What Information Drives Asset Prices?

Anisha Ghosh
Carnegie-Mellon University
anishagh@andrew.cmu.edu

George M. Constantinides
University of Chicago and NBER
gmc@chicagobooth.edu

Abstract

The market price-dividend ratio is highly correlated with several macroeconomic variables but not with aggregate consumption growth. We incorporate this observation in an exchange economy with learning about the economic regime from consumption history *and* a latent signal. The estimated model rationalizes the moments of consumption and dividend growth, stock market return, price-dividend ratio, and risk free rate and performs well in predictive regressions while a nested model with learning from consumption history alone does not. The intuition is that the beliefs process has high persistence and low variance because beliefs depend on consumption growth *and* the signal. The model fit remains largely intact when we replace the latent signal with the innovation in the CPI, hourly earnings, or a combination of macroeconomic variables highly correlated with the price-dividend ratio. The results highlight the informational role of macroeconomic variables and suggests that just one macroeconomic variable (CPI or earnings per hour), along with consumption growth, goes a long way towards proxying for the investors' relevant information set.

March 14, 2017

Keywords: macroeconomic news; CPI; hourly earnings; learning models; asset pricing

JEL classification: E3, G12, G14

We thank Rui Albuquerque, Daniel Andrei, Bruce Grundy, Lars Hansen, Zhiguo He, John Heaton, Michael Johannes, Christian Julliard, Anil Kashyap, Bryan Kelly, Ralph Koijen, Rajnish Mehra, Stefan Nagel, Lubos Pastor, Anna Pavlova, Amir Sufi, Allan Timmermann, Stijn Van Nieuwerburgh, Pietro Veronesi, Michael Weber, and other seminar participants at the University of Chicago, the London School of Economics, and McGill University for their helpful advice and feedback and we thank Yiran Fan for outstanding research assistance. Constantinides received financial support from the Center for Research in Security Prices, the University of Chicago and as trustee/director of the DFA group of funds, SW7 Holdings, Cook County Illinois Investment Policy Committee, and as member of the advisory board of FTSE-Russell.

1 Introduction

An overload of worldwide macroeconomic, business, and political news inundates investors. Little is known as to how investors cope with this vast amount of information, and, in particular, which subset of information they pay attention to. In the macro-finance literature, researchers typically model investors as focusing on the histories of a limited number of macroeconomic variables and applying a filter to rationally extract relevant information about the economy. Prominent examples include regime-switching models, where the regimes are latent and investors filter their beliefs about the current economic regime (hereafter referred to as the 'state') from the history of aggregate consumption alone. These models fare poorly in explaining several features of stock market data, including the high average level of the equity premium, the low level of the risk free rate, the high variability of the price-dividend ratio, and the low predictability of consumption growth by the price-dividend ratio.

Recent literature argues that a more challenging high-dimensional learning problem confronts investors where they need to learn not only about the current state but also about the true underlying model and its parameters and that such a learning problem plays an important role in enhancing the empirical performance of these models. While parameter and model uncertainty are undoubtedly important issues, in this paper we question the other central assumption of learning models, namely that investors form their beliefs about the economic state from consumption history alone. We show that expanding the information set of investors plays a central role in improving the performance of this class of models.

We highlight the advantages of our framework over standard learning models where investors are assumed to learn from the consumption history alone; long run risks models; and models with parameter and model uncertainty in addition to state uncertainty. We also discuss the importance of identifying the information set of investors for other applications, in addition to the application considered in this paper.

Reliable identification of the information set of investors is also important for a host of other applications. For example, conditional moments (e.g., means and variances) of returns are often modeled as projections onto a set of predetermined conditioning variables. Econometric considerations necessitate the choice of a small set of variables, introducing an element of arbitrariness in the modeling of expectations and may produce misleading estimates (see, Hansen and Richard (1987)). Therefore, the potential of summarizing the investors' information set with a small-dimensional set of variables is quite useful. Also recent research highlights how investors' subjective beliefs, or belief distortions, can be extracted from observed asset prices via the consumption Euler equations (see, Hansen, Hansen, and Mykland (2016)). Once again the econometric feasibility of this extraction crucially relies on being able to characterize the conditioning set underlying the Euler equations with a small number of variables. Our paper contributes towards identifying the contents of investors' information set. Our paper suggests that just one macroeconomic variable (CPI or earnings per hour), along with consumption growth, goes a long way towards proxying for the investors' relevant information set.

We find that two broad categories of publicly available macroeconomic information are the most highly correlated with the market-wide price-dividend ratio. The first category consists of price levels, including the Consumer Price index for all Urban Consumers (CPI-U) and the Producer Price Index (PPI). The second category consists of labor market variables, including the average hourly earnings, average hours of production, and numbers of employees in private non-farm payrolls in different sectors. These are the two classes of macro variables that, according to FactSet, Bloomberg users pay the most attention to (see, Ai and Bansal (2016)). On the other hand, contrary to the implications of learning models where investors are assumed to learn from the consumption history alone, the price-dividend ratio has negligible correlation with the contemporaneous consumption growth as well as a weighted average of current and lagged consumption growth rates.

These features of the data suggest that the poor empirical performance of standard learning models may be driven, at least in part, by the stringent assumption of learning from the consumption history alone. Therefore, in this paper, we consider a learning process where investors learn not only from the history of consumption but also from a latent signal. To shed light on the potential sources of the signal, we set the signal equal to judiciously chosen macroeconomic variables: innovation in the CPI, hourly earnings, or a combination of macroeconomic variables highly correlated with the price-dividend ratio.

We consider a representative-investor real exchange economy as in Lucas (1978). We isolate the role of broadening the investors' information set by abstracting from model and parameter uncertainty and assuming that the investors know the economic model and its parameters. The aggregate consumption and stock market dividend processes have different means in two latent economic regimes. Each period the investors rationally update the probability that the economy is in the first regime—their beliefs—by observing the updated history of aggregate consumption and an unspecified signal with innovations orthogonal to those of aggregate consumption and dividend growth. The signal is deliberately left unspecified at first in order to demonstrate that there exist potential signals over and above aggregate consumption, the commonly assumed signal in the literature. The investors are assumed to have recursive preferences as in Epstein and Zin (1989) and Weil (1990). We numerically solve for the equilibrium stock market price-dividend ratio, the conditional mean, volatility, and Sharpe ratio of the market return, and the risk free rate as functions of the investors' beliefs. We estimate the model parameters using the Generalized Method of Moments (GMM). The model provides a good fit to the sample moments of consumption growth and dividend growth, market return, market-wide price-dividend ratio, and risk free rate.

