

University of Mannheim / Department of Economics

Working Paper Series

Stock Price Booms and Expected Capital Gains

Klaus Adam Johannes Beutel

Albert Marcet

Working Paper 14-12

March 2014

Stock Price Booms and Expected Capital Gains¹

Klaus Adam
University of Mannheim & CEPR

Johannes Beutel
University of Mannheim

Albert Marcet
Institut d'Anàlisi Econòmica (CSIC), ICREA, UAB, MOVE,
Barcelona GSE & CEPR

March 3, 2014

¹This paper replaces an earlier paper titled 'Booms and Busts in Asset Prices' by Adam and Marcet, which appeared 2010 as Bank of Japan - IMES Discussion Paper No. 2010-E-2. Thanks go to Fernando Alvarez, Chryssi Giannitsarou, Vivien Lewis, Morten Ravn, Ricardo Reis, Mike Woodford, conference participants at the Banque de France and Chicago Fed Conference on Asset Price Bubbles, ESSIM 2010 in Tarragona, MONFISPOL network, and seminar participants at Columbia University, University College London, London School of Economics and London Business School for helpful comments and suggestions. Research assistance from Jeanine Baumert, Oriol Carreras, Dimitry Matveev and Sebastian Merkel is greatly appreciated. Klaus Adam thanks the Bank of Japan for the hospitality offered during early stages of this project. Albert Marcet acknowledges support from Programa de Excelencia del Banco de España, Plan Nacional (Ministry of Education), SGR (Generalitat de Catalunya). Both authors acknowledge support from the European Research Council under the EU 7th Framework Programme (FP/2007-2013), Starting Grant Agreement No. 284262 (Adam) and Advanced Grant Agreement No. 324048 (Marcet). Klaus Adam and Albert Marcet thank Fondation Banque de France for its support on this project. All errors remain ours.

Abstract

The booms and busts in U.S. stock prices over the post-war period can to a large extent be explained by fluctuations in investors' subjective capital gains expectations. Survey measures of these expectations display excessive optimism at market peaks and excessive pessimism at market troughs. Formally incorporating subjective price beliefs into an otherwise standard asset pricing model with utility maximizing investors, we show how subjective belief dynamics can temporarily delink stock prices from their fundamental value and give rise to asset price booms that ultimately result in a price bust. The model successfully replicates (1) the volatility of stock prices and (2) the positive correlation between the price dividend ratio and expected returns observed in survey data. We show that models imposing objective or 'rational' price expectations cannot simultaneously account for both facts. Our findings imply that large parts of U.S. stock price fluctuations are not due to standard fundamental forces, instead result from self-reinforcing belief dynamics triggered by these fundamentals.

JEL Class. No.: G12, D84

‘Bull-markets are born on pessimism, grow on skepticism, mature on optimism and die on euphoria’

Sir John Templeton, Founder of Templeton Mutual Funds

1 Motivation

Following the recent boom and bust cycles in a number of asset markets around the globe, there exists renewed interest in understanding better the forces contributing to the emergence of such drastic asset price movements. This paper argues that movements in investor optimism and pessimism, as measured by the movements in investors’ subjective expectations about future capital gains, are a crucial ingredient for understanding these fluctuations.

We present an asset pricing model that incorporates endogenous belief dynamics about expected capital gains. The model gives rise to sustained stock price booms and busts and is consistent with the behavior of investors’ capital gains expectations, as measured by survey data. The model suggests that more than half of the variance of the price dividend ratio in U.S. post-WWII data is due to movements in subjective expectations.

The standard approach in the consumption-based asset pricing literature consists of assuming that stock price fluctuations are fully efficient. Campbell and Cochrane (1999) and Bansal and Yaron (2004), for example, present models in which stock price fluctuations reflect the interaction of investor preferences and stochastic driving forces in a setting with optimizing investors who hold rational expectations.

The empirical evidence we present casts considerable doubt on the prevailing view that stock price fluctuations are efficient. Specifically, we show that the RE hypothesis gives rise to an important counterfactual prediction for the behavior of investors’ expectations. This counterfactual prediction is a model-independent implication of the RE hypothesis, but - as we explain below - key for understanding stock price volatility and its efficiency properties.

As previously noted by Fama and French (1988), the empirical behavior of asset prices implies that rational return expectations correlate *negatively* with the price dividend (PD) ratio.¹ Somewhat counter-intuitively, the RE hypothesis thus predicts that investors have been particularly pessimistic about future stock returns in the early part of the year 2000, when the tech stock boom and the PD ratio of the S&P500 reached its

¹The RE hypothesis implies also a negative correlation between the PD ratio and expected capital gains. Since most variation in returns is due to variation in capital gains, we tend to use both terms interchangeably.

all-time maximum. As we document, the available survey evidence implies precisely the opposite: all quantitative survey measures of investors' return (or capital gain) expectations available for the U.S. economy, unambiguously and unanimously correlate *positively* with the PD ratio; and perhaps not surprisingly, return expectations reached a temporary maximum rather than a minimum in the early part of the year 2000, i.e., precisely at the peak of the tech stock boom, a fact previously shown in Vissing-Jorgensen (2003). Using a formal test we confirm that the survey data is at odds with the RE hypothesis at any conventional significance level because survey expectations and RE covary differently with the PD ratio.

The positive comovement of stock prices and survey expectations suggests that price fluctuations are amplified by overly optimistic beliefs at market peaks and by overly pessimistic beliefs at market troughs. Furthermore, it suggests that investors' capital gains expectations are influenced - at least partly - by the capital gains observed in the past, in line with evidence presented by Malmendier and Nagel (2011). Indeed, a simple adaptive updating equation captures the time series behavior of the survey data and its correlation with the PD ratio very well.

Taken together, these observations motivate the construction of an asset pricing model in which investors hold subjective beliefs about the capital gains from stock investments.² We incorporate such beliefs into a Lucas (1978) asset pricing model, assuming that agents are uncertain about the capital gains process but invest optimally given their beliefs and update beliefs according to Bayes' law.

With this modification, the Lucas model with standard time separable preferences and standard stochastic driving processes becomes quantitatively consistent with the observed volatility of stock prices and the positive correlation between the PD ratio and subjective return expectations. Considering the same model under RE, produces - amongst other things - too little price volatility and the wrong sign for the correlation between the PD ratio and expected returns.

The strong improvement in the model's empirical performance arises because agents' attempts to improve their knowledge about price behavior can temporarily delink asset prices from their fundamental (RE) value and give rise to belief-driven boom and bust cycles in stock prices. This occurs because with imperfect information about the price process, optimal behavior dictates that agents use past capital gains observations to learn about the stochastic process governing the behavior of capital gains; this generates

²As is explained in Adam and Marcet (2011), the presence of subjective price beliefs reflects a lack of common knowledge about agents' beliefs and preferences.

a feedback between capital gain expectations and realized capital gains.

Suppose, in line with the empirical evidence, that agents become more optimistic about future capital gains whenever they are positively surprised by past capital gains.³ A positive surprise then increases asset prices further, whenever increased optimism leads to an increase in investors' asset demand. If this effect is sufficiently strong, then positive surprises trigger further positive surprises and thus further price increases. As we show analytically, stock prices in our model do increase with capital gain optimism whenever the substitution effect of increased optimism dominates the wealth effect of such belief changes. Asset prices in the model then display sustained price booms, similar to those observed in the data.

After a sequence of sustained increases, countervailing forces come into play that endogenously dampen the upward price momentum, eventually halt it and cause a reversal. Specifically, in a situation where increased optimism about capital gains has led to a stock price boom, stock prices make up for a larger share of agents' total wealth.⁴ As we show analytically, this eventually causes the wealth effect to become as strong as (or even stronger than) the substitution effect.⁵ Increases in optimism then cease to cause further increases in stock demand and thus stock prices, so that investors' capital gains expectations turn out to be too optimistic relative to the realized outcomes. This induces downward revision in beliefs, which gives rise to negative price momentum and an asset price bust.

The previous arguments show how belief dynamics can temporarily delink asset prices from their fundamental value. Clearly, these price dynamics are inefficient as they are not justified by innovations to preferences or other fundamentals.

We obtain these results even though we depart from the standard paradigm in a minimal way only. Specifically, we assume that investors are internally rational (IR) in the sense of Adam and Marcet (2011). This implies that all investors hold an internally consistent system of beliefs about variables that are exogenous to their decision problem and choose investment and consumption optimally. Although agents' beliefs do not fully capture the actual behavior of prices in equilibrium, in line with the survey evidence, agents' beliefs are broadly plausible given the behavior of equilibrium prices and the behavior of prices in the data. In particular, agents believe the average growth rate of

³Such positive surprises may be triggered by fundamental shocks, e.g., a high value for realized dividend growth.

⁴This occurs because stock prices are high, but also because agents discount other income streams, e.g., wage income, at a higher rate.

⁵With CRRA utility, this happens whenever the coefficient of relative risk aversion is larger than one.

stock prices to slowly drift over time, which is consistent with the presence of prolonged periods of price booms followed by price busts.

The current paper shows how the framework of internal rationality allows studying learning about market behavior in a model of intertemporal decision making, while avoiding some of the pitfalls of the adaptive learning literature, where agents' belief updating equations and choices are often not derived from individual maximization. We thus show how explicit microfoundations can guide modelling choices in settings featuring subjective beliefs about market outcomes, as is the case in settings imposing RE.

The remainder of the paper is structured as follows. Section 3 documents that there is a strong positive correlation between the PD ratio and survey measures of investors' return and capital gain expectations and that this is incompatible with the RE hypothesis. It then documents that from a purely statistical standpoint approximately two thirds of the variation in the PD ratio of S&P500 can potentially be accounted for by variations in expected capital gains. Section 4 presents our asset pricing model with subjective beliefs. For benchmark purposes, section 5 determines the RE equilibrium. Section 6 introduces a specific model for subjective price beliefs; it does so by relaxing agents' prior beliefs about price behavior relative to the RE equilibrium beliefs. This section also derives the resulting Bayesian updating equations characterizing belief dynamics over time, involving learning about the permanent component of stock price growth. After imposing market clearing in section 7, we present closed form solutions for the PD ratio in section 8 in the special case of vanishing uncertainty. We then explain how the interaction between belief updating dynamics and price outcomes can endogenously generate boom and bust dynamics in asset prices. Section 9 considers the model with empirically plausible amounts of uncertainty and documents its ability to replicate the time series behavior of the postwar US PD ratio and of the survey data. Section 10 documents that the model under learning replicates important asset pricing moments much better than under RE. A conclusion briefly summarizes and presents an outlook on future research avenues. Technical material and proofs can be found in the appendix.

2 Related Literature

The literature on adaptive learning previously studied the role of deviations from RE in asset pricing models. Work by Bullard and Duffy (2001) and Brock and Hommes (1998), for example, explores learning about price forecasting and shows that learning

dynamics can converge to complicated attractors that increase asset return volatility, if the RE equilibrium is unstable under learning dynamics.⁶ Lansing (2010) shows how near-rational bubbles can arise under learning dynamics when agents forecast a composite variable involving future price and dividends. Branch and Evans (2011) present a model where agents learn about risk and return and show how it gives rise to bubbles and crashes. Boswijk, Hommes and Manzan (2007) estimate a model with fundamentalist and chartist traders whose relative shares evolve according to an evolutionary performance criterion, showing that the model can generate a run-up in asset prices and subsequent mean-reversion to fundamental values. DeLong et al. (1990) show how the pricing effects of positive feedback trading survives or even get amplified by the introduction of rational speculators. Timmermann (1993, 1996) explores learning about dividend behavior but finds overall limited pricing implications. Cogley and Sargent (2008) have studied a model of robustness, where agents learn about fundamentals and behave according to max-min utility.

We contribute to this literature in three ways. First, we compare the implications of our model more closely to the data, both in terms of matching the time series of asset prices and survey data, as well as in terms of matching asset pricing moments.

Second, we specify proper microfoundations for agents' infinite horizon decision problem with subjective beliefs and derive agents' optimal consumption plans and belief updating equations from this problem. The subjective consumption plans are then used to price the stock market. Earlier work on infinite horizon models in the adaptive learning literature typically falls short of specifying proper optimization problems. As explained in section 2 in Adam and Marcet (2011), this leads to arbitrariness in the modeling of agents' behavior, which can affect model predictions and the resulting conclusions. Important progress has been made in recent work by Eusepi and Preston (2011, 2013), who derive choices from properly formulated optimization problems featuring subjective beliefs. Here we go a step further by jointly deriving the optimal decisions and the belief updating rules from the utility maximization problem, instead of making appeal to the anticipated utility framework in Kreps (1998), which implies that future belief revisions are abstracted from when deriving decisions.

Third, we are able to derive our main results using a closed-form solution. This provides clearer insights into the economic mechanisms driving the asset pricing results. We also discuss issues of existence and uniqueness of optimal plans in models with subjec-

⁶Stability under learning dynamics is defined in Marcet and Sargent (1989).

tive beliefs and conditions under which the optimal plan has a recursive representation. Furthermore, we explain why rational agents can hold separate subjective beliefs about prices and fundamentals.

Fuster, Herbert and Laibson (2011) present an asset pricing model where fundamentals exhibit momentum in the short-run and partial mean reversion in the long-run and where agents underparameterize the fundamental process, thereby missing the long-run mean reversion. They show how such a model can give rise to pro-cyclical excess optimism as in the present paper. Fundamentals in our model display neither momentum nor mean reversion, excess optimism and pessimism arise instead endogenously from the interaction between price outcomes and expectations.

Hassan and Mertens (2010) present a stock market model where investors deviate from fully rational behavior, as agents make small common errors in formulating expectations. They show how the market amplifies these errors and how this can have large welfare consequences by shifting investment away from domestic production opportunities into foreign safe bonds. The present model does not consider effects on output and welfare, instead derives empirical implications for stock price volatility and the behavior of expectations in a setting with fully optimal behavior and Bayesian updating, given imperfect knowledge of the economy.

Adam Marcet and Nicolini (2013) quantitatively evaluate the ability of models of learning to explain asset price volatility. To be able to formally estimate the model using the method of simulated moments, they rely on a number of short-cuts. In particular, they assume dividends to be a negligible part of total income, so that consumption equals exogenous labor income. As a result, the stochastic discount factor is exogenous. While being analytically convenient, this prevents the emergence of the wealth effects referred to in the introduction, requiring asset price booms to be stopped by exogenously imposing an upper bound on agents' beliefs.⁷ Clearly, this prevents a discussion of asset price booms and their end. They also do not discuss survey evidence.

The experimental and behavioral literature provides further evidence supporting the presence of subjective price beliefs. Asparouhova, Bossaerts, Roy and Zame (2013), for example, implement the Lucas asset pricing model in the experimental laboratory and document that there is excess volatility in prices that is unaccounted for by the rational expectations equilibrium and that likely arises from participants' expectations

⁷The performance of the model in terms of quantitatively replicating asset pricing moments is, however, robust to the precise value chosen for this upper bound, because the bound is binding only rarely along the equilibrium path.

about future prices. Furthermore, the type of learning employed in the present model is in line with evidence presented in Malmendier and Nagel (2011) who show that experienced returns affect beliefs about future asset returns.⁸

3 Stock Prices & Stock Price Expectations: Facts

This section explains how two important and widely accepted asset pricing facts imply a counterfactual behavior for the behavior of stock price expectations, whenever one imposes that agents hold rational price expectations. We present the evidence informally in section 3.1 and derive a formal statistical test in section 3.2. The test shows that the RE hypothesis is inconsistent with the survey data because it is incompatible with the way survey data covary with the PD ratio. Section 3.3 illustrates how simple adaptive prediction of prices, in line with Malmendier and Nagel (2011, 2013), quantitatively captures the relationship between survey expectations and the PD ratio. It also shows how, in a purely statistical sense, variations in expected capital gains can potentially account for up to two thirds of the variation of the U.S. PD ratio over the postwar period.

3.1 Survey Expectations and the PD Ratio

This section explains how the presence of boom and bust dynamics in stock prices, together with the unpredictability of dividend growth imply that investors' expectations about future stock returns should correlate *negatively* with the PD ratio whenever investors' hold rational expectations (RE) about future stock prices. It then documents that survey measures of investors' return expectations correlate instead *positively* with the PD ratio; this positive correlation is statistically significant and robust to considering different survey measures and data sources.

As is well known stock prices experience substantial price booms and price busts. Figure 1 illustrates this behavior for the post-WWII period for the United States, using the quarterly price dividend ratio (PD) of the S&P 500 index.⁹ The PD ratio displays persistent run-ups and reversals, with the largest one occurring around the year 2000. This shows that price growth can persistently outstrip dividend growth over a number of

⁸Nagel and Greenwood (2009) show that - in line with this hypothesis - young mutual fund managers displayed trend chasing behavior over the tech stock boom and bust around the year 2000.

⁹Quarterly dividend payments have been deseasonalized in a standard way by averaging them across the current and preceding 3 quarters. See appendix A.1 for details about the data used in this section.

periods, but that the situation eventually reverses. In fact, the quarterly autocorrelation of the PD ratio equals 0.98. Similar run-ups and reversals can be documented for other mature stock markets, e.g., for the European or Japanese markets.

