

Federal Reserve Bank of Minneapolis
Research Department

Stock Market Volatility and Learning*

Klaus Adam, Albert Marcet, and Juan Pablo Nicolini

Working Paper 720

February 2015

ABSTRACT

Consumption-based asset pricing models with time-separable preferences can generate realistic amounts of stock price volatility if one allows for small deviations from rational expectations. We consider rational investors who entertain subjective prior beliefs about price behavior that are not equal but close to rational expectations. 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. When estimating the model on U.S. stock price data using the method of simulated moments, we find that it can quantitatively account for the observed volatility of returns, the volatility and persistence of the price-dividend ratio, and the predictability of long-horizon returns. For reasonable degrees of risk aversion, the model generates up to one-half of the equity premium observed in the data. It 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.

Keywords: Asset pricing; Learning; Subjective beliefs; Internal rationality
JEL classification: G12, E44

*Adam: University of Mannheim and CEPR. Marcet: Institut d'Anàlisi Econòmica CSIC, ICREA, MOVE, Barcelona GSE, UAB, and CEPR. Nicolini: Federal Reserve Bank of Minneapolis and Universidad Di Tella. Thanks go to Bruno Biais, Peter Bossaerts, Tim Cogley, Davide Debortoli, Luca Dedola, George Evans, Katharina Greulich, Seppo Honkapohja, Bruce McGough, Bruce Preston, Tom Sargent, Ken Singleton, Hans-Joachim Voth, Ivan 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 those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

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.

—Shiller (2005, p. 56)

1 Introduction

The purpose of this paper is to show that a 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. We find it a striking observation that the quantitative asset pricing implications of the standard model are not robust to small departures from rational expectations and that this nonrobustness is empirically so encouraging.

We study a simple variant of the Lucas (1978) model with standard time-separable consumption 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 risk premium.

We stick to Lucas’ framework but relax the standard assumption that agents have perfect knowledge about the pricing function that maps each history of fundamental shocks into a market outcome for the stock price.¹ We assume instead that investors hold subjective beliefs about all payoff-relevant random variables that are beyond their control; this includes beliefs about model endogenous variables, such as 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” because they know all internal aspects of their individual decision problem and maximize utility given this knowledge. Furthermore, their system of beliefs is “internally consistent” in the sense

¹Lack of knowledge of the pricing function may arise from a lack of common knowledge of investors’ preferences, price beliefs and dividend beliefs, as explained in detail in the work of Adam and Marcet (2011).

that it specifies for all periods the joint distribution of all payoff-relevant variables (i.e., dividends, income, and stock prices), but these probabilities differ from those implied by the model in equilibrium. We then consider systems of beliefs implying only a small deviation from RE, as we explain further below.

We show that given the subjective beliefs we specify, subjective utility maximization dictates that agents update subjective expectations about stock price behavior using realized market outcomes. Consequently, 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. More specifically, the empirical success of the model emerges whenever agents learn about the growth rate of stock prices (i.e., about the capital gains from their investments) using past observations of capital gains.

We first demonstrate the ability of the model to produce datalike 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 encompassing most learning algorithms that have been used in the learning literature. 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.

We then consider a specific system of beliefs that allows for subjective prior uncertainty about the average growth rate of stock market prices, given the values for all exogenous variables. As we show, internal rationality (i.e., standard utility maximization given these beliefs) then dictates that agents' price growth expectations react to the realized growth rate of market prices. In particular, the subjective prior prescribes that agents should update conditional expectations of one-step-ahead risk-adjusted price growth using a

constant gain model of adaptive learning. This constant gain model belongs to the general class of learning rules that we studied analytically before and therefore displays momentum and mean reversion.

We document in several ways that the resulting beliefs represent only a small deviation from RE beliefs. First, we show that for the special case in which the prior uncertainty about price growth converges to zero, the learning rule delivers RE beliefs, and prices under learning converge to RE prices (see section 4.3). In our empirical section, we then find that the asset pricing facts can be explained with a small amount of prior uncertainty. Second, using an econometric test that exhausts the empirical implications of agents' subjective model of price behavior, we show that agents' price beliefs would not be rejected by the data (section 6.2). Third, using the same test but applying it to artificial data generated by the estimated model, we show that it is difficult to detect that price beliefs differ from the actual behavior of prices in equilibrium (section 6.2).

To quantitatively evaluate the learning model, we first consider how well it 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), adapting the results of Duffie and Singleton (1993). We find that the model can individually match all the asset pricing 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. Using t -statistics derived from asymptotic theory, we cannot reject that any of the individual model moments differs from the moments in the data in one of our estimated models (see Table 2 in section 5.2). The model also delivers an equity premium of up to one-half of the value observed in the data. All this is achieved even though we use time-separable CRRA preferences and a degree of relative risk aversion equal to 5 only.

We also perform a formal econometric test for the overall goodness of fit of our consumption-based asset pricing model. This is a considerably more stringent test than implied by individually matching asset pricing moments as in calibration exercises (e.g., Campbell and Cochrane (1999)) but is a natural one to explore given our MSM strategy. As it turns out, the overall goodness of fit test is much more stringent and rejects the model if one includes both the risk-free rate and the mean stock returns, but if we leave out the risk premium by excluding the risk-free rate from the estimation, the p -value of the model amounts to a respectable 7.1%.

The general conclusion we obtain is that for moderate risk aversion, the model can quantitatively account for all asset pricing facts, except for the equity premium. For a sufficiently high risk aversion as in Campbell and Cochrane (1999), the model can also replicate the equity premium, whereas under RE it explains only one-quarter of the observed value. This is a remarkable improvement relative to the performance of the model under RE and suggests that allowing for small departures from RE is 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 the MSM estimation and testing strategy that we use 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 Tables 2 and 3 in section 5.2. Section 6 investigates the robustness of our finding to a number of alternative modeling assumptions, as well as the degree to which agents could detect if they are making systematic forecast errors. A conclusion briefly summarizes the main findings.

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. Stock price behavior under Bayesian learning has been studied in the work of Timmermann (1993, 1996), Brennan and Xia (2001), Cecchetti, Lam, and Mark (2000), and Cogley and Sargent (2008), among others. Some papers in this vein study agents that have asymmetric information or asymmetric beliefs; examples include the work of Biais, Bossaerts, and Spatt (2010) and Dumas, Kurshev, and Uppal (2009). 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

²See the work of Campbell (2003) for an overview.

explained in the work of 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 these models stock prices typically represent redundant information given agents' assumed knowledge, 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. Branch and Evans (2011) study a model with learning about returns and return risk. Lansing (2010) shows how near-rational bubbles can arise under learning dynamics when agents forecast a composite variable involving future price and dividends. Boswijk, Hommes, and Manzan (2007) estimate a model with fundamentalist and chartist traders whose relative shares evolve according to an evolutionary performance criterion. 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 (2008) assume that agents know the mean stock price and find that learning does not then significantly alter

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

⁴Timmerman reports that this form of learning delivers even lower volatility than in 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*.

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.

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.

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 the work of Shiller (2005) for a nontechnical summary. We borrow from this literature an interest in deviating from RE, but we are keen on making only a minimal deviation from the standard approach: we assume that agents behave optimally given an internally consistent system of subjective beliefs that is close (but not equal) 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, and we reproduce them here as a point of reference using a single and updated database.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 price-dividend (PD) ratio in the United States, where the PD ratio is defined as the ratio of stock prices over quarterly dividend payments. The PD ratio displays very large fluctuations around its sample mean (the bold horizontal line in the graph): in the year 1932, the quarterly PD ratio takes on values below 30, whereas in the year 2000, values are close to 350. The standard deviation of

⁵Details on the data sources are provided in Appendix A.1.




Figure 1: Quarterly U.S. price-dividend ratio 1927:2-2012:2

the PD ratio (σ_{PD}) is approximately 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 is 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.

Related to the excessive volatility of prices is the observation that the volatility of stock returns (σ_{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.

Although 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 five-year horizon as

Fact 4 in Table 1.⁶

Fact 1	Volatility of PD ratio	E_{PD} σ_{PD}	123.91 62.43
Fact 2	Persistence of PD ratio	$\rho_{PD,-1}$	0.97
Fact 3	Excessive return volatility	σ_{r^s}	11.44
Fact 4	Excess return predictability	c_5^2 R_5^2	-0.0041 0.2102
Fact 5	Equity premium Quarterly real stock returns Quarterly real bond returns	E_{r^s} E_{r^b}	2.25 0.15
Dividend Behavior	Mean growth Std. dev. of growth	$E_{\frac{\Delta D}{D}}$ $\sigma_{\frac{\Delta D}{D}}$	0.41 2.88
PD ratio is price over quarterly dividend. Growth rates and returns are expressed in real terms and at quarterly rates of increase.			

Table 1: U.S. asset pricing facts 1927:2-2012:2

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**

⁶We focus on the five-year horizon for simplicity but obtain very similar results for other horizons. Our focus on a single horizon is justified because chapter 20 in the work of 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 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 five-year horizon to check the *quantitative* model implications.

5 in Table 1, which shows that the average quarterly real return on bonds $E_{r,b}$ is much lower than the corresponding quarterly 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.

4 The Model

We describe below a Lucas (1978) asset pricing model with agents who hold subjective prior beliefs about stock price behavior. We show that the presence of subjective uncertainty implies that utility-maximizing agents update their beliefs about stock price behavior using observed stock price realizations.⁷ 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. Section 4.3 also shows under what conditions this system of beliefs gives rise to small deviations from RE.

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 $E(\varepsilon_t^d) = 1$, $E_{\frac{\Delta D}{D}} \equiv E\left(\frac{D_t - D_{t-1}}{D_{t-1}}\right) = a - 1$ and $\sigma_{\frac{\Delta D}{D}}^2 \equiv var\left(\frac{D_t - D_{t-1}}{D_{t-1}}\right) = a^2 \left(e^{s_d^2} - 1\right)$. To capture the fact that the empirically observed consumption process is considerably less volatile than the dividend process and to replicate the correlation between dividend and consumption growth, we assume that each agent receives in addition an endowment Y_t of perishable consumption

⁷This draws on the results in the work of Adam and Marcet (2011).

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 aggregate consumption supply process⁸

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

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 the work of 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 $\rho_{c,d} = 0.2$.

Objective Function and Probability Space: Agent $i \in [0, 1]$ has a standard time-separable expected utility function⁹

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

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 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*, and C_t denotes the aggregate *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}) .

The probability measure \mathcal{P} specifies the joint distribution of $\{P_t, Y_t, D_t\}_{t=0}^{\infty}$ at all dates and is fixed at the outset. Although the measure is fixed, in-

⁸The process for Y_t is then implied by feasibility.

⁹We assume standard preferences so as to highlight the effect of learning on asset price volatility.

vestors' beliefs about unknown parameters describing the stochastic processes of these variables, as well as investors' conditional expectations of future values of these variables, will change over time in a way that is derived from \mathcal{P} and that will depend on realized data. This specification thus encompasses settings in which agents are learning about the stochastic processes describing $P_t, Y_t,$ and D_t . Moreover, unlike in the anticipated utility framework proposed in the work of Kreps (1998), agents are fully aware of the fact that beliefs in the future will get revised. Although the probability measure \mathcal{P} is not equal to the distribution of $\{P_t, Y_t, D_t\}_{t=0}^{\infty}$ implied by the model in equilibrium, it will be chosen in a way such that it is close to it in a sense that we make precise in sections 4.3 and 6.3.

