

Stock Market Volatility and Learning¹

Klaus Adam², Albert Marcet³ Juan Pablo Nicolini⁴

April 2, 2012

¹Thanks go to Davide Debortoli, Luca Dedola, George Evans, Katharina Greulich, Seppo Honkapohja, Ricardo Lagos, Erzo Luttmer, Bruce McGough, Bruce Preston, Jaume Ventura, Joachim Voth, Iván Werning and Raf Wouters for interesting comments and suggestions. We particularly thank Philippe Weil for a very interesting and stimulating discussion. Oriol Carreras, Sofía Bauducco and Davide Debortoli have supported us with outstanding research assistance. Klaus Adam acknowledges support from the ERC Grant No. 284262. Albert Marcet acknowledges support from CIRIT (Generalitat de Catalunya), Plan Nacional project ECO2008-04785/ECON (Ministry of Science and Education, Spain), CREI, Axa Research Fund, Programa de Excelencia del Banco de España 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 Federal Reserve Bank of Minneapolis or the Federal Reserve System.

²University of Mannheim and CEPR

³*Institut d'Anàlisi Econòmica CSIC, ICREA, Barcelona GSE, UAB and CEPR*

⁴Universidad Di Tella and Federal Reserve Bank of Minneapolis

Abstract

We study a standard consumption based asset pricing model with rational investors who entertain subjective prior beliefs about price behavior. Optimal behavior then dictates that investors learn about price behavior from past price observations. We show that this imparts momentum and mean reversion into the equilibrium behavior of the price dividend ratio, similar to what can be observed in the data. Estimating the model on U.S. stock price data using the method of simulated moments, we show that it can quantitatively account for the observed stock price volatility, the persistence of the price-dividend ratio, and the predictability of long-horizon returns. For reasonable degrees of risk aversion, the model also passes a formal statistical test for the overall goodness of fit, provided one excludes the equity premium from the set of moments to be matched.

JEL Class. No.: G12, D84

"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 quantitatively reproduce a variety of stylized asset pricing facts if one allows for slight deviations from rational expectations while maintaining the assumption of rational behavior.

We study a simple variant of the Lucas (1978) model with standard time separable preferences. It is well known that the asset pricing implications of this model under rational expectations (RE) are at odds with basic facts, such as the observed high persistence and volatility of the price dividend ratio, the high volatility of stock returns, the predictability of long horizon excess stock returns, and the high level of the risk premium.

We stick to Lucas' framework but allow for agents whose prior beliefs about stock price behavior deviate slightly from those assumed under RE. Investors nevertheless hold a consistent system of subjective beliefs about all payoff-relevant random variables that are beyond their control; this includes beliefs about model endogenous variables, such as competitive stock market prices, as well as model exogenous variables, such as the dividend and income processes. Given these subjective beliefs, investors maximize utility subject to their budget constraints. We call such agents 'internally rational', as they know all internal aspects of their individual decision problem and behave optimally given this knowledge. They just do not know the exact model generating stock prices or external variables more generally.

We employ this setup to relax the standard assumption that agents have perfect knowledge about how a certain history of fundamental shocks maps into a market outcome for the stock price.¹ We assume that agents express this lack of knowledge by specifying a subjective joint distribution for the behavior of stock prices, incomes and dividends over time. In such a setting,

¹Such uncertainty may arise from a lack of common knowledge of investors' price and dividend beliefs, as is explained in detail in Adam and Marcet (2011).

agents optimally update their subjective expectations about stock price behavior in the light of realized market outcomes, so that agents' stock price expectations influence stock prices and observed stock prices feed back into agents' expectations. This self-referential aspect of the model turns out to be key for generating stock price volatility of the kind that can be observed in the data.

We demonstrate the ability of the model to produce data-like behavior by deriving a number of analytical results about the behavior of stock prices that is implied by a general class of belief updating rules. Specifically, we show that learning from market outcomes imparts 'momentum' on stock prices around their RE value, which gives rise to sustained deviations of the price dividend ratio from its mean, as can be observed in the data. Such momentum arises because if agents' *expectations* about stock price growth increase in a given period, 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 through the feedback from outcomes to beliefs. At the same time, the model displays 'mean reversion' over longer horizons, so that even if subjective stock price growth expectations are very high (or very low) at some point in time, they will eventually return to fundamentals. The model thus displays price cycles of the kind described in the opening quote above.

For the quantitative evaluation of the learning model, we consider a specific system of subjective beliefs which allows for subjective prior uncertainty about the average growth rate of stock market prices. Internal rationality then dictates that agents' beliefs optimally react to the realization of market prices. In particular, agents optimally use a constant gain model of adaptive learning to update their conditional expectations of one step ahead risk-adjusted price growth, while no such updating occurs in the absence of prior uncertainty. In our empirical section we document that the quantitative fit with the data improves tremendously, if agents place a small but strictly positive weight on observed market information.

To evaluate the learning model we first consider how well the model matches asset pricing moments individually, just as many papers on stock price volatility have done. We use formal structural estimation based on the method of simulated moments (MSM) and we derive the asymptotic distribution for test statistics adapting results of Duffie and Singleton (1993). We find that the model can quantitatively match most of the moments we consider, including the volatility of stock market returns, the mean, persistence and volatility of the price dividend ratio and the evidence on excess return predictability over long-horizons. With low relative risk aversion (equal to

five) the t-statistics for these moments are all below two in one of our estimated models (see Table 4 in section 5.2).

We also perform a formal econometric test for the overall goodness of fit of our consumption based asset pricing model. As far as we know, this is a more stringent test than has been applied in previous papers matching stock price volatility, but a natural one to explore given our MSM strategy. When risk aversion is as high as in Campbell and Cochrane (1999) the model marginally fails the overall goodness of fit test at the 1% level. When considering a lower degree of relative risk aversion (equal to five) and when we exclude the risk premium by excluding the risk free rate from the estimation, the p-value of the model for the overall goodness of fit increases to more than 6%.

As we explain, this is a remarkable improvement relative to the performance of the model under RE, even though our model is extremely simple and represents a minimal deviation from the standard asset pricing framework. We find it a striking observation that the quantitative asset pricing implications of the standard model are not robust to such small departures from rational price expectations and that this non-robustness is empirically so encouraging. This suggests that allowing for such minimal departures from a strong assumption (RE) could be a promising avenue for research more generally.

The paper is organized as follows. In section 2 we discuss the related literature. Section 3 presents the stylized asset pricing facts we seek to match. We outline the asset pricing model in section 4 where we also derive analytic results showing how - for a general class of belief systems - our model can *qualitatively* deliver the stylized asset pricing facts described before in section 3. Section 5 presents our MSM estimation and testing strategy and documents that the model with subjective beliefs can *quantitatively* reproduce the stylized facts. Readers interested in obtaining a glimpse of the quantitative performance of our one parameter extension of the RE model may directly jump to Table 2 in section 5.2.

2 Related Literature

A large body of literature documents that the basic asset pricing model with time separable preferences and RE has great difficulties in matching the observed volatility of stock returns.² Models of learning have long been considered as a promising avenue to match stock price volatility.

²See Campbell (2003) for an overview.

Stock price behavior under Bayesian learning has been studied in Timmermann (1993, 1996), Brennan and Xia (2001), Cecchetti, Lam, and Mark (2000), Cogley and Sargent (2008) and Pastor and Veronesi (2003), among others. Agents in these papers learn about the dividend or income process and then set the asset price equal to the discounted expected sum of dividends. As explained in Adam and Marcet (2011), this amounts to assuming that agents know exactly how dividends and income map into prices, so that there is a rather asymmetric treatment of the issue of learning: while agents learn about how dividends and income evolve, they are assumed to know perfectly the stock price process, conditional on the realization of dividends and income. As a result, in all these models stock prices represent redundant information given what agents are assumed to know and there exists no feedback from market outcomes (stock prices) to beliefs. Since agents are then learning about exogenous processes only, their beliefs are ‘anchored’ by the exogenous processes and the volatility effects resulting from learning are generally limited when considering standard time separable preference specifications. In contrast, we largely abstract from learning about the dividend and income processes and focus on learning about stock price behavior. Price beliefs and actual price outcomes then mutually influence each other. It is precisely this self-referential nature of the learning problem that imparts momentum to expectations and is key in explaining stock price volatility.

A number of papers within the adaptive learning literature study agents who learn about stock prices. Bullard and Duffy (2001) and Brock and Hommes (1998) show that learning dynamics can converge to complicated attractors and that the RE equilibrium may be unstable under learning dynamics.³ Branch and Evans (2010) study a model where agents’ algorithm to form expectations switches depending on which of the available forecast models is performing best. Timmermann (1996) analyzes a case with self-referential learning, assuming that agents use dividends to predict future price.⁴ Marcet and Sargent (1992) also study convergence to RE in a model where agents use today’s price to forecast the price tomorrow in a stationary environment with limited information. Cárceles-Poveda and Giannitsarou (2007) assume that agents know the mean stock price and find that learning does then not significantly alter the behavior of asset prices. Chakraborty and Evans (2008) show that a model of adaptive learning can account for the forward premium puzzle in foreign exchange markets.

³Stability under learning dynamics is defined in Marcet and Sargent (1989).

⁴Timmerman reports that this form of learning delivers even lower volatility than a settings with learning about the dividend process only. It is thus crucial for our results that agents use information on past *price growth* behavior to predict future *price growth*.

We contribute relative to the adaptive learning literature by deriving the learning and asset pricing equations from internally rational investor behavior. In addition, we use formal econometric inference and testing to show that the model can quantitatively match the observed stock price volatility. Finally, our paper also shows that the key issue for matching the data is that agents learn about the mean *growth rate* of stock prices from past stock prices observations.

Other papers have studied stock prices under rational market expectations when agents have asymmetric information or asymmetric beliefs, examples include Biais, Bossaerts and Spatt (2010) and Dumas, Kurshev and Uppal (2009). In contrast to the RE literature, the behavioral finance literature seeks to understand the decision-making process of individual investors by means of surveys, experiments and micro evidence, exploring the intersection between economics and psychology, see Shiller (2005) for a non-technical summary. We borrow from this literature an interest in deviating from RE but are keen on making only a minimal deviation from the standard approach: we assume that agents behave optimally given a consistent system of subjective beliefs that is assumed to be close to the RE beliefs.

3 Facts

This section describes stylized facts of U.S. stock price data that we seek to replicate in our quantitative analysis. These observations have been extensively documented in the literature, we reproduce them here as a point of reference using a single and updated data base.⁵

Since the work of Shiller (1981) and LeRoy and Porter (1981) it has been recognized that the volatility of stock prices in the data is much higher than standard RE asset pricing models suggest, given the available evidence on the volatility of dividends. Figure 1 shows the evolution of the quarterly price dividend (PD) ratio in the United States. The PD ratio displays very large fluctuations around its sample mean (the bold horizontal line in the graph): in the year 1932, for example, the quarterly PD ratio takes on values below 30, while in the year 2000 values close to 350. The standard deviation of the PD ratio (σ_{PD}) is almost one half of its sample mean (E_{PD}). We report this feature of the data as **Fact 1** in Table 1.

Figure 1 also shows that the deviation of the PD ratio from its sample mean are very persistent, so that the first order quarterly autocorrelation of the PD ratio ($\rho_{PD,-1}$) is very high. We report this as **Fact 2** in Table 1.

⁵Details on the data sources are provided in Appendix 7.1.

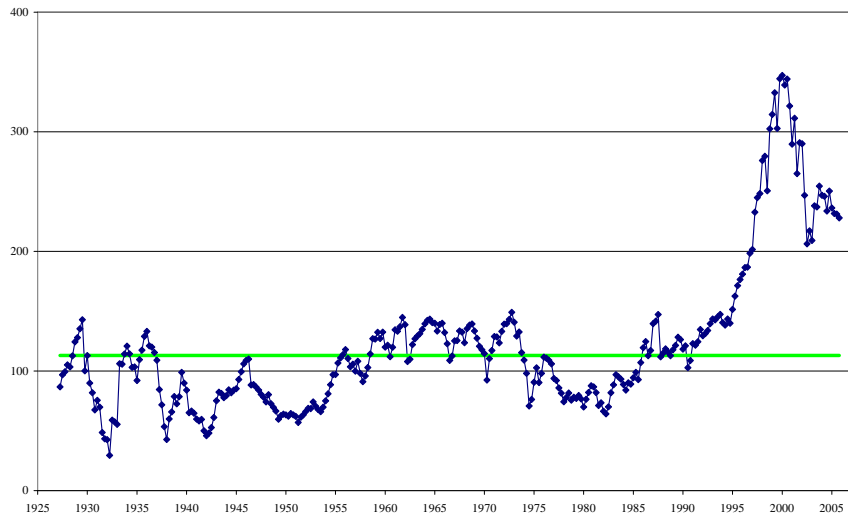


Figure 1: Quarterly U.S. price dividend ratio 1927:1-2005:4

Related to the excessive volatility of prices is the observation that the volatility of stock returns ($\sigma_{r,s}$) in the data is almost four times the volatility of dividend growth ($\sigma_{\Delta D/D}$). We report the volatility of returns as **Fact 3** in Table 1, and the mean and standard deviation of dividend growth at the bottom of the table.