Equally well-known is the fact that the growth rate of dividends is largely unpredictable, e.g., Campbell (2003). It is especially hard to predict using the PD ratio. The R^2 values of an in-sample predictive regression of cumulative dividend growth 1, 5 or 10 years ahead on a constant and the log PD ratio are rather small and amount to 0.03, 0.04, and 0.07, respectively, for the U.S. post-war data.¹⁰

Taken together the previous two facts imply that under RE one would expect that the PD ratio *negatively* predicts future stock market returns. To see this, let the asset return R_{t+1} be defined as

$$R_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t} = \frac{\frac{P_{t+1}}{D_{t+1}} + 1}{\frac{P_t}{D_t}} \frac{D_{t+1}}{D_t},$$

where P denotes the stock price and D dividends. Given a high value of P_t/D_t , we have - due to the mean reverting behavior of the PD ratio - that $P_{t+1}/D_{t+1} < P_t/D_t$ on average. Since D_{t+1}/D_t is unpredictable, it follows that a high PD ratio negatively predicts future returns.¹¹ A symmetric argument holds if P_t/D_t is low.

In the setup previously described expectations about future stock returns should covary negatively with the PD ratio if investors hold RE. In particular, rational expectations about stock returns should be very low at the height of the tech stock boom in the year 2000 when the PD ratio reached its historical peak.

Survey evidence on investors' return expectations displays instead a strong *positive* correlation between investors' expected returns and the PD ratio.¹² Figure 2 depicts this for our preferred survey, the UBS Gallup Survey, which is based on a representative sample of approximately 1.000 U.S. investors that own at least 10.000 US\$ in financial wealth.¹³ Figure 2 graphs the US PD ratio (the black line) together with measures of

¹⁰We use logPD as a regressor, in line with Campbell (2003). The R^2 values are unchanged when using the level of the PD ratio instead.

¹¹There may exist - at least in theory - other predictors of future returns which correlate negatively with the PD ratio and that overturn the negative relationship between PD ratio and expected stock returns emerging from the forces described above. We take these formally into account in our statistical test in section 3.2, where we show that the the PD ratio negatively predicts future returns in the data. The latter fact is related to the well-known predictability of *excess* stock returns using the PD ratio and the unpredictability of future risk free rates, see Campbell (2003) for example.

¹²The positive co-movement between return expectations and stock market valuation has previously been documented in somewhat different form by Vissing-Jorgensen (2003) and Bacchetta, Mertens, and Wincoop (2009).

¹³About 40% of respondents own more than 100.000 US\$ in financial wealth. As documented below,

the cross-sectional average of investors' one-year ahead expected real return.¹⁴ Return expectations are expressed in terms of quarterly real growth rates and the figure depicts two expectations measures: investors' expectations about the one year ahead stock market return, as well as their expectations about the one year ahead returns on their own stock portfolio. These measures behave very similarly over the period for which both series are available, but the latter is reported for a longer time period, so that we focus on it as our baseline. Figure 2 reveals that there is a strong positive correlation between the PD ratio and expected returns. The correlation between the expected own portfolio returns and the PD ratio is +0.70 and even higher for the expected stock returns (+0.82). Moreover, investors' return expectations were highest at the beginning of the year 2000, which is precisely the year the PD ratio reached its peak during the tech stock boom. At that time, investors expected annualized real returns of around 13% from stock investments. Conversely, investors were most pessimistic in the year 2003 when the PD ratio reached its bottom, expecting then annualized real returns of below 4%.

Table 1 shows that the strong positive correlation evident from figure 2 is robust to a number of alternative approaches for extracting expectations from the UBS survey, such as using the median instead of the mean expectation, when using inflation expectations from the Michigan survey to obtain real return expectations, when considering plain nominal returns instead of real returns, or when restricting attention to investors with more than 100.000 US\$ in financial wealth. The numbers reported in brackets in table 1 (and in subsequent tables) are autocorrelation robust p-values for the hypothesis that the correlation is smaller or equal to zero.¹⁵ The p-values for this hypothesis are all below the 5% significance level and in many cases below the 1% level.

A positive and statistically significant correlation is equally obtained when considering other survey data. Table 2 reports the correlations between the PD ratio and the stock price growth expectations from Bob Shiller's Individual Investors' Survey.¹⁶ The table shows that price growth expectations are also strongly positively correlated with the PD ratio, suggesting that the variation in expected returns observed in the UBS survey is

this subgroup does not behave differently.

¹⁴To be consistent with the asset pricing model presented in later sections we report expectations of real returns. The nominal return expectations from the survey have been transformed into real returns using inflation forecasts from the Survey of Professional Forecasters. Results are robust to using other approaches, see the subsequent discussion.

¹⁵The sampling width is four quarters, as is standard for quarterly data, and the test allows for contemporaneous correlation, as well as for cross-correlations at leads and lags. The p-values are computed using the result in Roy (1989)

¹⁶Shiller's price growth data refers to the Dow Jones Index. The table thus reports the correlation of the survey measure with the PD ratio of the Dow Jones.

due to variations in expected capital gains. Table 2 also shows that correlations seem to become stronger for longer prediction horizons.

Table 3 reports the correlations for the stock return expectations reported in the Chief Financial Officer (CFO) survey which surveys chief financial officers from large U.S. corporations. Again, one finds a strong positive correlation; it is significant at the 1% level in all cases.

Table 4 reports the correlations between the PD ratio and the realized real returns (or capital gains) in the data, using the same sample periods as are available for the surveys considered in tables 1 to 3, respectively. The point estimate for the correlation is negative in all cases, although the correlations fall short of being significant the 5% level due to the short sample length for which the survey data is available. Nevertheless, table 4 suggests that investors' expectations are most likely incompatible with RE. The next section investigates this issue more formally.

3.2 Survey Expectations versus Rational Expectations

Using a formal econometric test, this section shows that the RE assumption is incompatible with the behavior of survey expectations because RE and survey expectations covary differently with the PD ratio. The subsequent section presents a simple model of learning about prices that correctly captures the covariance between survey data and the PD ratio.

Let $E_t^{\mathcal{P}}$ denote the agents' subjective (and potentially less-than-fully-rational) expectations operator based on information up to time t , and $R_{t,t+N}$ the cumulative stock returns between period t and $t + N$.¹⁷ When these expectations are observable, say via survey data, one can write the *regression* equation

$$E_t^{\mathcal{P}} R_{t,t+N} = a^N + c^N \frac{P_t}{D_t} + u_t^N, \quad (1)$$

where

$$E(x_t u_t^N) = 0 \quad (2)$$

for $x_t' = (1, P_t/D_t)$.^{18,19} Importantly, the operator E in the orthogonality condition (2)

¹⁷In line with the survey data, we consider plain stock returns and not excess returns, which are the focus of traditional excess return predictability regressions.

¹⁸This regression is well-defined, as long as agents' expectations $E_t^{\mathcal{P}} R_{t,t+N}$ and the PD ratio P_t/D_t are stationary and have bounded second moments.

¹⁹Since the Shiller survey reports expectations about capital gains instead of returns, we interpret the variable $R_{t,t+N}$ as the real growth rate of stock prices between periods t and $t + N$ when using the Shiller

denotes the objective expectation for the *true* data generating process, independently of how agents' expectations are formed. The regression residual u_t^N captures the variation in agents' expectations that cannot be linearly attributed to the price-dividend ratio and summarizes all other information that agents believe to be useful in predicting $R_{t,t+N}$.²⁰ We let \hat{c}^N denote the OLS estimator of c^N in equation (1).²¹

In the special case with rational expectations ($E_t^P = E_t$) equation (1) implies

$$R_{t,t+N} = a^N + c^N \frac{P_t}{D_t} + u_t^N + \varepsilon_t^N \quad (3)$$

where $\varepsilon_t^N = R_{t,t+N} - E_t R_{t,t+N}$ is the prediction error arising from the *true* data-generating process and thus orthogonal to all past observations dated t or earlier. It satisfies

$$E [x_t (u_t^N + \varepsilon_t^N)] = 0, \quad (4)$$

so that an estimate of c^N that is consistent with the RE assumption can be derived by estimating (3) with OLS. We let $\hat{\hat{c}}^N$ denote this estimate.

The correlations reported in tables 1-4 imply - by construction - that $\hat{c}^N > 0$ and $\hat{\hat{c}}^N < 0$. The regression estimates are nevertheless useful because under the RE hypothesis \hat{c}^N and $\hat{\hat{c}}^N$ are consistent estimates of the *same* parameter c^N , allowing to formally test the RE hypothesis, i.e., $H_0 : \hat{c}^N = \hat{\hat{c}}^N$.^{22,23} Clearly, if the asset price and survey data were generated by a rational expectations model, say the models of Campbell and Cochrane (1999) or Bansal and Yaron (2004), this test would be accepted.

Table 5 reports the p-values for $H_0 : \hat{c}^N = \hat{\hat{c}}^N$ using the survey data from the previous section.²⁴ The point estimates all satisfy $\hat{c} > 0$ and $\hat{\hat{c}} < 0$, and the difference between the

survey.

²⁰The residual u_t^N is likely to be correlated with current and past observables (other than the PD ratio) and thus serially correlated.

²¹Sometimes the validity of surveys is questioned because survey responses are noisy data. While this point applies to many other macroeconomic variables, it is the case that under the standard assumption that the measurement error is uncorrelated with the regressor, the estimators and standard errors reported below remain valid.

²²Under the RE hypothesis, the correlations in tables 1-3 are not equal to the corresponding correlations reported in table 4, albeit both should have the same sign.

²³To obtain p-values for $H_0 : \hat{c}^N = \hat{\hat{c}}^N$, we stack up equations (1) and (3), create a SUR system of equations to find the *joint* distribution of \hat{c}^N and $\hat{\hat{c}}^N$ and build a t-test for H_0 . We use serial-correlation and heteroskedasticity robust asymptotic covariance matrix of the estimators, using 4 lags, results are robust to increasing the lag length to up to 12 lags. For each considered survey we use data on actual returns (or price growth) for the same time period for which survey data is available when computing the p-values. Further details of the test are described in appendix A.2.

²⁴Tests for 'own portfolio' expectations are not shown because we do not observe agents' returns on their own portfolio.

two estimates is highly statistically significant: the null hypothesis is always rejected at the 5% level and in all but two cases also at the 1% level. This formally shows that the RE hypothesis fails to explain the survey data because survey expectations and actual data covary differently with the PD ratio.

3.3 How Models of Learning May Help

This section illustrates that a simple ‘adaptive’ approach to forecasting stock prices is a promising alternative to explain the joint behavior of survey expectations and stock price data.

Figure 2 shows that the peaks and troughs of the PD ratio are located very closely to the peaks and troughs of investors’ return expectations. This suggests that agents become optimistic about future capital gains whenever they have observed capital gains in the past. Such behavior can be captured by models where agents expectations are influenced by past experience prompting us to assume for a moment that agents’ subjective conditional capital gain expectations $\tilde{E}_t [P_{t+1}/P_t]$ evolve according to the following adaptive prediction model

$$\tilde{E}_t [P_{t+1}/P_t] = \tilde{E}_{t-1} [P_t/P_{t-1}] + g \left(\frac{P_t}{P_{t-1}} - \tilde{E}_{t-1} [P_t/P_{t-1}] \right), \quad (5)$$

where $g > 0$ indicates how strongly capital gain expectations are updated in the direction of the forecast error. While equation (5) may appear ad-hoc, we show in section 6 how a very similar equation can be derived from Bayesian belief updating in a setting where agents estimate the persistent component of price growth from the data.

One can use equation (5) and feed into it the historical price growth data of the S&P 500 over the postwar period. Together with an assumption about capital gain expectations at the start of the sample this will deliver a time series of implied capital gain expectations $\tilde{E}_t [P_{t+1}/P_t]$ that can be compared to the expectations from the UBS survey.²⁵ Figure 3 reports the outcome of this procedure when assuming initial beliefs in Q1:1946 to be equal to -1.11% per quarter and $g = 0.02515$, which minimizes the sum of squared deviations from the survey evidence.²⁶ Figure 3 shows that the adaptive model

²⁵We transform the UBS survey measures of return expectations into a measure of price growth expectations using the identity $R_{t+1} = \frac{P_{t+1}}{P_t} + \frac{D_{t+1}}{P_t} = \frac{P_{t+1}}{P_t} + \beta^D \frac{D_t}{P_t}$ where β^D denotes the expected quarterly growth rate of dividends that we set equal to the sample average of dividend growth over Q1:1946-Q1:2012, i.e. $\beta^D = 1.0048$. Results regarding implied price growth are very robust towards changing β^D to alternative empirically plausible values.

²⁶The figure reports growth expectations in terms of quarterly real growth rates.

captures the behavior of UBS expectations extremely well: the correlation between the two series is equal to +0.89.

A similarly strong positive relationship between the PD ratio and the capital gains expectations implied by equation (50) exists over the entire postwar period, as figure 4 documents. The figure plots the joint distribution of the capital gains expectations (as implied by equation (50)) and the PD ratio in the data. When regressing the PD ratio on a constant and the expectations of the adaptive prediction model, one obtains an R^2 coefficient of 0.55; using also the square of the expectations, the R^2 rises further to 0.67. Variations in expected capital gains can thus account - in a purely statistical sense - for up to two thirds of the variability in the postwar PD ratio.²⁷

The previous findings suggest that an asset pricing model consistent with equation (5), which additionally predicts a positive relationship between the PD ratio and subjective expectations about future capital gains, has a good chance of replicating the observed positive co-movement between price growth expectations and the PD ratio. The next sections spell out the microfoundations of such a model. As we show, the model can simultaneously replicate the behavior of stock prices and stock price expectations.

4 A Simple Asset Pricing Model

Consider an endowment economy populated by a unit mass of infinitely lived agents $i \in [0, 1]$ with time-separable preferences. Agents trade one unit of a stock in a competitive stock market. They earn each period an exogenous non-dividend income $W_t > 0$ that we refer to as ‘wages’ for simplicity. Stocks deliver the dividend $D_t > 0$. Dividend and wage incomes take the form of perishable consumption goods.

The Investment Problem. Investor i solves

$$\begin{aligned} & \max_{\{C_t^i \geq 0, S_t^i \in \mathcal{S}\}_{t=0}^{\infty}} E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t u(C_t^i) & (6) \\ \text{s.t.} & S_t^i P_t + C_t^i = S_{t-1}^i (P_t + D_t) + W_t \quad \text{for all } t \geq 0 \end{aligned}$$

where $S_{-1}^i = 1$ and C^i denotes consumption, u the instantaneous utility of the consumer, assumed to be continuous, differentiable, increasing and strictly concave, S^i the agent’s stockholdings, chosen from some compact, non-empty and convex set \mathcal{S} such that $1 \in \mathcal{S}$,

²⁷Interestingly, the relationship between implied price growth expectations and the PD ratio seems to have shifted upwards after the year 2000, as indicated by the squared icons in figure 4. We will come back to this observation in section 9.

$P \geq 0$ the (ex-dividend) price of the stock, $D \geq 0$ an exogenous dividend, $W \geq 0$ the exogenous wage income, and \mathcal{P}^i the agent's subjective probability measure, which may or may not satisfy the rational expectations hypothesis. Further details of \mathcal{P}^i will be specified below.

Dividend and Wage Income. As standard in the literature, we assume that dividends grow at a constant rate and that dividend growth innovations are unpredictable

$$\ln D_t = \ln \beta^D + \ln D_{t-1} + \ln \varepsilon_t^D,$$

where $\beta^D \geq 1$ denotes gross mean dividend growth, $\ln \varepsilon_t^D$ an i.i.d. growth innovation described further below.

We also specify an exogenous wage income process W_t , which is chosen such that the resulting aggregate consumption process $C_t = W_t + D_t$ is empirically plausible. First, in line with Campbell and Cochrane (1999), we set the standard deviation of consumption growth to be 1/7 of the standard deviation of dividend growth. Second, again following these authors, we set the correlation between consumption and dividend growth equal to 0.2. Third, we choose a wage process such that the average consumption-dividend ratio in the model ($E[C_t/D_t]$) equals the average ratio of personal consumption expenditure to net dividend income, which equals approximately 22 in U.S. postwar data.²⁸ All this can be parsimoniously achieved using the following wage income process

$$\ln W_t = \ln \rho + \ln D_t + \ln \varepsilon_t^W,$$

where

$$\begin{pmatrix} \ln \varepsilon_t^D \\ \ln \varepsilon_t^W \end{pmatrix} \sim i.i.N \left(-\frac{1}{2} \begin{pmatrix} \sigma_D^2 \\ \sigma_W^2 \end{pmatrix}, \begin{pmatrix} \sigma_D^2 & \sigma_{DW} \\ \sigma_{DW} & \sigma_W^2 \end{pmatrix} \right) \quad (7)$$

and $E\varepsilon_t^D = E\varepsilon_t^W = 1$. Given the variance of dividend growth σ_D^2 , which can be estimated from dividend data, one can use σ_{DW} and σ_W^2 to impose the desired volatility of consumption growth and the desired correlation with dividend growth. Furthermore, one can choose $\rho = 22$ to obtain the targeted average consumption-dividend ratio. Appendix A.3 explains how this is achieved.

