

Equity Return Predictability, Time Varying Volatility and Learning About the Permanence of Shocks*

Daniel L Tortorice[†]

May 12, 2014

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.

JEL Codes: D83, D84, E21, G12

Keywords: Consumption, Savings, Asset Pricing, Learning, Expectations

*I thank Jens Hilscher, David Laibson, Blake LeBaron, Davide Pettenuzzo, and participants at the Brandeis University Economics Department Faculty Lunch 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.

[†]Department of Economics and International Business School, Brandeis University. MS 021. 415 South St. Waltham, MA 02452. tortoric@brandeis.edu. (781) 736 5208.

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 these 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 dividends. Optimal behavior implies that the price of the equity asset is related 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 and changes in 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 puts a substantial 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 as well. Furthermore, the model generates excess volatility of returns.

The model also makes several accurate predictions for consumption. 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.

My paper lands firmly in this literature.

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. However, my paper does share a motivation with Bansal and Yaron (2004). In their paper, the dividend growth rate has a very persistent but ultimately stationary component. They argue that it is difficult in small samples to distinguish this component from a purely i.i.d. process. Similarly in my paper it is hard in small samples to distinguish between a unit root and a near unit root process for dividends.

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 the most relevant papers. Barsky and De Long (1993) and Timmermann (1993) examine learning about the growth rate of dividends. Barsky and De Long (1993) assume agents take a weighted average of past growth rates to forecast future growth rates and show that this can explain excess volatility in the stock market. Timmermann (1993) examines a similar model using least squares learning to estimate the growth rate and shows that the model generates excess volatility and predictability of excess returns using the dividend yield. 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.

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. In their model, consumption growth fluctuates between high growth and low growth stages, and agents underestimate the persistence of these states. They show this assumption is helpful in explaining the equity premium and the predictability of excess returns. 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 (at times) overes-

timate the persistence of shocks. Additionally, Lam et al. (2000) do not consider learning. So 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 under-reaction 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, agent's beliefs in the incorrect model are exogenous and never change. 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.

The rest of the paper proceeds as follows: section 2 describes the data moments I attempt to explain, section 3 describes the consumption based asset pricing model and the formulation of beliefs, section 4 explains the model calibration and simulation, section 5 highlights the main mechanism of the model and the results from simulating the model, section 6 examines robustness of the results and the importance of changing beliefs for the key results and finally section 7 concludes.

2 Data and Data Moments

Many authors (e.g. Fama and French (1988); Lettau and Ludvigson (2010); Lakonishok et al. (1994); Poterba and Summers (1988); Campbell and Shiller (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 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.

2.1 Data

Data begin in 1929 and end in 2013. Consumption data is real per-capita consumption of non-durables and services and comes from the National Income and Product Accounts (NIPA) available from the Bureau of Economic Analysis (BEA).¹ 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.² 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.³

2.2 Return Predictability

To demonstrate and highlight the statistical significance of return predictability, I follow Fuster et al. (2012) and calculate these correlations using annual data beginning in 1929 and ending in 2013. The correlations are listed in Table 1a. The table first reports the correlation between the current excess return r_t and the cumulative return over the next 2 to 5 years: $r_{t+2} + \dots + r_{t+5}$. 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

¹The data are available at http://www.bea.gov/iTable/index_nipa.cfm. Consumption data come are in table 2.3.5 and price deflators are in table 2.3.4. Population data are in table 2.1.

²Data 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.

³Data 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 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)."

same cumulative return over the next 2 to 5 years: $r_{t+2} + \dots + r_{t+5}$. 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 $r_{t+2} + \dots + r_{t+5}$ is -0.34. Additionally, I find that future cumulative consumption growth $\Delta \ln c_{t+3} + \dots + \Delta \ln c_{t+6}$ 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 $\Delta \ln d_{t+2} + \dots + \Delta \ln d_{t+5}$ 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 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.^{4,5} 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.⁶ 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 1,000 times and 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

⁴I correct the coefficient on the PE for small sample bias using the methods of Andrews (1981).

⁵I 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.

⁶To initialize I begin at a randomly chosen observation for the PE ratio and consumption growth and the actual initial observations of the dividend series.

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.

2.3 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 (the fourth moment $\frac{E(x-\mu)^4}{\sigma^4}$) for quarterly returns and the PE ratio.⁷ In the data quarterly return kurtosis equals 4.1. In contrast, in returns were normal one would expected kurtosis equal to 3. The PE-ratio 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%.

