Journal of Economic Behavior and Organization 148 (2018) 315–343

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Equity return predictability, time varying volatility and learning about the permanence of shocks☆

Daniel L. Tortorice

Department of Economics and Accounting, College of the Holy Cross, 1 College St. Worcester, MA 01610, United States

ARTICLE INFO

Article history:
Received 22 June 2017
Revised 12 December 2017
Accepted 4 January 2018
Available online 2 February 2018

JEL classification:
D83
D84
E21
G12

Keywords:
Consumption
Savings
Asset pricing
Learning
Expectations

ABSTRACT

I consider a consumption based asset pricing model where the consumer does not know if shocks to dividends are stationary (temporary) or non-stationary (permanent). The agent uses a Bayesian learning algorithm with a bias towards recent observations to assign probability to each process. While the true process is stationary, the consumer's beliefs change as he misinterprets a drift in dividends from their steady state value as an increased likelihood that the dividend process is non-stationary. Belief changes result in large swings in asset prices which are subsequently reversed. The model then is consistent with a broad array of asset pricing puzzles. It predicts the negative correlation between current returns and future returns and the PE ratio and future returns. Consistent with the data, I also find that consumption growth negatively correlates with future returns and the PE ratio and consumption growth forecast future consumption growth. The model amplifies return volatility over the benchmark rational expectations case and exactly matches the standard deviation of consumption. Finally, the model generates time varying volatility consistent with the data on quarterly equity returns.

1. Introduction

One of the key challenges for macro-finance models is to explain observations on equity returns that appear to be at odds with simple rational expectations models. For example, the PE ratio exhibits mean reversion, returns appears to have predictive power for future returns, returns are much more volatile than dividends and we observe large increases and decreases in asset prices which are hard to justify with news on fundamentals. Much work has attempted to square rational expectations models with these asset pricing puzzles. However, there is still no completely satisfactory explanation for what has been observed in the data. Therefore in this paper, I propose a novel learning model which is consistent with the observed negative correlation and amplifies the volatility of asset prices and returns with respect to fundamentals. The model also endogenously generates time varying volatility consistent with quarterly U.S. return data.

I consider a dynamic economy where an agent chooses to consume and invest. For investment, he can invest in a risk free asset or an equity asset which pays a dividend. Optimal behavior implies that the price of the equity asset is related

☆ I thank Jens Hilscher, David Laibson, Blake LeBaron, Davide Pettenuzzo, the editor, two anonymous referees and participants at the American Economic Association Meetings, Eastern Economic Association Meetings, Eastern Finance Association Meetings, Midwest Finance Meetings, Southwest Finance Meetings, Midwest Macroeconomics Meetings, Society for Nonlinear Dynamics and Econometrics Conference and the Brandeis University Economics Department Internal Faculty Research Seminar for helpful comments and suggestions. I thank the Tomberg Fund and the Theodore and Jane Norman Fund at Brandeis University for financial support. As always, any errors are mine.

E-mail address: dtortori@holycross.edu

https://doi.org/10.1016/j.jebo.2018.01.003
0167-2681/

to the present discounted value of the dividends. However, the agent does not know if the true process for dividends is stationary (so shocks are temporary) or non-stationary (so shocks are permanent). As a result he must learn about the true process and his beliefs have significant effects on equilibrium asset prices.

Two results make this learning significant. First, if shocks are permanent there is a much larger impact of a shock on the present discounted value of dividends than if shocks are temporary: (see for example Deaton, 1992). Therefore the agent's beliefs have a large impact on the price of equity. Second, inference in the model is quite difficult. Though the two models have very different implications for the long run effect of shocks, distinguishing between them in samples the length of the US macroeconomic time series is quite difficult: (see for example Cochrane, 1988; Stock, 1991). Unit root and near unit root process often have very similar short run dynamics and tests using their long run dynamics have very lower power. As a result, it is hard to distinguish between stationary and non-stationary processes and the agent may often hold incorrect beliefs.

To model learning, I assume that the true process for dividends is stationary, but the agent does not know this. Instead they use the Bayesian learning model of Cogley and Sargent (2005). In this model, the learner updates both the parameters on his candidate models and the probability that each model is true. Since dividends are an exogenous process in my model, convergence to the true model is ensured. Consequently, I adapt the standard Bayesian learning model to overweight recent observations in a way analogous to constant gain learning in the least squares learning literature (Evans and Honkapohja, 2001).

The learning mechanism substantially affects the dynamics of asset prices. After a random series of shocks, where the dividend drifts away from its steady state value, the agent put a significant probability weight on the non-stationary process. This change in beliefs generates a large swing in asset prices that is subsequently reversed as the stationary dividend drifts back to its long run average. These dynamics allow the model to explain a broad array of asset pricing puzzles. The model explains the observed negative correlation between excess returns today and future excess returns. In the model, the price-to-earnings (PE) ratio is negatively correlated with future returns and dividend growth. Furthermore, the model generates excess volatility of returns.

The model also makes several accurate predictions for consumption, which is determined endogenously in the model in contrast to many papers in the macro asset pricing literature. Consumption growth is negatively correlated with future returns and future consumption growth. The PE ratio also predicts lower future consumption growth. Finally, the model exactly matches the observed volatility of consumption in the data.

Additionally, the model is consistent with the time varying volatility of returns observed in the data. In quarterly returns data, I present evidence of excess kurtosis, positive autocorrelation of squared returns, and significant GARCH estimates of time varying volatility. I show that the learning model is able to explain all of these observations. When the weight on the non-stationary model increases, the volatility of asset returns rises. In contrast, the rational expectations benchmark and a model which puts a non-zero but constant probability on both the stationary and non-stationary model are unable to explain these facts. The model also provides a new explanation for large swings in equity prices. Changes in beliefs concerning the permanence of shocks, driven by random changes in dividends, lead to large increases and decreases in equity prices.

The current paper relates to many strands of the literature. First it relates to the empirical literature on asset pricing puzzles. The ability of returns to forecast future returns is stressed by Fama and French (1988), Poterba and Summers (1988), and Lakonishok et al. (1994). The ability of the price to earnings ratio to forecast future returns is noted by Campbell and Shiller (1988) and the ability of consumption to forecast future returns and consumption is highlighted in the work of Lettau and Ludvigson (2010). The observation of excess volatility stems from the work of Shiller (1981). These papers, among others, have spawned an enormous amount of theoretical work aimed at explaining these observations.

While there are many attempts to explain these puzzles in a rational expectations framework, two of the best know are Campbell and Cochrane (1999) and Bansal and Yaron (2004). Like these papers, my paper is a consumption based asset pricing model that attempts to resolve these empirical puzzles. But my model also differs in important respects. In both these papers, the agents exactly understand the structure of the economy and base their expectations on the true structure of the economy. In my model, agents do not know the true dividend process and therefore are forming expectations based on incorrect beliefs. My approach is supported by survey evidence, discussed below, which suggests that investors do not correctly anticipate forecastable movements in returns.

Many papers examine the asset pricing implications of learning. While a complete list of all such papers would be out of place here – see Pastor and Veronesi (2009) for a survey – I will highlight some of the most relevant papers. Barsky and De Long (1993) and Timmermann (1993) examine learning about the growth rate of dividends. The present paper differs from these works in considering a consumption based asset pricing model and examining the implication of learning for a wider range of asset pricing puzzles. Additionally, the nature of learning is different. In their work agents learn about the growth rates of dividends while in my model agents are learning about the permanence of shocks.

Collin-Dufresne et al. (2016) examine the ability of learning to generate long run risks in a consumption based asset pricing framework with Epstein-Zin preferences. They show that learning amplifies return volatility and the equity premium. The current paper differs in that consumption here is determined endogenously and my focus is on return predictability and

time varying volatility.[1] In a similar vein, Adam et al. (2017, 2012, 2016) generate return predictability in an asset pricing model with learning about the growth rate of prices when agents believe there are permanent shocks to the growth rate of prices. The current paper differs in important ways. First in my model consumption is endogenous, and therefore I can analyze its predictions for correlations between consumption and future returns. Secondly, the model here endogenously generates time varying volatility. Finally, in the Adam et. al papers agents forecast only one step ahead. However, if they knew that the Euler equation holds in all periods (as it must in equilibrium) then the price again becomes the present discounted value of dividends and individuals would then need to form long-run forecasts about fundamentals suggesting an important relevance of the approach in this paper.

Veronesi (1999) studies learning about the growth rate of dividends in a consumption based asset pricing model. He finds the model generates clustering of volatility and time-varying expected returns. The current paper differs in several important ways. First, Veronesi presents a risk based explanation for asset pricing puzzles. Learning introduces an additional source of risk (revisions in beliefs) that generates time varying volatility and expected returns. In this paper, return predictability is driven by overreaction to news and not time varying risk premia. This distinction is important because much survey evidence [e.g. Fisher and Statman, 2002; Greenwood and Shleifer, 2014; Shiller, 2000; Vissing-Jorgensen, 2004] finds high price-to-earnings ratios appear correlated with, if anything, higher forecasts of expected returns. Additionally, time varying volatility in the current paper comes from learning across two models: one where the present discounted value of fundamentals is very volatile and one in which it is not. Indeed, in Section 7, I show that 70% of the variance in prices changes remains even absent changes in beliefs. Finally, the current paper is more ambitious in the sense that it aims to quantitatively match facts about return and consumption predictability and also the auto-correlation of squared returns and GARCH estimates of time varying volatility.[2]

At least three works have considered how incorrect beliefs can explain asset pricing puzzles. Lam et al. (2000) consider a model where the agent has mistaken beliefs concerning the growth rate of consumption. Like my paper, incorrect beliefs are important for explaining asset pricing puzzles. However, the nature of misspecification is quite different. In their model, agents underestimate the persistence of shocks, while in my model they overestimate the persistence of shocks. Additionally, Lam et al. (2000) do not consider learning, hence their agent exogenously believes in an incorrect model and never considers revising his beliefs. Barberis et al. (1998) consider a model where dividends follow a random walk but the agent believes dividends follow either a mean reverting model or an extrapolation model. They use this model to explain underreaction and overreaction to news. As in my paper agents use Bayes' rule to change their beliefs in the likelihood of each model. However, in my paper agents are allowed to put some probability weight on the true model and I also address a broader range of asset pricing puzzles. The paper that in some ways is the most similar to mine is Fuster et al. (2012). I modify their model and target some of the same moments they do. In their paper they explain asset pricing puzzles by assuming that dividends are non-stationary, but with some long run mean reversion, while agents believe in a non-stationary model without mean reversion. My paper differs though in at least two important ways. In their model, the agent's belief in the incorrect model is exogenous and never changes. While in my model they are the result of a clearly specified learning rule. As a result, I can be explicit about the magnitude of the mistake the agent is making and agents are able to revise their beliefs if the incorrect model seems highly unlikely. Additionally, because beliefs change, this adds an additional source of volatility that is useful in explaining consumption and return volatility and is necessary to endogenously generate time varying volatility which is absent in the Fuster et al. (2012) model. For similar reasons this paper differs from Barberis et al. (2015) who explain similar moments as Fuster et al. but with a heterogeneous agent model where some agents exogenously extrapolate returns.

The rest of the paper proceeds as follows: Section 2 describes the data moments I attempt to explain, Section 3 presents a simplified version of the model which provides intuition and contrasts the model with other approaches in the literature, Section 4 describes the consumption based asset pricing model and the formulation of beliefs, Section 5 explains the model calibration and simulation, Section 6 highlights the main mechanism of the model and the results from simulating the model, Section 7 examines robustness of the results and the importance of changing beliefs for the key results and finally Section 8 concludes.

2. Data and data moments

Many authors (e.g. Campbell and Shiller, 1988; Fama and French, 1988; Lakonishok et al., 1994; Lettau and Ludvigson, 2010; Poterba and Summers, 1988) have noted weak to moderate predictability of stock returns. This predictability manifests itself in multiple ways. First, annual returns are negatively correlated with returns over the next few years. Similarly, the PE ratio is also negatively correlated with future returns. These patterns also emerge when one considers aggregate consumption growth. Annual consumption growth is negatively correlated with returns over the next few years, and consumption growth and the PE ratio are negatively correlated with future consumption growth. Additionally, in quarterly returns data

[1] In related work, Johannes et al. (2016), explore the return predictability implications of the Collin-Dufresne et al. (2016) framework and show they can generate return predictability using the same mechanism as Timmermann (1993), namely ex-ante rational overestimation of the dividend growth rate.

[2] The connection between time varying volatility and learning is also studied by Johnson (2001) who examines the implications for time varying volatility in a model with learning and occasional permanent shocks to the *growth* rate of dividends and David and Veronesi (2013) who examine the complex relationship between time varying volatility and fundamentals in a model with learning about the growth rate of earnings, consumption and inflation.

Table 1
Data moments.

Table 1a: Return predictability			
		Bootstrap procedure	
Moment	Data	Mean	90% CI
corr(r_t, r_{t+2} + – + r_{t+5})	−0.2	−0.01	[−0.19 0.17]
corr($P/E_{10,t}$, r_{t+2} + – + r_{t+5})	−0.41	−0.13	[−0.43 0.2]
corr(Δlnc_t, r_{t+2} + – + r_{t+5})	−0.34	−0.01	[−0.27 0.25]
corr($P/E_{10,t}$, Δlnc_{t+3} + – + Δlnc_{t+6})	−0.16	−0.03	[−0.45 0.36]
corr(Δlnc_t, Δlnc_{t+3} + – + Δlnc_{t+6})	−0.23	0.01	[−0.24 0.26]
corr($P/E_{10,t,}$ $\Delta lndt+2$ + – + Δlnd_{t+5})	−0.25	−0.02	[−0.20 0.16]
$\sigma(r_t)$	20.50%	–	–
$\sigma(\Delta lnc_t)$	2%	–	–
Table 1b: Time varying volatility			
Kurtosis	Data	Confidence bounds	
r_t	4.1	[2.58 3.46]	
$P/E_{10,\,t}$	4.6	[1.67 3.50]	
$\%\|r_t\| > 1.96^*\sigma(r_t)$	6.2%	[3% 7%]	
Autocorrelation of squared returns	Data	Confidence bounds	
lag 1	0.079	[−0.028 0.186]	
lag 2	0.01	[−0.097 0.117]	
lag 3	0.47	[0.36 0.58]	
lag 4	0.14	[0.033 0.25]	
GARCH	Data	Confidence Bounds	
Garch	0.61	[0.46 0.76]	
Arch	0.29	[0.17 0.41]	
p-value Engle test	0.048		

Note: This table reports the correlation of returns, the PE ratio and consumption growth with future returns, the correlation of the PE ratio and consumption growth with future consumption growth and the correlation of the PE ratio with future dividend growth. It reports the results of a bootstrapping exercise designed to illustrate the statistical significance of these correlations. It then reports the standard deviation of excess returns and the standard deviation of consumption growth. The table also reports kurtosis of excess returns and the PE ratio, autocorrelation of squared returns, and GARCH estimate for the excess returns data. Standard errors and confidence bounds from a bootstrapping exercise are given as well. For the autocorrelation and GARCH estimates 90% confidence bounds are reported as +/−1.65 standard errors.

there is clear evidence of time varying volatility of returns. (Time varying volatility has also been documented by many authors in the literature, for example see French et al., 1987 and Schwert, 1989.) In this section, I present evidence on return predictability in US data and I also present evidence on excess kurtosis, positive autocorrelation of squared returns and significant GARCH effects.

Evidence on return predictability is presented in Table 1a.[3] The table first reports the correlation between the current excess return r_t and the cumulative return over the next 2 to 5 years: $r_{t+2} + \cdots + r_{t+5}$. It is −0.2. This result indicates mild mean reversion in stock returns. Next I report the correlation between the current PE ratio $P/E_{10,\,t}$ and the same cumulative return over the next 2 to 5 years: $r_{t+2} + \cdots + r_{t+5}$. This correlation is more negative: −0.41. Again this result indicates some predictability of excess returns and some mean reversion in stock prices. I also find that consumption growth negatively predicts stock returns. The correlation between consumption growth today $\Delta \ln c_t$ and the same cumulative return $r_{t+2} + \cdots + r_{t+5}$ is −0.34. Additionally, I find that future cumulative consumption growth $\Delta \ln c_{t+3} + \cdots + \Delta \ln c_{t+6}$ is correlated with the current PE ratio $P/E_{10,\,t}$ with a correlation coefficient of −0.16 and current consumption growth $\Delta \ln c_t$ with a coefficient of −0.23. Finally, the table also reports that the PE ratio negatively forecasts dividend growth over the next few years. The correlation between $P/E_{10,\,t}$ and $\Delta \ln d_{t+2} + \cdots + \Delta \ln d_{t+5}$ is −0.25. The table also reports that the standard deviation of excess returns equals 20.5% and the standard deviation of consumption growth equals 2%.