In contrast we show that an alternative nested model in which the investors learn from the consumption history alone fails along a number of dimensions: it generates an average equity premium of only 0.5% compared to its sample counterpart of 5.6% over the entire available sample period from 1929 to 2013, thereby failing to account for the equity premium puzzle; it implies less than 1% volatility of the market-wide price-dividend ratio which is two orders of magnitude smaller than the 45.5% value observed in the data, thereby failing to explain the excess volatility puzzle; it implies a high correlation of over 50% between consumption growth and the innovation in the risk free rate, in stark contrast to its sample counterpart of -1%; and it implies a high correlation, of over 85% in absolute terms, between consumption growth and the price-dividend ratio, while its sample counterpart is only 13.5%.

We provide further evidence in support of the main model by setting the, so far, latent signal equal to one of the macroeconomic variables that have high correlation with the market-wide price-dividend ratio. We find that the main model retains its good empirical fit when the signal is set equal to one of these macroeconomic variables compared to when the signal is left unspecified. In fact, not only does the model fare much better than the alternative model at matching the unconditional moments of asset prices and returns, it also generates a time series of the price-dividend ratio that lines up much more closely with the historical time series compared to the alternative model. In particular, in the main model, the implied time series of the price-dividend ratio has correlation 0.69 with the historical time series when the signal is set equal to the innovation in the CPI and correlation 0.70 with the historical time series when the signal is set equal to the innovation in the average hourly earnings of production on private nonfarm payrolls in the manufacturing sector. In contrast the alternative model generates a time series of the price-dividend ratio that has correlation close to zero with the historical series. These results provide strong support for the economic mechanism highlighted in the main model. However, we stress that the two particular macro variables considered here are intended only as an illustration of our central insight regarding the pivotal role of the information set and that we do not claim that the investors exclusively learn from these variables, in addition to consumption growth.

The success of the main model is a particularly strong result because it derives solely from expanding the information set of investors to include variables in addition to the history of consumption growth. The question, therefore, arises as to what drives the superior performance of the main model. The crux of the intuition lies in the stark contrast between the properties of the state and beliefs processes in the two models. In the main model the state process and the beliefs processes have high persistence and low variance while in the alternative model these processes have much lower persistence and high variance. In the main model the high persistence of the state process imparts high persistence in the investors' beliefs and renders the price-dividend ratio highly persistent, consistent with the data. Also, the current beliefs are highly informative about the future and, therefore, the price-dividend ratio is highly responsive to changes in beliefs.

This feature of the learning process explains why the price-dividend ratio is highly variable in this model, consistent with the data.

In the alternative model, on the other hand, the current beliefs provide little information about the future and this leads to the counterfactual predictions that the price-dividend ratio is essentially constant and the volatility of the market return equals the volatility of aggregate dividend growth. If the state process and, therefore, the beliefs process were highly persistent, then the mean consumption growth would be very different in the two states, thereby imparting a counterfactually high volatility in the consumption growth process.

The high persistence in the beliefs process in the main model, combined with the preference for early resolution of uncertainty, yields a high equity premium and low risk free rate, consistent with the data. In contrast, in the alternative model the low persistence of the beliefs process yields a low equity premium.

Finally, consistent with the data, the main model generates strong time-variation in the conditional mean and variance of the market return. Perhaps more impressive is the observation that it does so without relying on countercyclical heteroscedasticity of the consumption growth rate or the additional signal (the volatilities of consumption growth and the signal are assumed to be equal in the two states)—a phenomenon for which there is limited empirical evidence. Instead, the model generates time-variation in the conditional moments of the market return from the heteroscedasticity of the beliefs process. Other competing theories, such as the long-run risks model of Bansal and Yaron (2004), rely critically on heteroscedasticity of consumption growth in order to generate time-variation in the expected equity premium.

We further assess the empirical plausibility of the main and alternative models with predictive regressions for consumption growth, dividend growth and market return. For the main model, the predictive regressions yield slope coefficients and R^2 that are consistent with the data. In particular, the model generates the low predictive power of the price-dividend ratio for the one-year ahead as well as long horizon consumption growth rates, consistent with the data. This also stands in stark contrast with the long-run risks model that implies an order of magnitude higher R^2 in forecasting regressions of the one-year and three-year ahead consumption growth rates and two orders of magnitude higher R^2 in a forecasting regression of the five-year ahead consumption growth rate on the price-dividend ratio. The superior performance of our main model is rendered possible because it does not rely on the persistence of consumption growth to generate key asset pricing results, relying instead on the persistence of other macro variables in the information set of investors. For the alternative model, on the other hand, all regression coefficients have the wrong sign, further demonstrating that the alternative model does not fit the data.

The paper draws on several strands of the literature. It draws on the extant literature that focuses on learning about latent states or a single parameter as in Ai (2010), Ai and Bansal (2016), Bansal and Shaliastovich (2011), Croce, Lettau, and Ludvigson (2015), Drechsler (2013), Li (2005), Nieuwerburgh and Veldkamp (2006)), and Veronesi (2000). Pastor and Veronesi (2009) review learning models. Several papers highlight learning from macroeconomic variables. Andrei, Hasler, and Jeanneretz (2016) study learning from macroeconomic variables about the persistence of expected economic growth rather than about its level. Basak (2005) surveys the literature on learning by agents with heterogeneous beliefs. D'Acundo, Hoang, and Weber (2016) study announcements of future increases in consumption taxes to generate inflation expectations and accelerate consumption expenditure. David and Veronesi (2013) study learning from inflation shocks and address the money illusion and the covariance between stock and bond returns. Gurkaynak, Sack, and Swanson (2005) argue that investors adjust their expectations of the long-run level of inflation in response to macroeconomic and monetary policy surprises. Finally, Hall (2014) examines learning from unemployment and addresses discount rates.