To assess the statistical significance of these estimates I simulate 1,000 return and PE series of length equal to the data length and report the 5th and 95th percentiles of the statistics in the table. 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. 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 positive. The correlation at 1 lag is 0.08, its 0.01 at 2 lags, it increases to 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

⁷Quarterly data begin in 1927 and end in 2013. 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.

four.⁸

As a final check for time varying volatility of returns and the ability of the learning model to replicate this feature of the data, I estimate GARCH(1,1) models on the quarterly and annual 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 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 Model

3.1 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 .⁹ 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: $d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \dots + \rho_p^s d_{t-p} + \varepsilon_t^s$.¹⁰ However the agent does not know this. He puts some probability on the alternative that dividends follow the non-stationary (unit root process) $\Delta d_t = \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \dots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns}$. Since he does not know which process

⁸Standard errors are calculated as $\frac{1}{\sqrt{T}}$ the standard error under the null hypothesis of zero autocorrelation. See Hamilton (1994) pp. 111.

⁹The 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.

¹⁰I 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, for the results in the paper it is only key that the true process exhibits more mean reversion than the alternative considered process, not that the true process is trend stationary.

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.

3.2 Representative Agent Problem

The agent maximizes:

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

subject to:

$$w_t = -Rb_t + \Theta_{t-1}d_t + \Theta_{t-1}p_t \quad (2)$$

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

$$d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \dots + \rho_p^s d_{t-p} + \varepsilon_t^s \text{ with } p = p_{s,t}$$

$$\Delta d_t = \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \dots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns} \text{ with } p = 1 - p_{s,t}$$

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 assuming his future beliefs will be the same as his current beliefs. 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. 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: $d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \dots + \rho_p^s d_{t-p} + \varepsilon_t^s$ and with probability $p = 1 - p_{s,t}$ the dividend process is non-stationary $\Delta d_t = \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \dots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns}$.

For the utility function I use the following exponential utility function with habit formation: $u(c_t, c_{t-1}) = \frac{-1}{\alpha} \exp[-\alpha(c_{t+s} - \gamma c_{t+s-1})]$ (Caballero (1990); Alessie and Lusardi (1997)). 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 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. Secondly, it allows one to have realistic smooth adjustment in consumption without creating time varying risk aversion as in Campbell and Cochrane (1999).

3.3 Model Solution

The Bellman equation for the model is:

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

where the vector of state variables $z_t = \left[b_t \quad c_{t-1} \quad 1 \quad y \quad d_t \quad \vec{d}_t \right]'$. 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: $\left[1 \quad d_t \quad \Delta d_t \quad \dots \quad \Delta d_{t-p+1} \quad 1 \quad t \quad d_t \quad \dots \quad d_{t-p+1} \right]$.

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

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

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, the results of Fuster et al. (2012) imply the price of

the equity asset is given by:

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

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 $\hat{E}_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s} = p_{s,t} \left[E_t \sum_{s=1}^{\infty} \frac{d_{t+1}}{R^s} | S \right] + (1 - p_{s,t}) \left[E_t \sum_{s=1}^{\infty} \frac{d_{t+1}}{R^s} | NS \right]$. 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.

3.4 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. 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 are distributed normally according to:

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

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

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

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 $\Theta_{i,t-1}$, σ_i^2 is the estimate of the variance of the model residuals, s_{t-1} is an analogue to the sum of squared residuals, and v_{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_{t-1}/v_{t-1} . After observing the dividend d_t the agent's posterior beliefs are given by:

$$\begin{aligned} p(\Theta_{i,t} | \sigma_i^2, D^t) &= N(\Theta_{i,t}, \sigma_i^2 P_t^{-1}) \\ p(\sigma_i^2 | D^t) &= IG(s_t, v_t) \end{aligned}$$

Cogley and Sargent (2005) gives the following recursion to update the parameters of the beliefs:

$$\begin{aligned} P_t &= P_{t-1} + x_t x_t' \\ \theta_t &= P_t^{-1} (P_{t-1} \theta_{t-1} + x_t y_t) \\ s_t &= s_{t-1} + y_t^2 + \theta_{t-1}' P_{t-1} \theta_{t-1} - \theta_t' P_t \theta_t \\ v_t &= v_{t-1} + 1 \end{aligned}$$

Here x_t is the vector of right hand side variables for the model at time t and y_t is the left hand side variable for the model at time t . This recursion gives the parameters of each model. Now it is necessary to calculate the probability weight on each model.

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

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

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:

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

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

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

and therefore

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

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

the posterior mean in my calculations.

This recursion implies the following recursion for model weights.

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

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.

