\beta_t}$ and is likely to deliver more volatile stock returns (fact 3). As discussed in Cochrane (2002), a serially correlated and mean reverting PD ratio gives rise to excess return predictability (fact 4). The next section illustrates the previous findings with the help of simulations and shows that the risk-neutral model under learning and can also generate a sizable equity premium (fact 5).

3.2 The Risk Neutral Model with OLS

In this section we report simulations of the risk neutral model under rational expectations and with learning. The latter requires specifying a particular learning rule. This is done in the next section.

3.2.1 The learning rule

We follow most empirical applications in the bounded rationality literature and use the standard updating equation from the stochastic control literature

$$\beta_t = \beta_{t-1} + \frac{1}{\alpha_t} \left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) \quad (19)$$

for all $t \geq 1$, for a given sequence of α_t , and a given initial belief β_0 which is given outside the model.³¹ The sequence $1/\alpha_t$ is called the ‘gain’ sequence and dictates how strongly beliefs are updated in the direction of the last prediction error. In this section, we assume the simplest possible specification:

$$\begin{aligned} 1/\alpha_t &= 1/(\alpha_{t-1} + 1) & t \geq 2 \\ 1/\alpha_1 &\in [0, 1] & \text{given.} \end{aligned} \quad (20)$$

With these assumptions the model evolves as follows. In the first period, β_0 determines the first price P_0 and, given the historical price P_{-1} , the first observed growth rate $\frac{P_0}{P_{-1}}$, which is used to update beliefs to β_1 using (19); the belief β_1 determines P_1 and the process evolves recursively in this manner. As in any self-referential model of learning, prices enter in the determination of beliefs and vice versa. As we argued in the previous section, it is this feature of the model that generates the dynamics we are interested in.

Using simple algebra, equation (19) implies

$$\beta_t = \frac{1}{t + \alpha_1 - 1} \left(\sum_{j=0}^{t-1} \frac{P_j}{P_{j-1}} + (\alpha_1 - 1) \beta_0 \right).$$

For the case where α_1 is an integer, this expression shows that β_t is equal to the average sample growth rate, if - in addition to the actually observed prices - we would have $(\alpha_1 - 1)$ observations of a growth rate equal to β_0 . A low initial value for $1/\alpha_1$ thus indicates that agents possess a high degree of ‘confidence’ in their initial belief β_0 .

In a Bayesian interpretation, β_0 would be the prior mean of stock price growth, $(\alpha_1 - 1)$ the precision of the prior, and - assuming that the growth rate of prices is normally distributed and i.i.d. - the beliefs β_t would be equal to the posterior mean. One might thus be tempted arguing that β_t is effectively a Bayesian estimator. Obviously, this is only true for a ‘Bayesian’ placing probability one on $\frac{P_t}{P_{t-1}}$ being i.i.d.. Since learning causes price growth to deviate from i.i.d. behavior, such priors fail to contain the ‘grain of truth’ typically assumed to be present in Bayesian analysis. While the i.i.d. assumption will hold asymptotically (we will prove this later on), it is violated under the transition dynamics.³²

³¹In the long-run the particular initial value β_0 is of little importance.

³²In a proper Bayesian formulation agents would use a likelihood function with the property that if agents use it to update their posterior, it turns out to be the true likelihood of the

For the case $1/\alpha_1 = 1$ the belief β_t is given by the sample average of stock price growth, i.e., the ordinary least squares (OLS) estimate of the mean growth rate. The initial belief β_0 then matters only for the first period, but ceases to affect beliefs after the first piece of data has arrived. More generally, assuming a value for $1/\alpha_1$ close to 1 would spuriously generate a large amount of price fluctuations, simply due to the fact that initial beliefs are heavily influenced by the first few observations and thus very volatile. Also, pure OLS assumes that agents have no faith whatsoever in their initial belief and possess no knowledge about the economy in the beginning.

In the spirit of restricting equilibrium expectations in our learning model to being close to rational, we set initial beliefs equal to the value of the growth rate of prices under RE

$$\beta_0 = a$$

and choose as initial value for $1/\alpha_1$ close to zero. Thus, we assume that beliefs start at the RE value, and that the initial degree of confidence in the RE belief is high, but not perfect. Clearly, in the limit $1/\alpha_1 \rightarrow 0$, our learning model reduces to the RE model, so that the initial gain can be interpreted as a measure of how ‘close’ the learning model is to the rational expectations model. The maximum distance from RE is achieved for $1/\alpha_1 = 1$, i.e., pure OLS learning.

Finally, we need to introduce a feature that prevents perceived stock price growth from violating the upper inequality in (12). For simplicity, we follow Timmermann (1996) and Cogley and Sargent (2006) and apply a projection facility: if in some period the belief β_t as determined by (19) is larger than some constant $K \leq \delta^{-1}$, then set

$$\beta_t = \beta_{t-1} \tag{21}$$

in that period, otherwise we use (19). The interpretation is that if the observed price growth implies beliefs that are too high, agents realize that this would prompt a crazy action (infinite stock demand) and they decide to ignore this observation. The constant K will be chosen such that the implied PD ratio is less than a certain upper bound U^{PD} . In the simulations below this facility is binding only rarely and the properties of the learning model are not sensitive to the precise value we assign to U^{PD} .

3.2.2 Asymptotic Rationality

In this section we study the limiting behavior of the model under learning, drawing on results from the literature on least squares learning. This literature shows that the T -mapping defined in equation (16) is central to whether or not agents’ beliefs converge to the RE value.³³ It is now well established that in a

model in all periods. Since the ‘correct’ likelihood in each period depends on the way agents learn, it would have to solve a complicated fixed point. Finding such a truly Bayesian learning scheme is very difficult and the question remains how agents could have learned a likelihood that has such a special property. For these reasons Bray and Kreps (1987) concluded that models of self-referential Bayesian learning were unlikely to be a fruitful avenue of research.

³³See Marcat and Sargent (1989) and Evans and Honkapohja (2001)

large class of models convergence (divergence) of least squares learning to (from) RE equilibria is strongly related to stability (instability) of the associated o.d.e. $\dot{\beta} = T(\beta) - \beta$. Most of the literature considers models where the mapping from perceived to actual expectations does not depend on the change in perceptions, unlike in our case where T depends on $\Delta\beta_t$. Since for large t the gain $(\alpha_t)^{-1}$ is very small, we have that (19) implies $\Delta\beta_t \approx 0$. One could thus think of the relevant mapping for convergence in our paper as being $T(\cdot, 0) = a$ for all β . Asymptotically the T -map is thus flat and the differential equation $\dot{\beta} = T(\beta) - \beta = a - \beta$ stable. This seems to indicate that beliefs should converge to the RE equilibrium value $\beta = a$ relatively quickly. One might then conclude that there is not much to be gained from introducing learning into the standard asset pricing model.

Appendix A.6 shows in detail that the above approximations are partly correct. In particular, learning globally converges to the RE equilibrium in this model, i.e., $\beta_t \rightarrow a$ almost surely. The learning model thus satisfies ‘Asymptotic Rationality’ as defined in section III in Marcet and Nicolini (2003). It implies that agents using the learning mechanism will realize in the long run that they are using the best possible forecast, therefore, would not have incentives to change their learning scheme.

Yet, the remainder of this shows that our model behaves very different from RE during the transition to the limit. This occurs although agents are using an estimator that starts with strong confidence at the RE value, that converges to the RE value, and that will be the best estimator in the long run. The difference is so large that even this very simple version of the model together with the very simple learning scheme introduced in section 3.1 qualitatively matches the asset pricing facts much better than the model under RE.

3.2.3 Simulations

We now illustrate the previous discussion of the model under learning by reporting simulation results in a calibrated example. We compare outcomes with the RE solution to show in what dimensions the behavior of the model improves when learning is introduced.

We choose the parameter values for the dividend process (1) so as to match the mean and standard deviation of US dividends. Using the log-normality assumption we set

$$a = 1.0035, \quad s = 0.0298 \tag{22}$$

We bias results in favor of the RE version of the model by choosing the discount factor so that the RE model matches the average PD ratio we observe in the

data.³⁴ This amounts to choosing

$$\delta = 0.9877.$$

As we mentioned before, for the learning model we set

$$\beta_0 = a \quad \text{and} \quad 1/\alpha_1 = 0.02$$

These starting values are chosen to insure that the agents' expectations will not depart too much from rationality: initial beliefs are equal to the RE value and the first quarterly observation of stock price growth obtains a weight of just 2%, with the remaining weight of 98% being placed on the RE 'prior'. Similarly, the initial value α_1 implies that it takes more than twelve years for β_t to converge half way from the initial RE value towards the observed sample mean of stock price growth. Finally, we set the upper bound on β_t so that the quarterly PD ratio will never exceed 500.

Table 3 shows the moments from the data together with the corresponding model statistics.³⁵ The column labeled US data reports the statistics discussed in section 2. It is clear from table 3 that the RE model fails to explain key asset pricing moments. Consistent with our earlier discussion the RE equilibrium fails to match the equity premium, the volatility of the PD ratio, the variability of stock returns, and the predictability of excess returns.³⁶

Statistic	US Data	RE model	Learning Model
$E(r^s)$	2.41	1.24	2.04
$E(r^b)$	0.18	1.24	1.24
$E(PD)$	113.20	113.20	86.04
σ_{r^s}	11.65	3.01	8.98
σ_{PD}	52.98	0.00	40.42
$\rho_{PD,-1}$	0.92	-	0.91
c_2^5	-0.0048	-	-0.0070
R_5^2	0.1986	0.00	0.2793

Table 3: Data and model under risk neutrality

³⁴This differs from the latter part of the paper where we choose δ to match the risk free rate. For the risk neutral model this would $\beta\delta > 1$ in the rational expectations equilibrium and thus infinite stock price.

³⁵We compute model statistics as follows: for each model we use 5000 realizations of 295 periods each, i.e., the same length as the available data. The reported statistic is the average value of the statistics across simulations, which is a numerical approximation to the expected value of the statistic.