The Underlying Probability Space. Agents hold a set of subjective probability beliefs about all payoff-relevant variables that are beyond their control. In addition to fundamental variables such as dividends and wage income, agents also perceive compet-

²⁸See appendix A.3 for details.

itive stock prices to be beyond their control. Therefore, the belief system also specifies probabilities about prices. Formally, letting Ω denote the space of possible realizations for infinite sequences, a typical element $\omega \in \Omega$ is given by $\omega = \{P_t, D_t, W_t\}_{t=0}^\infty$. As usual, Ω^t then denotes the set of all (nonnegative) price, dividend and wage histories from period zero up to period t and ω^t its typical element. The underlying probability space for agents' beliefs is then given by $(\Omega, \mathcal{B}, \mathcal{P}^i)$ with \mathcal{B} denoting the corresponding σ -Algebra of Borel subsets of Ω , and \mathcal{P}^i a probability measure over (Ω, \mathcal{B}) .

The agents' plans will be contingent on the history ω^t , i.e., the agent chooses state-contingent consumption and stockholding functions

$$C_t^i : \Omega^t \rightarrow \mathcal{R}^+ \tag{8}$$

$$S_t^i : \Omega^t \rightarrow \mathcal{S} \tag{9}$$

The fact that C^i and S^i depend on price realizations is a consequence of optimal choice under uncertainty, given that agents consider prices to be exogenous random variables.

The previous setup is general enough to accommodate situations where agents learn about the stochastic processes governing the evolution of prices, dividends, and wages. For example, \mathcal{P}^i may arise from a stochastic process describing the evolution of these variables that contains unknown parameters about which agents hold prior beliefs. The presence of unknown parameters then implies that agents update their beliefs using the observed realizations of prices, dividends and wages. A particular example of this kind will be presented in section 6 when we discuss learning about stock price behavior.

The probability space defined above is more general than that specified in a RE analysis of the model, where Ω contains usually only the variables that are exogenous to the model (in this case D_t and W_t), but not variables that are endogenous to the model and exogenous to the agent only (in this case P_t). Under the RE hypothesis, agents are assumed to know the pricing function $P_t((D, W)^t)$ mapping histories of dividends and wages into a market price. In that case prices carry redundant information and can be excluded from the probability space without loss of generality. The more general formulation we entertain here allows us to consider agents who do not know exactly which price materializes given a particular history of dividends and wages; our agents do have a view about the distribution of P_t conditional on $(D, W)^t$, but in their minds this is a proper distribution, not a point mass as in the RE case. Much akin to academic economists, investors in our model have not converged on a single asset pricing model

that associates one market price with a given history of exogenous fundamentals.

Parametric Utility Function. To obtain closed-form solutions, we consider in the remaining part of the paper the utility function

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma} \quad \text{with } \gamma > 1, \quad (10)$$

and also consider agents who hold rational expectations about dividends and wages (\mathcal{P}^i incorporates knowledge of the process (7)), so as to be able to isolate the pricing effects arising from subjective capital gains beliefs. We furthermore assume that

$$\delta\beta^{RE} < 1, \quad (11)$$

where $\beta^{RE} \equiv (\beta^D)^{1-\gamma} e^{\gamma(\gamma-1)\sigma_D^2/2}$, which insures existence of an equilibrium under rational price expectations. Since solving the optimization problem (6) for general (potentially non-rational) price beliefs is non-standard, appendix A.4 discusses conditions guaranteeing existence of an optimum, sufficiency of first order conditions and the existence of a recursive solution. These conditions are all satisfied for the preference specification (10) and the subjective price beliefs introduced in the remaining part of the paper and guarantee that the optimal solution to (6) takes the form

$$S_t^i = S^i \left(S_{t-1}^i, \frac{P_t}{D_t}, \frac{W_t}{D_t}, m_t^i \right). \quad (12)$$

where m_t^i is a sufficient statistic characterizing the subjective distributions about future values of $\left(\frac{D_{t+j}}{D_{t+j-1}}, \frac{P_{t+j}}{D_{t+j}}, \frac{W_{t+j}}{D_{t+j}} \right)$ for $j > 0$.

5 Rational Expectations (RE) Equilibrium

As a point of reference, we determine the equilibrium stock price implied by the RE hypothesis. Appendix A.5 derives the following result:

Proposition 1 *If agents hold rational expectations and if price expectations satisfy the usual transversality condition (stated explicitly in appendix A.5), then RE equilibrium price is given by*

$$\frac{P_t^{RE}}{D_t} = (1 + \rho\varepsilon_t^W)^\gamma b \frac{\delta\beta^{RE}}{1 - \delta\beta^{RE}} \quad (13)$$

where $b \equiv E[(1 + \rho\varepsilon_t^W)^{-\gamma} (\varepsilon_t^D)^{1-\gamma}] e^{\gamma(1-\gamma)\sigma_D^2/2}$ and $\beta^{RE} \equiv (\beta^D)^{1-\gamma} e^{\gamma(\gamma-1)\sigma_D^2/2}$.

The PD ratio is an iid process under RE, thus fails to match the persistence of the PD ratio observed in the data. Moreover, since the volatility of ε_t^W tends to be small, it fails to match the large variability of stock prices. Furthermore, the RE equilibrium implies a *negative* correlation between the PD ratio and expected returns, contrary to what is evidenced by survey data. To see this note that (13) implies

$$\ln P_{t+1}^{RE} - \ln P_t^{RE} = \ln \beta^D + \ln \varepsilon_{t+1}^P, \quad (14)$$

where $\varepsilon_{t+1}^P \equiv \varepsilon_{t+1}^D (1 + \rho \varepsilon_{t+1}^W) / (1 + \rho \varepsilon_t^W)$, so that one-step-ahead price growth expectations covary negatively with the current price dividend ratio.²⁹ Since the dividend component of returns also covaries negatively with the current price, the same holds true for expected returns.

In the interest of deriving analytical solutions, we consider below the limiting case with vanishing uncertainty ($\sigma_D^2, \sigma_W^2 \rightarrow 0$). The RE solution then simplifies to the perfect foresight outcome

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

which has prices and dividends growing at the common rate β^D .

6 Learning about Capital Gains and Internal Rationality

Price growth in the RE equilibrium displays only short-lived deviations from dividend growth, with any such deviation being undone in the subsequent period, see equation (14). Price growth in the data, however, can persistently outstrip dividend growth, thereby giving rise to a persistent increase in the PD ratio and an asset price boom; conversely it can fall persistently short of dividend growth and give rise to a price bust, see figure 1. This behavior of actual asset prices suggests that it is of interest to relax the RE beliefs about price behavior. Indeed, in view of the behavior of actual asset prices in the data, agents may entertain a more general model of price behavior, incorporating the possibility that the growth rate of prices persistently exceeds/falls short of the growth rate of dividends. To the extent that the equilibrium asset prices implied by these beliefs display such data-like behavior, agents' beliefs will be generically validated.

²⁹The PD ratio under RE is proportional to $1 + \rho \varepsilon_t^W$, see equation (13), while ε_{t+1}^P depends inversely on $1 + \rho \varepsilon_t^W$.

Generalized Price Beliefs. In line with the discussion in the previous paragraph, we assume agents perceive prices evolving according to the process

$$\ln P_{t+1} - \ln P_t = \ln \beta_{t+1} + \ln \varepsilon_{t+1}, \quad (16)$$

where ε_{t+1} denotes a transitory shock to price growth and β_{t+1} a persistent price growth component that drifts slowly over time according to

$$\ln \beta_{t+1} = \ln \beta_t + \ln \nu_{t+1}. \quad (17)$$

This setup can capture periods with sustained increases in the PD ratio ($\beta_{t+1} > \beta^D$) or sustained decreases ($\beta_{t+1} < \beta^D$).³⁰ In the limiting case where the variance of the innovation $\ln \nu_{t+1}$ becomes small, the persistent price growth component behaves almost like a constant, as is the case in the RE solution.

For simplicity, we assume that agents perceive the innovations $\ln \varepsilon_{t+1}$ and $\ln \nu_{t+1}$ to be jointly normally distributed according to

$$\begin{pmatrix} \ln \varepsilon_{t+1} \\ \ln \nu_{t+1} \end{pmatrix} \sim iiN \left(\begin{pmatrix} -\frac{\sigma_\varepsilon^2}{2} \\ -\frac{\sigma_\nu^2}{2} \end{pmatrix}, \begin{pmatrix} \sigma_\varepsilon^2 & 0 \\ 0 & \sigma_\nu^2 \end{pmatrix} \right). \quad (18)$$

Since agents observe the change of the asset price, but do not separately observe the persistent and transitory elements driving it, the previous setup defines a filtering problem in which agents need to decompose observed price growth into the persistent and transitory subcomponents, so as to forecast optimally.

To emphasize the importance of learning about price behavior rather than learning about the behavior of dividends or the wage income process, which was the focus of much of an earlier literature on learning in asset markets, e.g., Timmermann (1993, 1996), we continue to assume that agents know the processes (7), i.e., hold rational dividend and wage expectations.

Internal Rationality of Price Beliefs. Among academics there appears to exist a widespread belief that rational behavior and knowledge of the fundamental processes (dividends and wages in our case) jointly *dictate* a certain process for prices and thus the

³⁰We deliberately do not incorporate any mean-reversion into price growth beliefs as we seek to determine model-endogenous forces that lead to a reversal of asset price booms and busts, rather than having these features emerge because they are hard-wired into beliefs. Incorporating such mean reversion in prices would not be difficult though. Furthermore, as we discuss below, return expectations display some degree of mean reversion even with the present specification.

price beliefs agents can rationally entertain.³¹ If this were true, then rational behavior would imply rational expectations, so that postulating subjective price beliefs as those specified in equation (16) would be inconsistent with the assumption of optimal behavior on the part of agents.

This view is correct in some special cases, for example when agents are risk neutral and do not face trading constraints. It fails to be true, however, more generally. Therefore, agents in our model are ‘internally rational’: their behavior is optimal given an internally consistent system of subjective beliefs about variables that are beyond their control, including prices.

To illustrate this point, consider first risk neutral agents with rational dividend expectations and ignore limits to stock holdings. Forward-iteration on the agents’ own optimality condition (41) then delivers the present value relationship

$$P_t = E_t \left[\sum_{i=1}^T \delta^i D_{t+i} \right] + \delta^T E_t^{\mathcal{P}^i} [P_{t+T}],$$

which is independent of the agents’ own choices. Provided agents’ price beliefs satisfy a standard transversality condition ($\lim_{T \rightarrow \infty} \delta^T E_t^{\mathcal{P}^i} [P_{t+T}] = 0$ for all i), then each rational agent would conclude that there must be a degenerate joint distribution for prices and dividends given by

$$P_t = E_t \left[\sum_{i=1}^{\infty} \delta^i D_{t+i} \right] \text{ a.s.} \quad (19)$$

Since the r.h.s of the previous equation is fully determined by dividend expectations, the beliefs about the dividend process deliver the price process compatible with optimal behavior. In such a setting, it would be plainly inconsistent with optimal behavior to assume the subjective price beliefs (16)-(17).³²

Next, consider a concave utility function $u(\cdot)$ satisfying standard Inada conditions. Forward iteration on (41) and assuming an appropriate transversality condition then delivers

$$P_t u'(C_t^i) = E_t^{\mathcal{P}^i} \left[\sum_{j=1}^{\infty} \delta^j D_{t+j} u'(C_{t+j}^i) \right] \text{ a.s.} \quad (20)$$

Unlike in equation (19), the previous equation depends on the agent’s current and future consumption. Equation (20) thus falls short of mapping beliefs about the dividend process

³¹We often received this reaction during seminar presentations.

³²See Adam and Marcet (2011) for a discussion of how in the presence of trading constraints, this conclusion breaks down, even with risk-neutral consumption preferences.

into a price outcome. Indeed, given *any* equilibrium price P_t , the agent will choose her consumption plans such that (20) holds, i.e., such that the price equals the discounted sum of dividends, discounting with her on internally rational consumption plan.³³ Equation (20) thus fails to deliver any restriction on what optimizing agents can possibly believe about the price process.

With the considered non-linear utility function, we can thus simultaneously assume that agents maximize utility, hold the subjective price beliefs (16)-(17) and rational beliefs about dividends and wages.

Learning about the Capital Gains Process. The beliefs (16) give rise to an optimal filtering problem. To obtain a parsimonious description of the evolution of beliefs we specify conjugate prior beliefs about the unobserved persistent component $\ln \beta_t$ at $t = 0$. Specifically, agent i 's prior is

$$\ln \beta_0 \sim N(\ln m_0^i, \sigma^2), \quad (21)$$

where prior uncertainty σ^2 is assumed to be equal to its Kalman filter steady state value, i.e.,

$$\sigma^2 \equiv \frac{-\sigma_v^2 + \sqrt{(\sigma_v^2)^2 + 4\sigma_v^2\sigma_\varepsilon^2}}{2}, \quad (22)$$

and the prior is also assumed independent of all other random variables at all times. Equations (16), (17) and (21), and knowledge of the dividend and wage income processes (7) then jointly specify agents' probability beliefs \mathcal{P}^i .

The optimal Bayesian filter then implies that the posterior beliefs following some history ω^t are given by³⁴

$$\ln \beta_t | \omega^t \sim N(\ln m_t^i, \sigma^2), \quad (23)$$

with

$$\ln m_t^i = \ln m_{t-1}^i - \frac{\sigma_v^2}{2} + g \left(\ln P_t - \ln P_{t-1} + \frac{\sigma_\varepsilon^2 + \sigma_v^2}{2} - \ln m_{t-1}^i \right) \quad (24)$$

$$g = \frac{\sigma^2}{\sigma_\varepsilon^2}. \quad (25)$$

Agents' beliefs can thus be parsimoniously summarized by a single state variable (m_t^i)

³³This follows directly from the fact that consumption plans must satisfy (41) at all contingencies.

³⁴See theorem 3.1 in West and Harrison (1997). Choosing a value for σ^2 different from the steady state value (22) would only add a deterministically evolving variance component σ_t^2 to posterior beliefs with the property $\lim_{t \rightarrow \infty} \sigma_t^2 = \sigma^2$, i.e., it would converge to the steady state value.

describing agents' degree of optimism about future capital gains. These beliefs evolve recursively according to equation (24) and imply that

$$E_t^{\mathcal{P}^i} \left[\frac{P_{t+1}}{P_t} \right] = e^{\ln m_t^i} e^{\sigma^2/2}, \quad (26)$$

which is - up to the presence of a log and exponential transformation and some variance correction terms - identical to the adaptive prediction model considered in section 3.3.

Nesting PF Equilibrium Expectations. The subjective price beliefs (16),(17) and (21) generate perfect foresight equilibrium price expectations in the special case in which prior beliefs are centered at the growth rate of dividends, i.e.,

$$\ln m_0^i = \ln \beta^D,$$

and when considering the limiting case with vanishing uncertainty, where $(\sigma_\varepsilon^2, \sigma_\nu^2, \sigma_D^2, \sigma_W^2) \rightarrow 0$. Agents' prior beliefs at $t = 0$ about price growth in $t \geq 1$ then increasingly concentrates at the perfect foresight outcome $\ln \beta^D$, see equations (16) and (17). With price and dividend expectations being at their PF value, the perfect foresight price $PD_0 = \delta\beta^{RE}/(1 - \delta\beta^{RE})$ becomes the equilibrium outcome at $t = 0$ in the limit. Importantly, it continues to be possible to study learning dynamics in the limit with vanishing risk: keeping the limiting ratio $\sigma_\nu^2/\sigma_\varepsilon^2$ finite and bounded from zero as uncertainty vanishes, the Kalman gain parameter g defined in (25), remains well-specified in the limit and satisfies $\lim \frac{\sigma_\nu^2}{\sigma_\varepsilon^2} = \lim \frac{g^2}{1-g}$. We will exploit this fact in section 8 when presenting analytical results.

7 Dynamics under Learning

This section explains how equilibrium prices are determined under the subjective beliefs introduced in the previous section and how they evolve over time.

Agents' stock demand is given by equation (12). Stock demand depends on the belief m_t^i , which characterizes agents' capital gains expectations. These beliefs evolve according to (24). As a benchmark, we shall now assume that all agents hold identical beliefs ($m_t^i = m_t$ for all i). While agents may initially hold heterogenous prior beliefs m_0^i , heterogeneity would asymptotically vanish because all agents observe the same price history. The asset dynamics derived under the assumption of identical beliefs thus describe the long-run outcome of the model.

Using this assumption and imposing market clearing in periods t and $t - 1$ in equation (12) shows that the equilibrium price in any period $t \geq 0$ solves

$$1 = S \left(1, \frac{P_t}{D_t}, \frac{W_t}{D_t}, m_t \right), \quad (27)$$

which exploits the fact that the total supply of stocks is equal to one.

The beliefs m_t and the price dividend ratio P_t/D_t are now simultaneously determined via equations (24) and (27). Unfortunately, this simultaneity could give rise to multiple market clearing price and belief pairs, due to a complementarity between realized capital gains and expected future capital gains.³⁵ While this multiplicity may be a potentially interesting avenue to explain asset price booms and busts, analyzing price dynamics within such a setting would require introducing non-standard features, such as an equilibrium selection device for periods in which there are multiple solutions to (24) and (27). Instead, we resort to a standard approach of using only lagged information for updating beliefs.