To assess the statistical significance of the first six negative correlations, I calculate the distribution of these statistics under the null hypothesis of no-predictability via a bootstrap exercise the details of which are presented in appendix A. I report the mean and 5% and 95% percentiles for the statistics. The mean statistics are all near zero with the exception of the correlation between the PE ratio and future returns which has a mean correlation of −0.13. I find that the correlation between consumption growth and future returns and the PE ratio and future dividends are outside the 90% confidence interval. The correlations between the PE ratio and future returns, returns and future returns, and consumption growth and future consumption growth are at the lower bound of the 90% confidence interval. The correlation of the PE ratio with future consumption is within but towards the lower half of the confidence interval. While the time series is not very short, because the statistics involve multiple overlapping observations, they emit wide confidence intervals. These results indicate that while one would not strongly reject the null of no correlation, the moderate negative correlations we observe are unlikely under the null of zero correlation.

[3] See Appendix A for details on the data and calculations.

2.1. Time varying volatility

Table 1b examines the presence of time varying volatility in the returns and PE ratio data. Here I use quarterly data as it is difficult to detect time varying volatility at an annual frequency. I first report kurtosis for quarterly returns and the PE ratio. In the data quarterly return kurtosis equals 4.1. In contrast, if returns were normal one would expected kurtosis equal to 3. The PEratio also exhibits kurtosis. In quarterly data, kurtosis of the PE ratio is 4.6. Finally I also look at the percent of absolute returns which are greater than 1.96 times the standard deviation of returns. If returns were normal this statistic should be 5%. However, in the data it is 6.2%.

I also present results from the bootstrapping exercise designed to assess the statistical significance of the results. I find that the observed kurtosis levels for the return and PE series are outside the 95% percentile of 3.5. However the upper bound on percent of returns greater than 1.96 is 7% compared to the 6.2% found in the data. Therefore, the excess kurtosis of return and the PE ratio is unlikely if shocks are normally distributed, however the percent of returns above 1.96 standard deviations is not outside the confidence bound.

In Table 1b, I also examine the autocorrelation of squared returns as a way to measure time varying volatility. If this correlation is positive, then large (in magnitude) returns are likely to be followed by more large returns. We see that in the quarterly data, up to four lags all these autocorrelations are positive. The correlation is 0.08 at 1 lag, 0.01 at 2 lags, 0.47 at three lags and falls to 0.14 at four lags. The standard error for these autocorrealtions is 0.065 implying statistically significant autocorrelation at lags three and four.[4]

As a final check for time varying volatility of returns, I estimate GARCH(1,1) models on the quarterly return series and compare the predictions of the rational expectations models and the constant probability model to the learning model. The GARCH(1,1) model is:

$$\sigma_t^2 = \kappa + \gamma_1 \sigma_{t-1}^2 + a_1 \varepsilon_{t-1}^2.$$

In this model the variance of $\varepsilon_t = r_t - E(r_t)$ is varying over time. If γ_1 and a_1 are positive then the model predicts periods of particularly high volatility. For the quarterly return data I estimate $\gamma_1 = 0.61$ and $a_1 = 0.29$. Both estimates are highly statistically significant. Furthermore the Engle test (Engle, 1982) rejects the null of no GARCH effects at the 95% confidence level.

To summarize, there is evidence of moderate return predictability and also time varying return volatility in U.S data. Additionally, the PE ratio predicts consumption growth and consumption growth is negatively autocorrelated over medium horizons. I next describe a consumption based asset pricing model with learning that is consistent with these facts.

3. Analytical model

This section uses a simple asset pricing model to illustrate key asset pricing puzzles. It then provides intuition for the learning model in this paper and contrasts it with other proposed solutions in the literature.

Consider the simple asset pricing model, where the price of the asset is the present discounted value of dividends:

$$P_t = E_t \sum_{s=1}^{\infty} \frac{D_{t+s}}{(1+r)^s}$$

If dividends grow at a constant rate g we get the standard result of Gordon (1959):

$$P_t = D_t \frac{(1+g)}{r-g}$$

From this formula we can easily see some of the challenges of matching key moments of asset pricing data. The first is the excess volatility puzzle. According to this formula, all unpredictable movements in stock price are due only to innovations in the dividend process:

$$P_t - E_{t-1}P_t = \frac{(1+g)}{r-g}[D_t - E_{t-1}D_t].$$

First, since the standard deviation in fundamental (dividend news) is small, the model generates little volatility in asset price changes. In fact, since prices are proportional to dividends, the standard deviation of returns just equals the standard deviation of dividend growth. Second, r is a constant and therefore the model can not match facts on return predictability. Finally, time varying volatility is a key feature of the data, but the model can only generate time varying volatility with exogenous time varying volatility of the dividend.

3.1. Learning about the growth rate

Timmermann (1993) proposes learning about the growth rate of the dividend as a solution to some of these puzzles. In his model, dividends are a random walk:

$$\ln D_t = \mu + \ln D_{t-1} + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma^2)$$

[4] Standard errors are calculated as $\frac{1}{\sqrt{T}}$, the standard error under the null hypothesis of zero autocorrealtion. See Hamilton (1994) pp. 111.

Investors estimate the parameters of the dividend process with a recursive least squares algorithm. Based on those estimates, they form an estimate of the growth rate of dividends $\hat{g}_t = \hat{\mu}_t + \frac{\hat{\sigma}_t^2}{2}$. Asset prices then are given by:

$$P_t = D_t \frac{(1+\hat{g})}{r-\hat{g}}.$$

This modification helps solve the excess volatility puzzle. Again, the asset price is a multiple of the dividend. However, now revisions in the term multiplying the dividend also generates volatility in asset prices. Additionally, the dividend yield will be able to predict returns (ex-post). This effect occurs because times when the dividend growth rate is over-estimated result in high prices. These prices will then fall as investors revise down their growth expectations. While this model is successful at addressing these two puzzles, its predictions for time-varying volatility are less clear and are not explored by Timmermann. Another potential drawback of this approach is that agents act as if there are permanent shifts in the growth rate of dividends, which is perhaps a stronger assumption than warranted by the data.

3.2. Extrapolation

Barberis et al. (2015) explore the role that allowing for non-rational investors who extrapolate returns has in resolving asset pricing puzzles. In their model the change in the **level** of dividends is given by[5]:

$$\Delta D_t = g^d + \varepsilon_t.$$

Given that agents have constant absolute risk aversion preferences, Barberis et al. (2015) show the rational expectations equilibrium price – with no non-rational investors – is given by:

$$P_t = \frac{D_t}{r} + \frac{g_d}{r^2} - \phi$$

where r is the risk free rate and ϕ is a constant related to the agent's degrees of risk aversion and the variance of ε_t. The first two terms of this formula are just the present discounted value of dividends, though it takes a different form than the Gordon growth model because the dividend process specified in constant changes as opposed to a constant growth rate.

In their model, extrapolators forecast stock market returns as a weighted average of past stock market price changes:

$$S_t = \beta \sum_{s=t}^{-\infty} e^{-\beta(t-s)}[P_s - P_{s-1}].$$

They then show that the presence of extrapolators influences equilibrium asset prices so that the price is now given by:

$$P_t = A + BS_t + \frac{D_t}{r}$$

where A and B are constants and S_t follows the mean reverting process:

$$\Delta S_t = -\alpha\left(S_t - \frac{g_d}{r}\right) + \eta_t$$

This addition of extrapolators helps the model fit the data better along two dimensions. First, volatility is now higher because changes in S_t (which the authors call sentiment) creates additional volatility. The model also creates return predictability because S_t is mean reverting. High values of sentiment today, increase prices, but lead to lower prices in the future as sentiment mean reverts. However this formulation does not generate time varying volatility.

3.3. Long run risks

The long-run risks model of Bansal and Yaron (2004) generates excess volatility in a rational expectations model. Again one can understand their approach using a model where the equity price is a present discounted value of dividends. However, for analytical tractability it is easier to work with the log-linear approximation of Campbell and Shiller (1988). In this set up, the stock price and dividends (in logs) satisfies:

$$p_t - d_t = \frac{k}{1-\delta} + E_t \sum_{j=0}^{\infty} \delta^j (\Delta d_{t+1+j} - r)$$

where k and δ are the standard Campbell-Shiller approximating constants and r is the expected stock return (which we take to be constant). In the long run risks model dividend growth is given by:

$$\Delta d_{t+1} = \mu^d + x_t + u_t$$
$$x_t = \rho x_{t-1} + e_t$$

[5] For ease of comparison I represent their model in discrete time.

Using this dividend process, we can show that

$$\ln\left(\frac{P_t}{D_t}\right) - E_{t-1}\ln\left(\frac{P_t}{D_t}\right) = u_t + \frac{1}{1-\delta\rho}e_t$$

The model, then, can amplify volatility by allowing for shocks to the growth rate. And as $\rho \to 1$, these shocks have large effects on equilibrium asset prices. In this sense the model is very close to Timmermann (1993), though in the long-run risks model the growth rate of dividends is actually changing, where in Timmermann they are constant and only the investor's estimates are changing.

As the long-run risks model stands, it is unable to generate return predictability or time varying volatility.[6] To achieve these objectives, Bansal and Yaron add time varying volatility of fundamentals. This immediately gives time varying volatility of asset prices and returns. Additionally, because they use Epstein-Zin utility, the stochastic discount factor depends not only on current consumption growth but the whole future path of consumption growth. Because the x_t factor influences future consumption growth there is a high correlation between the stochastic discount factor and returns on equity. This effect creates a realistic equity premium and importantly one that varies over time as the variance of the fundamental shocks changes.

3.4. Model learning

The starting point for the model in this paper is also one where the equity price is given by a present discounted value of dividends:

$$P_t = E_t \sum_{s=1}^{\infty} \frac{D_{t+s}}{(1+r)^s}$$

The departure from the previous literature is that the investor does not know the correct model to forecast dividends. They consider two possibilities a random walk and an AR(1) process (which is the true process).[7]

$$D_{t+1} = \mu^{NS} + D_t + \varepsilon_t^{NS}$$
$$D_{t+1} = \mu^{S} + \rho D_t + \varepsilon_t^{S}$$

In this setup the equity price is given by a weighted average of the present discounted value under the two possible dividend processes:

$$P_t = (1-p_{s,t})\frac{\mu^{NS}\left[\frac{1+r}{r}\right] + D_t}{r} + p_{s,t}\frac{\mu^{S}\left[\frac{1+r}{r}\right] + \rho D_t}{1+r-\rho}$$

Here $p_{s,\,t}$ is the agent's belief that the AR(1) process is the true process.

From this equation one can see how the model is able to generate amplified volatility, time-varying volatility, and predictability of returns. First calculate the unexpected change in the price[8]:

$$P_t - E_{t-1}P_t = (1-p_{s,t})\frac{\varepsilon_t^{NS}}{r} + p_{s,t}\frac{\rho\varepsilon_t^{S}}{1+r-\rho} + (p_{s,t} - p_{s,t-1})\left(\frac{\mu^{S}\left[\frac{1+r}{r}\right] + \rho D_{t-1}}{1+r-\rho} - \frac{\mu^{NS}\left[\frac{1+r}{r}\right] + D_{t-1}}{r}\right).$$

The unexpected change in the price level is a probability weighted average of forecast errors from the two models, plus a catch-up term. The catch-up term captures the revision in the price that comes from updating your probability and is equal to the change in the probability weight times the difference in the prices implied last period by the two models of the dividend process.

Now we can see how the model amplifies volatility. First, when dividends are a random walk, the shock to fundamentals is amplified more than under the true process, the AR(1) process. Secondly, revisions to beliefs create large changes in prices when the random walk and AR(1) model give very different asset prices. Volatility will also be time varying because the first two terms will imply higher volatility when $(1-p_{s,t})$ is bigger and the second term will imply higher volatility when dividends are large. Last, returns will be predictable. Since the true process is an AR(1) process, the agent is over-reacting to fundamentals. For example after a positive shock, the price will jump up. But fundamentals will eventually mean revert, resulting in negative values for ε_t^{NS}, which will result in lower prices and returns in the future.

This simple model also makes clear the difference between Fuster et al. (2012) and this paper. Their paper can be understood as treating $p_{s,\,t}$ as a constant equal to zero. The result is excess volatility and return predictability, but volatility does not vary over time.

To provide additional intuition for the model in this paper, Fig. 1 plots the (log) dividend impulse response functions to a one standard deviation shock to the dividend.[9] On impact the effect of the shock on the dividend is the same across the two

[6] See Bansal and Yaron (2004) pp. 1486 for more details.

[7] The full model in the paper will allow for an arbitrary number of lags.

[8] Note we assume the agent treats his probability belief as a constant – and not a random variable – so that $E_{t-1}p_{s,t} = p_{s,t-1}$.

[9] The impulse response is estimated with an AR(4) consistent with the main model in the paper.

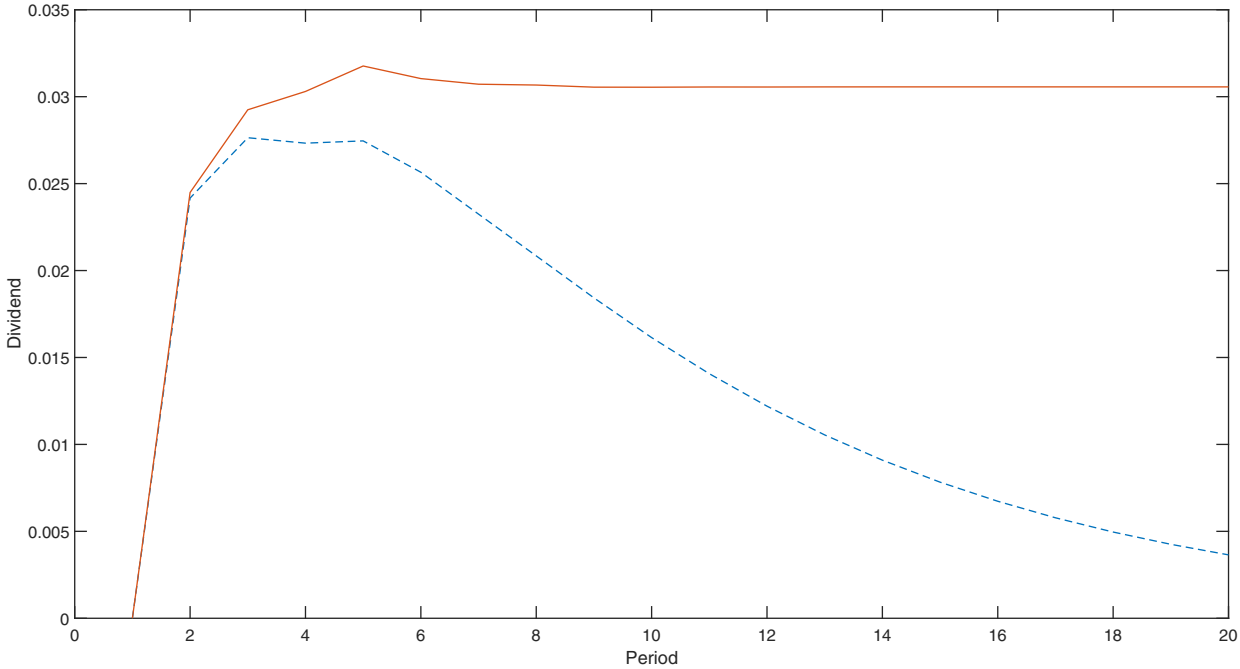

Fig. 1. Impulse responses: non-stationary model (solid) vs. stationary model (dashed).

models. One quarter out, both models also predict very similar effects on the dividend because the short term dynamics of the stationary model are very similar to the short run dynamics of the non-stationary model. However, after the first two quarters the paths start to diverge with the non-stationary model predicting an increasing and permanent effect on future dividends, while the stationary model predicts a temporary effect which dies out after about 20 quarters.

From these dynamics, we can deduce two effects. First, when the agent believes strongly in the non-stationary model the price will be quite volatile, as every shock generates a large change in the present value of future dividends. On the other hand, if the agent believes mostly in the stationary model, asset prices will be very smooth as shocks to fundamentals have only small effects on the present value of future dividends. Second, at times, it will be difficult to know for certain that the stationary model is true. Note that the model predictions a few quarters out are not very different. They are about one fifth of a standard deviation apart from each other. This implies that at certain times, given a random draw of the data generated by the stationary model, it will look as if dividends are more likely generated by the non-stationary model. This implication is a well know result from the literature on unit roots in macroeconomic time series (Cochrane, 1988; Deaton, 1992; Stock, 1991). While these models imply very different long run impacts of shocks, they are very hard to tell apart in time series the length of most macroeconomic series. Tests which construct the long run response of shocks lack statistical power in small samples. As a result, we rely on the short run dynamics of the parametric representations (e.g. an AR(4) in differences versus levels). But these models have similar short run dynamics and are therefore difficult to differentiate. It is this result that is at the heart of why learning matters in this model. These two processes are difficult to tell apart but also have very different implications as to how prices should respond to shocks to fundamentals.

To see the implication of learning about stationary versus non-stationary dividend processes, consider the sample price in Fig. 2. This price is based on a single simulation of the dividend process from the stationary model for 340 quarters. I plot the price from the learning model (solid line) versus the price from a model where the agent puts all the weight on the stationary model (dashed line). We see first that the learning model generates a more volatile price series then the stationary model since it puts some weight on the non-stationary model where shocks to dividends are permanent. More interestingly, we see two large price increases and subsequent crashes in the learning model around quarter 75 where prices rise cumulatively 10% and then subsequently crash and another around quarter 300 where prices rise 12% and then crash. In contrast the stationary model has almost no change in its price. These fluctuations look like fluctuations in price not driven by fundamentals (i.e. dividends).