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 also include price histories in the realization ω^t . Standard practice is to assume instead that agents know the exact mapping $P_t(\cdot)$ 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 us - 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 models such as those described in the second paragraph of section 2, and models incorporating robustness concerns. This 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 pricing function $P_t(\cdot)$. Although a convenient modeling device, assuming exact knowledge of this function is at the same time very restrictive: it assumes that agents have a very detailed knowledge of how prices are formed. This makes it of interest to study the implication of (slightly) relaxing the assumption that agents know the function $P_t(\cdot)$. Adam and Marcet (2011) show that rational behavior is indeed perfectly compatible with agents not knowing the *exact* form of the equilibrium pricing function $P_t(\cdot)$.¹⁰

¹⁰Specifically, they show that with incomplete markets (i.e., in the absence of state-contingent forward markets for stocks), agents cannot simply learn the equilibrium map-

Choice Set and Constraints: Agents make contingent plans for consumption C_t^i , bondholdings B_t^i and stockholdings S_t^i ; that is, they choose the functions

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

for all $t \geq 0$. Agents' 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 ensure the existence of a maximum, the following bounds are assumed to hold:

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

We only assume the bounds $\underline{S}, \overline{S}, \underline{B}, \overline{B}$ are finite and they satisfy $\underline{S} < 1 < \overline{S}$, $\underline{B} < 0 < \overline{B}$.

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 specified through 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} .

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

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

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

ping $P_t(\cdot)$ 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 their own optimization conditions.

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.

Before doing so, we briefly review the well-known result that under RE the model is 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 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 ($\rho_{c,d} = 0.2$).¹¹ The model thus also misses Fact 5 in Table 1.

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 .¹² However, outside of strict rational expectations we may have

¹¹Under 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 .

¹² $C_t^i = C_t$ follows from market clearing and the fact that all agents are identical.

$E_t^{\mathcal{P}} [C_{t+1}^i] \neq E_t^{\mathcal{P}} [C_{t+1}]$ even if in equilibrium $C_t^i = C_t$ holds ex post.¹³ 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, in general $E_t^{\mathcal{P}} [C_{t+1}^i] \neq E_t^{\mathcal{P}} [C_{t+1}]$, so that one cannot routinely substitute individual by aggregate consumption on the right-hand side of the 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 agents with beliefs \mathcal{P} realize in period t that $C_{t+1}^i / C_t^i \approx C_{t+1} / C_t$. This follows from the flow budget constraint for period $t+1$ and the fact that $S_t^i = 1$, $B_t^i = 0$, and $C_t^i = C_t$ in equilibrium in period t . One can then rely on the approximations

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

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

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

Assumption 1 We assume that Y_t is sufficiently large and that $E_t^{\mathcal{P}} P_{t+1} / D_t < \overline{M}$ for some $\overline{M} < \infty$ so that, given finite asset bounds $\underline{S}, \overline{S}, \underline{B}, \overline{B}$, the approximations (11) and (12) hold with sufficient accuracy.

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

¹³This 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.

$t + 1$.¹⁴ The bound on subjective price expectations imposed in Assumption 1 is justified by the fact that the price-dividend ratio will be bounded in equilibrium, so that the objective expectation $E_t P_{t+1}/D_t$ will also be bounded.¹⁵

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 in which 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 $\beta_t^D = \beta_t$ and (15) delivers the pricing implications derived in the Bayesian RE asset pricing literature, as surveyed in section 2.

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 .

¹⁵To see this, note that $P_{t+1}/D_{t+1} < \overline{PD}$ implies $E_t[P_{t+1}]/D_t < a\overline{PD} < \infty$ where a denotes the mean dividend growth rate.

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:

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

Under this assumption, the asset pricing equation (15) simplifies to¹⁶

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

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.

¹⁶Some 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.

To improve our understanding of the empirical performance of the model and to illustrate that the results we obtain in our empirical application do not depend on the specific belief system considered, we now derive analytical results for a general nonlinear belief-updating scheme.

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 upward (downward) 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 form¹⁷

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

for given nonlinear 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) , $\beta \in (0, \beta^U)$ 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

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

¹⁷Note that β_t is determined from observations up to period $t - 1$ only. This simplifies the analysis and 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 the work of Adam (2003).

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

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

where

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

Given initial conditions (Y_0, 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 two-dimensional phase diagram for the dynamics of (β_t, β_{t-1}) , which is shown in Figure 2 for the case in which the shocks $(\varepsilon_t^c)^{-\gamma} \varepsilon_t^d$ assume their unconditional mean value ρ_ε .¹⁸ 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.¹⁹ Since we have a difference equation rather than a differential equation, we cannot plot the evolution of expectations exactly, because 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 for some areas in the figure, these jumps will be correspondingly small, as we now explain.

Consider, for example, region A in the diagram. In this area $\beta_t < \beta_{t-1}$ and β_t keeps decreasing, showing that there is *momentum* in price changes. This holds true even if β_t is already at or below its fundamental value $a^{1-\gamma} \rho_\varepsilon$. Provided the updating gain is small, beliefs in region A will slowly move above the 45° line in the direction of the lower left corner of the graph. Yet, once they enter area B, β_t starts increasing, so that in the next period, beliefs will discretely jump into area C. In region C we have $\beta_t > \beta_{t-1}$ and β_t keeps increasing, so that beliefs then display upward momentum. This manifests

¹⁸Appendix A.2 explains in detail the construction of the phase diagram.

¹⁹The 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

itself in an upward and rightward move of beliefs over time, until these reach area D . There, beliefs β_t start decreasing, so that they move in one jump back into area A , thereby displaying *mean reversion*. The elliptic movements of beliefs around $a^{1-\gamma}\rho_\varepsilon$ imply that expectations (and thus the PD ratio) are likely to oscillate in sustained and persistent swings around the RE value.

The effect 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 upward. 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 the work of Cochrane (2005), a serially correlated and mean-reverting

PD ratio gives rise to excess return predictability, that is, it 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.*²⁰

Therefore, up to a linear approximation of the updating function f ,

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

whenever $\Delta\beta_t > 0$ and $\beta_t \leq a^{1-\gamma}\rho_\varepsilon$. Therefore, beliefs have a tendency to increase (decrease) further following an initial increase (decrease) whenever beliefs are at or below (above) 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.²¹

Mean reversion: *Consider an arbitrary initial belief $\beta_t \in (0, \beta^U)$. In the absence of further disturbances ($\varepsilon_{t+j}^d = \varepsilon_{t+j}^c = 0$ for all $j \geq 0$),*

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

Furthermore, if $\beta_t > a^{1-\gamma}$, there is a period $t' \geq t$ such that β_t is nondecreasing between t and t' and nonincreasing 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 (when $\limsup = \liminf$) or stay fluctuating around it forever (when

²⁰The 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.

²¹See Appendix A.3 for the proof under an additional technical assumption.

$\limsup > \liminf$). Any initial deviation, however, is eventually eliminated with the reversion process being monotonic. This result also implies that an upper bound on price beliefs cannot be an absorbing point: if beliefs β_t go up and they get close to the upper bound β^U , they will eventually bounce off this upper bound and go back down toward the RE value.

Summing up, the previous results show that for a general set of 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.

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. We show below a precise sense in which this system of beliefs represents a small deviation from RE.

In line with Assumption 2, we consider agents who hold rational expectations about the dividend and aggregate consumption processes. 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 price growth. 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

$$\begin{aligned} \left(\frac{C_t}{C_{t-1}} \right)^{-\gamma} \frac{P_t}{P_{t-1}} &= b_t + \varepsilon_t \\ b_t &= b_{t-1} + \xi_t \end{aligned} \tag{25}$$

for $\varepsilon_t \sim iiN(0, \sigma_\varepsilon^2)$, $\xi_t \sim iiN(0, \sigma_\xi^2)$, independent of each other and also jointly i.i.d. with ε_t^d and ε_t^c .²² The latter implies $E[(\varepsilon_t, \xi_t) | I_{t-1}] = 0$, where I_{t-1} includes all the variables in the agents' information set at $t-1$, including all prices, endowments, and dividends dated $t-1$ or earlier.

²²Notice that we use the notation $C_t = Y_t + D_t$, so that equation (25) contains only payoff-relevant variables that are beyond the agent's control.

The previous setup encompasses the RE equilibrium beliefs as a special case. Namely, when agents believe $\sigma_\xi^2 = 0$ and assign probability one to $b_0 = a^{1-\gamma}\rho_\varepsilon$, we have that $\beta_t = a^{1-\gamma}\rho_\varepsilon$ for all $t \geq 0$ and prices are as given by RE equilibrium prices in all periods.

In what follows we allow for a nonzero variance σ_ξ^2 , that is, 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 consists of optimally filtering out the persistent component of price growth b_t . Assuming that agents' prior beliefs 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

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

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.²³

These beliefs constitute a small deviation from RE beliefs in the limiting case with vanishing innovations to the random walk process ($\sigma_\xi^2 \rightarrow 0$).

²³In 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.

Agents' prior uncertainty then 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 , so that one recovers the RE equilibrium value for risk-adjusted price growth expectations. This shows that for a *given* distribution of asset price, agents' beliefs are close to RE beliefs for a small value of the gain parameter ($1/\alpha$).²⁴

For our empirical application, we need to slightly modify the updating equation (26) 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 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 (26) to²⁵

$$\beta_t = w \left(\beta_{t-1} + \frac{1}{\alpha} \left[\left(\frac{C_{t-1}}{C_{t-2}} \right)^{-\gamma} \frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right] \right), \quad (27)$$

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 (26), as long as they are below the threshold β^L , whereas for higher beliefs we have that $w(x) \leq x$. The modified algorithm (27) 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 $b_t > \beta^U$.^{26,27}

²⁴We show below that continues to be true when using the equilibrium distribution of asset prices generated by these beliefs.

²⁵The exact functional form for w that we use in the estimation is shown in appendix A.5.

²⁶The issue of bounding beliefs so as to ensure 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 using a projection facility, assuming that agents simply *ignore* observations that would imply updating beliefs beyond the required bound. See the work of Timmermann (1993, 1996), Marcet and Sargent (1989), or Evans and Honkapohja (2001). This approach has two problems. First, it does not arise from Bayesian updating. Second, it 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.

²⁷Adam, Beutel, and Marcet (2014) solve the model without Assumption 1, which is

We now show that for a small value of the gain ($1/\alpha$), agents' beliefs are close to RE beliefs when using the equilibrium distribution of prices generated by these beliefs. More precisely, the setup gives rise to a stationary and ergodic equilibrium outcome in which 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. 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 $1/\alpha \rightarrow 0$.

Stationarity, Ergodicity, and Small Deviations from RE: *Suppose agents' posterior beliefs evolve according to equation (27) and equilibrium prices are determined according to equation (16). Then β_t is geometrically ergodic for sufficiently large α . 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 the work of Duffie and Singleton (1993) and contained in appendix A.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. Geometric ergodicity is required for estimation by MSM.

We explore further in section 6 the connection between agents' beliefs and model outcomes.

5 Quantitative Model Performance

This section evaluates the quantitative performance of the 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 to focus on the ability of the model to explain the specific

computationally much more costly, so that one cannot use MSM to estimate the model. They use the optimal consumption plan that solves the individual investor's problem to price the asset and show how a projection facility then arises endogenously from the wealth effect in the standard case that relative risk aversion is larger than one.

moments of the data described in Table 1.²⁸ Regarding Fact 5 in Table 1, the model gives rise to an equity premium that is much higher than under RE, but for reasonable degrees of relative risk aversion, it tends to fall short of the risk premium observed in the data.

We first evaluate the model’s ability to explain the individual moments, which is the focus of much of the literature on matching stock price volatility. We find that the model can explain the individual moments well. Using t -statistics based on formal asymptotic distribution, we find that in some versions of the model, all t -statistics are at or below 2 in absolute value, even with a moderate relative risk aversion of $\gamma = 5$. Moreover, with this degree of risk aversion, the model can explain up to 50% of the equity premium.