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

Finally, to calculate the model probabilities, the consumer normalizes the weights to one, and therefore the weight on the stationary model is given by:

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

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 $L(\Theta_i, \sigma_i^2, D^t) = \prod_{s=1}^t p(y_s|x_s, \Theta_i, \sigma_i^2)^{(1-g)^s}$. 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 directly into 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.

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.

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.

4 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.

I set the gain parameter equal to 0.075. I find that I need a gain value of 0.05 or higher to have non-trivial perpetual learning (i.e. failure to converge to the true process). I choose a higher value here to better fit the correlation between the PE ratio and future returns. 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 higher gain parameter but of similar magnitude as to what is chosen in the Bayesian dynamic model averaging literature. I examine robustness to alternative choices for the gain, $g = 0.05$ and $g = 0.1$ in section 6.

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 1,000 periods and use only the last $340 = 4 \times (2013 - 1928)$ 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 ($d_t = \alpha^s + \gamma^s t + \rho_1^s d_{t-1} + \dots + \rho_p^s d_{t-p} + \varepsilon_t^s$) and a non-stationary process ($\Delta d_t = \alpha^{ns} + \rho_1^{ns} \Delta d_{t-1} + \dots + \rho_p^{ns} \Delta d_{t-p} + \varepsilon_t^{ns}$) by ordinary least squares using data on the net operating surplus of private enterprises.¹¹ 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, $\Theta_{s,0}$ and $\Theta_{ns,0}$ 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 diffuse prior. From section 3.4, 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.¹²

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 1,000 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.¹³

5 Results

In this section I describe the results from the model. First I graph the impulse response functions for the two models considered and plot a sample draw of dividends and the cor-

¹¹I 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.

¹²I 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.

¹³Following 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.

responding model probabilities and prices. I examine these plots first to build intuition for the results of the model. Then I consider the simulated statistics for the model and compare them to the data.

5.1 Model Intuition

Figure 1: Impulse Responses: Non-Stationary Model (Solid) vs. Stationary Model (Dashed)





Figure 1 plots the (log) dividend impulse response functions to a one standard deviation shock to the dividend. On impact the effect on the dividend shock is the same. 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 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. I calibrate the model with an AR(4). Note that the model predictions four 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 (Stock (1991); Deaton (1992); Cochrane (1988)). 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.


Figure 2: Sample Price: Learning Model (Solid) vs. Stationary Model (Dashed)



To see the implication of learning about stationary versus non-stationary dividend processes, consider the sample price in figure 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 than 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 175 where prices rise cumulatively 16% and then subsequently crash and another around quarter 300 where prices rise 25% 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).

Figure 3: Sample Probability on the (True) Stationary Model




To see what drives these price crashes examine figure 3. What one sees is that around time 175 the probability on the stationary model goes from 80% to below 10%. Similarly, around time 300 the probability of the stationary model falls from 60% to below 10%. 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. Figure 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 state value. We see that around time 175, the dividend begins a sharp rise, increasing cumulatively 17%. At time 300 there is a 25% increase in the dividend. While the true

process is the stationary model, this 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. When the dividend begins to mean revert towards the end of the sample, the beliefs are reevaluated again. The agent puts much more weight on the stationary model, and the price crashes.

Figure 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.

5.2 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.¹⁴ All the other correlation are of the wrong sign. For example, the PE ratio is positively correlated with future returns and consumption growth.¹⁵ 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 -0.21 versus -0.2 in the data. It predicts a correlation of the PE ratio and future returns of -0.24 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 -0.22 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 -0.34 vs -0.25 in the data. The model also significantly amplifies volatility. It explains 16% of the volatility of returns observed in the data, 25 times more than the rational expectations benchmark and it matches the volatility of consumption exactly.

The middle column of table 2 allows us to demonstrate the importance of learning about

¹⁴Small sample bias prevents these correlations from being exactly zero.

¹⁵Again these correlation go to zero as the simulated time series length is increased.

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 is key to explaining the ability of the model to match the data.

To summarize, the model generates 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. Eventually, when dividends begin to mean revert prices do as well explaining 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.

5.3 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 and a model which puts 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.¹⁶ 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 6.37. 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 and 2.9 respectively. The learning model is able to amplify kurtosis, though only slightly, predicting a kurtosis of 3.1. 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%. 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

¹⁶This probability is 0.66.

1.96 standard deviations.

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 1 lag was 0.08, it was 0.01 at 2 lags, it increased to 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 0.21 at one lag down to 0.1 at four lags.

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.¹⁷ 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 $\gamma_1 = 0.55$ and $a_1 = 0.25$ versus $\gamma_1 = 0.61$ and $a_1 = 0.29$ 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.