To see what drives these price crashes examine Fig. 3. What one sees is that around time 75 the probability on the stationary model goes from 70% to 45%. Similarly, around time 300 the probability of the stationary model falls from 70% to about 30%. It is these changes in beliefs that lead to price increases.

Now recall that beliefs are endogenously determined in the model. This movement in probability is not an exogenous shock to beliefs but an endogenous response to the change in dividends. Fig. 4 plots the dividend process. Even though it is driven by the stationary process, the process is close enough to a unit root to occasionally have large drifts from its steady

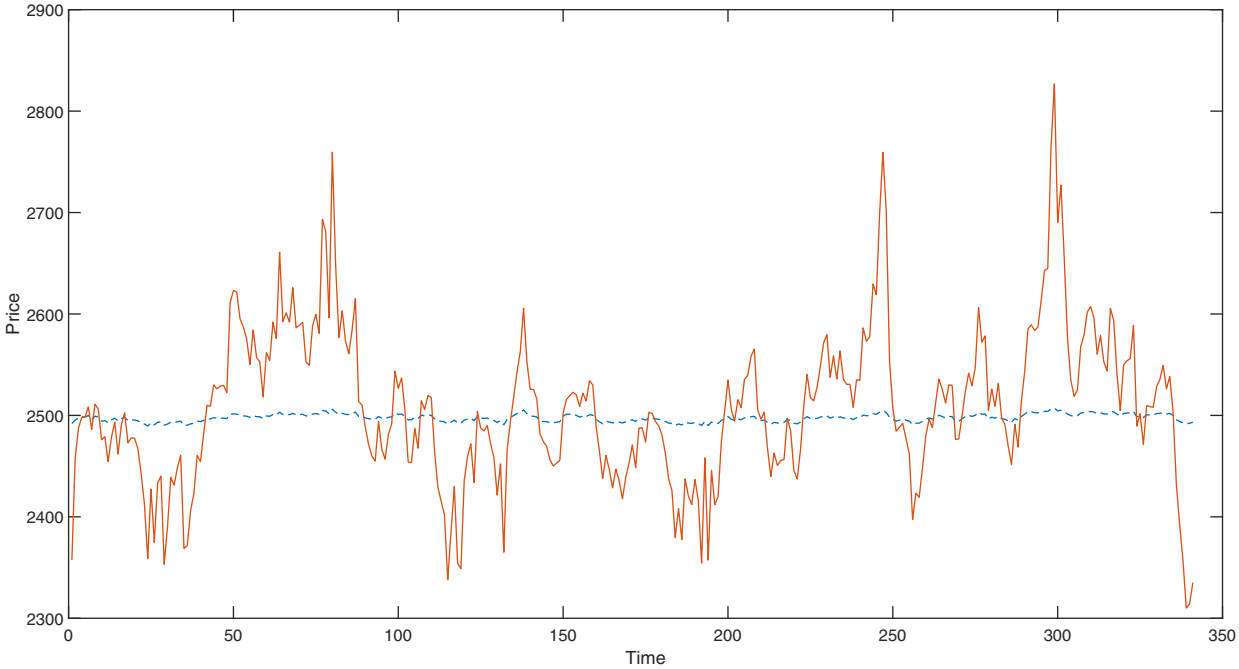

Fig. 2. Sample price: learning model (solid) vs. stationary model (dashed).

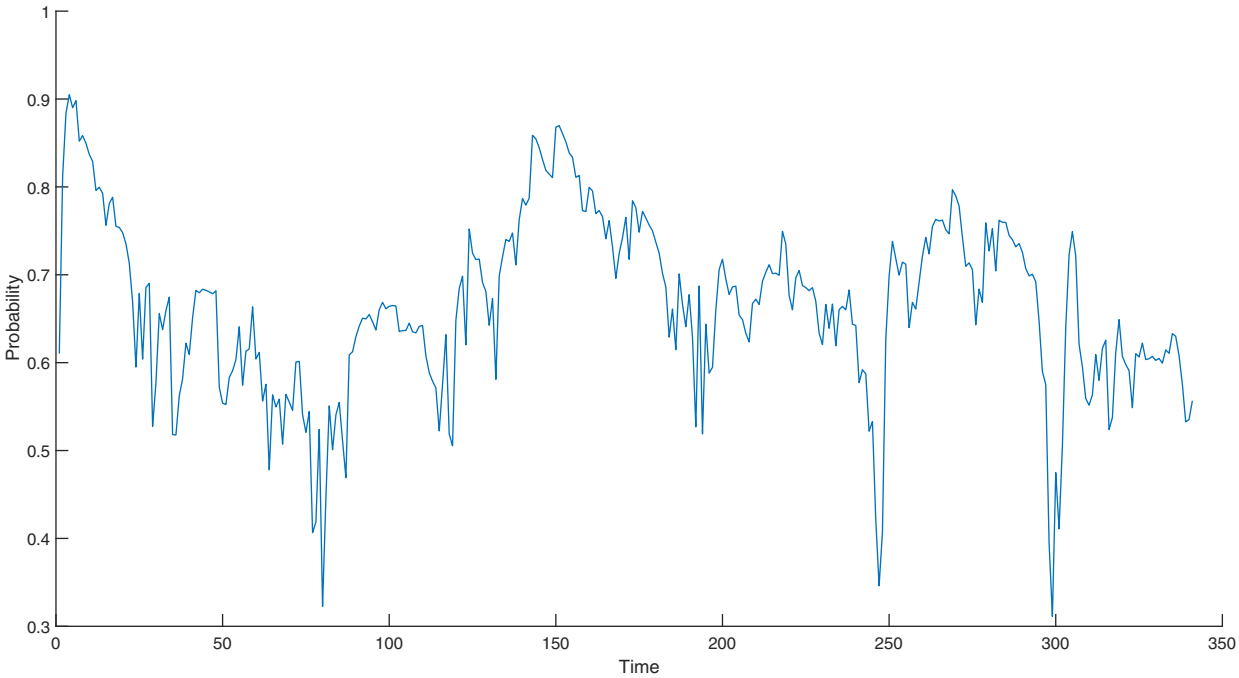

Fig. 3. Sample probability on the (true) stationary model.

state value. We see that around time 75, the dividend begins a sharp rise, increasing cumulatively 20%. At time 300 there is a similar increase in the dividend. While the true process is the stationary model, this level of the dividend is unlikely given the stationary model, and so the agent begins to think that the non-stationary model is true. These dividend changes lead to a massive reevaluation of beliefs and revaluation of price.

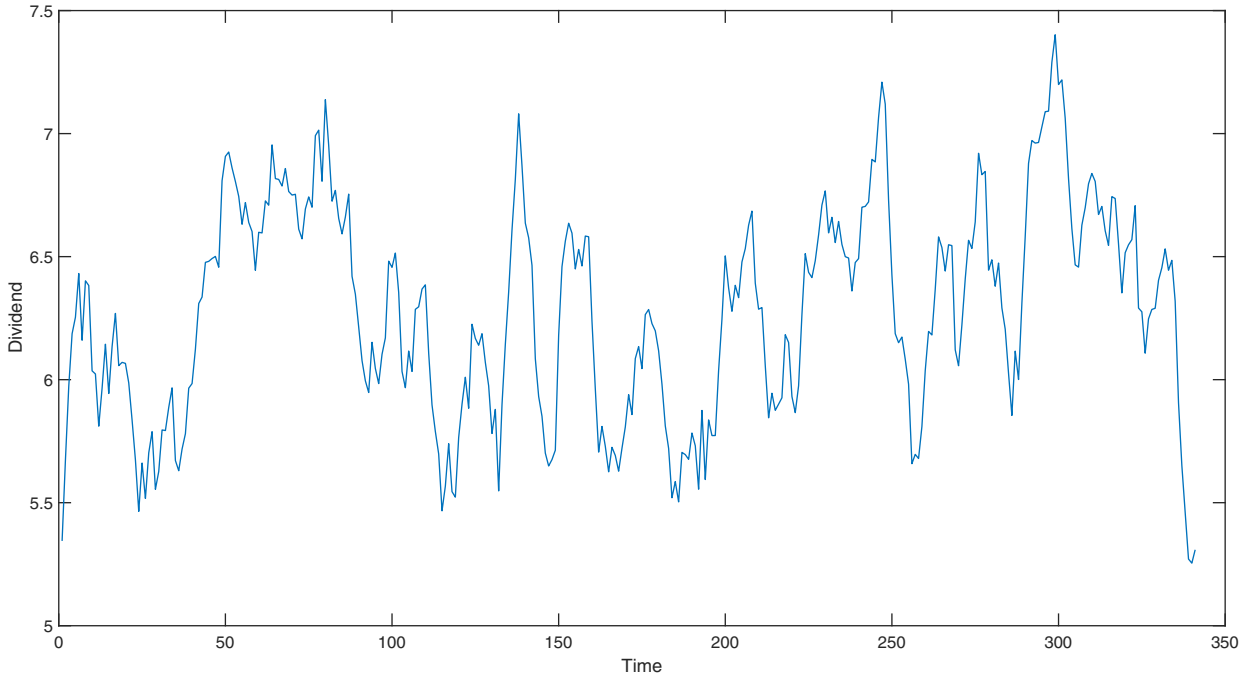

Fig. 4. Sample dividend .

While this is just one random draw of dividends, and we rely on the simulation statistics to evaluate the model, it explains how the model may be able to match the data. Agents misinterpret random movements in dividends generated by the stationary model. This misrepresentation leads to large swings in prices that are subsequently reversed when dividends begin to mean revert. Consider a large increase in the price due to this misinterpretation. We will see a positive return, an increase in the PE ratio and an increase in consumption as wealth goes up. However, when dividends subsequently mean revert, these positive increases will be followed by declines in the price resulting in negative returns and consumption growth. Additionally, learning should increase volatility and therefore lead to a higher standard deviation of consumption and returns.

3.5. Comparison with constant gain learning

Finally, one can contrast this set-up with learning about the parameters of the stationary process, μ^S and ρ but knowing that $p_{s,t} = 1$. For clarity imagine that learning involves only the intercept term. Then I can calculate[10]:

$$P_t - E_{t-1}P_t = \frac{\rho\varepsilon_t^S}{1+r-\rho} + \frac{\mu_t^S - \mu_{t-1}^S}{1+r-\rho}\left[\frac{1+r}{r}\right]$$

Comparing this expression to the model learning expression we can see two important differences between model learning and parameter learning. First, since the agent knows the stationary model is true shocks to fundamentals affect the price less. This result occurs because the term that multiplies the shock is less than the corresponding term for the non-stationary model.[11] On the other hand, the second term, the revision in beliefs is different than the belief revision term for the model learning model. This term creates additional volatility in price changes, which can at times be larger than the model learning model. Learning about the model parameters can also generate predictability of returns, because when μ^s is higher than the true value, prices will be high, and then they will tend to fall over time as μ reverts to its true value. Finally, there is no clear reason why the model would generate time varying volatility, but if the changes in μ^S are bigger at certain times than other times the model may generate time varying volatility. It therefore becomes an empirical exercise to compare constant gain learning to model learning which we analyze in Section 6.

[10] Note we assume that $E_{t-1}\mu_t^S = \mu_{t-1}^S$.
[11] The model is restricted to be stationary so that $\rho < 1$ always.

4. Full consumption based asset pricing model

4.1. Model description

The model consists of an infinitely lived representative agent who receives utility from consumption. The agent can choose to borrow or invest in a risk free asset with fixed (gross) return R.[12] He can also purchase claims to a risky asset (equity) which pays a stochastic dividend d_t. To price this asset we will assume that in equilibrium the agent will hold the fixed, one unit supply of the asset. Dividends follow a trend stationary process AR(p) process: $d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \ldots + \rho_p^s d_{t-p} + \varepsilon_t^s$.[13] However the agent does not know this. He puts some probability on the alternative that dividends follow the non-stationary (unit root process) $\Delta d_t = \alpha^{ns} + \rho_1^{ns}\Delta d_{t-1} + \ldots + \rho_p^{ns}\Delta d_{t-p} + \varepsilon_t^{ns}$. Since he does not know which process is true, he will put some weight on both processes, and these weights will evolve over time according to the likelihood of each model as described below.

4.2. Representative agent problem

The agent maximizes:

$$\max_{c_{t+s}} \hat{E}_t \sum_{s=0}^{\infty} \delta^s u(c_{t+s}, c_{t+s-1}) \tag{1}$$

subject to:

$$w_t = -Rb_t + \Theta_{t-1} d_t + \Theta_{t-1} p_t \tag{2}$$

$$b_{t+1} = c_t + \Theta_t p_t - w_t - y \tag{3}$$

$$\begin{aligned} d_t &= \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \ldots + \rho_p^s d_{t-p} + \varepsilon_t^s \text{ with } p = p_{s,t} \\ \Delta d_t &= \alpha^{ns} + \rho_1^{ns}\Delta d_{t-1} + \ldots + \rho_p^{ns}\Delta d_{t-p} + \varepsilon_t^{ns} \text{ with } p = 1 - p_{s,t} \end{aligned}$$

The agent maximizes lifetime expected utility with a discount factor δ. I use the hat notation on the expectations operator to denote that the agent does not know the true process for dividends and therefore this expectation is taken with respect to his beliefs at time t concerning the dividend process. Importantly I make a standard assumption from the learning literature, that of anticipated utility (Kreps (1998)), i.e. the agent makes decisions treating his beliefs as if they are constants and not random variables. However, beliefs can and do change in the future.

The agent's wealth evolves according to $w_t = -Rb_t + \Theta_{t-1} d_t + \Theta_{t-1} p_t$. Here b_t is beginning of period debt on which the agent pays interest $R - 1$. He receives dividend income $\Theta_{t-1} d_t$ where Θ_{t-1} are share purchases of the risky asset last period and d_t is the dividend payment from the risky asset. The value of the claim to the risky asset is $\Theta_{t-1} p_t$ where p_t is the price of the risky asset at time t. Debt evolves according to $b_{t+1} = c_t + \Theta_t p_t - w_t - y$ where consumption c_t, and share purchases beyond current wealth $\Theta_t p_t - w_t$, increase debt and income y decreases debt.[14] Finally, the agent does not know the true dividend process so his expectation is taken with respect to the following beliefs: with probability $p = p_{s,t}$ the dividend process is stationary: $d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \ldots + \rho_p^s d_{t-p} + \varepsilon_t^s$ and with probability $p = 1 - p_{s,t}$ the dividend process is non-stationary $\Delta d_t = \alpha^{ns} + \rho_1^{ns}\Delta d_{t-1} + \ldots + \rho_p^{ns}\Delta d_{t-p} + \varepsilon_t^{ns}$.

For the utility function I use the following, CARA exponential utility function with habit formation: $u(c_t, c_{t-1}) = \frac{-1}{\alpha} exp[-\alpha(c_{t+s} - \gamma c_{t+s-1})]$ (Alessie and Lusardi, 1997; Caballero, 1990). As noted by, Fuster et al. (2012) this choice of utility is useful for two reasons. First, it allows for a closed form solution to the consumption problem. This solution is helpful because it allows me to consider moderately complicated dynamics for the true dividend process and not be restricted to simply AR(1) processes. Additionally, I can solve for consumption endogenously increasing the set of important macroeconomic moments I can explain. Secondly, it allows one to have realistic smooth adjustment in consumption without creating time varying risk aversion as in Campbell and Cochrane (1999). It is unlikely that the results here are driven by the choice of utility function. In the context of pricing a housing asset, Tortorice (2015) shows that a similar model can generate the same effects as in this paper with CRRA utility. However, that version of the model does not endogenize consumption.

[12] The agent can borrow and lend unlimitedly at the risk free rate. I make this assumption so that consumption is not completely linked to the dividend outcome.

[13] I choose a stationary process as the true process because it aligns the consumer's problem with a common problem in macroeconomic time series analysis, i.e. identifying the presence of a unit root. However, Section 7 shows one obtains similar results with a non-stationary process as the true process as long as the true process exhibits more mean reversion than the alternative considered process.

[14] While the model could tractably include stochastic income, to focus on the effects on consumption of learning about dividends I set income deterministically equal to zero in the simulations.

4.3. Model solution

The Bellman equation for the model is:

$$V(z_t) = \frac{-1}{\alpha} exp[-\alpha(c_t - \gamma c_{t-1})] + \delta \hat{E}_t V(z_{t+1}) \tag{4}$$

where the vector of state variables $z_t = \begin{bmatrix} b_t & c_{t-1} & 1 & y & d_t & \vec{d_t} \end{bmatrix}'$. Here b_t is beginning of period debt, c_{t-1} is last period's consumption, 1 allows for a constant term in the consumption function, y is labor income, d_t is today's dividend income and $\vec{d_t}$ is the forecast vector for the non-stationary and stationary models given by: $\begin{bmatrix} 1 & d_t & \triangle d_t & \dots & \Delta d_{t-p+1} & 1 & t & d_t & \dots & d_{t-p+1} \end{bmatrix}$.