While stock returns are difficult to predict at short horizons, the PD ratio helps to predict future excess stock returns in the long run. More precisely, estimating the regression

$$X_{t,n} = c_n^1 + c_n^2 PD_t + u_{t,n} \quad (1)$$

where $X_{t,n}$ is the observed real excess return of stocks over bonds from quarter t to quarter t plus n years, and $u_{t,n}$ the regression residual, the estimate c_n^2 is found to be negative, significantly different from zero, and the absolute value of c_n^2 and the R -square of this regression, denoted R_n^2 , increase with n . We choose to include the OLS regression results for the 5-year horizon as **Fact 4** in Table 1.⁶

⁶We focus on the 5-year horizon for simplicity, but obtain very similar results for other horizons. Our focus on a single horizon is justified because chapter 20 in Cochrane (2005) shows that Facts 1, 2 and 4 are closely related: up to a linear approximation, the presence of return predictability and the increase in the R_n^2 with the prediction horizon n

U.S. asset pricing facts, 1927:2-2005:4 (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	$\rho_{PD,-1}$	0.92
Fact 3	Excessive return volatility	σ_{r^s}	11.65
Fact 4	Excess return predictability	c_5^2 R_5^2	-0.0048 0.1986
Fact 5	Equity premium	E_{r^s} E_{r^b}	2.41 0.18
Dividend Behavior	Mean Growth Std of Growth	$E_{\frac{\Delta D}{D}}$ $\sigma_{\frac{\Delta D}{D}}$	0.0035 2.98

Table 1: Stylized asset pricing facts

Finally, it is well known that through the lens of standard models real stock returns tend to be too high relative to short-term real bond returns, a fact often referred to as the equity premium puzzle. We report it as **Fact 5** in Table 1, which shows that the average quarterly real return on bonds E_{r^b} is much lower than the corresponding return on stocks E_{r^s} .

Table 1 reports ten statistics. As we show in section 5, we can replicate these statistics using a model that has only four free parameters.

are *qualitatively* a joint consequence of persistent PD ratios (Fact 2) and *i.i.d.* dividend growth. It is not surprising, therefore, that our model also reproduces the increasing size of c_n^2 and R_n^2 with n . We match the regression coefficients at the 5-year horizon to check the *quantitative* model implications.

4 The Model

We describe below a Lucas (1978) asset pricing model with rationally investing agents who hold subjective prior beliefs about stock price behavior. The presence of subjective uncertainty implies that agents optimally update their beliefs about stock price behavior from observed stock price realizations. Using a generic updating mechanism, section 4.2 shows that such learning gives rise to oscillations of asset prices around their fundamental value and qualitatively contributes to reconciling the Lucas asset pricing model with the empirical evidence. Section 4.3 then introduces a specific system of prior beliefs that gives rise to constant gain learning and that we employ in our empirical work in section 5.

4.1 Model Description

The Environment: Consider an economy populated by a unit mass of infinitely-lived investors, endowed with one unit of a stock that can be traded on a competitive stock market and that pays dividend D_t consisting of a perishable consumption good. Dividends evolve according to

$$\frac{D_t}{D_{t-1}} = a \varepsilon_t^d \quad (2)$$

for $t = 0, 1, 2, \dots$, where $\log \varepsilon_t^d \sim ii\mathcal{N}(-\frac{s_d^2}{2}, s_d^2)$ and $a \geq 1$. This implies $E(\varepsilon_t^d) = 1$, $E_{\frac{\Delta D}{D}} \equiv E\left(\frac{D_t}{D_{t-1}}\right) = a$ and $\sigma_{\frac{\Delta D}{D}}^2 \equiv var\left(\frac{D_t}{D_{t-1}}\right) = e^{s_d^2} - 1$. To capture the fact that the empirically observed consumption process is considerably less volatile than the dividend process and to replicate the weak correlation between dividend and consumption growth, we assume that each agent receives in addition an endowment Y_t of perishable consumption goods. Total supply of consumption goods in the economy is then given by the feasibility constraint $C_t = Y_t + D_t$. Following the consumption-based asset pricing literature, we impose assumptions directly on the consumption supply process⁷

$$\frac{C_t}{C_{t-1}} = a \varepsilon_t^c, \quad (3)$$

where $\log \varepsilon_t^c \sim ii\mathcal{N}(-\frac{s_c^2}{2}, s_c^2)$ and $(\log \varepsilon_t^c, \log \varepsilon_t^d)$ jointly normal. In our empirical application, we follow Campbell and Cochrane (1999) and choose $s^c = \frac{1}{7}s^d$ and the correlation between $\log \varepsilon_t^c$ and $\log \varepsilon_t^d$ equal to $\rho_{c,d} = 0.2$.

⁷The process for Y_t is then implied by feasibility.

Objective Function and Probability Space: Agent $i \in [0, 1]$ has a standard time-separable expected utility function⁸

$$E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \delta^t \frac{(C_t^i)^{1-\gamma}}{1-\gamma}$$

where $\gamma \in (0, \infty)$ and C_t^i denotes consumption demand of agent i . The expectation is taken using a subjective probability measure \mathcal{P} that assigns a consistent set of probabilities to all external variables, i.e., all payoff-relevant variables that are beyond the agent's control. Importantly, C_t^i denotes the agent's consumption *demand*, while C_t denotes the total *supply* of consumption goods in the economy.

The competitive stock market assumption and the exogeneity of the dividend and income processes imply that investors consider the process for stock prices $\{P_t\}$ and the income and dividends processes $\{Y_t, D_t\}$ as exogenous to their decision problem. The underlying sample (or state) space Ω thus consists of the space of realizations for prices, dividends and income. Specifically, a typical element $\omega \in \Omega$ is an infinite sequence $\omega = \{P_t, Y_t, D_t\}_{t=0}^{\infty}$. As usual, we let Ω^t denote the set of histories from period zero up to period t and ω^t its typical element. The underlying probability space is thus given by $(\Omega, \mathcal{B}, \mathcal{P})$ with \mathcal{B} denoting the corresponding σ -Algebra of Borel subsets of Ω , and \mathcal{P} is the agent's subjective probability measure over (Ω, \mathcal{B}) . Expected utility is then defined as

$$E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \delta^t \frac{(C_t^i)^{1-\gamma}}{1-\gamma} \equiv \int_{\Omega} \sum_{t=0}^{\infty} \delta^t \frac{C_t^i(\omega^t)^{1-\gamma}}{1-\gamma} d\mathcal{P}(\omega). \quad (4)$$

Our specification of the probability space is more general than the one used in other modeling approaches because we include also price histories in the realization ω^t . Standard practice is to assume instead that agents know the exact mapping from a history of income and dividends to equilibrium asset prices $P_t(Y^t, D^t)$, so that market prices carry only redundant information. This allows - without loss of generality - to exclude prices from the underlying state space. This practice is standard in models of rational expectations, models with rational bubbles, in Bayesian RE model, and models incorporating robustness concerns. The standard practice amounts to imposing a *singularity* in the joint density over prices, income and dividends, which is equivalent to assuming that agents know *exactly* the equilibrium

⁸We assume standard preferences so as to highlight the effect of learning on asset price volatility.

pricing function $P_t(\cdot)$. While being a convenient modeling device, assuming exact knowledge of this function is at the same time very restrictive, which makes it of interest to study the implication of (slightly) relaxing it. Adam and Marcet (2011) show that rational behavior is indeed perfectly compatible with agents not knowing the exact form of the equilibrium pricing function $P_t(\cdot)$.⁹

Choice Set and Constraints: Agents make contingent plans for consumption C_t^i , bondholdings B_t^i and stockholdings S_t^i , i.e., choose the functions

$$(C_t^i, S_t^i, B_t^i) : \Omega^t \rightarrow R^3 \quad (5)$$

for all $t \geq 0$. Agent's choices are subject to the budget constraint

$$C_t^i + P_t S_t^i + B_t^i \leq (P_t + D_t) S_{t-1}^i + (1 + r_{t-1}) B_{t-1}^i + Y_t \quad (6)$$

for all $t \geq 0$, where r_{t-1} denotes the real interest rate on riskless bonds issued in period $t - 1$ and maturing in period t . The initial endowments are given by $S_{-1}^i = 1$ and $B_{-1}^i = 0$, so that bonds are in zero net supply. To avoid Ponzi schemes and to insure existence of a maximum the following bounds must hold

$$\begin{aligned} \underline{S} &\leq S_t^i \leq \bar{S} \\ \underline{B} &\leq B_t^i \leq \bar{B} \end{aligned} \quad (7)$$

for some finite bounds with the property $\underline{S} < 1 < \bar{S}$, $\underline{B} < 0 < \bar{B}$.

Maximizing Behavior (Internal Rationality): The investor's problem then consists of choosing the sequence of functions $\{C_t^i, S_t^i, B_t^i\}_{t=0}^\infty$ to maximize (4) subject to the budget constraint (6) and the asset limits (7), where all constraints have to hold for all t almost surely in \mathcal{P} . Later on, the probability measure \mathcal{P} will be constructed by specifying some perceived law of motion describing the agent's view about the evolution of (P, Y, D) over time, together with a prior distribution about the parameters governing this law of motion. Optimal behavior will then entail learning about these parameters, in the sense that agents update their posterior beliefs about the unknown parameters in the light of new price, income and dividend observations. For the moment, this learning problem remains 'hidden' in the belief structure \mathcal{P} .

⁹Specifically, with incomplete markets, i.e., in the absence of state contingent forward markets for stocks, agents cannot simply learn the equilibrium mapping $P_t(\cdot)$ by observing market prices. Furthermore, if the preferences and beliefs of agents in the economy fail to be common knowledge, then agents cannot deduce the equilibrium mapping from what they know.

Optimality Conditions: Since the objective function is concave and the feasible set is convex, the agent's optimal plan is characterized by the first order conditions

$$(C_t^i)^{-\gamma} P_t = \delta E_t^{\mathcal{P}} \left[(C_{t+1}^i)^{-\gamma} P_{t+1} \right] + \delta E_t^{\mathcal{P}} \left[(C_{t+1}^i)^{-\gamma} D_{t+1} \right] \quad (8)$$

$$(C_t^i)^{-\gamma} = \delta(1+r_t)E_t^{\mathcal{P}} \left[(C_{t+1}^i)^{-\gamma} \right] \quad (9)$$

These conditions are standard except for the fact that the conditional expectations are taken with respect to the subjective probability measure \mathcal{P} .

4.2 Asset Pricing Implications: Analytical Results

This section presents analytical results that explain why the asset pricing model with subjective beliefs can explain the asset pricing facts presented in Table 1.

As is well known, under RE the model is completely at odds with these asset pricing facts. A routine calculation shows that the unique RE solution of the model is given by

$$P_t^{RE} = \frac{\delta a^{1-\gamma} \rho_\varepsilon}{1 - \delta a^{1-\gamma} \rho_\varepsilon} D_t \quad (10)$$

where

$$\begin{aligned} \rho_\varepsilon &= E \left[(\varepsilon_{t+1}^c)^{-\gamma} \varepsilon_{t+1}^d \right] \\ &= e^{\gamma(1+\gamma)\frac{s_c^2}{2}} e^{-\gamma\rho_{c,d} s_c s_d}. \end{aligned}$$

The PD ratio is then constant, return volatility equals approximately the volatility of dividend growth and there is no (excess) return predictability, so that the model misses Facts 1 to 4 listed in Table 1. This holds independently of the parameterization of the model. Furthermore, even for very high degrees of relative risk aversion, say $\gamma = 80$, the model implies a fairly small risk premium. This emerges because of the low correlation between the innovations to consumption growth and dividend growth in the data ($\rho_{c,d} = 0.2$).¹⁰ The model thus also misses Fact 5 in Table 1.

¹⁰Under RE, the risk free rate is given by $1+r = \left(\delta a^{-\gamma} e^{\gamma(1+\gamma)\frac{s_c^2}{2}} \right)^{-1}$ and the expected equity return equals $E_t[(P_{t+1} + D_{t+1})/P_t] = (\delta a^{-\gamma} \rho_\varepsilon)^{-1}$. For $\rho_{c,d} = 0$ there is thus no equity premium, independently of the value for γ .

We now characterize the equilibrium outcome under learning. One may be tempted to argue that C_{t+j}^i can be substituted by C_{t+j} for $j = 0, 1$ in the first order conditions (8) and (9), simply because $C_t^i = C_t$ holds in equilibrium for all t .¹¹ However, outside of strict rational expectations we may have $E_t^{\mathcal{P}} [C_{t+1}^i] \neq E_t^{\mathcal{P}} [C_{t+1}]$ even if in equilibrium $C_t^i = C_t$ holds ex-post.¹² To understand how this arises, consider the following simple example: suppose agents know the aggregate process for D_t and Y_t . In this case, $E_t^{\mathcal{P}} [C_{t+1}]$ is a function only of the exogenous variables (Y^t, D^t) . At the same time, $E_t^{\mathcal{P}} [C_{t+1}^i]$ is generally a function of price realizations also, since in the eyes of the agent, optimal future consumption demand depends on future prices and, therefore, also on today's prices whenever agents are learning about price behavior. As a result, we have $E_t^{\mathcal{P}} [C_{t+1}^i] \neq E_t^{\mathcal{P}} [C_{t+1}]$, so that one cannot routinely substitute individual by aggregate consumption on the right-hand side of agent's first order conditions (8) and (9).

Nevertheless, if in any given period t the optimal plan for period $t+1$ from the viewpoint of the agent is such that $(P_{t+1}(1 - S_{t+1}^i) - B_{t+1}^i) / (Y_t + D_t)$ is expected to be small according to the agent's expectations $E_t^{\mathcal{P}}$, then C_{t+1}/C_t is very close to C_{t+1}^i/C_t^i . Hence, in this case one can rely on the approximations¹³

$$E_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right] \simeq E_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}^i}{C_t^i} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right] \quad (11)$$

$$E_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right] \simeq E_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}^i}{C_t^i} \right)^{-\gamma} \right], \quad (12)$$

The following assumption provides sufficient conditions for this to be the case:

Assumption 1 Given some asset holding constraints $\underline{S} < 1 < \bar{S}$ and $\underline{B} < 0 < \bar{B}$, we assume Y_t to be high enough so that the approximations (11) and (12) hold with sufficient accuracy.