We then turn to the more demanding task of testing if all the moments are accepted jointly by computing chi-square test statistics. Due to their stringency, such test statistics are rarely reported in the consumption-based asset pricing literature. A notable exception is the work of Bansal, Kiku, and Yaron (2013), who test the overidentifying restrictions of a long-run risk model. Different from our approach, they test equilibrium conditions instead of matching statistics. Also, they use a diagonal weighting matrix instead of the optimal weighting matrix in the objective function (29) introduced below.

We find that with a relative risk aversion of $\gamma = 5$, the model fails to pass an overall goodness of fit as long as one includes the equity premium. Yet, the test reaches a moderate p -value of 2.5% once we exclude the risk-free rate from the set of moments to be matched, confirming that it is the equity premium that poses a quantitative challenge to the model.²⁹ With a relative risk aversion of $\gamma = 3$, the p -value increases even further to 7.1% when we again exclude the risk-free rate.

²⁸A 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 (2004) 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.

²⁹The literature suggests a number of other model ingredients, that - once added - would allow generating a higher equity premium. See, for example, ambiguity aversion as in Collard et al. (2011), initially pessimistic expectations as in Cogley and Sargent (2008), or habits in consumption preferences.

Finally, we allow for a very high risk aversion coefficient. Specifically, we set $\gamma = 80$, which is the steady state value of relative risk aversion used in Campbell and Cochrane (1999).³⁰ The model then replicates all moments in Table 1, including the risk premium. In particular, the model generates a quarterly equity premium of 2.0%, slightly below the 2.1% per quarter observed in U.S. data, while still replicating all other asset pricing moments.

The next section explains the 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. This is a simple adaptation of standard MSM to include matching of statistics that are functions of simple moments by using the delta method (see appendix A.6 for details).

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

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

The *four* parameters will be chosen so as to match some or all of the *ten* sample moments in Table 1:³¹

$$\left(\widehat{E}_{r^s}, \widehat{E}_{PD}, \widehat{\sigma}_{r^s}, \widehat{\sigma}_{PD}, \widehat{\rho}_{PD,-1}, \widehat{c}_2^5, \widehat{R}_5^2, \widehat{E}_{r^b}, \widehat{E}_{\frac{\Delta D}{D}}, \widehat{\sigma}_{\frac{\Delta D}{D}} \right). \quad (28)$$

Let $\widehat{\mathcal{S}}_N \in R^s$ denote the subset of sample moments in (28) that will be matched in the estimation, with N denoting the sample size and $s \leq 10$.³²

³⁰This value is reported on page 244 in their paper.

³¹Many elements listed in (28) are not sample moments, but they are nonlinear functions of sample moments. For example, the R^2 coefficient is a function of sample moments. This means we have to use the delta method to adapt standard MSM (see appendix A.6). It would be more precise to refer to the elements in (28) as “sample statistics,” as we do in the appendix. For simplicity, we avoid this terminology in the main text.

³²As discussed before, we exclude the risk premium from some estimations; in those cases, $s < 10$.

Furthermore, let $\tilde{\mathcal{S}}(\theta)$ denote the moments implied by the model for some parameter value θ . The MSM parameter estimate $\hat{\theta}_N$ is defined as

$$\hat{\theta}_N \equiv \arg \min_{\theta} \left[\hat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\theta) \right]' \hat{\Sigma}_{\mathcal{S},N}^{-1} \left[\hat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\theta) \right], \quad (29)$$

where $\hat{\Sigma}_{\mathcal{S},N}$ is an estimate of the variance-covariance matrix of the sample moments $\hat{\mathcal{S}}_N$. The MSM estimate $\hat{\theta}_N$ chooses the model parameter such that the model moments $\tilde{\mathcal{S}}(\theta)$ fit the observed moments $\hat{\mathcal{S}}_N$ as closely as possible in terms of a quadratic form with weighting matrix $\hat{\Sigma}_{\mathcal{S},N}^{-1}$. We estimate $\hat{\Sigma}_{\mathcal{S},N}$ from the data in a standard way. Adapting standard results from MSM, one can prove that for a given list of moments included in $\hat{\mathcal{S}}_N$, the estimate $\hat{\theta}_N$ is consistent and is the best estimate among 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

$$\widehat{W}_N \equiv N \left[\hat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\hat{\theta}_N) \right]' \hat{\Sigma}_{\mathcal{S},N}^{-1} \left[\hat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\hat{\theta}_N) \right] \rightarrow \chi_{s-4}^2 \text{ as } N \rightarrow \infty, \quad (30)$$

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

In our application we find a nearly singular $\hat{\Sigma}_{\mathcal{S},N}$. As shown in appendix A.6, asymptotic results require this matrix to be invertible. The near-singularity indicates that one statistic is nearly redundant (i.e., carries practically no additional information). The appendix describes a procedure for selecting the redundant statistic; it suggests that we drop the coefficient from the five-year-ahead excess return regression \hat{c}_2^5 from the estimation. 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 2. This happens even though information about \hat{c}_2^5 has not been used in the estimation.

5.2 Estimation Results

Table 2 reports estimation outcomes when assuming $\gamma = 5$. The second and third columns in the table report the asset pricing moments from the data

and the estimated standard deviation for each of these moments, respectively. Columns 4 and 5 then show the model moments and the t -statistics, respectively, when estimating the model using all asset pricing moments (except for c_2^5 , which has been excluded for reasons explained in the previous section). All estimations impose the restriction $\delta \leq 1$.

The estimated model reported in columns 4 and 5 of Table 2 quantitatively replicates the volatility of stock returns ($\sigma_{r,s}$), the large volatility and high persistence of the PD ratio ($\sigma_{PD}, \rho_{PD,-1}$), as well as the excess return predictability evidence (c_5^2, R_5^2). This is a remarkable outcome given the assumed time-separable preference structure. The model has some difficulty in replicating the mean stock return and dividend growth, but t -statistics for all other moments have an absolute value well below 2, and more than half of the t -statistics are below 1.

	U.S. data		Estimated model (c_2^5 not included)		Estimated model ($c_2^5, E_{r,s}$ not included)	
	Data Moment $\widehat{S}_{N,i}$	Std. dev. $\widehat{\sigma}_{\widehat{S}_i}$	Model moment $\widetilde{S}_i(\widehat{\theta})$	t -stat.	Model moment $\widetilde{S}_i(\widehat{\theta})$	t -stat.
Quarterly mean stock return $E_{r,s}$	2.25	0.34	1.27	2.70	1.49	2.06
Quarterly mean bond return $E_{r,b}$	0.15	0.19	0.39	-1.27	0.49	-1.78
Mean PD ratio E_{PD}	123.91	21.36	122.50	0.07	119.05	0.23
Std. dev. stock return $\sigma_{r,s}$	11.44	2.71	10.85	0.22	11.60	-0.06
Std. dev. PD ratio σ_{PD}	62.43	17.60	67.55	-0.29	69.59	-0.41
Autocorrel. PD ratio $\rho_{PD,-1}$	0.97	0.01	0.95	0.62	0.95	0.84
Excess return reg. coefficient c_5^2	-0.0041	0.0014	-0.0066	1.79	-0.0067	1.90
R^2 of excess return regression R_5^2	0.2102	0.0825	0.2132	-0.04	0.1995	0.13
Mean dividend growth $E_{\Delta D/D}$	0.41	0.17	0.00	2.79	0.10	1.82
Std. dev. dividend growth $\sigma_{\Delta D/D}$	2.88	0.82	2.37	0.61	2.45	0.52
Discount factor $\widehat{\delta}_N$			0.9959		1.0000	
Gain coefficient $1/\widehat{\alpha}_N$			0.0073		0.0076	
Test statistic \widehat{W}_N			82.6		62.6	
p -value \widehat{W}_N			0.0%		0.0%	

Table 2: Estimation outcome for $\gamma = 5$

The last two columns in Table 2 report the estimation outcome when

dropping the mean stock return $E_{r,s}$ from the estimation and restricting δ to 1, which tends to improve the ability of the model to match individual moments. All t -statistics are then close to or below 2, including the t -statistics for the mean stock return and for c_2^5 that have not been used in the estimation, and the majority of the t -statistics are below 1. This estimation outcome shows that the subjective beliefs model successfully matches individual moments with a relatively low degree of risk aversion. The model also delivers an equity premium of 1% per quarter, that is, nearly half of the value observed in U.S. data (2.1% per quarter).

The measure for the overall goodness of fit \widehat{W}_N and its p -value are reported in the last two rows of Table 2. The statistic is computed using all moments that are included in the estimation. The reported values of \widehat{W}_N are off the chart of the χ^2 distribution, implying that the overall fit of the model is rejected, even if all moments are matched individually.³³ This indicates that some of the joint deviations observed in the data are unlikely to happen given the observed second moments. It also shows that the overall goodness of fit test is considerably more stringent.

To show that the equity premium is indeed the source of the difficulty for passing the overall test, columns 4 and 5 in Table 3 report results obtained when we repeat the estimation excluding the risk-free rate $E_{r,b}$ instead of the stock returns $E_{r,s}$ from the estimation. The estimation imposes the constraint $\widehat{\delta}_N \leq 1$, since most economists believe that values above 1 are unacceptable. This constraint turns out to be binding. The t -statistics for the individual moments included in the estimation are then quite low, but the model fails to replicate the low value for the bond return $E_{r,b}$, which has not been used in the estimation. Despite larger t -statistics, the model now comfortably passes the overall goodness of fit test at the 1% level, as the p -value for the reported $\widehat{W}_N = 12.87$ statistic is 2.5%. The last two columns in Table 3 repeat the estimation when imposing $\gamma = 3$ and $\widehat{\delta}_N = 1$. The performance in terms of matching the moments is then very similar with $\gamma = 5$, but the p -value of the \widehat{W}_N statistic increases to 7.1%.

Figure 3 shows realizations of the time series outcomes for the PD ratio generated from simulating the estimated model from Table 3 with $\gamma = 5$, for

³³The χ^2 distribution has 5 degrees of freedom for the estimations in Table 2, where the last two columns drop a moment but also fix $\delta = 1$. For the estimation in Table 3, we exclude c_2^5 and $E_{r,b}$ from the estimation, but the constraint $\widehat{\delta} \leq 1$ is either binding or imposed, so that we also have 5 degrees of freedom. Similarly, we have 5 degrees of freedom for the estimation in Table 4.

the same number of quarters as numbers of observations in our data sample. The simulated time series display price booms and busts, similar to the ones displayed in Figure 1 for the actual data, so that the model also passes an informal “eyeball test.”

The estimated gain coefficients in Tables 2 and 3 are fairly small. The estimate in Table 3 implies that agents’ risk-adjusted return expectations respond only 0.7% in the direction of the last observed forecast error, suggesting that the system of price beliefs in our model indeed represents only a small deviation from RE beliefs. Under strict RE, the reaction to forecast errors is 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 $R_5^2 = 0$.

