Asset prices under recency-biased learning

Pauline Gandré[*]

January 2017

Abstract: This paper presents a consumption-based asset pricing model in which fluctuations in asset prices are persistently driven by time-varying expectations because of learning on the fundamental process. Agents estimate the actual parameter of the fundamental process as bayesian econometricians, except that they weight the precision of recent observations more heavily relative to earlier ones (the so-called recency bias), due to cognitive limitations when processing sequential information over time. The stylized model proves able to replicate the boom-and-bust episode on the US S&P 500 stock market in the run-up to the recent global financial crisis along with features of subjective expectations documented in survey data. A comparison with the rational expectations model and the fully bayesian learning model enables an assessment of the specific role of the behavioral recency bias for explaining observed stock market excess volatility.

JEL classification: D83, D84, G12, G15.

Keywords: Asset Prices, Excess Volatility, Booms and Busts, Learning, Recency Bias, Survey data.

[*]gandre@gate.cnrs.fr, Ecole Normale Supérieure de Lyon, F-69007 Lyon, CNRS, GATE Lyon-St Etienne, UMR 5824, F-69130, Ecully, and Université Paris 1 Panthéon-Sorbonne, F-75005, Paris. I am grateful to Julien Azzaz, Sophie Béreau, Lise Clain-Chamosset, Camille Cornand, Frédéric Dufourt, Jean-Yves Gnabo, Frédéric Jouneau-Sion, Meguy Kuete, François Legrand, Valérie Mignon, Fabien Tripier, Mirko Wiederholt and Bruno Ziliotto for very helpful comments or discussions at distinct stages of this project, as well as to participants at the 12th Macrofi workshop, at the XVth RIEF Meetings in International Trade and Finance, at the 10th BIGSEM Workshop on Economic Theory, at the 5th EconomiX conference in International Economics and Financial Econometrics and at the 65th AFSE Annual Congress. I also wish to thank the scientific committee of the XVth RIEF Meetings for granting me the Best paper prize. All remaining errors are mine.

1 Introduction

Standard asset price theory fails to account for the long-standing excess volatility puzzle in stock prices relative to dividends, as is well-known since the seminal work of Shiller (1981) and LeRoy and Porter (1981). In addition, the standard stock pricing model cannot explain the existence of short-lived episodes of persistently high prices followed by a sudden steep decrease – the so-called booms and busts – as observed in the early 2000s in the US stock market in the run-up to the subprime financial crisis. In the face of this empirical evidence, Williams (2014) highlights that "one way to make sense of these patterns in asset prices is to abandon the assumption of full information assumed in the standard asset price model with rational expectations. Instead, investors make do with the limited information at hand when judging likely future discounted dividend payments and the future price of the asset."

Indeed, it is now well established that the rational expectations assumption does not withstand empirical scrutiny in many economic fields (Lovell, 1986) and that learning can generate excess volatility (see for instance Bullard and Duffy (2001)). More specifically, in what regards the stock market, several empirical studies emphasize the existence of extrapolative patterns in expected future returns, relying on survey data. Thus, Shiller (2000), using the results of a questionnaire submitted to US institutional investors in the 1990s, shows that expected future returns tend to co-move with the stock market. Recently, Greenwood and Shleifer (2014) find evidence of the fact that expected stock returns are positively correlated with both the price-dividend ratio and past returns, based on six distinct surveys of investor expectations.

Drawing upon recent developments in the literature on the impact of learning in macro-finance, the present paper provides simple theoretical foundations for explaining stock price fluctuations during the specific boom-and-bust episode that occurred in the US S&P 500 market in the 2000s. Our explanation relies on the relaxation of the rational expectations assumption and the introduction of imperfect attention in what regards the processing of sequential information over time. Relaxing the rational expectations assumption simply means that agents do not know with certainty the parameters of the laws of motion governing economic processes. More specifically, in our setting, agents learn the location parameter of the law of motion of the dividends' growth rate over time by inferring it from the history of past observations in a standard consumption-based asset pricing general equilibrium model (Lucas, 1978).

The departure from optimal rational behavior that is introduced in this setting is that agents are recency-biased, in the sense that they are constrained by limited cognitive ability to process information over time and thus pay gradually decreasing attention to previous data. Therefore, the recency-biased behavior is reflected in recursive discounting of the precision of past observations, in the context of bayesian inference. In this setting, a closed-form solution for stock prices can be derived, which enables to make explicit the dependence of the price-dividend ratio on expectations. Furthermore, this closed-form solution has microfoundations, because it is directly derived from the representative investor's maximization program.

Focusing on one specific boom-and-bust episode in the US stock market, i.e., from the early 2000s to the global financial crisis, it appears that the parsimonious model with recency-biased learning on dividends allows for the replication of the fluctuations in the price-dividend ratio and for a better fit to a number of additional features of expected and realized stock returns in comparison with the rational expectations and the bayesian learning benchmark models. Under recency-biased learning, the precision of agents' beliefs never reaches infinity, regardless of the precision of prior beliefs and the number of observations, thus generating persistent fluctuations in stock prices.

This paper relates to the growing strand of literature that incorporates learning schemes into standard asset pricing models, including Timmermann (1993), Timmermann (1996), Guidolin and Timmermann (2007), Koulovatianos and Wieland (2011), Adam et al. (2016a) and Adam et al. (2016b). This class of models is an alternative with strong intuitive appeal to rational expectations models, aiming at generating more realistic asset prices dynamics under time-separable utility, while simultaneously modeling subjective expectations, in accordance with growing evidence provided by the recent development of survey data on expectations. In comparison with the long-run risk literature (Bansal and Yaron, 2004), in the learning literature, the dynamics of expectations are more consistent with empirical evidence and fluctuations in asset prices intuitively result from variations in investors' perception without the need to assume ex-ante small long-run predictable component driving consumption and fluctuating economic uncertainty.

The learning mechanism investigated in our paper is however distinct from those investigated in the literature. First, agents take into account the uncertainty about their estimate of the parameter of the underlying fundamental process when forming their beliefs as if they were full bayesian learners, unlike in adaptive learning specifications such as Timmermann (1993) and Timmermann (1996).

Second, unlike Koulovatianos and Wieland (2011), who investigate the impact of learning on the probability of a rare disaster, we introduce learning on the mean parameter of the dividend's growth rate, which allows for the generation of fluctuations in the price-dividend ratio that are not smooth, consistent with the non-monotonic fluctuations observed in the data and allowing for the replication of a number of additional empirical features.

Third, agents rationally derive the equilibrium price as a function of their beliefs regarding the dividend process, which implies no restriction on the agents' knowledge of the mapping of fundamentals onto prices and no additional assumption on the agents' perceived law of motion of prices,[1] unlike Adam et al. (2016a), Adam et al. (2016b) and Nakov and Nuno (2015). Therefore, this paper focuses on the US stock price fluctuations in the run-up to the subprime financial crisis that other models with learning, specified and calibrated to target long-term moments, have had more difficulties to explain, suggesting that this period exhibits specific features. Even though learning on prices proves very successful in replicating the long-run moments and historical evolution of the US price-dividend ratio because of feedback effects (Adam et al., 2016a), it is less helpful in explaining recent evolution (see also Nakov and Nuno (2015)). Thus, during the 2000s, the price-dividend ratio exhibited a much higher and positive correlation with the dividend growth rate in comparison with the previous period, in particular in comparison with the 'dot-com' boom-and-bust episode over the 1995-2002 period (Table 1).

1960-1994	1995-2002	2003-2009
0.14	-0.47	0.64

Table 1: Contemporaneous correlation between the price-dividend ratio and the dividend's growth rate

This suggests that learning on dividends, as opposed to learning on prices, might still be useful to explain some specific boom-and-bust episodes in the stock market that seem to be more closely

[1] In models in which agents learn the stock price process rather than (or in addition to) the dividend process, the perceived law of motion of prices is distinct from the true one in the general case, not only regarding the parameters values but also regarding its general specification, and is thus specified exogenously.

correlated with the evolution of fundamentals than others.

Fourth, in our model, expectation dynamics are affected by agents' recency bias, which plays a major role in driving the model's dynamics. Indeed, the recency bias affects the mean value of beliefs, the uncertainty associated with these beliefs over time and their degree of persistence. The recency bias, which is a key element of the model mechanism, is supported by growing empirical evidence (De Bondt and Thaler (1990), Cheung and Friedman (1997), Agarwal et al. (2013), Erev and Haruvy (2013), Gallagher (2014)). Even the recent development of new risk assessment models in financial markets did not seem to prevent this bias. Thus, Andrew Haldane, a Chief Economist with the Bank of England, highlighted that the sophistication of financial risk assessment models in the banking sector in the early 2000s did not help identify anomalous patterns in financial markets because the regulation authorities did not pay attention to long-term data while using these models. Sophisticated stress tests models were only fed with macroeconomic and financial data from the most recent decade, characterized by lower variance, even though the sub-sample distribution was very distinct from the long-run historical distribution (Haldane, 2009).

In the face of empirical evidence, the assumption of decreasing weights allocated to past data has become rather standard in theoretical literature on financial markets dynamics as well (Bansal and Shaliastovich (2010), Shaliastovich (2015), Nakov and Nuno (2015)) and allows for the generation of persistent fluctuations in subjective expectations, never converging to rational expectations, and thus in prices, consistently with what is observed in the data. Recency-biased learning in our setting differs from alternative specifications, such as constant-gain learning schemes, to the extent that agents evaluate the uncertainty on their forecast as full bayesian learners and discount the precision of past information rather than past forecast errors. Consequently, the recency bias impacts the whole distribution of beliefs and notably the degree of confidence that agents have on their own estimate, generating gradually decreasing – but persistent – uncertainty. Therefore, the interpretation of recency bias in our model is that it relates to the limited ability of agents to process sequential information – with new information being available in each period through new realizations of dividends –, and in particular to the tendency to pay lower attention to information received earlier when revising beliefs, in accordance with the experimental literature (Ashton and Ashton (1988), Tuttle et al. (1997)). Decreasing attention is assimilated with increasing imprecision, and is thus modeled through recursive weights on the precision of past information.

The remaining of the paper is structured as follows. Section 2 presents the standard rational expectations model for benchmark purposes and then derives the subjective expectations model with bayesian learning and with recency-biased learning regarding the location parameter of the dividend growth rate process. Section 2.2 investigates the specific role of recency bias in driving asset prices dynamics, showing that it induces persistent volatility relative to unbiased learning. Section 3 calibrates the model on the US S&P 500 stock market in the recent period to provide quantitative results on the ability of the model to replicate the boom-and-bust episode of the 2000s and on the role of recency bias. Section 4 concludes and discusses the policy implications of our results.

2 Theoretical framework

We first characterize the model's rational expectations equilibrium. Such an equilibrium can be interpreted as the efficient benchmark equilibrium, in the sense that the rational expectations equilibrium arises when the economy is not subject to any type of friction, particularly informational friction.