An equilibrium in this model is conditional on beliefs about the dividend process. Given beliefs about the parameters of each model represented by$\Upsilon_{i,\,t}$ and a belief about the probability the true process is the stationary process $p_{s,\,t}$ we can state the equilibrium as follows. The equilibrium is given by a consumption function $c(z_t, p_t)$, an asset demand function $\Theta(z_t, p_t)$ and an equity pricing equation $p_t(z_t)$ such that c and Θ maximize utility (1) subject to the wealth evolution constraints (2) and (3) and the supply of the asset equals demand for the asset θ, i.e. $\Theta(z_t, p_t(z_t)) = 1$, at the market clearing price $p_t(z_t)$. Beliefs are then updated after the equilibrium values are realized.[15]

Following the derivation in Fuster et al. (2012), Appendix B shows that the optimal solution for consumption is:

$$\frac{\gamma}{R} c_{t-1} - \frac{1}{R-1}\left[\frac{1}{\alpha}\ln(\delta R) + \frac{\alpha}{2}\sigma_c^2\right] + \left(1 - \frac{\gamma}{R}\right)\frac{R-1}{R}\left[-Rb_t + d_t + \hat{E}_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s}\right] \tag{5}$$

here σ_c^2 equals $p_{s,t}\sigma_{c,s}^2 + (1 - p_{s,t})\sigma_{c,ns}^2$ where $\sigma_{c,s}^2$ is the conditional variance of consumption growth under the assumption that the stationary model is true and $\sigma_{c,ns}^2$ is the conditional variance of consumption growth under the assumption that the non-stationary model is true. Because of habit formation, consumption depends on the previous period's consumption, where γ is the degree of habit formation. The second term in the consumption function is a downward shift in consumption that represents the consumers degree of patience captured by δR and the precautionary savings motive. The last term represents an annuity value of wealth that is most important in determining consumption.

Given this consumption function, following Fuster et al. (2012) Appendix B shows that the price of the equity asset is given by[16]:

$$p_t = \hat{E}_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s} - \frac{R\alpha\sigma_c^2}{(1 - \frac{\gamma}{R})(R-1)^2} \tag{6}$$

The equity price equals the present discounted value of dividends minus a penalty related to the riskiness of the asset which is proportional to the variance of consumption. Here $\hat{E}_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s} = p_{s,t}[E_t \sum_{s=1}^{\infty} \frac{d_{t+1}}{R^s}|S] + (1 - p_{s,t})[E_t \sum_{s=1}^{\infty} \frac{d_{t+1}}{R^s}|NS]$. This equation is the weighted average of the conditional expectations of the discounted future sum of dividends (conditional on which dividend process is true), where the weights are the agent's belief that each model is true. Again, as in Cogley and Sargent (2005) the agents evaluate these expectations as if their model parameters are constants even though they are treated as random variables for the purpose of their estimation.

4.4. Beliefs

I use the methods of Cogley and Sargent (2005) to calculate the parameters of each model and the probability weights on the stationary and non-stationary model. Their model uses Bayesian methods to recursively update the parameters on each model and then uses the likelihood of each model to calculate a probability weight on each model. I present a brief summary of their setup here. Details of their approach are given in Appendix C. For a given model (i.e. the stationary or non-stationary), let $i = \{s, ns\}$ index the model and let the vector of model parameters be given by $\Upsilon_{i,\,t}$. Updating of model parameters ais given by:

$$\begin{aligned}
P_{i,t} &= P_{i,t-1} + x_{i,t}x_{i,t}' \\
\Upsilon_{i,t} &= P_{i,t}^{-1}(P_{i,t-1}\Upsilon_{i,t-1} + x_{i,t}y_{i,t}) \\
s_{i,t} &= s_{i,t-1} + y_{i,t}^2 + \Upsilon_{i,t-1}' P_{i,t-1}\Upsilon_{i,t-1} - \Upsilon_{i,t}' P_{i,t}\Upsilon_{i,t} \\
\nu_{i,t} &= \nu_{i,t-1} + 1
\end{aligned}$$

Here $x_{i,\,t}$ is the vector of right hand side variables for the model at time t and $y_{i,\,t}$ is the left hand side variable for the model at time t, P_{t-1} is the precision matrix that captures the estimate of $\Upsilon_{i,t-1}$, σ_i^2 is the estimate of the variance of the model residuals, $s_{i,t-1}$ is an analogue to the sum of squared residuals, and $\nu_{i,t-1}$ is a measure of the degrees of freedom to calculate the residual variance such that the point estimate of $\sigma_{i,t-1}^2$ is given by $s_{i,t-1}/\nu_{i,t-1}$.

[15] The first order condition for the maximization problem are $\frac{\partial V}{\partial c_t} = 0$ and $\frac{\partial V_t}{\partial \Theta_t} = 0$ and given in the appendix.

[16] The market clearing price is the one at which where demand for the asset θ = supply (1).

Additionally, weights on each model follow the recursion:

$$\frac{w_{s,t+1}}{w_{ns,t+1}} = \frac{m_{s,t+1}/m_{s,t}}{m_{ns,t+1}/m_{ns,t}} \frac{w_{s,t}}{w_{ns,t}} \tag{7}$$

where $m_{i,\ t}$ represents the marginalized likelihood of each model. Finally, to calculate the model probabilities, the consumer normalizes the weights to one, and therefore the weight on the stationary model is given by:

$$p_{s,t} = \frac{1}{1 + w_{ns,t}/w_{s,t}} \tag{8}$$

Since dividends are an exogenous process, the model will eventually put all the weight on the true process. To allow for perpetual learning, I adapt the concept of constant gain learning from the least squares learning literature to the current setup. I introduce a gain parameter (g) that over-weights current observations. The gain probability can be interpreted as the probability of a structural break in the economy, such that the history of the dividend process no longer has any bearing on the current process generating dividends, hence the previous weight ratio is set to one. So then with probability $1 - g$ there is no structural break and the probability is given by equation (8) with the weights given by equation (7) and with probability g there is a structural break and the probability is given by equation (8) with the weights given by equation (7) but with $\frac{w_{s,t}}{w_{ns,t}}$ set to 1.

$$p_{s,t} = (1-g)\frac{1}{1 + w_{ns,t}/w_{s,t}} + g\frac{1}{1 + \frac{m_{s,t+1}/m_{s,t}}{m_{ns,t+1}/m_{ns,t}}}$$

In introducing the gain, I am motivated by the literature on constant gain learning (Evans and Honkapohja, 2001), however the gain here has a different function as we are learning about models as opposed to parameters. In this way, the gain is closer to the forgetting parameter used in the literature on Bayesian dynamic model averaging (Koop and Korobilis, 2012; Raftery et al., 2010). This literature takes the likelihood to be an exponentially weighted average of past prediction errors, weighing recent observations more heavily. That is to say $L(\Upsilon_i, \sigma_i^2, D^t) = \prod_{s=1}^{t} p(y_s|x_s, \Upsilon_i, \sigma_i^2)^{(1-g)^s}$. In order to preserve the analytic and recursive structure of my model, I introduce the weighting as a probability of a structural break instead of down-weighting the likelihood. However the overall effect is the same – to overemphasize the more recent observations in calculating the likelihood.

In addition to creating perpetual learning, there are two economic motivations for considering the gain parameter. The first is that the agent may believe that there is a possibility of a structural break in the economy. In that case, the agent would wish to guard against this possibility by over-weighting more recent observations. One may wonder why this concern for robustness to a structural break in the economy does not extend to the parameters of each model as well. In fact, I have explored allowing for constant gain learning in the estimation of the model parameters along with the model weights. I have found that adding this dimension for learning does not quantitatively affect the model learning results in this paper and therefore I omit constant gain learning of parameters in the interest of clarity.

Additionally, much psychological evidence indicates that individuals tend to overweight more recent observations. Tversky and Kahneman (1973) document the availability bias which causes agents to overweight more readily available information when forming forecasts of future events. For example, after a plane crash is in the news, individuals think a plane crash is more likely. Rabin (2002) calls this bias the "law of small numbers" where individuals use a recent string of random numbers to incorrectly infer the nature of an underlying statistical process. Rabin and Vayanos (2010) use this approach to explain various puzzles in financial markets. In the current model, the gain functions to overweight recent observations consistent with the psychological evidence that individuals tend to overweight the most readily accessible information.

There is also substantial evidence that agents overreact to news about fundamentals. For example, De Bondt and Thaler (1985, 1987, 1990) all find that professional security analysts overreact to fundamentals and have forecasts that are two optimistic after positive news. Similar evidence is obtained by Ertan et al. (2016). Additionally, Beshears et al. (2013) in an experimental setting find that subjects fail to anticipate mean reversion in a slowly mean reverting process. Adding a gain to a model generates overreaction and insufficient expectations of mean reversion.

The gain also serves to capture the potential of agents to succumb to new era stories during periods of large run ups in asset prices, for example, believing in the 1990s that the internet was a revolutionary new technology that has fundamentally changed the determination of asset prices. Shiller (2005) and Reinhart and Rogoff (2009) argue that this dynamic is an important driver of asset price bubbles and subsequent financial panics.

To conclude the discussion of the belief section, I discuss two issues. The first is why it is key to have learning versus just exogenous extrapolation. The second is the reasons for setting up the learning problem and the information structure in the way I have. It is worth noting that the non-stationary dividend model is an extrapolation model. A variety of authors have explored the presence of extrapolative agents in economic models, mostly in models of finance: for references see, Barberis et al. (2015). There are at least three reasons the learning model here is preferable to just assuming extrapolation exogenously. First learning provides a basis for explaining the behavior of extrapolaters, i.e. they extrapolate because recent data supports extrapolation. Second, in equity markets one expects extrapolaters to influence asset prices more when extrapolation better fits the data, both because extrapolaters will have become wealthier and more individuals will start to extrapolate. The current model captures this effect. Finally, endogenous extrapolation is necessary here to explain time varying volatility.

To additionally motivate this learning model, I contrast this approach with other potential approaches. The first is a Markov switching model where the dividend process switches between the stationary and non-stationary formulations. I do not take this approach because it does not generate significant variation in long-run forecasts. If transition probabilities are high then initial beliefs do not matter much for where you end up in the long run, if transition probabilities are low then you do not observe many transitions in simulation samples. The second approach would have a dividend process with permanent and transitory shocks. The agent's learning problem could be solved with the Kalman filter. However in this model agents always react the same way to shocks (as a linear combination of a permanent and transitory shocks with weights given by the relative variances of the two shocks). This model does not then generate endogenous time varying volatility. A third approach would combine the two previous approaches: a Markov switching model where in state one the economy's shocks are permanent and in state two the economy's shocks are only temporary. However, this model is not tractable. The whole history of shocks and time varying probabilities of all past states would be necessary to make forecasts. My approach in this paper tries to capture the dynamics of this last approach in a straightforward, tractable way.

5. Calibration and simulation

Time is quarterly and I set the risk free rate $(R-1)$ equal to 0.0025 implying a 1% annual risk free rate. I set δ, the rate of time preference, to $\frac{1}{R}$. I set $\gamma = 0.7$ to match the standard deviation of consumption and set the risk aversion parameter α to 0. I use no risk aversion so that asset price movements are driven by changes in expectations about future dividends only and not time varying risk assessments.[17] Robustness to the choice of the risk free rates is presented in Section 7.

I calibrate the gain parameter to minimize the distance between the learning model implied forecasts and forecasts of corporate profit growth from the Survey of Professional Forecasters. I choose the corporate profits forecasts because it is the variable closest to dividends in the professional forecast data.[18] The calibrated value of the gain is 0.0389. My gain is analogous to $1-f$ where f is the forgetting parameter used in Koop and Korobilis (2012). They consider a range of f from 0.99 to 0.95. Therefore, I have a gain parameter of similar magnitude as to what is chosen in the Bayesian dynamic model averaging literature. I examine robustness to alternative choices for the gain from $g = 0.01$ and $g = 0.05$ in Section 7.

I also need to assign initial beliefs for the candidate models. However, these initial beliefs do not impact the results because I simulate the model for 1000 periods and use only the last 340 = 4*(2013-1928) observations to correspond to the length of my data set. I begin with an initial prior on the stationary model $p_{s,t} = 0.5$. To set beliefs for the two possible dividend processes I estimate a stationary process ($d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \ldots + \rho_p^s d_{t-p} + \varepsilon_t^s$) and a non-stationary process ($\Delta d_t = \alpha^{ns} + \rho_1^{ns}\Delta d_{t-1} + \ldots + \rho_p^{ns}\Delta d_{t-p} + \varepsilon_t^{ns}$) by ordinary least squares using data on the net operating surplus of private enterprises.[19] The data come from the National Income and Product Accounts (Table 1.10 line 12), Bureau of Economic Analysis and are deflated with the GDP deflator. Data are quarterly, begin in 1947 and end in 2012. I let the number of AR lags $p = 4$. I then set the initial beliefs about the parameters, $Y_{s,\,0}$ and $Y_{ns,\,0}$ to the estimated parameters of the dividend process. I set the precision matrices to: $P_0 = 0.01 * I_p$ which allows for a fairly defuse prior. From Section 4.4, we see that this prior gives a standard error for the initial coefficient equal to $\sqrt{100} = 10$ times the standard deviation of the regression residual which in this case leads to a standard error equal to 20% of the dividend. The initial (sum of squared residuals)s_0 is set equal to the estimated residual variance of each model and the initial degrees of freedom are set equal to one.[20]

I then simulate the model assuming the true dividend process is the stationary dividend process and ε_t^s is distributed $N(0, \sigma_s^2)$ where σ_s^2 is estimated using the sample variance of the regression residuals. I simulate the model 500 times using a simulation length of 1000 quarters. I keep only the last 340 = 4*(2013-1928) and report median statistics for the model. The model is calibrated at a quarterly frequency, so I construct an annual data series using year end prices, the cumulative return over the four quarters in the year, and the quarter four to quarter four change in log consumption.[21]

I choose to make the true process stationary given the evidence of overreaction in the data, i.e. the negative correlation among these variables at medium run horizons. In this way, belief in the non-stationary model is a mistake and the agent's correction of this mistake allows the model to be able to match the data moments. The model therefore gives an intuitive explanation for these asset pricing puzzles, one in which agents overact to news about dividends. However, this overreaction is endogenous as they are more likely to be mistaken when the data tends to look non-stationary. I also allow the true process to be stationary because it is consistent with the view that people overemphasize the current situation and ignore factors which may lead to long run mean reversion. For example, emphasizing current profitability and ignoring increased

[17] I implement this calibration by taking the limit as $\alpha \to 0$ in equations (5) and (6). This sets the second term in these equations to zero. Note that since $\delta = \frac{1}{R}$, $\ln(\delta R) = 0$ the second term in equation (5) approaches zero as $\alpha \to 0$.

[18] See Appendix D for exact details.

[19] I use these data as opposed to the Shiller dividend series because they are seasonally adjusted and do not exclude profits shareholders receive through share buybacks. However, I obtain similar results using the Shiller dividend series.

[20] I estimate the stationary process with a time trend and log dividends and the non-stationary process with an intercept. To get the model process (which is in levels) I use the estimated coefficients and set the time trend to zero for the stationary model and the intercept to zero for the non-stationary model. I estimate the initial sample variance of residuals by taking the sample variance and multiplying by the steady state value of the dividend from the model.

[21] Following Fuster et al. (2012), I use cumulative gains as opposed to excess returns to calculate the correlation statistics. These are obtained by taking the excess return R_t and multiplying by p_{t-1}, though I find the difference to be unimportant. Note also the true dividend process here does not result in simulations with negative dividends because the true process is stationary.

Table 2
Model results – return predictability.

Moment	Data	Rational expectations	Parameter learning	Model learning	Constant gain learning
$\text{corr}(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2	−0.03	−0.04	−0.22	−0.09
$\text{corr}(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41	0.11	0.09	−0.23	−0.003
$\text{corr}(\Delta \text{lnc}_t, r_{t+2} + \ldots + r_{t+5})$	−0.34	−0.04	−0.06	−0.26	−0.14
$\text{corr}(P/E_{10,t}, \Delta \text{lnc}_{t+3} + \ldots + \Delta \text{lnc}_{t+6})$	−0.16	0.11	0.1	−0.21	−0.005
$\text{corr}(\Delta \text{lnc}_t, \Delta \text{lnc}_{t+3} + \ldots + \Delta \text{lnc}_{t+6})$	−0.23	−0.03	−0.05	−0.23	−0.14
$\text{corr}(P/E_{10,t}, \Delta \text{lnd}_{t+2} + \ldots + \Delta \text{lnd}_{t+5})$	−0.25	0.15	0.15	−0.33	−0.22
$\sigma(r_t)$	20.50%	0.13%	0.15%	3.1%	1.50%
$\sigma(\Delta \text{lnc}_t)$	2%	0.1%	0.1%	2%	1.0%

Note: This table reports the correlation of returns, dividend growth, the PE ratio and consumption growth from the learning model described in the paper. The rational expectations version comes from setting the probability on the stationary model equal to one and the stationary parameters to the true model for all time periods. The parameter learning column allows learning about the stationary parameters but keeps the probability on the stationary model equal to one. The constant gain learning column reports results from the constant gain learning algorithm with a known stationary dividend process. The model learning column has learning both about the parameters and the correct model. I report median statistics obtained by simulating the model for 500 trials.

competition which may lower profit margins in the future or observing an increase in house prices now and ignoring how supply may increase in the future lowering prices. This type of overreaction is consistent with the "this time is different analysis" of Reinhart and Rogoff (2009) and the tendency of agents to justify temporary movement with new-era stories as described in Shiller (2005).