	U.S. data	Estimated model $\gamma = 5$ (c_2^5 , E_{r^b} not included)		Estimated model $\gamma = 3$ (c_2^3 , E_{r^b} not included)	
	Data moment $\widehat{S}_{N,i}$	Model moment $\widetilde{S}_i(\widehat{\theta})$	t -stat.	Model moment $\widetilde{S}_i(\widehat{\theta})$	t -stat.
Quarterly mean stock return E_{r^s}	2.25	1.32	2.50	1.51	2.00
Quarterly mean bond return E_{r^b}	0.15	1.09	-4.90	1.30	-5.98
Mean PD ratio E_{PD}	123.91	109.66	0.69	111.28	0.58
Std. dev. stock return σ_{r^s}	11.44	5.34	2.25	5.10	2.33
Std. dev. PD ratio σ_{PD}	62.43	40.09	1.33	39.11	1.31
Autocorrel. PD ratio $\rho_{PD,-1}$	0.97	0.96	0.30	0.96	0.23
Excess return reg. coefficient c_5^2	-0.0041	-0.0050	0.64	-0.0050	0.60
R^2 of excess return regression R_5^2	0.2102	0.2282	-0.22	0.2302	-0.24
Mean dividend growth $E_{\Delta D/D}$	0.41	0.22	1.14	0.43	-0.09
Std. dev. dividend growth $\sigma_{\Delta D/D}$	2.88	1.28	1.95	1.23	2.00
Discount factor $\widehat{\delta}_N$		1.0000		1.0000	
Gain coefficient $1/\widehat{\alpha}_N$		0.0072		0.0071	
Test statistic \widehat{W}_N		12.87		11.07	
p -value of \widehat{W}_N		2.5%		7.1%	

Table 3: Estimation outcome for $\gamma = 5$ and $\gamma = 3$

To further examine what it takes to match the risk premium and to compare more carefully our results with the performance of other models in the

literature, we now assume a high degree of risk aversion of $\gamma = 80$, in line with the steady state degree of risk aversion assumed in the work of Campbell and Cochrane (1999). Furthermore, we use all asset pricing moments listed in equation (28) for estimation, except for c_2^5 . The estimation results are reported in Table 4. The learning model then successfully replicates all moments in the data, including the risk premium: all the t -statistics for the individual moments are below 2 in absolute value, with most of them even assuming values below 1. For sufficiently high risk aversion, we thus match all individual moments, so that the model performance is comparable to that of Campbell and Cochrane (1999) but achieved with a time-separable preference specification. Yet, the p -value for the test statistic \widehat{W}_N in Table 4 is again off the charts, implying that the model fails the overall goodness of fit test. This highlights that the \widehat{W}_N test statistic is a much stricter test than imposed by matching moments individually.

	U.S. data	Estimated model (c_2^5 not included)	
	Data moment $\widehat{S}_{N,i}$	Model moment $\widetilde{S}_i(\widehat{\theta})$	t -stat.
Quarterly mean stock return E_{r^s}	2.25	2.11	0.40
Quarterly mean bond return E_{r^b}	0.15	0.11	0.21
Mean PD ratio E_{PD}	123.91	115.75	0.38
Std. dev. stock return σ_{r^s}	11.44	16.31	-1.80
Std. dev. PD ratio σ_{PD}	62.43	71.15	-0.50
Autocorrel. PD ratio $\rho_{PD,-1}$	0.97	0.95	1.13
Excess return reg. coefficient c_5^2	-0.0041	-0.0061	1.39
R^2 of excess return regression R_5^2	0.2102	0.2523	-0.51
Mean dividend growth $E_{\Delta D/D}$	0.41	0.16	1.50
Std. dev. dividend growth $\sigma_{\Delta D/D}$	2.88	4.41	1.86
Discount factor $\widehat{\delta}_N$		0.998	
Gain coefficient $1/\widehat{\alpha}_N$		0.0021	
Test statistic \widehat{W}_N		28.8	
p -value of \widehat{W}_N		0.0%	

Table 4: Estimation outcome for $\gamma = 80$




Figure 3: Simulated PD ratio, estimated model from Table 3 ($\gamma = 5$)

Interestingly, the learning model gives rise to a significantly larger risk premium than its RE counterpart.³⁴ For the estimated parameter values in Table 4, the quarterly real risk premium under RE is less than 0.5%, which falls short of the 2.0% emerging in the model with learning.³⁵ Surprisingly, the model generates a small, positive ex post risk premium for stocks even when investors are risk neutral ($\gamma = 0$). This finding may be surprising, since we did not introduce any feature in the model to generate a risk premium. To understand why this occurs, 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 and 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 1 in our simulations of the learning model, whereas it is slightly larger than 1 under RE.³⁶ 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 Robustness of Results

This section discusses the robustness of our findings with regard to different learning specifications and parameter choices (section 6.1), analyzes in detail the extent to which agents' forecasts could be rejected by the data or the

³⁴The RE counterpart is the model with the same parameterization, except for $1/\alpha = 0$.

³⁵The learning model and the RE model imply the same risk-free rate, because we assumed that agents have objective beliefs about the aggregate consumption and dividend process.

³⁶For the learning model, we choose the RE-PD ratio as our starting value.

equilibrium outcomes of the model (section 6.2), and finally offers a discussion of the rationality of agents' expectations about their own future choices (section 6.3).

6.1 Different Parameters and Learning Specifications

We explored the robustness of the model along a number of dimensions. Performance turns out to be robust as long as agents are learning in some way about price growth using past price growth observations. For example, Adam, Beutel, and Marcet (2014) use a model in which agents learn directly about price growth (without risk adjustment) using observations of past price growth; they document a very similar quantitative performance. Adam and Marcet (2010) considered learning about returns using past observations of returns, showing how this leads to asset price booms and busts. Furthermore, within the setting analyzed in the present paper, results are robust to relaxing Assumption 2. For example, the asset pricing moments are virtually unchanged when considering agents who also learn about risk-adjusted dividend growth, using the same weight $1/\alpha$ for the learning mechanism as for risk-adjusted price growth rates. Indeed, given the estimated gain parameter, adding learning about risk-adjusted dividend growth contributes close to nothing to replicating stock price volatility. We also explored a model of learning about risk-adjusted price growth that switches between ordinary least squares learning and constant gain learning, as in the work of Marcet and Nicolini (2003). Again, model performance turns out to be robust. Taken together, these findings suggest that the model continues to deliver an empirically appealing fit, as long as expected capital gains are positively affected by past observations of capital gains.

The model fails to deliver a good fit with the data if one assumes that agents learn only about the relationship between prices and dividends, say about the coefficient in front of D_t in the RE pricing equation (10), using the past observed relationship between prices and dividends (see the work of Timmermann (1996)). Stock price volatility then drops significantly below that observed in the data, illustrating that the asset pricing results are sensitive to the kind of learning introduced in the model. Our finding is that introducing uncertainty about the growth rate of prices is key for understanding asset price volatility.

Similarly, for lower degrees of relative risk aversion around 2, we find that the model continues to generate substantial volatility in stock prices but not

enough to quantitatively match the data.

At the same time, it is not difficult to obtain an even better fit than the one reported in section 5.2. For example, we imposed the restriction $\widehat{\delta}_N \leq 1$ in the estimations reported in Table 3. Yet, in a setting with output growth and uncertainty, values above 1 are easily compatible with a well-defined model and positive real interest rates. Reestimating Table 3 for $\gamma = 5$ without imposing the restriction on the discount factor, one obtains $\widehat{\delta}_N = 1.0094$ and a p -value of 4.3% for the overall fit instead of the 2.5% reported. The fit could similarly be improved by changing the parameters of the projection facility. Choosing $(\beta^L, \beta^U) = (200, 400)$ for the estimation in Table 3 with $\gamma = 5$ instead of the baseline values $(\beta^L, \beta^U) = (250, 500)$ raises the p -value from 2.5% to 3.1%.³⁷

6.2 Testing for the Rationality of Price Expectations

In section 4.3 we presented limiting results that guarantee that agents' beliefs constitute only a small deviation from RE, in the sense that for an arbitrarily small gain, agents' beliefs are close to the beliefs of agents in a rational expectations model. This section studies to what extent agents could discover that their system of beliefs is not exactly correct by observing processes for (P_t, D_t, C_t) .³⁸ We study this for the beliefs implied by the estimated models from section 5.1.

In a first step, we derive a set of testable restrictions implied by agents' beliefs system (2), (3), and (25). Importantly, under standard assumptions, any process satisfying these testable restrictions can - in terms of its auto-covariance function - be generated by the postulated system of beliefs. The set of derived restrictions thus fully characterizes the second-moment implications of the beliefs system.

In a second step, we test the derived restrictions against the data. We show that the data uniformly accept all testable second-order restrictions. This continues to be the case when we consider certain higher-order or non-linear tests that go beyond second-moment implications. Based on this, we can conclude that the agents' belief system is reasonable: given the behavior of actual data, the belief system is one that agents could have entertained.

³⁷Choosing $(\beta^L, \beta^U) = (300, 600)$ causes the p -value to decrease to 1.8%.

³⁸Here, C_t denotes aggregate consumption, that is, $C_t = Y_t - D_t$, which agents take as given.

In a third step, we test the derived restrictions against simulated model data. Again, we find that the restrictions are often accepted in line with the significance level of the test, although for some of the models and some of the tests, we obtain more rejections than implied by the significance level, especially when considering longer samples of artificial data. Since the testable implications are accepted by the actual data, rejections obtained from simulated data indicate areas in which the asset pricing model could be improved further.

6.2.1 Testable Restrictions

In order to routinely use asymptotic theory, we transform the variables into stationary ones and consider the joint implications of the belief system (2), (3), and (25) for the vector $x_t = (e_t, D_t/D_{t-1}, C_t/C_{t-1})$, where

$$e_t \equiv \Delta \left(\frac{C_t}{C_{t-1}} \right)^{-\gamma} \frac{P_t}{P_{t-1}}, \quad (31)$$

with Δ denoting the difference operator.³⁹ The following proposition presents a set of testable restrictions about $\{x_t\}$.⁴⁰

Proposition 1 (Necessity of Restrictions 1-4) *If $\{x_t\}$ follows the system of beliefs (2), (3), and (25), then*

- Restriction 1* : $E(x_{t-i}e_t) = 0$ for all $i \geq 2$
- Restriction 2* : $E \left(\left(\frac{D_t}{D_{t-1}} + \frac{D_{t-1}}{D_{t-2}}, \frac{C_t}{C_{t-1}} + \frac{C_{t-1}}{C_{t-2}} \right) e_t \right) = 0$
- Restriction 3* : $b'_{DC} \Sigma_{DC} b_{DC} + E(e_t e_{t-1}) < 0$
- Restriction 4* : $E(e_t) = 0$,

³⁹One might be tempted to test (31) using an augmented Dickey-Fuller (ADF) test, which involves running a regression with a certain number of lags, and test if the residual is serially correlated. This approach is problematic in our application: as shown in appendix A.7, we have $e_t = \varepsilon_t - \varepsilon_{t-1} + \xi_t$; since the gain is small in the estimates in Tables 2 - 4, we also have that $\sigma_\xi^2/\sigma_\varepsilon^2$ is small, so that the MA representation of e_t has a near unit root. In this case, the true autoregressive representation of $\Delta \left(\frac{C_t}{C_{t-1}} \right)^{-\gamma} \frac{P_t}{P_{t-1}}$ has coefficients that decay very slowly with the lag length. The ADF test then does not work: if we introduce only a few lags into the regression, then the error would be serially correlated and the test would be asymptotically invalid; if we introduce many lags, then the test has little power for reasonable sample lengths.

⁴⁰Appendix A.7 provides a proof.

where $\Sigma_{DC} \equiv \text{var} \left(\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right)$ and $b_{DC} \equiv \Sigma_{DC}^{-1} E \left[\left(\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right)' e_t \right]$.

Given standard assumptions entertained in the asset pricing literature, it turns out that Restrictions 1-4 in the previous proposition are also sufficient for $\{x_t\}$ to be consistent with the belief system in terms of second-moment implications. In particular, suppose the following holds:

Assumption 3 (i) x_t is second-order stationary; (ii) $\left(\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right)$ is serially uncorrelated and $E \left(\frac{D_t}{D_{t-1}} \right) = E \left(\frac{C_t}{C_{t-1}} \right)$; (iii) $\left(\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right)$ is uncorrelated with e_{t-j} for all $j > 1$.