6 Robustness and Further Analysis

This section examines the robustness of the results on return predictability to different parameter choices. 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

¹⁷This procedure is necessary because under the null of no GARCH the likelihood function is flat and I am unable to identify the GARCH parameters.

to the model where the agent puts a constant non-zero probability on the non-stationary model.

6.1 Robustness

The model is fairly tightly parametrized. However, I did need to set a risk free rate, gain parameters 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 2 I find correlations that are a little bit smaller, but most are within 0.04 of the main results. I still am able to generate the negative correlations in the data and get very 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 6 or 8 versus 4. These lag lengths all generate very similar negative correlations and standard deviations of consumption growth and returns.

For the gain parameter I find a very similar result. Lowering the gain from 0.075 to 0.05 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 0.1 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.

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 3.2% in the baseline case and the volatility of consumption is 2.4% versus 2% in the baseline case. This result occurs because higher risk free rates lowers consumption and prices and therefore changes in dividends have a larger proportional impact.

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. Although, one caveat remains. The gain parameter must be high enough

to allow for 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.

6.2 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_t^S , and the equilibrium price if the individual believed the non-stationary model p_t^{NS} was true.¹⁸

$$p_t^L = (1 - p_{s,t})p_t^{NS} + p_{s,t}p_t^S$$

Writing this in changes we find

$$\begin{aligned}\Delta p_t^L &= p_{s,t}\Delta p_t^{NS} + (1 - p_{s,t})\Delta p_t^{NS} + (p_{s,t} - p_{s,t-1}) * (p_{t-1}^S - p_{t-1}^{NS}) \\ \Delta p_t^L &= dp_t^1 + dp_t^2\end{aligned}$$

where $dp_t^1 = p_{s,t}\Delta p_t^{NS} + (1 - p_{s,t})\Delta p_t^{NS}$ 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:

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

Table 5a reports the result of this variance decomposition, taking the median results across the 500 trials. I find that 69% of the variance in the learning price comes from the first term, but 31% comes from the second term. This result indicates that more than 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 5.3. 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

¹⁸Because 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.

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. -0.24 for the full model) and the correlation of the PE ratio with future consumption growth (-0.16 vs. -0.21 for the full model). However this second correlation is the same as the data value which is -0.16. The constant probability model does a 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 30% less volatility in returns and 25% less volatility of 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.

7 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.

References

- Rob Alessie and Annamaria Lusardi. Consumption, saving and habit formation. *Economics Letters*, 55(1):103–108, August 1997.
- Donald W.K. Andrews. Exactly median-unbiased estimation of first order autoregressive/unit root models. *Econometrica*, 61(1):139–65, January 1981.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 08 2004.
- Nicholas Barberis, Andrei Shleifer, and Robert Vishny. A model of investor sentiment1. *Journal of Financial Economics*, 49(3):307–343, 1998.
- Robert B Barsky and J Bradford De Long. Why does the stock market fluctuate? *The Quarterly Journal of Economics*, 108(2):291–311, May 1993.
- Ricardo J. Caballero. Consumption puzzles and precautionary savings. *Journal of Monetary Economics*, 25(1):113–136, January 1990.
- John Y. Campbell and John H. Cochrane. By force of habit: A consumption based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107:205–251, 1999.
- John Y Campbell and Robert J Shiller. Stock prices, earnings, and expected dividends. *Journal of Finance*, 43(3):661–76, July 1988.
- John H Cochrane. How big is the random walk in gnp? *Journal of Political Economy*, 96(5):893–920, October 1988.

- Timothy Cogley and Thomas J. Sargent. The conquest of us inflation: Learning and robustness to model uncertainty. *Review of Economic Dynamics*, 8(2):528–563, April 2005.
- Angus Deaton. *Understanding Consumption*. Oxford University Press, 1992.
- Robert F Engle. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4):987–1007, July 1982.
- George Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- Eugene F Fama and Kenneth R French. Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–73, April 1988.
- Kenneth R. French, G. William Schwert, and Robert F. Stambaugh. Expected stock returns and volatility. *Journal of Financial Economics*, 19(1):3–29, September 1987.
- Andreas Fuster, Benjamin Hebert, and David Laibson. Natural expectations, macroeconomic dynamics, and asset pricing. *NBER Macroeconomics Annual*, 26(1):1 – 48, 2012.
- James Hamilton. *Time Series Analysis*. Princeton University Press, 1994.
- Gary Koop and Dimitris Korobilis. Forecasting Inflation Using Dynamic Model Averaging. *International Economic Review*, 53(3):867–886, 08 2012.
- David Kreps. *Anticipated Utility and Dynamic Choice, a 1997 Schwartz Lecture*. Frontiers of Research in Economic Theory, 1998.
- Josef Lakonishok, Andrei Shleifer, and Robert W Vishny. Contrarian investment, extrapolation, and risk. *Journal of Finance*, 49(5):1541–78, December 1994.
- Pok-Sang Lam, Stephen G. Cecchetti, and Nelson C. Mark. Asset pricing with distorted beliefs: Are equity returns too good to be true? *American Economic Review*, 90(4):787–805, September 2000.
- Martin Lettau and Sydney Ludvigson. Measuring and modelling variation in the risk-return trade-off. *Handbook of Financial Econometrics*, 1:617–690, 2010.
- Lubos Pastor and Pietro Veronesi. Learning in Financial Markets. *Annual Review of Financial Economics*, 1(1):361–381, November 2009.