The paper also draws on the empirical evidence in Albuquerque, Eichenbaum, and Rebelo (2012), Duffee (2005), Greenwald, Lettau, and Ludvigson (2015), and Lettau and Ludvigson (2013) that the correlation between the stock market return and aggregate consumption growth is weak.

The paper relates to models of ambiguity by Collin-Dufresne, Johannes, and Lochstoer (2016), Epstein and Schneider (2003), Hansen and Sargent (2001), Johannes, Lochstoer, and Mou (2016), Klibanoff, Marinacci, and Mukerji (2005), and Maccheroni, Marinacci, and Rustichini (2006). Specifically, Collin-Dufresne, Johannes, and Lochstoer (2016) and Johannes, Lochstoer, and Mou (2016) argue that introducing a high-dimensional learning problem where the investors need to learn not only about the latent state(s), but also about the true underlying model and its parameters, plays an important role in enhancing the empirical performance of these models. Johannes, Lochstoer, and Mou (2016) assume that the investors learn either from consumption history alone or from a combination of consumption and GDP histories. In all cases, their model overstates the mean risk free rate by a factor of two and understates the volatility of the market-wide price-dividend ratio. Our paper, on the other hand, abstracts from parameter and model uncertainty while expanding the information set of the investors to accommodate learning from judiciously chosen macroeconomic variables, in addition to consumption growth. Our results suggest that this simple modification to the investors' information set greatly improves the empirical performance of the model.

Finally, the paper draws on the long-run risks literature by Bansal and Yaron (2004) and Hansen, Heaton, and Li (2008) who argue for the presence of a small predictable component in aggregate consumption and dividend growth. Note that our model also implies the presence of a predictable component in aggregate consumption and dividend growth. Beeler and Campbell (2012) and Constantinides and Ghosh

(2011) point out that long-run risks models imply too high autocorrelations of the aggregate consumption growth rate and, therefore, excessive predictability of consumption growth by the price-dividend ratio. By contrast, our main model implies an autocorrelation function for consumption growth that is close to zero at all lags, consistent with the data. This is rendered possible because, in our model, although consumption growth is not very persistent, the additional signal is and this enables the investors to learn about the current economic state as well as forecast the future based on the history of the signal.

The remainder of the paper is organized as follows. In Section 2 we present the main model and derive its pricing implications. We discuss the data in Section 3. The empirical methodology and estimation results are presented in Section 4, along with a comparison between the main model and the alternative one. In Section 5 we launch a systematic investigation of the sources of macroeconomic information that drive aggregate prices. In Section 6 we provide further evidence in support of the main model by setting the signal equal to macroeconomic variables most highly correlated with the market-wide price-dividend ratio. In Section 7 we discuss the economic intuition underlying the results. In Section 8 we present predictive regressions for consumption growth, dividend growth, and market return. We conclude in Section 9.

2 The Model and Solution

2.1 Model Description

We consider a representative-investor real exchange economy. The aggregate consumption and dividend processes are exogenously specified with their parameters estimated from the data. We assume that the investors know the economic model and its parameters. We model the aggregate consumption and dividend growth rates as having different means across two latent regimes, $s_t = 1, 2$, as

$$\Delta c_{t+1} = \mu_{c,s_{t+1}} + \sigma_c \varepsilon_{c,t+1} \tag{1}$$

and

$$\Delta d_{t+1} = \mu_{d,s_{t+1}} + \sigma_d \varepsilon_{d,t+1}, \tag{2}$$

where c_t and d_t are the logarithms of aggregate consumption and stock market dividend, respectively, in period *t*. The means of consumption and dividend growth rates, $\mu_{c,s_{t+1}}$ and $\mu_{d,s_{t+1}}$, respectively, are generally different in the two regimes and are driven by a two-valued scalar state variable, $s_t = 1, 2$, that denotes the economic regime. The volatilities of consumption growth, σ_c, and dividend growth, σ_d, are assumed to be constant across the two regimes, motivated by the observation that strong countercyclical heteroscedasticity is not a feature of macroeconomic data.[1] The shocks $\varepsilon_{c,t+1}$ and $\varepsilon_{d,t+1}$ are *i.i.d.* standard normal with cross-correlation ρ.

We assume that s_t is an exogenous Markov process with transition probability matrix

$$\begin{pmatrix} \pi_1 & 1-\pi_2 \\ 1-\pi_1 & \pi_2 \end{pmatrix}, \tag{3}$$

where $\pi_i \equiv prob\left(s_t = i \mid s_{t-1} = i\right)$ and $0 < \pi_i < 1$ for $i = 1, 2$. The parameters are estimated from the data. The unconditional probability of $s_t = 1$ is $\left(1-\pi_2\right)/\left(2-\pi_1-\pi_2\right)$ and its expected duration is $\left(1-\pi_1\right)^{-1}$ years. In the empirical section we interpret the state $s_t = 1$ as the state of economic expansion and the state $s_t = 2$ as the state of economic recession or slow economic growth.

We assume that the investors do not observe the regime at time t but learn from a history of signals, $\mathbb{F}\left(t\right)$. We assume that the investors' history of signals is $\mathbb{F}\left(t\right) = \left\{c_\tau, x_\tau\right\}_{\tau=-\infty,\ldots,t}$, where c_t is consumption over $\left[t-1, t\right]$ and x_t is a scalar reflecting additional variables over $\left[t-1, t\right]$ that investors rely on to form beliefs about the economic regime at time *t*.

Most of the existing literature typically assumes that the investors' history of signals is simply the consumption history, $\mathbb{F}\left(t\right) = \left\{c_\tau\right\}_{\tau=-\infty,\ldots,t}$. These models fare poorly in explaining the high observed level of the equity premium, the low level of the risk free rate, and the excess volatility of asset prices relative to fundamentals. The crux of our model lies in allowing investors to form their beliefs not only from the history of consumption but also from other publicly available macroeconomic variables—a natural modeling choice given the multitude of publicly available information—and exploring the resulting improvement in the empirical performance of the model.