³⁶Since PD is constant under RE, the coefficient c_2^5 of the predictability equation is undefined. This is not the case for the R^2 values.

In table 3, the learning model shows a considerably higher volatility of stock returns, high volatility and high persistence of the PD ratio, and the coefficients and R^2 of the excess predictability regressions all move strongly in the direction of the data. This is consistent with our earlier discussion about the price dynamics implied by learning. Clearly, the statistics of the learning model do not match the moments in the data quantitatively, but the purpose of the table is to show that allowing for small departures from rationality substantially improves the outcome. This is surprising, given that the model adds only one free parameter ($1/\alpha_1$) relative to the RE model; additional simulations we conducted show that the qualitative improvements provided by the model are very robust to changes in $1/\alpha_1$, as long as it is neither too close to zero - in which case the model behaves as the RE model - nor too large - in which case the model delivers too much volatility.

Table 3 also shows that the learning model delivers a positive equity premium. To understand why this occurs it proves useful to re-express the stock return between period 0 and period T as the product of three terms

$$\prod_{t=1}^T \frac{P_t + D_t}{P_{t-1}} = \underbrace{\prod_{t=1}^T \frac{D_t}{D_{t-1}}}_{=R_1} \cdot \underbrace{\left(\frac{PD_T + 1}{PD_0} \right)}_{=R_2} \cdot \underbrace{\prod_{t=1}^{T-1} \frac{PD_t + 1}{PD_t}}_{=R_3}.$$

The first term (R_1) is independent of the way expectations are formed, thus cannot contribute to explaining the emergence of an equity premium in the learning model. The second term (R_2) can potentially generate an equity premium if the terminal price dividend ratio in the learning model (PD_T) is on average higher than under rational expectations.³⁷ Yet, our simulations show that the opposite is the case: under learning the terminal price dividend ratio is lower (on average) than under rational expectations; this term thus generates a negative premium under learning. The equity premium must thus be due the behavior of the last component (R_3). This term gives rise to an equity premium via a *mean* effect and a *volatility* effect.

The mean effect emerges if agents' beliefs β_t tend to converge 'from below' towards their rational expectations value. Less optimistic expectations about stock price growth during the convergence process imply lower stock prices and thereby higher dividend yields, i.e., higher ex-post stock returns. Our simulations show that the mean effect is indeed present and that on average the price dividend ratio under learning is lower than under rational expectations. This explanation for the equity premium resembles the one advocated by Cogley and Sargent (2006).

Besides this mean effect, there exists also a volatility effect, which emerges from the convexity of R_3 in the price dividend ratio. It implies that the ex-post equity premium is higher under learning since the price dividend ratio

³⁷The value of PD_0 is the same under learning and rational expectations since initial expectations in the learning model are set equal to the rational expectations value.

has a higher variance than under rational expectations.³⁸ The volatility effect suggests that the inability to match the variability of the price dividend ratio and the equity premium are not independent facts and that models that generate insufficient variability of the price dividend ratio also tend to generate an insufficiently high equity premium.

4 Baseline model with risk aversion

The remaining part of the paper shows that the learning model can also *quantitatively* account for the moments in the data, once one allows for moderate degrees of risk-aversion, and that this finding is robust to a number of alternative specifications. Here we present the baseline model with risk aversion, our simple baseline specification for the learning rule (OLS), and the baseline calibration procedure. The quantitative results for this baseline specification are discussed in section 5, while section 6 illustrates the robustness of our quantitative findings to a variety of changes in the learning rule and the calibration procedure.

4.1 Learning under risk aversion

We now present the baseline learning model with risk aversion and show that the insights from the risk neutral model extend in a natural way to the case with risk aversion.

The investor's first-order conditions (3) together the assumption that agents know the conditional expectations of dividends deliver the asset pricing equation under learning:³⁹

$$P_t = \delta \tilde{E}_t \left(\left(\frac{D_t}{D_{t+1}} \right)^\sigma P_{t+1} \right) + \delta E_t \left(\frac{D_t^\sigma}{D_{t+1}^{\sigma-1}} \right) \quad (23)$$

With rational expectations about future price, the equilibrium stock price is⁴⁰

$$P_t^{RE} = \frac{\delta \beta^{RE}}{1 - \delta \beta^{RE}} D_t \quad (24)$$

$$\beta^{RE} = a^{1-\sigma} e^{-\sigma(1-\sigma)\frac{s^2}{2}} \quad (25)$$

³⁸The data suggest that this convexity effect is only moderately relevant: for the US data 1927:2-2000:4, it is at most 0.16% in quarterly real terms, thus explains about 8% of the equity premium.

³⁹As in section 2, we impose the market clearing condition $C_t = D_t$ and will associate consumption with dividends in the data. This is not entirely innocuous as dividend growth in the data is considerably more volatile than consumption growth, as discussed by Campbell and Cochrane (1999) in detail. Section 6 will illustrate the robustness of our quantitative findings when allowing for the fact that $C_t \neq D_t$.

⁴⁰Assuming $E_t \left(\left(\frac{D_t}{D_{t+1}} \right)^\sigma P_{t+1}^{RE} \right) = \beta^{RE} P_t^{RE}$ in (3), one obtains (24). Using the latter to substitute for prices in the assumed relationship above delivers (25).

implying

$$E_t \left(\left(\frac{D_t}{D_{t+1}} \right)^\sigma P_{t+1}^{RE} \right) = \beta^{RE} P_t^{RE}.$$

In close analogy to the risk-neutral case, we consider learning agents whose expectations in (23) are of the form

$$\tilde{E}_t \left(\left(\frac{D_t}{D_{t+1}} \right)^\sigma P_{t+1} \right) = \beta_t P_t \quad (26)$$

where β_t denotes agents' best estimate of $E(D_t/D_{t+1})^\sigma P_{t+1}/P_t$, i.e., their expectations of *risk-adjusted* stock price growth. For $\beta_t = \beta^{RE}$ agents have rational expectations and if $\beta_t \rightarrow \beta^{RE}$ beliefs will satisfy Asymptotic Rationality.

As a baseline specification, we consider again the case where agents use OLS to formulate their expectations of future (risk-adjusted) stock price growth

$$\beta_t = \beta_{t-1} + \frac{1}{\alpha_t} \left[\left(\frac{D_{t-2}}{D_{t-1}} \right)^\sigma \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right] \quad (27)$$

where the gain sequence $1/\alpha_t$ continues to be described by (20).

In the spirit of allowing for only small deviations from rationality, we restrict initial beliefs to be rational ($\beta_0 = \beta^{RE}$) and assume that agents have high but less than full confidence in this belief ($1/\alpha_1$ close to but not equal to zero). Appendix A.6 shows that learning will cause beliefs to globally converge to RE, i.e., $\beta_t \rightarrow \beta^{RE}$ and $|P_t^{RE} - P_t| \rightarrow 0$ almost surely. The learning scheme thus satisfies Asymptotic Rationality.

For $\sigma = 0$ the setup above reduces to the risk-neutral model with learning studied in section 3. For $\sigma > 0$ the setup is isomorphic with that under risk neutrality, except that 1. the beliefs β now have an interpretation as risk-adjusted stock price growth rather than simple stock price growth; 2. The term $(D_{t-2}/D_{t-1})^\sigma$ now enters the belief updating formula (27). Since for sufficiently large σ the variance of realized risk-adjusted stock price growth under RE increases with σ , the latter implies that larger risk aversion is likely to generate more volatility in beliefs and, therefore, of actual prices under learning.⁴¹ This will improve the ability of the learning model to match the moments in the data.

As in the risk-neutral case we need to impose a projection facility to insure that beliefs satisfy the inequality (12). To facilitate model calibration, described in the next section, we impose a differentiable projection facility. The details are describe in appendix A.5.3. As before, the projection facility insures that the PD ratio will never exceed a value of 500.

Finally, we show that beliefs continue to display momentum and mean-reversion, similar to the case with risk-neutrality. Using equations (26), (23),

⁴¹The formula for the variance risk adjusted stock price growth under rational expectations is

$$VAR \left(\left(\frac{D_{t-2}}{D_{t-1}} \right)^\sigma \frac{P_{t-1}^{RE}}{P_{t-2}^{RE}} \right) = a^{2(1-\sigma)} e^{(-\sigma)(1-\sigma)\frac{s^2}{2}} (e^{(1-\sigma)^2 s^2} - 1)$$

This variance reaches a minimum for $\sigma = 1$.

and the fact that $E_t(D_t^\sigma D_{t+1}^{1-\sigma}) = \beta^{RE} D_t$ shows that stock prices under learning are given by

$$P_t = \frac{\delta \beta^{RE}}{1 - \delta \beta_t} D_t \quad (28)$$

$$\frac{P_t}{P_{t-1}} = \left(1 + \frac{\delta \Delta \beta_t}{1 - \delta \beta_t}\right) a \varepsilon_t \quad (29)$$

From equations (27) and (29) follows that the expected dynamics of beliefs in the risk averse model can be described by

$$E_{t-1} \Delta \beta_{t+1} = \frac{1}{\alpha_{t+1}} (T(\beta_t, \Delta \beta_t) - \beta_{t-1}) \quad (30)$$

where

$$T(\beta_t, \Delta \beta_t) = \left(\beta^{RE} + \frac{\beta^{RE} \delta \Delta \beta_t}{1 - \delta \beta_t} \right) \quad (31)$$

The updating equation (30) has the same structure as equation (18) and the T-map (31) is identical to (16), which has been studied before.⁴² The implications regarding momentum and mean reversion from section 3.1 thus directly apply to the expected belief dynamics in the model with risk-aversion.

The only difference is that risk aversion $\sigma > 0$ changes the value of the limit point β^{RE} relative to the asymptote δ^{-1} . It is well known that, for σ sufficiently large, β^{RE} as well as the variance of realized risk-adjusted stock price growth under RE are increasing with σ .⁴³ Since β_t tends to be around β^{RE} , this means that $\delta \beta_t$ is going to be closer to 1 and equation (16) then implies that volatility under learning is even higher, as this equation has an asymptot at $\delta \beta_t = 1$.