Conditions (i)-(iii) in the previous assumption hold true in our asset pricing model. We do not question their validity when testing the belief system using actual data, because they are working assumptions maintained by much of the consumption-based asset pricing literature. Appendix A.7 then proves the following result:

Proposition 2 (Sufficiency of Restrictions 1-4) *Suppose the stochastic process $\{x_t\}$ satisfies Assumption 3. If this process also satisfies Restrictions 1-4 stated in Proposition 1, then there exists a belief system of the form (2), (3), and (25), whose autocovariance function is identical to that of $\{x_t\}$.*

The proposition shows that - conditional on Assumption 3 being satisfied - any process satisfying Restrictions 1-4 in Proposition 1 can - in terms of its second-moment implications - be generated from the belief system.

One can derive further higher moment implications from the belief system, based on the observation that e_t in equation (31) has an MA(1) structure, as we show in appendix A.7, and that all variables in the belief system are jointly normally distributed. Under normality, the absence of serial correlation implies independence, so that we have

$$E[z_{t-i}e_t] = 0, \tag{32}$$

where z_{t-i} can be any stationary nonlinear transformation of variables contained in the $t - i$ information set of agents with $i \geq 2$. Obviously, due to the large number of possible instruments z_{t-i} , it is impossible to provide an exhaustive test of (32). We thus simply report tests of (32) based on some natural instruments z_{t-i} .

We test the moment restrictions from Proposition 1 in a standard way, as described in detail in appendix A.8. Testing Restriction 1 involves an infinite number of variables and therefore requires some discretionary choice regarding the set of instruments. We proceed by running separate tests with each of the three elements in x_{t-2} , also including a constant and three lags of the considered element.⁴¹

6.2.2 Testing Beliefs Against Actual Data

Table 5 reports the test statistics when testing Restrictions 1-4 from Proposition 1 using actual data. We compute risk-adjusted consumption growth in the data assuming $\gamma = 5$ (second column) and $\gamma = 80$ (third column).⁴² The 5% critical value of the test statistic is reported in the last column of Table 5. The table shows that the test statistic is in all cases below its critical value and often so by a wide margin. It then follows from Proposition 2 that agents find the observed asset pricing data, in terms of second moments, to be compatible with their belief system.

	Test statistic $\gamma = 5$	Test statistic $\gamma = 80$	5% critical value
Restriction 1 using $\frac{D_{t-i}}{D_{t-i-1}}$	6.69	3.10	9.48
Restriction 1 using $\frac{C_{t-i}}{C_{t-i-1}}$	3.47	0.80	9.48
Restriction 1 using $\Delta \left(\frac{C_{t-i}}{C_{t-i-1}} \right)^{-\gamma} \frac{P_{t-i}}{P_{t-i-1}}$	6.97	1.38	9.48
Restriction 2	0.28	4.31	5.99
Restriction 3	-7.15	-2.96	1.64
Restriction 4	0.01	0.11	3.84

Table 5: Empirical test of \mathcal{P} based on Proposition 1

Table 6 presents further tests based on equation (32) using natural non-linear transformations of the variables x_{t-i} , namely, past PD ratio and past

⁴¹We also performed joint tests that include as instruments a constant, the entire vector x_{t-2} , and three lags of the vector. This leads to very similar conclusions, but in case of a rejection is less informative about which element in x delivers the rejection.

⁴²We use the consumption data provided by Campbell and Cochrane (1999), which is available for the period 1947-1994.

price growth. As before, tests include a constant and three lags of the stated variable. Test statistics are again below the 5% critical value in all cases. Taken together with the evidence from Table 5, this shows that agents' belief system is a reasonable one to hold, given the way the data actually behaved.

Instrument	Test statistic $\gamma = 5$	Test statistic $\gamma = 80$	5% critical value
$\frac{P_{t-i}}{D_{t-i}}$	6.33	2.90	9.48
$\frac{P_{t-i}}{P_{t-i-1}}$	4.68	4.50	9.48

Table 6: Empirical test of \mathcal{P} , additional instruments

6.2.3 Testing Beliefs Against Simulated Data

Table 7 reports the rejection frequencies for Restriction 1 when using simulated model data.⁴³ Specifically, the table reports the likelihood of rejecting Restriction 1 at the 5% significance level, using simulated data based on the point estimates from Table 3 ($\gamma = 5$) and the point estimates from Table 4 ($\gamma = 80$). Rejection frequencies are shown for different instruments, different sample lengths T of simulated quarterly data, and different numbers of lags in the tests. The longest sample length corresponds to the length of the data sample.

Obviously, one cannot expect that this test is never rejected. Even the correct model would be rejected because of Type I errors (i.e., about 5% of the time). One can evaluate agents' subjective beliefs within the model by checking whether the rejection frequencies exceed the 5% significance level.

Table 7 shows that with 60 to 100 quarters of simulated data, the rejection frequencies for the different tests considered are about as many times below the 5% level as they are above this level. With 200 or 340 quarters of data, the rejection rates are higher for the dividend growth instruments using parameters from Table 3 and for risk-adjusted price and consumption growth for the model from Table 4. This indicates that these variables may help to improve agents' forecasts within the model. Yet, given that these same

⁴³The rejection frequencies are obtained from simulating 1,000 random samples of the specified length.

tests are not rejected when using actual data, these rejection rates suggest dimensions along which the model could be further improved.

Instrument	lags\T:	Model from Table 3, $\gamma = 5$				Model from Table 4			
		60	100	200	340	60	100	200	340
$\frac{D_{t-i}}{D_{t-i-1}}$	1	5.0%	6.9%	13.5%	19.6%	1.7%	2.1%	2.5%	1.4%
	3	4.9%	9.7%	18.5%	26.2%	3.6%	3.7%	3.4%	2.8%
$\frac{C_{t-i}}{C_{t-i-1}}$	1	2.3%	3.7%	6.0%	5.8%	10.3%	18.7%	29.0%	44.0%
	3	5.3%	5.7%	9.9%	11.3%	12.0%	21.3%	38.0%	56.7%
$\Delta \left(\frac{C_{t-i}}{C_{t-i-1}} \right)^{-\gamma} \frac{P_{t-i}}{P_{t-i-1}}$	1	1.8%	1.9%	1.3%	0.8%	5.8%	10.9%	13.8%	18.5%
	3	4.3%	4.6%	9.0%	16.0%	9.0%	12.6%	25.1%	36.5%

Rejection frequency at 5% significance level
Table 7: Test of Restriction 1 using simulated data

Table 8 reports the rejection frequencies for the additional nonlinear instruments P_{t-i}/D_{t-i} and P_{t-i}/P_{t-i-1} . Table 8 shows that with regard to these additional variables, there is a tendency to reject the null more often than 5% when considering the model estimates from Table 3, but rejection frequencies are in line with the significance level for the model from Table 4.

Instrument	# of lags\T:	Model from Table 3, $\gamma = 5$				Model from Table 4			
		60	100	200	340	60	100	200	340
$\frac{P_{t-2}}{D_{t-2}}$	1	3.6%	5.7%	33.8%	69.6%	2.7%	1.8%	2.0%	1.4%
	3	4.4%	8.6%	20.6%	35.1%	4.9%	5.0%	5.0%	5.3%
$\frac{P_{t-2}}{P_{t-3}}$	1	8.3%	17.3%	16.5%	29.1%	3.6%	2.7%	3.2%	5.2%
	3	5.2%	11.8%	19.3%	39.1%	6.0%	5.9%	6.0%	6.8%

Rejection frequency at 5% significance level
Table 8: Tests using additional instrument

A similar outcome can be documented when testing Restrictions 2-4 from Proposition 1 on simulated data, as reported in Table 9. Although the model

from Table 4 comfortably passes these restrictions, the model from Table 3 generates too many rejections for Restrictions 2 and 3. Again, with these tests being accepted in the actual data, these findings suggest that the model from Table 3 could be further improved.

Overall, we conclude that it will not be easy for agents to reject their beliefs upon observing the model-generated data. Although some tests reject too often relative to the significance level, others reject too little. Clearly, upon diagnosing a rejection, agents may choose to reformulate their forecasting model, possibly by including additional regressors in the belief system (25). Although investigating the implications of such belief changes is of interest, the fact that agents' beliefs are compatible with the actual data (see the previous section) shows that some of the results from Tables 7-9 indicate dimensions along which the asset pricing model can be further improved in order to match more closely the behavior of the actual asset pricing data. We leave this issue to further research.

T:	Model from Table 3, $\gamma = 5$				Model from Table 4			
	60	100	200	340	60	100	200	340
Restriction 2	57.4%	61%	72.7%	85.1%	2.8%	2.2%	1.1%	1.4%
Restriction 3	72.6%	75.7%	97.0%	100%	6.3%	3.6%	1.3%	0.2%
Restriction 4	0.1%	0.0%	0.0%	0.0%	0.3%	0.0%	0.0%	0.0%

Rejection frequency at 5% significance level

Table 9: Test of Restrictions 2-4 on simulated data

6.3 Subjective versus Objective Plans

This section discusses to what extent agents' expectations about their own future consumption and stockholding choices coincide with the objective expectations of future choices.

It is important to note that agents hold the correct perception regarding their own choices *conditional* on the realizations of the future values of the variables P , Y , and D . This is the case because agents choose an optimal plan (C_t^i, S_t^i, B_t^i) satisfying (5) and make decisions according to this plan, so that $E^{\mathcal{P}}[C_t^i | \omega^t] = E[C_t^i | \omega^t]$ for all t and ω^t . Nevertheless, the fact that agents hold expectations about ω^t that are not exactly equal to those realized within the model means that expectations about (C_t^i, S_t^i, B_t^i) that

condition on less information may differ from the true expectations given such a reduced information set. This fact highlights that discrepancies between agents' subjective expectations about their own actions and objective expectations about these actions indicate the presence of subjective beliefs about contingencies (i.e., prices), as explored in the previous section.

We first show that the gap between subjective and objective consumption growth expectations is approximately zero. This gap can be expressed as⁴⁴

$$\begin{aligned}
& E_t^{\mathcal{P}} [C_{t+1}^i/C_t] - E_t [C_{t+1}^i/C_t] \\
= & E_t^{\mathcal{P}} [C_{t+1}^i/C_t] - E_t [C_{t+1}^i/C_t] \\
= & E_t^{\mathcal{P}} \left[\frac{P_{t+1}(1 - S_{t+1}^i) - B_{t+1}^i + D_{t+1} + Y_t}{D_t + Y_t} \right] - E_t [C_{t+1}^i/C_t] \\
= & E_t^{\mathcal{P}} \left[\frac{P_{t+1}(1 - S_{t+1}^i) - B_{t+1}^i}{D_t + Y_t} \right].
\end{aligned}$$

Since the choices for S_{t+1}^i and B_{t+1}^i are bounded, we have that $P_{t+1}(1 - S_{t+1}^i) - B_{t+1}^i$ is bounded whenever beliefs about future values of the price P_{t+1} are bounded. Assumption 1 in section 4.2 then ensures that the subjective expectations in the last line of the preceding equation will be approximately equal to zero for every state, so that the gap between the subjective and the objective consumption expectations vanishes under the maintained assumptions.