The case where the investors' history of signals is the consumption and dividend histories, $\mathbb{F}\left(t\right) = \left\{c_\tau, d_\tau\right\}_{\tau=-\infty,\ldots,t}$, leads to as poor results as the case where the investors' history of signals is the consumption history alone for two reasons. First, dividend growth is not very persistent and, in fact, it has an even smaller first-order autocorrelation coefficient than the consumption growth rate. As we shall see

[1] See Campbell, Pflueger, and Viceira (2015)). In addition to there being limited empirical evidence supporting heteroscedasticity of macro data, our modeling choice of constant volatility of consumption and dividend growth in the two states is also made to highlight that our key results obtain even in the absence such heteroscedasticity. Allowing for heteroscedasticity is likely to further enhance the performance of the model as it provides an additional source of risk premia.

later on, it is important that the signal be persistent. Moreover, similar to the conclusions obtained for the consumption growth rate, a regression of the market-wide price-dividend ratio on the contemporaneous dividend growth produces a statistically insignificant slope coefficient and $R^2 = 1\%$ only; and regressions of the price-dividend ratio on an exponentially-weighted moving average of the contemporaneous and lagged dividend growth rates also produce small $R^2 = 3.4\%$ and $R^2 = 11.4\%$, respectively, when five and ten lags of the dividend growth are included in the regressor. As we show in Section 5, these are substantially smaller than those obtained with some other macro variables such as inflation and labor market variables.

We assume that the scalar signal x_{t+1} has constant volatility but different mean across the two regimes as

$$x_{t+1} = \mu_{x,s_{t+1}} + \sigma\varepsilon_{x,t+1}, \tag{4}$$

where $\varepsilon_{x,t+1}$ is *i.i.d.* standard normal and orthogonal to $\varepsilon_{c,t+1}$ and $\varepsilon_{d,t+1}$. Without loss of generality we set $\mu_{x,1} = 0$ and $\sigma = 1$. Furthermore we simplify the notation by denoting $\mu_{x,2}$ by μ.

We denote the mean of the vector $u_{t+1} = \left(\Delta c_{t+1}, \Delta d_{t+1}, x_{t+1}\right)$, conditional on $s_{t+1} = i, i = 1,2$, as μ_i and the variance-covariance matrix as Σ, where

$$\mu_i = \begin{bmatrix} \mu_{c,i} \\ \mu_{d,i} \\ \mu_{x,i} \end{bmatrix} \tag{5}$$

and

$$\Sigma = \begin{pmatrix} \sigma_c^2 & \rho\sigma_c\sigma_d & 0 \\ \rho\sigma_c\sigma_d & \sigma_d^2 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{6}$$

The investors assign probability p_t that the economy is in the first regime at date *t*:

$$p_t \equiv prob\left(s_t = 1 \,|\, \mathbb{F}(t)\right). \tag{7}$$

The joint probability density function of u_{t+1}, conditional on the information available at time *t*, is

$$
\begin{aligned}
&g\left(u_{t+1} \mid \mathbb{F}(t)\right) \\
&=\frac{f\left(p_t\right)}{(2\pi)^{3/2}|\Sigma|}\exp\left(-\frac{1}{2}\left(u_{t+1}-\mu_1\right)^T \Sigma^{-1}\left(u_{t+1}-\mu_1\right)\right)+\frac{1-f\left(p_t\right)}{(2\pi)^{3/2}|\Sigma|}\exp\left(-\frac{1}{2}\left(u_{t+1}-\mu_2\right)^T \Sigma^{-1}\left(u_{t+1}-\mu_2\right)\right),
\end{aligned} \tag{8}
$$

where $f\left(p_t\right)$ is the probability that the investors assign that the state in the next period *t+1* is the expansion state, conditional on p_t:

$$
\begin{aligned}
&f\left(p_t\right) \equiv prob\left(s_{t+1}=1 \mid \mathbb{F}(t)\right) \\
&= prob\left(s_{t+1}=1 \mid s_t=1\right) \times prob\left(s_t=1 \mid \mathbb{F}(t)\right)+prob\left(s_{t+1}=1 \mid s_t=2\right) \times prob\left(s_t=2 \mid \mathbb{F}(t)\right) \\
&=\pi_1 p_t+\left(1-\pi_2\right)\left(1-p_t\right) \\
&=1-\pi_2+\left(\pi_1+\pi_2-1\right) p_t .
\end{aligned} \tag{9}
$$

Conditional on $\left(\Delta c_{t+1}, x_{t+1}, \mathbb{F}(t)\right)$, the value of p_{t+1} is updated via Bayes' rule as

$$
p_{t+1} \mid \left(\Delta c_{t+1}, x_{t+1}, \mathbb{F}(t)\right)=\frac{f\left(p_t\right) e^{-\frac{\left(\Delta c_{t+1}-\mu_{c,1}\right)^2}{2\sigma_c^2}-\frac{x_{t+1}^2}{2}}}{2\pi\sigma_c g\left(\Delta c_{t+1}, x_{t+1} \mid \mathbb{F}(t)\right)}, \tag{10}
$$

where $g\left(\Delta c_{t+1}, x_{t+1} \mid \mathbb{F}(t)\right)$ is the marginal density of $\left(\Delta c_{t+1}, x_{t+1}\right)$, conditional on the information available at time *t*. Therefore the conditional expectation of p_{t+1} is $f\left(p_t\right)$:

$$
\begin{aligned}
&E\left[p_{t+1} \mid \mathbb{F}(t)\right] \\
&=\iint\left\{p_{t+1} \mid \left(\Delta c_{t+1}, x_{t+1}, \mathbb{F}(t)\right)\right\} g\left(\Delta c_{t+1}, x_{t+1} \mid \mathbb{F}(t)\right) d\Delta c_{t+1} dx_{t+1} \\
&=\iint \frac{f\left(p_t\right) e^{-\frac{\left(\Delta c_{t+1}-\mu_{c,1}\right)^2}{2\sigma_c^2}-\frac{x_{t+1}^2}{2}}}{2\pi\sigma_{c,1} g\left(\Delta c_{t+1}, x_{t+1} \mid \mathbb{F}(t)\right)} g\left(\Delta c_{t+1}, x_{t+1} \mid \mathbb{F}(t)\right) d\Delta c_{t+1} dx_{t+1} \\
&=f\left(p_t\right) \iint \frac{e^{-\frac{\left(\Delta c_{t+1}-\mu_{c,1}\right)^2}{2\sigma_c^2}-\frac{x_{t+1}^2}{2}}}{2\pi\sigma_c} d\Delta c_{t+1} dx_{t+1} \\
&=f\left(p_t\right)
\end{aligned} \tag{11}
$$