Intuitively, for high enough income Y_t , the agent's asset trading decisions matter little for the agents' stochastic discount factor $\left(\frac{C_{t+1}^i}{C_t^i} \right)^{-\gamma}$, allowing us

¹¹ $C_t^i = C_t$ follows from market clearing and the fact that all agents are identical.

¹² This is the case because the preferences and beliefs of agents are not assumed to be common knowledge, so that agents do not know that $C_t^i = C_t$ must hold in equilibrium.

¹³ Equations (11) and (12) follow from using the budget constraint (6) to express C_{t+1}^i and by exploiting the fact that market clearing in period t implies that equilibrium choices are such that $C_t^i = C_t$, $S_t^i = 1$ and $B_t^i = 0$.

to approximate individual consumption in $t + 1$ by aggregate consumption in $t + 1$.¹⁴

With assumption 1, the risk-free interest rate solves

$$1 = \delta(1 + r_t)E_t^{\mathcal{P}} \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right], \quad (13)$$

Furthermore, defining the subjective expectations of risk-adjusted stock price growth

$$\beta_t \equiv E_t^{\mathcal{P}} \left(\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \frac{P_{t+1}}{P_t} \right) \quad (14)$$

and subjective expectations of risk-adjusted dividend growth

$$\beta_t^D \equiv E_t^{\mathcal{P}} \left(\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \frac{D_{t+1}}{D_t} \right),$$

the first order condition for stocks (8) implies that the equilibrium stock price under subjective beliefs is given by

$$P_t = \frac{\delta\beta_t^D}{1 - \delta\beta_t} D_t, \quad (15)$$

provided $\beta_t < \delta^{-1}$. The equilibrium stock price is thus increasing in (subjective) expected risk-adjusted dividend growth and also increasing in expected risk-adjusted price growth.

For the special case when agents know the RE growth rates $\beta_t = \beta_t^D = a^{1-\gamma}\rho_\varepsilon$ for all t , equation (15) delivers the RE price outcome (10). Furthermore, when agents hold subjective beliefs about risk-adjusted dividend growth but objectively rational beliefs about risk-adjusted price growth, then with β_t^D and β_t denoting the respective posterior means of these beliefs, equation (15) delivers the pricing implications derived in the Bayesian RE asset pricing literature, as surveyed at the beginning of section 2. To highlight the fact that the improved empirical performance of the present asset pricing model derives exclusively from the presence of subjective beliefs about risk-adjusted price growth, we shall entertain assumptions that are orthogonal to those made in the Bayesian RE literature. Specifically, we assume that agents know the true process for risk-adjusted dividend growth:

¹⁴Note that independent from their tightness, the asset holding constraints never prevent agents from marginally trading or selling securities in any period t along the equilibrium path, where $S_t^i = 1$ and $B_t^i = 0$ holds for all t .

Assumption 2 Agents know the process for risk adjusted dividend growth, i.e., $\beta_t^D \equiv a^{1-\gamma}\rho_\varepsilon$ for all t .

Under this assumption the asset pricing equation (15) simplifies to:¹⁵

$$P_t = \frac{\delta a^{1-\gamma}\rho_\varepsilon}{1 - \delta\beta_t} D_t. \quad (16)$$

4.2.1 Stock Price Behavior under Learning

We now derive a number of analytical results regarding the behavior of asset prices over time. We start out with a general observation about the volatility of prices and thereafter derive results about the behavior of prices over time for a general belief updating scheme.

The asset pricing equation (16) implies that fluctuations in subjective price expectations can contribute to the fluctuations in actual prices. As long as the correlation between β_t and the last dividend innovation ε_t^d is small (as occurs for the updating schemes for β_t that we consider in this paper), equation (16) implies

$$\text{var} \left(\ln \frac{P_t}{P_{t-1}} \right) \simeq \text{var} \left(\ln \frac{1 - \delta\beta_{t-1}}{1 - \delta\beta_t} \right) + \text{var} \left(\ln \frac{D_t}{D_{t-1}} \right). \quad (17)$$

The previous equation shows that even small fluctuations in subjective price growth expectations can significantly increase the variance of price growth, and thus the variance of stock price returns, if β_t fluctuates around values close to but below δ^{-1} .

To determine the behavior of asset prices over time, one needs to take a stand on how the subjective price expectations β_t are updated over time. The optimal updating rules are thereby dictated by the probability measure \mathcal{P} . In section 4.3 we present a specific probability measure \mathcal{P} that we employ in our empirical application. Yet, to illustrate that the results we obtain in our empirical application do not depend on the specific measure assumed and to improve our understanding for the empirical performance of the model, we derive the analytical results in this section for a more general nonlinear belief updating scheme.

¹⁵Some readers may be tempted to believe that entertaining subjective price beliefs while entertaining objective beliefs about the dividend process is inconsistent with individual rationality. Adam and Marcet (2011) show, however, that there exists no such contradiction, as long as the preferences and beliefs of agents in the economy are not common knowledge.

Given that β_t denotes the subjective one-step-ahead expectation of risk adjusted stock price growth, it appears natural to assume that the measure \mathcal{P} implies that rational agents revise β_t upwards (downwards) if they under-predicted (overpredicted) the risk adjusted stock price growth ex-post. This prompts us to consider measures \mathcal{P} that imply updating rules of the form¹⁶

$$\Delta\beta_t = f_t \left(\left(\frac{C_{t-1}}{C_{t-2}} \right)^{-\gamma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}; \beta_{t-1} \right) \quad (18)$$

for given non-linear updating functions $f_t : R^2 \rightarrow R$ with the properties

$$f_t(0; \beta) = 0 \quad (19)$$

$$f_t(\cdot; \beta) \text{ increasing} \quad (20)$$

$$0 < \beta + f_t(x; \beta) < \beta^U \quad (21)$$

for all (t, x, β) and for some constant $\beta^U \in (a^{1-\gamma}\rho_\varepsilon, \delta^{-1})$. Properties (19) and (20) imply that β_t is adjusted in the same direction as the last prediction error, where the strength of the adjustment may depend on the current level of beliefs, as well as on calendar time (e.g., on the number of observations available to date). Property (21) is needed to guarantee that positive equilibrium prices solving (16) always exist.

In section 4.3 below we provide an explicit system of beliefs \mathcal{P} in which agents optimally update beliefs according to a special case of equation (18). Updating rule (18) is more general and nests also other systems of beliefs, as well as a range of learning schemes considered in the literature on adaptive learning, such as the widely used least squares learning or switching gains learning used by Marcet and Nicolini (2003).

To derive the equilibrium behavior of price expectations and price realizations over time, we first use (16) to determine realized price growth

$$\frac{P_t}{P_{t-1}} = \left(a + \frac{a\delta}{1 - \delta\beta_t} \Delta\beta_t \right) \varepsilon_t^d \quad (22)$$

Combining the previous equation with the belief updating rule (18) one obtains

$$\Delta\beta_{t+1} = f_{t+1} \left(T(\beta_t, \Delta\beta_t) (\varepsilon_t^c)^{-\gamma} \varepsilon_t^d - \beta_t; \beta_t \right) \quad (23)$$

¹⁶Note that β_t is determined from observations up to period $t - 1$ only. This simplifies the analysis and it avoids simultaneity of price and forecast determination. This lag in the information is common in the learning literature. Difficulties emerging with simultaneous information sets in models of learning are discussed in Adam (2003).

where

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

Given initial conditions (D_0, P_{-1}) , and initial expectations β_0 , equation (23) completely characterizes the equilibrium evolution of the subjective price expectations β_t over time. Given that there is a one-to-one relationship between β_t and the PD ratio, see equation (16), the previous equation also characterizes the evolution of the equilibrium PD ratio under learning. High (low) price growth expectations are thereby associated with high (low) values for the equilibrium PD ratio.

The properties of the second order difference equation (23) can be illustrated in a 2-dimensional phase diagram for the dynamics of (β_t, β_{t-1}) , which is shown in Figure 2 for the case where the shocks $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d$ assume their unconditional mean value ρ_ε .¹⁷ The effects of different shock realizations for the dynamics will be discussed separately below.

The arrows in Figure 2 indicate the direction in which the vector (β_t, β_{t-1}) evolves over time according to equation (23), and the solid lines indicate the boundaries of these areas.¹⁸ Since we have a difference equation rather than a differential equation, we cannot plot the evolution of expectations exactly, as the difference equation gives rise to discrete jumps in the vector (β_t, β_{t-1}) over time. Yet, if agents update beliefs only relatively weakly in response to forecast errors, as will be the case for our estimated model later on, then these jumps will be correspondingly small. The arrows then suggest that for the case where the shocks assume their average values, the expectations are likely to move in ellipses around the rational expectations equilibrium $(\beta_t, \beta_{t-1}) = (a^{1-\gamma}\rho_\varepsilon, a^{1-\gamma}\rho_\varepsilon)$.

Consider, for example, point A in the diagram. At this point β_t is already below its fundamental value $a^{1-\gamma}\rho_\varepsilon$, but the phase diagram indicates that expectations will fall further. This shows that there is *momentum* in price changes: the fact that agents at point A have become less optimistic relative to the previous period ($\beta_t < \beta_{t-1}$) implies that price growth optimism and prices will fall further. Expectations move, for example, to point B over time where they will start to revert direction and move on to point C, then display upward momentum and move to point D, thereby displaying *mean reversion*. The elliptic movements imply that expectations (and thus the PD ratio) are likely to oscillate in sustained and persistent swings around the RE value $a^{1-\gamma}\rho_\varepsilon$.

¹⁷ Appendix 7.2 explains in detail the construction of the phase diagram.

¹⁸ The vertical solid line close to δ^{-1} is meant to illustrate the restriction $\beta < \delta^{-1}$.

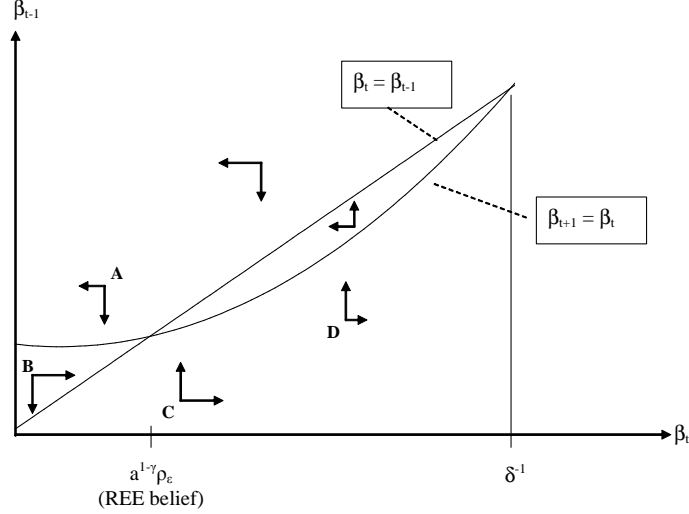


Figure 2: Phase diagram illustrating momentum and mean-reversion

The effects of the stochastic disturbances $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d$ is to shift the curve labeled " $\beta_{t+1} = \beta_t$ " in Figure 2. Specifically, for realizations $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d > \rho_\varepsilon$ this curve is shifted upwards. As a result, beliefs are more likely to increase, which is the case for all points below this curve. Conversely, for $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d < \rho_\varepsilon$ this curve shifts downward, making it more likely that beliefs decrease from the current period to the next.

The previous results show that learning causes beliefs and the PD ratio to stochastically oscillate around its RE value. Such behavior will be key in explaining the observed volatility and the serial correlation of the PD ratio, i.e., Facts 1 and 2 in Table 1. Also, from the discussion around equation (17) it should be clear that such behavior makes stock returns more volatile than dividend growth, which contributes to replicating Fact 3. As discussed in Cochrane (2005), a serially correlated and mean reverting PD ratio gives rise to excess return predictability, i.e., contributes to matching Fact 4.

The momentum of changes in beliefs around the RE value of beliefs, as well as the overall mean reverting behavior can be more formally captured in the following results:

Momentum: If $\Delta\beta_t > 0$ and

$$\beta_t \leq a^{1-\gamma} (\varepsilon_t^c)^{-\gamma} \varepsilon_t^d, \quad (24)$$

then $\Delta\beta_{t+1} > 0$. This also holds if all inequalities are reversed.

The result follows from the fact that condition (24) implies that the first argument in the f function on the right-hand side of equation (23) is positive. Note that the expected value of random variables appearing on the right-hand side of condition (24) is equal $a^{1-\gamma}\rho_\varepsilon$, which is the RE value of risk-adjusted stock price growth. Therefore, if the updating function f is sufficiently close to linear in its first argument, the previous result implies that

$$E_{t-1}[\Delta\beta_{t+1}] > 0$$

whenever $\Delta\beta_t > 0$ and $\beta_t \leq a^{1-\gamma}\rho_\varepsilon$, so that beliefs have a tendency to increase further following an initial increase, whenever beliefs are at or below the RE value. By the same token, beliefs have a tendency to decrease further following an initial decrease, provided that beliefs are above or at the RE value.

The following result shows formally that stock prices would eventually return to their (deterministic) RE value in the absence of further disturbances, and that such reverting behavior occurs monotonically:¹⁹

Mean reversion: Consider an arbitrary initial belief $\beta_t \in (0, \beta^U)$. In the absence of further disturbances

$$\limsup_{t \rightarrow \infty} \beta_t \geq a^{1-\gamma} \geq \liminf_{t \rightarrow \infty} \beta_t$$

Furthermore, if $\beta_t > a^{1-\gamma}$, there is a period $t' \geq t$ such that β_t is non-decreasing between t and t' and non-increasing between t' and t'' , where t'' is the first period where $\beta_{t''}$ is arbitrarily close to $a^{1-\gamma}$. Symmetrically, if $\beta_t < a^{1-\gamma}$.