Finally, it's worth noting that the empirical evidence on the stationarity of the dividend process is mixed (conditional on a deterministic time trend). While tests for the existence of a unit root in the dividend series fail to reject the null, these tests have notoriously low power against the alternative of a near-unit root model. Many authors have noted this possibility, for example Cochrane (1988, 1991), Stock (1991) in examining several macroeconomic time-series. Additionally, Fama and French (1988) argue that longer run mean reversion in returns make it difficult to convincingly establish the existence of a unit root in the stock price.

Still, given the common assumption of non-stationary dividends in the literature I provide more support for the assumption of the model in the robustness section, Section 7. I show that running the actual dividend data through the featured learning process does not give convincing evidence that dividends are non-stationary. I also show the main results of the paper can be replicated with a model where the true process is non-stationary but has more mean reversion than an alternative candidate model.

6. Results

6.1. Return predictability

Table 2 gives the simulation statistics concerning return predictability. First for the data, and then for three simulations: a rational expectations benchmark where the agent knows the true dividend process is stationary and knows the parameters of the process, a parameter learning model where the agent knows the true model is stationary but not the parameters of the process, and a model learning simulation where the agent learns about both the parameters of the dividend process and whether the true dividend model is stationary or non-stationary. First we see that the rational expectations benchmark performs poorly. The correlation between returns and future returns is close to zero, as is the correlation between consumption growth and future returns and consumption growth and future consumption growth.[22] All the other correlation are of the wrong sign. For example, the PE ratio is positively correlated with future returns and consumption growth.[23] Additionally, the PE ratio positively forecasts future dividend growth. The model generates 1/200th the volatility of returns in the data and 1/20th the observed volatility of consumption. It is worth noting that one could improve these volatility statistics by using a non-stationary process for the true dividend. However, I choose a stationary benchmark of a specific reason. One challenge in explaining equity returns is how to amplify the effect of changes in fundamentals. What I am able to show is that this model can take very smooth fundamentals and substantially amplify the volatility of returns.

In contrast to the rational expectations benchmark, the model with learning about the non-stationary model does substantially better. It predicts a correlation between current returns and future returns of −0.22 versus −0.2 in the data. It predicts a correlation of the PE ratio and future returns of −0.23 versus −0.41 in the data. Consumption growth is negatively correlated with future returns. The model predicts a consumption growth future return correlation of −0.26 versus −0.34 in the data. Similarly, the PE ratio negatively predicts consumption growth with a correlation of −0.21 versus −0.16 in the data. The correlation of consumption growth in the model with future consumption growth matches the data exactly at −0.23. Finally, the model predicts a correlation between the PE ratio and future dividend growth of −0.33 vs −0.25 in the data. The model also significantly amplifies volatility. It explains 15% of the volatility of returns observed in the data, 24 times more than the rational expectations benchmark and it matches the volatility of consumption exactly.

[22] Small sample bias prevents these correlations from being exactly zero.
[23] Again these correlation go to zero as the simulated time series length is increased.

Table 3
Kurtosis, autocorrelation of squared returns, GARCH.

Kurtosis	Data	Confidence bounds	RE	Learning	CG learning	Constant Prob
r_t	4.1	[2.58 3.46]	2.95	6.33	12.05	2.96
$P/E_{10,\,t}$	4.6	[1.67 3.50]	2.78	3.03	2.99	2.82
$\%\lvert r_t\rvert > 1.96^{*}s(r_t)$	6.2%	[3% 7%]	5.0%	5.6%	4.4%	5.0%
Autocorrelation of Squared Returns	**Data**	**Confidence Bounds**	**RE**	**Learning**	**CG Learning**	**Constant Prob**
lag 1	0.079	[−0.028 0.186]	−0.008	0.2	0.25	−0.006
lag 2	0.01	[−0.097 0.117]	−0.003	0.17	0.16	−0.004
lag 3	0.47	[0.36 0.58]	−0.008	0.14	0.11	0.001
lag 4	0.14	[0.033 0.25]	−0.002	0.11	0.1	−0.006
GARCH	**Data**	**Confidence Bounds**	**RE**	**Learning**	**CG Learning**	**Constant Prob**
Garch	0.61	[0.46 0.76]	0	0.57	0.62	0
Arch	0.29	[0.17 0.41]	0	0.25	0.23	0
p-value Engle test	0.048					

Note: This table reports kurtosis, autocorrelation of squared returns, and GARCH estimate for the excess returns data and the simulated model data. CG Learning refers to the constant gain learning algorithm. Constant Prob refers to a model with constant probability on the non-stationary model. For the autocorrelation and GARCH estimates 90% confidence bounds are reported as +/− 1.65 standard errors.

The middle column of Table 2 allows us to demonstrate the importance of learning about the likelihood of the non-stationary model. We see that allowing the agent to learn about the parameters of the stationary process while assuming he knows the true process is stationary makes little difference. The predicted moments are similar to the rational expectations benchmark. I obtain this result because I run the simulation for many periods before I select the data used for calculating the statistics, and by then the parameters have mostly converged to their true values. Therefore, model learning then becomes key to explaining the ability of the model to match the data.

I also explore the role that constant gain learning can have in matching these moments. For information on the constant gain learning algorithm see Appendix E. For this algorithm the agent knows the true model is stationary, but the constant gain allows the coefficients on the model to vary more over time and never converge to their true value.[24] The gain value is calibrated at 0.0219 again by minimizing the distance between model forecasts and professional forecasts. Constant gain learning creates some of the negative correlations we see in the data. The correlation between returns and future returns equals −0.09 and the correlation with consumption growth and future returns is −0.14. Consumption growth is also negatively correlated with future consumption growth and the PE ratio is negatively correlated with future dividend growth just as in the data. Constant gain learning generates about half the return volatility generated by the model learning. However, constant gain learning, for this calibration, fails to reproduce the negative correlations between the PE ratio and future returns and the PE ratio and future consumption growth.

To summarize, the main model in this paper generates a negative correlation between returns when agents misinterpret random movements in dividends generated by the stationary process. They believe the true dividend process may be non-stationary and this misinterpretation leads to a large change in the equity price. When dividends eventually begin to mean revert, prices do as well. This reversion explains the negative correlation. I showed through a simulation that the model with learning about the true dividend process is able to explain the negative correlations we observe in equity markets and a substantial amount of volatility in consumption and returns.

6.2. Time varying volatility

Table 3 examines the ability of the model to generate time varying volatility in returns and the PE ratio compared to the benchmark rational expectations model, a constant gain learning model and a model with constant non-zero probability on the non-stationary model. As the constant probability, I use the median probability across time and trials for the learning model.[25] Here I focus on quarterly data as it is difficult to detect time varying volatility at the annual frequency.

Recall, in the data quarterly return kurtosis equals 4.1. Both the rational expectations benchmark and the constant probability model imply quarterly returns should look normal with a kurtosis of 3. However, the learning model is able to amplify kurtosis predicting a kurtosis of 6.33. The constant gain learning model tends to generate normally mild but occasionally extreme returns, this pattern results in a large kurtosis value equal to 12.05. The PE-ratio also exhibits kurtosis. In quarterly data, kurtosis of the PE-ratio is 4.6. The RE and constant probability models predict kurtosis of 2.8. The learning model is able to amplify kurtosis, though only slightly, predicting a kurtosis of 3.03 similar to the constant gain learning value of 2.99. Finally I also look at the percent of absolute returns which are greater than 1.96 times the standard deviation of returns.

[24] When updated beliefs result in a non-stationary parameterization of the dividend process, beliefs are set to the value from the previous period to ensure stationarity.

[25] This probability is 0.67.

If returns were normal this statistic should be 5%. However, in the data it is 6.2%. Both the rational expectations and the constant probability model predict 5% of returns should be greater than 1.96 standard deviations. The learning model better matches the data predicting that 5.6% of returns should be above 1.96 standard deviations. For the constant gain learning model this statistic is 4.4%

In the second panel of Table 3 we examine the autocorrelation of squared returns as a way to measure time varying volatility. We found that in the quarterly data, up to four lags all these autocorrelations were positive. The correlation at one lag was 0.08, 0.01 at two lags, 0.47 at three lags and fell to 0.14 at four lags. The rational expectations benchmark model and the constant probability model do not predict any autocorrelation in squared returns. All estimated autocorrelation coefficients are near zero. However the learning model does predict positive autocorrelation of squared returns. It predicts an autocorrelation of 0.22 at one lag down to 0.11 at four lags. The constant gain learning model also generates autocorrelation of squared returns ranging from 0.25 to 0.11.

Finally, I examine the ability of the model to explain the GARCH effects found in the data. For the GARCH(1,1):

$$\sigma_t^2 = \kappa + \gamma_1 \sigma_{t-1}^2 + a_1 \varepsilon_{t-1}^2$$

I estimated $\gamma_1 = 0.61$ and $a_1 = 0.29$. Both estimates were highly statistically significant. To examine the ability of the models to generate these facts, I first simulate return data from the models. Then I run an Engle test with the null hypothesis of no conditional heteroscedasticity. If the test rejects I estimate the GARCH parameters, otherwise I assign zeros for the parameters.[26] I then report the median statistics across the simulations. We first see that the rational expectations model and the constant probability model predict no GARCH effects. However the learning model predicts $\gamma_1 = 0.57$ and $a_1 = 0.25$ versus $\gamma_1 = 0.61$ and $a_1 = 0.29$ in the data. Statistics were similar for the constant gain learning model with $\gamma_1 = 0.62$ and $a_1 = 0.23$

The common approach for explaining time varying volatility is to assume time varying volatility in the dividend process. While this certainly may be the case, I believe the current setup should be preferred for a variety of reasons. First, exogenous time varying volatility is an additional assumption that must be added into the baseline model to generate time varying volatility while it is orthogonal to the other data moments of interest. Here the learning mechanism improves the baseline model on this dimension and many other dimensions. Secondly, a key empirical challenge is to explain return volatility with a relatively smooth dividend process. To put it another way the goal is to generate amplification from small dividend shocks. In line with this motivation, it seems important to explore how small, constant volatility shocks can result in return series with changing volatility. Thirdly, very few shocks are truly exogenous. Even the dividend process is an outcome of the firms corporate finance decisions and the economic, firm and industry level conditions that determine profitability. Therefore to the extent that we wish to understand time varying volatility it is important to endogenize it. Relatedly, it is difficult to determine if changes in earnings volatility cause changes in return volatility of if another unobserved variable impacts both earning and returns. In fact, efforts to isolate truly exogenous movements in fundamentals that correspond to observed large changes in returns have been quite unsuccessful (see for example Cutler et al., 1989; Roll, 1984). These results suggest the need for models which generate periods of high return volatility absent changes in the underlying volatility of fundamentals.

Of course, it would be straightforward to add in time varying volatility to the true dividend process. This would only make it easier for the model to match the level of time varying volatility in the data.

There is clear evidence of time varying volatility in the quarterly returns data. The learning model endogenously generates this as the agent's belief that the world is non-stationary is changing over time. Periods where the agent increases his belief that the non-stationary model is true are periods when the volatility of returns increases. The benchmark models without changes in beliefs cannot endogenously generate this time varying volatility.

7. Robustness and further analysis

This section examines the robustness of the results on return predictability and time varying volatility to different parameter choices. It provides support for the assumption that dividends are stationary and shows that this is not a key assumption to generate the results. I also give a suggestive welfare analysis that supports the idea that these belief mistakes do not result in large welfare losses. It then quantifies the importance of changing beliefs for the variance of price changes and the model's predictions of return predictability. I do this by decomposing changes in the equity price and comparing the return predictability results of the main model to the model where the agent puts a constant non-zero probability on the non-stationary model.

7.1. Robustness

The model is fairly tightly parametrized. However, I did need to set a risk free rate, gain parameter and AR lag length. Table 4 gives the model results varying one of these parameters, while keeping all other parameters constant.

The choice of lag length does not matter much for the results. Using a lag length equal to two I find correlations that are a little bit smaller, but almost all are within 0.05 of the main results. The model still generates the negative correlations

[26] This procedure is necessary because under the null of no GARCH the likelihood function is flat and I am unable to identify the GARCH parameters.

Table 4
Robustness.

Moment	Data	AR	2	4	6	8
$\text{corr}(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2		−0.18	−0.22	−0.17	−0.21
$\text{corr}(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41		−0.19	−0.23	−0.16	−0.22
$\text{corr}(\Delta \ln c_t, r_{t+2} + \ldots + r_{t+5})$	−0.34		−0.23	−0.26	−0.23	−0.27
$\text{corr}(P/E_{10,t}, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.16		−0.18	−0.21	−0.16	−0.2
$\text{corr}(\Delta \ln c_t, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.23		−0.21	−0.23	−0.22	−0.26
$\text{corr}(P/E_{10,t}, \Delta \ln d_{t+2} + \ldots + \Delta \ln d_{t+5})$	−0.25		−0.28	−0.33	−0.25	−0.3
$\sigma(r_t)$	20.50%		3.1%	3.1%	2.5%	2.6%
$\sigma(\Delta \ln c_t)$	2%		2.0%	2%	1.7%	1.70%
$\Delta(r_t^2, r_{t-1}^2)$	0.079		0.17	0.2	0.17	0.18
Moment	**Data**	**Gain**	**0.01**	**0.02**	**0.0389**	**0.05**
$\text{corr}(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2		−0.2	−0.21	−0.22	−0.22
$\text{corr}(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41		−0.12	−0.18	−0.23	−0.25
$\text{corr}(\Delta \ln c_t, r_{t+2} + \ldots + r_{t+5})$	−0.34		−0.24	−0.25	−0.26	−0.26
$\text{corr}(P/E_{10,t}, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.16		−0.1	−0.15	−0.21	−0.23
$\text{corr}(\Delta \ln c_t, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.23		−0.22	−0.22	−0.23	−0.24
$\text{corr}(P/E_{10,t}, \Delta \ln d_{t+2} + \ldots + \Delta \ln d_{t+5})$	−0.25		−0.18	−0.26	−0.33	−0.35
$\sigma(r_t)$	20.50%		2.0%	2.6%	3.1%	3.30%
$\sigma(\Delta \ln c_t)$	2%		1.30%	2%	2%	2.10%
$\Delta(r_t^2, r_{t-1}^2)$	0.079		0.28	0.23	0.2	0.26
Moment	**Data**	**R**	**1.0025**	**1.005**	**1.01**	**1.02**
$\text{corr}(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2		−0.22	−0.21	−0.22	−0.21
$\text{corr}(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41		−0.23	−0.22	−0.23	−0.24
$\text{corr}(\Delta \ln c_t, r_{t+2} + \ldots + r_{t+5})$	−0.34		−0.26	−0.25	−0.26	−0.25
$\text{corr}(P/E_{10,t}, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.16		−0.21	−0.2	−0.21	−0.22
$\text{corr}(\Delta \ln c_t, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.23		−0.23	−0.23	−0.24	−0.23
$\text{corr}(P/E_{10,t}, \Delta \ln d_{t+2} + \ldots + \Delta \ln d_{t+5})$	−0.25		−0.33	−0.33	−0.34	−0.38
$\sigma(r_t)$	20.50%		3.1%	3.1%	3.3%	3.6%
$\sigma(\Delta \ln c_t)$	2%		2%	2.0%	2.2%	2.4%
$\Delta(r_t^2, r_{t-1}^2)$	0.079		0.2	0.19	0.18	0.17
Moment	**Data**		**Two NS**	**S vs NS**		
$\text{corr}(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2		−0.19	−0.22		
$\text{corr}(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41		−0.28	−0.23		
$\text{corr}(\Delta \ln c_t, r_{t+2} + \ldots + r_{t+5})$	−0.34		−0.21	−0.26		
$\text{corr}(P/E_{10,t}, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.16		−0.24	−0.21		
$\text{corr}(\Delta \ln c_t, \Delta \ln c_{t+3} + \ldots + \Delta \ln c_{t+6})$	−0.23		−0.18	−0.23		
$\text{corr}(P/E_{10,t}, \Delta \ln d_{t+2} + \ldots + \Delta \ln d_{t+5})$	−0.25		−0.53	−0.33		
$\sigma(r_t)$	20.50%		3.70%	3.1%		
$\sigma(\Delta \ln c_t)$	2%		2.40%	2%		
$\Delta(r_t^2, r_{t-1}^2)$	0.079		0.08	0.2		

Note: This table shows robustness of the results from varying key parameters and considering a model where the true divdend process is non-stationary.

in the data and obtains similar results for the standard deviation of consumption and returns. Increasing the lag length has a similar effect. There is very little difference between the statistics generated with a lag length of six or eight versus four. These lag lengths all generate very similar negative correlations and standard deviations of consumption growth and returns. For each lag length, we still see substantial autocorrelation of squared returns.

For the gain parameter I find a very similar result. Lowering the gain to 0.02 or 0.01 has no real effect on the correlation of returns with future returns and consumption growth with future returns and future consumption growth. The correlation of the PE ratio with future returns and consumption falls slightly. We also see a small fall in the standard deviation of returns and consumption growth. In contrast, increasing the gain to 0.05 increases the PE correlations slightly while having little effect on the other correlations. It also increases the volatility of returns and consumption. But in either case the results from changing the gain are very similar to the results from the baseline calibration. The predicted autocorrelation of squared returns is only slightly affected.