The situation is different when considering subjective stockholding plans. Assuming an interior solution for stockholdings in period t , the agent's first-order condition satisfies in equilibrium

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

With C_{t+1}^i/C_t converging to C_{t+1}/C_t under Assumption 1, this equation recovers our pricing equation. Yet, if this equation were to hold almost surely for each period and each contingency in the future under the agent's subjective plan, then one could iterate forward on this equation and obtain (under

⁴⁴The first equality uses the fact that under the objective beliefs $E [C_{t+1}^i] = E [C_t]$ in equilibrium, the second equality the investor's budget constraint and the fact that $S_t^i = 1$ and $B_t^i = 0$ in equilibrium in t , and the last equality the fact agents have rational expectations about the processes for aggregate consumption.

a suitable transversality condition) that the equilibrium price must equal the present value of dividends. Since the equilibrium price is different, it must be the case that the agent expects that with positive probability, either $S_t^i = \underline{S}$ or $S_t^i = \overline{S}$ will hold in the future. These expectations, however, will not be fulfilled in equilibrium, because the agent will never buy or sell, as we have $S_t^i = 1$ for all t along the equilibrium path. It remains to be explored to what extent this gap between agents' expectations and equilibrium outcomes can be reduced by introducing agent heterogeneity in terms of preferences or beliefs.

7 Conclusions and Outlook

A 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 are internally rational, in the sense that they formulate their doubts about market outcomes using a consistent set of subjective beliefs about prices and maximize expected utility given this set of beliefs. The system of beliefs is internally consistent in the sense that it specifies a proper joint distribution of prices and fundamental shocks at all dates. Furthermore, the perceived distribution of price behavior, although different from the true distribution, is nevertheless close to it and the discrepancies are hard to detect.

In such a setting, optimal behavior 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 (RE) value emerge, even though all agents act rationally in light of their beliefs.

We also submit our consumption-based asset pricing model to a formal econometric test based on the method of simulated moments. The model performs remarkably well, despite its simplicity. Although the model gives rise to a significant equity premium, it fails to match the empirical premium found for reasonable degrees of relative risk aversion. When risk aversion is as high as in some of the previous work, the 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 in a setting with time-separable preferences and RE, our results suggest that learning about price behavior may be a crucial ingredient in understanding stock price volatility. 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.

A natural question arising within our setting is to what extent the present theory could be used to also price other assets, say, the term structure of interest rates or the cross section of stock returns. Exploring these pricing implications appears to be an interesting avenue for further research.

The finding that large asset price fluctuations can result from optimizing agents with subjective beliefs 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., J. BEUTEL, AND A. MARCET (2014): “Stock Price Booms and Expected Capital Gains,” CEPR Discussion Paper No. 9988.
- ADAM, K., AND A. MARCET (2010): “Booms and Busts in Asset Prices,” IMES Discussion Paper Series 10-E-02, Institute for Monetary and Economic Studies, Bank of Japan.
- (2011): “Internal Rationality, Imperfect Market Knowledge and Asset Prices,” *Journal of Economic Theory*, 146, 1224–1252.
- ARELLANO, M., AND S. BOND (1991): “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies*, 58, 277–297.
- BANSAL, R., D. KIKU, AND A. YARON (2013): “Risks for the Long Run: Estimation with Time Aggregation,” University of Pennsylvania Working Paper.
- 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. (2004): “Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory,” *Review of Economic Studies*, 71, 63–86.
- BOSWIJK, P., C. H. HOMMES, AND S. MANZAN (2007): “Behavioral Heterogeneity in Stock Prices,” *Journal of Economic Dynamics and Control*, 31, 1938–1970.
- BRANCH, W., AND G. W. EVANS (2010): “Asset Return Dynamics and Learning,” *Review of Financial Studies*, 23, 1651–1680.
- BRANCH, W. A., AND G. EVANS (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3, 159–191.

- 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 (2008): “Asset Pricing with Adaptive Learning,” *Review of Economic Dynamics*, 11, 629–651.
- 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*, rev. ed. Princeton University Press, Princeton.
- 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.
- COLLARD, F., S. MUKERJI, K. SHEPPARD, AND J.-M. TALLON (2011): “Ambiguity and the Historical Equity Premium,” University of Bern mimeo.
- 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 Models,” *Econometrica*, 50, 1269–1286.
- KREPS, D. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge University Press, Cambridge.
- LANSING, K. J. (2010): “Rational and Near-Rational Bubbles Without Drift,” *Economic Journal*, 120, 1149–1174.
- LEROY, S. F., AND R. D. PORTER (1981): “The Present-Value Relation: Tests Based on Implied Variance Bounds,” *Econometrica*, 49, 555–574.
- LUCAS, R. E. JR. (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.
- 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*, 2nd ed. Princeton University Press, Princeton.

TIMMERMANN, A. G. (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.

A Appendix

A.1 Data Sources

Our data are for the United States and have been downloaded from the Global Financial Database (<http://www.globalfinancialdata.com>). The period covered is 1925:4-2012:2. For the subperiod 1925:4-1998:4, our data set corresponds very closely to Campbell's (2003) handbook data set available at <http://scholar.harvard.edu/campbell/data>.

In the estimation part of the paper, we use moments that are based on the same number of observations as we have data points. 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 2007:1. For consistency, the effective sample end for all other moments reported in Table 1 has been shortened by five years to 2007:1. In addition, due to the seasonal adjustment procedure for dividends described below and the way we compute the standard errors for the moments described in appendix A.6, the effective starting date was 1927:2. The names of the data series used are reported below.

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 to 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 the work of Campbell (2003), dividends have been deseasonalized by taking averages of the actual dividend payments over the current and preceding three quarters.

A.2 Details on the Phase Diagram

The second-order difference equation (23) describes the evolution of beliefs over time and allows us 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 as 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 labeled “ $\beta_{t+1} = \beta_t$ ” 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 at $x_1 = x_2 = a^{1-\gamma} \rho_\varepsilon$, which is the RE value and also 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 upward (downward), as indicated by the function $x_2 = \frac{1}{\delta} - \frac{x_1(1-\delta x_1)}{a^{1-\gamma} \delta (\varepsilon^c)^{-\gamma} \varepsilon^d}$.

A.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 adds a grain of truth to 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

$$\Delta\beta_{\tilde{t}} \geq 0 \text{ for all } \tilde{t} \text{ such that } t < \tilde{t} < t', \text{ and} \quad (33)$$

$$\Delta\beta_{t'} < 0. \quad (34)$$

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 $\bar{t} \geq t$ 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 (33) and (34) hold if we take $t' = \bar{t}$.

In case *ii)* β_t cannot 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 (33) and (34) hold for $t' = \bar{t} + 1$.

This shows that (33) and (34) 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

$$\Delta\beta_{t'+j} < 0 \quad \text{and} \quad (35)$$

$$\beta_{t'+j} > a + \eta \quad (36)$$

then

$$\begin{aligned} \Delta\beta_{t'+j+1} &= f_{t'+j+1} \left(a + \frac{\Delta\beta_{t'+j}}{1 - \delta\beta_{t'+j}} - \beta_{t'+j}, \beta_{t'+j} \right) < f_{t'+j+1} (a - \beta_{t'+j}, \beta_{t'+j}) \\ &< f_{t'+j+1} (-\eta, \beta_{t'+j}) \leq -\eta \mathcal{D}_{t'+j+1} \leq 0, \end{aligned} \quad (37)$$

where the first inequality follows from (35), the second inequality from (36), and the third from the mean value theorem, $\eta > 0$ and $\mathcal{D}_{t'+j+1} \geq 0$. Assume, toward a contradiction, that (36) holds for all $j \geq 0$. Since (35) holds for $j = 0$, it follows by induction that $\Delta\beta_{t'+j} \leq 0$ for all $j \geq 0$ and, therefore, that (37) 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 (36) cannot 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.

A.4 Proof of Geometric Ergodicity

Defining $\eta_t \equiv (\varepsilon_t^c)^{-\gamma} \varepsilon_t^d$ and using (22) and (27), we can write the learning algorithm that gives the dynamics of β_t as

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

where the first element of F , denoted F_1 , is given by the right side of (27) 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 page 942 in the work of 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

$$\lambda_t^+ \leq B_t + \frac{f_{A_t}(1 - \frac{1}{\alpha})}{2} < B_t + 1 - \frac{1}{\alpha} \quad \text{for } \Delta\beta_{t-1} \geq 0, \quad (38)$$

where we used

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

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

$$\lambda_t^+ \leq \frac{1}{2}f_{A_t}\left(1 - \frac{1}{\alpha} + B_t\right) \leq \frac{f_{A_t}\left(1 - \frac{1}{\alpha}\right)}{2} < 1 - \frac{1}{\alpha} \quad \text{for } \Delta\beta_{t-1} < 0. \quad (39)$$

From (27) we have

$$\Delta\beta_t \leq \frac{1}{\alpha} \left(\eta_{t-1} a^{1-\gamma} \left[1 + \frac{\Delta\beta_{t-1}}{1 - \delta\beta_{t-1}} \right] - \beta_{t-1} \right). \quad (40)$$

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 (39) and using the inequality for (38), 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 the “ 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

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

which is smaller than one for large enough α .

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

Now, adding $a^{1-\gamma}\eta_{t-1}$ to both sides of (40) and taking expectations at the ergodic distribution, we have

$$E(\beta_{t-1} - \eta_{t-1}a^{1-\gamma}) \leq E\left(\frac{\Delta\beta_{t-1}}{1 - \delta\beta_{t-1}}\eta_{t-1}a^{1-\gamma}\right). \quad (41)$$

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 (40) holds as an equality with an arbitrarily large probability. Taking expectations on both sides for the realizations where this holds as an 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. ■

A.5 Differentiable Projection Facility

The function w used in the differentiable projection facility is

$$w(x) = \begin{cases} x & \text{if } x \leq \beta^L \\ \beta^L + \frac{x - \beta^L}{x + \beta^U - 2\beta^L}(\beta^U - \beta^L) & \text{if } \beta^L < x. \end{cases} \quad (42)$$

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. Although the projection facility might suggest that profitable

trading rules could be devised, this is true only if one assumes that the parameters β^U and β^L are fixed and unchanging over time, as we do here for simplicity. In a slightly more realistic model, it would be difficult for agents to time stock purchases and stock sales to exploit the projection facility.

A.6 Details on the MSM Procedure

The estimation method and the proofs adapt the results from a standard MSM estimation. Online appendix B to this paper contains a much more detailed account of these results. We use the definitions introduced at the beginning of section 5.

Let N be the sample size, $(\mathbf{y}_1, \dots, \mathbf{y}_N)$ the observed data sample, with \mathbf{y}_t containing m variables. Define sample moments $\widehat{M}_N \equiv \frac{1}{N} \sum_{t=1}^N h(\mathbf{y}_t)$ for a given *moment function* $h : R^m \rightarrow R^q$. Sample statistics $\widehat{\mathcal{S}}_N$ shown in (28) are not exactly moments of the data. For example, \widehat{M}_N contains the sample moments $\widehat{var}(P_t/D_t)$ and $\widehat{cov}(P_t/D_t, P_{t-1}/D_{t-1})$, and $\widehat{\mathcal{S}}_N$ contains the serial correlation of P_t/D_t which is a function of these moments. The sample statistics can be written as $\widehat{\mathcal{S}}_N \equiv \mathcal{S}(\widehat{M}_N)$ for a *statistics function* $\mathcal{S} : R^q \rightarrow R^s$. The h, \mathcal{S} mappings in our application are written explicitly in online appendix B.

Let $y_t(\theta)$ be the series generated by our model of learning for parameter values θ and some realization of the underlying shocks. Denote the true parameter value θ_0 . Let $M(\theta) \equiv E [h(y_t(\theta))]$ be the moments for parameter values θ at the stationary distribution of $y_t(\theta)$, let $M_0 \equiv M(\theta_0)$ be the true moments, and $\tilde{\mathcal{S}}(\theta) \equiv \mathcal{S}(M(\theta))$ the statistics for parameter θ . Denote by M_0^j the true j -th autocovariance

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

Define $S_w \equiv \sum_{j=-\infty}^{\infty} M_0^j$. A consistent estimator $\widehat{S}_{w,N} \rightarrow S_w$ is found by using standard Newey-West estimators. The variance for the sample statistics $\widehat{\mathcal{S}}_N$ reported in the second column of Table 1 is given by

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

Note that the model is not needed for this estimator; we use only observed data. The exact form of $\frac{\partial \mathcal{S}(M)}{\partial M'}$ is in online appendix B.