The previous result implies that - absent any shocks - β_t cannot stay away from the RE value forever. Beliefs either converge to the (deterministic) RE value or stay fluctuating around it forever. Any initial deviation, however, is eventually eliminated with the reversion process being monotonic.

Summing up, the previous results show that for a general set of nonlinear belief updating rules, stock prices and beliefs fluctuate around their RE values in a way that helps to qualitatively account for Facts 1 to 4 listed in Table 1.

¹⁹See Appendix 7.3 for the proof under an additional technical assumption.

4.3 Optimal Belief Updating: Constant Gain Learning

We now introduce a fully specified probability measure \mathcal{P} and derive the optimal belief updating equation it implies. We employ this belief updating equation in our empirical work in section 5. As we prove below, it guarantees that prices are geometrically ergodic, a condition that is required for the MSM estimation approach to be applicable, see Duffie and Singleton (1993). In addition, we show that this system of beliefs represents a small deviation from RE.

We parameterize the measure \mathcal{P} such that the implied expectations about risk-adjusted stock price and dividend growth can be chosen to lie arbitrarily close to the RE beliefs of these variables. We shall even assume that agents know the objective distributions for the dividend and aggregate consumption processes (or alternatively for the dividend and income processes). In line with Assumption 2, agents then hold objectively rational expectations about risk-adjusted dividend growth. At the same time, we allow for subjective beliefs about risk-adjusted stock price growth by allowing agents to entertain the possibility that risk-adjusted price growth may contain a small and persistent time-varying component. This is motivated by the observation that in the data there are periods in which the PD ratio increases persistently, as well as periods in which the PD ratio falls persistently, see figure 1. In an environment with unpredictable innovations to dividend growth, this implies the existence of persistent and time-varying components in stock market returns. For this reason, we consider agents who think that the process for risk-adjusted stock price growth is the sum of a persistent component b_t and of a transitory component ε_t

$$\begin{aligned} \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \frac{P_t}{P_{t-1}} &= b_t + \varepsilon_t \\ b_t &= b_{t-1} + \xi_t \end{aligned}$$

where the persistent process b_t follows a random walk and $\varepsilon_t \sim iiN(0, \sigma_\varepsilon^2)$, $\xi_t \sim iiN(0, \sigma_\xi^2)$, independent of each other.

Besides the empirical appeal of the model, the previous setup is of theoretical interest because there is a limiting case in which it gives rise to RE. This happens when agents believe $\sigma_\xi^2 = 0$ and assign probability one to $b_0 = a^{1-\gamma}\rho_\varepsilon$. Expectations about risk-adjusted stock price growth then equals the RE value $a^{1-\gamma}\rho_\varepsilon$ in all periods.

In what follows we relax these beliefs and allow for a non-zero variance σ_ξ^2 , i.e., for the presence of a persistent time-varying component in price growth. The setup then gives rise to a learning problem because agents

observe only the realizations of risk-adjusted price growth, but not the persistent and transitory component separately. The learning problem thus consist of optimally filtering out the persistent component of price growth. Assuming that agents prior beliefs about b_0 are centered at the RE value and given by

$$b_0 \sim N(a^{1-\gamma}\rho_\varepsilon, \sigma_0^2)$$

and setting σ_0^2 equal to the steady state Kalman filter uncertainty about b_t , which is given by

$$\sigma_0^2 = \frac{-\sigma_\xi^2 + \sqrt{(\sigma_\xi^2)^2 + 4\sigma_\xi^2\sigma_\varepsilon^2}}{2},$$

agents' posterior beliefs at any time t are given by

$$b_t \sim N(\beta_t, \sigma_0)$$

with optimal updating implying that β_t , defined in equation (14), recursively evolves according to

$$\beta_t = \beta_{t-1} + \frac{1}{\alpha} \left(\left(\frac{C_t}{C_{t-1}} \right)^{-\gamma} \frac{P_t}{P_{t-1}} - \beta_{t-1} \right) \quad (25)$$

The optimal (Kalman) gain is given by $1/\alpha = (\sigma_0^2 + \sigma_\xi^2) / (\sigma_0^2 + \sigma_\xi^2 + \sigma_\varepsilon^2)$ and captures the strength with which agents optimally update their posteriors in response to surprises.²⁰

In the limiting case with vanishing innovations to the random walk process ($\sigma_\xi^2 \rightarrow 0$), agents' prior uncertainty vanishes ($\sigma_0^2 \rightarrow 0$) and the optimal gain converges to zero ($1/\alpha \rightarrow 0$). As a result, $\beta_t \rightarrow a^{1-\gamma}\rho_\varepsilon$ in distribution for all t for any *given* distribution of prices. We prove below that this convergence result also applies when we consider instead the equilibrium price distribution emerging from the beliefs with vanishing gain.

For our empirical application, we need to slightly modify the updating equation (25) to guarantee that the bound $\beta_t < \beta^U$ holds for all periods and equilibrium prices always exist. The exact way in which this bound is imposed matters little for our empirical result, because the moments we compute do not change much as long as β_t is close to β^U only rarely over

²⁰In line with equation (18) we incorporate information with a lag, so as to eliminate the simultaneity between prices and price growth expectations. The lag in the updating equation could be justified by a specific information structure where agents observe some of the lagged transitory shocks to risk-adjusted stock price growth.

the sample length considered. To impose this bound, we consider in our empirical application a concave, increasing and differentiable function $w : R_+ \rightarrow (0, \beta^U)$ and modify the belief updating equation (25) to²¹

$$\beta_t = w \left(\beta_{t-1} + \frac{1}{\alpha} \left[\left(\frac{C_{t-1}}{C_{t-2}} \right)^{-\gamma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right] \right) \quad (26)$$

where

$$w(x) = x \text{ if } x \in (0, \beta^L)$$

for some $\beta^L \in (a^{1-\gamma}\rho_\varepsilon, \beta^U)$. Beliefs thus continue to evolve according to (25), as long as they are below the threshold β^L , while for higher beliefs we have that $w(x) \leq x$. The modified algorithm (26) satisfies the constraint (21) and can be interpreted as an approximate implementation of a Bayesian updating scheme where agents have a truncated prior that puts probability zero on $b_t > \beta^U$.²² The learning setup will then give rise to a stationary and ergodic equilibrium outcome and in the limit $1/\alpha \rightarrow 0$ to small deviations from rational expectations:

Stationarity, Ergodicity, and Small Deviations from RE: *Suppose agents' posterior beliefs evolve according to equation (26) and equilibrium prices are determined according to equation (16). Then β_t is geometrically ergodic for α sufficiently large. Furthermore, as $1/\alpha \rightarrow 0$, we have $E[\beta_t] \rightarrow a^{1-\gamma}\rho_\varepsilon$ and $VAR(\beta_t) \rightarrow 0$.*

The proof is based on results from Duffie and Singleton (1993) and contained in appendix 7.4. Geometric ergodicity implies the existence of a unique stationary distribution for β_t that is ergodic and that is reached from any initial condition. This justifies the MSM approach to estimation pursued in section 5. The previous result also shows that risk-adjusted stock price growth expectations have a distribution that is increasingly centered at the RE value $a^{1-\gamma}\rho_\varepsilon$ as the gain parameter becomes vanishingly small.

²¹The exact functional form for w can be found in appendix 7.5.

²²The issue of bounding beliefs so as to insure that expected utility remains finite is present in many applications of both Bayesian and adaptive learning to asset prices. The literature has typically dealt with this issue by assuming that agents simply *ignore* observations that would imply updating beliefs beyond the required bound, see Timmermann (1993, 1996), Marcet and Sargent (1989), or Evans and Honkapohja (2001). This approach, however, introduces a discontinuity in the simulated moments and creates difficulties for our MSM estimation in section 5, prompting us to pursue the differentiable approach to bounding beliefs described above.

From equation (16) it then follows that actual equilibrium prices also become increasingly concentrated at their RE value, so that the difference between beliefs and outcomes becomes vanishingly small as $1/\alpha \rightarrow 0$. This proves that for small values of $1/\alpha$ agent’s conditional expectations about risk-adjusted price growth deviate only slightly from RE.

5 Quantitative Model Performance

This section evaluates the quantitative performance of our asset pricing model with subjective price beliefs and shows that it can robustly replicate Facts 1 to 4 listed in Table 1. We formally estimate and test the model using the method of simulated moments (MSM). This approach to structural estimation and testing helps us focusing on the ability of the model to explain the specific moments of the data described in Table 1.²³

While the model gives rise to an equity premium (Fact 5 in Table 1), it tends to fall short of replicating the observed magnitude of this premium for reasonable degrees of relative risk aversion. This occurs even though the predicted equity premium is much higher than under RE. To document that the ability of the model to replicate Facts 1 to 4 in Table 1 is not sensitive to the assumed degree of risk aversion, we consider a setting with a very high degree of risk aversion, asking the model then to replicate also the risk premium, and a setting with a lower value, in which case we leave the risk premium out of the estimation targets. For the high value of relative risk aversion we choose $\gamma = 80$, which is the steady state value of relative risk aversion used in Campbell and Cochrane (1999).²⁴ The model then replicates all moments in Table 1, including the risk premium, but marginally fails a formal statistical test for the overall goodness of fit. For low risk aversion we choose $\gamma = 5$. The model then replicates Facts 1 to 4 in Table 1. It also passes a formal statistical test for the overall goodness of fit when we do not include the risk free rate, and it explains all individual moments when we exclude stock returns instead.

²³A popular alternative approach in the asset pricing literature has been to test if agents’ first order conditions hold in the data. Hansen and Singleton (1982) pioneered this approach for RE models and Bossaerts (2003) provides an approach that can be applied to models of learning. We pursue the MSM estimation approach here because it naturally provides additional information on how the formal test for goodness of fit of the model relates to the model’s ability to match the moments of interest. The results are then easily interpretable, they point out which parts of the model fit well and which parts do not, thus providing intuition about possible avenues for improving the model fit.

²⁴This value is reported on p.244 in their paper.

The next section explains the method of simulated moments (MSM) approach for estimating the model and the formal statistical test for evaluating the goodness of fit. The subsequent section reports on the estimation and test outcomes.

5.1 MSM Estimation and Statistical Test

This section outlines the MSM approach and the formal test for evaluating the fit of the model. All technical details are contained in appendix 7.6.

For a given value of the coefficient of relative risk aversion, there are four free parameters left in the model, comprising the discount factor δ , the gain parameter $1/\alpha$, and the mean and standard deviation of dividend growth, denoted by a and $\sigma_{\frac{\Delta D}{D}}$, respectively. We summarize these in the parameter vector

$$\theta \equiv \left(\delta, 1/\alpha, a, \sigma_{\frac{\Delta D}{D}} \right).$$

Throughout the paper we restrict consideration to discount factors satisfying $\delta < 1$. The *four* parameters will be chosen so as to match the *ten* sample moments in Table 1:²⁵

$$\left(\widehat{E}_{rs}, \widehat{E}_{PD}, \widehat{\sigma}_{rs}, \widehat{\sigma}_{PD}, \widehat{\rho}_{PD,-1}, \widehat{c}_2^5, \widehat{R}_5^2, \widehat{E}_{rb}, \widehat{E}_{\frac{\Delta D}{D}}, \widehat{\sigma}_{\frac{\Delta D}{D}} \right) \quad (27)$$

Let $\widehat{\mathcal{S}}_N \in R^s$ denote the subset of sample moments in (27) that will be matched in the estimation, with N denoting the sample size and $s \leq 10$.²⁶ Furthermore, let $\widetilde{\mathcal{S}}(\theta)$ denote the moments implied by the model for some parameter value θ .²⁷ The MSM parameter estimate $\widehat{\theta}_N$ is defined as

$$\widehat{\theta}_N \equiv \arg \min_{\theta} \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\theta) \right]' \widehat{\Sigma}_{\mathcal{S},N}^{-1} \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\theta) \right] \quad (28)$$

where $\widehat{\Sigma}_{\mathcal{S},N}$ is an estimate of the variance-covariance matrix of the sample moments $\widehat{\mathcal{S}}_N$. The MSM estimate $\widehat{\theta}_N$ chooses the model parameter such that the model moments $\widetilde{\mathcal{S}}(\theta)$ fit the observed moments $\widehat{\mathcal{S}}_N$ as close as possible in terms of a quadratic form with weighting matrix $\widehat{\Sigma}_{\mathcal{S},N}^{-1}$. Appendix 7.6 explains how we estimate $\widehat{\Sigma}_{\mathcal{S},N}$ from the data. Adapting standard results

²⁵Many elements listed in (27) are not sample moments but they are non-linear functions of sample moments. This generates some technical complications, which are discussed in appendix 7.6. It would be more precise to refer to the elements in (27) as ‘sample statistics’, as we do in the appendix. For simplicity we avoid this terminology in the main text.

²⁶As discussed before, we sometimes exclude the risk premium from the estimation.

²⁷Note that the smooth bounding function w in equation (26) guarantees that a Monte-Carlo approximation to $\widetilde{\mathcal{S}}(\theta)$ is differentiable.

from MSM, we have the following results: for a given the list of moments included in $\widehat{\mathcal{S}}_N$, the estimate $\widehat{\theta}_N$ is consistent and the best estimate amongst those obtained with different weighting matrices.

The MSM estimation approach also provides an overall test of the model. Under the null hypothesis that the model is correct, we have

$$\widehat{W}_N \equiv N \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\widehat{\theta}_N) \right]' \widehat{\Sigma}_{\mathcal{S},N}^{-1} \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\widehat{\theta}_N) \right] \rightarrow \chi_{s-4}^2 \text{ as } N \rightarrow \infty \quad (29)$$

where convergence is in distribution. Furthermore, we obtain a proper asymptotic distribution for each element of the deviations $\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\widehat{\theta}_N)$, so that we can build t-statistics that indicate which moments are better matched in the estimation, see appendix 7.6 for details.