- James M. Poterba and Lawrence H. Summers. Mean reversion in stock prices : Evidence and implications. *Journal of Financial Economics*, 22(1):27–59, October 1988.
- Matthew Rabin. Inference By Believers In The Law Of Small Numbers. *The Quarterly Journal of Economics*, 117(3):775–816, August 2002.
- Matthew Rabin and Dimitri Vayanos. The Gambler’s and Hot-Hand Fallacies: Theory and Applications. *Review of Economic Studies*, 77(2):730–778, 2010.
- Adrian Raftery, Miroslav Karny, and Pavel Ettl. Online Prediction Under Model Uncertainty via Dynamic Model Averaging: Application to a Cold Rolling Mill. *Technometrics*, 52(1):52–66, 2010.
- Carmen Reinhart and Kenneth Rogoff. *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press, 2009.
- G William Schwert. Why Does Stock Market Volatility Change over Time? *Journal of Finance*, 44(5):1115–53, December 1989.
- Robert Shiller. *Irrational Exuberance*. Princeton University Press, 2005.
- Robert J Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–36, June 1981.
- James H. Stock. Confidence intervals for the largest autoregressive root in u.s. macroeconomic time series. *Journal of Monetary Economics*, 28(3):435–459, December 1991.
- Allan G Timmermann. How learning in financial markets generates excess volatility and predictability in stock prices. *The Quarterly Journal of Economics*, 108(4):1135–45, November 1993.
- Amos Tversky and Daniel Kahneman. Availability: A Heuristic for Judging Frequency and Probabilities. *Cognitive Psychology*, 5:207–232, 1973.

A Consumption Rule and Asset Price

The agent maximizes:

function $c_t = P'z_t$. The evolution of the state vector satisfies: $z_t = Mz_{t-1} + C\varepsilon_t$ where

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

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

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 $\tilde{P} = P - \gamma \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \end{bmatrix}$ so $\tilde{P}'z_t = c_t - \gamma c_{t-1}$. Then the Bellman equation becomes

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

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

Now to evaluate

$$\begin{aligned}
& E_t \frac{-\Psi}{\alpha} \exp[-\alpha(\tilde{P}' z_{t+1})] \\
& E_t \frac{-\Psi}{\alpha} \exp[-\alpha(\tilde{P}' M z_t + \tilde{P}' C \varepsilon_{t+1})] \\
& \frac{-\Psi}{\alpha} \exp[-\alpha(\tilde{P}' M z_t + \frac{\alpha^2}{2} C' \tilde{P} \tilde{P}' C)] \\
& \frac{-\Psi}{\alpha} \exp(-\alpha[(\tilde{P}' M z_t - \frac{\alpha}{2} C' \tilde{P} \tilde{P}' C)])
\end{aligned}$$

So the Bellman equation becomes:

$$\frac{-\Psi}{\alpha} \exp[-\alpha(\tilde{P}' z_t)] = \frac{-1}{\alpha} \exp[-\alpha(\tilde{P}' z_t)] - \delta \frac{\Psi}{\alpha} \exp(-\alpha[(\tilde{P}' M z_t - \frac{\alpha}{2} C' \tilde{P} \tilde{P}' C)])$$

This gives

$$\Psi = 1 + \delta \Psi \exp[-\alpha((\tilde{P}' M - I) z_t - \frac{\alpha}{2} C' \tilde{P} \tilde{P}' C)] \quad (7)$$

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

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

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

From (7) we have

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

We know proceed by guessing the policy function.