Finally I consider increasing the quarterly, gross risk free rate from 1.0025 to 1.02. This change increases the annual rate from 1% to 8% per year. The choice of the risk free rate does not matter for the correlation coefficients. I get almost identical results for the different choices of R. I do find that volatility increases when R is increased. When R = 1.02, the standard deviation is 3.6% for returns versus 3.1% in the baseline case.

The last panel of Table 4 presents the results from a model where the true dividend process is non-stationary.

$$\Delta d_t = \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \ldots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns}$$

I allow the true process to be an AR(40) in differences as in Fuster et al. (2012) and the agent considers two possible processes for the differenced dividend process: an AR(4) and an AR(40). While non-stationary, the AR(40) model exhibits

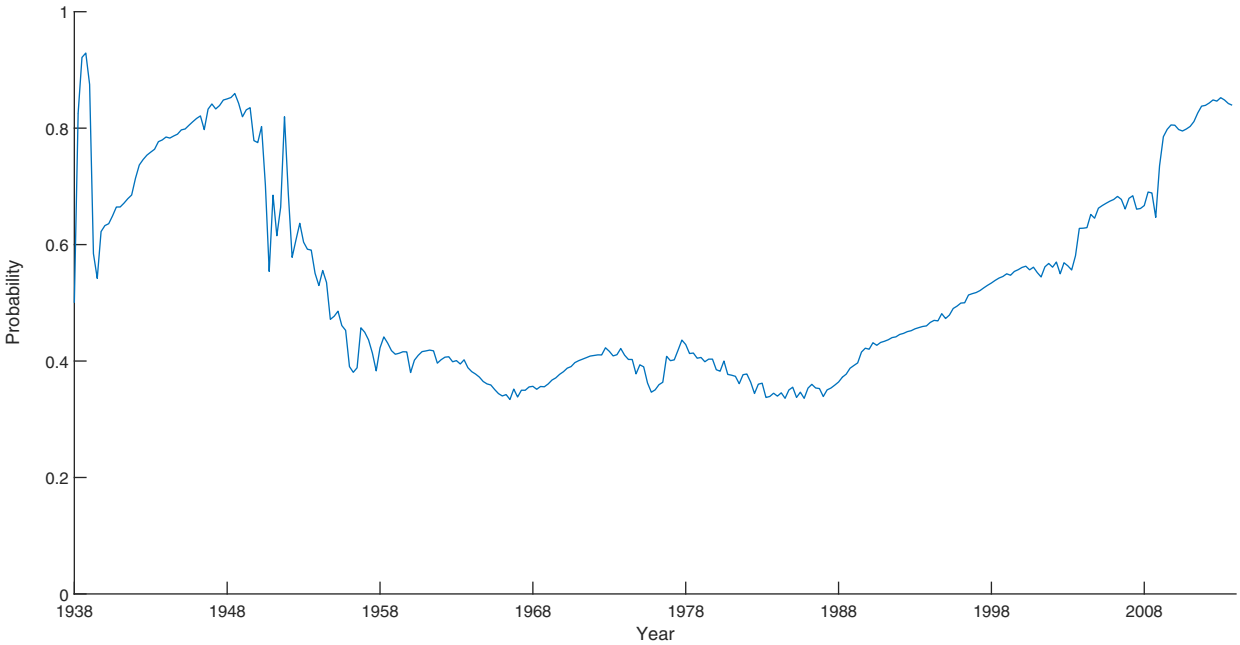

Fig. 5. Probability on stationary model (Shiller data).

substantially more mean reversion than the AR(4) model. Results from this version of the learning model show that the assumption of stationarity is not key to generate the main results of the paper. Indeed all of the main correlation coefficients are close to the benchmark case. The model with two non-stationary processes does generate higher volatility overall ($\sigma(r) = 3.7\%$ vs. 3.1% in the benchmark case). But it generates less time varying volatility, the autocorrelation of squared returns equals 0.08 versus 0.2 in the original model.

To provide further evidence on the possibility of a trend stationary dividend process I run the Bayesian learning model in the paper using actual Shiller dividend data on dividends accruing to the S&P 500. I allow for a gain of zero here but results were very similar for higher gains. The results are given in Fig. 5. The probability weight on the stationarity model ranges from 0.4 to 0.8 and the data indicate that the trend stationary dividend is as likely as the non-stationary dividend. To get some intuition for this result examine Fig. 6 which plots the log of the dividend data along with a time trend. What one sees is that dividends do look like they tend to revert to this time trend. Moreover, we put the least weight on the stationary model when dividends are far from this trend, e.g. 1968 and 1988. And we put increasing more weight on the stationary model as dividends return to this trend.

As a final check on the reasonableness of the results, I consider the welfare cost of using the learning rule versus using the rational expectations rule. I calculate consumption assuming the representative agent knows the true process (c^s) versus if he needs to learn about the true process (c^l).The welfare cost τ is given by:

$$E_t \frac{-1}{\alpha} \sum_{s=0}^{\infty} \delta^s exp[-\alpha(c^l_{t+s} - \gamma c^l_{t+s-1})] - exp[-\alpha((1-\tau)c^s_{t+s} - \gamma(1-\tau)c^s_{t+s-1})] = 0.$$

Hence the welfare is the fraction of per-period consumption that the rational agent would need to loose in order to be indifferent between being in the learning economy and the rational economy. I find that these losses are quite small. I consider, to amplify the losses, a coefficient of CRRA equal to 10.[27] Even then welfare losses are only 0.45% of consumption. This analysis is only suggestive however because in a world with heterogeneous agents there may be additional welfare gains to rational agents of exploiting learning agents' mistaken beliefs.

In summary, the results are very robust to different choices of the AR lag, the gain parameter and the risk free rate. The results are not in anyway dependent on the choices for these parameters. Welfare losses of shifting to learning seem to be small and it is not necessary to assume a stationary dividend process in order to generate the model results. One caveat remains however. The gain parameter must be high enough to allow for non-trivial perpetual learning. With a lower value of the gain, I would need to shorten the sample size to generate similar results, making the results potentially more dependent on the initial choice of priors.

[27] Fuster et al. (2012) show that given a CRRA coefficient σ then the corresponding CARA coefficient is $\alpha = \frac{\sigma}{d(1-\frac{\gamma}{R})}$ where d is the mean dividend.

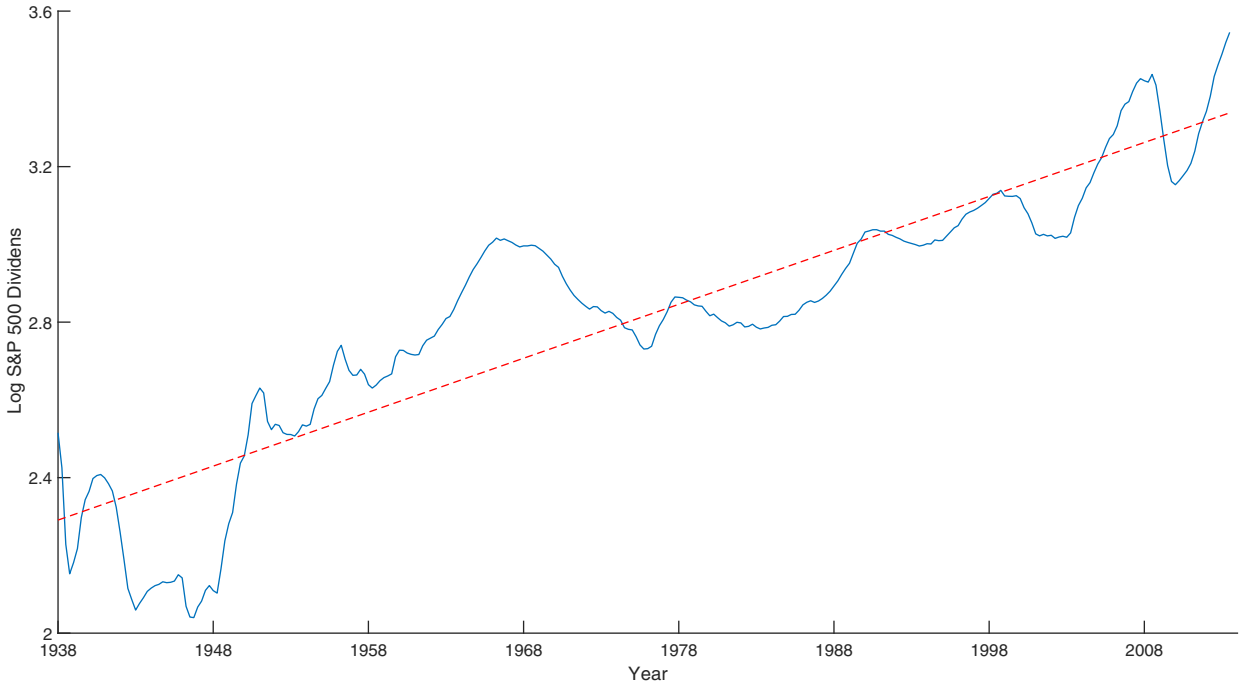

Fig. 6. Trend Line and Log Shiller Dividend Data.

Table 5
Impact of model learning.

5a Price change decomposition	$cov(\Delta p^L_t, dp_1)/var(\Delta p^L_t)$ 70.0%	$cov(\Delta p^L_t, dp_2)/var(\Delta p^L_t)$ 30.0%	
5b Moments at median probability			
moment	Data	Model	Constant probability
$corr(r_t, r_{t+2} + \ldots + r_{t+5})$	−0.2	−0.22	−0.26
$corr(P/E_{10,t}, r_{t+2} + \ldots + r_{t+5})$	−0.41	−0.23	−0.18
$corr(\Delta lnc_t, r_{t+2} + \ldots + r_{t+5})$	−0.34	−0.26	−0.3
$corr(P/E_{10,t}, \Delta lnc_{t+3} + \ldots + \Delta lnc_{t+6})$	−0.16	−0.21	−0.15
$corr(\Delta lnc_t, \Delta lnc_{t+3} + \ldots + \Delta lnc_{t+6})$	−0.23	−0.23	−0.25
$corr(P/E_{10,t}, \Delta lnd_{t+2} + \ldots + \Delta lnd_{t+5})$	−0.25	−0.33	−0.23
$\sigma(r_t)$	20.50%	3.1%	2.1%
$\sigma(\Delta lnc_t)$	2%	2%	1.4%

Note: This table examines the importance of changing beliefs about the true model. Panel A decomposes the change in price into a weighted average of price changes of each model $dp_1 = p_{S,t} * \Delta p^S_t + (1-p_{S,t}) * \Delta p^{NS}_t$ and the model belief change times the difference in the model predictions for price $dp_2 = (p_{S,t} - p_{S,t-1}) * (p^S_{t-1} - p^{NS}_{t-1})$. Panel B examines the model predictions if the probability was constant at its median (across time and trials).

7.2. Quantitative importance of belief changes

Table 5 assesses the importance of changes in model beliefs for generating price volatility and return predictability. For the learning model in the paper we can calculate the price as the probability weighted average of the equilibrium price if the individual believed the stationary model was true: p^S_t, and the equilibrium price if the individual believed the non-stationary model was true p^{NS}_t:[28]

$$p^L_t = (1 - p_{s,t})p^{NS}_t + p_{s,t}p^S_t$$

Writing this in changes we find

$$\Delta p^L_t = p_{s,t}\Delta p^{NS}_t + (1 - p_{s,t})\Delta p^{NS}_t + (p_{s,t} - p_{s,t-1}) * (p^S_{t-1} - p^{NS}_{t-1})$$
$$\Delta p^L_t = dp^1_t + dp^2_t$$

[28] Because the price is risk adjusted using the variance of consumption, this formula does not hold exactly when risk aversion does not equal zero. However, since I use no risk aversion it holds exactly.

where $dp_t^1 = p_{s,t}\Delta p^{NS}{}_t + (1 - p_{s,t})\Delta p_t^{NS}$ and $dp_t^2 = (p_{s,t} - p_{s,t-1}) * (p_{t-1}^S - p_{t-1}^{NS})$. dp_t^1 represents the response of the learning price to news about dividends. It is a probability weighted average of the response under the two models. dp_t^2 represents the response of the learning price to changes in beliefs. This representation allows for the following variance decomposition:

$$1 = \frac{cov(dp_t^1, \Delta p_t^L)}{var(\Delta p_t^L)} + \frac{cov(dp_t^2, \Delta p_t^L)}{var(\Delta p_t^L)}$$

Table 5a reports the result of this variance decomposition, taking the median results across the 500 trials. I find that 70% of the variance in the learning price comes from the first term. That the model does not rely solely on belief changes to explain time varying volatility is a significant difference from Veronesi (1999). However belief changes are not irrelevant: 30% of the variance comes from the second term. This result indicates that 30% of the variance in the price comes from revisions in beliefs.

As a further check on the importance of changes in beliefs for the results, I examine the return predictability predictions of the constant probability model of Section 6.2. Recall, this model kept a constant probability weight on the non-stationary model. This version of the model has overreaction to news from believing the non-stationary model, but it is not time varying as the weight on the non-stationary model is constant. I find that the model with constant probability still generates the negative correlations, but it understates the correlation of the PE ratio with future returns (−0.18 vs. −0.23 for the full model) and the correlation of the PE ratio with future consumption growth (−0.15 vs. −0.21 for the full model). The constant probability model does a slightly worse job matching the negative correlation of returns over time and the correlation of consumption growth with future consumption growth. Though the model does slightly better in matching the correlation of consumption growth with future returns. The constant probability model generates noticeably less volatility of returns and consumption growth.

To summarize, changes in beliefs are important for the model predictions. Thirty percent of the variance in price changes comes from changes in beliefs and changes in beliefs help the model explain the correlation of the PE ratio with future returns.

8. Conclusion

In this paper I examine a novel explanation for asset pricing puzzles. Namely, that the agent is unable to determine if the dividend process is stationary (and shocks are temporary) or non-stationary (shocks are permanent). I embed this uncertainty into a consumption based asset pricing model using Bayesian learning (with a bias towards current observations) to generate probability weights on the two processes. While the true dividend process is stationary, after a sequence of random shocks which results in the dividend series being far from its mean, the agent begins to think the non-stationary model is more likely. This leads to large swings in the equilibrium price of the equity asset which are subsequently reversed when the stationary series reverts back to its mean.

As a result the model is able to explain many asset pricing puzzles. First current returns and the PE ratio are negatively correlated with future excess returns. Consumption growth negatively correlates with future returns and both the PE ratio and consumption growth are negatively correlated with future consumption growth. The model matches the standard deviation of consumption and greatly amplifies returns over the benchmark rational expectations model. Finally, the model generates time varying volatility consistent with what is observed in the quarterly return series.

This paper has focused on asset pricing puzzles which are averages over the whole time series. One of the interesting features of the learning model is that the magnitude of incorrect beliefs is not constant over time but varies with the history of dividends. This observation suggests that the model may have implications for other less conventional statistics. For example, it may be useful in understanding infrequent but large changes in asset prices. It may also be helpful in explaining why the asset pricing correlations are so weak: as there are periods in time when the agent believes the true model and these correlations would go away, and periods of time when they believe the incorrect model, and these correlations would show up more strongly. I hope to explore these implications in future work.

Appendix A. Data and bootstrap calculations

Data begin in 1929 and end in 2013. Consumption data is real per-capita consumption of non-durables and services and come from the National Income and Product Accounts (NIPA) available from the Bureau of Economic Analysis (BEA).[29] The price to earnings ratio is the price of the S&P 500 index divided by average annual earning for S&P 500 companies over the current and previous 9 years. Dividend data are dividends accruing to the index.[30] Returns are excess returns measured as the value-weighted return on all NYSE, AMEX, and NASDAQ listed firms minus the one month T-bill rate.[31]

[29] The data are available at http://www.bea.gov/iTable/index_nipa.cfm. Consumption data are in table 2.3.5 and price deflators are in table 2.3.4. Population data are in Table 2.1.

[30] Data are come from the website of Robert Shiller: http://aida.wss.yale.edu/~shiller/data.htm. The PE ratio is the year end data from the monthly PE series. The dividend data are yearly averages of the monthly dividend data.

[31] Data are available from Kenneth R. French's online data library (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html) and are described as: "Rm-Rf, the excess return on the market, value-weight return of all CRSP firms incorporated in the US and listed on the NYSE, AMEX, or NASDAQ

To assess the statistical significance of the first six negative correlations, I run the following bootstrap exercise. I generate a simulated time series for excess returns (r_t), consumption growth ($\Delta \ln c_t$), the PE ratio ($P/E_{10,\,t}$) and future dividends d_t. I do this first by estimating AR(1) models for consumption growth and the PE ratio and an AR(4) model with a time trend for the log dividend process.[32,33] Then I sample with replacement from the excess return series and the residual series for the consumption, PE and dividend regressions. I draw one random year and use the excess return and residuals that correspond to that year. I then use the residuals to calculate the current period value of consumption growth, PE ratio, and dividend using the residual and last period's value.[34] I continue this way until I have a series of length 85 (the length of the original data set). I then calculate the correlation statistics as above. I repeat this process 1000 times and report the mean and 5% and 95% percentiles for the statistics.