All asymptotic distribution results hinge on the assumption that $\widehat{\Sigma}_{\mathcal{S},N}$ is invertible, as the inverse of this matrix appears in equations (28) and (29). In our practical implementation of the MSM approach, we find that the estimate $\widehat{\Sigma}_{\mathcal{S},N}$ obtained from the data happens to be nearly singular. The distribution of \widehat{W}_N in small samples is then not well approximated by the asymptotic chi-square statistic, as some moments are almost exact linear combinations of other moments. These moments carry almost no additional information, while the estimation procedure would ask them to be fitted exactly. To eliminate this near singularity in $\widehat{\Sigma}_{\mathcal{S},N}$ we drop one additional moment from the list of moments (27), using a formal procedure described in detail in appendix 7.6. The procedure picks the moment that is most highly correlated with the remaining moments and suggests that we drop the coefficient from the five year ahead excess return regression \widetilde{c}_2^5 . In the empirical section below, the value of the regression coefficient implied by the estimated model is always such that the t-statistic for this moment remains below one. This happens despite the fact that information about \widetilde{c}_2^5 has not been used in the estimation.

5.2 Estimation Results

This section presents the estimation results and test outcomes from using the MSM approach described in the preceding section.

Table 2 reports the estimation outcomes when we assume a high degree of risk aversion ($\gamma = 80$, in line with Campbell and Cochrane (1999)). We then use for estimation all asset pricing moments listed in equation (27), with the exception of the coefficient c_2^5 from the excess return regression.²⁸

²⁸ As discussed in the previous section, c_2^5 is excluded from the estimation to avoid the near-singularity of the weighting matrix.

The results in table 2 show that the estimated learning model successfully replicates all moments in the data, including the risk-premium. All the t-statistics for the individual moments are below two, with most of them even assuming values below one. The model also quantitatively replicates the value of the regression coefficient c_2^5 , which has not been included in the estimation.²⁹

Figure 3 shows some realizations of the time series outcomes for the PD ratio generated from simulating the estimated model for the same number of quarters as numbers of observations in our data sample. The time series looks similar to that of the data, see figure 1, so that the model also passes an informal ‘eyeball test’.

The discount factor estimated in Table 2 is close to but below one and the estimated gain coefficient is small. The latter implies that agents’ risk-adjusted return expectations respond only 0.21% in the direction of the last observed forecast error, suggesting that the system of price beliefs in our model represents indeed only a small deviation from RE beliefs. Under strict RE the reaction to forecast errors is assumed to be zero, but the model then provides a very bad match with the data. It counterfactually implies $\sigma_{r^s} \approx \sigma_{\Delta D/D}$, $\sigma_{PD} = 0$, and $R_5^2 = 0$.³⁰ The model then also cannot simultaneously match E_{r^s} , $E_{\Delta D/D}$ and E_{PD} , unlike is the case with the learning model, see Table 2. This follows from the fact that the PD ratio is constant under RE so that the mean asset return can be expressed as

$$E\left[\frac{P_{t+1} + D_{t+1}}{P_t}\right] = \frac{PD + 1}{PD} E\left[\frac{D_{t+1}}{D_t}\right]$$

If the RE model replicates the sample averages of the PD ratio and dividend growth in the data, the previous equation implies an average quarterly real stock return of 0.35%, while the data average for this moment is 2.41%.³¹

Interestingly, the learning model also gives rise to a significantly larger

²⁹The model also behaves well along other dimensions that have not been included in the estimation, e.g., the model implied autocorrelation of quarterly returns is 0.14. The corresponding value in the data equals -0.06, which is estimated with a standard deviation of 0.13, so that the t-statistic for this moment equals 1.59.

³⁰Due to the constant PD ratio, the autocorrelation of the PD ratio and the regression coefficient c_2^5 would not be well defined under RE.

³¹Exploiting the fact that the average PD ratio and average dividend growth rate are estimated with uncertainty does not improve the situation much: assuming a model implied RE PD ratio that is two standard deviations below and a mean dividend growth rate that is two standard deviations above their respective data averages, delivers an upper bound for the quarterly real stock return of 0.74%, which still falls significantly short of the sample average.

risk premium than its RE counterpart.³² For the estimated parameter values in Table 2, the quarterly real risk premium under RE is just 0.5%, which falls short of the 2.08% emerging in the model with learning.³³ The learning model generates a higher equity premium than the RE model independently of the assumed degree of risk aversion. We provide an explanation for this outcome at the end of this section.

	US data		Estimated model (c_2^5 not included)	
	Data moment $\widehat{S}_{N,i}$	Std. dev. $\widehat{\sigma}_{\widehat{S}_i}$	Model moment $\widetilde{S}_i(\widehat{\theta})$	t-statistic
Mean stock return E_{r^s}	2.41	0.45	2.16	0.57
Mean bond return E_{r^b}	0.18	0.23	0.08	0.47
Mean PD ratio E_{PD}	113.20	15.15	115.85	-0.18
Std.dev. stock return σ_{r^s}	11.65	2.88	16.99	-1.89
Std.dev. PD ratio σ_{PD}	52.98	16.53	70.93	-1.09
Autocorrel. PD ratio $\rho_{PD,-1}$	0.92	0.02	0.95	-1.52
Excess return reg. coefficient c_5^2	-0.0048	0.002	-0.0063	0.73
R ² of excess return regression R_5^2	0.1986	0.082	0.2516	-0.06
Mean dividend growth $E_{\Delta D/D}$	0.35	0.19	0.18	0.98
Std. dev. dividend growth $\sigma_{\Delta D/D}$	2.98	0.84	4.67	-2.02
Discount factor $\widehat{\delta}_N$			0.9986	
Gain coefficient $1/\widehat{\alpha}_N$			0.0021	
Test statistic \widehat{W}_N			17.23	

Table 2: Estimation outcome for $\gamma = 80$

Our MSM approach naturally provides a measure for the overall goodness of fit via the test statistic \widehat{W}_N . To the best of our knowledge, previous papers seeking to describe stock price volatility do not consider such a formal test. While the learning model is able to fit all moments individually and also significantly improves upon the empirical performance of model under RE, the model marginally fails an overall test of the goodness of fit. Table

³²The RE counterpart is the model with the same parameterization, except for $1/\alpha = 0$.

³³The learning model and the RE model imply the same risk free rate, as we assumed that agents have objective beliefs about the aggregate consumption and dividend process.

2 reports the value for the test statistic \widehat{W}_N . It has an χ_5^2 asymptotic distribution and the p-value for the reported value is 0.41%, so that the model would marginally fail a test at the 1% significance level. Therefore, even if each asset pricing moment can be matched individually, some of the joint deviations observed in the data are unlikely to happen if the model is true. This shows that the \widehat{W}_N test statistic is a much stricter test than that imposed by matching moments individually.

We now show that the model's ability to replicate many aspects of the behavior of stock prices does not hinge on the assumed high degree of risk aversion. Table 3 below reports estimation outcomes when assuming $\gamma = 5$. The fourth and fifth column in the table report the outcome when we again seek to match all moments, except for c_2^5 . With the exception of the average stock return, the t-statistics for the individual moments are then again all below two and half of the t-statistics are below one. The failure of the model to match the equity premium causes the \widehat{W}_N statistic to increase significantly relative to the case with a high degree of risk aversion. To show that the equity premium is indeed the source of the difficulty for passing the overall goodness of fit test, columns 6 and 7 in Table 3 report the estimation outcome obtained when excluding the risk-free rate $E_{r,b}$ from the estimation. The t-statistics for the majority of included moments then decreases and the model comfortably passes the overall goodness of fit test: the p-value for the reported \widehat{W}_N statistic is 6.31%.³⁴ We conclude that the model gives a very good fit of all moments in Table 1 even for low risk aversion with the exception of the equity premium. The model nevertheless generates a sizeable equity premium, but has difficulties to quantitatively match this aspect of the data precisely.

³⁴The asymptotic distribution of \widehat{W}_N is now χ_4^2 , as there is one moment less in the estimation and the statistic is computed using only moments that are included in the estimation.

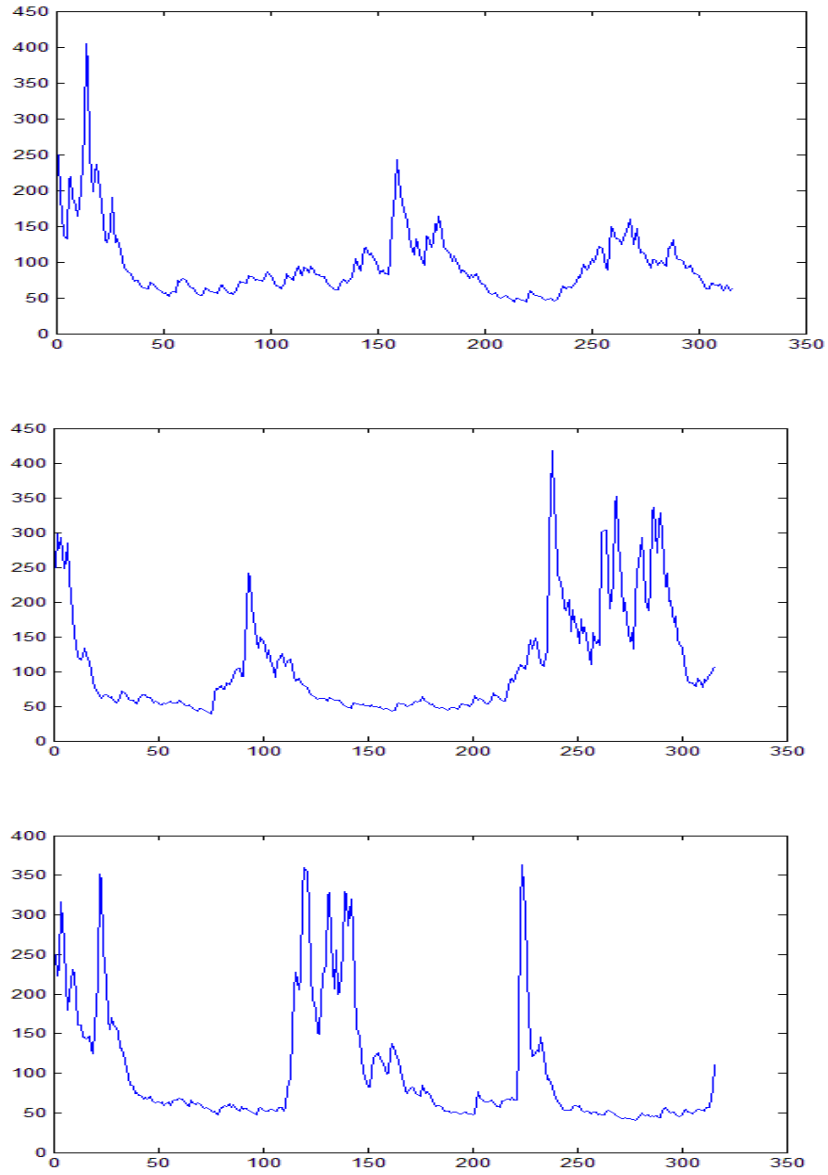


Figure 3: Simulated paths for the PD ratio, estimated model ($\gamma = 80$)

	US data		Estimated model (c_2^5 not included)		Estimated model (c_2^5, E_{r^b} not included)	
	Data moment $\widehat{S}_{N,i}$	Std. dev. $\widehat{\sigma}_{\widehat{S}_i}$	Model moment $\widetilde{S}_i(\widehat{\theta})$	t-statistic	Model moment $\widetilde{S}_i(\widehat{\theta})$	t-statistic
E_{r^s}	2.41	0.45	1.25	2.58	1.34	2.38
E_{r^b}	0.18	0.23	0.42	-1.07	0.88	-3.07
E_{PD}	113.20	15.15	123	-0.68	116	-0.20
σ_{r^s}	11.65	2.88	10.59	0.37	7.86	1.32
σ_{PD}	52.98	16.53	67.23	-0.86	53.90	-0.06
$\rho_{PD,-1}$	0.92	0.02	0.96	-1.84	0.95	-1.68
c_5^2	-0.0048	0.002	-0.0067	0.94	-0.0055	0.36
R_5^2	0.1986	0.082	0.23	-0.36	0.21	-0.18
$E_{\Delta D/D}$	0.35	0.19	0.03	1.73	0.17	0.95
$\sigma_{\Delta D/D}$	2.98	0.84	2.16	0.99	2.10	1.04
Discount factor $\widehat{\delta}_N$			0.9972		0.9999	
Gain coefficient $1/\widehat{\alpha}_N$			0.0073		0.0070	
Test statistic \widehat{W}_N			57.02		8.92	

Table 3: Estimation results for $\gamma = 5$

We have performed a range of robustness tests using a number of variations of the model and of the estimation strategy. For example, we used the mean bond return in the estimation and dropped the mean stock return instead, assuming as in Table 3 that $\gamma = 5$. Results for this case are shown in Table 4. All t-statistics then assume values well below 2, including the t-statistic for the mean stock return, which has not been used in the estimation. This estimation outcome performs better than that in many previous papers matching moments on stock price volatility, as it successfully matches individual moments with a relatively low degree of risk aversion. Nevertheless, the more stringent test statistic \widehat{W}_N increases to a value of 35 in this

setting, implying that the overall fit of the model is rejected.