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

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

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

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

Now to validate the guess for P we show that $\tilde{P}'(M - I)z_t$ is constant for this choice of

P . First note that $M - I = \bar{M} - I + NP' =$

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

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

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

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

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

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

$$\begin{aligned} & -(R-1)\left(1 - \frac{\gamma}{R}\right)Q + \left(\frac{\gamma}{R} - \gamma\right)Q \\ & \left(-R - \frac{\gamma}{R} + 1 + \gamma\right)Q + \left(\frac{\gamma}{R} - \gamma\right)Q \\ & (1-R)Q \end{aligned}$$

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

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

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

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

Now note that

$$\begin{aligned}\frac{\Phi}{R} \frac{\Phi}{R} &= \frac{\Phi}{R} - \frac{\Phi}{R} + \frac{\Phi}{R} \frac{\Phi}{R} \\ \frac{\Phi}{R} \frac{\Phi}{R} &= (I - \frac{\Phi}{R}) \frac{-\Phi}{R} + \frac{\Phi}{R} \\ (I - \frac{\Phi}{R})^{-1} \frac{\Phi}{R} \frac{\Phi}{R} &= (I - \frac{\Phi}{R})^{-1} \frac{\Phi}{R} - \frac{\Phi}{R}\end{aligned}$$

So

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

Now we can solve for Q from the first order condition

$$\begin{aligned}1 + \delta\Psi\tilde{P}'N \exp[-\alpha([\tilde{P}'M - I]z_t - \frac{\alpha}{2}C'\tilde{P}\tilde{P}'C)] &= 0 \\ 1 + \delta\Psi\tilde{P}'N \exp[-\alpha([Q(1-R) - \frac{\alpha}{2}C'\tilde{P}\tilde{P}'C)] &= 0\end{aligned}$$

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

$$\begin{aligned}1 - \delta R \exp[-\alpha([Q(1-R) - \frac{\alpha}{2}C'\tilde{P}\tilde{P}'C)] &= 0 \\ \ln(\delta R) - \alpha Q(1-R) + \frac{\alpha^2}{2}C'\tilde{P}\tilde{P}'C &= 0 \\ Q = \frac{1}{R-1} \left[\frac{-1}{\alpha} \ln(\delta R) - \frac{\alpha}{2}C'\tilde{P}\tilde{P}'C \right]\end{aligned}$$

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

$$\begin{aligned} & \tilde{P}'_5 C_5 + \tilde{P}'_7 C_7 + \tilde{P}'_8 C_8 + \tilde{P}'_{8+p+2} C_{8+p+2} \\ & \tilde{P}'_5 (p_{s,t} \sigma^s + (1 - p_{s,t}) \sigma^{ns}) + \tilde{P}'_7 \sigma^{ns} + \tilde{P}'_8 \sigma^{ns} + \tilde{P}'_{8+p+2} \sigma^s \\ & \frac{(R-1)(1-\frac{\gamma}{R})}{R} (p_{s,t} \sigma^s + (1 - p_{s,t}) \sigma^{ns}) + \left[\frac{(R-1)(1-\frac{\gamma}{R})}{R} \tilde{e}' (I - \frac{\Phi}{R})^{-1} \frac{\Phi}{R} \right] \begin{bmatrix} 0 \\ \sigma^{ns} \\ \sigma^{ns} \\ \vec{0}_{p+1} \\ \sigma^s \\ \vec{0}_{p-1} \end{bmatrix} \\ & \frac{(R-1)(1-\frac{\gamma}{R})}{R} \left(p_{s,t} (I - \frac{\Phi}{R})^{-1}_{p+5,p+5} \sigma^s + (1 - p_{s,t}) (I - \frac{\Phi}{R})^{-1}_{2,2} \sigma^{ns} + (1 - p_{s,t}) \left[(I - \frac{\Phi}{R})^{-1} \frac{\Phi}{R} \right]_{2,3} \sigma^{ns} \right) = \sigma_c \end{aligned}$$

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 $c_t =$

$$\begin{aligned} & -(R-1)(1-\frac{\gamma}{R})b_t + \frac{\gamma}{R}c_{t-1} + Q + \frac{(R-1)(1-\frac{\gamma}{R})}{R}d_t + \left[\frac{(R-1)(1-\frac{\gamma}{R})}{R} \tilde{e}' (I - \Phi)^{-1} \frac{\Phi}{R} \right] \vec{d}_t \\ & \frac{\gamma}{R}c_{t-1} - \frac{1}{R-1} \left[\frac{1}{\alpha} \ln(\delta R) + \frac{\alpha}{2} \sigma_c^2 \right] + (1 - \frac{\gamma}{R}) \frac{R-1}{R} \left[-Rb_t + d_t + E_t \sum_{s=1}^{\infty} \frac{d_{t+s}}{R^s} \right] \end{aligned}$$