2.1 The benchmark rational expectations model

The theoretical setting features an endowment economy, drawing upon the canonical Lucas (1978)' tree model. In each period, a representative risk-averse agent with a CRRA utility function decides what to consume, what to invest in a risky asset (stock) that pays exogenous perishable dividends (the output) and what to borrow or lend through risk-free bonds that pay an endogenous interest rate. All quantities are expressed in units of the single consumption good, which price serves as the numéraire. The representative agent's maximization program is as follows:

$$\max E_0^P \sum_{t=0}^{J} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{1}$$

s.t.

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

where P in the expectation operator E^P represents the objective probability measure, β is the discount factor, $\gamma > 0$ is the relative risk aversion coefficient, C_t is consumption in period t, P_t is the stock price, S_t the quantity of stocks, D_t the dividends earned on stock holdings S_{t-1} at the beginning of period t, r_{t-1} the real interest rate on bonds which is known already at the end of period $t-1$ and B_t the net amount of bonds traded in period t. As in the Lucas' model, stocks are in one unit exogenous supply, $S_{-1} = 1$, and bonds are in zero net supply.[2]

The dividend growth rate follows a normal process with parameters d and σ, as is standard in the literature (e.g. Pesaran et al. (2007) and Adam et al. (2016a)):

$$\log\left(\frac{D_t}{D_{t-1}}\right) = d + \sigma\varepsilon_t, \tag{2}$$

with $d > 0$, $\sigma > 0$ and ε_t white noise with unit variance. Given the resource constraints and the standard Euler equation, one gets:

$$D_t^{-\gamma} = \beta E_t\left[D_{t+1}^{-\gamma}\left(\frac{P_{t+1}+D_{t+1}}{P_t}\right)\right]. \tag{3}$$

Given the dividend growth process specification, an explicit expression for the price of stocks can be derived:

$$P_t = \delta_t D_t, \tag{4}$$

where $\delta_t = \frac{\beta\theta-(\beta\theta)^{J-t+1}}{1-\beta\theta}$ and $\theta = \exp\left(d(1-\gamma)+\frac{(1-\gamma)^2\sigma^2}{2}\right)$. See proof in Appendix A. One can immediately observe that the price-dividend ratio in period t (that is, δ_t) only depends on the model's fundamental parameters and on time t. Therefore, fluctuations in prices only reflect dividend shocks (and time-varying consumption decisions caused by the changing distance to the last period). Prices are not affected by any other variable in the model. Therefore, non-fundamental booms and busts cannot arise, which is even more obvious when $t << J$ and $\beta\theta < 1$ (which is the case in the calibrated version of the model in section 3), because $\delta_t \to \frac{\beta\theta}{1-\beta\theta}$ and thus the

[2]Note that we restrict the maximization program under rational expectations to be solved by finitely but very long-lived agents to remain close to a configuration involving infinitely lived agents. Following Pesaran et al. (2007), this assumption relates to the fact that in the standard bayesian learning case with infinitely lived agents, under general conditions, stock prices do not converge. The fact that stock prices do not converge under an infinite horizon can be an issue, because prices then increase with the planning horizon when the latter becomes very high. However, in the simulation exercise in Section 4, we verify ex-post that asset prices are insensitive to the choice of J for a very large range of values for J, through many iterated simulations.

price-dividend ratio tends to be constant.

In addition, under the rational expectations assumption, volatility in stock returns stems only from unpredictable dividend shocks, because:

$$\frac{P_t + D_t}{P_{t-1}} = \frac{1+\delta_t}{\delta_{t-1}} \frac{D_t}{D_{t-1}} = \frac{1}{\beta\theta} \exp(d + \sigma\varepsilon_t). \tag{5}$$

Consequently, expected returns are constant:

$$E_t\left[\frac{P_{t+1} + D_{t+1}}{P_t}\right] = \frac{1}{\beta\theta} \exp(d + 0.5\sigma^2). \tag{6}$$

Furthermore, the risk-free rate is constant in the rational expectations model:

$$1 + r_t = \frac{1}{\beta \exp(-\gamma d + 0.5\gamma^2\sigma^2)}. \tag{7}$$

Appendix B shows how all these implications of the rational expectations standard asset pricing model are at odds with several features of the data regarding the US S&P 500 in the run-up to the Great Recession.

2.2 The subjective expectations model

We now introduce one specific type of informational friction and assume that agents no longer know the true location parameter of the dividend growth process d and learn it over time through recency-biased updating.[3] Thus, the informational efficiency hypothesis is relaxed because the latter necessarily implies that agents hold rational expectations.

2.2.1 Belief dynamics

Agents observe the change in dividends, that is, the realization of the random variable $y_t = \log\left(\frac{D_t}{D_{t-1}}\right)$, but they do not separately observe its permanent fixed component d and its transitory component ε_t. Agents take d as a random variable and they evaluate the uncertainty about their estimate in each time period. When entering the stock market, agents have prior beliefs

[3] Assuming that agents know the true precision parameter allows for the direct identification of the impact of the evolution of one specific belief on asset prices, because agents only learn one parameter. In addition, with standard conjugate priors, the posterior distribution for the variance does not exist, as shown in Pesaran et al. (2007).

on the distribution of d that they update following the new realization of the dividend growth rate, according to Bayes' rule. Thereafter, they make their consumption-portfolio decision. The departure from rational behavior that we introduce is that agents face cognitive constraints to fully process information regarding dividend realizations received sequentially.[4] They are recency-biased, in the sense that they have limited ability to pay full attention to previous data. Less recent data is thus more imprecise for the agents. Therefore, the precision of previous observations is discounted relative to more recent ones with an informational discount factor $0 \leq \alpha < 1$. Thus, in period t, the precision of observations going back to period $t-k$ is discounted by α^k. Because α is a discount rate, the higher α is, the lower the degree of recency bias becomes. The case where $\alpha = 1$ reduces to standard bayesian learning.

Bayes' rule applied in period t is as follows:

$$P\left(d \mid I_t, \sigma\right) \propto L\left(y^t \mid d, \sigma\right) P\left(d \mid I_0, \sigma\right), \tag{8}$$

with $P\left(d \mid I_t, \sigma\right)$ the posterior distribution, $P\left(d \mid I_0, \sigma\right)$ the prior distribution and $L\left(y^t \mid d, \sigma\right)$ the likelihood function. I_t is the information set available on date t, which includes the history of past and current realizations of the dividend growth rate $y^t = \{y_1, y_2, ..., y_{t-1}, y_t\}$. Recency bias modifies the bayesian inference process only to the extent that the precision of past realizations is discounted. Thus, only parameters are affected.

Given that the dividend growth process follows a normal distribution, a natural prior distribution for d is the normal conjugate prior that allows for the analytical derivation of the posterior distribution, because it has the same form as the prior distribution. The prior distribution can thus be expressed as $P_0 \sim N(m_0, \sigma_0)$. Therefore, applying Bayes' rule under recency-biased learning ($\alpha < 1$) yields the posterior distribution at the end of period t, $d \backsim N(m_t, \sigma_t)$ with:

$$m_t = \frac{\tau \sum_{k=0}^{t-1} y_{t-k}\alpha^k + \tau_0 * m_0 * \alpha^t}{\frac{\tau(1-\alpha^t)}{1-\alpha} + \alpha^t \tau_0}, \tag{9}$$

where $\tau_0 = \frac{1}{\sigma_0^2}$ is the prior precision and $\tau = \frac{1}{\sigma^2}$ is the known precision of the log-normal process

[4] Agents are also assumed to be adaptive learners, in the sense that in each period t they condition their expectations on their beliefs in period t but do not take into account the possibility that their future beliefs might change following new realizations of the data (which are still unknown in period t). For a comparison between adaptive and rational bayesian learning, see Koulovatianos and Wieland (2011).

2. The posterior mean belief is the average of the discounted prior belief and of past and current observations weighted by their respective discounted precision in time t (the precision associated with each observation or with the prior evolves in each period because of the recency bias). We also have:

$$\sigma_t^2 = \frac{1-\alpha}{\tau(1-\alpha^t)+\alpha^t(1-\alpha)\tau_0}. \tag{10}$$

The uncertainty about the posterior estimate of d at time t is 10 (for a derivation of the expressions of m_t and σ_t, see Appendix C). The higher σ_t^2 is, the lower the confidence of agents about their mean belief becomes. Unsurprisingly, this uncertainty is higher when the precision τ of the fundamental process is lower and when the precision of the prior τ_0 is lower.

2.2.2 Closed-form solution for stock prices under recency-biased learning

The agent's maximization program is:

$$\max E_0^P \sum_{t=0}^{J} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{11}$$

s.t.

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

where P is now the representative agent's subjective probability measure.

From the first-order and market clearing conditions, we can derive the subjective expectations equilibrium stock price:

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j E_t^P \left[\left(\frac{D_{t+j}}{D_t} \right)^{1-\gamma} \right], \tag{12}$$

(see details in Appendix D). Hence, conditionally on d and σ and information up to period t, we have:

$$E^P \left[\left(\frac{D_{t+j}}{D_t} \right)^{1-\gamma} \mid I_t, d, \sigma \right] = \exp\left[(1-\gamma) j d + 0.5(1-\gamma)^2 j \sigma^2\right]. \tag{13}$$

Therefore, when d is no longer known, agents integrate the previous expression over the whole distribution of d:

$$E^P\left[\exp\left((1-\gamma)jd + 0.5(1-\gamma)^2\sigma^2 j^2\right) \mid I_t, \sigma\right]$$
$$= \exp\left(0.5(1-\gamma)^2\sigma^2 j\right) E^P\left[\exp((1-\gamma)jd) \mid I_t, \sigma\right]. \tag{14}$$

Because d is believed to follow a normal distribution with parameters m_t and σ_t in period t, $\exp((1-\gamma)jd) \sim Log - N((1-\gamma)jm_t, (1-\gamma)j\sigma_t)$. Therefore,

$$E^P\left[\left(\frac{D_{t+j}}{D_t}\right)^{1-\gamma} \mid I_t, \sigma\right] = \exp(0.5(1-\gamma)^2\sigma^2 j)\exp((1-\gamma)jm_t + 0.5(1-\gamma)^2 j^2\sigma_t^2). \tag{15}$$

Eventually, we get the following pricing function.

Under recency-biased learning, the stock price is equal to:

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)). \tag{16}$$

One can immediately observe that under learning, the price-dividend ratio now depends on the hyperparameters m_t and σ_t, which are updated every period and are therefore time-varying. Fluctuations in the price-dividend ratio over time – that is, fluctuations in prices that do not directly reflect changes in dividend realizations – are driven by changes in expectations in response to stochastic realizations of the dividend growth rate. The dynamics of subjective expectations following fundamental shocks thus amplify and propagate those shocks. The learning process therefore generates excess volatility in asset prices. Comparative statics show that the price-dividend ratio in t increases in the mean belief m_t if and if only $\gamma < 1$ and increases in the discount factor β, whereas its relation to the relative risk aversion coefficient γ is ambiguous (see Appendix E).