	US data		Estimated model (c_2^5, E_{r^s} not included)	
	Data moment	Std. dev.	Model moment	t-statistic
	$\widehat{S}_{N,i}$	$\widehat{\sigma}_{\widehat{S}_i}$	$\widetilde{S}_i(\widehat{\theta})$	
E_{r^s}	2.41	0.45	1.61	1.79
E_{r^b}	0.18	0.23	0.45	-1.19
E_{PD}	113.20	15.15	113.74	-0.04
σ_{r^s}	11.65	2.88	12.92	-0.44
σ_{PD}	52.98	16.53	72.32	-1.17
$\rho_{PD,-1}$	0.92	0.02	0.94	-1.08
c_5^2	-0.0048	0.002	-0.0067	0.96
R_5^2	0.1986	0.082	0.1681	0.36
$E_{\Delta D/D}$	0.35	0.19	0.06	1.55
$\sigma_{\Delta D/D}$	2.98	0.84	2.89	0.10
Discount factor $\widehat{\delta}_N$			0.9986	
Gain coefficient $1/\widehat{\alpha}_N$			0.0080	
Test statistic \widehat{W}_N			35.7	

Table 4: Matching bond returns, $\gamma = 5$

Furthermore, we considered a setting where agents also learn about risk-adjusted dividend growth, using the same weight $1/\alpha$ for the learning mechanism about dividend and price growth rates. The moments of the model and t-statistics are then virtually unchanged relative to a setting without learning about dividend growth. We find the same outcome when we use a model of learning that switches between ordinary least squares learning and constant gain learning, as in Marcet and Nicolini (2003). We also used different values for the coefficient of relative risk aversion γ . For lower values of relative risk aversion around 2, we find that the model continues to generate volatile stock prices, but not enough to quantitatively match the data.

We now briefly discuss why our model is also able to generate a sizable risk premium for stocks. Surprisingly, the model generates a small ex-post risk premium for stocks even when investors are risk neutral ($\gamma = 0$). To understand this feature, note that the realized gross stock return between

period 0 and period N can be written as the product of three terms

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

The first term (R_1) is independent of the way prices are formed, thus cannot contribute to explaining the emergence of an equity premium in the model with learning. The second term (R_2), which is the ratio of the terminal over the starting value of the PD ratio could potentially generate an equity premium but is on average below one in our simulations of the learning model, while it is slightly larger than one under RE.³⁵ The equity premium in the learning model must thus be due to the last component (R_3). This term is convex in the PD ratio, so that a model that generates higher volatility of the PD ratio (but the same mean value) will also give rise to a higher equity premium. Therefore, because our learning model generates a considerably more volatile PD ratio, it also gives rise to a larger ex-post risk premium.

6 Conclusions and Outlook

A very simple consumption based asset pricing model is able to quantitatively replicate a number of important asset pricing facts, provided one slightly relaxes the assumption that agents perfectly know how stock prices are formed in the market. We assume that agents formulate their doubts about market outcomes using a consistent set of subjective beliefs about prices which is close - but not equal - to the RE prior beliefs routinely assumed in the literature. Optimal behavior then dictates that agents learn about the equilibrium price process from past price observations. This gives rise to a self-referential model of learning about prices that imparts momentum and mean reversion behavior into the price dividend ratio. As a result, sustained departures of asset prices from their fundamental value emerge, even though all agents act rationally in the light of their beliefs.

We submit our consumption based asset pricing model also to a formal econometric test based on the method of simulated moments. The model performs remarkably well despite its simplicity. While the model gives rise to a significant equity premium, it fails to match quantitatively the premium found in the data for reasonable degrees of relative risk aversion. When risk aversion is sufficiently high, and as high as in some of the previous work, the

³⁵For the learning model we choose the RE PD ratio as our starting value.

model can also replicate the equity premium, but we leave a full treatment of this issue to future research.

Given the difficulties documented in the empirical asset pricing literature in accounting for stock price volatility under RE, our results suggest that models of learning may be economically more relevant than previously thought. Indeed, the most convincing case for models of learning can be made by explaining facts that appear ‘puzzling’ from the RE viewpoint, as we attempt to do in this paper.

The finding that large asset price fluctuations can result from individually rational fluctuations in investor optimism and pessimism is also relevant from a policy perspective. The desirability of policy responding to asset price fluctuations will depend to a large extent on whether or not asset price fluctuations are fundamentally justified.

References

- ADAM, K. (2003): “Learning and Equilibrium Selection in a Monetary Overlapping Generations Model with Sticky Prices,” *Review of Economic Studies*, 70, 887–908.
- ADAM, K., AND A. MARCET (2011): “Internal Rationality, Imperfect Market Knowledge and Asset Prices,” *Journal of Economic Theory*, 146, 1224–1252.
- BIAIS, B., P. BOSSAERTS, AND C. SPATT (2010): “Equilibrium Asset Pricing and Portfolio Choice Under Asymmetric Information,” *Review of Financial Studies*, 23, 1503–1543.
- BOSSAERTS, P. (2003): “Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory,” *Review of Economic Studies*, 70, 1–24.
- BRANCH, W., AND G. W. EVANS (2010): “Asset Return Dynamics and Learning,” *Review of Financial Studies*, 23, 1651–1680, University of Oregon mimeo.
- 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 (2007): “Asset Pricing with Adaptive Learning,” *Review of Economic Dynamics* (forthcoming).

- 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.
- CHAKRABORTY, A., AND G. EVANS (2008): “Can Perpetual Learning Explain the Forward Premium Puzzle?,” *Journal of Monetary Economics*, 55(3), 477–490.
- COCHRANE, J. H. (2005): *Asset Pricing (Revised Edition)*. Princeton University Press, Princeton, USA.
- COGLEY, T., AND T. J. SARGENT (2008): “The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?,” *Journal of Monetary Economics*, 55, 454–476.
- DUFFIE, D., AND K. J. SINGLETON (1993): “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 61, 929–952.
- DUMAS, B., A. KURSHEV, AND R. UPPAL (2009): “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility,” *Journal of Finance*, 64, 579–629.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- HANSEN, L. P., AND K. J. SINGLETON (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Model,” *Econometrica*, 50, 1269–1286.
- LEROY, S. F., AND R. PORTER (1981): “The Present-Value Relation: Test Based on Implied Variance Bounds,” *Econometrica*, 49, 555–574.
- LUCAS, R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 1426–1445.
- 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.
- (1992): “The Convergence of Vector Autoregressions to Rational Expectations Equilibria,” in *Macroeconomics: A Survey of Research*

Strategies, ed. by A. Vercelli, and N. Dimitri, pp. 139–164. Oxford University Press, Oxford.

PASTOR, L., AND P. VERONESI (2003): “Stock Valuation and Learning About Profitability,” *Journal of Finance*, 58(5), 1749–1789.

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.

7 Appendix

7.1 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 data set corresponds very closely to Campbell’s (2003) handbook data set available at <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 (c_5^2 and R_5^2) we can only use data points up to 2000:4. For consistency the effective sample end for all other moments 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 7.6, 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.

7.2 Details on the phase diagram

The second order difference equation (23) describes the evolution of beliefs over time and allows to construct the directional dynamics in the (β_t, β_{t-1}) plane, as shown in Figure 2 for the case $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d = 1$. Here we show the algebra leading to the arrows displayed in this figure as well the effects of realizations $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d \leq 1$. Define $x_t' \equiv (x_{1,t}, x_{2,t}) \equiv (\beta_t, \beta_{t-1})$. The dynamics can then be described by

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

The points in Figure 2 where there is no change in each of the elements of x are the following: we have $\Delta x_2 = 0$ at points $x_1 = x_2$, so that the 45° line gives the point of no change in x_2 , and $\Delta x_2 < 0$ above this line. We have $\Delta x_1 = 0$ for $x_2 = \frac{1}{\delta} - \frac{x_1(1-\delta x_1)}{a^{1-\gamma} \delta (\varepsilon^c)^{-\gamma} \varepsilon^d}$. For $(\varepsilon^c)^{-\gamma} \varepsilon^d = \rho_\varepsilon$ this is the curve labelled " $\beta_{t+1} = \beta_t$ " in Figure 2, and we have $\Delta x_1 > 0$ below this curve. So for $(\varepsilon^c)^{-\gamma} \varepsilon^d = \rho_\varepsilon$, the zeroes for Δx_1 and Δx_2 intersect are at $x_1 = x_2 = a^{1-\gamma} \rho_\varepsilon$ which is the REE value and, interestingly, at $x_1 = x_2 = \delta^{-1}$ which is the limit of rational bubble equilibria. These results give rise to the directional dynamics shown in figure 2. Finally, for $(\varepsilon^c)^{-\gamma} \varepsilon^d > \rho_\varepsilon$ ($(\varepsilon^c)^{-\gamma} \varepsilon^d < \rho_\varepsilon$) the curve " $\beta_{t+1} = \beta_t$ " in Figure 2 is shifted upwards (downwards) as indicated by the function $x_2 = \frac{1}{\delta} - \frac{x_1(1-\delta x_1)}{a^{1-\gamma} \delta (\varepsilon^c)^{-\gamma} \varepsilon^d}$.

7.3 Proof of mean reversion

To prove mean reversion for the general learning scheme (18) we need the following additional technical assumption on the updating function f_t :

Assumption A1 *There is a $\bar{\eta} > 0$ such that $f_t(\cdot, \beta)$ is differentiable in the interval $(-\bar{\eta}, \bar{\eta})$ for all t and all β*

Furthermore, letting

$$\mathcal{D}_t \equiv \inf_{\bar{\Delta} \in (-\bar{\eta}, \bar{\eta}), \beta \in (0, \beta^U)} \frac{\partial f_t(\bar{\Delta}, \beta)}{\partial \Delta},$$

we have

$$\sum_{t=0}^{\infty} \mathcal{D}_t = \infty$$

This is satisfied by all the updating rules considered in this paper and by most algorithms used in the stochastic control literature. For example, it is guaranteed in the OLS case where $\mathcal{D}_t = 1/(t + \alpha_1)$ and in the constant gain where $\mathcal{D}_t = 1/\alpha$ for all t, β . The assumption would fail and $\sum \mathcal{D}_t < \infty$, for example, if the weight given to the error in the updating scheme is $1/t^2$. In that case beliefs could get ‘stuck’ away from the fundamental value simply because updating of beliefs ceases to incorporate new information for t large enough. In this case, the growth rate would be a certain constant but agents would forever believe that the growth rate is another constant, different from the truth. Hence in this case agents would make systematic mistakes forever. Therefore, assumption A1 is likely to be satisfied by any system of beliefs that puts a "grain of truth" on the RE equilibrium.

The statement about limsup is equivalent to saying that if $\beta_t > a$ in some period t , then for any $\eta > 0$ sufficiently small, there is a finite period $t' > t$ such that $\beta_{t'} < a + \eta$.

Fix $\eta > 0$ such that $\eta < \min(\bar{\eta}, (\beta_t - a)/2)$ where $\bar{\eta}$ is as in assumption A1.

We first prove that there exists a finite $t' \geq t$ such that

$$\Delta\beta_{\tilde{t}} \geq 0 \text{ for all } \tilde{t} \text{ such that } t < \tilde{t} < t', \text{ and} \quad (30)$$

$$\Delta\beta_{t'} < 0 \quad (31)$$

To prove this, choose $\epsilon = \eta(1 - \delta\beta^U)$. Since $\beta_t < \beta^U$ and $\epsilon > 0$ it is impossible that $\Delta\beta_{\tilde{t}} \geq \epsilon$ for all $\tilde{t} > t$. Let $\bar{t} \geq t$ to be the *first* period where $\Delta\beta_{\bar{t}} < \epsilon$.

There are two possible cases: either *i)* $\Delta\beta_{\bar{t}} < 0$ or *ii)* $\Delta\beta_{\bar{t}} \geq 0$.

In case *i)* we have that (30) and (31) hold if we take $t' = \bar{t}$.

In case *ii)* β_t can not decrease between t and \bar{t} so that

$$\beta_{\bar{t}} \geq \beta_t > a + \eta$$

Furthermore, we have

$$\begin{aligned} T(\beta_{\bar{t}}, \Delta\beta_{\bar{t}}) &= a + \frac{\Delta\beta_{\bar{t}}}{1 - \delta\beta_{\bar{t}}} < a + \frac{\epsilon}{1 - \delta\beta_{\bar{t}}} \\ &< a + \frac{\epsilon}{1 - \delta\beta^U} = a + \eta \end{aligned}$$

where the first equality follows from the definition of T in the main text. The previous two relations imply

$$\beta_{\bar{t}} > T(\beta_{\bar{t}}, \Delta\beta_{\bar{t}})$$

Therefore

$$\Delta\beta_{\bar{t}+1} = f_{\bar{t}+1}(T(\beta_{\bar{t}}, \Delta\beta_{\bar{t}}) - \beta_{\bar{t}}; \beta_{\bar{t}}) < 0$$

and in case *ii*) we have that (30) and (31) hold for $t' = \bar{t} + 1$.

This shows that (30) and (31) hold for a finite t' as in the first part of the statement of Mean Reversion in the text. Now we need to show that beliefs eventually fall below $a + \eta$ and do decrease monotonically.

Consider η as defined above. First, notice that given any $j \geq 0$, if

$$\Delta\beta_{t'+j} < 0 \quad \text{and} \quad (32)$$

$$\beta_{t'+j} > a + \eta \quad (33)$$

then

$$\begin{aligned} \Delta\beta_{t'+j+1} &= f_{t'+j+1} \left(a + \frac{\Delta\beta_{t'+j}}{1 - \delta\beta_{t'+j}} - \beta_{t'+j}, \beta_{t'+j} \right) < f_{t'+j+1} (a - \beta_{t'+j}, \beta_{t'+j}) \\ &< f_{t'+j+1} (-\eta, \beta_{t'+j}) \leq -\eta \mathcal{D}_{t'+j+1} \leq 0 \end{aligned} \quad (35)$$

where the first inequality follows from (32), the second inequality from (33) and the third from the mean value theorem, $\eta > 0$ and $\mathcal{D}_{t'+j+1} \geq 0$. Assume, towards a contradiction, that (33) holds for all $j \geq 0$. Since (32) holds for $j = 0$, it follows by induction that $\Delta\beta_{t'+j} \leq 0$ for all $j \geq 0$ and, therefore, that (35) would hold for all $j \geq 0$ hence

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

for all $j > 0$. Assumption A1 above would then imply $\beta_t \rightarrow -\infty$ showing that (33) can not hold for all j . Therefore there is a finite j such that $\beta_{t'+j}$ will go below $a + \eta$ and β is decreasing from t' until it goes below $a + \eta$.