We conclude that, qualitatively, the main features of the model under learning are likely to remain after risk aversion is introduced, but that the model now has an even larger chance to generate high volatility.

4.2 Baseline calibration procedure

This section describes and discusses our preferred calibration procedure. The parameter vector of our baseline learning model is $\theta \equiv (\delta, \sigma, 1/\alpha_1, a, s)$, where δ is the discount factor, σ the coefficient of relative risk aversion, $1/\alpha_1$ the gain parameter controlling agents' confidence in the rational expectations value, and a and s the mean and standard deviation of log dividend growth, respectively.

We choose the parameters (a, s) and δ to match moments that are not directly related to stock prices. In particular, we identify (a, s) directly with the mean and standard deviation of log dividend growth in the data, i.e., use the values displayed in equation (22), and choose δ such that the model matches the average real interest rate in the data.⁴⁴ As in Campbell and Cochrane (1999),

⁴²The latter holds because $\beta^{RE} = a$ in the case with risk-neutrality.

⁴³For the parameter values of this paper, β^{RE} increases with σ as long as $\sigma \gtrsim 3$.

⁴⁴More precisely, we choose δ to match the evidence in table 1, i.e., such that $r^B = \delta E_t \left(\frac{D_t^\sigma}{D_{t+1}^\sigma} \right) = 1.0018$, given the values for a, s and σ .

our model exhibits a constant real interest rate, and since agents are assumed to know the dividend process, the interest rate is the same as with rational expectations. The baseline learning model thus can not improve fit along this dimension.

This leaves us with *two* free parameters $(\sigma, 1/\alpha_1)$ and *seven* remaining stock price moments from table 1

$$\widehat{\mathcal{S}}' \equiv \left(\widehat{E}(r^s), \widehat{E}(PD), \widehat{\sigma}_{r^s}, \widehat{\sigma}_{PD}, \widehat{\rho}_{PD_t, PD_{t-1}}, \widehat{c}_2^5, \widehat{R}_{5,}^2, \right)$$

As discussed in detail in section 2, these moments quantitatively capture the stock price facts we seek to explain.

The parameters $(\sigma, 1/\alpha_1)$ have no immediate link to any of the moments in $\widehat{\mathcal{S}}$, i.e., it is far from obvious which of the moments to choose for calibration. In addition, some of these moments have a rather large standard deviation in the data (see table 4 below). Matching any of them exactly thus appears arbitrary, as one obtains rather different parameters estimates depending on which moment is chosen for calibration. For these reasons, we depart from the usual calibration practice and choose the values for (σ, α_1) so as to fit all seven statistics in the vector $\widehat{\mathcal{S}}$ as good as possible.

As measures of goodness-of-fit we focus on the t-ratios

$$\frac{\widehat{\mathcal{S}}_i - \mathcal{S}_i(\theta^c)}{\widehat{\sigma}_{\mathcal{S}_i}} \quad (32)$$

where $\widehat{\mathcal{S}}_i$ denotes the i -th moment from the data, $\mathcal{S}_i(\theta)$ the corresponding moment implied by the model, and $\widehat{\sigma}_{\mathcal{S}_i}$ the standard deviation of the moment. We then choose the parameterization $(\sigma, 1/\alpha_1)$ that minimizes the sum of squared t-ratios, where the sum is taken over for all seven moments. This implies that moments with a larger standard deviation receive less weight, i.e., are matched less precisely, but also that the calibration result is invariant to a potential rescaling of the moments. The details of the procedure are defined and explained in appendix A.5.

In the calibration literature it is standard to construct $\widehat{\sigma}_{\mathcal{S}_i}$ in equation (32) by looking at the model implied standard deviation of the considered moment. It is also common to conclude that the model's fit is satisfactory if the t-ratios are less than, say, two or three in absolute value. In our application, this practice has a number of problems. First, by choosing a model parameterization that gives rise to large standard deviations for the considered moment, i.e., a parameterization implying very unsharp predictions, one appears to improve 'fit'. For example, a model predicting implausibly large volatility of stock prices will 'fit' very well, simply because the standard deviation of most moments will be very high and the t-ratios (32) correspondingly low. Second, we wish to assess model performance across a variety of alternative model specifications. Using model implied standard deviations would then cause the goodness-of-fit criterion to vary across models, an aspect we do not feel comfortable with.

To avoid these problems we prefer to find an estimate of the standard deviation of each statistic $\widehat{\sigma}_{\mathcal{S}_i}$ that is based purely on data sources. We show in

appendix A.5 how to obtain consistent estimates of these standard deviations from the data. With these estimates we then follow the calibration literature, i.e., use these resulting t-ratios as our measure of ‘fit’ for each sample statistic and claim to have a good fit if this ratio is below two or three.

The procedure just described may appear somewhat like estimation by the method of simulated moments, and using the t-ratios as measures of fit may appear like using test statistics. In appendix A.5 we describe how this interpretation could be made, but we do not wish to interpret our procedure as a formal econometric method. The distribution of the parameters and test statistics for these formal estimation methods relies on asymptotics, but asymptotically our baseline learning model is indistinguishable from RE. Therefore, one would have to rely on short-sample statistics, but short-sample econometric theory is underdeveloped and developing it is a full research agenda in itself.

Summing up, we interpret the method just described as a way to pick the parameter (σ, α_1) in a systematic way, such that the model has a good chance to meet the data, but where the model could also be rejected. In this way we attempt to convince the reader that the learning model is able to quantitatively account for key features of stock price data.

5 Quantitative results

We now evaluate the quantitative performance of the baseline learning model when using the baseline calibration approach described in the previous section.

Our results are summarized in table 4 below. The second and third column of the table report the asset pricing moments from the data that we seek to match and the estimated standard deviation of the moments, respectively. The table shows that some of the reported moments, e.g., σ_{PD} , are estimated rather imprecisely; our calibration procedure will automatically place less emphasis on matching these moments.

The calibrated parameters values of the learning model are reported at the bottom of the table and appear reasonable on a priori grounds. The coefficient of relative risk aversion is well within the ranges used in previous studies. Moreover, the gain parameter $1/\alpha_1$ is close to zero, reflecting the tendency of the data to give large (but less than full) weight to the RE prior about stock price growth. As has been explained before, high values of the gain $1/\alpha_1$ tend to give rise to too much volatility because beliefs are then very volatile. The calibrated gain reported in the table implies that when updating beliefs in the initial period, the RE prior receives a weight of approximately 98,4% and the first quarterly observation a weight of about 1.6%.

Statistics	US Data		OLS	
		std		t-ratio
$E(r^s)$	2.41	0.45	2.39	0.05
$E(PD)$	113.20	15.15	109.01	0.28
σ_{r^s}	11.65	2.88	12.91	-0.44
σ_{PD}	52.98	16.53	65.47	-0.76
$\rho_{PD_t, PD_{t-1}}$	0.92	0.02	0.94	-1.14
c_2^5	-0.0048	0.002	-0.0079	1.54
R_5^2	0.1986	0.083	0.2972	-1.19
Parameters:	$\sigma = 4.47$		$1/\alpha_1 = 0.0162$	

Table 4: Moments and parameters, baseline model and baseline calibration

The fourth column in table 4 reports the moments implied by the calibrated learning model. The learning model performs remarkably well. In particular, the model with risk aversion maintains the high variability and serial correlation of the PD ratio and the variability of stock returns, as in section 3. In addition, it now succeeds in matching the mean of the PD ratio and also matches the equity premium.⁴⁵ As discussed in section 2 before, it is natural that the excess return regressions can be explained reasonably well once the serial correlation of the PD is matched.

Clearly, the point estimate of some model moments does not match exactly the observed moment in the data, but this tends to occur for moments that, in the short sample, have a large variance. This is shown in the last column of table 4 which reports the goodness-of-fit measures (t-ratios) for each considered moment. The t-ratios are all well below two and thus well within what is a 95% confidence interval, if this were be a formal econometric test (which it is not).

In summary, the results of table 4 show that introducing learning substantially improves the fit of the model relative to the case with RE and is overall very successful in quantitatively accounting for the empirical evidence described in section 2. We find this result remarkable, given that we used the simplest version of the asset pricing model and combined it with the simplest available learning mechanism.

6 Robustness

This section shows that the quantitative performance of the model is robust to many extensions. We start by exploring alternative models of learning, then consider alternative calibration procedures. In all cases, the model explains the stylized facts described in section 2.

⁴⁵Recall that the discount factor has been chosen so as to match the sample mean of the risk-free real interest rate. Matching the mean stock return, therefore, implies that the model also matches the risk premium.

Learning about dividends. In the baseline model we assume agents know the conditional expectation of dividends. This is done to simplify the exposition and because learning about dividends has been considered in previous papers.⁴⁶ Since it may appear inconsistent to assume that agents know the dividend growth process but do not know how to forecast stock prices, we consider a model where agents learn about dividend growth and stock price growth. In Appendix A.4 we lay out the model and show that, while the analysis is more involved, the basic results do not change. Table 5 below shows the quantitative results with learning about dividends using the baseline calibration procedure described in section 5. It shows that introducing dividend learning does not lead to significant changes.

Constant gain learning. An undesirable feature of the OLS learning scheme is that volatility of asset prices decreases over time, which may seem counterfactual. Therefore, we introduce a learning scheme with a constant gain, i.e., assume that the weight on the forecast error in the learning scheme is constant: $\alpha_t = \alpha_1$ for all t .⁴⁷ Beliefs then respond to forecast errors in the same way throughout the sample and asset price volatility does not decrease. As discussed extensively in the learning literature, such a learning scheme improves forecasting performance when there are changes in the environment. Agents who suspect the existence of switches in the average growth rate of prices or fundamentals, for example, may be reasonably expected to use such a scheme. Table 5 reports the quantitative results for the constant gain model using the baseline calibration approach. For obvious reasons, stock prices are now more volatile, even if the initial gain is substantially lower than in the baseline case. Overall, the fit of the model improves, especially regarding the evidence on excess return predictability.⁴⁸

Switching weights. We now introduce a learning model in which the gain switches over time, as in Marcet and Nicolini (2003).⁴⁹ The idea is to combine constant gain with OLS, using the former in periods in which forecast errors are large and the latter when the forecast error is low. We report the quantitative results in Table 5, which are very similar to those with pure constant gain learning. The latter occurred because the model was frequently in ‘constant gain mode’.