Appendix A.6 shows that the simultaneity can be overcome by slightly modifying the information structure. The modification is relatively straightforward and consists of assuming that agents observe at any time t information about the lagged temporary price growth component ε_{t-1} entering equation (16). The appendix then shows that Bayesian updating implies that

$$\ln m_t = \ln m_{t-1} - \frac{\sigma_v^2}{2} + g \left(\ln P_{t-1} - \ln P_{t-2} + \frac{\sigma_\varepsilon^2 + \sigma_v^2}{2} - \ln m_{t-1} \right) + g \ln \varepsilon_t^1, \quad (28)$$

where updating now occurs using only lagged price growth (even though agents do observe current prices) and where $\ln \varepsilon_t^1 \sim iiN\left(\frac{-\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right)$ is a time t innovation to agent's information set (unpredictable using information available to agents up to period $t - 1$), which reflects the information about the transitory price growth component ε_{t-1} received in period t .

With this slight modification, agents' beliefs m_t are now pre-determined at time t , so that the economy evolves according to a uniquely determined recursive process: equation (27) determines the market clearing price for period t given the beliefs m_t and equation (28) determines how time t beliefs are updated following the observation of the new

³⁵Intuitively, a higher PD ratio implies higher realized capital gains and thus higher expectations of future gains via equation (24). Higher expected future gains may in turn induce a higher willingness to pay for the asset, thereby justifying the higher initial PD ratio.

market clearing price.³⁶

8 Equilibrium: Analytic Findings

This section derives a closed form solution for the equilibrium asset price for the special case where all agents hold the same subjective beliefs \mathcal{P} and where these beliefs imply no (or vanishing) uncertainty about future prices, dividends and wages. While the absence of uncertainty is unrealistic from an empirical standpoint, it allows deriving key insights into how the equilibrium price depends on agents' beliefs, as well as on how prices and beliefs evolve over time.³⁷ The empirically more relevant case with uncertainty will be considered in section 9 using numerical solutions.

The next section provides a closed form expression for the equilibrium PD ratio as a function of agents' subjective expectations about future stock market returns. Section 8.2 then discusses the pricing implications of this result for the subjective capital gains beliefs introduced in section 6. Finally, section 8.3 shows how the interaction between asset price behavior and subjective belief revisions can temporarily de-link asset prices from their fundamental value, i.e., give rise to a self-feeding boom and bust in asset prices along which subjective expected returns rise and fall.

8.1 Main Result

The following proposition summarizes our main finding:³⁸

Proposition 2 *Suppose $u(C) = C^{1-\gamma}/(1-\gamma)$, agents' beliefs \mathcal{P} imply no uncertainty about future prices, dividends and wages, and*

$$\lim_{T \rightarrow \infty} E_t^{\mathcal{P}} R_T > 1 \text{ and } \lim_{T \rightarrow \infty} E_t^{\mathcal{P}} \left(\sum_{j=1}^T \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right) W_{t+j} \right) < \infty, \quad (29)$$

³⁶There could still be an indeterminacy arising from the fact that $S(\cdot)$ is non-linear, so that equation (27) may not have a unique solution, but we have encountered such problems neither in our analytical solution nor when numerically solving the model.

³⁷In the absence of uncertainty one can evaluate more easily the expectations of nonlinear functions of future variables showing up in agents' FOCs.

³⁸The proof can be found in appendix A.7.

then the equilibrium PD ratio in period t is given by

$$\begin{aligned} \frac{P_t}{D_t} &= \left(1 + \frac{W_t}{D_t}\right) \sum_{j=1}^{\infty} \left(\left(\delta^{\frac{1}{\gamma}}\right)^j \left(E_t^{\mathcal{P}} \prod_{i=1}^j \frac{1}{R_{t+i}} \right)^{\frac{\gamma-1}{\gamma}} \right) \\ &\quad - \frac{1}{D_t} E_t^{\mathcal{P}} \left(\sum_{j=1}^{\infty} \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right) W_{t+j} \right) \end{aligned} \quad (30)$$

Conditions (29) insure that the infinite sums in the pricing equation (30) converge.³⁹ Under the additional assumption that agents hold rational wage and dividend expectations, equation (30) simplifies further to

$$\begin{aligned} \frac{P_t}{D_t} &= (1 + \rho) \sum_{j=1}^{\infty} \left(\left(\delta^{\frac{1}{\gamma}}\right)^j \left(E_t^{\mathcal{P}} \prod_{i=1}^j \frac{1}{R_{t+i}} \right)^{\frac{\gamma-1}{\gamma}} \right) \\ &\quad - \rho \left(\sum_{j=1}^{\infty} (\beta^D)^j \left(E_t^{\mathcal{P}} \prod_{i=1}^j \frac{1}{R_{t+i}} \right) \right). \end{aligned} \quad (31)$$

We now discuss the implications of equation (31), focusing on the empirically relevant case where $\rho > 0$ and $\gamma > 1$.

Consider first the upper term on the r.h.s. of equation (31), which is decreasing in the expected asset returns. This emerges because for $\gamma > 1$ the wealth effect of a change in return expectations then dominates the substitution effect, so that expected asset demand and therefore the asset price has a tendency to decrease as return expectations increase. The negative wealth effect thereby increases in strength if the ratio of wage to dividend income (ρ) increases. This is the case because higher return expectations also reduce the present value of wage income.

Next, consider the lower term on the r.h.s. of equation (31), including the negative sign pre-multiplying it. This term depends positively on the expected returns and captures a substitution effect that is associated with increased return expectations. This substitution effect only exists if $\rho > 0$, i.e., only in the presence of non-dividend income, and it is increasing in ρ . It implies that increased return expectations are associated with increased stock demand and thus with a higher PD ratio in equilibrium. It is this term that allows the model to match the positive correlation between expected returns and the PD ratio.

³⁹These are satisfied, for example, for the expectations associated with the perfect foresight RE solution. Equation (30) then implies that the PD ratio equals the perfect foresight PD ratio (15), as is easily verified. Conditions (29) are equally satisfied for the subjective beliefs defined in section 6, when considering the case with vanishing uncertainty $(\sigma_\varepsilon^2, \sigma_\nu^2, \sigma_D^2, \sigma_W^2) \rightarrow 0$.

This substitution effect is present even in the limiting case with log consumption utility ($\gamma \rightarrow 1$). The upper term on the r.h.s. of equation (31) then vanishes because the substitution and wealth effects associated with changes in expected returns cancel each other, but the lower term still induces a positive relationship between prices and return expectations. The substitution effect is also present for $\gamma > 1$ and can then dominate the negative wealth effect arising from the upper term on the r.h.s. of (31). Consider, for example, the opposite limit with $\gamma \rightarrow \infty$. Equation (31) then delivers

$$\frac{P_t}{D_t} = \sum_{j=1}^{\infty} \left(1 + \rho \sum_{j=1}^{\infty} \left(1 - (\beta^D)^j \right) \right) \left(E_t^{\mathcal{P}} \prod_{i=1}^j \frac{1}{R_{t+i}} \right).$$

Since $\beta^D > 1$, there is a positive relationship between prices and expected asset returns, whenever ρ is sufficiently large. The two limiting results ($\gamma \rightarrow 1$ and $\gamma \rightarrow \infty$) thus suggest that for sufficiently large ρ the model can generate a positive relationship between return expectations and the PD ratio, in line with the evidence obtained from survey data.

8.2 PD Ratio and Expected Capital Gains

We now consider the implications of equation (31) for the subjective capital gains beliefs introduced in section 6.⁴⁰ Equation (31) implies a non-linear relationship between the PD ratio and the subjective capital gain expectations m_t , but one cannot obtain a closed-form solution for the PD ratio as a function of the capital gains expectations.⁴¹ Figure 5 depicts the relationship between the PD ratio and m_t using the parameterization employed in our quantitative application in section 9, but abstracting from future uncertainty.⁴²

Figure 5 shows that there is a range of price growth beliefs in the neighborhood of the perfect foresight value ($m_t = \beta^D$) over which the PD ratio depends positively on expected

⁴⁰ Appendix A.8 proves that condition (29) is then satisfied for all beliefs $m_t > 0$.

⁴¹ More precisely, with vanishing uncertainty the beliefs from section 6 imply

$$E_t^{\mathcal{P}} [P_{t+i}] = (m_t)^i P_t,$$

which together with perfect foresight about dividends allows expressing agents' expectations of future inverse returns as a function of m_t and the current PD ratio:

$$E_t^{\mathcal{P}} \frac{1}{R_{t+i}} = \frac{E_t^{\mathcal{P}} P_{t+i-1}}{E_t^{\mathcal{P}} P_{t+i} + E_t^{\mathcal{P}} D_{t+i}} = \frac{(m_t)^{i-1} \frac{P_t}{D_t}}{(m_t)^i \frac{P_t}{D_t} + (\beta^D)^i}.$$

Substituting this into (31) one can solve numerically for P_t/D_t as a function of m_t .

⁴² The parameterization assumes a moderate degree of risk aversion $\gamma = 2$, a quarterly discount factor of $\delta = 0.995$, quarterly real dividend growth equal to the average postwar growth rate of real dividends $\beta^D = 1.0048$, and $\rho = 22$ to match the average dividend-consumption ratio in the U.S. over 1946-2011, see section 9 for further details.

price growth, similar to the positive relationship between expected returns and the PD ratio derived analytically in the previous section. Over this range, the substitution effect dominates the wealth effect because our calibration implies that dividend income finances only a small share of total consumption (approximately 4.3%). As a result, stock market wealth is only a small share of the total present value of household wealth (the same 4.3%) when beliefs assume their perfect foresight value ($m_t = \beta^D$).

Figure 5 also reveals that there exists a capital gains belief beyond which the PD ratio starts to decrease. Mathematically, this occurs because if $m_t \rightarrow \infty$, expected returns also increase without bound⁴³, so that $E_t^{\mathcal{P}} \prod_{i=1}^j \frac{1}{R_{t+i}} \rightarrow 0$. From equation (31) one then obtains $\frac{P_t}{D_t} \rightarrow 0$.

The economic intuition for the existence of a maximum PD ratio is as follows: for higher m_t the present value of wage income is declining, as increased price growth optimism implies higher expected returns⁴⁴ and therefore a lower discount factor. This can be seen by noting that the FOC (41) can alternatively be written as

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

which implies that increased return expectations $E_t^{\mathcal{P}} R_{t+1}$ imply a lower discount factor $\delta E_t^{\mathcal{P}} [(C_{t+1}/C_t)^{-\gamma}]$.⁴⁵ With increased optimism, the present value of wage income thus falls. At the same time, stock market wealth initially increases strongly. Indeed, at the maximum PD ratio, stock market wealth amounts to approximately 4.5 times the value it assumes in the perfect foresight solution, see figure 5. This relative wealth shift has the same effect as a decrease in the wage to non-wage income ratio ρ . As argued in section 8.1, for sufficiently small values of ρ the income effect starts to dominate the substitution effect, so that prices start to react negatively to increased return optimism.

8.3 Endogenous Boom and Bust Dynamics

We now explain how the interplay between price realizations and belief updating can temporarily de-link asset prices from their fundamental values. This process emerges endogenously and takes the form of a sustained asset price boom along which expected returns rise and that ultimately results in a price bust along which expected returns fall.

⁴³This follows from $E_t^{\mathcal{P}} R_{t+i+1} = E_t^{\mathcal{P}} \frac{P_{t+i+1} + D_{t+i+1}}{P_{t+i}} > E_t^{\mathcal{P}} \frac{P_{t+i+1}}{P_{t+i}} = m_t$.

⁴⁴This is shown in appendix A.9, which depicts the relationship between expected capital gains and expected returns at various forecast horizons.

⁴⁵This holds true under the maintained assumption of no or vanishing uncertainty.

This feature allows the model to generate volatile asset prices and to capture the positive correlation between expected returns and the PD ratio.

Consider figure 5 and a situation in which agents become optimistic, in the sense that their capital gains expectations m_t increase slightly above the perfect foresight value $m_{t-1} = \beta^D$ entertained in the previous period.⁴⁶ Figure 5 shows that this increase in expectations leads to an increase in the PD ratio, i.e., $P_t/D_t > P_{t-1}/D_{t-1}$. Moreover, due to the relatively steep slope of the PD function, realized capital gains will strongly exceed the initial increase in expected capital gains. The belief updating equation (28) then implies further upward revisions in price growth expectations and thus further capital gains, leading to a sustained asset price boom in which the PD ratio and return expectations jointly move upward.

The price boom comes to an end when expected price growth reaches a level close to where the PD function in figure 5 reaches its maximum.⁴⁷ At this point, stock prices grow at most at the rate of dividends (β^D), but agents hold considerably more optimistic expectations about future capital gains ($m_t > \beta^D$). Investors' high expectations will thus be disappointed, which subsequently leads to a reversal.

The previous dynamics are also present in a stochastic model considered in the next sections. They introduce low frequency movements in the PD ratio, allowing the model to replicate boom and bust dynamics and thereby to empirically plausible amounts of asset price volatility, despite assuming standard consumption preferences. These dynamics also generate a positive correlation between the PD ratio and expected returns.⁴⁸

9 Historical PD Ratio and Survey Evidence

This section considers the asset pricing model with subjective beliefs and uncertainty; it shows that the model can successfully replicate the low-frequency movements in the postwar U.S. PD ratio, as well as the available survey evidence.

Solving the non-linear asset pricing model with subjective beliefs is computationally costly, which prevents us from pursuing formal estimation or moment matching. We thus

⁴⁶In the model with uncertainty, such upward revisions can be triggered by fundamentals, e.g., by an exceptionally high dividend growth realization in the previous period, which is associated with an exceptionally high price growth realization.

⁴⁷In the model with noise, fundamental shocks, e.g., a low dividend growth realization, can cause the process to end well before reaching this point.

⁴⁸While the arguments above only show that expected capital gains correlate positively with the PD ratio, Appendix A.9 shows that expected capital gains and expected returns comove positively, so that expected returns also comove positively with the PD ratio.

resort to calibration.

Table 6 reports the calibrated parameters and the calibration targets.⁴⁹ The mean and standard deviation of dividend growth (β^D and σ_D) are chosen to match the corresponding empirical moments of the U.S. dividend process. The ratio of non-dividend to dividend income (ρ) is chosen to match the average dividend-consumption ratio in the U.S. for 1946-2011.⁵⁰ The standard deviation of wage innovations (σ_W) and the covariance between wage and dividend innovations (σ_{DW}) are chosen, in line with Campbell and Cochrane (1999), such that the correlation between consumption and dividend growth is 0.2 and the standard deviation of consumption growth is one seventh of the the standard deviation of dividend growth.⁵¹ The perceived uncertainty in stock price growth (σ_ε) is set equal to the empirical standard deviation of stock price growth.⁵²

This leaves us with four remaining parameters: the belief updating parameter g , the initial price growth belief $m_{Q1:1946}$, the time discount factor δ and the risk aversion parameter γ . We choose $g = 0.02515$ and $m_{Q1:1946} = -1.11\%$, in line with the values employed in constructing figure 3, which allowed matching the UBS survey expectations. We then assume risk aversion of $\gamma = 2$ and choose the quarterly discount factor δ , so as to obtain a good match between the model-implied and the empirical PD ratio over the postwar period. It turns out that $\delta = 0.995$ achieves a good fit.

Figure 6 depicts the equilibrium PD ratio obtained from numerically solving the asset pricing model with uncertainty, together with the equilibrium PD ratio in the absence of uncertainty analyzed in the previous section.⁵³ While the presence of price, dividend, and wage risk lowers the equilibrium PD ratio compared to a setting without risk, the functional form of the relationship remains qualitatively unchanged. The findings obtained

⁴⁹The targets are chosen to match features of the fundamental processes emphasized in the asset pricing literature.

⁵⁰See appendix A.3 for further details.

⁵¹For details on how this can be achieved, see appendix A.3.

⁵²Since the gain parameter g will be small, the contribution of σ_ν^2 in (16) is negligible.

⁵³The numerical solution is obtained by numerically determining the stock demand function (12) solving the FOC (41) under the subjectively perceived dividend, wage and price dynamics, where agents understand that their beliefs evolve according to (28). The PD ratio as a function of m_t depicted in figure 6 is determined from the market clearing condition (27) assuming $W_t/D_t = \rho$, to be comparable with the value this variable assumes in the vanishing risk limit. We verified that in the limiting case without uncertainty, our numerical solution algorithm recovers the analytical solution derived in proposition 2. Furthermore, in the case with uncertainty, we insure the accuracy of the numerical solution by verifying that the Euler equation errors are in the order of 10^{-5} over the relevant area of the state space. Insuring this requires a considerable amount of adjustment by hand of the grid points and grid size used for spanning the model's state space. This prevents us from formally estimating the model, as the model cannot be solved with sufficient accuracy using an automated procedure. Further details of the solution approach are described in appendix A.10. The MatLab code used for solving the model is available upon request.

in the previous section thus continue to apply in the presence of quantitatively realistic amounts of uncertainty.