In addition, under subjective expectations, stock returns are likely to display greater volatility than under rational expectations. Indeed,

$$\frac{P_t + D_t}{P_{t-1}} = \frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}\frac{D_t}{D_{t-1}} = \frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}\exp(d + \sigma\varepsilon_t), \tag{17}$$

with $\delta_{t,SE} = \sum_{j=1}^{J-t}\beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2))$. Unlike in the rational expectations case, volatility in stock returns does not stem only from volatility in the dividend growth rate, because the ratio $\frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}$ is now time-varying and reflects the additional impact of dividend shocks on expectations.

Regarding expected returns, they are now given by:

$$E_t\left[\frac{P_{t+1}+D_{t+1}}{P_t}\right] = \frac{1}{\delta_{t,SE}}\Big(\sum_{j=1}^{J-t-1}[\beta^j \exp(m_t((1-\gamma)j+1)+0.5((1-\gamma)j+1)^2\sigma_t^2 \\ +0.5((1-\gamma)^2 j+1)\sigma^2)] + \exp(m_t + 0.5(\sigma_t^2+\sigma^2))\Big),$$

(see Proof in Appendix F).

The first term of the product is the inverse of the current price-dividend ratio whereas the second term of the product is positively correlated with the current price-dividend ratio. Indeed, similarly to the current price-dividend ratio, the latter term increases in σ_t and increases in m_t when $\gamma < 1$. This situation creates the possibility for a positive correlation between expected returns and the price-dividend ratio, as evidenced by survey data on expectations about future stock returns (see Appendix B). The intuition behind such a positive correlation under learning is the following: stock returns $\frac{P_{t+1}+D_{t+1}}{P_t}$ can be rewritten as $\frac{\frac{P_{t+1}}{D_{t+1}}+1}{\frac{P_t}{D_t}}\frac{D_{t+1}}{D_t}$, where $\frac{D_{t+1}}{D_t}$ is unpredictable. Therefore, under learning, unlike in the rational expectations case, when the price-dividend ratio is higher in period t, it is expected to be higher in period $t+1$, because current beliefs are extrapolated to the future.

In addition, the risk-free rate now writes:

$$1 + r_t = \frac{1}{\beta \exp(-\gamma m_t + 0.5\gamma^2\sigma^2 + 0.5\gamma^2\sigma_t^2)}.$$

3 Recency bias, expectations and asset price fluctuations

3.1 Recency bias and belief dynamics

To begin, we present analytical results on the impact of a marginal decrease in the degree of recency bias (that is, the impact of a marginal increase in the informational discount factor α) on the dynamics of beliefs, which is then transmitted to stock prices.

First, a marginal decrease in the degree of recency bias has an ambiguous impact on the mean belief m_t as the sign of the following quantity is ambiguous:

$$\frac{\partial m_t}{\partial \alpha} = \frac{(\tau \sum_{k=0}^{t-1} k y_{t-k} \alpha^{k-1} + \tau_0 m_0 t \alpha^{t-1})(\tau \frac{1-\alpha^t}{1-\alpha} + \alpha^t \tau_0)}{(\tau \frac{1-\alpha^t}{1-\alpha} + \alpha^t \tau_0)^2} - \frac{(\tau_0 t \alpha^{t-1} + \tau \frac{1+\alpha^t(t-1)-t\alpha^{t-1}}{(1-\alpha)^2})(\tau \sum_{k=0}^{t-1} y_{t-k} \alpha^k + \tau_0 m_0 \alpha^t)}{(\tau \frac{1-\alpha^t}{1-\alpha} + \alpha^t \tau_0)^2}.$$

Indeed, a lower recency bias generates lower discounting of the precision of past observations, and thus a higher weight is allocated to the precision of past observations when deriving the mean belief. When past observations are higher relative to recent observations, a lower degree of recency bias may cause the mean belief to increase. When past observations are lower relative to recent observations, a lower degree of recency bias may cause the mean belief to decrease. The overall impact of a decrease in the degree of recency bias is therefore ambiguous and depends on the sequence of realizations available to the agent.

Second, unsurprisingly, one can immediately observe that the dispersion parameter of the posterior distribution in time t, $\sigma_t^2 = \frac{1-\alpha}{\tau(1-\alpha^t)+\alpha^t\tau_0(1-\alpha)} = \frac{1}{\alpha^t(\tau_0 - \frac{\tau}{1-\alpha}) + \frac{\tau}{1-\alpha}}$, decreases with the informational discount rate α (which means that it increases with the degree of recency bias). This effect is related to the fact that when the weight allocated to the precision of the information regarding the true parameter provided by each past realization of the data is higher, the overall precision is higher. Note however that the dispersion parameter of the posterior distribution σ_t converges faster to its minimal value $\frac{1-\alpha}{\tau}$ when the recency bias is higher (that is, when α is lower). Indeed, the quantity $\alpha^t(\tau_0 - \frac{\tau}{1-\alpha})$ converges faster to 0 when α is lower.[5]

The stock price level depends on the degree of recency bias in two ways. First, it depends on α through the mean belief m_t, which is a function of α. Second, it depends on α through the dispersion parameter of the posterior distribution σ_t, which is also a function of α. Therefore, the impact of a marginal increase in α on the price level is also ambiguous:

$$\frac{\partial P_t}{\partial \alpha} = D_t \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma) m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2))((1-\gamma) j \frac{\partial m_t}{\partial \alpha} + (1-\gamma)^2 j^2 \sigma_t \frac{\partial \sigma_t}{\partial \alpha}).$$

[5]Notice that the finite forecast horizon set-up does not prevent us from analyzing the convergence properties of beliefs when the number of past observations becomes very high, because the size of the sample of past observations is not related to the planning horizon.

A higher α decreases the uncertainty parameter σ_t, which decreases the stock price. However, a higher α also has an ambiguous impact on the mean belief m_t, depending on the sequence of past and current observations. When the dividend growth rate is above its mean d for some time, agents start to overestimate the unknown parameter d, which raises the stock price. However, whether agents overestimate the unknown parameter more or less when the degree of recency bias decreases depends on how high recent realizations of the dividend growth rate are relative to less recent realizations. If recent realizations are, for some time, much higher on average than past realizations, a higher degree of recency bias will tend to decrease the mean belief m_t, thus decreasing the stock price because σ_t also decreases.

If the reverse occurs, m_t will increase, thus increasing the stock price if the increase in m_t is high enough relative to the decrease in σ_t.

In addition, beliefs' convergence properties to rational expectations are affected by the recency bias.

Subjective expectations converge to rational ones in the limit if and only if $\alpha = 1$.

Indeed,

$$\lim_{t \to +\infty} \sigma_t^2 = \frac{1}{\sum_{k=0}^{t-1} \alpha^k * \tau + \alpha^t \tau_0} = 0 \Longleftrightarrow \alpha = 1. \tag{18}$$

Under bayesian learning ($\alpha = 1$), when the sample of observations increases, uncertainty in the posterior estimate of d decreases and converges to zero. Thus, in the limit, agents are no longer uncertain regarding their mean belief, and the law of large numbers ensures that the bayesian estimate of d is equal to $\lim_{t \to +\infty} m_t = \lim_{t \to +\infty} \frac{\sum_{k=0}^{t-1} y_{t-k}}{t} = d$. The posterior estimate of d then converges to its true value when the number of past observations tends towards infinity (it is unbiased).

Under recency-biased learning ($\alpha \neq 1$), beliefs never converge to rational expectations ones. Indeed, $\lim_{t \to \infty} \frac{1}{\frac{\tau(1-\alpha^t)}{1-\alpha} + \alpha^t \tau_0} = \frac{1-\alpha}{\tau}$. Thus, the precision of the estimate is never infinite, and agents' estimates continue to evolve following new realizations of dividend growth even in the limit, with $\lim_{t \to +\infty} m_t = \sum_{k=0}^{t-1} y_{t-k} \alpha^k (1-\alpha)$.

3.2 Recency bias and excess volatility in asset prices

Recency-biased learning plays a crucial role in generating excess volatility in asset prices compared with unbiased learning. First, the quantity $\frac{\partial m_t}{\partial y_t}$ measures how the mean belief reacts to the current realization of the data, depending on the degree of recency bias α and on the history of past realizations. It writes:

$$\frac{\partial m_t}{\partial y_t} = \frac{\tau}{\tau \frac{1-\alpha^t}{1-\alpha} + \alpha^t \tau_0} \geq 0. \tag{19}$$

In accord with intuition, the higher the realization of the growth rate of the dividends is, the higher the mean belief on the mean of the fundamental process becomes. For $\alpha = 1$ (that is, in the case of unbiased learning), the following expression holds:

$$\frac{\partial m_t}{\partial y_t} = \frac{\tau}{\tau t + \tau_0} \geq 0. \tag{20}$$

In both cases, $\frac{\partial^2 m_t}{\partial y_t \partial \alpha}$ is negative. Therefore, the lower the degree of recency bias is, the lower the reaction of the mean belief to new realizations of the data becomes. In particular, the variation is clearly higher when $\alpha < 1$ compared with the variation observed when $\alpha = 1$, given that $\frac{1-\alpha^t}{1-\alpha} \leq t$ for any $t > 1$.

This disparity is even stronger when t tends to infinity. Thus, $\frac{\partial m_t}{\partial y_t} \to 1 - \alpha$ when $\alpha < 1$ and $\frac{\partial m_t}{\partial y_t} \to 0$ when $\alpha = 1$, reflecting the fact that in the absence of recency bias, the uncertainty about the model's parameter vanishes in the limit. In addition, it appears that there are significant non-linearities in the impact of a marginal decrease in recency bias on the volatility of beliefs. These non-linearities are obvious when the number of observations becomes very high, as $\frac{\partial^2 m_t}{\partial y_t \partial \alpha} \to -1$ when $\alpha < 1$, whereas $\frac{\partial^2 m_t}{\partial y_t \partial \alpha} \to \frac{-1}{2}$ when $\alpha = 1$, which implies that a marginally lower recency bias more strongly mitigates the volatility in the mean belief when the recency bias does not tend to zero.

The impact of recency bias on the volatility of the dispersion parameter of the posterior distribution σ_t is more ambiguous, except in the limit, in which the variance of the posterior distribution becomes constant at $\frac{1-\alpha}{\tau}$ when $\alpha < 1$ and at 0 when $\alpha = 1$, because of convergence to the true parameter (the uncertainty about the actual parameter is null).

This result implies that under unbiased learning, the price-dividend ratio is constant in the limit because agents are perfectly confident that their estimate is equal to the actual unknown parameter. By contrast, under recency-biased learning, even when uncertainty is constant after a large number of periods, agents keep updating their beliefs because uncertainty does not converge to zero. This effect triggers persistent volatility in the price-dividend ratio, driven by persistent fluctuations in the mean belief m_t, even when the uncertainty parameter σ_t is constant.