⁴⁶E.g., Timmermann (1993, 1996).

⁴⁷The derivations for this model are as in section 4 and require only changing the evolution of α .

⁴⁸We do not use constant gain as our main learning scheme because β_t does then not converge, i.e., we lose asymptotic rationality. Nevertheless, constant gain would generate better forecasts than OLS in a setup where there are trend switches in fundamentals. Since there is some evidence of such switches in the data, we believe that a model with constant gain will eventually generate better forecasts in a model with trend switches and fit the data better. We leave this for future research.

⁴⁹The advantage of using switching gains, relative to using simple constant gain, is that under certain conditions, the learning model converges to RE, i.e., Asymptotic Rationality is indeed preserved.

Statistic	US Data		Learning about Div.		Constant gain		Switching weights	
				t-ratio		t-ratio		t-ratio
$E(r^s)$	2.41	2.38	0.07		2.19	0.50	2.18	0.51
$E(PD)$	113.20	105.04	0.54		117.94	-0.31	117.72	-0.30
σ_{r^s}	11.65	12.31	-0.23		13.91	-0.78	13.62	-0.68
σ_{PD}	52.98	59.38	-0.39		73.59	-1.25	73.04	-1.21
$\rho_{PD_t, PD_{t-1}}$	0.92	0.94	-1.26		0.93	-0.24	0.93	-0.20
c_2^5	-0.0048	-0.0087	1.9446		-0.0063	0.745	-0.0064	0.7767
R_5^2	0.1986	0.3276	-1.5580		0.2278	-0.352	0.2273	-0.3465
Parameters:								
σ		4.64			3.31		3.03	
$1/\alpha_1$		0.0163			0.0069		0.0070	

Table 5: Robustness, Part I

Consumption data. Throughout the paper we made the simplifying assumption $C_t = D_t$ and calibrated this process to dividend data, mainly because in studying stock price volatility the data on dividends has to be brought out. In the data, consumption and dividend growth display rather different volatilities. Therefore, we now allow for $C_t \neq D_t$ and calibrate the consumption and dividend process separately to the data.⁵⁰ While dividends are as before, we set

$$\frac{C_{t+1}}{C_t} = a\varepsilon_{t+1}^c \quad \text{for } \ln \varepsilon_t^c \sim iiN\left(-\frac{s_c^2}{2}; s_c^2\right)$$

The presence of two shocks modifies the equations for the RE version of the model in a well known way and we do not describe it in detail here. We calibrate the model by following Campbell and Cochrane (1999) and set $s_c = \frac{s}{7}$ and $\rho(\varepsilon^c, \varepsilon) = .2$.⁵¹ We also have to slightly modify our calibration approach because choosing δ to match exactly the average interest rate will result in $\delta\beta^{RE} > 1$ for reasonable values of σ , while we require $\delta\beta^{RE} < 1$ to have a finite stock price. Therefore, we simply include the risk-free interest rate in our set of moments to be matched, and δ joins σ and $1/\alpha_1$ in the minimization of the sum of squared t-ratios, which now includes also t-ratio for the interest rate.

The quantitative results are reported in table 6 below, which shows that we fit the data very well for the moments reported in the table. Yet, it turns out that the value for δ implies $E(r^b) = 0.83$ in the model, while in the data we have $\widehat{E}(r^b) = 0.18$ with a standard deviation of $\widehat{\sigma}_{r^b} = 0.23$. This results in a t-ratio equal to -2.86. Therefore, this particular model does not explain the risk-free interest rate or, equivalently, does not explain the equity premium

⁵⁰This would require changing the model described in section 4 to one with an exogenous endowment that is added to the budget constraint of the agent and to the resource constraint. The modification is obvious and we omit the details.

⁵¹We take these ratios and values from table 1 in Campbell and Cochrane (1999), which is based on a slightly shorter sample than the one used in this paper.

puzzle. Nevertheless, it appears that we are still close to explaining it: if this were a valid econometric test for the moment $\widehat{E}(r^b)$ (which it is not), it would be explained at the 99% confidence level. We do not wish to make much of this quasi-rejection, as we do not see the objective of this paper as matching perfectly all moments. In particular, the equity premium is not the main focus of this paper and the literature suggests many reasonable extensions that would further improve the fit of the model in this dimension. Our conclusion from this robustness exercise is that it confirms the finding that the introduction of learning improves considerably the ability to match the data on stock prices.

Full weighting matrix We now investigate the robustness of our findings to changes in the calibration procedure. In an econometric MSM setup, one would have to find the parameters that minimize a quadratic form with a full instead of a diagonal weighting matrix, see the discussion around equation (44) in appendix A.5. As table 6 shows, we still obtain a good model fit when using such a full weighting matrix. The degree of relative risk aversion increases and the ability to match the serial correlation of the PD ratio deteriorates. We are not sure why this happens, but it may be related to the fact that the full weighting matrix is nearly singular, a situation that causes the estimates to be nearly ill-defined. We do not pursue this issue further, but simply wanted illustrate to a reader interested in econometrics: first, results do not change much if one uses more econometric-based calibration techniques; second, it is not immediate how to formulate an econometric test of the model with standard MSM tools, see also the discussion in appendix A.5.

Model-generated standard deviations As a final exercise we demonstrate the robustness of our finding to using t-ratios based on model-generated standard deviations. We do this to show to readers who are partial to calibration that results are largely unchanged if we use their preferred criterion of fit. As discussed before, we do not use this criterion in the remaining part of the paper because the comparison across models is awkward, and because some of the robustness exercises (the exercise with a separate process for consumption data and the exercise with the full weighting matrix) could not be performed using model-generated variances. The quantitative results reported in the last two columns of table 6 show that the model performs very well when using model-generated standard deviations and that calibrated parameter values remain largely unchanged.

Statistics	US Data	$C_t \neq D_t$		OLS, Full matrix		$\hat{\sigma}_{S_i}$ from Model	
			t-ratio		t-ratio		t-ratio
$E(r^s)$	2.41	2.01	0.90	2.14	0.60	2.33	0.30
$E(PD)$	113.20	114.51	-0.09	105.35	0.52	109.64	0.12
σ_{r^s}	11.65	10.43	0.43	12.84	-0.41	12.68	-0.27
σ_{PD}	52.98	61.57	-0.52	63.70	-0.65	64.94	-0.61
$\rho_{PD_t, PD_{t-1}}$	0.92	0.95	-1.42	0.97	-2.40	0.95	-0.52
$c_{\frac{5}{2}}$	-0.0048	-0.0091	2.112	-0.0061	0.6373	-0.0077	1.045
$R_{\frac{5}{2}}^2$	0.1986	0.2422	-0.526	0.3442	-1.7584	0.3020	-0.563
Parameters:							
σ		3.23		7.08		4.51	
$1/\alpha_1$		0.0248		0.0089		0.0154	

Table 6: Robustness, Part II

The previous robustness exercises allowed for deviations from the baseline model and baseline calibration one step at a time. We also performed a variety of other robustness tests that combine more than one of the features of the robustness test described above. We found that the model continues to explain the moments well in all these cases.

7 Conclusions and Outlook

A one parameter learning extension of the most standard asset pricing model strongly improves the ability of the model to quantitatively account for a number of asset pricing facts. This outcome is remarkable, given the difficulties documented in the empirical asset pricing literature in accounting for these facts. The difficulties of rational expectations models suggests that learning processes may be more relevant for explaining empirical phenomena than previously thought. Indeed, it seems that the most convincing case for models of learning can be made by explaining facts that appear ‘puzzling’ from the rational expectations viewpoint. This is what this paper does.

The simple setup presented in this paper could be extended in a number of interesting ways and also applied to study other substantive questions. One avenue that we currently explore is to ask whether learning processes can account also for the otherwise puzzling behavior of exchange rates. Clearly, the ability of simple models of learning to explain puzzling empirical phenomena in more than one market would further increase confidence in that learning-induced small deviations from rationality are indeed economically relevant.

A Appendix

A.1 Data Sources

Our data is for the United States and has been downloaded from ‘The Global Financial Database’ (<http://www.globalfinancialdata.com>). The period covered is 1925:4-2005:4. For the subperiod 1925:4-1998:4 our dataset corresponds very closely to Campbell’s (2003) data (<http://kuznets.fas.harvard.edu/~campbell/data.html>).

In the calibration part of the paper we use moments that are based on the same number of observations. Since we seek to match the return predictability evidence at the five year horizon, the effective sample end for the means and standard deviations reported in table 1 has been shortened by five years to 2000:4. In addition, due to the seasonal adjustment procedure for dividends described below and the way we compute the standard errors for the moments described in appendix A.5, the effective starting date was 1927:2.

To obtain real values, nominal variables have been deflated using the ‘USA BLS Consumer Price Index’ (Global Fin code ‘CPUSAM’). The monthly price series has been transformed into a quarterly series by taking the index value of the last month of the considered quarter.

The nominal stock price series is the ‘SP 500 Composite Price Index (w/GFD extension)’ (Global Fin code ‘_SPXD’). The weekly (up to the end of 1927) and daily series has been transformed into quarterly data by taking the index value of the last week/day of the considered quarter. Moreover, the series has been normalized to 100 in 1925:4.

As nominal interest rate we use the ‘90 Days T-Bills Secondary Market’ (Global Fin code ‘ITUSA3SD’). The monthly (up to the end of 1933), weekly (1934-end of 1953), and daily series has been transformed into a quarterly series using the interest rate corresponding to the last month/week/day of the considered quarter and is expressed in quarterly rates, i.e., not annualized.