We now evaluate the ability of the model to replicate the postwar time series of the PD ratio. We do so by first feeding the historical capital gains into our model-based belief updating equation (28), so as to obtain a model-implied value for expected capital gains.⁵⁴ The resulting series is shown in figure 7. It displays a strong rise and fall of price growth expectations around the year 2000, as well as relatively low capital gains expectations from the mid 1970's to the mid 1980's. In a second step, we use the model-implied equilibrium PD function to derive a model-implied time series for the PD ratio associated with the model-based beliefs. Figure 8 depicts this model-implied PD series and compares it with historical PD series.⁵⁵

Figure 8 reveals that the model captures a lot of the low-frequency variation in the historically observed PD ratio. It captures particularly well the variations before the year 2000, including the strong run-up in the PD ratio from the mid 1990's to year 2000. The model also predicts a strong decline of the PD ratio after the year 2000, but overpredicts the decline relative to the data. For the period up to and including the year 2000, the variance of the gap between the model predicted PD and the PD in the data amounts to just 20.1% of the overall variance of the PD in the data. In this sense, the subjective belief model is capable of capturing approximately 79.9% of the variation of the PD ratio in the data. Since the fit deteriorates some time after the year 2000, it explains - using the same measure - about 52.5% of the variance for the full sample. We find this to be a remarkable result.

Figure 9 depicts the model-implied price growth expectations and those implied by the UBS survey.⁵⁶ While the model fits the survey data overall well, the model predicts after the year 2003 considerably lower capital gains expectations, which partly explains why the model underpredicts the PD ratio in figure 8 towards the end of the sample period. Yet, the expectations gap in figure 9 narrows considerably after the year 2004, while this fails to be the case in figure 8. Underprediction of expected price growth thus explains only partly the deterioration of the fit of the PD ratio towards the end of the sample period .

⁵⁴We thereby shut down all other sources of information about price growth, i.e., set $\ln \varepsilon_t^1 = 0$ for all t in equation (28).

⁵⁵When computing the model-implied PD ratio, we set the non-dividend to dividend income ratio equal to its steady state value ($W_t/D_t = \rho$), so as to obtain only pricing effects due to variation in subjective capital gains expectations. The effects of fundamental shocks to wages and dividends will be considered in section 10.

⁵⁶See footnote 25 for how to compute price growth expectations from the UBS survey.

The gap after the year 2000 emerging in figure 8 is hardly surprising, given the empirical evidence presented in figure 4, which shows that the relationship between the PD ratio and the expectations implied by equation (5) has shifted upward in the data following the year 2000. While we can only speculate about potential reasons causing this shift, the exceptionally low real interest rates implemented by the Federal Reserve following the reversal of the tech stock boom and following the collapse of the subsequent housing boom may partly contribute to the observed discrepancy. Formally incorporating the effects of monetary policy decisions - while of interest - is beyond the scope of the present paper.

10 Model Simulations

The previous section evaluated to what extent subjective belief updating dynamics alone can explain the behavior of the PD ratio in the data, but it ignored the role of the fundamental dividend and wage processes as ultimate drivers of asset price and belief dynamics. This section evaluates the ability of the model to replicate key asset pricing moments, using model simulations with dividend and wage shocks as fundamental drivers.

To do so, we compare the asset pricing moments in the data to those obtained from simulating the model, considering both the model with subjective beliefs as well as the RE model. We use the parameters from table 6 to simulate the model and formally evaluate the model fit by reporting t-statistics for a number of asset pricing moments.⁵⁷

Table 7 reports the data moments (column 2 of the table), the moments of the subjective beliefs models and the implied t-statistic (columns 3 and 4), as well as the moments and t-statistics of the RE version of the model (columns 5 and 6).⁵⁸ The first eight asset pricing moments listed in the table are those considered in Adam, Marcet, and Nicolini (2013); we augment these by the correlation between the PD ratio and expected stock returns, as implied by the UBS survey data.

The model with subjective beliefs turns out to be able to quantitatively account for many asset pricing moments, even though parameters have not been chosen to maximize the fit.⁵⁹ The RE version of the model performs rather poorly. Besides generating insuffi-

⁵⁷The t-statistic is based on an estimate of the standard deviation of the data moment as a measure of uncertainty, where we estimate the standard deviation of the moment in the data using standard procedures. This delivers an asymptotically valid t-test given the parameter values.

⁵⁸All variables are reported in terms of quarterly real values.

⁵⁹While this would be desirable, numerically solving the model with high accuracy is rather time-consuming.

cient asset price volatility, it wrongly predicts a negative sign for the correlation between the PD ratio and investors' expected returns.

Table 7 reveals that the subjective belief model quantitatively replicates 7 of the 9 considered moments at the 1% significance level, while the RE version matches only 3 of the 9 moments. The learning model replicates particularly well the mean of the PD ratio (denoted by $E[PD]$ in the table), the high autocorrelation of the PD ratio ($\text{Corr}[PD_t, PD_{t-1}]$), the regression coefficient obtained from regressing 5 year ahead excess returns on the current PD ratio (c), as well as the R^2 of that regression (R^2).⁶⁰ The model similarly matches the mean of the stock returns ($E[r^s]$) and the positive correlation between the PD ratio and expected returns ($\text{Corr}[PD_t, E_t^P R_{t+1}]$). It generates a somewhat too high value for the standard deviation of the PD ratio ($\text{Std}[PD]$) and - as a result - predicts a too high value for the standard deviation of stock returns ($\text{Std}[r^s]$). The learning model also misses the equity premium, although it produces about half of the premium observed in the data. This is a considerable success, given the low degree of risk aversion assumed ($\gamma = 2$).

Since none of these moments have been targeted when calibrating the model, the ability of the subjective belief model to quantitatively replicate the data moments is surprisingly good. This is especially true when compared to the performance of the model under RE. Comparing the last column in table 7 to column 4 in the same table shows that the t-statistics all increase (in absolute terms) when imposing RE, with some increases being quite dramatic.

The RE version of the model produces insufficient asset price volatility, i.e., too low values for the standard deviation of the PD ratio and of stock returns. It also produces a tiny equity premium only and gets the sign of the correlation between the PD ratio and expected stock returns wrong. This highlights the strong quantitative improvement in the empirical performance obtained by incorporating subjective belief dynamics. It also highlights that - according to our model - asset price volatility is to a large extent due to subjective belief dynamics.

11 Conclusions

We present a model with rationally investing agents that gives rise to market failures in the sense that the equilibrium stock price deviates from its fundamental value. These

⁶⁰The regression also includes a constant, which is statistically insignificant and whose value is not reported.

deviations take the form of asset price boom and bust cycles that are fueled by the belief-updating dynamics of investors who behave optimally given their imperfect knowledge of the world. Investors update beliefs about market behavior using observed market outcomes and Bayes' law, causing their subjective expectations about future capital gains to comove positively with the price-dividend ratio, consistent with the evidence available from investor surveys. As we argue, this feature cannot be replicated within asset pricing models that impose rational price expectations.

We relax slightly the RE assumption but maintain full rationality of investors. The fact that a fairly small deviation from a standard asset pricing model significantly improves the empirical fit of the model strongly suggests that issues of learning are important when accounting for stock price fluctuations. Indeed, our empirical analysis shows that more than half of the observed variation of the S&P500 PD ratio over the post-war period can be accounted for by variations in subjective beliefs.

If asset price dynamics are to a large extent influenced by investors' subjective belief dynamics, i.e., by subjective optimism and pessimism, then the asset price fluctuations observed in the data are to considerable extent inefficient. Due to a number of simplifying assumptions, this did not yield adverse welfare implications within the present setup.⁶¹ For more realistic models incorporating investor heterogeneity, endogenous output or endogenous stock supply, such fluctuations can give rise to significant distortions that affect welfare. Exploring these within a setting that gives rise to quantitatively credible amounts of asset price fluctuations appears to be an interesting avenue for further research. Such research will in turn lead to further important questions, such as whether policy can and should intervene with the objective to stabilize asset prices.

⁶¹This is true if one evaluates welfare using ex-post realized consumption.

UBS Gallup	Nominal		Real Ret. Exp.		Real Ret. Exp.	
	Return Exp.		(SPF)		(Michigan)	
	Average	Median	Average	Median	Average	Median
Own portfolio, >100k US\$	0.80 (0.01)	0.78 (0.01)	0.79 (0.01)	0.77 (0.01)	0.84 (0.01)	0.83 (0.01)
Own portfolio, all investors	0.80 (0.01)	0.76 (0.02)	0.79 (0.01)	0.75 (0.02)	0.84 (0.01)	0.80 (0.01)
Stock market, >100k US\$	0.90 (0.03)	0.89 (0.04)	0.90 (0.03)	0.88 (0.03)	0.91 (0.03)	0.88 (0.03)
Stock market, all investors	0.90 (0.03)	0.87 (0.04)	0.90 (0.03)	0.87 (0.04)	0.91 (0.03)	0.88 (0.03)

Table 1: Correlation between PD ratio and 1-year ahead expected return measures (UBS Gallup Survey, robust p-values in parentheses)

Shiller Survey	Nominal Capital Gain Exp.		Real Capital Gain. Exp. (SPF)		Real Capital Gain Exp. (Michigan)	
	Average	Median	Average	Median	Average	Median
Horizon						
1 month	0.46 (0.01)	0.48 (0.01)	0.45 (0.01)	0.47 (0.01)	0.46 (0.01)	0.49 (0.01)
3 months	0.57 (0.01)	0.64 (0.00)	0.54 (0.01)	0.61 (0.00)	0.56 (0.01)	0.62 (0.01)
6 months	0.58 (0.01)	0.75 (0.01)	0.54 (0.02)	0.70 (0.01)	0.56 (0.02)	0.71 (0.01)
1 year	0.43 (0.03)	0.69 (0.01)	0.38 (0.05)	0.62 (0.01)	0.42 (0.04)	0.64 (0.02)
10 years	0.74 (0.01)	0.75 (0.01)	0.66 (0.02)	0.71 (0.01)	0.71 (0.02)	0.75 (0.01)

Table 2: Correlation between PD ratio and expected stock price growth (Shiller's Individual Investors' Survey, robust p-values in parentheses)

CFO Survey	Nominal Return Exp.		Real Return Exp. (SPF)		Real Return Exp. (Michigan)	
	Average	Median	Average	Median	Average	Median
1 year	0.71 (0.00)	0.75 (0.00)	0.62 (0.00)	0.69 (0.00)	0.67 (0.00)	0.72 (0.00)

Table 3: Correlation between PD ratio and 1-year ahead expected stock return measures (CFO Survey, robust p-values in parentheses)

Variables	Time Period	Stock Index	Correlation
PD, 1 year-ahead real return	UBS Gallup sample (stock market exp.)	S&P 500	-0.66 (0.08)
PD, 1 year-ahead real price growth	Shiller 1 year sample	Dow Jones	-0.42 (0.06)
PD, 10 year-ahead real price growth	Shiller 10 year sample	Dow Jones	-0.88 (0.16)
PD, 1 year-ahead real return	CFO sample	S&P 500	-0.46 (0.06)

Table 4: Correlation between PD and actual real returns/capital gains
(robust p-value in parentheses)

Survey measure	$\hat{c} \cdot 10^3$	$\hat{c} \cdot 10^3$	p -value $H_0 : \hat{c} = \hat{c}$	$\hat{c} \cdot 10^3$	$\hat{c} \cdot 10^3$	p -value $H_0 : \hat{c} = \hat{c}$
	S&P 500, real returns					
	Survey Average			Survey Median		
UBS*, >100k, 1 yr, SPF	0.56	-2.93	0.0000	0.44	-2.93	0.0000
UBS*, >100k, 1 yr, Michigan	0.55	-2.93	0.0000	0.43	-2.93	0.0000
UBS*, all, 1 yr, SPF	0.54	-2.93	0.0000	0.45	-2.93	0.0000
UBS*, all, 1 yr, Michigan	0.53	-2.93	0.0000	0.44	-2.93	0.0000
CFO, 1 yr, SPF	0.20	-1.88	0.0004	0.24	-1.74	0.0366
CFO, 1 yr, Michigan	0.26	-1.88	0.0002	0.32	-1.74	0.0252
	Dow Jones, real price growth					
	Survey Average			Survey Median		
Shiller, 1 yr, SPF	0.23	-1.48	0.0000	0.23	-1.48	0.0000
Shiller, 1 yr, Michigan	0.28	-1.48	0.0000	0.29	-1.48	0.0000
Shiller, 10 yrs, SPF	4.11	-6.48	0.0000	5.49	-6.48	0.0000
Shiller, 10 yrs, Michigan	3.51	-6.48	0.0000	4.89	-6.48	0.0000

*stock market return expectations

Table 5: Forecast rationality test using the PD ratio

Parameter	Value	Calibration Target
β^D	1.0048	average quarterly real dividend growth
σ_D	0.0192	std. deviation quarterly real dividend growth
ρ	22	average consumption-dividend ratio
σ_{DW}	$-3.74 \cdot 10^{-4}$	<i>jointly chosen</i> s.t. $corr_t(C_{t+1}/C_t, D_{t+1}/D_t) = 0.2$
σ_W	0.0197	and $std_t(C_{t+1}/C_t) = \frac{1}{7}std_t(D_{t+1}/D_t)$
σ_ε	0.0816	std. deviation of quarterly real stock price growth

Table 6: Model calibration

	U.S. Data	Subj. Beliefs		RE	
	Moment	Moment	t-stat	Moment	t-stat
$E[PD]$	139.7	122.2	0.70	105.5	1.37
$Std[PD]$	65.3	97.3	-2.17	3.94	4.15*
$Corr[PD_t, PD_{t-1}]$	0.98	0.98	0.54	-0.0058	>100*
$Std[r^s]$	8.01	9.44	-3.57*	4.23	9.50*
c	-0.0041	-0.0049	0.67	-0.0126	7.08*
R^2	0.24	0.18	0.47	0.12	0.93
$E[r^s]$	1.89	1.93	-0.09	1.50	0.84
$E[r^b]$	0.13	0.97	-5.10*	1.50	-8.28*
UBS Survey Data:					
$Corr[PD_t, E_t^p R_{t+1}]$	0.79	0.85	-0.79	-0.99	24.86*

* indicates rejection at the 1% level

Table 7: Asset pricing moments



Figure 1: Quarterly PD Ratio of the S&P 500

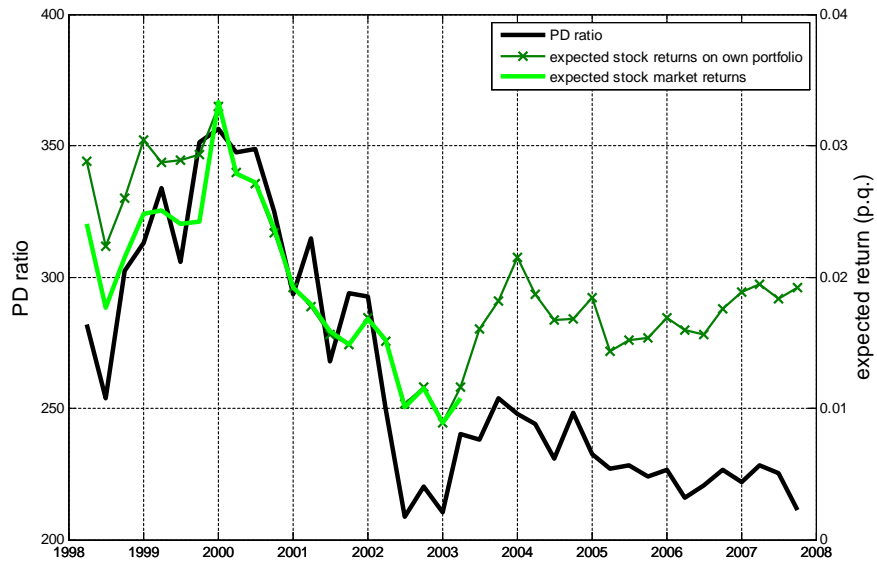


Figure 2: PD ratio and investors' expected returns (UBS Gallup Survey)