4 The boom-and-bust episode in the US S&P 500 stock market in the run-up to the global financial crisis

4.1 Simulation results

To assess the role of recency bias in explaining the 2000s boom-and-bust episode, we now calibrate the model on US S&P 500 monthly data (Robert Shiller's dataset[6]). We conduct a horse race of the rational expectations model, the bayesian learning model and the recency-biased learning model.

In the following simulation, agents enter the market in the early 2000s but they have access to information on the past behavior of stock prices. Agents are assumed to rely on historical information on dividend realizations from 1960 onwards. Indeed, Pesaran et al. (2007) identify a structural break in dividends around the year 1960.

Parameters of the real dividend growth process are chosen so as to match those in the data over the period from January 1960 to June 2009 (Table 2). June 2009 corresponds to the period immediately following the climax of the bust on the US stock market. Subsequent data are likely to present structural breaks, and the learning process agents relied on in the former period is likely to have suffered strong changes, in terms of both underlying parameters and learning scheme. We indeed observe a decreasing correlation between dividends and stock prices following the bust.

Because the prior distribution $d \sim N(m_0, \sigma_0)$ summarizes the initial information the agent has regarding the estimated parameter, the parameters m_0 and σ_0 are derived from pre-sample dividend realizations. The prior parameters are retrieved from dividend realizations over the period 1960-2003 in a way that is consistent with the in-sample agents' learning process: m_0 is the weighted average of past observations relying on the informational discount rate α, with the realization in January 1960 being discounted the most and the realization in December 2002 not being discounted, and σ_0^2 is the inverse of the discounted sum of the precision of the information provided by each data point, i.e., the inverse of the discounted sum of the precision of the fundamental process τ. Finally, the model is fed with the exact same dividend realizations as those observed

[6]Monthly data on the US S&P 500 stock market are retrieved from Shiller's dataset available online.

in the data in the period of interest when deriving the model-implied price-dividend ratio, following Adam et al. (2016a). Such a strategy enables for the assessment of the ability of the model to replicate actual stock price behavior from actual observed dividends (as opposed to random dividends).[7]

The three remaining parameters: the monthly discount rate β, the monthly informational discount rate α, and the constant relative risk aversion parameter γ are chosen to minimize the absolute distance between the model-generated price-dividend ratio and the actual price-dividend ratio in the run-up to the Great Recession as follows:

$$\sum_{t=2003M1}^{2009M6} \mid \delta_{t,SE} - \delta_{t,data} \mid,$$

where $\delta_{t,SE}$ is the model-generated price-dividend ratio and $\delta_{t,data}$ is the actual price-dividend ratio. Such a minimization procedure is standard in the literature focusing on the replication of specific boom-and-bust episodes in financial markets (Kuang, 2014). [8] The minimization procedure yields that $\beta = 0.998$, which is consistent with values in the literature (Nakov and Nuno, 2015). The minimization procedure also implies that $\alpha = 0.90$, which is similar to the value found in empirical studies that have evaluated the informational discount factor in distinct markets (see in particular Agarwal et al. (2013)), and that $\gamma = 0.93$, locating it in the upper part of the range over which we minimized the objective function, while still being relatively low. Therefore, the model calibration allows us to show that recency-biased learning per se is sufficient to replicate the fluctuations in asset prices observed in the 2000s, even at low degrees of risk aversion. Eventually, the future horizon J is taken inside the range for which the value for J does not affect the simulation results. Iterated simulations reveal that this range is very large, ranging roughly from 4000 to 180 000 months, i.e. from 333 to 15000 years.

[7]Note that agents' prior beliefs, based on pre-sample dividend information, are such that agents undervalue the actual parameter (Table 3). The simulation results are robust to relying on the median value of the observed dividends distribution rather than on the mean value.

[8]Note that α is constrained to be in the $[0.9, 1]$ range such that the minimization generates plausible values for the degree of recency bias and γ (which is also the inverse of the intertemporal elasticity of substitution because of CRRA preferences) is constrained to be less than 1 such that the substitution effect exceeds the wealth effect and counter-intuitive features that would contradict the data are excluded.

Parameter	Calibrated value
Mean of the dividend process d	0.0011
Standard error of the dividend process σ	0.0056
Monthly discount rate β	0.998
Relative risk aversion γ	0.93
Informational discount factor α	0.90
Prior mean m_0	-0.0010
Prior standard error σ_0	0.0018
Planning horizon J	$\in [4000, 180000]$

Table 2: Calibration

The subjective expectations model with recency bias performs strikingly much better than the rational expectations benchmark and the full bayesian learning model (Figure 1).

First, the model replicates the boom period in the US stock market in the aftermath of the dot-com bubble bust, with the price-dividend ratio increasing significantly and remaining persistently above its rational expectations value – which can be identified as a boom episode. Second, the model replicates the deep decrease in the price-dividend ratio, reaching values well below the rational expectations value – the bust. The bust in the price-dividend ratio results from the transmission of a strong negative dividend shock – which may reflect exogenous deteriorating financial and economic conditions in the context of the subprime crisis – to beliefs about the actual fundamental process, because of persistent parameter uncertainty. This negative effect is then transmitted to demand for the risky asset. Indeed, because the substitution effect dominates, stock prices collapse.

The extent of the bust is stronger given that the actual parameter of the fundamental process was overvalued in the previous periods, because of a series of positive fundamental shocks. Indeed, the strong negative shock thus induces a higher reassessment of the parameter estimate, which is caused by higher uncertainty, making beliefs more sensitive to current data realizations. The simulation exercise thus clarifies how the alternation of phases of overvaluation and undervaluation of stocks' fundamental value generates significant boom-and-bust episodes in the stock market. The model is also able to explain at least part of additional features of the stock market over the period of interest.

Thus, first, the model well replicates the strong autocorrelation in the price-dividend ratio be-

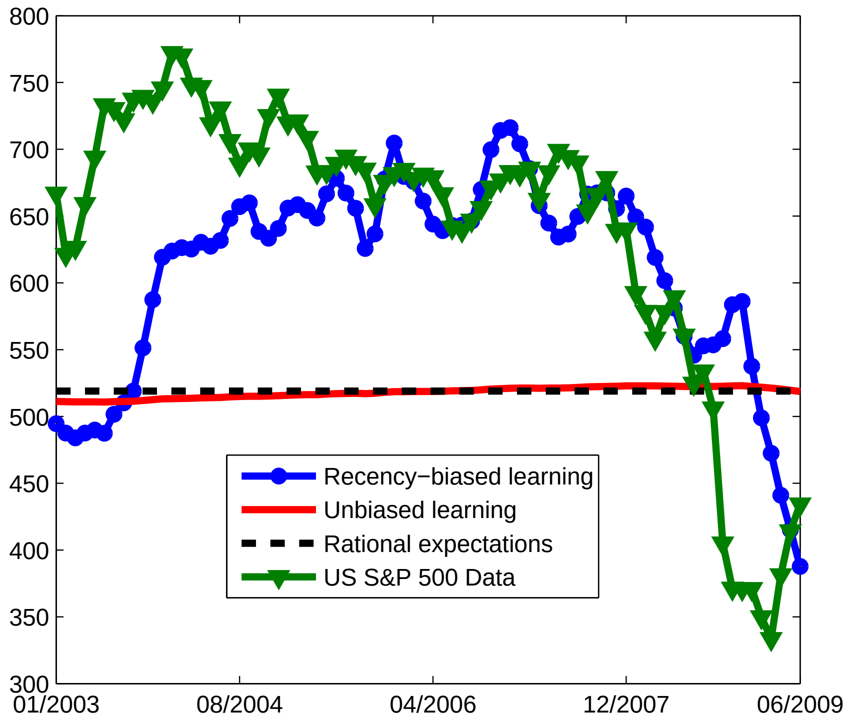

Figure 1: 2003-2009 Monthly price-dividend ratio

havior over the period (Table 3), although this is not specific to the learning scheme with recency bias. Second, the dynamics of the realized monthly stock returns are more consistent with those observed empirically than those generated by the rational expectations benchmark model and by the bayesian learning model (Figure 2). The recency-biased model misses part of the volatility of the stock returns, but it however generates almost 5 times more volatility in the returns than the two other models (Table 3). Similarly, the recency-biased model fails to replicate the full equity premium but it explains significantly more of it than the rational expectations model and the unbiased learning model. Thus, unbiased learning helps better explain part of the equity premium puzzle at low levels of risk aversion. Third, the recency-biased subjective expectations model accounts for several features of stock returns expectations, as measured by the CFO Survey. Indeed, firstly, the model generates a positive correlation between expected stock returns and the price-dividend ratio, which is consistent with survey data on expected returns (Table 3),[9] even though it tends

[9]Number in brackets are P-values. All returns in Table 3 are gross monthly returns which are annualized. Model-implied one-year ahead expected returns are simply annualized monthly returns. The CFO survey being published quarterly and the model time lapse being a month, for comparability issues, we linearly interpolated quarterly expected

to overestimate it, and so does the unbiased learning model. Secondly, the recency-biased learning model closely replicates the volatility in forecast errors (measured as the difference between realized stock returns and expected returns in the previous period) for stock returns in the CFO survey.

	Rational Expectations	Bayesian learning	Recency-biased learning	Data
$Corr(\frac{P_t}{D_t}, \frac{P_{t-1}}{D_{t-1}})$	NaN	0.98	0.90	0.95
$E[Z_t]$	1.068	1.072	1.132	1.492
$\sigma(Z_t)$	0.11	0.12	0.52	0.80
$E[1+r_t]$	1.037	1.036	1.079	1.050
Correlation between 1-year ahead expected returns and the P/D ratio				
Monthly	NaN	1 (0.00)	0.99 (0.00)	0.74 (0.00)
Quarterly	NaN	1 (0.00)	0.99 (0.00)	0.67 (0.00)
Forecast errors for gross stock returns				
Mean	0.04	0.04	0.06	0.05
Standard-error	0.10	0.10	0.51	0.52

Table 3: Simulation results

Finally, to clarify what drives the dynamics of asset prices in the model, Figure 3 shows the joint dynamics of model-implied beliefs and the price-dividend ratio. The figure shows how the mean belief m_t (left-hand scale) fluctuates around the actual parameter d. First, it displays a sustained period of optimism – in which $m_t > d$ – and second, it displays a sudden peak of pessimism – in which m_t goes far below d –, which drives the dynamics of the price-dividend ratio (right-hand scale). The uncertainty parameter σ_t being almost constant in the simulation results,[10] it is clear that the dynamics of the mean belief parameter directly translate into the dynamics of the price-dividend ratio.

4.2 The role of the degree of recency bias

We now assess the quantitative implications of the degree of recency bias on asset price volatility. Figure 4 presents the evolution of the price-dividend ratio for distinct degrees of recency bias α. At

returns in the CFO survey to get monthly data. To obtain an additional result that is independent of interpolation methods, we compare quarterly data as well by taking model-implied annualized monthly expected returns of the last month of each quarter as a proxy for each quarter one-year ahead expected returns. For details on the CFO survey, see Appendix B.