Nominal dividends have been computed as follows

$$D_t = \left(\frac{I^D(t)/I^D(t-1)}{I^{ND}(t)/I^{ND}(t-1)} - 1 \right) I^{ND}(t)$$

where I^{ND} denotes the ‘SP 500 Composite Price Index (w/GFD extension)’ described above and I^D is the ‘SP 500 Total Return Index (w/GFD extension)’ (Global Fin code ‘_SPXTRD’). We first computed monthly dividends and then quarterly dividends by adding up the monthly series. Following Campbell (2003), dividends have been deseasonalized by taking averages of the actual dividend payments over the current and preceding three quarters.

A.2 Proof of mean reversion

To proof mean reversion of beliefs we need two additional technical assumptions:

1. We slightly strengthen (12) to $0 < \beta_t < \beta^U < \delta^{-1}$; 2. Letting

$$D_t \equiv \inf_{\Delta} \frac{\partial f_t}{\partial \Delta}(\bar{\Delta})$$

we assume $\sum_t D_t = \infty$, which is standard in the stochastic control literature. It is satisfied by all the learning mechanisms considered in this paper and is needed to avoid beliefs getting stuck away from the fundamental value simply because updating ceases too quickly.

We start proving mean reversion for the case $\beta_t > a + \eta$. Fix η and choose $\varepsilon = \eta(1 - \delta\beta^U)$. Since $\varepsilon > 0$ it cannot be that $\Delta\beta_{t'} \geq \varepsilon$ for all $t' > t$ as this would imply $\beta_t \rightarrow \infty$, violating the upper bound on β . Therefore, for some $t' > t$ we have $\Delta\beta_{t'} < \varepsilon$.

Take $t' > t$ to be the first period where $\Delta\beta_{t'} < \varepsilon$. At this point $\beta_{t'} \geq \beta_t$ and therefore

$$\beta_{t'} > a + \eta$$

and

$$T(\beta_{t'}, \Delta\beta_{t'}) = a + \frac{\Delta\beta_{t'}}{1 - \delta\beta_{t'}} < a + \frac{\varepsilon}{1 - \delta\beta^U} = a + \eta$$

where the inequality uses $\Delta\beta_{t'} < \varepsilon$ and that β^U is an upper bound for β ; the equality on the right follows from the way ε is chosen above. The previous two relations imply

$$\beta_{t'} > T(\beta_{t'}, \Delta\beta_{t'})$$

which together with (18) and the fact that f_t is increasing and $f_t(0) = 0$ gives

$$\Delta\beta_{t'+1} = f_{t'+1}(T(\beta_{t'}, \Delta\beta_{t'}) - \beta_{t'}) < 0 \quad (33)$$

There are two possibilities: either

- i) $\beta_{t'+j} \leq a + \eta$ eventually for some $j > 0$, or
- ii) $\beta_{t'+j} > a + \eta$ for all $j > 0$.

In case i) we are done. We now prove that case ii) is impossible by setting up an inductive argument that leads to a contradiction. For any $j > 0$, $\Delta\beta_{t'+j} < 0$ implies $a + \frac{\Delta\beta_{t'+j}}{1 - \delta\beta_{t'+j}} < a$ and if ii) holds $T(\beta_{t'+j}, \Delta\beta_{t'+j}) < \beta_{t'+j}$. From the properties of f_t it follows $\Delta\beta_{t'+j+1} < 0$. Thus, $\Delta\beta_{t'+j} < 0$ implies $\Delta\beta_{t'+j+1} < 0$ for any $j > 0$ with the initial condition given by (33). The negativity of $\Delta\beta_{t'+j}$ for $j > 0$ implies the first inequality below

$$a + \frac{\Delta\beta_{t'+j}}{1 - \delta\beta_{t'+j}} - \beta_{t'+j} < a - \beta_{t'+j} < a - (a + \eta) = -\eta \quad \text{for all } j > 0 \quad (34)$$

with the second following from ii). We thus have

$$\Delta\beta_{t'+j+1} = f_{t'+j+1} \left(a + \frac{\Delta\beta_{t'+j}}{1 - \delta\beta_{t'+j}} - \beta_{t'+j} \right) < f_{t'+j+1}(-\eta) \leq -\eta D_{t'+j+1}$$

where the first inequality follows from (34) and the second from the mean value theorem and $D_t \geq 0$. The previous result implies

$$\beta_{t'+j} = \sum_{i=1}^j \Delta\beta_{t'+i} + \beta_{t'} \leq -\eta \sum_{i=1}^j D_{t'+i}$$

for all $j > 0$. The assumption $\sum_t D_t = \infty$ would then imply $\beta_t \rightarrow -\infty$ which violates our assumptions that f_t is such that $\beta_t > 0$ at all times.

For the case $\beta_t < a - \eta$ we choose $\varepsilon = \eta$ and we can use a symmetric argument to make the proof.

A.3 Details on the phase diagram

The second order difference equation (18) describing the deterministic evolution of beliefs allows to construct non-linear first-order learning dynamics in the (β_t, β_{t-1}) plane. For clarity, we define $x'_t \equiv (x_{1,t}, x_{2,t}) \equiv (\beta_t, \beta_{t-1})$, whose dynamics are given by

$$x_{t+1} = \begin{pmatrix} x_{1,t} + f_{t+1} \left(a + \frac{a\delta(x_{1,t} - x_{2,t})}{1 - \delta x_{1,t}} - x_{1,t} \right) \\ x_{1,t} \end{pmatrix}$$

The zeros of the phase diagram are $\Delta x_2 = 0$ at points $x_1 = x_2$ and $\Delta x_1 = 0$ for $x_2 = \frac{1}{\delta} - \frac{x_1(1 - \delta x_1)}{a\delta}$. So the zeroes for Δx_1 and Δx_2 intersect are at $x_1 = x_2 = a$ which is the REE and, interestingly, at $x_1 = x_2 = \delta^{-1}$ which is the limit of the rational bubble equilibrium. Moreover, as is easy to verify $\Delta x_2 > 0$ for $x_1 > x_2$ and $\Delta x_1 > 0$ for $x_2 < \frac{1}{\delta} - \frac{x_1(1 - \delta x_1)}{a\delta}$. These results give rise to the phase diagram shown in figure 2.

A.4 Model with learning about dividends

This section considers agents who learn to forecast future dividends in addition to forecast future price. We make the arguments directly for the general model with risk aversion from section 4. Equation (23) then becomes

$$P_t = \delta \tilde{E}_t \left(\left(\frac{C_t}{C_{t+1}} \right)^\sigma P_{t+1} \right) + \delta \tilde{E}_t \left(\frac{D_t^\sigma}{D_{t+1}^{\sigma-1}} \right)$$

Under RE one has

$$\begin{aligned} E_t \left(\frac{D_t^\sigma}{D_{t+1}^{\sigma-1}} \right) &= E_t \left(\frac{D_{t+1}^{1-\sigma}}{D_t^{-\sigma}} \right) = E_t \left(\left(\frac{D_{t+1}}{D_t} \right)^{1-\sigma} \right) D_t \\ &= E_t \left((a\varepsilon)^{1-\sigma} \right) D_t \\ &= \beta^{RE} D_t \end{aligned}$$

This justifies that learning agents will forecast future dividends according to

$$\tilde{E}_t \left(\frac{D_{t+1}^{1-\sigma}}{D_t^{-\sigma}} \right) = \gamma_t D_t$$

where γ_t is agents's best estimate of $\tilde{E}_t \left(\left(\frac{D_{t+1}}{D_t} \right)^{1-\sigma} \right)$, which can be interpreted as risk-adjusted dividend-growth. In close analogy to the learning setup for

future price we assume that agents' estimate evolves according to

$$\gamma_t = \gamma_{t-1} + \frac{1}{\alpha_t} \left(\left(\frac{D_{t-1}}{D_{t-2}} \right)^{1-\sigma} - \gamma_{t-1} \right) \quad (35)$$

which can be given a proper Bayesian interpretation. In the spirit of allowing for only small deviations from rationality, we assume that the initial belief is correct

$$\gamma_0 = \beta^{RE}.$$

Moreover, the gain sequence α_t is the same as the one used for updating the estimate for β_t . Learning about β_t remains to be described by equation (27). With these assumptions realized price and price growth are given by

$$P_t = \frac{\delta \gamma_t}{1 - \delta \beta_t} D_t$$

$$\frac{P_t}{P_{t-1}} = \frac{\gamma_t}{\gamma_{t-1}} \left(1 + \frac{\delta \Delta \beta_t}{1 - \delta \beta_t} \right) a \varepsilon_t$$

The map T from perceived to actual expectations of the risk-adjusted price growth $\frac{P_{t+1}}{P_t} \left(\frac{D_t}{D_{t+1}} \right)^\sigma$ in this more general model is given by

$$T(\beta_{t+1}, \Delta \beta_{t+1}) \equiv \frac{\gamma_{t+1}}{\gamma_t} \left(\beta^{RE} + \frac{\beta^{RE} \delta \Delta \beta_{t+1}}{1 - \delta \beta_{t+1}} \right) \quad (36)$$

which differs from (31) only by the factor $\frac{\gamma_{t+1}}{\gamma_t}$. From (35) it is clear that $\frac{\gamma_{t+1}}{\gamma_t}$ evolves exogenously and that $\lim_{t \rightarrow \infty} \frac{\gamma_{t+1}}{\gamma_t} = 1$ since $\lim_{t \rightarrow \infty} \gamma_t = \beta^{RE}$ and $\alpha_t \rightarrow \infty$. Thus, for medium to high values of α_t and initial beliefs not too far from the RE value, the T-maps with and without learning about dividends are very similar.

For the deterministic setting with risk-neutrality considered in section 3, one has $\gamma_t = \gamma_0 = a$ and $\beta^{RE} = a$ so that (36) becomes identical to (16).

A.5 Calibration procedure

This section describes the details of our calibration approach and explains how we estimate the standard deviation of the sample statistics reported in table 4.