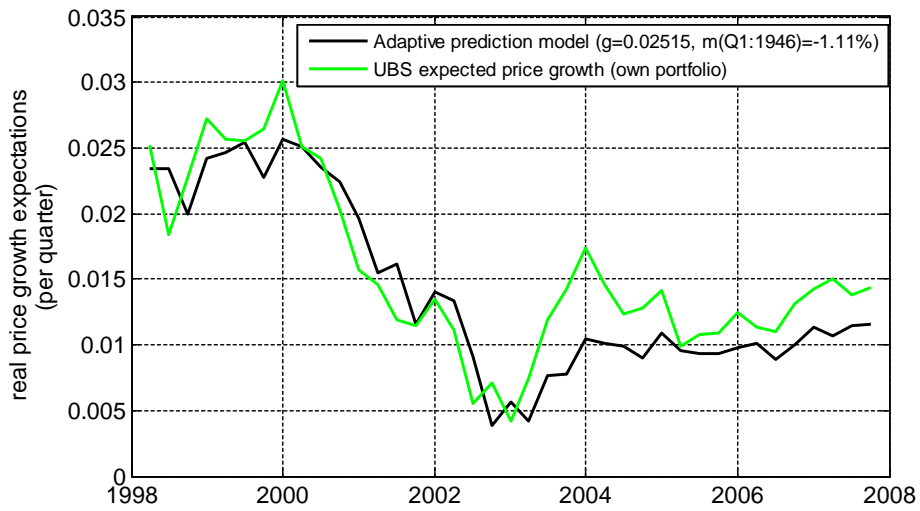


Figure 3: UBS survey expectations versus adaptive prediction model

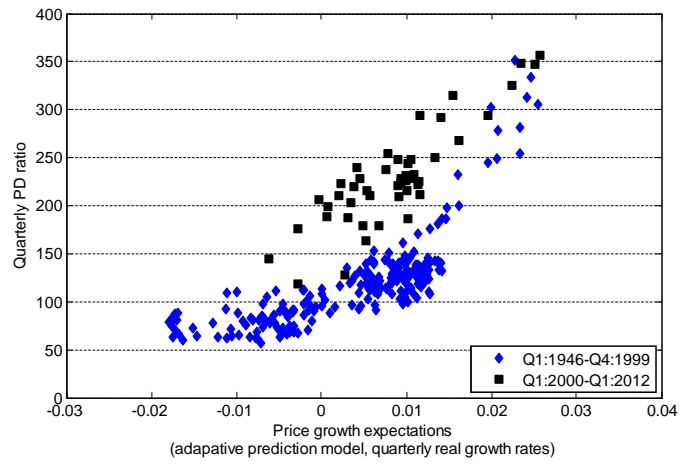


Figure 4: PD ratio S&P 500 vs. adaptive price growth predictions

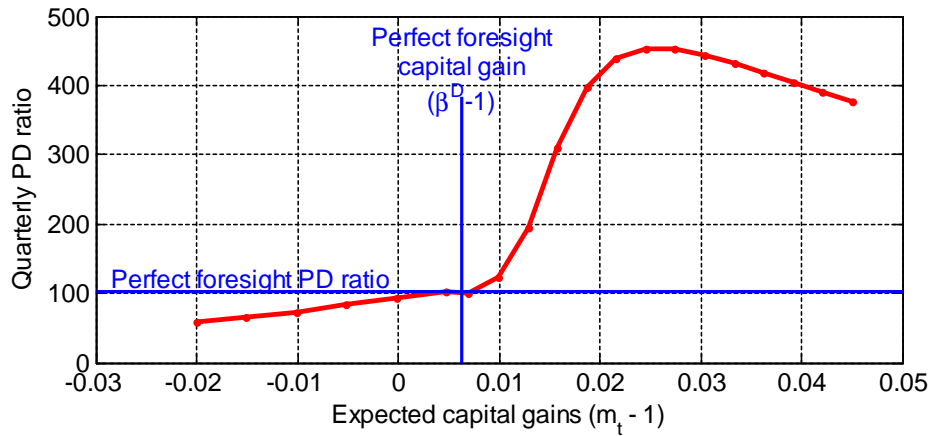


Figure 5: PD ratio and expected capital gains (vanishing noise)

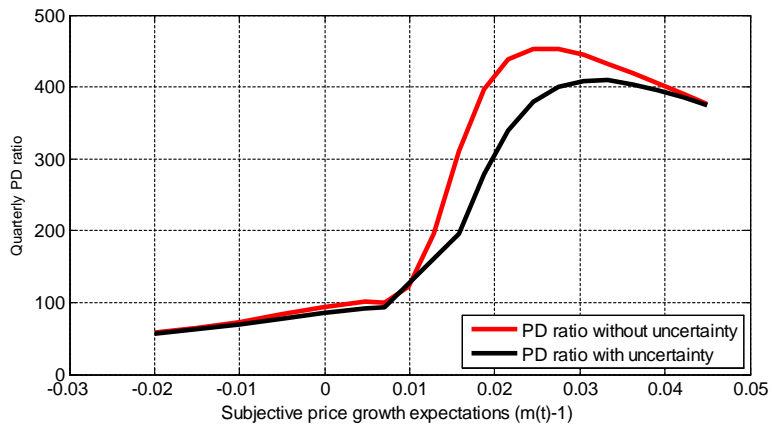


Figure 6: The effects of uncertainty on the equilibrium PD ratio

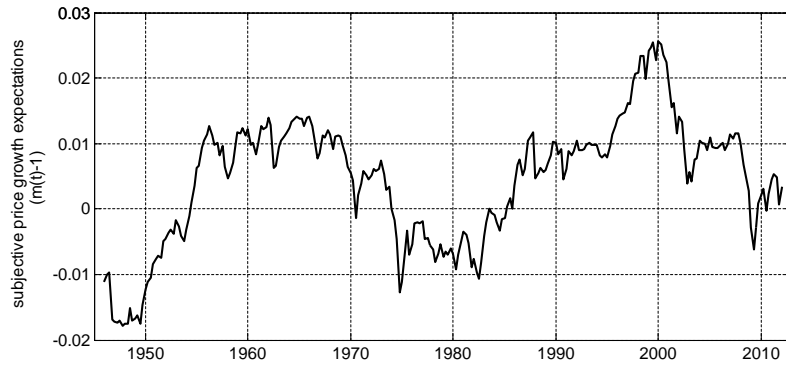


Figure 7: Price growth expectations implied by Bayesian updating and historical price growth information

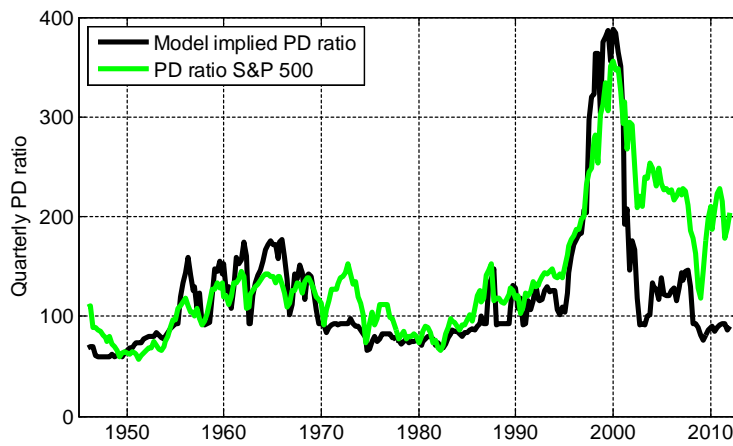


Figure 8: PD ratio - model vs. data

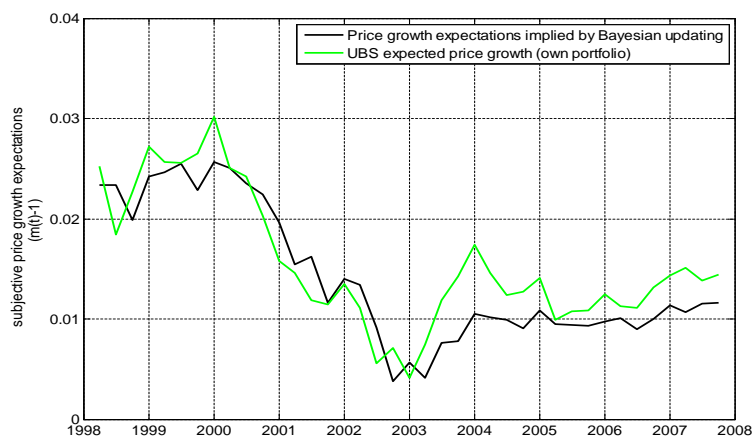


Figure 9: Price growth expectations: UBS survey vs. Bayesian updating model

A Appendix (not for publication)

A.1 Data Sources

Stock price data: our stock price data is for the United States and has been downloaded from ‘The Global Financial Database’ (<http://www.globalfinancialdata.com>). The period covered is Q1:1946-Q1:2012. The nominal stock price series is the ‘SP 500 Composite Price Index (w/GFD extension)’ (Global Fin code ‘_SPXD’). The daily series has been transformed into quarterly data by taking the index value of the last day of the considered quarter. 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. Nominal dividends have been computed as follows

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

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

Stock market survey data: The UBS survey is the UBS Index of Investor Optimism, which is available (against a fee) at

http://www.ropercenter.uconn.edu/data_access/data/datasets/ubs_investor.html.

The quantitative question on stock market expectations has been surveyed over the period Q2:1998-Q4:2007 with 702 responses per month on average and has thereafter been suspended. For each quarter we have data from three monthly surveys, except for the first four quarters and the last quarter of the survey period where we have only one monthly survey per quarter. The Shiller survey data covers individual investors over the period Q1:1999Q1-Q4:2012 and has been kindly made available to us by Robert Shiller at Yale University. On average 73 responses per quarter have been recorded for the question on stock price growth. Since the Shiller data refers to the Dow Jones, we used the PD ratio for the Dow Jones, which is available at <http://www.djaverages.com/>, to compute correlations. The CFO survey is collected by Duke University and CFO magazine and collects responses from U.S. based CFOs over the period Q3:2000-Q4:2012 with on average 390 responses per quarter, available at <http://www.cfosurvey.org/>.

Inflation expectations data: The Survey of Professional Forecasters (SPF) is available from the Federal Reserve Bank of Philadelphia at <http://www.phil.frb.org/research-and-data/real-time-center/survey-of-professional-forecasters/>. The Michigan Surveys of Consumers are collected by Thomson Reuters/University of Michigan

(<http://www.sca.isr.umich.edu/>).

A.2 Details of the t-Test in Section 3.2

Under the RE hypothesis, equations (1) and (3) both hold for the same parameters a^N, c^N , given any horizon N . These two equations define a standard SUR model. Dependent variables are $E_t^P R_{t+N}^N$ and R_{t+N}^N , where the latter is the N -period rate of return

and $E_t^P R_{t+N}^N$ is the observed survey expectation at time t , explanatory variables in both equations are $x_t = (1, \frac{P_t}{D_t})$, satisfying the orthogonality conditions (2)-(4). For expositional clarity we relabel the true parameters in equation (3) as (\bar{a}^N, \bar{c}^N) . The aim is to design efficient estimators of the true parameters $\beta_0^N \equiv (a^N, c^N, \bar{a}^N, \bar{c}^N)$ and to test the hypothesis $H_0 : c^N = \bar{c}^N$.

As is standard in SUR models, without any additional assumption on the distribution of $u, \varepsilon, P/D$, the OLS estimator equation by equation β_T defined by

$$\beta_T \equiv \begin{bmatrix} a_T^N \\ c_T^N \\ \bar{a}_T^N \\ \bar{c}_T^N \end{bmatrix} = \left(\sum_{t=1}^T x_t x_t' \otimes I_2 \right)^{-1} \sum_{t=1}^T x_t \begin{bmatrix} E_t^P R_{t+N}^N \\ R_{t+N}^N \end{bmatrix},$$

where I_2 is a 2×2 identity matrix, is consistent and efficient among the set of estimators using only orthogonality conditions (2)-(4).

To simplify on notation we now drop the superscripts N in the remaining part of this appendix. As is well known, with stationarity, strong ergodicity and bounded second moments, the estimator is consistent and its asymptotic distribution as $T \rightarrow \infty$ is given by

$$\sqrt{T}(\beta_T - \beta_0) \rightarrow N \left(0, [E(x_t x_t') \otimes I_2]^{-1} S_w [E(x_t x_t') \otimes I_2]^{-1} \right), \quad (32)$$

where

$$\begin{aligned} S_w &= \Gamma_0 + \sum_{k=1}^{\infty} \Gamma_k + \Gamma_k' \\ \Gamma_k &= E \left(\begin{bmatrix} u_t \\ u\varepsilon_t \end{bmatrix} [u_{t-k}, u\varepsilon_{t-k}] \otimes x_t x_{t-k}' \right), \end{aligned}$$

where $u\varepsilon_t \equiv u_t + \varepsilon_{t+N}$. To build the test-statistic, we now only need to find an estimator for var-cov matrix in (32).

We can estimate $E(x_t x_t')$ by $\frac{1}{T} \sum_{t=1}^T x_t x_t'$. To estimate the Γ_k terms, we exploit the special form of the error $u\varepsilon_t$. In particular, partition each Γ_k into four 2×2 matrices, with $\Gamma_{ij,k}$ denoting the (i, j) -th element of this partition. Then, letting \hat{u}_t and $\hat{u\varepsilon}_t$ denote the calculated errors of each equation, we use standard estimators

$$\begin{aligned} \Gamma_{11,k,T} &= \frac{1}{T-k} \sum_{t=1}^{T-k} \hat{u}_t \hat{u}_{t-k} x_t x_{t-k}' \\ \Gamma_{12,k,T} &= \frac{1}{T-k} \sum_{t=1}^{T-k} \hat{u}_t \hat{u\varepsilon}_{t-k} x_t x_{t-k}' \end{aligned}$$

Since u_t is not a forecasting error, there is no reason why $\Gamma_{11,k}$ should be zero for any k . We deal with this by using Newey-West weights to truncate the infinite sum in S_w .

Since ε_{t+N} is a forecast error using information up to t we have

$$\Gamma_{21,k} = E([u_t + \varepsilon_{t+N}] u_{t-k} x_t x_{t-k}') = \Gamma_{11,k} \text{ for all } k \geq 0,$$

so estimated $\Gamma_{11,k,T}$ is an estimate of $\Gamma_{21,k,T}$. Furthermore, we have

$$\begin{aligned}\Gamma_{22,k} &= E(u\varepsilon_t u\varepsilon_{t-k} x_t x'_{t-k}) \\ &= \Gamma_{12,k} + E(\varepsilon_{t+N} \varepsilon_{t+N-k} x_t x'_{t-k}) \quad \text{for all } k\end{aligned}$$

where the second equality follows from $E(\varepsilon_{t+N} u_{t-k} x_t x'_{t-k}) = 0$. Moreover, since ε_{t+N} is orthogonal to $\varepsilon_{t+N-k} x_t x'_{t-k}$ for $k \geq N$ we have $\Gamma_{22,k} = \Gamma_{12,k}$ for $k \geq N$. Therefore, we can use the relationship

$$\begin{aligned}\Gamma_{22,k,T} &= \Gamma_{21,k,T} + \frac{1}{T-k} \sum_{t=1}^{T-k} (\widehat{u\varepsilon}_t - \widehat{u}_t) (\widehat{u\varepsilon}_{t-k} - \widehat{u}_{t-k}) x_t x'_{t-k} \quad \text{for } k < N \\ &= \Gamma_{21,k,T} \quad \text{for } k \geq N\end{aligned}$$

which allows using the estimated $\Gamma_{21,k,T}$ as our estimate for $\Gamma_{22,k}$.