[10]This is caused by the large number of pre-sample periods included in the prior calculation. However, because of informational discounting, even though more data are accumulated by agents over time, the precision of information is not improved over the period because the earliest data are allocated zero weight. For the given model's parameter values, the constant precision of agents' information regarding the unknown parameter is equal to $3.1 * 10^5$ (that is, the uncertainty about this parameter is equal to $3.2 * 10^{-6}$). Therefore, the results show that even for a very low value of uncertainty, belief updating still generates significant volatility in the price-dividend ratio.

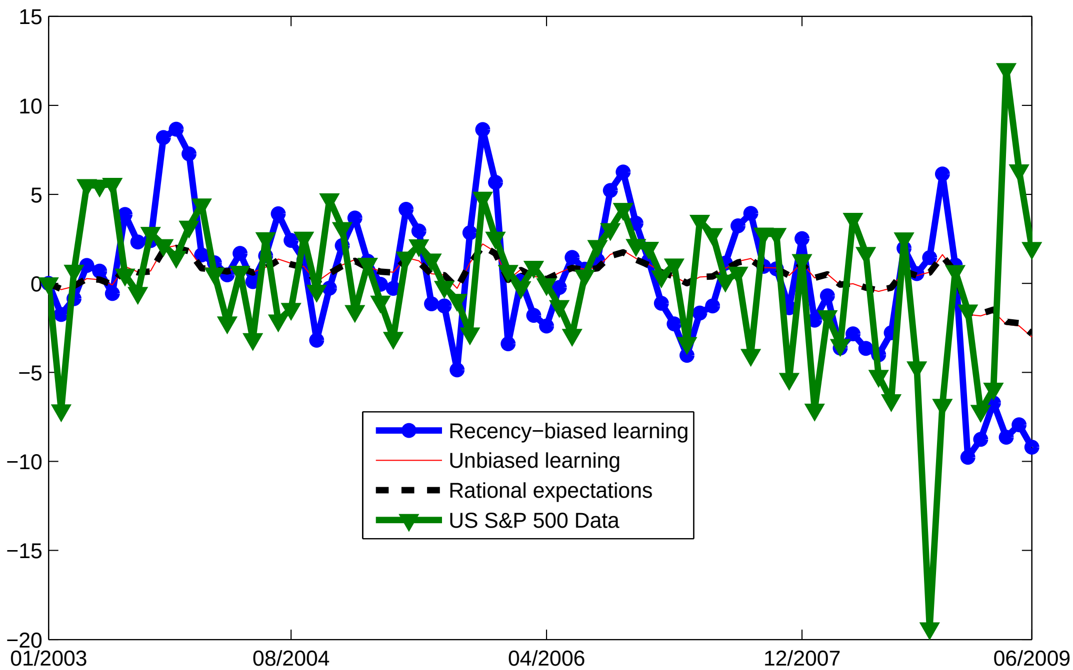

Figure 2: Monthly net returns on stocks (%)

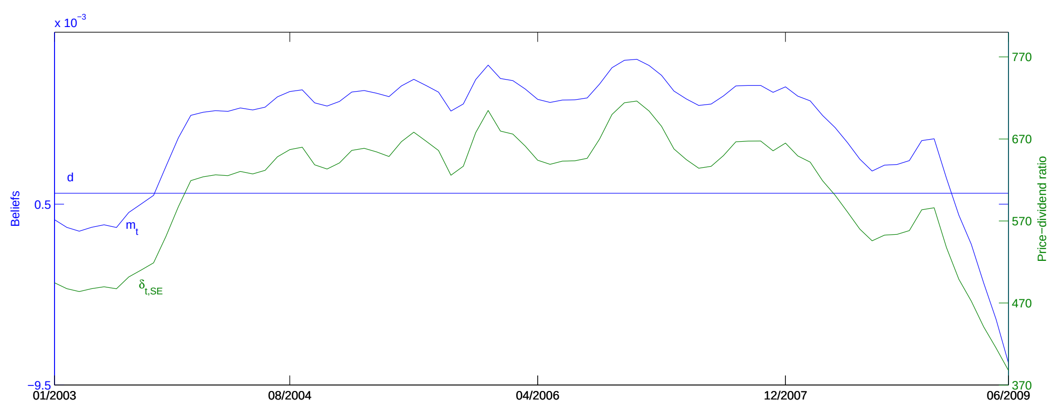

Figure 3: Joint dynamics of beliefs and the price-dividend ratio

first glance, it is striking that the degree of recency bias significantly affects asset price dynamics. First, the volatility in the price-dividend ratio tends to zero only when the representative agent is not recency-biased, in accord with analytical results. Even for a very small informational discount rate of $\alpha = 0.98$, the price-dividend ratio remains significantly far from its rational expectations 'fundamental' value. However, when the degree of recency bias decreases, volatility always decreases. Second, it appears that the marginal impact of the recency bias on the price-dividend ratio varies with the informational discount rate, reflecting the existence of non-linearities.

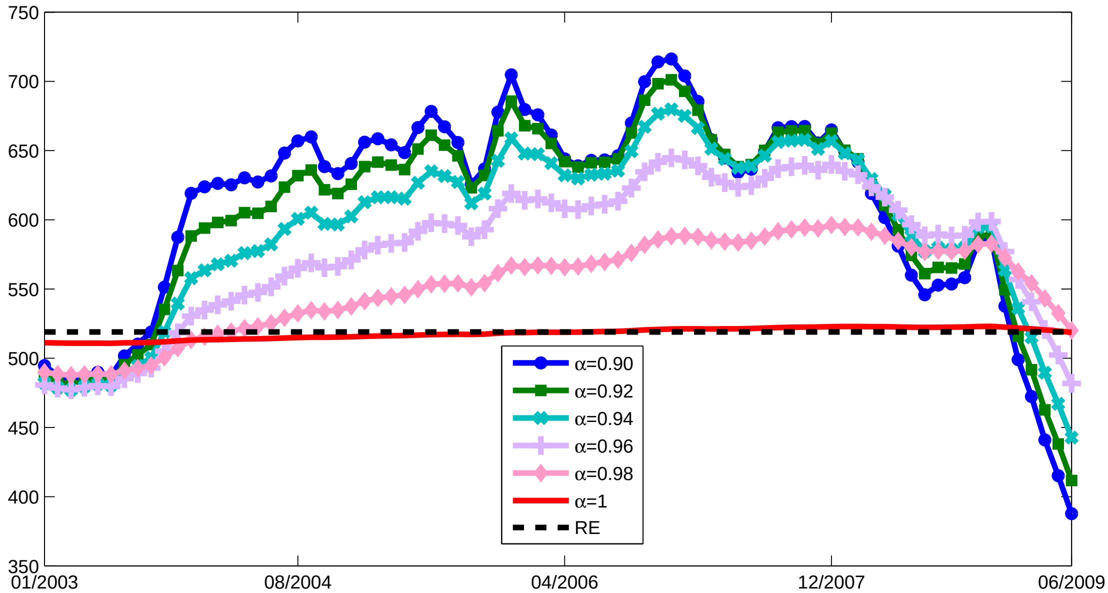

Figure 4: The impact of the degree of recency bias on the volatility of the price-dividend ratio

To complement this finding, Figure 5 presents the variance of the subjective expectations price-dividend ratio and the mean of the squared distance of the subjective expectations price-dividend ratio to its rational expectations value as functions of the informational discount factor. First, it appears that a lower degree of recency bias monotonically reduces the variance of the subjective expectations price-dividend ratio and its average squared distance to the rational expectations value. However, both variables tend to zero only when the precision of the mean belief under subjective expectations is maximal. Second, a marginal increase in the precision of the estimate (caused by lower recency bias) more strongly decreases the volatility in the price-dividend ratio when the informational discount rate is higher, except when it is already very close to 1.

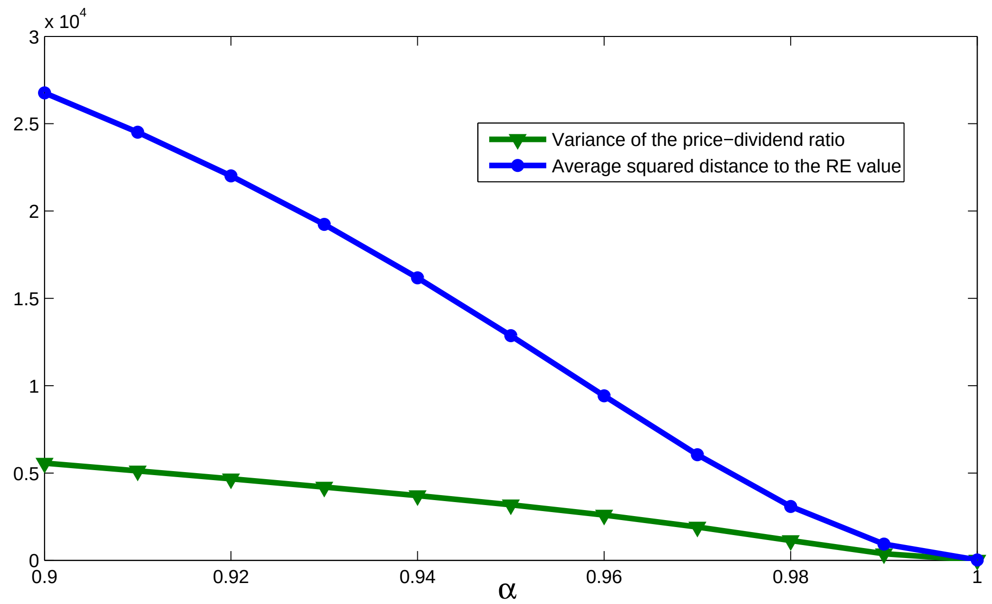

Figure 5: Statistical features of the price-dividend ratio depending on the degree of recency bias α

Finally, we assess under what degree of recency bias only small deviations from the rational expectations value arise.[11] The simulation results are displayed in Table 4.

Objective	Minimal α (order -4)
Standard error of P/D: <5% of the mean	0.9886
Standard error of P/D: <10% of the mean	0.9399
Squared root of the mean squared distance of P/D to its RE value: <5% of the RE value	0.9908
Squared root of the mean squared distance of P/D to its RE value: <10% of the RE value	0.9791

Table 4: Minimal degree of recency bias required to reduce the gap with the frictionless rational expectations benchmark

It appears that very small degrees of recency bias are required to observe a significant decrease in the price-dividend ratio volatility and in the average distance of the price-dividend ratio to its fundamental value.

Note that the prior distribution specified above displays significant precision despite the recency bias, because it is based on a rather large prior sample. However, with an uninformative prior – corresponding to the opposite case in which prior precision is null and thus in which the uncertainty parameter fluctuates over the period of study –, the volatility would still decrease

[11]For comparison, with $\alpha = 0.90$, the standard deviation of the price-dividend ratio over time is equal to 12.24% of its mean, and the square root of the mean squared distance of the price-dividend ratio to its rational expectations value is equal to 22.55% of the rational expectations value.

gradually with the degree of recency bias, as evidenced by the analytical results discussed in section 3.2. In addition, the biased learning case would still display significant differences relative to the unbiased learning case. However, unsurprisingly, the price-dividend ratio volatility would remain significantly higher, even when agents are not recency-biased, compared with the case with a more informative prior.