Let N be the sample size, $(\mathbf{y}_1, \dots, \mathbf{y}_N)$ the observed data sample, with \mathbf{y}_t containing m variables. We consider the sample statistics $\mathcal{S}(M_N)$ which are a function of the sample moment M_N with $\mathcal{S} : R^q \rightarrow R^s$ being a *statistic function* that maps moments into the considered statistics. The statistic function is required because some of the statistics we seek to match, see table 4, are not moments but functions of moments.⁵² Let $h : R^m \rightarrow R^q$ denote the function

⁵²This is the case, for example, for the correlation and the R^2 coefficients. The standard case of matching moments is obtained when \mathcal{S} is the identity.

determining the sample moments on which our statistics are based, i.e., $M_N \equiv \frac{1}{N} \sum_{t=1}^N h(\mathbf{y}_t)$. The explicit expressions for $h(\cdot)$ and $\mathcal{S}(\cdot)$ for our particular application are stated in Appendix A.5.1 below.

Given the sample statistics, we now explain how we compute the corresponding model statistics for a given model parameterization $\theta \in R^n$. Let ω^s denote a realization of shocks and $(y_1(\theta, \omega^s), \dots, y_N(\theta, \omega^s))$ the random variables corresponding to a history of length N generated by the model for shock realization ω^s . Furthermore, let

$$M_N(\theta, \omega^s) \equiv \frac{1}{N} \sum_{t=1}^N h(y_t(\theta, \omega^s))$$

denote the model moment for realization ω^s and

$$\mathcal{S}(M_N(\theta, \omega^s))$$

the corresponding model statistic. The model statistics we wish to report are the expected value of the statistic across possible shock realizations:

$$\mathcal{S}(M_N(\theta)) = E[\mathcal{S}(M_N(\theta, \omega^s))]$$

One can obtain a numerical approximation to the theoretical model statistic $\mathcal{S}(M_N(\theta))$ by averaging (for a given a parameter vector θ) across a large number of simulations of length N the statistics $\mathcal{S}(M_N(\theta, \omega^s))$ implied by each simulation. We report this average in the tables of the main text.

Now that we have explained how to compute statistics in the data and the model, we explain how we calibrate the parameters so as to match the model statistics to the statistics of the data. Let $\mathcal{S}_i(M_N)$ denote the i -th statistic from the data, $\mathcal{S}_i(M_N(\theta))$ the corresponding statistic from the model, and $\hat{\sigma}_{\mathcal{S}_i}$ an estimate for the standard deviation of the i -th statistic. How we obtain $\hat{\sigma}_{\mathcal{S}_i}$ will be explained in detail below. The baseline parameter choice $\hat{\theta}_N$ is then found as follows

$$\hat{\theta}_N \equiv \arg \min_{\theta} \sum_{i=1}^7 \left(\frac{\mathcal{S}_i(M_N) - \mathcal{S}_i(M_N(\theta))}{\hat{\sigma}_{\mathcal{S}_i}} \right)^2 \quad (37)$$

subject to the restrictions on a, s, δ that have been described in the text. Our procedure thus tries to match the model statistics as closely as possible to the data statistics, but gives less weight to statistics with a larger standard deviation. Notice that the calibration result is invariant to a rescaling of the variables of interest and that with the available two parameters (α_1, σ) , it is generally not possible to match all seven moments numerically, unless the model is closely related to the data.

To be able to use standard numerical procedures to solve the minimization problem (37), we slightly modify the projection facility described in (21) to

insure that it is continuously differentiable. Appendix A.5.3 describes this in detail.

We now explain how we obtain the estimate for the standard deviation of the i -th statistic $\widehat{\sigma}_{\mathcal{S}_i}$. We start by discussing desirable asymptotic properties of such an estimate and then explain how it has been constructed.

We would like to have an estimator $\widehat{\sigma}_{\mathcal{S}_i}$ that converges to the standard deviation of the statistic sufficiently fast, so that asymptotically the t-ratio has a standard normal distribution, i.e., we require

$$\sqrt{N} \frac{\mathcal{S}_i(M_N) - \mathcal{S}_i(M_N(\theta_0))}{\widehat{\sigma}_{\mathcal{S}_i}} \rightarrow \mathcal{N}(0, 1) \quad \text{in distribution} \quad (38)$$

as $N \rightarrow \infty$ and where θ_0 denotes the true value of the parameters and $\mathcal{S}_i(M_N(\theta_0))$ the corresponding true value of the i -th model statistic. Once we have such an estimator, it is possible to interpret t-ratios as goodness of fit measures that could be given an asymptotic interpretation.

We now explain how to find an estimate for the full covariance matrix $\widehat{\Sigma}_{\mathcal{S}, N}$ of model statistics from a sample of N observations such that

$$\widehat{\Sigma}_{\mathcal{S}, N} \rightarrow \Sigma_{\mathcal{S}} \quad \text{almost surely, for} \quad (39)$$

$$\sqrt{N} [\mathcal{S}(M_N) - \mathcal{S}(M_N(\theta_0))] \rightarrow \mathcal{N}(0, \Sigma_{\mathcal{S}}) \quad (40)$$

as $N \rightarrow \infty$. Note that the $\widehat{\sigma}_{\mathcal{S}_i}^2$ in (37) and (38) are simply the diagonal entries of $\widehat{\Sigma}_{\mathcal{S}, N}$. Let $M_0 = E[h(y_t(\theta_0, \omega^s))]$ denote the theoretical moment at the true parameter value. Assume y to be stationary and ergodic, \mathcal{S} to be continuously differentiable at M_0 , and

$$S_w \equiv \sum_{j=-\infty}^{\infty} E[(h(y_t) - M_0)(h(y_{t-j}) - M_0)'] < \infty \quad (41)$$

We then have $\Sigma_{\mathcal{S}}$ in (39) given by

$$\Sigma_{\mathcal{S}} = \frac{\partial \mathcal{S}(M_0)}{\partial M'} S_w \frac{\partial \mathcal{S}'(M_0)}{\partial M} \quad (42)$$

This follows from standard arguments: by the mean value theorem

$$\sqrt{N} [\mathcal{S}(M_0) - \mathcal{S}(M_N)] = \frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} \sqrt{N} [M_0 - M_N] \quad (43)$$

where \widetilde{M}_N is some convex combination of M_N and M_0 .⁵³ Under stationarity and ergodicity of y , we have $M_N \rightarrow M_0$ a.s. by the ergodic theorem. Since \widetilde{M}_N

⁵³As is well known, a different \widetilde{M}_N is needed for each row of \mathcal{S} but this is inconsequential for the proof and we ignore it here.

is between M_N and M_0 , this implies $\widetilde{M}_N \rightarrow M_0$ a.s. and, since $\frac{\partial \mathcal{S}(\cdot)}{\partial M'}$ has been assumed continuous at M_0 we have that

$$\frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} \rightarrow \frac{\partial \mathcal{S}(M_0)}{\partial M'} \text{ a.s.}$$

From the central limit theorem

$$\sqrt{N}(M_0 - M_N) \rightarrow N(0, S_w) \text{ in distribution}$$

Plugging the previous two relationships into (43) shows (42).

Therefore, by taking an estimate $S_{w,N}$ that converges a.s. to S_w and by taking

$$\widehat{\Sigma}_{\mathcal{S},N} \equiv \frac{\partial \mathcal{S}(M_N)}{\partial M'} S_{w,N} \frac{\partial \mathcal{S}'(M_N)}{\partial M}$$

one obtains property (38). Explicit expression for $\partial \mathcal{S}(M_N)/\partial M'$ are given in appendix A.5.2. It now only remains to find the estimates $S_{w,N}$ from the data. We follow standard practice and employ the Newey West estimator, which truncates the infinite sum in (41) and weighs the autocovariances in a particular way. This is standard and we do not describe the details here.

Our baseline procedure for choosing parameter values described above can be thought of as a hybrid between the method of simulated moments (MSM) and calibration.

We differ from fully-fledged MSM (described below) because we do not perform any formal estimation, we do not attempt to use an optimal weighting matrix, and we do not adjust for the degrees of freedom in the asymptotic distribution of the goodness of fit tests, and because we do not think of this as an exercise in accepting or rejecting the model. Instead, our procedure is simply a way of systematically choosing parameter values that allows us to display the behavior of the model and to interpret the t-ratios as giving a measure of closeness.

We also differ from calibration because we do not pin down each parameter with a given moment and use the remaining moments to test the model. Instead, we let the algorithm find the parameters (σ, α_1) that best fit the seven considered moments of the data. Moreover, in our procedure the standard deviation of the moment $\widehat{\sigma}_{\mathcal{S}_i}$ is computed from the data. This is useful for a number of reasons. First, the standard practice in calibration of using model-determined standard deviations gives an incentive to the researcher to generate models with high standard deviations, i.e., unsharp predictions, as these appear to improve fit. Second, if $\widehat{\sigma}_{\mathcal{S}_i}$ is determined from the data, it is constant across alternative models allows for model comparisons in a meaningful way.

In addition to the baseline calibration procedure above, we engage in a robustness exercise, reported under the heading ‘full matrix’ in table 6, which is closer to MSM. In particular, we choose parameters to solve

$$\widehat{\theta}_N \equiv \arg \min_{\theta} [\mathcal{S}(M_N(\theta)) - \mathcal{S}(M_N)]' \widehat{\Sigma}_{\mathcal{S},N}^{-1} [\mathcal{S}(M_N(\theta)) - \mathcal{S}(M_N)] \quad (44)$$

This way of fitting the model is less intuitive but generally has the advantage that $\widehat{\Sigma}_{\mathcal{S},N}^{-1}$ is an optimal weighting matrix so the estimate should be closer to the true model parameter. One problem we encountered is that $\widehat{\Sigma}_{\mathcal{S},N}$ is nearly singular and it is well known that in this case the weighting matrix in short samples does not produce good results. While the literature suggests ways to address this problem, this is clearly beyond the scope of this paper.