A.3 Parameterization of the Wage Process

To calibrate ρ we compute the average dividend-consumption share in the U.S. from 1946-2011, using the ‘Net Corporate Dividends’ and the ‘Personal Consumption Expenditures’ series from the Bureau of Economic Analysis. This delivers an average ratio of $\rho = 22$. Following Campbell and Cochrane (1999) we then choose the standard deviation of one-step-ahead consumption growth innovations to be 1/7 of that of one-step-ahead dividend growth innovations, i.e.,

$$\sqrt{\frac{\text{var}_t(\ln C_{t+1} - \ln C_t)}{\text{var}_t(\ln D_{t+1} - \ln D_t)}} = \frac{1}{7},$$

and the correlation between one-step-ahead consumption and dividend growth to be equal to 0.2, i.e.

$$\frac{\text{cov}_t(\ln C_{t+1} - \ln C_t, \ln D_{t+1} - \ln D_t)}{\sqrt{\text{var}_t(\ln C_{t+1} - \ln C_t) \text{var}_t(\ln D_{t+1} - \ln D_t)}} = 0.2$$

To achieve this we need to compute the required variance and covariances. We have

$$\begin{aligned}\text{var}_t(\ln D_{t+1} - \ln D_t) &= \sigma_D^2 \\ \text{var}_t(\ln C_{t+1} - \ln C_t) &= \text{var}_t(\ln(D_{t+1} + W_{t+1}) - \ln(D_t + W_t)) \\ &= \text{var}_t(\ln(D_{t+1} + \rho D_{t+1} \varepsilon_{t+1}^W)) \\ &= \text{var}_t(\ln D_{t+1} + \ln(1 + \rho \varepsilon_{t+1}^W)) \\ &= \text{var}_t(\ln D_{t+1}) + 2\text{cov}_t(\ln D_{t+1}, \ln(1 + \rho \varepsilon_{t+1}^W)) + \text{var}_t(\ln(1 + \rho \varepsilon_{t+1}^W)) \\ &= \sigma_D^2 + 2\text{cov}_t(\ln \varepsilon_{t+1}^D, \ln(1 + \rho \varepsilon_{t+1}^W)) + \text{var}_t(\ln(1 + \rho \varepsilon_{t+1}^W)) \quad (33)\end{aligned}$$

and

$$\begin{aligned}\text{cov}_t(\ln C_{t+1} - \ln C_t, \ln D_{t+1} - \ln D_t) &= \text{cov}_t(\ln C_{t+1}, \ln \varepsilon_{t+1}^D) \\ &= \text{cov}_t(\ln(D_{t+1} + W_{t+1}), \ln \varepsilon_{t+1}^D) \\ &= \text{cov}_t(\ln D_{t+1} + \ln(1 + \rho \varepsilon_{t+1}^W), \ln \varepsilon_{t+1}^D) \\ &= \text{cov}_t(\ln \varepsilon_{t+1}^D + \ln(1 + \rho \varepsilon_{t+1}^W), \ln \varepsilon_{t+1}^D) \\ &= \sigma_D^2 + \text{cov}_t(\ln(1 + \rho \varepsilon_{t+1}^W), \ln \varepsilon_{t+1}^D) \quad (34)\end{aligned}$$

Linearly approximating $\ln(1 + \rho\varepsilon_{t+1}^W)$ around the unconditional mean $\varepsilon^W = 1$ delivers

$$\ln(1 + \rho\varepsilon_{t+1}^W) \approx c + \frac{\rho}{1 + \rho} \ln \varepsilon_{t+1}^W + O(2)$$

where c is a constant and $O(2)$ a second order approximation error. Using this approximation we have

$$\text{var}_t(\ln C_{t+1} - \ln C_t) \approx \sigma_D^2 + 2\frac{\rho}{1 + \rho}\sigma_{DW} + \left(\frac{\rho}{1 + \rho}\right)^2 \sigma_W^2 \quad (35)$$

So that

$$\sqrt{\frac{\text{var}_t(\ln C_{t+1} - \ln C_t)}{\text{var}_t(\ln D_{t+1} - \ln D_t)}} \approx \sqrt{1 + 2\frac{\rho}{1 + \rho}\frac{\sigma_{DW}}{\sigma_D^2} + \left(\frac{\rho}{1 + \rho}\right)^2 \frac{\sigma_W^2}{\sigma_D^2}} = \frac{1}{7} \quad (36)$$

Using the approximation we also have

$$\begin{aligned} \frac{\text{cov}_t(\ln C_{t+1} - \ln C_t, \ln D_{t+1} - \ln D_t)}{\sqrt{\text{var}_t(\ln C_{t+1} - \ln C_t)\text{var}_t(\ln D_{t+1} - \ln D_t)}} &\approx \\ \frac{\sigma_D^2 + \frac{\rho}{1 + \rho}\sigma_{WD}}{\sqrt{\left(\sigma_D^2 + 2\frac{\rho}{1 + \rho}\sigma_{WD} + \left(\frac{\rho}{1 + \rho}\right)^2 \sigma_W^2\right)\sigma_D^2}} &= 0.2 \end{aligned} \quad (37)$$

Using (36) to substitute the root in the denominator in (37) we get

$$\frac{\sigma_D^2 + \frac{\rho}{1 + \rho}\sigma_{WD}}{\frac{1}{7}\sigma_D^2} = 0.2 \iff \sigma_{WD} = -\frac{68}{70}\frac{1 + \rho}{\rho}\sigma_D^2 \quad (38)$$

Using (36) we then get

$$\begin{aligned} \sigma_W^2 &= -\frac{48}{49}\left(\frac{1 + \rho}{\rho}\right)^2 \sigma_D^2 - 2\frac{1 + \rho}{\rho}\sigma_{WD} \\ &= \frac{236}{245}\left(\frac{1 + \rho}{\rho}\right)^2 \sigma_D^2. \end{aligned} \quad (39)$$

A.4 Existence of Optimum, Sufficiency of FOCs, Recursive Solution

Existence of Optimum & Sufficiency of FOCs. The choice set in (6) is compact and non-empty. The following condition then insures existence of optimal plans:

Condition 1 The utility function $u(\cdot)$ is bounded above and for all $i \in [0, 1]$

$$E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t u(W_t + D_t) > -\infty. \quad (40)$$

The expression on the left-hand side of condition (40) is the utility associated with never trading stocks ($S_t^i = 1$ for all t). Since this policy is always feasible, condition

(40) guarantees that the objective function in (6) is also bounded from below, even if the flow utility function $u(\cdot)$ is itself unbounded below. The optimization problem (6) thus maximizes a bounded continuous utility function over a compact set, which guarantees existence of a maximum.

Under the assumptions made in the main text (utility function given by (10), knowledge of (7) and $\delta\beta^{RE} < 1$), condition 1 holds, as can be seen from the following derivation:

$$\begin{aligned}
& E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t u(W_t + D_t) \\
&= E_0 \sum_{t=0}^{\infty} \delta^t u(W_t + D_t) \\
&= E_0 \sum_{t=0}^{\infty} \delta^t ((1 + \rho\varepsilon_t^W)D_t)^{1-\gamma} \\
&= ((1 + \rho\varepsilon_0^W)D_0)^{1-\gamma} + E_0 \sum_{t=1}^{\infty} \delta^t ((1 + \rho\varepsilon_t^W)D_t)^{1-\gamma} \\
&= ((1 + \rho\varepsilon_0^W)D_0)^{1-\gamma} + E_0 \sum_{t=1}^{\infty} \left(\delta (\beta^D)^{1-\gamma} \right)^t \left((1 + \rho\varepsilon_t^W) \varepsilon_t^D \prod_{k=1}^{t-1} \varepsilon_k^D \right)^{1-\gamma} \\
&= ((1 + \rho\varepsilon_0^W)D_0)^{1-\gamma} + E \left[((1 + \rho\varepsilon^W)\varepsilon^D)^{1-\gamma} \right] \cdot E_0 \sum_{t=1}^{\infty} \left(\delta (\beta^D)^{1-\gamma} \right)^t \left(\prod_{k=1}^{t-1} \varepsilon_k^D \right)^{1-\gamma} \\
&= ((1 + \rho\varepsilon_0^W)D_0)^{1-\gamma} + E \left[((1 + \rho\varepsilon^W)\varepsilon^D)^{1-\gamma} \right] \cdot \sum_{t=1}^{\infty} \left(\delta (\beta^D)^{1-\gamma} \right)^t \left(e^{\frac{\sigma_D^2}{2}\gamma(\gamma-1)} \right)^{t-1} \\
&= ((1 + \rho\varepsilon_0^W)D_0)^{1-\gamma} + \frac{E \left[((1 + \rho\varepsilon^W)\varepsilon^D)^{1-\gamma} \right]}{e^{\frac{\sigma_D^2}{2}\gamma(\gamma-1)}} \cdot \sum_{t=1}^{\infty} (\delta\beta^{RE})^t
\end{aligned}$$

Since (6) is a strictly concave maximization problem the maximum is unique. With the utility function being differentiable, the first order conditions

$$u'(C_t^i) = \delta E_t^{\mathcal{P}^i} \left[u'(C_{t+1}^i) \frac{P_{t+1} + D_{t+1}}{P_t} \right] \quad (41)$$

plus a standard transversality condition are necessary and sufficient for the optimum.

Recursive Solution. We have a recursive solution whenever the optimal stockholding policy can be written as a time-invariant function $S_t^i = S^i(x_t)$ of some state variables x_t . We seek a recursive solution where x_t contains appropriately rescaled variables that do not grow to infinity. With this in mind, we impose the following condition:

Condition 2 The flow utility function $u(\cdot)$ is homogeneous of degree $\eta \geq 0$. Furthermore, the beliefs \mathcal{P}^i imply that $\theta_t \equiv \left(\frac{D_t}{D_{t-1}}, \frac{P_t}{D_t}, \frac{W_t}{D_t} \right)$ has a state space representation,

i.e., the conditional distribution $\mathcal{P}^i(\theta_{t+1}|\omega^t)$ can be written as

$$\mathcal{P}^i(\theta_{t+1}|\omega^t) = \mathcal{F}^i(m_t^i) \quad (42)$$

$$m_t^i = \mathcal{R}^i(m_{t-1}^i, \theta_t) \quad (43)$$

for some finite-dimensional state vector $\ln m_t^i$ and some time-invariant functions \mathcal{F}^i and \mathcal{R}^i .

Under Condition 2 problem (6) can then be re-expressed as

$$\max_{\{S_t^i \in \mathcal{S}\}_{t=0}^{\infty}} E_0^{\mathcal{P}^i} \sum_{t=0}^{\infty} \delta^t \mathcal{D}_t u \left(S_{t-1}^i \left(\frac{P_t}{D_t} + 1 \right) - S_t^i \frac{P_t}{D_t} + \frac{W_t}{D_t} \right), \quad (44)$$

given $S_{-1}^i = 1$, where \mathcal{D}_t is a time-varying discount factor satisfying $\mathcal{D}_{-1} = 1$ and

$$\mathcal{D}_t = \mathcal{D}_{t-1} (\beta^D \varepsilon_t^D)^\eta.$$

The return function in (44) depends only on the exogenous variables contained in the vector θ_t . Since the beliefs \mathcal{P}^i are assumed to be recursive in θ_t , standard arguments in dynamic programming guarantee that the optimal solution to (44) takes the form (12). This formulation of the recursive solution is useful, because scaling P_t and W_t by the level of dividends eliminates the trend in these variables, as desired. This will be useful when computing numerical approximations to $S^i(\cdot)$. The belief systems \mathcal{P}^i introduced in section 6 will satisfy the requirements stated in condition 2.

A.5 Proof of Proposition 1

In equilibrium $S_t^i = 1$ for all $t \geq 0$, so that the budget constraint implies

$$C_t^i = D_t + W_t = (1 + \rho \varepsilon_t^W) D_t.$$

Substituting into the agent's first order condition delivers

$$P_t = \delta E_t \left[\left(\frac{(1 + \rho \varepsilon_{t+1}^W) D_{t+1}}{(1 + \rho \varepsilon_t^W) D_t} \right)^{-\gamma} (P_{t+1} + D_{t+1}) \right]. \quad (45)$$

Assuming that the following transversality condition holds

$$\lim_{j \rightarrow \infty} E_t \left[\delta^j \left\{ \left(\frac{1 + \rho \varepsilon_{t+j}^W}{1 + \rho \varepsilon_t^W} \right) \frac{D_{t+j}}{D_t} \right\}^{-\gamma} P_{t+j} \right] = 0, \quad (46)$$

one can iterate forward on (45) to obtain

$$\frac{P_t}{D_t} = E_t \left[\sum_{j=1}^{\infty} \delta^j \left(\frac{1 + \rho \varepsilon_{t+j}^W}{1 + \rho \varepsilon_t^W} \right)^{-\gamma} \left(\frac{D_{t+j}}{D_t} \right)^{1-\gamma} \right],$$

Using $D_{t+j}/D_t = (\beta^D)^j \prod_{k=1}^j \epsilon_{t+k}^D$ one has

$$\begin{aligned}
\frac{P_t}{D_t} &= (1 + \rho \varepsilon_t^W)^\gamma \sum_{j=1}^{\infty} (\delta(\beta^D)^{1-\gamma})^j E_t \left[(1 + \rho \varepsilon_{t+j}^W)^{-\gamma} \left(\prod_{k=1}^j \epsilon_{t+k}^D \right)^{1-\gamma} \right] \\
&= (1 + \rho \varepsilon_t^W)^\gamma \sum_{j=1}^{\infty} (\delta(\beta^D)^{1-\gamma})^j E_t \left[(1 + \rho \varepsilon_{t+j}^W)^{-\gamma} (\epsilon_{t+j}^D)^{1-\gamma} \right] E_t \left[\left(\prod_{k=1}^{j-1} \epsilon_{t+k}^D \right)^{1-\gamma} \right] \\
&= (1 + \rho \varepsilon_t^W)^\gamma E_t \left[(1 + \rho \varepsilon_{t+j}^W)^{-\gamma} (\epsilon_{t+j}^D)^{1-\gamma} \right] e^{\gamma(1-\gamma)\sigma_D^2/2} \frac{\delta \beta^{RE}}{1 - \delta \beta^{RE}},
\end{aligned}$$

as claimed in proposition 1.

A.6 Bayesian Foundations for Lagged Belief Updating

We now present a slightly modified information structure for which Bayesian updating gives rise to the lagged belief updating equation (28). Specifically, we generalize the perceived price process (16) by splitting the temporary return innovation $\ln \varepsilon_{t+1}$ into two independent subcomponents:

$$\ln P_{t+1} - \ln P_t = \ln \beta_{t+1} + \ln \varepsilon_{t+2}^1 + \ln \varepsilon_{t+1}^2$$

with $\ln \varepsilon_{t+2}^1 \sim iiN(-\frac{\sigma_{\varepsilon_1}^2}{2}, \sigma_{\varepsilon_1}^2)$, $\ln \varepsilon_{t+1}^2 \sim iiN(-\frac{\sigma_{\varepsilon_2}^2}{2}, \sigma_{\varepsilon_2}^2)$ and

$$\sigma_\varepsilon^2 = \sigma_{\varepsilon_1}^2 + \sigma_{\varepsilon_2}^2.$$

We then assume that in any period t agents observe the prices, dividends and wages up to period t , as well as the innovations ε_t^1 up to period t . Agents' time t information set thus consists of $I_t = \{P_t, D_t, W_t, \varepsilon_t^1, P_{t-1}, D_{t-1}, W_{t-1}, \varepsilon_{t-2}^1, \dots\}$. By observing the innovations ε_t^1 , agents learn - with a one period lag - something about the temporary components of price growth. The process for the persistent price growth component $\ln \beta_t$ remains as stated in equation (17), but we now denote the innovation variance by σ_v^2 instead of σ_ε^2 . As before, $\ln m_t$ denotes the posterior mean of $\ln \beta_t$ given the information available at time t . We prove below the following result:

Proposition 3 Fix $\sigma_\varepsilon^2 > 0$ and consider the limit $\sigma_{\varepsilon_2}^2 \rightarrow 0$ with $\sigma_v^2 = \sigma_{\varepsilon_2}^2 g^2 / (1 - g)$. Bayesian updating then implies

$$\ln m_t = \ln m_{t-1} - \frac{\sigma_v^2}{2} + g \left(\ln P_{t-1} - \ln P_{t-1} + \frac{\sigma_\varepsilon^2 + \sigma_v^2}{2} - \ln m_{t-1} \right) - g \ln \varepsilon_t^1 \quad (47)$$

The modified information structure thus implies that only lagged price growth rates enter the current state estimate, so that beliefs are predetermined, precisely as assumed in equation (28). Intuitively, this is so because lagged returns become infinitely more informative relative to current returns as $\sigma_{\varepsilon_2}^2 \rightarrow 0$, which eliminates the simultaneity problem. For non-vanishing uncertainty $\sigma_{\varepsilon_2}^2$ the weight of the last observation actually remains positive but would still be lower than that given to the lagged return observation, see equation (50) in the proof below and the subsequent discussion for details.

We now sketch the proof of the previous proposition. Let us define the following augmented information set $\tilde{I}_{t-1} = I_{t-1} \cup \{\varepsilon_t^1\}$. The posterior mean for β_t given \tilde{I}_{t-1} , denoted $\ln m_{t|\tilde{I}_{t-1}}$ is readily recursively determined via

$$\ln m_{t|\tilde{I}_{t-1}} = \ln m_{t-1|\tilde{I}_{t-2}} - \frac{\sigma_v^2}{2} + \tilde{g} \left(\ln P_{t-1} - \ln P_{t-2} - \ln \varepsilon_t^1 + \frac{\sigma_v^2 + \sigma_{\varepsilon 2}^2}{2} - \ln m_{t-1|t-1} \right) \quad (48)$$

and the steady state posterior uncertainty and the Kalman gain by

$$\begin{aligned} \sigma^2 &= \frac{-\sigma_v^2 + \sqrt{(\sigma_v^2)^2 + 4\sigma_v^2\sigma_{\varepsilon 2}^2}}{2} \\ \tilde{g} &= \frac{\sigma^2}{\sigma_{\varepsilon 2}^2} \end{aligned} \quad (49)$$

Standard updating formulas for normal distributions then imply that the posterior mean of $\ln \beta_t$ using information set I_t can be derived by updating the posterior mean based on \tilde{I}_{t-1} according to

$$\ln m_{t|I_t} = \ln m_{t|\tilde{I}_{t-1}} + \frac{\sigma^2}{\sigma^2 + \sigma_{\varepsilon 1}^2 + \sigma_{\varepsilon 2}^2 + \sigma_v^2} \left(\ln P_t - \ln P_{t-1} + \frac{\sigma_{\varepsilon 1}^2 + \sigma_{\varepsilon 2}^2 + \sigma_v^2}{2} - \ln m_{t|\tilde{I}_{t-1}} \right) \quad (50)$$

Since $\frac{\sigma^2}{\sigma^2 + \sigma_{\varepsilon 1}^2 + \sigma_{\varepsilon 2}^2 + \sigma_v^2} < \frac{\sigma^2}{\sigma_{\varepsilon 2}^2} = \tilde{g}$, the weight of the price observation dated t is reduced relative to the earlier observation dated $t-1$ because it is ‘noisier’. Now consider the limit $\sigma_{\varepsilon 2}^2 \rightarrow 0$ and along the limit choose $\sigma_{\varepsilon 1}^2 = \sigma_{\varepsilon}^2 - \sigma_{\varepsilon 2}^2$ and $\sigma_v^2 = \frac{g^2}{1-g} \sigma_{\varepsilon 2}^2$, as assumed in the proposition. Equation (50) then implies that $\ln m_{t|I_t} = \ln m_{t|\tilde{I}_{t-1}}$, i.e., the weight of the last observation price converges to zero. Moreover, from $\sigma_v^2 = \frac{g^2}{1-g} \sigma_{\varepsilon 2}^2$ and (49) we get $\tilde{g} = g$. Using these results, equation (48) can exactly be written as stated by equation (47) in the main text.