For the case $\beta_t < a - \eta$, choosing $\epsilon = \eta$ one can use a symmetric argument to construct the proof.

7.4 Proof of Geometric Ergodicity

Defining $\eta_t \equiv (\varepsilon_t^c)^{-\gamma} \varepsilon_t^d$ and using (22) and (26) we can write the learning algorithm that gives the dynamics of β_t as

$$\begin{bmatrix} \beta_t \\ \Delta\beta_t \end{bmatrix} = F \begin{bmatrix} \beta_{t-1} \\ \Delta\beta_{t-1}, \eta_{t-1} \end{bmatrix}$$

where the first element of F , denoted F_1 , is given by the right side of (26) and $F_2(\beta, \Delta\beta, \eta_{t-1}) \equiv F_1(\beta, \Delta\beta, \eta_{t-1}) - \beta$. Therefore,

$$F'_t \equiv \frac{\partial F(\cdot, \eta_{t-1})}{\partial \begin{bmatrix} \beta_t & \Delta\beta_t \end{bmatrix}} = w'_t \cdot \begin{bmatrix} A_t, & 1 - \frac{1}{\alpha} + B_t \\ A_t, & -\frac{1}{\alpha} + \frac{1}{\alpha} B_t \end{bmatrix}$$

for $A_t = \frac{1}{\alpha} \frac{a\delta\eta_{t-1}}{1-\delta\beta_{t-1}}$, $B_t = \frac{1}{\alpha} \frac{a\delta \Delta\beta_{t-1}\eta_{t-1}}{(1-\delta\beta_{t-1})^2}$, with w'_t denoting the derivative of

w at period t . The eigenvalues of the matrix in brackets are

$$\lambda_t^+, \lambda_t^- = \frac{A_t + 1 - \frac{1}{\alpha} + B_t \pm \sqrt{(A_t + 1 - \frac{1}{\alpha} + B_t)^2 - 4A_t}}{2}$$

Since $A_t, B_t \rightarrow 0$ for large α we have that λ_t^+ is the larger eigenvalue in modulus and that the radicand is positive. We wish to find a uniform bound for λ_t^+ , because given that $|w'_t| < 1$ this will be a uniform bound for the largest eigenvalue of F'_t . Such a bound will play the role of $\rho_\theta(\varepsilon_t)$ in the definition of the " L^2 unit circle condition" on p. 942 in Duffie and Singleton (1993) (henceforth DS).

Consider the function $f_a(x) = x + a + \sqrt{(x+a)^2 - 4a}$ for some constant $a > 0$ and x large enough for the radicand to be positive. For $\varepsilon > 0$ the mean value theorem implies

$$f_a(x + \varepsilon) \leq \left(1 + \frac{x}{\sqrt{(x+a)^2 - 4a}} \right) \varepsilon + f_a(x)$$

Evaluating this expression at $a = A_t$, $\varepsilon = B_t$ and $x = 1 - 1/\alpha$ we have

$$\lambda_t^+ \leq B_t + \frac{f_{A_t}(1 - \frac{1}{\alpha})}{2} < B_t + 1 - \frac{1}{\alpha} \quad \text{for } \Delta\beta_{t-1} \geq 0 \quad (36)$$

where we used

$$f_{A_t}(1 - \frac{1}{\alpha}) < A_t + 1 - \frac{1}{\alpha} + \sqrt{\left(A_t + 1 - \frac{1}{\alpha}\right)^2 - 4A_t(1 - \frac{1}{\alpha})} = 2 \left(1 - \frac{1}{\alpha}\right)$$

Since $f_{A_t}(\cdot)$ is monotonic, using the expression for B_t we have

$$\lambda_t^+ \leq \frac{1}{2} f_{A_t}(1 - \frac{1}{\alpha} + B_t) \leq \frac{f_{A_t}(1 - \frac{1}{\alpha})}{2} < 1 - \frac{1}{\alpha} \quad \text{for } \Delta\beta_{t-1} < 0 \quad (37)$$

From (26) we have

$$\Delta\beta_t \leq \frac{1}{\alpha} \left(\eta_{t-1} a^{1-\gamma} \left[1 + \frac{\Delta\beta_{t-1}}{1 - \delta\beta_{t-1}} \right] - \beta_{t-1} \right) \quad (38)$$

So, if $\Delta\beta_{t-1} \geq 0$, using $\beta_{t-1} > 0$

$$\Delta\beta_t \leq \frac{1}{\alpha} \eta_{t-1} a^{1-\gamma} \left[1 + \frac{|\Delta\beta_{t-1}|}{1 - \delta\beta_{t-1}} \right]$$

Therefore, adding the right side of this inequality to (37), and using the inequality for (36) we have that for all $\Delta\beta_{t-1}$

$$\begin{aligned} \lambda_t^+ &\leq \frac{1}{\alpha} \frac{a^{1-\gamma} \delta}{(1 - \delta\beta_{t-1})^2} \left(\frac{a^{1-\gamma}}{\alpha} \eta_{t-2} \left[1 + \frac{|\Delta\beta_{t-2}|}{1 - \delta\beta_{t-2}} \right] \right) \eta_{t-1} + 1 - \frac{1}{\alpha} \\ &\leq \frac{1}{\alpha^2} \tilde{K} \eta_{t-2} \eta_{t-1} + 1 - \frac{1}{\alpha} \end{aligned}$$

for a constant $0 < \tilde{K} < \infty$, where we used $|\Delta\beta_{t-2}|, \beta_{t-1}, \beta_{t-2} < \beta^U$.

Since $w' \leq 1$, it is clear from the mean value theorem that $\frac{\tilde{K}}{\alpha^2} \eta_{t-2} \eta_{t-1} + 1 - \frac{1}{\alpha}$ plays the role of $\rho_\theta(\varepsilon_t)$ in the Definition of " L^2 unit circle condition" of DS, where our α plays the role of θ , and $\eta_{t-1} \eta_{t-2}$ the role of ε_t in DS. Therefore, we need to check that $E \left(\frac{\tilde{K}}{\alpha^2} \eta_t \eta_{t-1} + 1 - \frac{1}{\alpha} \right)^2 < 1$ for α large enough. A routine calculation shows that

$$E \left(\frac{\tilde{K}}{\alpha^2} \eta_t \eta_{t-1} + 1 - \frac{1}{\alpha} \right)^2 = 1 - \frac{1}{\alpha} - \frac{1}{\alpha} \left[1 - \frac{1}{\alpha} - 2 \left(1 - \frac{1}{\alpha} \right) \frac{\tilde{K}}{\alpha} E(\eta_t \eta_{t-1}) - \frac{\tilde{K}^2}{\alpha^3} E(\eta_t^2 \eta_{t-1}^2) \right]$$

For α large the term in brackets is positive. Therefore,

$$E \left(\frac{\tilde{K}}{\alpha^2} \eta_t \eta_{t-1} + 1 - \frac{1}{\alpha} \right)^2 < 1 - \frac{1}{\alpha} < 1$$

This proves that for large α the variable β_t satisfies DS' L^2 unit-circle condition, hence it satisfies DS' AUC condition, and Lemma 3 in DS guarantees that β_t is geometrically ergodic.

Now, adding $a^{1-\gamma} \eta_{t-1}$ on both sides of (38) and taking expectations at the ergodic distribution we have

$$E(\beta_{t-1} - \eta_{t-1} a^{1-\gamma}) \leq E \left(\frac{\Delta\beta_{t-1}}{1 - \delta\beta_{t-1}} \eta_{t-1} a^{1-\gamma} \right) \quad (39)$$

Our previous argument shows that the right side is arbitrarily small for α large, therefore $E\beta_{t-1} \leq E\eta_{t-1}a^{1-\gamma}$. A similar argument shows that $\text{var}\beta_t$ goes to zero as $\alpha \rightarrow \infty$. Therefore for α large $\beta_t \leq \beta^L$ with arbitrarily large probability so that (38) holds as equality with arbitrarily large probability. Taking expectations on both sides for the realizations where this holds as equality, we have that $E\beta_t \rightarrow E\eta_{t-1}a^{1-\gamma} = \beta^{RE}$ as $\alpha \rightarrow \infty$ which completes the proof. ■

7.5 Differentiable projection facility

The function w used in the differentiable projection facility is

$$w(x) = U \begin{cases} x & \text{if } x \leq \beta^L \\ \beta^L + \frac{x - \beta^L}{x + \beta^U - 2\beta^L}(\beta^U - \beta^L) & \text{if } \beta^L < x \leq \beta^U \end{cases} \quad (40)$$

$$\beta^L + \frac{1}{2}(\beta^U - \beta^L)$$

Clearly w is continuous, the only point where continuity is questionable is at $x = \beta^L$, but it is easy to check that

$$\begin{aligned} \lim_{x \nearrow \beta^L} w(x) &= \lim_{x \searrow \beta^L} w(x) = \beta^L \\ \lim_{x \nearrow \beta^L} w'(x) &= \lim_{x \searrow \beta^L} w'(x) = 1 \\ \lim_{x \rightarrow \infty} w(x) &= \beta^U \end{aligned}$$

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. Therefore, this dampening is applied in few observations.

7.6 Details of the MSM procedure

7.6.1 Estimation and testing

Here we provide an estimator for the covariance matrix of the statistics $\widehat{\Sigma}_{S,N}$, we give some more details about consistency of the estimator defined in (28), and derive the needed asymptotic distribution results.

Let N be the sample size, $(\mathbf{y}_1, \dots, \mathbf{y}_N)$ the observed data sample, with \mathbf{y}_t containing m variables. The standard version of the method of simulated

moments (MSM) is to find parameter values that make the moments of the structural model close to sample moments $\widehat{M}_N \equiv \frac{1}{N} \sum_{t=1}^N h(\mathbf{y}_t)$ for a given *moment function* $h : R^m \rightarrow R^q$. However, many of the statistics that we wish to match in Table 1 are not of this form, but they are functions of moments. Formally, the statistics in table 1 can be written as $\widehat{S}_N \equiv \mathcal{S}(\widehat{M}_N)$ for a *statistic function* $\mathcal{S} : R^q \rightarrow R^s$ mapping sample moments \widehat{M}_N into the considered statistics. Explicit expressions for $h(\cdot)$ and $\mathcal{S}(\cdot)$ in our particular application are stated in section 7.6.2. In the text we talked about "moments" as describing all statistics to be matched in (27). In this appendix we properly use the term "statistic" as possibly different from "moment".

We propose to base our MSM estimates and tests on matching the statistics \widehat{S}_N . Since this deviates from standard MSM we need to adapt standard proofs and asymptotic distribution results. The proofs follow standard steps so we provide an outline of the argument and the derivations. The statistics to be matched \widehat{S}_N can be all or a subset of the statistics in (27), the reason we sometimes consider a subset will be discussed in detail after result "Asymptotic Distribution of MSM" below. Since we have endogenous state variables in the model (namely past beliefs β_{t-1}, β_{t-2}) the asymptotic theory result needed is from Duffie and Singleton (1993) (DS).

Let $y_t(\theta)$ be the series generated by the structural model at hand for parameter values θ and some realization of the underlying shocks. All the results below are derived under the null hypothesis that the model is true, more specifically, that the observed data is generated by the structural model at hand for a true parameter value θ_0 . Let $M(\theta) \equiv E[h(y_t(\theta))]$ be the true moments for parameter values θ at the stationary distribution of $y_t(\theta)$, hence $M_0 \equiv M(\theta_0)$ are the true moments, and let $\widehat{S}(\theta) \equiv \mathcal{S}(M(\theta))$ be the true statistics when the model parameter is θ . Denote by M_0^j the j -th autocovariance of the moment function at the true parameter, that is

$$M_0^j \equiv E[h(y_t(\theta_0)) - M_0][h(y_{t-j}(\theta_0)) - M_0]'$$

Define

$$S_w \equiv \sum_{j=-\infty}^{\infty} M_0^j \tag{41}$$

We use the following estimate of the variance for the sample statistics \widehat{S}_N

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

The asymptotic properties of this estimate are given in the following result:

Variance of $\widehat{\mathcal{S}}_N$ Suppose that

- a) $S_w < \infty$, and we have consistent estimators of this matrix $\widehat{S}_{w,N} \rightarrow S_w$ a.s. as $N \rightarrow \infty$
- b) \mathcal{S} is continuously differentiable at M_0 .
- c) the observed process $\{\mathbf{y}_t\}$ is stationary and ergodic

Then we have that

$$\widehat{\Sigma}_{\mathcal{S},N} \rightarrow \Sigma_{\mathcal{S}} \equiv \frac{\partial \mathcal{S}(M_0)}{\partial M'} S_w \frac{\partial \mathcal{S}(M_0)'}{\partial M} \quad (42)$$

and that $\Sigma_{\mathcal{S}}$ is the asymptotic covariance matrix of $\widehat{\mathcal{S}}_N$:

$$E \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M_0) \right] \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M_0) \right]' \rightarrow \Sigma_{\mathcal{S}} \quad (43)$$

both limits occurring a.s. as $N \rightarrow \infty$.

Therefore, $\widehat{\Sigma}_{\mathcal{S},N}$ is a consistent estimator of the asymptotic variance of the sample statistics.