Table 1: Data Moments

Table 1a: Return Predictability

<u>Moment</u>	<u>Data</u>	<u>Bootstrap Procedure</u>	
		<u>Mean</u>	<u>90% CI</u>
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2	-0.01	[-0.19 0.17]
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41	-0.13	[-0.43 0.2]
$\text{corr}(\Delta \text{Inc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34	-0.01	[-0.27 0.25]
$\text{corr}(P/E_{10,t}, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.16	-0.03	[-0.45 0.36]
$\text{corr}(\Delta \text{Inc}_t, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.23	0.01	[-0.24 0.26]
$\text{corr}(P/E_{10,t}, \Delta \text{Ind}_{t+2} + \dots + \Delta \text{Ind}_{t+5})$	-0.25	-0.02	[-0.20 0.16]
$\sigma(r_t)$	20.50%	--	--
$\sigma(\Delta \text{Inc}_t)$	2%	--	--

Table 1b: Time Varying Volatility

Kurtosis

	<u>Data</u>	<u>Confidence Bounds</u>
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

	<u>Data</u>	<u>Standard Error</u>
lag 1	0.079	0.065
lag 2	0.01	0.065
lag 3	0.47	0.065
lag 4	0.14	0.065

GARCH

	<u>Data</u>	<u>Standard Error</u>
Garch	0.61	0.09
Arch	0.29	0.07
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.

Table 2: Model Results -- Return Predictability

<u>Moment</u>	<u>Data</u>	<u>Rational Expectations</u>	<u>Parameter Learning</u>	<u>Model Learning</u>
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2	-0.02	-0.03	-0.21
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41	0.12	0.09	-0.24
$\text{corr}(\Delta \text{hnc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34	-0.04	-0.05	-0.26
$\text{corr}(P/E_{10,t}, \Delta \text{hnc}_{t+3} + \dots + \Delta \text{hnc}_{t+6})$	-0.16	0.13	0.1	-0.22
$\text{corr}(\Delta \text{hnc}_t, \Delta \text{hnc}_{t+3} + \dots + \Delta \text{hnc}_{t+6})$	-0.23	-0.03	-0.04	-0.23
$\text{corr}(P/E_{10,t}, \Delta \text{hnd}_{t+2} + \dots + \Delta \text{hnd}_{t+5})$	-0.25	0.16	0.13	-0.34
$\sigma(r_t)$	20.50%	0.13%	0.19%	3.29%
$\sigma(\Delta \text{hnc}_t)$	2%	0.1%	0.1%	2%

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 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.

Table 3: Kurtosis, Autocorrelation of Squared Returns, GARCH

<u>Kurtosis</u>						
	<u>Data</u>	<u>Confidence Bounds</u>	<u>RE</u>	<u>Learning</u>	<u>Constant Prob</u>	
r_t	4.1	[2.58 3.46]	3.00	6.37	2.98	
$P/E_{10,t}$	4.6	[1.67 3.50]	2.79	3.09	2.86	
$\% r_t > 1.96*\sigma(r_t)$	6.2%	[3% 7%]	5.0%	5.6%	5.0%	
<u>Autocorrelation of Squared Returns</u>						
	<u>Data</u>	<u>Standard Error</u>	<u>RE</u>	<u>Learning</u>	<u>Constant Prob</u>	
lag 1	0.079	0.065	0	0.2	-0.007	
lag 2	0.01	0.065	-0.002	0.16	-0.003	
lag 3	0.47	0.065	0.001	0.13	-0.002	
lag 4	0.14	0.065	0.004	0.1	-0.006	
<u>GARCH</u>						
	<u>Data</u>	<u>Standard Error</u>	<u>RE</u>	<u>Learning</u>	<u>Constant Prob</u>	
Garch	0.61	0.09	0	0.56	0	
Arch	0.29	0.07	0	0.25	0	
<u>p-value Engle test</u>	<u>0.048</u>					

Note: This table reports kurtosis, autocorrelation of squared returns, and GARCH estimate for the excess returns data and the simulated model data. Constant Prob refers to a model with constant probability on the non-stationary model.