and the unconditional mean of p_t is

$$
\bar{p}=\frac{1-\pi_2}{2-\pi_1-\pi_2}. \tag{12}
$$

Our model nests the version typically assumed in the literature where investors learn from the consumption history alone: setting $\mu = 0$ results in the latter specification. We estimate the main model, featuring learning from both consumption growth and signal histories, and the alternative model, featuring learning from the consumption history alone. Note that the main model has just one additional parameter, μ, compared to the alternative model.

The investors have recursive preferences as in Epstein and Zin (1989) and Weil (1990)

$$U_t = \left\{ (1-\delta)(C_t)^{1-1/\psi} + \delta \left(E_t \left[(U_{t+1})^{1-\gamma} \right] \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{1/(1-1/\psi)}, \tag{13}$$

where δ is the subjective discount factor, γ is the RRA coefficient, ψ is the EIS, and $\theta \equiv \dfrac{1-\gamma}{1-1/\psi}$. As shown in Epstein and Zin (1989), the stochastic discount factor (SDF) is

$$SDF_{t+1} = \delta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} R_{c,t+1}^{\theta-1}, \tag{14}$$

where $R_{c,t+1}$ is the return on the wealth portfolio, a portfolio that pays dividend each period equal to aggregate consumption.

2.2 Model Solution

We solve the model numerically through value function iteration. The highly nonlinear solutions for equilibrium asset prices—including the market-wide price-dividend ratio and the conditional mean and volatility of the market return obtained in Section 7—highlight the importance of eschewing the Campbell and Shiller (1988) log-linearization or any other form of approximation in solving the model.

We numerically calculate the wealth-consumption ratio, $z(p_t) \equiv P_t / C_t$, as a function of p_t as follows. The Euler equation for the consumption claim is

$$P_{c,t} = E \left[\delta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} R_{c,t+1}^{\theta-1} \left(P_{c,t+1} + C_{t+1} \right) | \, p_t \right]. \tag{15}$$

We divide both sides of equation (15) by C_t and write

$$\frac{P_{c,t}}{C_t} = E\left[\delta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}} \left(\frac{P_{c,t+1}+C_{t+1}}{P_{c,t}}\right)^{\theta-1} \left(\frac{P_{c,t+1}}{C_{t+1}}+1\right)\left(\frac{C_{t+1}}{C_t}\right) | \, p_t\right]$$

or

$$\left(\frac{P_{c,t}}{C_t}\right)^\theta = E\left[\delta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}+\theta} \left(\frac{P_{c,t+1}}{C_{t+1}}+1\right)^\theta | \, p_t\right] \tag{16}$$

$$= \iint \delta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}+\theta} \left(\frac{P_{c,t+1}}{C_{t+1}}+1\right)^\theta g\left(\Delta c_{t+1}, x_{t+1} \,|\, \mathbb{F}(t)\right) d\left(\Delta c_{t+1}\right) dx_{t+1}.$$

We solve for the function $z(p)$ by iteration as follows:

(1) We initially set $z(p) = 30 \times p$.

(2) We set $P_{c,t+1} / C_{t+1} = z(p_{t+1})$.

(3) For each value of p_t, we draw $(\Delta c_{t+1}, x_{t+1})$ from the distribution $g\left(\Delta c_{t+1}, x_{t+1} \,|\, \mathbb{F}(t)\right)$ N times and calculate the left-hand side of equation (16), $P_{c,t} / C_t$ as the conditional expectation on the right-hand side of equation (16).

(4) We consider a grid of values of p_t and repeat step (3) to obtain an updated function for the price-consumption ratio, $z(p) = P_{c,t} / C_t$ as a function of *p*.

(5) We repeat steps (2)-(4) until the function $z(p)$ converges.

The return on the consumption claim is

$$R_{c,t+1} = \frac{P_{c,t+1}+C_{t+1}}{P_{c,t}} = \left(\frac{P_{c,t+1}}{C_{t+1}}+1\right)\frac{C_{t+1}}{C_t}\frac{C_t}{P_{c,t}} = \frac{z(p_{t+1})+1}{z(p_t)}\frac{C_{t+1}}{C_t} \tag{17}$$

and is a known function of $p_t, p_{t+1},$ and C_{t+1} / C_t. Therefore the SDF in equation (14) becomes a known function of $p_t, p_{t+1},$ and C_{t+1} / C_t.

The risk free rate is obtained by numerically integrating the Euler equation

$$\frac{1}{R_{f,t}} = \iint \delta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\theta}{\psi}+\theta-1} \left(\frac{z(p_{t+1})+1}{z(p_t)}\right)^{\theta-1} g\left(\Delta c_{t+1}, x_{t+1} \,|\, \mathbb{F}(t)\right) d\left(\Delta c_{t+1}\right) dx_{t+1}. \tag{18}$$

We calculate the unconditional mean, variance, and auto-correlation of the risk free rate and the correlation of the risk free rate and consumption growth via simulation.

We follow a similar procedure to numerically calculate the equilibrium price-dividend ratio, $z_M\left(p_t\right)$ and, therefore, the unconditional mean, variance, and auto-correlation of the price-dividend ratio, the correlation of the price-dividend ratio and consumption growth, and the unconditional mean, variance, and auto-correlation of the market return. Finally we obtain the conditional mean, volatility, and Sharpe ratio of the market return as a function of p_t. We repeat all of the above numerical procedures in the alternative model where the investors' information set includes only the consumption history.