For time varying volatility calculations data are the same as for predictability calculations, with the exception that for kurtosis estimates I use only data beginning in 1937. This is because in the first decade of the data (1932:Q3 and 1933:Q2) there are 76% and 88% excess returns that if included would make quarterly kurtosis equal to 18. The bootstrap exercise for the time varying volatility estimates is based on simulating 1000 return and PE series of length equal to the data length. To simulate the return series I draw from a normal distribution with the same mean and standard deviation as the return series in the data. To simulate the PE series I use the bootstrapping procedure outlined above but draw residuals from a normal distribution with the same standard deviation of the regression residuals as opposed to the regression residuals directly.

Appendix B. Consumption rule and asset price

The agent maximizes:

$$\max_{c_{t+s}} \sum_{s=0}^{\infty} \delta^s \frac{-1}{\alpha} exp[-\alpha(c_{t+s} - \gamma c_{t+s-1})]$$

subject to:

$$\begin{aligned} w_t &= -Rb_t + \Theta_{t-1} d_t + \Theta_{t-1} p_t \\ b_{t+1} &= c_t + \Theta_t p_t - w_t - y \\ d_t &= \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \ldots + \rho_p^s d_{t-p} + \varepsilon_t^s \text{ with } p = p_{s,t} \\ \Delta d_t &= \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \ldots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns} \text{ with } p = 1 - p_{s,t} \end{aligned}$$

The forecasting vector for future dividends is given by:

$$\overrightarrow{d_t} = \begin{bmatrix} 1 & d_t & \triangle d_t & \ldots & \Delta d_{t-p+1} & 1 & t & d_t & \ldots & d_{t-p+1} \end{bmatrix}'$$

Then $\vec{d}_{t+1} = \Phi \vec{d}_t$ where

$$\Phi = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & & & & & & \\ \alpha^{ns} & 1 & \rho_1^{ns} & \rho_2^{ns} & \cdots & \rho_p^{ns} & & & & & & \\ \alpha^{ns} & 0 & \rho_1^{ns} & \rho_2^{ns} & \ldots & \rho_p^{ns} & & & & 0 & & \\ 0 & 0 & 1 & 0 & 0 & 0 & & & & & & \\ 0 & 0 & 0 & \ddots & 0 & 0 & & & & & & \\ 0 & 0 & 0 & 0 & 1 & 0 & & & & & & \\ & & & & & & 1 & 0 & 0 & 0 & 0 & 0 \\ & & & & & & 1 & 1 & 0 & 0 & 0 & 0 \\ & & 0 & & & & \alpha^s & \gamma^s & \rho_1^s & \ldots & \ldots & \rho_p^s \\ & & & & & & 0 & 0 & 1 & 0 & 0 & 0 \\ & & & & & & 0 & 0 & 0 & \ddots & 0 & 0 \\ & & & & & & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Therefore $E_t d_{t+s} = E_t \tilde{e}' \vec{d}_{t+s}$ where $\tilde{e}' = [0 \quad 1 - p_{s,t} \quad \vec{0}_p \quad 0 \quad 0 \quad p_{s,t} \quad \vec{0}_{p-1}]$

that have a CRSP share code of 10 or 11 at the beginning of month t, good shares and price data at the beginning of t, and good return data for t minus the one-month Treasury bill rate (from Ibbotson Associates)."

[32] I correct the coefficient on the PE regression for small sample bias using the methods of Andrews (1981).

[33] I use a trend stationary model for dividends to be consistent with the model in the paper, but I get similar results here using a difference stationary model.

[34] To initialize I begin at a randomly chosen observation for the PE ratio and consumption growth and the actual initial observations of the dividend series.

We will define the state vector $z_t = \begin{bmatrix} b_t & c_{t-1} & 1 & y & d_t & \vec{d_t} \end{bmatrix}$ and guess a linear policy function $c_t = P'z_t$. The evolution of the state vector satisfies: $z_t = Mz_{t-1} + C\varepsilon_t$ where

$$M = \begin{bmatrix} R & 0 & 0 & -1 & -1 & \vec{0}'_{2p+4} \\ 0 & 0 & 0 & 0 & 0 & \vec{0}'_{2p+4} \\ 0 & 0 & 1 & 0 & 0 & \vec{0}'_{2p+4} \\ 0 & 0 & 0 & 1 & 0 & \vec{0}'_{2p+4} \\ 0 & 0 & 0 & 0 & 0 & \tilde{e}'\Phi \\ & & 0_{2p+4,5} & & & \Phi \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vec{0}_{2p+4} \end{bmatrix} P'$$

$$C = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ p_{s,t}\sigma^s + (1-p_{s,t})\sigma^{ns} \\ 0 \\ \sigma^{ns} \\ \sigma^{ns} \\ \vec{0}_{p-1} \\ 0 \\ 0 \\ \sigma^s \\ \vec{0}_{p-1} \end{bmatrix}$$

which implies $z_t = \bar{M}z_{t-1} + Nc_{t-1} + C\varepsilon_t$.

We will guess a solution to the Bellman equation $V(z_t) = \frac{-\Psi}{\alpha} exp[-\alpha(c_t - \gamma c_{t-1})]$. Now let $\widetilde{P} = P - \gamma \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \end{bmatrix}$ so $\widetilde{P}'z_t = c_t - \gamma c_{t-1}$. Then the Bellman equation becomes

$$V(z_t) = \frac{-1}{\alpha} exp[-\alpha(c_t - \gamma c_{t-1})] + \delta E_t V(z_{t+1})$$

$$\frac{-\Psi}{\alpha} exp[-\alpha(c_t - \gamma c_{t-1})] = \frac{-1}{\alpha} exp[-\alpha(\widetilde{P}'z_t)] + \delta E_t \frac{-\Psi}{\alpha} exp[-\alpha(\widetilde{P}'z_{t+1})]$$

Now to evaluate

$$E_t \frac{-\Psi}{\alpha} exp[-\alpha(\widetilde{P}'z_{t+1})]$$

$$E_t \frac{-\Psi}{\alpha} exp[-\alpha(\widetilde{P}'Mz_t + \widetilde{P}'C\varepsilon_{t+1})]$$

$$\frac{-\Psi}{\alpha} exp\left[-\alpha\left(\widetilde{P}'Mz_t + \frac{\alpha^2}{2}C'\widetilde{P}\widetilde{P}'C\right)\right]$$

$$\frac{-\Psi}{\alpha} exp\left(-\alpha\left[\left(\widetilde{P}'Mz_t - \frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right)\right]\right)$$

So the Bellman equation becomes:

$$\frac{-\Psi}{\alpha} exp[-\alpha(\widetilde{P}'z_t)] = \frac{-1}{\alpha} exp[-\alpha(\widetilde{P}'z_t)] - \delta\frac{\Psi}{\alpha} exp\left(-\alpha\left[\left(\widetilde{P}'Mz_t - \frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right)\right]\right)\Bigg]$$

This gives

$$\Psi = 1 + \delta\Psi exp\left[-\alpha\left(\left[(\widetilde{P}'M - I)z_t - \frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]\right)\right] \tag{B.1}$$

Since this must hold for all $z_t$we can conclude that $(\widetilde{P}'M - I)z_t$ is a constant.

We can also derive the first order condition for the optimal consumption choice

$$
\begin{aligned}
&exp[-\alpha(\widetilde{P}'z_t)]\frac{d\widetilde{P}'z_t}{dc_t} + \delta\Psi E_t exp[-\alpha(\widetilde{P}'z_{t+1})]\frac{d\widetilde{P}'z_{t+1}}{dc_t} = 0\\
&exp[-\alpha(\widetilde{P}'z_t)] + \delta\Psi\widetilde{P}'NE_t exp[-\alpha(\widetilde{P}'z_{t+1})] = 0\\
&exp[-\alpha(\widetilde{P}'z_t)] + \delta\Psi\widetilde{P}'Nexp\left(-\alpha\left[\left(\widetilde{P}'Mz_t - \frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right)\right]\right) = 0\\
&1 + \delta\Psi\widetilde{P}'Nexp\left[-\alpha\left(\left[(\widetilde{P}'M - I)z_t - \frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]\right)\right] = 0
\end{aligned}
$$

From (B.1) we have

$$
\begin{aligned}
&1 + \delta\Psi\widetilde{P}'N\frac{\Psi - 1}{\delta\Psi} = 0\\
&1 + \widetilde{P}'N(\Psi - 1) = 0\\
&\Psi = 1 - \frac{1}{\widetilde{P}'N}
\end{aligned}
$$

We know proceed by guessing the policy function.

$$
P = \begin{bmatrix} -(R-1)\left(1 - \frac{\gamma}{R}\right)\\ \frac{\gamma}{R}\\ \frac{R-\gamma}{R}\\ Q\\ \phi\\ \left[K\tilde{e}'(I-\Phi)^{-1}\frac{\Phi}{R}\right]' \end{bmatrix}
$$

Here Q and ϕ are constants to be determined. This guess implies:

$$
\tilde{P} = \begin{bmatrix} -(R-1)\left(1 - \frac{\gamma}{R}\right)\\ \frac{\gamma}{R} - \gamma\\ Q\\ \frac{R-\gamma}{R}\\ \phi\\ \left[K\tilde{e}'\left(I - \frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]' \end{bmatrix}
$$

Now note that $\widetilde{P}'N = -(R-1)(1-\frac{\gamma}{R}) + \frac{\gamma}{R} - \gamma = (1-R)$ so $\Psi = 1 - \frac{1}{1-R} = \frac{R}{R-1}$.

Now to validate the guess for P we show that $\widetilde{P}'(M-I)z_t$ is constant for this choice of P. First note that $M - I = \bar{M} - I + NP' =$

$$
\begin{bmatrix}
R-1 & 0 & 0 & -1 & -1 & \vec{0}'_{2p+4}\\
0 & -1 & 0 & 0 & 0 & \vec{0}'_{2p+4}\\
0 & 0 & 0 & 0 & 0 & \vec{0}'_{2p+4}\\
0 & 0 & 0 & 0 & 0 & \vec{0}'_{2p+4}\\
0 & 0 & 0 & 0 & -1 & \tilde{e}'\Phi\\
 & & 0_{2p+4,5} & & & \Phi - I
\end{bmatrix}
+
\begin{bmatrix}
-(R-1)\left(1-\frac{\gamma}{R}\right) & \frac{\gamma}{R} & Q & \frac{R-\gamma}{R} & \phi & \left[K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]\\
-(R-1)\left(1-\frac{\gamma}{R}\right) & \frac{\gamma}{R} & Q & \frac{R-\gamma}{R} & \phi & \left[K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]\\
 & & & 0_{2p+7,2p+9} & & \\
\end{bmatrix}
$$

For the first element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)(R-1)\frac{\gamma}{R} - \left(\frac{\gamma}{R}-\gamma\right)(R-1)\left(1-\frac{\gamma}{R}\right) = 0$$

For the second element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)\frac{\gamma}{R} + \left(\frac{\gamma}{R}-\gamma\right)\left(\frac{\gamma}{R}-1\right)$$
$$\left(\frac{\gamma}{R}-\gamma\right)\left(1-\frac{\gamma}{R}\right) + \left(\frac{\gamma}{R}-\gamma\right)\left(\frac{\gamma}{R}-1\right) = 0$$

For the third element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)Q + \left(\frac{\gamma}{R}-\gamma\right)Q$$
$$\left(-R-\frac{\gamma}{R}+1+\gamma\right)Q + \left(\frac{\gamma}{R}-\gamma\right)Q$$
$$(1-R)Q$$

For the fourth element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)\frac{-\gamma}{R} + \left(\frac{\gamma}{R}-\gamma\right)\left(1-\frac{\gamma}{R}\right) = 0$$

For the fifth element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)(\phi-1) + \left(\frac{\gamma}{R}-\gamma\right)\phi - \phi$$
$$-\left(R-\gamma-1+\frac{\gamma}{R}\right)(\phi-1) + \left(\frac{\gamma}{R}-\gamma\right)\phi - \phi$$
$$(1-R)\phi - (1-R)\left(1-\frac{\gamma}{R}\right) - \phi$$
$$\phi = \frac{(R-1)(1-\frac{\gamma}{R})}{R}$$

For the sixth element of $(\widetilde{P}'M - I)z_t$ we have

$$-(R-1)\left(1-\frac{\gamma}{R}\right)\left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right) + \left(\frac{\gamma}{R}-\gamma\right)\left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right) + \phi\tilde{e}'\Phi + \left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right)(\Phi - I)$$
$$(1-R)\left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right) + \frac{(R-1)(1-\frac{\gamma}{R})}{R}\tilde{e}'\Phi + \left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right)(\Phi - I)$$
$$-R\left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right) + \frac{(R-1)\left(1-\frac{\gamma}{R}\right)}{R}\tilde{e}'\Phi + R\left(K\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right)\frac{\Phi}{R}$$

Now note that

$$\frac{\Phi}{R}\frac{\Phi}{R} = \frac{\Phi}{R} - \frac{\Phi}{R} + \frac{\Phi}{R}\frac{\Phi}{R}$$
$$\frac{\Phi}{R}\frac{\Phi}{R} = \left(I-\frac{\Phi}{R}\right)\frac{-\Phi}{R} + \frac{\Phi}{R}$$
$$\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\frac{\Phi}{R} = \left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R} - \frac{\Phi}{R}$$

So

$$\frac{(R-1)(1-\frac{\gamma}{R})}{R}\tilde{e}'\Phi + -K\tilde{e}'\Phi = 0$$
$$K = \frac{(R-1)(1-\frac{\gamma}{R})}{R}$$

Now we can solve for Q from the first order condition

$$1+\delta\Psi\widetilde{P}'Nexp\left[-\alpha\left(\left[(\widetilde{P}'M-I)z_t-\frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]\right)\right]=0$$

$$1+\delta\Psi\widetilde{P}'Nexp\left[-\alpha\left(\left[Q(1-R)-\frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]\right)\right]=0$$

From above $\Psi\widetilde{P}'N=-R$ so

$$1-\delta Rexp\left[-\alpha\left(\left[Q(1-R)-\frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]\right)\right]=0$$

$$\ln(\delta R)-\alpha Q(1-R)+\frac{\alpha^2}{2}C'\widetilde{P}\widetilde{P}'C=0$$

$$Q=\frac{1}{R-1}\left[\frac{-1}{\alpha}\ln(\delta R)-\frac{\alpha}{2}C'\widetilde{P}\widetilde{P}'C\right]$$

Now note that $\widetilde{P}'C=$

$$\widetilde{P}_5'C_5+\widetilde{P}_7'C_7+\widetilde{P}_8'C_8+\widetilde{P}_{8+p+2}'C_{8+p+2}$$

$$\widetilde{P}_5'(p_{s,t}\sigma^s+(1-p_{s,t})\sigma^{ns})+\widetilde{P}_7'\sigma^{ns}+\widetilde{P}_8'\sigma^{ns}+\widetilde{P}_{8+p+2}'\sigma^s$$

$$\frac{(R-1)(1-\frac{\gamma}{R})}{R}(p_{s,t}\sigma^s+(1-p_{s,t})\sigma^{ns})+\left[\frac{(R-1)(1-\frac{\gamma}{R})}{R}\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]\begin{bmatrix}0\\ \sigma^{ns}\\ \sigma^{ns}\\ \vec{0}_{p+1}\\ \sigma^s\\ \vec{0}_{p-1}\end{bmatrix}$$

$$\frac{(R-1)(1-\frac{\gamma}{R})}{R}\left(p_{s,t}\left(I-\frac{\Phi}{R}\right)^{-1}_{p+5,p+5}\sigma^s+(1-p_{s,t})\left(I-\frac{\Phi}{R}\right)^{-1}_{2,2}\sigma^{ns}+(1-p_{s,t})\left[\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]_{2,3}\sigma^{ns}\right)=\sigma_c$$

Therefore

$$Q=\frac{1}{R-1}\left[\frac{-1}{\alpha}\ln(\delta R)-\frac{\alpha}{2}\sigma_c^2\right]$$

Now to calculate the rule for $c_t=P'z_t$ we have $c_t=$

$$-(R-1)\left(1-\frac{\gamma}{R}\right)b_t+\frac{\gamma}{R}c_{t-1}+Q+\frac{(R-1)(1-\frac{\gamma}{R})}{R}d_t+\left[\frac{(R-1)(1-\frac{\gamma}{R})}{R}\tilde{e}'\left(I-\frac{\Phi}{R}\right)^{-1}\frac{\Phi}{R}\right]\vec{d}_t$$

$$\frac{\gamma}{R}c_{t-1}-\frac{1}{R-1}\left[\frac{1}{\alpha}\ln(\delta R)+\frac{\alpha}{2}\sigma_c^2\right]+\left(1-\frac{\gamma}{R}\right)\frac{R-1}{R}\left[-Rb_t+d_t+E_t\sum_{s=1}^{\infty}\frac{d_{t+s}}{R^s}\right]$$