Duffie and Singleton (1993) show that to apply standard MSM asymptotics, one needs geometric ergodicity, and we showed that this holds in our model in section 4.3. Note that the smooth bounding function w in equation (27) guarantees that a Monte Carlo approximation to $\tilde{\mathcal{S}}$ is differentiable, as is required for an MSM asymptotic distribution.

Letting $\Sigma_{\mathcal{S}}$ be the asymptotic variance-covariance matrix of the sample statistics, under standard assumptions it can be shown that

$$\widehat{\Sigma}_{\mathcal{S},N} \rightarrow \Sigma_{\mathcal{S}} \text{ and } \widehat{\theta}_N \rightarrow \theta_0 \text{ a.s. as } N \rightarrow \infty. \quad (43)$$

Also, letting $\mathcal{B}_0 \equiv \frac{\partial M'(\theta_0)}{\partial \theta} \frac{\partial \mathcal{S}'(M_0)}{\partial M}$, it can be shown that

$$\sqrt{N} \left[\widehat{\theta}_N - \theta_0 \right] \rightarrow \mathcal{N}(0, (\mathcal{B}_0 \Sigma_{\mathcal{S}}^{-1} \mathcal{B}_0)^{-1}) \quad (44)$$

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

$$\widehat{W}_N \equiv N \left[\widehat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\widehat{\theta}) \right]' \widehat{\Sigma}_{\mathcal{S},N}^{-1} \left[\widehat{\mathcal{S}}_N - \tilde{\mathcal{S}}(\widehat{\theta}) \right] \rightarrow \chi_s^2 \quad (46)$$

in distribution as $N \rightarrow \infty$. Also, the weighting matrix $\widehat{\Sigma}_{\mathcal{S},N}^{-1}$ is optimal among all weighting matrices of the statistics. The t -statistics in Tables 2-4 use variances from (45), and the p -values for \widehat{W}_N are based on (46).

As can be seen from the above formula, we need to invert $\widehat{\Sigma}_{\mathcal{S},N}$ and its limit $\Sigma_{\mathcal{S}}$. A nearly singular $\widehat{\Sigma}_{\mathcal{S},N}$ presents problems because the distribution of \widehat{W}_N is nearly ill-defined, and the distribution in short samples is not close to a χ^2 . This singularity occurs because one of the statistics in $\widehat{\mathcal{S}}_N$ is nearly perfectly correlated with all the others. This only means that this is a redundant statistic, so we can drop it from $\widehat{\mathcal{S}}_N$ in the estimation. To select the redundant statistic, we predict each element of $\widehat{\mathcal{S}}_N$ with all the others according to $\widehat{\Sigma}_{\mathcal{S},N}$ and drop the statistic for which the R^2 is less than 1%. As it turns out, this occurs only for \widehat{c}_2^5 with an $R^2 = 0.006$.

A.7 Proof of Propositions 1 and 2

Proof of Proposition 1: Note that the system of beliefs implies that e_t , defined in equation (31), is given by

$$e_t = \varepsilon_t + \xi_t - \varepsilon_{t-1}, \quad (47)$$

so that Restriction 1 holds.

We also have $E\left(\frac{D_t}{D_{t-1}}e_t\right) = E(\varepsilon_t^D \varepsilon_t) = E(\varepsilon_{t-1}^D \varepsilon_{t-1}) = -E\left(\frac{D_{t-1}}{D_{t-2}}e_t\right)$.⁴⁵ Together with the analogous derivation for consumption growth, this delivers Restriction 2. From (47) we get

$$E(\varepsilon_t^2) = -E(e_t e_{t-1}) = E(\varepsilon_{t-1}^2). \quad (48)$$

Let $Proj(X|Y)$ denote the linear projection of a random variable X on a random vector Y . Then $Proj(\varepsilon_t|\varepsilon_t^D, \varepsilon_t^C) = (\varepsilon_t^D, \varepsilon_t^C)b_{DC}$ and using properties of linear projections, we have

$$E(\varepsilon_t^2) > var\left(Proj(\varepsilon_t|\varepsilon_t^D, \varepsilon_t^C)\right) = b'_{DC} \Sigma_{DC} b_{DC}.$$

Together with (48), this implies Restriction 3. Restriction 4 follows directly from (25).

Proof of Proposition 2: Consider a process $\{x_t\}_{-\infty}^{\infty} = \{e_t, D_t/D_{t-1}, C_t/C_{t-1}\}_{-\infty}^{\infty}$ satisfying Assumption 3 and Restrictions 1-4 from Proposition 1, where Assumption 3 (i) ensures that well-defined second moments exist. We then show how to construct a stationary process $\{\tilde{x}_t\}_{-\infty}^{\infty} = \{\tilde{e}_t, \widetilde{D_t/D_{t-1}}, \widetilde{C_t/C_{t-1}}\}_{-\infty}^{\infty}$ consistent with the belief system (2), (3), and (25) that has the same autocovariance function as $\{x_t\}_{-\infty}^{\infty}$. In particular, let $\left\{\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C, \tilde{\xi}_t, \tilde{\eta}_t\right\}_{-\infty}^{\infty}$ denote a white noise sequence, in which $(\tilde{\xi}_t, \tilde{\eta}_t)$ are uncorrelated contemporaneously with each other and with $(\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C)$, and $var(\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C) = \Sigma_{DC}$. The variances of $\tilde{\xi}_t$ and $\tilde{\eta}_t$ are determined from observable moments as follows:

$$\begin{aligned} \sigma_{\tilde{\xi}}^2 &= \sigma_e^2 + 2\sigma_{e,-1} \\ \sigma_{\tilde{\eta}}^2 &= -\sigma_{e,-1} - b'_{DC} \Sigma_{DC} b_{DC}, \end{aligned}$$

where $\sigma_{e,-1} \equiv E(e_t e_{t-1})$ and $\sigma_e^2 = E[e_t^2]$.

⁴⁵We slightly abuse notation in the proof because dividend and consumption growth follow the form

$$\begin{aligned} D_t/D_{t-1} &= a + a(\varepsilon_t^D - 1) \\ C_t/C_{t-1} &= a + a(\varepsilon_t^C - 1), \end{aligned}$$

where the last terms are mean zero innovations. When writing ε_t^D and ε_t^C in the proof, we actually mean $a(\varepsilon_t^D - 1)$ and $a(\varepsilon_t^C - 1)$, respectively.

Since x satisfies Restriction 3, it follows that $\sigma_{\eta}^2 > 0$. To see that $\sigma_{\xi}^2 \geq 0$ holds, note that Restriction 1 implies that the observed univariate process e_t is MA(1). Hence, we can write

$$e_t = u_t - \theta u_{t-1} \quad (49)$$

for some constant $|\theta| \leq 1$ and some white noise u_t . We thus have

$$\sigma_e^2 = \sigma_u^2(1 + \theta^2) \geq 2\theta\sigma_u^2 = -2\sigma_{e,-1},$$

where the last equality holds because (49) implies $\sigma_{e,-1} = -\theta\sigma_u^2$. Hence, $\sigma_{\xi}^2 \geq 0$. This proves that under the assumptions of this proposition, one can build a process $\{\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C, \tilde{\xi}_t, \tilde{\eta}_t\}_{-\infty}^{\infty}$ satisfying all the properties we have assumed about this process.

In line with (2) and (3), we then let

$$\begin{aligned} \widetilde{D_t/D_{t-1}} &= E[D_t/D_{t-1}] + \tilde{\varepsilon}_t^D \\ \widetilde{C_t/C_{t-1}} &= E[C_t/C_{t-1}] + \tilde{\varepsilon}_t^C. \end{aligned}$$

Part (ii) of Assumption 3 then implies that $\{D_t/D_{t-1}, C_t/C_{t-1}\}_{t=-\infty}^{\infty}$ and $\{\widetilde{D_t/D_{t-1}}, \widetilde{C_t/C_{t-1}}\}_{t=-\infty}^{\infty}$ have the same autocovariance functions. All that remains to be shown is that for some process \tilde{e}_t consistent with the system of beliefs, the covariances of this process with leads and lags of itself and of $(\widetilde{D_t/D_{t-1}}, \widetilde{C_t/C_{t-1}})$ are the same as in the autocovariance function of $\{x_t\}_{-\infty}^{\infty}$.

We construct \tilde{e}_t ; in line with (25), we let

$$\tilde{e}_t = \tilde{\varepsilon}_t + \tilde{\xi}_t - \tilde{\varepsilon}_{t-1}, \quad (50)$$

where

$$\tilde{\varepsilon}_t = (\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C) b_{DC} + \tilde{\eta}_t.$$

It is easy to check that the autocovariance function of $\{\tilde{e}_t\}_{-\infty}^{\infty}$ is - by construction - identical to that of $\{e_t\}_{-\infty}^{\infty}$, since both of them are MA(1) with the same variance and autocovariance.

In a final step, we verify that

$$E \left(\left(\left(\frac{\widetilde{D_{t-i}}}{\widetilde{D_{t-1-i}}}, \frac{\widetilde{C_{t-i}}}{\widetilde{C_{t-1-i}}} \right) \tilde{e}_t \right) \right) = E \left(\left(\left(\frac{D_{t-i}}{D_{t-1-i}}, \frac{C_{t-i}}{C_{t-1-i}} \right) e_t \right) \right) \quad (51)$$

for all $i \leq 0$. Clearly, since $\left\{ \tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C, \tilde{\xi}_t, \tilde{\eta}_t \right\}_{t=-\infty}^{\infty}$ are serially uncorrelated, these covariances are zero for all $i \geq 2$.

For $i = 0$, we have

$$\begin{aligned} E \left(\left(\frac{\widetilde{D}_t}{\widetilde{D}_{t-1}}, \frac{\widetilde{C}_t}{\widetilde{C}_{t-1}} \right) \tilde{e}_t \right) &= E \left(\left(\frac{\widetilde{D}_t}{\widetilde{D}_{t-1}}, \frac{\widetilde{C}_t}{\widetilde{C}_{t-1}} \right) (\varepsilon_t^D, \varepsilon_t^C) b_{DC} \right) \\ &= \Sigma_{DC} b_{DC} = E \left[\left(\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right) e_t \right], \end{aligned}$$

where the first equality follows from (50), $\left\{ \tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C, \tilde{\xi}_t, \tilde{\eta}_t \right\}_{t=-\infty}^{\infty}$ being serially uncorrelated, and $\tilde{\eta}_t$ being uncorrelated with $(\tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C)$.

For $i = 1$, the arguments used in the second paragraph of the proof of proposition 1 show that Restriction 2 also holds for $(\tilde{e}_t, \widetilde{D}_t/\widetilde{D}_{t-1}, \widetilde{C}_t/\widetilde{C}_{t-1})$. Having proved (51) for $i = 0$, Assumption 3 (ii) then gives (51) for $i = 1$.

Now it only remains to verify (51) for $i \leq -1$: since $\left\{ \tilde{\varepsilon}_t^D, \tilde{\varepsilon}_t^C, \tilde{\xi}_t, \tilde{\eta}_t \right\}_{t=-\infty}^{\infty}$ are serially uncorrelated, the left-hand side of (51) is zero; from Assumption 3 (iii) it follows that the right-hand side of (51) is also zero, which completes the proof.