A.5.1 The statistic and moment functions

This section gives explicit expressions for the statistic function $\mathcal{S}(\cdot)$ and the moment functions $h(\cdot)$ introduced in appendix A.5. The 7 statistics we consider can be expressed as function of eight sample moments $M_{N,i}$ ($i = 1, \dots, 8$).⁵⁴

$$\mathcal{S}(M_N) \equiv \begin{bmatrix} E(r_t^s) \\ E(PD_t) \\ \sigma_{r_t^s} \\ \sigma_{PD_t} \\ \rho_{PD_t, PD_{t-1}} \\ c_2^5 \\ R_5^2 \end{bmatrix} = \begin{bmatrix} M_{N,1} \\ M_{N,2} \\ \sqrt{M_{N,3} - (M_{N,1})^2} \\ \sqrt{M_{N,4} - (M_{N,2})^2} \\ \frac{M_{N,5} - (M_{N,2})^2}{M_{N,4} - (M_{N,2})^2} \\ c_2^5(M_N) \\ R_5^2(M_N) \end{bmatrix}$$

where the functions $c_2^5(M)$ and $R_5^2(M)$ defining the OLS and R^2 coefficients of the excess returns regressions, respectively, are

$$c^5(M_N) \equiv \begin{bmatrix} 1 & M_{N,2} \\ M_{N,2} & M_{N,4} \end{bmatrix}^{-1} \begin{bmatrix} M_{N,6} \\ M_{N,8} \end{bmatrix}$$

$$R_5^2(M_N) \equiv 1 - \frac{M_{N,7} - [M_{N,6}, M_{N,8}] c^5(M_N)}{M_{N,7} - (M_{N,6})^2}$$

The underlying sample moments are

$$M_N \equiv \begin{bmatrix} M_{1,N} \\ \cdot \\ \cdot \\ \cdot \\ M_{8,N} \end{bmatrix} \equiv \frac{1}{N} \sum_{t=1}^N h(y_t)$$

⁵⁴As in the main text, we slightly abuse notation and let, for example, $E[r^S]$ denote the sample mean of stock returns.

where $h(\cdot)$ and y_t are defined as

$$h(\mathbf{y}_t) \equiv \begin{bmatrix} r_t^s \\ PD_t \\ (r_t^s)^2 \\ (PD_t)^2 \\ PD_t PD_{t-1} \\ r_{t-20}^{s,20} \\ (r_{t-20}^{s,20})^2 \\ r_{t-20}^{s,20} PD_{t-20} \end{bmatrix}, \quad \mathbf{y}_t \equiv \begin{bmatrix} PD_t \\ D_t/D_{t-1} \\ PD_{t-1} \\ D_{t-1}/D_{t-2} \\ \vdots \\ PD_{t-19} \\ D_{t-19}/D_{t-20} \\ PD_{t-20} \end{bmatrix}$$

where $r_t^{s,20}$ denotes the stock return over 20 quarters, which can be computed using from y_t using $(PD_t, D_t/D_{t-1}, \dots, PD_{t-19}, D_{t-19}/D_{t-20})$.

A.5.2 Derivatives of the statistic function

This appendix gives explicit expressions for $\partial \mathcal{S}(\cdot)/\partial M^i$ using the statistic and moment functions stated in appendix A.5.1. Straightforward but tedious algebra shows

$$\begin{aligned} \frac{\partial \mathcal{S}_i}{\partial M_i} &= 1 && \text{for } i = 1, 2 \\ \frac{\partial \mathcal{S}_i}{\partial M_i} &= \frac{1}{2\mathcal{S}_i(M)} && \text{for } i = 3, 4 \\ \frac{\partial \mathcal{S}_i}{\partial M_j} &= \frac{-M_j}{\mathcal{S}_i(M)} && \text{for } (i, j) = (3, 1), (4, 2) \\ \frac{\partial \mathcal{S}_5}{\partial M_2} &= \frac{2M_2(M_5 - M_4)}{(M_4 - M_2^2)^2}, && \frac{\partial \mathcal{S}_5}{\partial M_5} = \frac{1}{M_4 - M_2^2}, && \frac{\partial \mathcal{S}_5}{\partial M_4} = -\frac{M_5 - M_2^2}{(M_4 - M_2^2)^2} \\ \frac{\partial \mathcal{S}_6}{\partial M_j} &= \frac{\partial c_2^5(M)}{\partial M_j} && \text{for } i = 2, 4, 6, 8 \\ \frac{\partial \mathcal{S}_7}{\partial M_j} &= \frac{[M_6, M_8] \frac{\partial c^5(M)}{\partial M_j}}{M_7 - M_6^2} && \text{for } j = 2, 4 \\ \frac{\partial \mathcal{S}_7}{\partial M_6} &= \frac{\left[c_1^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_6} \right] (M_7 - M_6^2) - 2M_6 [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2} \\ \frac{\partial \mathcal{S}_7}{\partial M_7} &= \frac{M_6^2 - [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2} \\ \frac{\partial \mathcal{S}_7}{\partial M_8} &= \frac{c_2^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_8}}{M_7 - M_6^2} \end{aligned}$$

Using the formula for the inverse of a 2x2 matrix

$$c^5(M) = \frac{1}{M_4 - M_2^2} \begin{bmatrix} M_4 M_6 - M_2 M_8 \\ M_8 - M_2 M_6 \end{bmatrix}$$

we have

$$\begin{aligned}\frac{\partial c^5(M)}{\partial M_2} &= \frac{1}{M_4 - M_2^2} \left(2M_2 c^5(M) - \begin{bmatrix} M_8 \\ M_6 \end{bmatrix} \right) \\ \frac{\partial c^5(M)}{\partial M_4} &= \frac{1}{M_4 - M_2^2} \left(-c^5(M) + \begin{bmatrix} M_6 \\ 0 \end{bmatrix} \right) \\ \frac{\partial c^5(M)}{\partial M_6} &\equiv \frac{1}{M_4 - M_2^2} \begin{bmatrix} M_4 \\ -M_2 \end{bmatrix} \\ \frac{\partial c^5(M)}{\partial M_8} &\equiv \frac{1}{M_4 - M_2^2} \begin{bmatrix} -M_2 \\ 1 \end{bmatrix}\end{aligned}$$

All remaining terms $\partial \mathcal{S}_i / \partial M_j$ not listed above are equal to zero.

A.5.3 Differentiable projection facility

As discussed in the main text, we need to introduce a feature that prevents perceived stock price growth from being higher than δ^{-1} , so as to insure a finite asset price. In addition, it is convenient for our calibration exercises if the learning scheme is a continuous and differentiable function, see the discussion in appendix A.5. We thus modify the simple discontinuous projection facility described in (21) by one that ‘phases in’ more gradually. We define

$$\beta_t^* = \beta_{t-1} + \frac{1}{\alpha_t} \left[\left(\frac{D_{t-1}}{D_{t-2}} \right)^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right] \quad (45)$$

and modify the updating scheme (27) to

$$\beta_t = \begin{cases} w(\beta_t^*) \beta^U + \beta_t^* & \text{if } \beta_t^* \leq \beta^L \\ w(\beta_t^*) \beta^U + [1 - w(\beta_t^*)] \beta^L & \text{otherwise} \end{cases} \quad (46)$$

where β^U is the upper bound on beliefs, chosen to insure that the implied PD ratio is always less than a certain upper bound U^{PD} , $\beta^L < \beta^U$ is some arbitrary level of beliefs above which the projection facility starts to operate, and $w(\cdot) : \mathbb{R}^+ \rightarrow [0, 1]$ is a weighting function. We choose the weighting function to guarantee that the updating scheme is continuously differentiable w.r.t. β_t^* and that $\lim_{\beta_t^* \rightarrow \infty} \beta_t = \beta^U$. In particular, we define

$$w(\beta_t^*) = \left[1 - \frac{\beta^U - \beta^L}{\beta_t^* + \beta^U - 2\beta^L} \right].$$

With this weighting function

$$\begin{aligned}\lim_{\beta_t^* \rightarrow (\beta^L)^-} \beta_t &= \lim_{\beta_t^* \rightarrow (\beta^L)^+} \beta_t = \beta^L \\ \lim_{\beta_t^* \rightarrow (\beta^L)^-} \frac{\partial \beta_t}{\partial \beta_t^*} &= \lim_{\beta_t^* \rightarrow (\beta^L)^+} \frac{\partial \beta_t}{\partial \beta_t^*} = 1\end{aligned}$$

In our numerical applications we choose β^U so that the implied PD ratio never exceeds $U^{PD} = 500$ and $\beta^L = \delta^{-1} - 2(\delta^{-1} - \beta^U)$, which implies that the dampening effect of the projection facility starts to come into effect for values of the PD ratio above 250.

In table 7 below we show that the modification of the projection facility to a continuous function does not significantly affect our quantitative results. Using the baseline model calibration from table 4, we report in table 7 the model moments if instead of the continuous projection facility one uses the projection facility (21) with $U^{PD} = 500$. The main difference to the baseline results consists of an increase in the volatility of the PD ratio (σ_{PD}). Clearly, this emerges because the continuous projection facility dampens beliefs already before they reach the upper bound implied by U^{PD} . Obviously, by choosing alternative values of σ and $1/\alpha_1$ one can obtain moments for the model with a non-continuous projection facility that are even closer to the moments in the data.