3 Description of the Data

For our main empirical results, we use US annual data over the entire available sample period from 1929 to 2013. The starting date of 1929 is dictated by the availability of annual data on per capita consumption of non-durables and services. The asset menu consists of the market return and the risk free rate. Our market proxy is the Center for Research in Security Prices (CRSP) value-weighted index of all stocks on the NYSE, AMEX, and NASDAQ. The proxy for the real risk free rate is obtained as follows: the quarterly nominal yield on three-month Treasury bills is deflated using the realized growth in the Consumer Price Index to obtain the ex-post real three-month Treasury-bill rate. The ex-ante quarterly risk free rate is then obtained as the fitted value from the regression of the ex-post three-month Treasury-bill rate on the three-month nominal yield and the realized growth in the Consumer Price Index over the previous year. The ex-ante quarterly risk free rate at the beginning of the year is annualized to obtain the ex-ante annual risk free rate. The equity premium is the difference in average log returns on the market and the risk free rate.

Also used in the empirical analysis are the price-dividend ratio and dividend growth rate of the market portfolio. These two time-series are computed using the monthly returns with and without dividends on the market portfolio obtained from the CRSP files. The monthly dividend payments within a year are added to obtain the annual aggregate dividend, i.e., we do not reinvest dividends either in T-bills or in the stock market. The annual price-dividend ratio is computed as the ratio of the price at the end of each calendar year to the annual aggregate dividends paid out during that year.

The consumption data consists of the per capita personal consumption expenditure on nondurable goods and services obtained from the Bureau of Economic Analysis. All nominal quantities are converted to real, using the Consumer Price Index (CPI).

We obtain panel data over 1964-2011 on 106 macroeconomic variables from Sydney Ludvigson's web site, based on the *Global Insights Basic Economics Database* and *The Conference Board's Indicators Database*. The variables cover six broad categories of macroeconomic data: output, labor market, housing sector, orders and inventories, money and credit, and price levels. We refer the reader to Ludvigson's website for a detailed description of these variables. Many of the macroeconomic time series are revised ex post. Gilbert (2016) argues that the market understands the subsequent revisions but revisions still matter.

4 Estimation Methodology and Preliminary Results

4.1 Estimation Methodology

The main model has thirteen parameters: two parameters of the regime transition matrix (π_1 and π_2); one parameter of the signal distribution (μ); seven parameters of the time-series processes of aggregate consumption and dividend growth ($\mu_{c,1}, \mu_{c,2}, \mu_{d,1}, \mu_{d,1}, \sigma_c, \sigma_d$); and three preference parameters ($\delta, \gamma,$ and ψ). We estimate the parameters using GMM to match the following seventeen sample moments: the unconditional mean, variance, and first-order autocorrelation of consumption growth, dividend growth, market return, market-wide price-dividend ratio, and risk free rate; the correlation between consumption and dividend growth rates; and the correlation between the consumption growth rate and the price-dividend ratio. Therefore, we have an over-identified system of seventeen moment restrictions and thirteen parameters. The weighting matrix used in the estimation is a diagonal matrix with unit entries corresponding to all the moments except for weights of 100 on the following eight moments: The mean and standard deviation of the risk free rate; the mean of the market return; the mean and standard deviation of the price-dividend ratio; the auto-correlation of dividend growth; the standard deviation of consumption growth; and the correlation between consumption growth and the price-dividend ratio. The parameter estimates and model fit remain largely unchanged if the efficient weighting matrix is used instead; these results are available from the authors upon request.

4.2 Estimation Results in the Main Model (learning from both consumption and an unspecified signal)

We demonstrate that investors learn from signals over and above consumption history by introducing an unspecified signal with unconstrained informativeness. The estimation results are presented in Table 1. As expected the model estimates a signal that is nearly fully informative. Recall that the signal is standardized to have volatility equal to one in both regimes and mean zero in the first regime by construction. In the third panel of Table

1, the estimated mean of the signal is –6.73 in the second regime. This implies that the signal is nearly fully informative as it is drawn from two very different distributions in the two regimes. In Section 6 we constrain the signal to equal a particular macroeconomic variable, for example, the change in CPI or the earnings per hour. We re-estimate the model and find that the signal is only partly informative but nevertheless explains the data well.

[Table 1 here]

In the first two panels of Table 1 we display the sample moments and the model-generated moments of the consumption and dividend growth rates, market return, risk free rate, and market-wide price-dividend ratio. In the "Data" row we report the moments computed from the data along with standard errors (Newey-West (1987) corrected using two lags) in parentheses. In the "Model" row we present the model-generated moments along with the 95% confidence intervals in square brackets. The model-generated moments are calculated analytically whenever analytical solutions are available and from a single long simulation of length one million otherwise. The 95% confidence intervals are obtained from 10,000 simulations of eighty-four years each, the same size as the historical sample.

The model does a good job in matching the unconditional first and second moments of aggregate consumption growth: the unconditional mean, volatility, and first-order autocorrelation of consumption growth are 0.019, 0.022, and 0.475, respectively, in the data, while their model-implied values are 0.005, 0.030, and 0.025, respectively.[2]

Note that the model underestimates the first-order autocorrelation of consumption growth. Although the first-order autocorrelation of consumption growth is quite high in the data at 0.475, it has been argued in the literature that the high first-order autocorrelation is an artifact of the data, driven by measurement error and temporal aggregation. This view is further supported by the following two observations. First, the first-order autocorrelation of consumption growth is quite sensitive to the precise sample period and the measure of consumption used. It takes the value 0.48 over the 1929 to 2013 sample period when nondurables and services consumption is used as the measure of consumption expenditures, while it is close to zero over the longer 1890 to 2009 period where total consumption is the measure of consumption expenditures. Second, even though the first-order autocorrelation of consumption growth is high over the 1929 to 2013 period, the higher-order autocorrelations drop to zero quickly: the second-, third-, and fifth-order autocorrelations are 0.18, -0.06, and -0.009, respectively. Models that feature a first-order autocorrelation of consumption growth comparable to its sample counterpart typically imply higher-order

[2] The confidence interval is sometimes not centered at the point estimate because of finite sample bias.

autocorrelations much higher than those observed in the data. Our parameter estimates, on the other hand, imply higher-order autocorrelations close to their sample counterparts.