A.8 Test Statistics from Section 6.2

We test moment restrictions of the form $E[e_t q_t] = 0$ in Proposition 1 for different instruments q using the test statistic

$$\widehat{Q}_T \equiv T \left(\frac{1}{T} \sum_{t=0}^T e_t q_t \right)' \widehat{S}_w^{-1} \left(\frac{1}{T} \sum_{t=0}^T e_t q_t \right) \rightarrow \chi_n^2,$$

where convergence is in distribution as $T \rightarrow \infty$, n denotes the dimension of q . Using the MA(1) property of e_t and independence of the shocks, we have $S_w = \sum_{i=-1}^{+1} E_t(q_{t+i} e_{t+i} q_t' e_t)$. This allows us to test Restrictions 1, 2, and 4. This test is an off-the-shelf application of a differences in differences test proposed by Arellano and Bond (1991) in the panel data context.⁴⁶

⁴⁶In our setting, differencing is useful to remove the random walk that is present under the agents' null hypothesis, whereas in the panel context it is used to remove fixed effects, but the test statistic is the same.

To test the inequality implied by Restriction 3 in Proposition 1, we estimate $E(e_t e_{t-1})$ and compare it with the estimates of b_{DC} and Σ_{DC} , which requires the joint distributions of these estimators. We obtain these from a GMM test. In particular, define the orthogonality conditions

$$\begin{aligned} g_1(\alpha, b_{DC}; x_t, x_{t-1}) &\equiv \left(\left[\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right] b_{DC} - e_{t-1} \right) e_t - \alpha \\ g_2(\alpha, b_{DC}; x_t, x_{t-1}) &\equiv \left[\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right]' \left(\left[\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right] b_{DC} - e_t \right) \end{aligned}$$

and let $g = (g_1, g_2)$. We then obtain from

$$E(g(\alpha, b_{DC}; x_t, x_{t-1})) = 0$$

three orthogonality conditions to estimate the three parameters (α, b_{DC}) . GMM sets $\hat{b}_{DC,T}$ to the OLS estimator of a regression of e_t on $\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}}$ and $\hat{\alpha}_T$ estimates $b'_{DC} \Sigma_{DC} b_{DC} - E(e_t e_{t-1})$. Therefore, Restriction 3 in Proposition 1 calls for testing the null hypothesis $H_0 : \alpha < 0$. Standard asymptotic distribution gives

$$\begin{aligned} \sqrt{T} \begin{bmatrix} \hat{\alpha}_T - \alpha \\ \hat{b}_{DC,T} - b_{DC} \end{bmatrix} &\rightarrow N(0, B^{-1} S_w (B')^{-1}) \text{ as } T \rightarrow \infty \\ B &= \begin{bmatrix} -1 & E \left(\left[\frac{D_t}{D_{t-1}}, \frac{C_t}{C_{t-1}} \right] e_t \right) \\ 0 & \Sigma_{DC} \end{bmatrix} \\ S_w &= \sum_{j=-\infty}^{\infty} E(g(b_{DC}, \alpha; x_t, x_{t-1}) g(b_{DC}, \alpha; x_{t-j}, x_{t-1-j})'). \end{aligned}$$

Substituting all moments by sample moments delivers the distribution for $\hat{\alpha}_T$.

B Online Appendix - Not for Publication

B.1 MSM Estimation and Testing

Here we provide an estimator for the covariance matrix of the statistics $\widehat{\Sigma}_{\mathcal{S},N}$, give some more details about the consistency of the estimator defined in (29), 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 $\widehat{\mathcal{S}}_N \equiv \mathcal{S}(\widehat{M}_N)$ 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 B.2. In the text we talked about “moments” as describing all statistics to be matched in (28). 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 $\widehat{\mathcal{S}}_N$. Since this deviates from a 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 $\widehat{\mathcal{S}}_N$ can be all or a subset of the statistics in (28). The reason why we sometimes consider a subset will be discussed in detail after result for “Asymptotic Distribution” 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 the work of Duffie and Singleton (1993) (hereafter 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 are 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 $\widetilde{\mathcal{S}}(\theta) \equiv \mathcal{S}(M(\theta))$ 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

$$S_w \equiv \sum_{j=-\infty}^{\infty} M_0^j. \quad (52)$$

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

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

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

$$\widehat{\Sigma}_{\mathcal{S},N} \rightarrow \Sigma_{\mathcal{S}} \equiv \frac{\partial \mathcal{S}(M_0)}{\partial M'} S_w \frac{\partial \mathcal{S}(M_0)'}{\partial M} \quad (53)$$

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

$$E \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M_0) \right] \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M_0) \right]' \rightarrow \Sigma_{\mathcal{S}} \quad (54)$$

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 (53).

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

$$\mathcal{S}(M_N) - \mathcal{S}(M_0) = \frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} [M_N - M_0] \quad (55)$$

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 (54). ■

Conditions *a)*, *c)* are standard minimal assumptions used in time series asymptotic results, and condition *b)* is clearly satisfied in our application (see the expression for \mathcal{S} in section B.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 the following.

Consistency *Let $\widehat{\theta}_N$ be the estimator defined in (29), 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 ensure 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 A.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 $B_0 \Sigma_S^{-1} B_0'$ is invertible. Then*

$$\sqrt{N} \left[\widehat{\theta}_N - \theta_0 \right] \rightarrow \mathcal{N}(0, (\mathcal{B}_0 \Sigma_S^{-1} \mathcal{B}_0')^{-1}) \quad (56)$$

$$\sqrt{N} \left[\widehat{\mathcal{S}}_N - \mathcal{S}(M(\widehat{\theta})) \right] \rightarrow \mathcal{N}(0, \Sigma_S - \mathcal{B}_0' (\mathcal{B}_0 \Sigma_S^{-1} \mathcal{B}_0')^{-1} \mathcal{B}_0) \quad (57)$$

$$\widehat{W}_N \rightarrow \chi_{s-n}^2 \quad (58)$$

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 (55) imply

$$\sqrt{N} [\mathcal{S}(M_N) - \mathcal{S}(M_0)] = \frac{\partial \mathcal{S}(\widetilde{M}_N)}{\partial M'} \sqrt{N} [M_N - M_0] \rightarrow \mathcal{N}(0, \Sigma_S) \quad (59)$$

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

$$\begin{aligned} \mathcal{S}(M(\widehat{\theta}_N)) - \mathcal{S}(M(\theta_0)) &= \mathcal{B}(\widetilde{\theta}, \widetilde{M}) \left[\widehat{\theta}_N - \theta_0 \right] \\ \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S},N}^{-1} \left[\mathcal{S}(M(\widehat{\theta})) - \mathcal{S}(M(\theta_0)) \right] &= \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S},N}^{-1} \mathcal{B}(\widetilde{\theta}, \widetilde{M})' \left[\widehat{\theta} - \theta_0 \right] \\ \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} \left[\mathcal{S}(M_N) - \mathcal{S}(M(\theta_0)) \right] &= \mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} \mathcal{B}(\widetilde{\theta}, \widetilde{M})' \left[\widehat{\theta} - \theta_0 \right] \end{aligned} \quad (60)$$

where the last equality follows because $\mathcal{B}(\widehat{\theta}_N, \widehat{M}_N) \widehat{\Sigma}_{\mathcal{S}}^{-1} \left[\mathcal{S}(M_N) - \mathcal{S}(\widehat{\theta}_N) \right] = 0$ at the maximum of (29). This implies (56).

To obtain (57) we use the mean value theorem and (60) to conclude

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

This gives (57).

(58) follows from (57) and that $\left(\Sigma_S - \mathcal{B}'_0 \left[\mathcal{B}_0 \Sigma_S^{-1} \mathcal{B}'_0 \right]^{-1} \mathcal{B}_0 \right)$ is an idempotent matrix. ■

We mentioned in the main text and earlier in this appendix that we drop some statistics from \widehat{S}_N when $\widehat{\Sigma}_{S,N}$ is nearly singular. The reason is that, as stated above, we need invertibility of Σ_S for both consistency and asymptotic distribution. In practice, a nearly singular $\widehat{\Sigma}_{S,N}$ 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 cannot 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 $\widehat{\Sigma}_{S,N}$. 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.⁴⁷ As we explain in our discussion of 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

⁴⁷An alternative to dropping one moment would be to match the linear combinations corresponding to the principal components of \widehat{S}_w . Although 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, whereas each individual moment does have a clear economic interpretation.

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 $\tilde{\mathcal{S}}(\theta)$ equal to

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

for large K . In other words, $\tilde{\mathcal{S}}(\theta)$ 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 the 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.

B.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 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

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

where $r_{t-20}^{s,20}$ denotes the excess stock return over 20 quarters, which can be computed using from \mathbf{y}_t using $(PD_t, D_t/D_{t-1}, r_t^b, \dots, PD_{t-19}, D_{t-19}/D_{t-20}, r_{t-19}^b, PD_{t-20})$.⁴⁸

The ten statistics we consider can be expressed as functions 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, or more precisely,

⁴⁸The excess return is defined as $r_{t-20}^{s,20} \equiv \left(\prod_{j=1}^{20} (1 + r_{t-20+j}^s) \right) - \left(\prod_{j=1}^{20} (1 + r_{t-20+j}^b) \right)$, where $1 + r_t^s = \frac{P_t + D_t}{P_{t-1}} = \frac{PD_{t+1}}{PD_{t-1}} \frac{D_t}{D_{t-1}}$ is the one-period quarterly stock return and a function of \mathbf{y}_t .

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

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

B.3 Derivatives of the Statistic Function

This appendix gives explicit expressions for $\partial\mathcal{S}/\partial M'$ using the statistic function stated in appendix B.2. Straightforward but tedious algebra shows

$$\begin{aligned} \frac{\partial\mathcal{S}_i}{\partial M_j} &= 1 && \text{for } (i, j) = (1, 1), (2, 2), (8, 9), (9, 10) \\ \frac{\partial\mathcal{S}_i}{\partial M_j} &= \frac{1}{2\mathcal{S}_i(M)} && \text{for } (i, j) = (3, 3), (4, 4), (10, 11) \\ \frac{\partial\mathcal{S}_i}{\partial M_j} &= \frac{-M_j}{\mathcal{S}_i(M)} && \text{for } (i, j) = (3, 1), (4, 2), (10, 10) \\ \frac{\partial\mathcal{S}_5}{\partial M_2} &= \frac{2M_2(M_5 - M_4)}{(M_4 - M_2^2)^2}, && \frac{\partial\mathcal{S}_5}{\partial M_5} = \frac{1}{M_4 - M_2^2}, && \frac{\partial\mathcal{S}_5}{\partial M_4} = -\frac{M_5 - M_2^2}{(M_4 - M_2^2)^2} \\ \frac{\partial\mathcal{S}_6}{\partial M_j} &= \frac{\partial c_2^5(M)}{\partial M_j} && \text{for } i = 2, 4, 6, 8 \\ \frac{\partial\mathcal{S}_7}{\partial M_j} &= \frac{[M_6, M_8] \frac{\partial c^5(M)}{\partial M_j}}{M_7 - M_6^2} && \text{for } j = 2, 4 \\ \frac{\partial\mathcal{S}_7}{\partial M_6} &= \frac{\left[c_1^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_6} \right] (M_7 - M_6^2) - 2M_6 [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2} \\ \frac{\partial\mathcal{S}_7}{\partial M_7} &= \frac{M_6^2 - [M_6, M_8] c^5(M)}{(M_7 - M_6^2)^2} \\ \frac{\partial\mathcal{S}_7}{\partial M_8} &= \frac{c_2^5(M) + [M_6, M_8] \frac{\partial c^5(M)}{\partial M_8}}{M_7 - M_6^2} \end{aligned}$$

Using the formula for the inverse of a 2x2 matrix

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

we have

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

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