Proof. Assumptions a), c) imply

$$\widehat{M}_N \rightarrow M_0 \text{ a.s. as } N \rightarrow \infty$$

and assumption b) gives (42).

Assumptions a), c) imply

$$E [M_N - M_0] [M_N - M_0]' \rightarrow S_w \text{ a.s. as } N \rightarrow \infty$$

The mean value theorem implies that

$$\mathcal{S}(M_N) - \mathcal{S}(M_0) = \frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} [M_N - M_0] \quad (44)$$

for a $\widetilde{M}_N \rightarrow M_0$ a.s. as $N \rightarrow \infty$. Taking expectations of $[\mathcal{S}(M_N) - \mathcal{S}(M_0)] [\mathcal{S}(M_N) - \mathcal{S}(M_0)]'$ we have (43). ■

Conditions a), c) are standard minimal assumptions used in time series asymptotic results, condition b) is clearly satisfied in our application, see the expression for \mathcal{S} in section 7.6.2. We choose consistent estimates $\widehat{S}_{w,N}$ applying the Newey West estimator using only the data. Hence the estimator $\widehat{\Sigma}_{\mathcal{S},N}$ can be found purely from data, without using the model or its parameter estimates. We now turn to

Consistency Let $\widehat{\theta}_N$ be the estimator defined in (28), where the maximization is over a set $\Theta \subset R^n$. Assume

- a) Θ is compact, the process $\{y_t(\theta)\}$ is well defined for all $\theta \in \Theta$, $\widetilde{\mathcal{S}}$ is continuous in Θ , and $\theta_0 \in \Theta$.
- b) $\{y_t(\theta)\}$ is geometrically ergodic for all $\theta \in \Theta$
- c) $\Sigma_{\mathcal{S}}$ is invertible
- d)

$$\left[\widetilde{\mathcal{S}}(\theta_0) - \widetilde{\mathcal{S}}(\theta) \right]' \Sigma_{\mathcal{S}}^{-1} \left[\widetilde{\mathcal{S}}(\theta_0) - \widetilde{\mathcal{S}}(\theta) \right] > 0 \text{ for all } \theta \neq \theta_0$$

Then

$$\widehat{\theta}_N \rightarrow \theta_0 \text{ a.s. as } N \rightarrow \infty.$$

The proof is easily obtained by adapting the consistency result from DS.

Condition a) is standard in GMM applications, the set Θ should be large enough to insure that it contains admissible values of the true parameter values. DS emphasize that a strong form of ergodicity is needed as in condition b), we showed in appendix 7.4 that this holds for α large enough, therefore b) is guaranteed if Θ is restricted to large α . Conditions c) and d) are standard identification requirements that the statistics selected are sufficient to identify the true parameter values. A necessary condition for d) is that the number of parameters is less than the number of statistics s .

Let

$$\mathcal{B}_0 \equiv \frac{\partial M'(\theta_0)}{\partial \theta} \frac{\partial \mathcal{S}'(M_0)}{\partial M}$$

Asymptotic Distribution In addition to all the assumptions in the above results, assume that $\mathcal{B}_0 \Sigma_{\mathcal{S}}^{-1} \mathcal{B}_0'$ is invertible. Then

$$\sqrt{N} \left[\widehat{\theta}_N - \theta_0 \right] \rightarrow \mathcal{N}(0, (\mathcal{B}_0 \Sigma_{\mathcal{S}}^{-1} \mathcal{B}_0)^{-1}) \quad (45)$$

$$\sqrt{N} \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M(\widehat{\theta})) \right] \rightarrow \mathcal{N}(0, \Sigma_{\mathcal{S}} - \mathcal{B}_0' (\mathcal{B}_0 \Sigma_{\mathcal{S}}^{-1} \mathcal{B}_0)^{-1} \mathcal{B}_0) \quad (46)$$

$$\widehat{W}_N \rightarrow \chi_{s-n}^2 \quad (47)$$

in distribution as $N \rightarrow \infty$, where $\widehat{W}_N \equiv N \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\widehat{\theta}) \right]' \widehat{\Sigma}_{\mathcal{S},N}^{-1} \left[\widehat{\mathcal{S}}_N - \widetilde{\mathcal{S}}(\widehat{\theta}) \right]$

Proof. The central limit theorem and (44) imply

$$\sqrt{N} [\mathcal{S}(M_N) - \mathcal{S}(M_0)] = \frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} \sqrt{N} [M_N - M_0] \rightarrow \mathcal{N}(0, \Sigma_{\mathcal{S}}) \quad (48)$$

in distribution. Letting

$$\mathcal{B}(\theta, M) \equiv \frac{\partial M'(\theta)}{\partial \theta} \frac{\partial \mathcal{S}'(M)}{\partial M}$$

The asymptotic distribution of the parameters is derived as

$$\begin{aligned} \mathcal{S}(M(\widehat{\theta}_N)) - \mathcal{S}(M(\theta_0)) &= \mathcal{B}(\widetilde{\theta}, \widetilde{M}) [\widehat{\theta}_N - \theta_0] \\ \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}, N}^{-1} [\mathcal{S}(M(\widehat{\theta})) - \mathcal{S}(M(\theta_0))] &= \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}, N}^{-1} \mathcal{B}(\widetilde{\theta}, \widetilde{M})' [\widehat{\theta} - \theta_0] \\ \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} [\mathcal{S}(M_N) - \mathcal{S}(M(\theta_0))] &= \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} \mathcal{B}(\widetilde{\theta}, \widetilde{M})' [\widehat{\theta} - \theta_0] \end{aligned} \quad (49)$$

where the last equality follows because $\mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} [\mathcal{S}(M_N) - \mathcal{S}(\widehat{\theta}_N)] = 0$ at the maximum of (28). This implies (45).

To obtain (46) we use mean value theorem and (49) to conclude

$$\begin{aligned} \mathcal{S}(\widehat{M}_N) - \mathcal{S}(M(\widehat{\theta}_N)) &= \mathcal{S}(\widehat{M}_N) - \mathcal{S}(M(\theta_0)) + \mathcal{B}(\widetilde{\theta}, \widetilde{M}) [\theta_0 - \widehat{\theta}] = \\ &= \left(I - [\mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \Sigma_{\mathcal{S}}^{-1} \mathcal{B}(\widetilde{\theta}, \widetilde{M})']^{-1} \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}, N}^{-1} \right) [\mathcal{S}(\widehat{M}_N) - \mathcal{S}(M(\theta_0))] \end{aligned}$$

This gives (46).

(47) follows from (46) and that $(\Sigma_{\mathcal{S}} - \mathcal{B}'_0 [\mathcal{B}_0 \Sigma_{\mathcal{S}}^{-1} \mathcal{B}'_0]^{-1} \mathcal{B}_0)$ is an idempotent matrix. ■

We mentioned in the main text and earlier in this appendix that we drop some statistics from $\widehat{\mathcal{S}}_N$ when $\widehat{\Sigma}_{\mathcal{S}, N}$ is nearly singular. The reason is that, as stated above, we need invertibility of $\Sigma_{\mathcal{S}}$ both for consistency and asymptotic distribution. In practice, a nearly singular $\widehat{\Sigma}_{\mathcal{S}, N}$ creates many problems. First, the results for the test W change very much with small changes in the model or testing procedure and the maximization algorithm is nearly unstable, making it difficult to find a maximum numerically. This happens because in this case the formula for \widehat{W}_N nearly divides zero by zero, hence the objective function is nearly undefined, and the asymptotic distribution is not necessarily a good approximation to the true distribution of the test statistic. But this is not a bad situation for an econometrician, it just means that one of the statistics is redundant, so it makes sense to simply drop one statistic from the estimation.

To decide which statistic to drop we compute the variability of each statistic that can not be explained by a linear combination of the remaining

statistics. This is analogous to the R^2 coefficient of running a regression of each statistic on all the other statistics when the regression coefficients are computed from $\widehat{\Sigma}_{\mathcal{S},N}$. We drop the statistics for which this R^2 is less than 1%, as it turns out this only occurs for \widehat{c}_2^5 with an R^2 equal to 0.006. After we drop \widehat{c}_2^5 the estimation results become sufficiently stable.³⁶ As we explain in our discussion around Table 2, the model is in any case able to match \widehat{c}_2^5 even when we drop it from the statistics used in the estimation.

There are various ways to compute the moments of the model $\widetilde{\mathcal{S}}(\theta)$ for a given $\theta \in R^n$. We use the following Monte-Carlo procedure. Let ω^i denote a realization of shocks drawn randomly from the known distribution that the underlying shocks are assumed to have and $(y_1(\theta, \omega^i), \dots, y_N(\theta, \omega^i))$ the random variables corresponding to a history of length N generated by the model for shock realization ω^i and parameter values θ . Furthermore, let

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

denote the model moment for realization ω^i . We set the model statistics $\widetilde{\mathcal{S}}(\theta)$ equal to

$$\frac{1}{K} \sum_{i=1}^K \mathcal{S}(M_N(\theta, \omega^i))$$

for large K . In other words, $\widetilde{\mathcal{S}}(\theta)$ is an average across a large number of simulations of length N of the statistics $\mathcal{S}(M_N(\theta, \omega^i))$ implied by each simulation. We use K of the order of 1000, therefore the model moments are computed with KN observations. These are the averages reported as model moments in tables 2 to 4 of the main text.

Many papers on MSM emphasize the dependence of the estimates on the ratio of number of observations in simulations to N . Since this is 1000 in our application this adds a negligible factor to the asymptotic variance-covariance matrices computed and we entirely ignore it in our results.

7.6.2 The statistic and moment functions

This section gives explicit expressions for the statistic function $\mathcal{S}(\cdot)$ and the moment functions $h(\cdot)$ that map our estimates into the framework just

³⁶ An alternative to dropping one moment would be to match the linear combinations corresponding to the principal components of \widehat{S}_w . While this has some advantages from the point of view of asymptotic theory, we find it less attractive from the economic point of view, since the economic meaning of these principal components would be unclear, while each individual moment does have clear economic interpretation.

discussed in this appendix.

The underlying sample moments needed to construct the statistics of interest are

$$M_N \equiv \frac{1}{N} \sum_{t=1}^N h(\mathbf{y}_t)$$

where $h(\cdot) : R^{42} \rightarrow R^{11}$ and \mathbf{y}_t are defined as

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

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

The ten statistics we consider can be expressed as function of the moments as follows:

$$\mathcal{S}(M) \equiv \begin{bmatrix} E(r_t^s) \\ E(PD_t) \\ \sigma_{r_t^s} \\ \sigma_{PD_t} \\ \rho_{PD_t, -1} \\ c_2^5 \\ R_5^2 \\ E(r_t^b) \\ E(\Delta D_t/D_{t-1}) \\ \sigma_{D_t/D_{t-1}} \end{bmatrix} = \begin{bmatrix} M_1 \\ M_2 \\ \sqrt{M_3 - (M_1)^2} \\ \sqrt{M_4 - (M_2)^2} \\ \frac{M_5 - (M_2)^2}{M_4 - (M_2)^2} \\ c_2^5(M) \\ R_5^2(M) \\ M_9 \\ M_{10} \\ \sqrt{M_{11} - (M_{10})^2} \end{bmatrix}$$

where M_i denotes the i -th element of M and the functions $c_2^5(M)$ and $R_5^2(M)$ define the OLS and R^2 coefficients of the excess returns regressions, more precisely

$$c^5(M) \equiv \begin{bmatrix} 1 & M_2 \\ M_2 & M_4 \end{bmatrix}^{-1} \begin{bmatrix} M_6 \\ M_8 \end{bmatrix}$$

$$R_5^2(M) \equiv 1 - \frac{M_7 - [M_6, M_8] c^5(M)}{M_7 - (M_6)^2}$$

7.6.3 Derivatives of the statistic function

This appendix gives explicit expressions for $\partial\mathcal{S}/\partial M'$ using the statistic function stated in appendix 7.6.2. Straightforward but tedious algebra shows

$$\frac{\partial\mathcal{S}_i}{\partial M_j} = 1 \quad \text{for } (i, j) = (1, 1), (2, 2), (8, 9), (9, 10)$$

$$\frac{\partial\mathcal{S}_i}{\partial M_j} = \frac{1}{2\mathcal{S}_i(M)} \quad \text{for } (i, j) = (3, 3), (4, 4), (10, 11)$$

$$\frac{\partial\mathcal{S}_i}{\partial M_j} = \frac{-M_j}{\mathcal{S}_i(M)} \quad \text{for } (i, j) = (3, 1), (4, 2), (10, 10)$$

$$\frac{\partial\mathcal{S}_5}{\partial M_2} = \frac{2M_2(M_5 - M_4)}{(M_4 - M_2^2)^2}, \quad \frac{\partial\mathcal{S}_5}{\partial M_5} = \frac{1}{M_4 - M_2^2}, \quad \frac{\partial\mathcal{S}_5}{\partial M_4} = -\frac{M_5 - M_2^2}{(M_4 - M_2^2)^2}$$

$$\frac{\partial\mathcal{S}_6}{\partial M_j} = \frac{\partial c_2^5(M)}{\partial M_j} \quad \text{for } i = 2, 4, 6, 8$$

$$\frac{\partial\mathcal{S}_7}{\partial M_j} = \frac{[M_6, M_8] \frac{\partial c^5(M)}{\partial M_j}}{M_7 - M_6^2} \quad \text{for } j = 2, 4$$

$$\frac{\partial\mathcal{S}_7}{\partial M_6} = \frac{\left[c_1^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_6} \right] (M_7 - M_6^2) - 2M_6 [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2}$$

$$\frac{\partial\mathcal{S}_7}{\partial M_7} = \frac{M_6^2 - [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2}$$

$$\frac{\partial\mathcal{S}_7}{\partial M_8} = \frac{c_2^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_8}}{M_7 - M_6^2}$$

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.