4.3 Robustness to alternative calibration strategies

We now assess the sensitivity of our results to alternative choice of parameters. First, we apply the minimization procedure described above for the recency-biased learning model to the rational expectations model and the bayesian learning model. We can thus separately derive the values of the monthly discount rate β and of the constant relative risk aversion parameter γ that minimize the distance between the realized $S\&P\ 500$ price-dividend ratio and the model-generated price-dividend ratio over the period of interest in both cases. Therefore, we can assess whether the failure of both models to replicate the 2000s boom-and-bust episode relates to the choice of parameters that are calibrated for the recency-biased learning case. Table 5 presents the results of the minimization procedure for the rational expectations model (RE) and the bayesian learning model (BL).

Parameter	Calibrated value RE	Calibrated value BL
Monthly discount rate β	0.998	0.998
Relative risk aversion γ	0.51	0.57

Table 5: Alternative calibration for the rational expectations model and the bayesian learning model

The calibrated monthly discount rate is similar for the three models. The discount rate has a strong impact on the mean price-dividend ratio and a similar value of the discount rate is required to generate a consistent order of magnitude of the mean price-dividend ratio in the three distinct models. As for the calibrated constant relative risk aversion, it is very low in both the rational expectations model and the bayesian learning model relative to the recency-biased learning model. Figure 6 presents the price-dividend ratio generated in the three models, based on the separate calibration procedure, and that in the data. Despite the specific calibration strategy, the

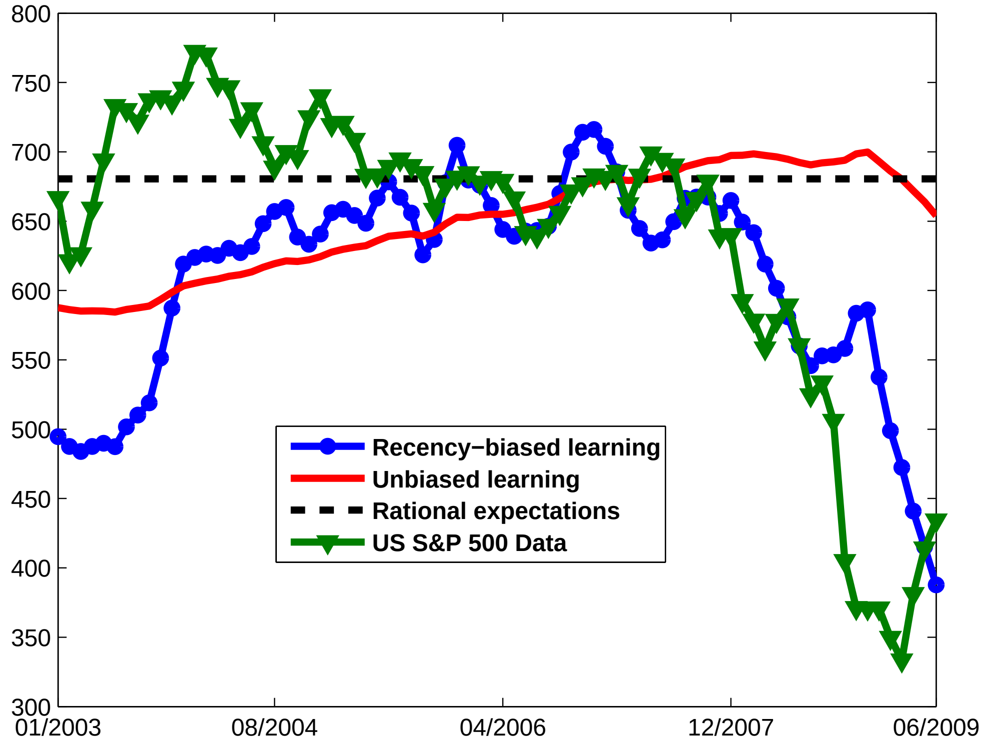

Figure 6: Price-dividend ratio with a minimization procedure specific to each model

rational expectations model and the unbiased learning model still prove unable to replicate the US stock market boom-and-bust episode in the 2000s. Note however that the bayesian learning model generates slightly more volatility in the price-dividend ratio under the new calibration.

Second, we investigate the sensitivity of the results to distinct out-of-sample calibration for the prior parameters, that is, the mean parameter m_0 and the uncertainty parameter σ_0. Figure 7 presents the price-dividend ratio for several values of the prior parameters, depending on the size of the out-of-sample data taken into account in the calibration.

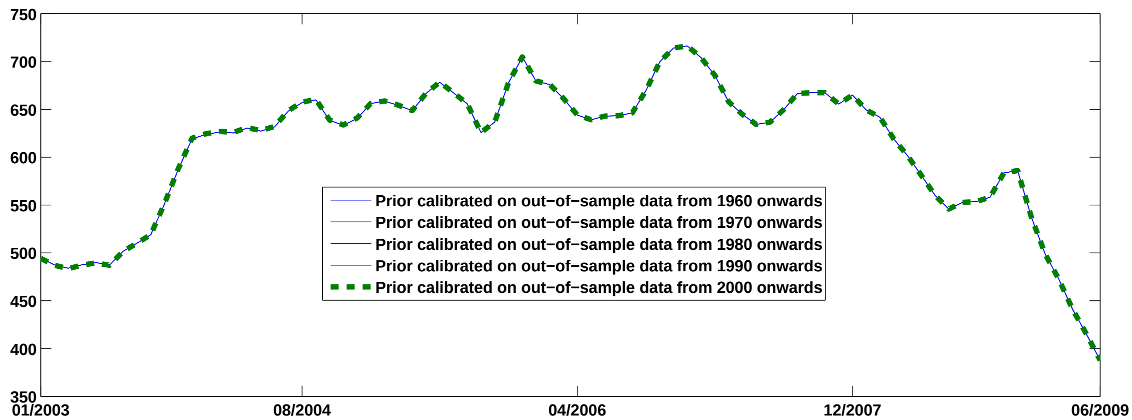

Figure 7: Price-dividend ratio for several samples of calibration of the prior distribution (recency-biased learning)

The result of the robustness check is striking : due to informational discounting of oldest observations under recency-bias learning, reducing the size of the pre-sample data used in the prior calibration has no impact on the price-dividend ratio. Very small differences in the price-dividend ratio only appear when the prior is calibrated on the data from 2000 onwards but they are hardly visible. By contrast, absent the recency bias, the choice of the prior parameters slightly affects the dynamics of the price-dividend ratio (Figure 8).

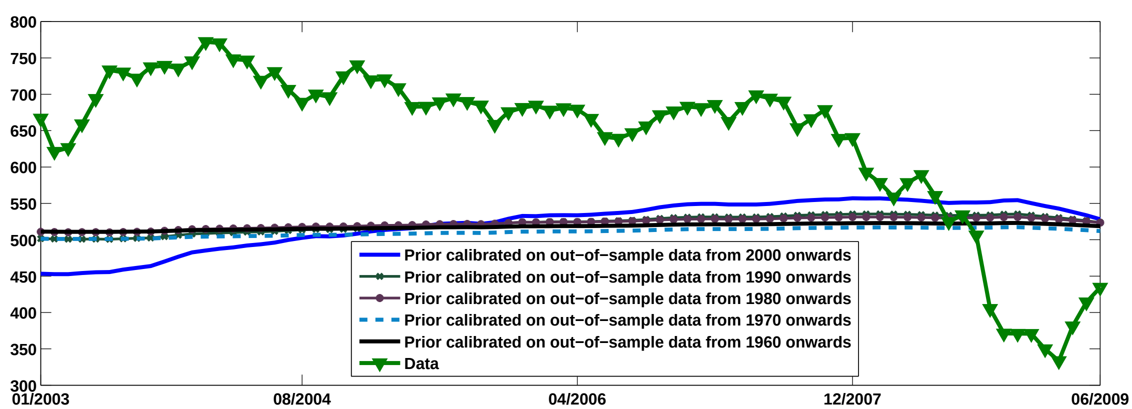

Figure 8: Price-dividend ratio for several samples of calibration of the prior distribution (unbiased bayesian learning)

Using a smaller pre-sample dataset for calibrating the prior slightly increases the volatility of the price-dividend ratio under fully bayesian learning because the prior precision is then smaller. Therefore, beliefs react more strongly to new data realizations. However, even when the out-of-sample dataset is very small (starting in 2000), the volatility of the price-dividend ratio under unbiased learning remains very low relative to the data and to the recency-biased learning case.

5 Conclusion and discussion

A parsimonious standard consumption-based asset pricing model in which agents learn the location parameter of the dividend growth process through recency-biased inference, providing microfoundations to investors' decisions, allows for the derivation of a closed-form solution for stock prices. This solution clarifies how the latter depends on investors' expectations and how this relationship triggers fluctuations in the price-dividend ratio and thus generates the potential for non-fundamental booms and busts. The specificity of the model is that the size and the persistence of these fluctuations over time are caused by the representative investor's recency bias, which relates in this setting to cognitive limitations to pay full attention to past data over time, as suggested by growing empirical evidence.

Even with a small degree of parameter uncertainty, the model proves able to replicate the boom-and-bust episode in the US stock market in the run-up to the subprime crisis. The price-dividend ratio shows significant volatility over time and evolves according to surprise effects – thus showing a steep decrease at the beginning of 2008 –, it is strongly auto-correlated, and it is positively correlated with expected future returns. The model also replicates quantitative features of the dynamics of forecast errors when it comes to predicting future stock returns.

Therefore, this paper shows that a seemingly simple assumption, supported by thorough empirical evidence, can help reconcile theoretical predictions of the most standard and simplest possible consumption-based asset pricing model with both qualitative and quantitative features of the data in the recent boom-and-bust period in the US stock market.

Because they emphasize the role of expectations, our results suggest that communication policies about the true fundamental process might be necessary to mitigate fluctuations in asset prices. Public information disclosure aimed at warning market participants against asset price misalignments has become a new financial risk management tool to mitigate expectations-driven fluctuations. Thus, from mid-2002 onwards, officials from the Reserve Bank of Australia made public statements highlighting the risk that the steep increase in asset prices and leverage ratios observed in the early 2000s was not driven by fundamental factors. Similarly, the Swedish Riskbank and the Bank of Canada recently expressed concerns about the increase in house prices. Embedding our learning model into a production economy to investigate the welfare effects of non-fundamental fluctuations in asset prices is therefore an important avenue for future research that would allow for the study of the impact of such recent communication policies on non-fundamental asset price fluctuations.

References

Adam, K., Beutel, J., and Marcet, A. (2016a). Stock Price Booms and Expected Capital Gains. *Mimeo*.