A.7 Proof of Proposition 2

The proof relies on the fact that in a situation without uncertainty the expectation of a non-linear function of ‘random’ variables is identical to the non-linear function of the expectation of these random variables, i.e., for some continuous non-linear function $f(\cdot, \cdot)$ and some random variables X_{t+j}, Y_{t+j} we have under the stated assumptions $E_t^{\mathcal{P}} f(X_{t+j}, Y_{t+j}) = f(E_t^{\mathcal{P}} X_{t+j}, E_t^{\mathcal{P}} Y_{t+j})$. Simplifying notation (and slightly abusing it) we let $X_{t+j} = E_t^{\mathcal{P}} X_{t+j}$ for all $j \geq 1$, so that X_{t+j} below denotes the subjective expectation conditional on information at time t of the variable X at time $t+j$. The first order conditions (41) can then be written as

$$\begin{aligned} 1 &= \left(\frac{C_{t+1+j}}{C_{t+j}} \right)^{-\gamma} \delta R_{t+1+j} \iff \\ \frac{C_{t+1+j}}{P_{t+1+j} + D_{t+1+j}} &= \delta^{\frac{1}{\gamma}} (R_{t+1+j})^{\frac{1-\gamma}{\gamma}} \frac{C_{t+j}}{P_{t+j}} \end{aligned} \quad (51)$$

for all $j \geq 0$. The budget constraint implies

$$S_{t-1}(P_t + D_t) = C_t - W_t + S_t P_t \implies S_{t-1} = \frac{C_t - W_t}{P_t + D_t} + \frac{P_t}{P_t + D_t} S_t$$

Iterating forward on the latter equation gives

$$S_{t-1} = \frac{C_t - W_t}{P_t + D_t} + \frac{P_t}{P_t + D_t} \frac{C_{t+1} - W_{t+1}}{P_{t+1} + D_{t+1}} + \frac{P_t}{P_t + D_t} \frac{P_{t+1}}{P_{t+1} + D_{t+1}} \frac{C_{t+2} - W_{t+2}}{P_{t+2} + D_{t+2}} + \dots$$

Repeatedly using equation (51) gives

$$\begin{aligned} S_{t-1} &= \frac{C_t - W_t}{P_t + D_t} + \frac{P_t}{P_t + D_t} \left(\delta^{\frac{1}{\gamma}} (R_{t+1})^{\frac{1-\gamma}{\gamma}} \frac{C_t}{P_t} - \frac{W_{t+1}}{P_{t+1} + D_{t+1}} \right) \\ &\quad + \frac{P_t}{P_t + D_t} \frac{P_{t+1}}{P_{t+1} + D_{t+1}} \left(\delta^{\frac{1}{\gamma}} (R_{t+2})^{\frac{1-\gamma}{\gamma}} \frac{C_{t+1}}{P_{t+1}} - \frac{W_{t+2}}{P_{t+2} + D_{t+2}} \right) + \dots \\ &= \frac{C_t}{P_t + D_t} + \delta^{\frac{1}{\gamma}} (R_{t+1})^{\frac{1-\gamma}{\gamma}} \frac{C_t}{P_t + D_t} \\ &\quad + \frac{P_t}{P_t + D_t} \delta^{\frac{1}{\gamma}} (R_{t+2})^{\frac{1-\gamma}{\gamma}} \frac{C_{t+1}}{P_{t+1} + D_{t+1}} + \dots \\ &\quad - \frac{W_t}{P_t + D_t} - \frac{W_{t+1}}{P_t + D_t} \frac{1}{R_{t+1}} - \frac{W_{t+2}}{P_t + D_t} \frac{1}{R_{t+1} R_{t+2}} - \dots \\ &= \frac{C_t}{P_t + D_t} + \delta^{\frac{1}{\gamma}} (R_{t+1})^{\frac{1-\gamma}{\gamma}} \frac{C_t}{P_t + D_t} \\ &\quad + \left(\delta^{\frac{1}{\gamma}} \right)^2 (R_{t+2} R_{t+1})^{\frac{1-\gamma}{\gamma}} \frac{C_t}{P_t + D_t} + \dots \\ &\quad - \frac{1}{P_t + D_t} \left(\sum_{j=0}^{\infty} W_{t+j} \prod_{i=1}^j \frac{1}{R_{t+i}} \right) \\ &= \frac{D_t}{P_t + D_t} + \frac{C_t}{P_t + D_t} \sum_{j=1}^{\infty} \left(\left(\delta^{\frac{1}{\gamma}} \right)^j \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right)^{\frac{\gamma-1}{\gamma}} \right) \\ &\quad - \frac{1}{P_t + D_t} \left(\sum_{j=1}^{\infty} \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right) W_{t+j} \right) \end{aligned} \tag{52}$$

Imposing on the previous equation $S_{t-1} = 1$ (the market clearing condition for period $t - 1$ if $t > 1$, or the initial condition for period $t = 0$) and $C_t = D_t + W_t$ (the market clearing condition for period $t \geq 0$) one obtains the result stated in the proposition under the convention that $R_{t+i} = E_t^P R_{t+i}$.

A.8 Verification of Conditions (29)

For the vanishing noise limit of the beliefs specified in section 6 we have

$$\begin{aligned} E_t^P [P_{t+j}] &= (m_t)^j P_t \\ E_t^P [D_{t+j}] &= (\beta^D)^j D_t \\ E_t^P [W_{t+j}] &= (\beta^D)^j W_t. \end{aligned}$$

We first verify the inequality on the l.h.s. of equation (29). We have

$$\lim_{T \rightarrow \infty} E_t^{\mathcal{P}}[R_T] = m_t + \lim_{T \rightarrow \infty} \left(\frac{\beta^D}{m_t} \right)^{T-1} \beta^D \frac{D_t}{P_t},$$

so that for $m_t > 1$ the limit clearly satisfies $\lim_{T \rightarrow \infty} E_t^{\mathcal{P}}[R_T] > 1$ due to the first term on the r.h.s.; for $m_t < 1$ the second term on the r.h.s. increases without bound, due to $\beta^D > 1$, so that $\lim_{T \rightarrow \infty} E_t^{\mathcal{P}}[R_T] > 1$ also holds.

In a second step we verify that the inequality condition on the r.h.s. of equation (29) holds for all subjective beliefs $m_t > 0$. We have

$$\begin{aligned} \lim_{T \rightarrow \infty} E_t^{\mathcal{P}} \left(\sum_{j=1}^T \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right) W_{t+j} \right) &= \lim_{T \rightarrow \infty} W_t E_t^{\mathcal{P}} \left(\sum_{j=1}^T (\beta^D)^j \left(\prod_{i=1}^j \frac{1}{R_{t+i}} \right) \right) \\ &= \lim_{T \rightarrow \infty} W_t \sum_{j=1}^T X_j \end{aligned} \quad (53)$$

where

$$X_j = \frac{(\beta^D)^j}{\prod_{i=1}^j \left(m_t + \left(\frac{\beta^D}{m_t} \right)^{i-1} \beta^D \frac{D_t}{P_t} \right)} \geq 0 \quad (54)$$

A sufficient condition for the infinite sum in (53) to converge is that the terms X_j are bounded by some exponentially decaying function. The denominator in (54) satisfies

$$\begin{aligned} &\prod_{i=1}^j \left(m_t + \left(\frac{\beta^D}{m_t} \right)^{i-1} \beta^D \frac{D_t}{P_t} \right) \\ &\geq (m_t)^j + \left(\frac{\beta^D}{m_t} \right)^{j \left(\frac{j-1}{2} \right)} \beta^D \frac{D_t}{P_t}, \end{aligned} \quad (55)$$

where the first term captures the the pure products in m_t , the second term the pure products in $\left(\frac{\beta^D}{m_t} \right)^{i-1} \beta^D \frac{D_t}{P_t}$, and all cross terms have been dropped. We then have

$$\begin{aligned} X_j &= \frac{(\beta^D)^j}{\prod_{i=1}^j \left(m_t + \left(\frac{\beta^D}{m_t} \right)^{i-1} \beta^D \frac{D_t}{P_t} \right)} \\ &\leq \frac{(\beta^D)^j}{(m_t)^j + \left(\frac{\beta^D}{m_t} \right)^{j \left(\frac{j-1}{2} \right)} \beta^D \frac{D_t}{P_t}} \\ &= \frac{1}{\left(\frac{m_t}{\beta^D} \right)^j + \left(\frac{\beta^D}{m_t} \right)^{j \left(\frac{j-1}{2} \right)} \frac{1}{(\beta^D)^{j-1}} \frac{D_t}{P_t}}, \end{aligned}$$

where all terms in the denominator are positive. For $m_t \geq \beta^D > 1$ we can use the first

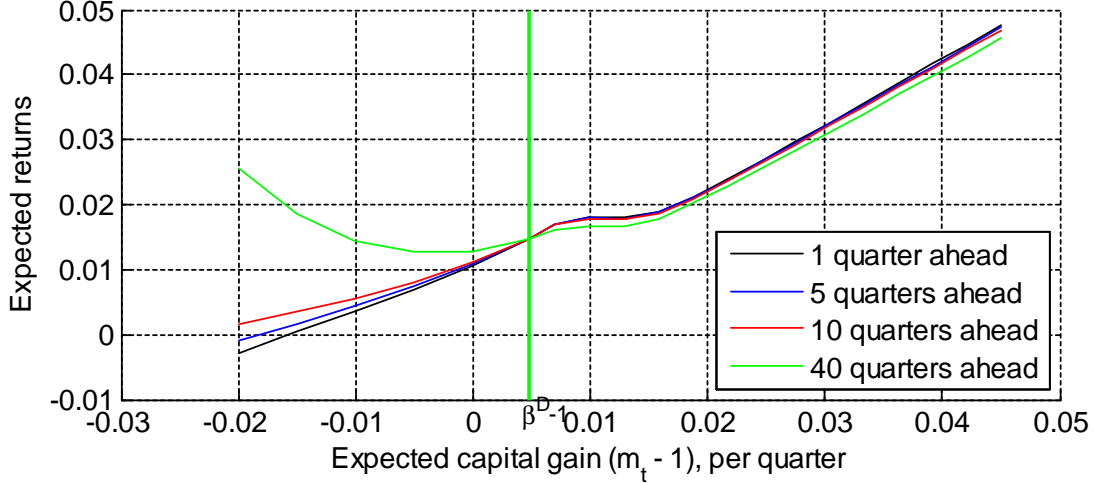


Figure 10: Expected return as a function of expected capital gain

term in the denominator to exponentially bound X_j , as $X_j \leq \left(\frac{\beta^D}{m_t}\right)^j$; for $m_t < \beta^D$ we can use the second term:

$$X_j \leq \frac{1}{\left(\frac{\beta^D}{m_t}\right)^{j\left(\frac{j-1}{2}\right)} \frac{1}{(\beta^D)^{j-1}} \frac{D_t}{P_t}} = \frac{1}{\left(\left(\frac{\beta^D}{m_t}\right)^{\frac{j}{2}} \frac{1}{\beta^D}\right)^{j-1} \frac{D_t}{P_t}}$$

Since $m_t < \beta^D$ there must be a $J < \infty$ such that

$$\left(\frac{\beta^D}{m_t}\right)^{\frac{j}{2}} \frac{1}{\beta^D} \geq \frac{\beta^D}{m_t} > 1$$

for all $j \geq J$, so that the X_j are exponentially bounded for all $j \geq J$.

A.9 Capital Gains Expectations and Expected Returns: Further Details

Figure 10 depicts how expected returns at various horizons depend on agent's expected price growth expectations using the same parameterization as used in figure 5. It shows that expected returns covary positively with capital gains expectations for $m_t \geq \beta^D$, as has been claimed in the main text. The flatish part at around $m_t - 1 \approx 0.01$ arises because in that area the PD ratio increases strongly, so that the dividend yield falls. Only for pessimistic price growth expectations ($m_t < \beta^D$) and long horizons of expected returns we find a negative relationship. The latter emerges because with prices expected to fall, the dividend yield will rise and eventually result in high return expectations.

A.10 Numerical Solution Algorithm

Algorithm: We solve for agents' state-contingent, time-invariant stockholdings (and consumption) policy (12) using time iteration in combination with the method of en-

ogenous grid points. Time iteration is a computationally efficient, e.g., Aruoba et al. (2006), and convergent solution algorithm, see Rendahl (2013). The method of endogenous grid points, see Carroll (2006), economizes on a costly root finding step which speeds up computations further.

Evaluations of Expectations: Importantly, agents evaluate the expectations in the first order condition (41) according to their subjective beliefs about future price growth and their (objective) beliefs about the exogenous dividend and wage processes. Expectations are approximated via Hermite Gaussian quadrature using three interpolation nodes for the exogenous innovations.

Approximation of Optimal Policy Functions: The consumption/stockholding policy is approximated by piecewise linear splines, which preserves the nonlinearities arising in particular in the PD dimension of the state space. Once the state-contingent consumption policy has been found, we use the market clearing condition for consumption goods to determine the market clearing PD ratio for each price-growth belief m_t .

Accuracy: Carefully choosing appropriate grids for each belief is crucial for the accuracy of the numerical solution. We achieve maximum (relative) Euler errors on the order of 10^{-3} and median Euler errors on the order of 10^{-5} (average: 10^{-4}).

Using our analytical solution for the case with vanishing noise, we can assess the accuracy of our solution algorithm more directly. Setting the standard deviations of exogenous disturbances to 10^{-16} the algorithm almost perfectly recovers the equilibrium PD ratio of the analytical solution: the error for the numerically computed equilibrium PD ratio for any price growth belief m_t on our grid is within 0.5 % of the analytical solution.

References

- ADAM, K., AND A. MARCET (2011): “Internal Rationality, Imperfect Market Knowledge and Asset Prices,” *Journal of Economic Theory*, 146, 1224–1252.
- ADAM, K., A. MARCET, AND J. P. NICOLINI (2013): “Stock Market Volatility and Learning,” University of Mannheim Mimeo.
- ARUOBA, B., J. FERNANDEZ-VILLAVARDE, AND J. RUBIO-RAMIREZ (2006): “Comparing solution methods for dynamic equilibrium economies,” *Journal of Economic Dynamics and Control*, 30.
- ASPAROUHOVA, E., P. BOSSAERTS, N. ROY, AND W. ZAME (2013): “Experiments on the Lucas Asset Pricing Model,” *NBER Macroeconomics Annual 2013 (forthcoming)*.
- BACCHETTA, P., E. MERTENS, AND E. VAN WINCOOP (2009): “Predictability in Financial Markets: What Do Survey Expectations Tell Us?,” *Journal of International Money and Finance*, 28, 406–426.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59, 1481–1509.
- 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. A., AND G. EVANS (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3, 159–191.
- 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.
- CARROLL, C. D. (2006): “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems,” *Economics Letters*, 91, 312–320.
- 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.
- DELONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. J. WALDMANN (1990): “Positive Feedback Investment Strategies and Destabilizing Rational Speculation,” *Journal of Finance*, 45, 379–395.
- EUSEPI, S., AND B. PRESTON (2011): “Expectations, Learning, and Business Cycle Fluctuations,” *American Economic Review*, 101, 2844–2872.
- (2013): “Fiscal Foundations of Inflation: Imperfect Knowledge,” Monash University mimeo.
- FAMA, E. F., AND K. R. FRENCH (1988): “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, 22, 3–25.
- FUSTER, A., B. HERBERT, AND D. LAIBSON (2011): “Natural Expectations, Macroeconomic Dynamics and Asset Prices,” in *NBER Macroeconomics Annual 2011*. NBER.
- HASSAN, T. A., AND T. M. MERTENS (2010): “The Social Cost of Near-Rational Investment,” NYU Working Paper 2451/29843.
- 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.

- LUCAS, R. E. (1978): “Asset Prices in an Exchange Economy,” *Econometrica*, 46, 1426–1445.
- MALMENDIER, U., AND S. NAGEL (2011): “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking,” *Quarterly Journal of Economics*, 126, 373–416.
- (2013): “Learning from Inflation Experiences,” Berkeley University Working Paper.
- MARCEY, 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.
- NAGEL, S., AND R. GREENWOOD (2009): “Inexperienced Investors and Bubbles,” *Journal of Financial Economics*, 93(2), 239–258.
- RENDAHL, P. (2013): “Inequality Constraint and Euler Equation Based Solution Methods,” *Economic Journal* (*forthcoming*).
- ROY, R. (1989): “Asymptotic covariance structure of serial correlations in multivariate time series,” *Biometrika*, 76, 824–7.
- TIMMERMANN, A. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1996): “Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.
- VISSING-JORGENSEN, A. (2003): “Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions,” in *2003 Macroeconomics Annual*, Boston. NBER.
- WEST, M., AND J. HARRISON (1997): *Bayesian Forecasting and Dynamic Models*. Springer, Berlin.