If we were to take at face value the first-order autocorrelation of consumption growth in the data (0.475) we would conclude that consumption growth is an informative signal. In Section 4.3 we re-estimate the model when the investors learn from the consumption history alone and find that this alternative model performs poorly.

Consumption growth has low correlation with the price-dividend ratio in the data. If a learning model features learning from consumption growth and a high autocorrelation of consumption growth, then the current realization of consumption growth contains a lot of information about the future and, therefore, the price-dividend ratio responds sharply to it. This generates the counterfactual prediction that consumption growth is highly correlated with the price-dividend ratio. This is why, in order to match the low observed correlation of the consumption growth with the price-dividend ratio, our main model sets the autocorrelation of consumption growth to be close to zero.

The unconditional mean, volatility, and first-order autocorrelation for the dividend growth rate are 0.012, 0.113, and 0.183, respectively, in the data and 0.001, 0.207, and 0.172, respectively, in the model. The model does a good job in matching these moments, albeit somewhat overestimating the volatility of dividends.[3] The correlation between the consumption and dividend growth rates is 0.567 in the data and is slightly higher than the model-implied value of 0.489. The correlation between the consumption growth rate and the price-dividend ratio is 0.135 in the data and is close to the model-implied value of 0.168.

In the second panel of Table 1 the model-implied mean risk free rate is 0.005, consistent with the sample value of 0.005. The model-implied volatility of the risk free rate is lower than the value in the data and the autocorrelation is higher than in the data.

The model-implied mean equity return is 0.050, close to the sample value of 0.061. The model-implied volatility of the market return is 0.287, slightly higher than the value of 0.200 in the data. The autocorrelation of the market return is essentially zero both in the data and the model.

The model-implied mean, volatility, and autocorrelation of the price-dividend ratio are 3.057, 0.434, and 0.902, respectively, in excellent agreement with the values of 3.393, 0.455, and 0.862, respectively, in the data.

[3] A large literature in corporate finance discusses motives for dividend smoothing. Therefore, overstating the volatility of dividend growth and, consequently, that of the market return should not necessarily be viewed as a shortcoming of the model.

Overall the model rationalizes the high mean market return, the volatility of the market return, and the low mean risk free rate observed in the data. Therefore it offers an explanation of the equity premium and risk free rate puzzles. It also rationalizes the mean and, more importantly, the volatility of the market-wide price-dividend ratio, thereby accounting for the excess volatility puzzle. To summarize, the model fits well the moments of consumption and dividend growth and returns, and does so without requiring implausible dynamics of the consumption and dividend growth processes.

In the third panel of Table 1 we display the point estimates of the parameters, along with the associated standard errors (Newey-West (1987) corrected using two lags) in parentheses. The point estimates of the coefficient of RRA, 13.9, and the EIS, 0.558, suggest preference for early resolution of uncertainty. The EIS is substantially lower than is typically required in models with recursive preferences. For example, in the Bansal and Yaron (2004) long-run risks model, the calibrated value of the RRA coefficient is 10 and of the EIS is 1.5; and in the Nakamura, Steinsson, Barro, and Ursua (2013) rare disasters model, the calibrated value of the RRA coefficient is 6.5 and of the EIS is 2.0.

The point estimates of the transition probabilities strongly suggest the existence of at least two regimes, since π_1 is very different from $1-\pi_2$. The consumption and dividend growth rates have higher mean in the first regime than in the second one. Therefore we identify the first regime as the regime of economic expansion and the second regime as the regime of economic contraction or recovery.

The signal is highly persistent with a first-order autocorrelation coefficient of 0.78. At first sight, it may seem puzzling that the signal has persistence that is an order of magnitude higher than that of consumption growth (0.78 versus 0.02), despite both of them being driven by the same underlying state s_t. However, closer look reveals that, for given values of the transition probabilities, the first-order autocorrelation coefficient of each of these variables is increasing in the ratio of the squared difference between the means in the two states to the constant variance, $\left(\mu_{y,1}-\mu_{y,2}\right)^2/\sigma_y^2$. For the consumption growth, this ratio is 0.14, while for the signal *x* this ratio is two orders of magnitude bigger at 36.8. This explains the substantially higher persistence of the signal compared to that of the consumption growth rate. We argue, in Section 7, that this feature of the signal plays a key role in generating the asset pricing results.

4.3 Estimation Results in the Alternative Model (learning only from consumption history)

In the previous section we estimated the model when the investors learn both from the consumption history and the unspecified signal. In this section we re-estimate the model when the investors learn from the consumption history alone. The results are reported in Table 2.

[Table 2 here]

The model fails along several dimensions. The model-implied volatility of the price-dividend ratio is essentially zero, at odds with the high volatility of 0.455 observed in the data. The model-implied correlation of consumption growth with the price-dividend ratio is -0.756, at odds with the correlation of 0.135 observed in the data. Finally the model-implied auto-correlation of the price-dividend ratio is 0.213, at odds with the auto-correlation of 0.862 observed in the data. These results demonstrate that the alternative model simply does not fit the data. Some shortcomings of this model have earlier been highlighted in Johannes, Lochstoer, and Mou (2016) and our results confirm their findings.

5 What Are Potential Signals?

At the core of consumption-based asset pricing is the insight that asset prices reflect investors' beliefs about macroeconomic outcomes. This naturally raises a number of central questions. Given the large amount of macroeconomic information available to investors, which subset of information do investors pay the most attention to? How do they process this information to form beliefs about macroeconomic outcomes? In the consumption-based asset pricing literature, learning models typically hypothesize that investors learn about the latent economic regime from the history of consumption alone. These models perform poorly in explaining several stylized facts of asset market data including the equity premium, risk free rate, and excess volatility puzzles. Our results in Section 4 suggest that the poor performance of these models may potentially be attributed to the restrictive assumption that the information set of investors is the history of aggregate consumption alone. In this section we present further evidence supporting this argument.