Adam, K., Marcet, A., and Nicolini, J. (2016b). Stock Market Volatility and Learning. *Journal of Finance*, 71(1): 33-82.

Agarwal, S., Driscoll, J., Gabaix, X., and Laibson, D. (2013). Learning in the Credit Card Market. *NBER Working Paper*, 13822.

Ashton, A. and Ashton, R. (1988). Sequential Belief Revision in Auditing. *The Accounting Review*, 33(4): 18-42.

Bansal, R. and Shaliastovich, I. (2010). Confidence Risk and Asset Prices. *American Economic Review*, 100(2): 537-541.

Bansal, R. and Yaron, A. (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance*, 59(4): 1481-1509.

Bullard, J. and Duffy, J. (2001). Learning and Excess Volatility. *Macroeconomic Dynamics*, 5(2): 272-302.

Cheung, Y. and Friedman, D. (1997). Individual Learning in Normal Form Games: Some Laboratory Results. *Games and Economic Behavior*, 19(1): 46-76.

De Bondt, W. and Thaler, R. (1990). Do Security Analysts Overreact? *American Economic Review Papers and Proceedings*, 80(2): 52-57.

Erev, I. and Haruvy, E. (2013). *The Handbook of Experimental Economics*, volume 2, chapter Learning and the Economics of Small Decisions. Roth, A.E. & Kagel, J. Eds.

Gallagher, J. (2014). Learning from an Infrequent Event: Evidence from Flood Insurance Take-up in the United States. *American Economic Journal: Applied Economics*, 6(3): 206-233.

Greenwood, R. and Shleifer, A. (2014). Expectations of Returns and Expected Returns. *Review of Financial Studies*, 27(3): 714-746.

Guidolin, M. and Timmermann, A. (2007). Properties of Equilibrium Asset Prices Under Alternative Learning Schemes. *Journal of Economic Dynamics and Control*, 31(1): 161-217.

Haldane, A. (2009). Why Banks Failed the Stress Test ? In *Speech given at the Marcus-Evans Conference on Stress-Testing*.

Koulovatianos, K. and Wieland, W. (2011). Asset Pricing Under Rational Learning About Rare Disaster Risk. *CEPR Discussion Working Paper*, 8514.

Kuang, P. (2014). A Model of Housing and Credit Cycles with Imperfect Market Knowledge. *European Economic Review*, 70: 419-437.

LeRoy, S. and Porter, R. (1981). The Present Value Relation: Tests Based on Implied Variance Bounds. *Econometrica*, 49(3): 555-574.

Lovell, C. (1986). Tests of the Rational Expectations Hypothesis. *American Economic Review*, 76(1): 110-124.

Lucas, R. (1978). Asset Prices in an Exchange Economy. *Econometrica*, 46(6): 1429-1445.

Nakov, A. and Nuno, G. (2015). Learning from Experience in the Stock Market. *Journal of Economic Dynamics and Control*, 52: 224-239.

Pesaran, H., Pettenuzzo, D., and Timmermann, A. (2007). Learning, Structural Instability, and Present Value Calculations. *Econometric Reviews*, 26: 253-288.

Shaliastovich, I. (2015). Learning, Confidence and Option Prices. *Journal of Econometrics*, 187(1): 18-42.

Shiller, R. (1981). Do Stock Prices Move too much to be Justified by Subsequent Changes in Dividends? *American Economic Review*, 71(3): 421-436.

Shiller, R. (2000). Measuring Bubble Expectations and Investor Confidence. *Journal of Psychology and Financial Markets*, 1(1): 49-60.

Timmermann, A. (1993). How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices? *Quarterly Journal of Economics*, 108(4): 1135-1145.

Timmermann, A. (1996). Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning. *Review of Economic Studies*, 63(4): 523-557.

Tuttle, B., Coller, M., and Burton, F. (1997). An Examination of Market Efficiency: Information Order Effects in a Laboratory Market. *Accounting, Organizations and Society*, 22(1): 89-103.

Williams, J. (2014). Financial Stability and Monetary Policy: Happy Mariage or Untenable Union? *FRBSF Economic Letters*, 17.

A Proof of equation 2.1

The first Euler equation (with respect to the quantity of stocks S_t), including the market-clearing condition, is:

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

for $1 \leq t \leq J-1$. Isolating P_t on the left hand side yields:

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

Substituting P_{t+1} by its expression in the iterated forward version of the previous equation leads to:

$$P_t = \beta E_t \left[\left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} \beta E_{t+1} \left[\left(\frac{D_{t+2}}{D_{t+1}} \right)^{-\gamma} (P_{t+2} + D_{t+2}) \right] + \left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} D_{t+1} \right].$$

Applying the law of iterated expectations with nested conditioning sets ($E_t \left[E_{t+1}(X) \right] = E_t[X]$) and iterating forward again yields:

$$P_t = E_t \left[\beta^{J-t} \left(\frac{D_J}{D_t} \right)^{-\gamma} P_J \right] + E_t \left[\beta \left(\frac{D_{t+1}}{D_t} \right)^{-\gamma} D_{t+1} + ... + \beta^{J-t} \left(\frac{D_J}{D_t} \right)^{-\gamma} D_J \right].$$

In the last period J, under the non-bequest assumption – according to which all remaining wealth at the beginning of period J is consumed in J – stocks are no longer traded and therefore $E_t \left[\beta^{J-t} \left(\frac{D_J}{D_t} \right)^{-\gamma} P_J \right] = 0$. Finally,

$$P_t = E_t \left[\sum_{j=1}^{J-t} \beta^j \left(\frac{D_{t+j}}{D_t} \right)^{1-\gamma} D_t \right].$$

Thus,

$$E_t \left[\left(\frac{D_{t+j}}{D_{t+j-1}} * \frac{D_{t+j-1}}{D_{t+j-2}} * ... * \frac{D_{t+1}}{D_t} \right)^{1-\gamma} \right]$$
$$= E_t \left[E_{t+1} \left[\left(\frac{D_{t+1}}{D_t} \right)^{1-\gamma} * ... * E_{t+j-1} \left[\left(\frac{D_{t+j}}{D_{t+j-1}} \right)^{1-\gamma} \right] \right] \right].$$

Hence, when d and σ are known:

$$E_{t+j-1}\left[\left(\frac{D_{t+j}}{D_{t+j-1}}\right)^{1-\gamma}\right] = E_{t+j-2}\left[\left(\frac{D_{t+j-1}}{D_{t+j-2}}\right)^{1-\gamma}\right] = E_t\left[\left(\frac{D_{t+1}}{D_t}\right)^{1-\gamma}\right]$$
$$= \exp(d(1-\gamma) + \frac{(1-\gamma)^2\sigma^2}{2}) = \theta.$$

Therefore,

$$E_t\left[\left(\frac{D_{t+j}}{D_t}\right)^{1-\gamma}\right] = \theta^j,$$

and

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j \theta^j.$$

The stock price is the sum of the J-t first terms of a geometric sequence with common ratio $\beta\theta$ and first term $\beta\theta$. Therefore,

$$P_t = \frac{\beta\theta - (\beta\theta)^{J-t+1}}{1-\beta\theta} D_t.$$

B Inconsistent features of the rational expectations model

Constant versus highly volatile price-dividend ratio Figure 3.B.1 shows the US S&P 500 monthly price-dividend ratio over the recent period and the price-dividend ratio generated by the model with rational expectations, which is (roughly) constant (the parameters values used here are the same as those presented in the simulation exercise in Section 4). In contrast, over the period, the price-dividend ratio shows large variations and thus excess volatility relative to the prediction of the rational expectations model, suggesting that non-fundamental fluctuations in asset prices do arise.

Volatile stock returns Figure 3.B.2 presents realized monthly returns on US S&P 500 stocks. The returns in the rational expectations model exhibit very low volatility, which strikingly contradicts the data.

Positive correlation between one-year ahead time-varying expected stock returns (CFO Survey) and the current price-dividend ratio Figure 3.B.3 shows that expected returns from the CFO

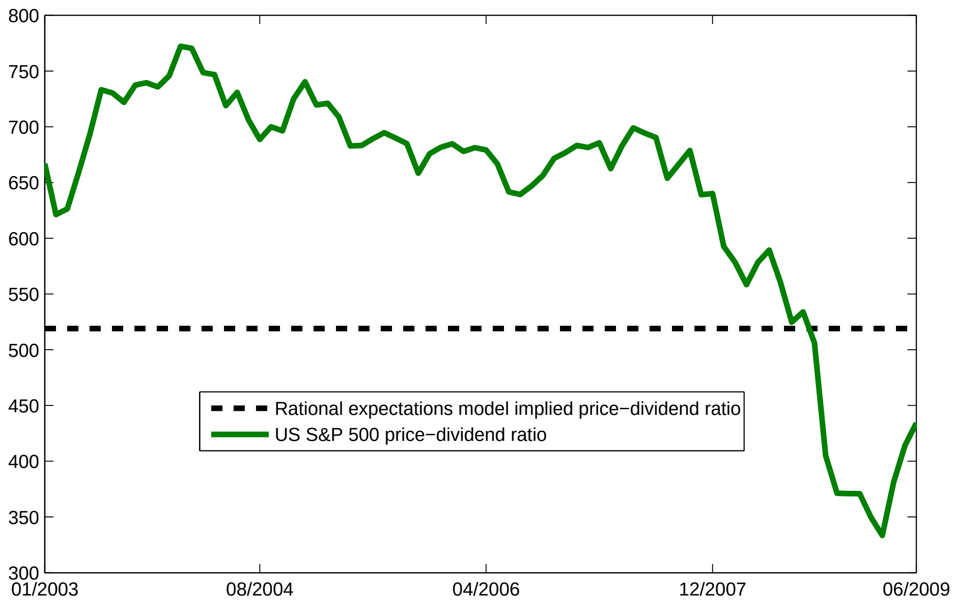

Figure 3.B.1: Monthly US S&P 500 price-dividend ratio versus price-dividend ratio in the rational expectations benchmark model

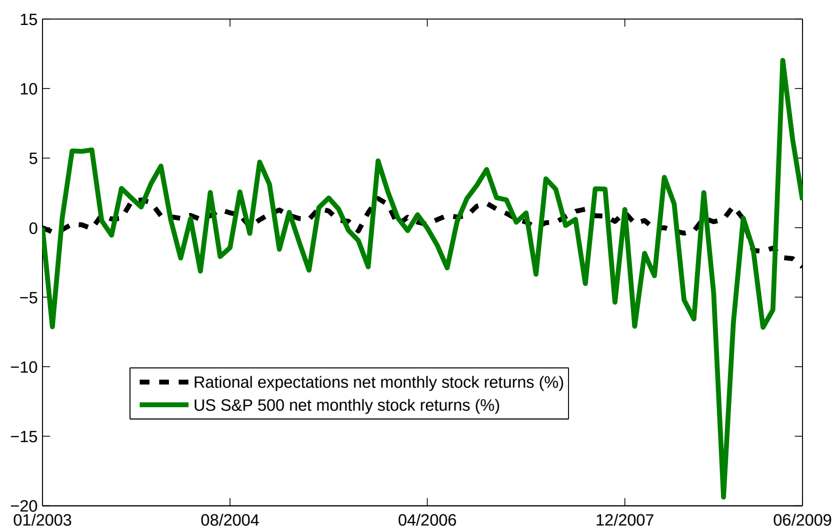

Figure 3.B.2: Monthly stock market returns: data versus rational expectations model

survey are not constant over time and are positively correlated with the price-dividend ratio. The CFO survey relies on a questionnaire on a broad range of issues, delivered quarterly to US senior financial executives in the public and the private sector, and encompasses a wide range of indus-