Statistics	Non-continuous projection facility	
	t-ratio	
$E(r^s)$	2.09	0.72
$E(PD)$	119.51	-0.42
σ_{r^s}	12.96	-0.45
σ_{PD}	81.09	-1.70
$\rho_{PD_t, PD_{t-1}}$	0.94	-1.20
c_2^5	-0.0064	0.81
R_5^2	0.2981	-1.20
Parameters:	$\sigma = 4.47$	$1/\alpha_1 = 0.0162$

Table 7: Robustness to projection facility

A.6 Convergence of least squares to RE

We show convergence directly for the general learning model with risk aversion from section 4. To obtain convergence we need bounded shocks. In particular, we assume existence of some $U^\varepsilon < \infty$ such that

$$\begin{aligned} \text{Prob}(\varepsilon_t < U^\varepsilon) &= 1 \\ \text{Prob}(\varepsilon_t^{1-\sigma} < U^\varepsilon) &= 1 \end{aligned}$$

Furthermore, we assume that the projection facility is not binding in the RE equilibrium

$$\frac{P_t^{RE}}{D_t} = \frac{\delta \beta^{RE}}{1 - \delta \beta^{RE}} < U^{PD}$$

where $\beta^{RE} = E[(a\varepsilon_t)^{1-\sigma}]$ and P_t^{RE} is the price in the RE equilibrium.

We first show that the projection facility will almost surely cease to be binding after some finite time. In a second step, we prove that β_t converges to β^{RE} from that time onwards.

The projection facility implies

$$\beta_t = \begin{cases} \beta_{t-1} + \alpha_t^{-1} \left((a\varepsilon_{t-1})^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) & \text{if } \frac{\delta a}{1 - \delta (\beta_{t-1} + \alpha_t^{-1} ((a\varepsilon_{t-1})^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}))} < U^{PD} \\ \beta_{t-1} & \text{otherwise} \end{cases} \quad (47)$$

If the lower equality applies one has $(a\varepsilon_{t-1})^{-\sigma} \frac{P_{t-1}}{P_{t-2}} \geq \beta_{t-1}$ and this gives rise to the following inequalities

$$\beta_t \leq \beta_{t-1} + \alpha_t^{-1} \left((a\varepsilon_{t-1})^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) \quad (48)$$

$$|\beta_t - \beta_{t-1}| \leq \alpha_t^{-1} \left| \left((a\varepsilon_{t-1})^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) \right| \quad (49)$$

which hold for all t . Substituting recursively in (48) for past β 's delivers

$$\begin{aligned} \beta_t &\leq \frac{1}{t-1+\alpha_1} \left(\sum_{j=0}^{t-1} (a\varepsilon_j)^{-\sigma} \frac{P_j}{P_{j-1}} + (\alpha_1 - 1) \beta_0 \right) \\ &= \underbrace{\frac{t}{t-1+\alpha_1} \left(\frac{1}{t} \sum_{j=0}^{t-1} (a\varepsilon_j)^{1-\sigma} + \frac{\alpha_1 - 1}{t} \beta_0 \right)}_{=T_1} + \underbrace{\frac{1}{t-1+\alpha_1} \left(\sum_{j=0}^{t-1} \frac{\delta \Delta\beta_j}{1 - \delta\beta_j} (a\varepsilon_j)^{1-\sigma} \right)}_{=T_2} \end{aligned} \quad (50)$$

where the second line follows from (29). Since $T_1 \rightarrow \beta^{RE}$ for $t \rightarrow \infty$ a.s., β_t will eventually be bounded away from its upper bound if we can establish $|T_2| \rightarrow 0$ a.s.. This is achieved by noting that

$$\begin{aligned} |T_2| &\leq \frac{1}{t-1+\alpha_1} \sum_{j=0}^{t-1} \frac{\delta (a\varepsilon_j)^{1-\sigma}}{1 - \delta\beta_j} |\Delta\beta_j| \\ &\leq \frac{U^\varepsilon}{t-1+\alpha_1} \sum_{j=0}^{t-1} \frac{a^{1-\sigma} \delta |\Delta\beta_j|}{1 - \delta\beta_j} \\ &\leq \frac{U^\varepsilon}{t-1+\alpha_1} \frac{a^{1-\sigma} U^{PD}}{\beta^{RE}} \sum_{j=0}^{t-1} |\Delta\beta_j| \end{aligned} \quad (51)$$

where the first inequality results from the triangle inequality and the fact that both ε_j and $\frac{1}{1-\delta\beta_j}$ are positive, the second inequality follows from the a.s. bound on ε_j , and the third inequality from the bound on the price dividend ratio insuring that $\delta\beta^{RE} (1 - \delta\beta_j)^{-1} < U^{PD}$. Next, observe that

$$(a\varepsilon_j)^{-\sigma} \frac{P_t}{P_{t-1}} = \frac{1 - \delta\beta_{t-1}}{1 - \delta\beta_t} (a\varepsilon_t)^{1-\sigma} < \frac{(a\varepsilon_t)^{1-\sigma}}{1 - \delta\beta_t} < \frac{a^{1-\sigma} U^\varepsilon U^{PD}}{\delta\beta^{RE}} \quad (52)$$

where the equality follows from (28), the first inequality from $\beta_{t-1} > 0$, and the second inequality from the bounds on ε and PD . Using result (52), equation (49) implies

$$|\beta_t - \beta_{t-1}| \leq \alpha_t^{-1} \left| (a\varepsilon_t)^{-\sigma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right| \leq \alpha_t^{-1} \left(\frac{a^{1-\sigma} U^\varepsilon U^{PD}}{\delta \beta^{RE}} + \delta^{-1} \right)$$

where the second inequality follows from the triangle inequality and the fact that $\beta_{t-1} < \delta^{-1}$. Since $\alpha_t \rightarrow \infty$ this establishes that $|\Delta\beta_t| \rightarrow 0$ and, therefore, $\frac{1}{t-1+\alpha_1} \sum_{j=0}^{t-1} |\Delta\beta_j| \rightarrow 0$. Then (51) implies that $|T_2| \rightarrow 0$ a.s. as $t \rightarrow \infty$. By taking the lim sup on both sides on (50), it follows from $T_1 \rightarrow \beta^{RE}$ and $|T_2| \rightarrow 0$ that

$$\limsup_{t \rightarrow \infty} \beta_t \leq \beta^{RE}$$

a.s.. The projection facility is thus operative infinitely often with probability zero. Therefore, there exists a set of realizations ω with measure one and a $\bar{t} < \infty$ (which depends on the realization ω) such that the projection facility does not operate for $t > \bar{t}$.

We now proceed with the second step of the proof. Consider, for a given realization ω , a \bar{t} for which the projection facility is not operative after this period. Then the upper equality in (47) holds for all $t > \bar{t}$ and simple algebra gives

$$\begin{aligned} \beta_t &= \frac{1}{t - \bar{t} + \alpha_{\bar{t}}} \left(\sum_{j=\bar{t}}^{t-1} (a\varepsilon_j)^{-\sigma} \frac{P_j}{P_{j-1}} + \alpha_{\bar{t}} \beta_{\bar{t}} \right) \\ &= \frac{t - \bar{t}}{t - \bar{t} + \alpha_{\bar{t}}} \left(\frac{1}{t - \bar{t}} \sum_{j=\bar{t}}^{t-1} (a\varepsilon_j)^{1-\sigma} + \frac{1}{t - \bar{t}} \sum_{j=\bar{t}}^{t-1} \frac{\delta \Delta\beta_j}{1 - \delta\beta_j} (a\varepsilon_j)^{1-\sigma} + \frac{\alpha_{\bar{t}}}{t - \bar{t}} \beta_{\bar{t}} \right) \end{aligned} \quad (53)$$

for $t > \bar{t}$. Equations (48) and (49) now hold with equality for $t > \bar{t}$. Similar operations as before then deliver

$$\frac{1}{t - \bar{t}} \sum_{j=\bar{t}}^{t-1} \frac{\delta \Delta\beta_j}{1 - \delta\beta_j} (a\varepsilon_j)^{1-\sigma} \rightarrow 0$$

a.s. for $t \rightarrow \infty$. Finally, taking the limit on both sides of (53) establishes

$$\beta_t \rightarrow \beta^{RE}$$

a.s. as $t \rightarrow \infty$. ■

References

- ABEL, A. B. (1990): “Asset Prices under Habit Formation and Catching Up with the Joneses,” *American Economic Review*, 80, 38–42.
- ADAM, K. (2003): “Learning and Equilibrium Selection in a Monetary Overlapping Generations Model with Sticky Prices,” *Review of Economic Studies*, 70, 887–908.
- (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *Economic Journal*, 117, 603–636.
- ALLEN, F., S. MORRIS, AND H. S. SHIN (2006): “Beauty Contest and Iterated Expectations in Asset Markets,” *Review of Financial Studies*, 19, 719–752.
- BRANCH, W., AND G. W. EVANS (2006): “Asset Return Dynamics and Learning,” University of Oregon mimeo.
- BRAY, M. M., AND D. M. KREPS (1987): “Rational Learning and Rational Expectations,” in *Arrow and the Ascent of Modern Economic Theory*, ed. by G. Feiwel, pp. 597–625. New York University Press, New York.
- BRENNAN, M. J., AND Y. XIA (2001): “Stock Price Volatility and Equity Premium,” *Journal of Monetary Economics*, 47, 249–283.
- 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.
- BULLARD, J., AND J. DUFFY (2001): “Learning and Excess Volatility,” *Macroeconomic Dynamics*, 5, 272–302.
- CAMPBELL, J. Y. (2003): “Consumption-Based Asset Pricing,” in *Handbook of Economics and Finance*, ed. by G. M. Constantinides, M. Harris, and R. Stulz, pp. 803–887. Elsevier, Amsterdam.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107, 205–251.
- CARCELES-POVEDA, E., AND C. GIANNITSAROU (2006): “Asset Pricing with Adaptive Learning,” SUNY Stony Brook and Cambridge University mimeo.
- 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.
- COGLEY, T., AND T. J. SARGENT (2006): “The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?,” *New York University mimeo*.

- LEROY, S. F., AND R. PORTER (1981): “The Present-Value Relation: Test Based on Implied Variance Bounds,” *Econometrica*, 49, 555–574.
- MANKIW, G., D. ROMER, AND M. D. SHAPIRO (1985): “An Unbiased Reexamination of Stock Market Volatility,” *Journal of Finance*, 40(3), 677–687.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93, 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 71, 421–436.
- (2005): *Irrational Exuberance: Second Edition*. Princeton University Press, Princeton, USA.
- TIMMERMANN, A. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1996): “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.