To calculate the market clearing asset price we follow Fuster et al. (2012) and write the Belman equation as

$$V(z_t)=\max_{c_t,\theta_t}u(c_t,c_{t-1})+\delta E_tV(z_{t+1})$$

where $z_t=[b_t \quad c_{t-1} \quad 1 \quad y \quad d_t \quad \vec{d}_t]'$. We can also write the optimal consumption choice as: $c_t=\frac{\gamma}{R}c_{t-1}+(1-\frac{\gamma}{R})x_t-\psi$ where $x_t=\frac{R-1}{R}[-Rb_t+d_t+E_t\sum_{s=1}^{\infty}\frac{d_{t+s}}{R^s}]$ and $\psi=\frac{1}{R-1}[\frac{1}{\alpha}\ln(\delta R)+\frac{\alpha}{2}\sigma_c^2]$. We solve for the optimal price by allowing for $\theta\neq1$ and then find the price for which supply equal demand, i.e. $\theta=1$. If the agent chooses $\theta\neq1$ the annuity value of wealth x_t becomes $x_t=\frac{R-1}{R}[-Rb_t-(\theta-1)p_t+d_t+\theta E_t\sum_{s=1}^{\infty}\frac{d_{t+s}}{R^s}]$. He pays p_t for any shares above 1 (his holdings last period) and expects θ times the discounted present value of dividends in future dividend income. The first order condition for optimal choice of θ is

$$\frac{dV}{d\theta}=\frac{\partial V}{\partial c}\left[\frac{\partial c_t}{\partial x}\frac{\partial x_t}{\partial\theta}+\frac{\partial c_t}{\partial\psi}\frac{\partial\psi}{\partial\theta}\right]=0$$

The term in brackets must be zero at $\theta=1$. Therefore

$$\left(1-\frac{\gamma}{R}\right)\frac{R-1}{R}\left[-p_t+E_t\sum_{s=1}^{\infty}\frac{d_{t+s}}{R^s}\right]=-\frac{1}{R-1}\left[\alpha\sigma_c^2\right].$$

Which gives

$$p_t = E_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s} - \frac{R\alpha\sigma_c^2}{(1-\frac{\gamma}{R})(R-1)^2}$$

Appendix C. Belief formation

Following Cogley and Sargent (2005) for a given model (i.e. the stationary or non-stationary) indexed by $i = \{s, ns\}$, and a dividend history D^{t-1}, we assume that agents prior beliefs about the model parameters $\Upsilon_{i,\,t}$ are distributed normally according to:

$$p(\Upsilon_{i,t-1}|\sigma_i^2, D^{t-1}) = N(\Upsilon_{i,t-1}, \sigma_i^2 P_{i,t-1}^{-1})$$

and their prior beliefs concerning the model residual variance are given by:

$$p(\sigma_{i,t-1}^2|D^{t-1}) = IG(s_{i,t-1}, \nu_{i,t-1})$$

Here N represents the normal distribution function and IG represents the inverse-gamma distribution function. P_{t-1} is the precision matrix that captures the confidence the agent has in his belief for $\Upsilon_{i,t-1}$, σ_i^2 is the estimate of the variance of the model residuals, $s_{i,t-1}$ is an analogue to the sum of squared residuals, and $\nu_{i,t-1}$ is a measure of the degrees of freedom to calculate the residual variance such that the point estimate of $\sigma_{i,t-1}^2$ is given by $s_{i,t-1}/\nu_{i,t-1}$. After observing the dividend d_t the agent's posterior beliefs are given by:

$$p(\Upsilon_{i,t}|\sigma_i^2, D^t) = N(\Upsilon_{i,t}, \sigma_i^2 P_{i,t}^{-1})$$
$$p(\sigma_i^2|D^t) = IG(s_{i,t}, \nu_{i,t})$$

where the distribution parameters follow the recursion presented in the main body of the paper.

Given a set of model parameters: $\{\Upsilon_i, \sigma_i\}$ we can calculate the conditional likelihood of the model as:

$$L(\Upsilon_i, \sigma_i^2, D^t) = \prod_{s=1}^{t} p(y_s|x_s, \Upsilon_i, \sigma_i^2)$$

where y_s and x_s are the left and right hand side variables of the model at time s and D^t is the dividend history up to time t. Based on this likelihood, one can write the marginalized likelihood of the model by integrating over all possible parameters:

$$m_{it} = \iint L(\Upsilon_i, \sigma_i^2, D^t)\, p(\Upsilon_i, \sigma_i^2)\, d\Upsilon_i d\sigma_i^2$$

Then we have the probability of the model given the observed data $p(M_i|D^t) \propto m_{i,t}\, p(M_i) \equiv w_{i,t}$. Here we have defined the weight on model i, $w_{i,\,t}$ and $p(M_i)$ is the prior probability on model i. Cogley and Sargent (2005) show that Bayes's rule implies

$$m_{it} = \frac{L(\Upsilon_i, \sigma_i^2, D^t)\, p(\Upsilon_i, \sigma_i^2)}{p(\Upsilon_i, \sigma_i^2|D_t)}$$

and therefore

$$\frac{w_{i,t+1}}{w_{i,t}} = \frac{m_{i,t+1}}{m_{i,t}} = p(y_{i,t+1}|x_{i,t}, \Upsilon_i, \sigma_i^2)\frac{p(\Upsilon_i, \sigma_i^2|D_t)}{p(\Upsilon_i, \sigma_i^2|D_{t+1})}$$

We assume that regression residuals are normally distributed allowing us to use the normal p.d.f to calculate $p(y_{i,t+1}|x_{i,t}, \Upsilon_i, \sigma_i^2)$. Cogley and Sargent (2005) show that $p(\Upsilon_i, \sigma_i^2|D_t)$ is given by the normal-inverse gamma distribution and provide the analytical expressions for this probability distribution. Any choice of Υ_i, σ_i^2 will give the same ratio of weights; I use the posterior mean in my calculations.

Appendix D. Gain calibration

To calibrate the gain value I use data from the Survey of Professional Forecasters conducted by the Federal Reserve Bank of Philadelphia on the median forecasts of corporate profits after tax. Data are available at: (https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters). Letting c_i be the forecast of corporate profits i quarters ahead, I can calculate annualized growth rate forecasts as: $g_i^f = 100[(\frac{4}{(i+1)}\Delta \ln(\frac{c_i}{c_{-1}})]$. Here c_{-i} is the previous quarter's value of corporate profits which is known to the participants when they forecast corporate profits. I then calculate the analogue of these forecasts using data on corporate profits and the learning algorithms in the paper, the model learning algorithm and the constant gain algorithm. Priors are set to the OLS estimates using data from 1955 to 1968:Q3 the quarter before forecast data are available. Forecasts begin in 1968:Q3 and end in 2017:Q3. Model forecasts are based on the corporate profits after tax without inventory valuation adjustment (IVA) and capital consumption adjustment (CCAdj) until 2006:Q1

when forecasters are then asked to forecast the variable with IVA and CCAdj. Then to calibrate the gain, I follow Branch and Evans (2006) to minimize the distance between the model growth rate forecast predictions and the data from the SPF. Here I sum over the forecasts 1 quarter ahead to 4 quarters ahead:

$$\min_{g} \sum_{i=1}^{4} sqrt\left(\frac{1}{T}\sum_{t}\left[g_i^f(t;SPF) - g_i^f(t;Model)\right]^2\right)$$

Appendix E. Constant gain learning

To facilitate comparison to the main learning model in the paper, when considering constant gain learning, I use the following recursive formulation:

$$P_{i,t} = (1-g)P_{i,t-1} + x_{i,t}x'_{i,t}$$
$$\Upsilon_{i,t} = P_{i,t}^{-1}((1-g)P_{i,t-1}\Upsilon_{i,t-1} + x_{i,t}y_{i,t})$$

This is equivalent to the standard constant gain learning algorithm (see Evans and Honkapohja (2001) pp. 334).

$$\Upsilon_{i,t} = \Upsilon_{i,t-1} + gR_t^{-1}x_{i,t}(y_{i,t} - \Upsilon'_{i,t-1}x_{i,t})$$
$$R_t = R_{t-1} + g(x_t x'_t - R_{t-1})$$

To see this, let $P_{i,t} = \frac{1}{g}R_t$, then

$$P_{i,t} = (1-g)P_{i,t-1} + x_{i,t}x'_{i,t}$$
$$P_{i,t} = P_{i,t-1} - gP_{i,t-1} + x_{i,t}x'_{i,t}$$
$$gP_{i,t} = gP_{i,t-1} + g(x_{i,t}x'_{i,t} - gP_{i,t-1})$$
$$R_t = R_{t-1} + g(x_t x'_t - R_{t-1})$$

And

$$\Upsilon_{i,t} = P_{i,t}^{-1}((1-g)P_{i,t-1}\Upsilon_{i,t-1} + x_{i,t}y_{i,t})$$
$$\Upsilon_{i,t} = (1-g)P_{i,t}^{-1}P_{i,t-1}\Upsilon_{i,t-1} + P_{i,t}^{-1}x_{i,t}y_{i,t}$$
$$\Upsilon_{i,t} = [I - P_{i,t}^{-1}x_{i,t}x'_{i,t}]\Upsilon_{i,t-1} + P_{i,t}^{-1}x_{i,t}y_{i,t}$$
$$\Upsilon_{i,t} = \Upsilon_{i,t-1} + P_{i,t}^{-1}x_{i,t}(y_{i,t} - x'_{i,t}\Upsilon_{i,t-1})$$
$$\Upsilon_{i,t} = \Upsilon_{i,t-1} + gR_t^{-1}x_{i,t}(y_{i,t} - \Upsilon'_{i,t-1}x_{i,t})$$

References

Adam, K., Kuang, P., Marcet, A., 2012. House price booms and the current account. NBER Macroecon. Annu. 26 (1), 77–122.
Adam, K., Marcet, A., Nicolini, J.P., 2016. Stock market volatility and learning. J. Finance 71 (1), 33–82.
Adam, K., Marcet, A., Beutel, J., 2017. stock price booms and expected capital gains. Am. Econ. Rev. 107 (8), 2352–2408.
Alessie, R., Lusardi, A., 1997. Consumption, saving and habit formation. Econ. Lett. 55 (1), 103–108.
Andrews, D.W., 1981. Exactly median-unbiased estimation of first order autoregressive/unit root models. Econometrica 61 (1), 139–165.
Bansal, R., Yaron, A., 2004. Risks for the long run: a potential resolution of asset pricing puzzles. J. Finance 59 (4), 1481–1509.
Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2015. X-CAPM: an extrapolative capital asset pricing model. J. Financ. Econ. 115 (1), 1–24.
Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment1. J. Financ. Econ. 49 (3), 307–343.
Barsky, R.B., De Long, J.B., 1993. Why does the stock market fluctuate? Q. J. Econ. 108 (2), 291–311.
Beshears, J., Choi, J.J., Fuster, A., Laibson, D., Madrian, B.C., 2013. What goes up must come down? Experimental evidence on intuitive forecasting. Am. Econ. Rev. 103 (3), 570–574.
Branch, W.A., Evans, G.W., 2006. A simple recursive forecasting model. Econ. Lett. 91 (2), 158–166.
Caballero, R.J., 1990. Consumption puzzles and precautionary savings. J. Monet. Econ. 25 (1), 113–136.
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption based explanation of aggregate stock market behavior. J. Polit. Econ. 107, 205–251.
Campbell, J.Y., Shiller, R.J., 1988. Stock prices, earnings, and expected dividends. J. Finance 43 (3), 661–676.
Cochrane, J.H., 1988. How big is the random walk in gnp? J. Polit. Econ. 96 (5), 893–920.
Cochrane, J.H., 1991. A critique of the application of unit root tests. J. Econ. Dyn. Contr. 15 (2), 275–284.
Cogley, T., Sargent, T.J., 2005. The conquest of us inflation: learning and robustness to model uncertainty. Rev. Econ. Dyn. 8 (2), 528–563.
Collin-Dufresne, P., Johannes, M., Lochstoer, L.A., 2016. Parameter learning in general equilibrium: The asset pricing implications. Am. Econ. Rev. 106 (3), 664–698.
Cutler, D.M., Poterba, J., Summers, L., 1989. What moves stock prices? J. Portfolio Manage. 15, 4–12.
David, A., Veronesi, P., 2013. What ties return volatilities to price valuations and fundamentals? J. Polit. Econ. 121 (4), 682–746.
De Bondt, W.F.M., Thaler, R., 1985. Does the stock market overreact? J. Finance 40 (3), 793–805.
De Bondt, W.F.M., Thaler, R.H., 1987. Further evidence on investor overreaction and stock market seasonalit y. J. Finance 42 (3), 557–581.
De Bondt, W.F.M., Thaler, R.H., 1990. Do security analysts overreact? Am. Econ. Rev. 80 (2), 52–57.
Deaton, A., 1992. Understanding Consumption. Oxford University Press.
Engle, R.F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. Econometrica 50 (4), 987–1007.
Ertan, A., Karolyi, S.A., Kelly, P., Stoumbos, R.C., 2016. Individual Investor Over-Extrapolation. In: 7th Miami Behavioral Finance Conference. Technical Report
Evans, G., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.
Fama, E.F., French, K.R., 1988. Permanent and temporary components of stock prices. J. Polit. Econ. 96 (2), 246–273.

Fisher, K., Statman, M., 2002. Blowing bubbles. J. Psychol. Financial Markets 3, 53–65.
French, K.R., Schwert, G.W., Stambaugh, R.F., 1987. Expected stock returns and volatility. J. Financ. Econ. 19 (1), 3–29.
Fuster, A., Hebert, B., Laibson, D., 2012. Natural expectations, macroeconomic dynamics, and asset pricing. NBER Macroecon Annu 26 (1), 1–48.
Gordon, M.J., 1959. Dividends, earnings, and stock prices. Rev. Econ. Stat. 41 (2), 99–105.
Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. Rev. Financial Stud. 27 (3), 714–746.
Hamilton, J., 1994. Time Series Analysis. Princeton University Press.
Johannes, M., Lochstoer, L.A., Mou, Y., 2016. Learning about consumption dynamics. J. Finance 71 (2), 551–600.
Johnson, T.C., 2001. Return dynamics when persistence is unobservable. Math. Finance 11 (4), 415–445.
Koop, G., Korobilis, D., 2012. Forecasting inflation using dynamic model averaging. Int. Econ. Rev. 53 (3), 867–886.
Kreps, D., 1998. Anticipated Utility and Dynamic Choice, A 1997 Schwartz Lecture. Frontiers of Research in Economic Theory.
Lakonishok, J., Shleifer, A., Vishny, R.W., 1994. Contrarian investment, extrapolation, and risk. J. Finance 49 (5), 1541–1578.
Lam, P.-S., Cecchetti, S.G., Mark, N.C., 2000. Asset pricing with distorted beliefs: are equity returns too good to be true? Am. Econ. Rev. 90 (4), 787–805.
Lettau, M., Ludvigson, S., 2010. Measuring and modelling variation in the risk-return trade-off. Handb. Financial Econometrics 1, 617–690.
Pastor, L., Veronesi, P., 2009. Learning in financial markets. Annu. Rev. Financial Econ. 1 (1), 361–381.
Poterba, J.M., Summers, L.H., 1988. Mean reversion in stock prices : evidence and implications. J. Financ. Econ. 22 (1), 27–59.
Rabin, M., 2002. Inference by believers in the law of small numbers. Q. J. Econ. 117 (3), 775–816.
Rabin, M., Vayanos, D., 2010. The Gambler's and hot-Hand fallacies: theory and applications. Rev. Econ. Stud. 77 (2), 730–778.
Raftery, A., Karny, M., Ettler, P., 2010. Online prediction under model uncertainty via dynamic model averaging: application to a cold rolling mill. Technometrics 52 (1), 52–66.
Reinhart, C., Rogoff, K., 2009. This Time is Different: Eight Centuries of Financial Folly. Princeton University Press.
Roll, R., 1984. Orange juice and weather. Am. Econ. Rev. 74 (5), 861–880.
Schwert, G.W., 1989. Why does stock market volatility change over time? J. Finance 44 (5), 1115–1153.
Shiller, R., 2000. Measuring bubble expectatations and investor confidence. J. Psychol. Financial Markets 1, 49–60.
Shiller, R., 2005. Irrational Exuberence. Princeton University Press.
Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? Am. Econ. Rev. 71 (3), 421–436.
Stock, J.H., 1991. Confidence intervals for the largest autoregressive root in u.s. macroeconomic time series. J. Monetary Econ. 28 (3), 435–459.
Timmermann, A.G., 1993. How learning in financial markets generates excess volatility and predictability in stock prices. Q. J. Econ. 108 (4), 1135–1145.
Tortorice, D., 2015. Long Run Expectations, Learning and the U.S. Housing Market. Working Papers 85. Brandeis University, Department of Economics and International Businesss School.
Tversky, A., Kahneman, D., 1973. Availability: a Heuristic for judging frequency and probabilities. Cogn. Psychol. 5, 207–232.
Veronesi, P., 1999. Stock market overreaction to bad news in good times: a rational expectations equilibrium model. Rev. Financial Stud. 12 (5), 975–1007.
Vissing-Jorgensen, A., 2004. Perspectives on behavioral finance: does irrationality disappear with wealth? Evidence from expectations and actions. In: NBER Macroeconomics Annual 2003, Volume 18. National Bureau of Economic Research, Inc, pp. 139–208. NBER Chapters