Table 4: Robustness

<u>Moment</u>	<u>Data</u>	<u>AR</u>	<u>2</u>	<u>4</u>	<u>6</u>	<u>8</u>
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2		-0.18	-0.21	-0.18	-0.19
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41		-0.2	-0.24	-0.18	-0.22
$\text{corr}(\Delta \text{Inc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34		-0.22	-0.26	-0.23	-0.26
$\text{corr}(P/E_{10,t}, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.16		-0.18	-0.22	-0.16	-0.22
$\text{corr}(\Delta \text{Inc}_t, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.23		-0.2	-0.23	-0.22	-0.26
$\text{corr}(P/E_{10,t}, \Delta \text{Ind}_{t+2} + \dots + \Delta \text{Ind}_{t+5})$	-0.25		-0.27	-0.34	-0.25	-0.32
$\sigma(r_t)$	20.50%		3.2%	3.2%	2.6%	2.7%
$\sigma(\Delta \text{Inc}_t)$	2%		2.0%	2%	1.7%	1.8%
<u>Moment</u>	<u>Data</u>	<u>Gain</u>	<u>0.05</u>	<u>0.075</u>	<u>0.1</u>	
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2		-0.21	-0.21	-0.22	
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41		-0.19	-0.24	-0.27	
$\text{corr}(\Delta \text{Inc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34		-0.25	-0.26	-0.27	
$\text{corr}(P/E_{10,t}, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.16		-0.17	-0.22	-0.24	
$\text{corr}(\Delta \text{Inc}_t, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.23		-0.22	-0.23	-0.24	
$\text{corr}(P/E_{10,t}, \Delta \text{Ind}_{t+2} + \dots + \Delta \text{Ind}_{t+5})$	-0.25		-0.27	-0.34	-0.37	
$\sigma(r_t)$	20.50%		2.70%	3.2%	3.4%	
$\sigma(\Delta \text{Inc}_t)$	2%		1.7%	2%	2.1%	
<u>Moment</u>	<u>Data</u>	<u>R</u>	<u>1.0025</u>	<u>1.005</u>	<u>1.01</u>	<u>1.02</u>
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2		-0.21	-0.21	-0.21	-0.21
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41		-0.24	-0.24	-0.24	-0.23
$\text{corr}(\Delta \text{Inc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34		-0.26	-0.26	-0.26	-0.26
$\text{corr}(P/E_{10,t}, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.16		-0.22	-0.21	-0.21	-0.2
$\text{corr}(\Delta \text{Inc}_t, \Delta \text{Inc}_{t+3} + \dots + \Delta \text{Inc}_{t+6})$	-0.23		-0.23	-0.22	-0.23	-0.23
$\text{corr}(P/E_{10,t}, \Delta \text{Ind}_{t+2} + \dots + \Delta \text{Ind}_{t+5})$	-0.25		-0.34	-0.34	-0.35	-0.37
$\sigma(r_t)$	20.50%		3.2%	3.2%	3.3%	3.6%
$\sigma(\Delta \text{Inc}_t)$	2%		2%	2.0%	2.2%	2.4%

Note: This table reports the results from table 2 but varying some of the parameter choices.

Table 5: Impact of Model Learning

5a Price Change Decomposition			
	$\text{cov}(\Delta p_{t,1}^L, dp_1)/\text{var}(\Delta p_{t,1}^L)$	$\text{cov}(\Delta p_{t,2}^L, dp_2)/\text{var}(\Delta p_{t,1}^L)$	
	69.0%	31.0%	
5b Moments at Median Probability			
Moment	Data	Model	Constant Probability
$\text{corr}(r_t, r_{t+2} + \dots + r_{t+5})$	-0.2	-0.21	-0.26
$\text{corr}(P/E_{10,t}, r_{t+2} + \dots + r_{t+5})$	-0.41	-0.24	-0.18
$\text{corr}(\Delta \text{hnc}_t, r_{t+2} + \dots + r_{t+4})$	-0.34	-0.26	-0.3
$\text{corr}(P/E_{10,t}, \Delta \text{hnc}_{t+3} + \dots + \Delta \text{hnc}_{t+6})$	-0.16	-0.22	-0.16
$\text{corr}(\Delta \text{hnc}_t, \Delta \text{hnc}_{t+3} + \dots + \Delta \text{hnc}_{t+6})$	-0.23	-0.23	-0.26
$\text{corr}(P/E_{10,t}, \Delta \text{hnd}_{t+2} + \dots + \Delta \text{hnd}_{t+5})$	-0.25	-0.34	-0.26
$\sigma(r_t)$	20.50%	3.2%	2.2%
$\sigma(\Delta \text{hnc}_t)$	2%	2%	1.5%

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 $dp_1 = p_{s,t}^S * \Delta p_{t,1}^S + (1 - p_{s,t}^S) * \Delta p_{t,1}^{NS}$ and the model belief change times the difference in the model predictions for price $dp_2 = (p_{s,t}^S - p_{s,t-1}^S) * (p_{t-1}^S - p_{t-1}^{NS})$. Panel B examines the model predictions if the probability was constant at its median (across time and trials).