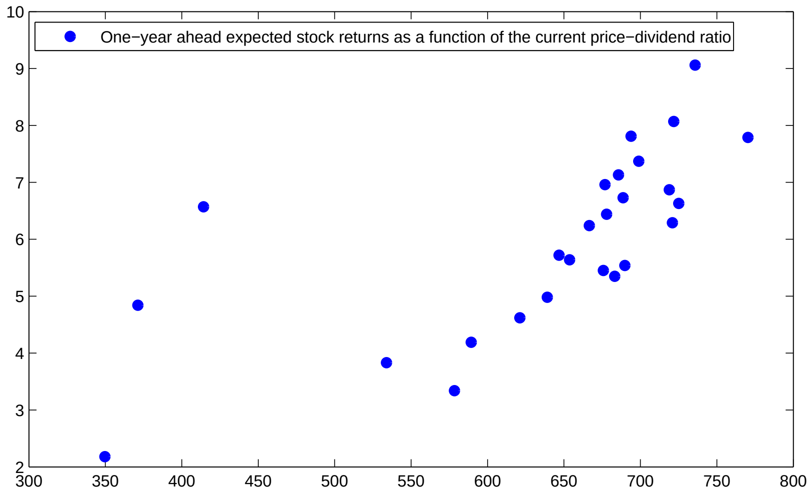

Figure 3.B.3: Expected stock returns (%) as a function of the price-dividend ratio

tries, geographic sectors and revenues. In particular, financial executives are asked a 1-year ahead expected stock return. There exists alternative survey data of expected stock returns. However, Greenwood and Shleifer (2014) show that all measures are highly positively correlated.

C Derivation of the posterior distribution under recency-biased learning

Under recency-biased learning, in period t, the precision of the prior distribution is discounted by α^t. Thus, the normal prior distribution $P(d \mid I_0, \sigma)$ is:

$$P(d \mid I_0, \sigma) = \sqrt{\frac{\alpha^t \tau_0}{2\pi}} \exp(-\frac{1}{2}\alpha^t \tau_0 (d - m_0)^2).$$

Similarly, the precision of the realization of the data y_k in period $k < t$ is discounted by α^k in the joint likelihood:

$$L(y^t \mid d, \sigma) = \prod_{k=0}^{t-1} \sqrt{\frac{\alpha^k \tau}{2\pi}} \exp(-\frac{1}{2}\alpha^k \tau (y_{t-k} - d)^2).$$

Therefore, the posterior distribution $P(d \mid I_t, \sigma)$ is:

$$P(d \mid I_t, \sigma) \propto \sqrt{\frac{\alpha^t \tau_0}{2\pi}} \prod_{k=0}^{t-1} \sqrt{\frac{\alpha^k \tau}{2\pi}} \exp(-\frac{1}{2}\alpha^t \tau_0 (d - m_0)^2)$$
$$\exp\left(-\frac{1}{2}\tau \sum_{k=0}^{t-1} \alpha^k \left((y_{t-k} - \bar{y_W})^2 + \sum_{k=0}^{t-1} \alpha^k (\bar{y_W} - d)^2\right)\right),$$

where $\bar{y_W} = \frac{\sum_{k=0}^{t-1} \alpha^k y_{t-k}}{\sum_{k=0}^{t-1} \alpha^k}$ is the weighted average of past observations. This yields:

$$P(d \mid I_t, \sigma) \propto \exp\left(-\frac{1}{2}\left(\alpha^t \tau_0 (d - m_0)^2 + \tau \sum_{k=0}^{t-1} \alpha^k (d - \bar{y_W})^2\right)\right).$$

Completing the square thus leads to:

$$P(d \mid I_t, \sigma) \propto \exp\left(-\frac{1}{2}\left(\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k\right)\left(d - \frac{\alpha^t \tau_0 m_0 + \tau \sum_{k=0}^{t-1} \alpha^k \bar{y_W}}{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k}\right)^2\right).$$

Therefore,

$$P(d \mid I_t, \sigma) \sim N\left(\frac{\alpha^t \tau_0 m_0 + \tau \sum_{k=0}^{t-1} \alpha^k y_{t-k}}{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k}, \frac{1}{\sqrt{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k}}\right),$$

i.e.,

$$P(d \mid I_t, \sigma) \sim N\left(\frac{\alpha^t \tau_0 m_0 + \tau \sum_{k=0}^{t-1} \alpha^k y_{t-k}}{\alpha^t \tau_0 + \tau \frac{1-\alpha^t}{1-\alpha}}, \frac{1}{\sqrt{\alpha^t \tau_0 + \tau \frac{1-\alpha^t}{1-\alpha}}}\right)$$

if and only if $\alpha < 1$.

D Derivation of equation 12

Starting from the Euler equation with respect to stocks under the subjective probability measure, isolating P_t on the left-hand side and iterating forward as in Proof of Proposition 1 yields:

$$P_t = \sum_{j=1}^{J-t} \beta^j D_t E_t^P \left[\left(\frac{D_{t+1}}{D_t}\right)^{1-\gamma} * E_{t+1}^P \left[\left(\frac{D_{t+2}}{D_{t+1}}\right)^{1-\gamma} * ... * E_{t+j-1}^P \left[\left(\frac{D_{t+j}}{D_{t+j-1}}\right)^{1-\gamma}\right]\right]\right],$$

$$= \sum_{j=1}^{J-t} \beta^j D_t E_t^P \left[E_{t+1}^P \left[E_{t+2}^P \left[...E_{t+j-1}^P \left[\left(\frac{D_{t+1}}{D_t}\right)^{1-\gamma} * \left(\frac{D_{t+2}}{D_{t+1}}\right)^{1-\gamma} * ... * \left(\frac{D_{t+j}}{D_{t+j-1}}\right)^{1-\gamma}\right]\right]\right]\right],$$

for $1 \leq t \leq J-1$. Under adaptive bayesian learning, i.e., when agents take into account the uncertainty about their estimates but not the probability distributions of future beliefs (agents do not take into account the fact that their current beliefs may change following future realizations of the dividend growth rate), the law of iterated expectations can still be applied: $E_t[E_{t+1}(X)] = E_t[X]$.

Eventually, this yields:

$$P_t = \sum_{j=1}^{J-t} \beta^j D_t E_t^P \left[\left(\frac{D_{t+j}}{D_t} \right)^{1-\gamma} \right].$$

E Comparative statics

We showed that $P_t = D_t \sum_{j=1}^{J-t} \beta^j \exp\left[(1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)\right]$.
For convenience purposes, we now write $\beta^j \exp[(1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)] = x_t$.

$$\frac{\partial x_t}{\partial \beta} = j\beta^{j-1} \exp[(1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)] > 0.$$

When the discount rate is higher – that is, preference for current consumption is lower –, the demand for stock prices increases (because more stock holdings allow for an increase in future consumption) and thus the equilibrium stock price increases relative to dividends.

$$\frac{\partial x_t}{\partial \gamma} = \beta^j j \exp[(1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2))(-m_t j - (1-\gamma)j(\sigma_t^2 j + \sigma^2)].$$

$$\frac{\partial x_t}{\partial \gamma} < 0 \Leftrightarrow m_t > (\gamma - 1)(\sigma_t^2 j + \sigma^2).$$

The sign of the derivative of each term of the sum with respect to the relative risk aversion coefficient (or equivalently with respect to the inverse of the intertemporal elasticity of substitution) depends on the hyperparameters m_t and σ_t in period t. When relative risk aversion increases (that is, when the intertemporal elasticity of substitution decreases), each term of the sum decreases if and only if the expected mean of the growth rate of the payoff on stocks is high enough.

$$\frac{\partial x_t}{\partial m_t} = \beta^j j \exp[(1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)](1-\gamma).$$

$\frac{\partial x_t}{\partial m_t} > 0 \Leftrightarrow \gamma < 1$. When the substitution effect exceeds the wealth effect, demand for stocks increases when the expected mean of the growth rate of the payoff on stocks increases.

F Expected Returns

Conditionally on information up to time t and on the value of the dividend process parameters d and σ, expected returns write:

$$E_t\left[\frac{P_{t+1}+D_{t+1}}{P_t}\mid d,\sigma\right]=\frac{1}{P_t}\left(E_t\left[P_{t+1}\mid d,\sigma\right]+D_tE_t\left[\frac{D_{t+1}}{D_t}\mid d,\sigma\right]\right).$$

In particular,

$$E_t[P_{t+1}\mid d,\sigma]=E_t\left[D_{t+1}\sum_{j=1}^{J-t-1}\beta^jE_{t+1}\left[\left(\frac{D_{t+1+j}}{D_{t+1}}\right)^{1-\gamma}\mid d,\sigma\right]\mid d,\sigma\right]$$

$$=D_t\sum_{j=1}^{J-t-1}\beta^jE_t\left[\frac{D_{t+1}}{D_t}E_{t+1}\left[\left(\frac{D_{t+1+j}}{D_{t+1}}\right)^{1-\gamma}\mid d,\sigma\right]\mid d,\sigma\right]$$

$$=D_t\sum_{j=1}^{J-t-1}\beta^j\exp[d((1-\gamma)j+1)+0.5((1-\gamma)^2j+1)\sigma^2].$$

Let's now assume that the parameter d is no longer known and is believed to follow a normal distribution with parameters m_t and σ_t in period t. In this case,

$$E_t^P[P_{t+1}\mid\sigma]=D_t\sum_{j=1}^{J-t-1}\beta^j\exp[m_t((1-\gamma)j+1)+0.5((1-\gamma)j+1)^2\sigma_t^2+0.5((1-\gamma)^2j+1)\sigma^2],$$

when assuming that the law of iterated expectations still applies, because agents are adaptive bayesian learners.

Eventually,

$$E_t^P\left[\frac{P_{t+1}+D_{t+1}}{P_t}\mid\sigma\right]\frac{D_t}{P_t}[\sum_{j=1}^{J-t-1}[\beta^j\exp(m_t((1-\gamma)j+1)+0.5((1-\gamma)j+1)^2\sigma_t^2$$
$$+0.5((1-\gamma)^2j+1)\sigma^2)]+\exp(m_t+0.5(\sigma_t^2+\sigma^2))].$$