We relate the market-wide price-dividend ratio to a broad cross-section of publicly available macroeconomic variables. The 106 macroeconomic variables we consider are obtained from the Global Insights Basic Economic Database and may be broadly classified into the following six categories: (1) output and income, which includes personal income, industrial production index (total as well as disaggregated by the type of product), and capacity utilization measures; (2) labor markets, which includes the unemployment rate, unemployment insurance claims, and employees in different sectors; (3) housing, which includes the number of authorized permits to build houses and the number of new housing constructions started in different geographical regions of the US; (4) consumption, orders, and inventories, which includes real personal consumption expenditures, the Index of Consumer Expectations, manufacturing and trade sales, manufacturing and trade inventories, ratio of manufacturing and trade inventories to sales, and new orders for different types of goods (consumer goods and materials, durable goods, nondefense capital goods); (5)

money and credit, including M1, M2, currency held by the public, commercial and industrial loans outstanding, consumer credit outstanding, and the ratio of consumer installment credit to personal income; and (6) prices, which includes the producer price index for different types of goods, the consumer price index for different goods and services, and the commodity prices index.

The price-dividend ratio is persistent but stationary with first-order autocorrelation coefficient 0.86 which is statistically smaller than one over the entire available sample period from 1929 to 2013. We apply transformations to some of the macroeconomic variables to make them stationary. Each of the transformed variables has first-order autocorrelation coefficient less than 0.9, thereby suggesting that it can appropriately be characterized as stationary. We refer the reader to Ludvigson's website for a detailed description of the macroeconomic variables and the transformations applied to make them stationary (http://www.econ.nyu.edu/user/ludvigsons/Data&ReplicationFiles.zip).

We relate the price-dividend ratio to the macroeconomic variables by running univariate regressions of the price-dividend ratio on the 106 transformed macroeconomic variables. In figure 1 we present the R^2 of these regressions. The figure reveals that prices and labor markets are two of the broad categories of macro variables that are strongly related to the price-dividend ratio. We obtain a similar ranking when we regress the price-dividend ratio on an exponentially-weighted moving average of the current and lagged values of the macro variables with five and ten lags.

[Figure 1 here]

The price variables include the Consumer Price Index (CPI), the Producer Price Index (PPI), and the implicit price deflator for personal consumption expenditures. The regression of the price-dividend ratio on the CPI growth for all Urban Consumers (CPI-U) for all items has $R^2 = 48.2\%$. The regressions on the CPI-U growth for disaggregated expenditure categories all have high R^2 as well: 66.0% for apparel and upkeep, 61.5% for medical care, 58.8% for durables, 49.4% for services, 40.4% for commodities, and 19.4% for transportation. Regressions on the growth of the implicit price deflator for personal consumption expenditures on durables, nondurables, and services have high R^2: 70.5%, 27.3%, and 63.0%, respectively. Finally regressions on the PPI growth variables have somewhat smaller, but still substantial, R^2: 27.4% for finished goods, 18.8% for finished consumer goods, and 15.9% for intermediate materials supplies and components.

The second category of variables strongly related to the price-dividend ratio corresponds to the labor market. In particular regressions on the growth of average hourly earnings of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing sector, goods-producing sector, and construction sector have R^2 42.4%, 36.8%, and 20.7%, respectively. Regressions on the growth of average weekly hours of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing and goods-producing sectors have R^2 30.8% and 24.6%, respectively. Finally, regressions on the growth of the number of employees on nonfarm payrolls in the financial sector, retail trade sector, wholesale trade, and trade, transportation, and utilities sector have R^2 13.4%, 13.3%, 11.4%, and 10.8%, respectively.

Note that, in the consumption-based asset pricing literature, most existing learning models assume that investors learn from the consumption history alone. Therefore, the equilibrium price-dividend ratio is a function of the investors' beliefs which, in turn, depend on the history of consumption growth alone. However, a regression of the market-wide price-dividend ratio on the contemporaneous consumption growth produces a statistically insignificant slope coefficient and $R^2 = 1.7\%$; and regressions of the price-dividend ratio on an exponentially-weighted moving average of the contemporaneous and lagged consumption growth rates produce even smaller $R^2 = 0.72\%$ and 0.05%, respectively, when five and ten lags of the consumption growth are included in the regressor.

The above results suggest that a substantial fraction of the variation in the price-dividend ratio can indeed be explained by variations in macroeconomic variables, but not by the one variable commonly assumed in the literature, namely the aggregate consumption growth. Of course, if the information set of the investors were expanded to include additional macroeconomic variables, then their beliefs and, therefore, the equilibrium price-dividend ratio would depend on these variables in addition to consumption growth, a feature that seems to be supported by the data.

Note that most macroeconomic variables are either pro-cyclical or counter-cyclical. However this does not necessarily imply that they are strongly correlated with the price-dividend ratio or serve as informative signals in the context of our model. Figure 2 displays the time series of GDP growth, consumption growth, change in CPI, and change in hourly earnings; and figure 3 displays the time series of the first three principal components of the covariance matrix of the 106 macroeconomic variables. All these variables have a business-cycle pattern but GDP growth, consumption growth, and the first and third principal components of the covariance matrix are neither strongly correlated with the price-dividend ratio nor serve as informative signals in the context of our model.

[Figures 2 and 3 here]

6 The Information that Drives Asset Prices

In Section 4 the signal is deliberately left unspecified in order to show that consumption growth is an incomplete signal and investors' beliefs are driven by a latent signal in addition to consumption growth. In this section we demonstrate that the model fit remains largely intact when we replace the latent signal with the innovation in the CPI, hourly earnings, or a combination of macroeconomic variables highly correlated with the price-dividend ratio. The results highlight the informational role of macroeconomic variables and suggests that just one macroeconomic variable (CPI or earnings per hour), along with consumption growth, goes a long way towards proxying for the investors' relevant information set.

In the first example, we constrain the signal to equal the change in CPI for all urban consumers (CPI-U). Thus equation (4) that describes the dynamics of the signal is replaced by the following equation:

$$\Delta CPI_{t+1} = \mu_{\Delta CPI, s_{t+1}} + \sigma_{\Delta CPI}\varepsilon_{\Delta CPI,t+1}, \tag{19}$$

where ΔCPI_{t+1} is the change in CPI over one period and $\varepsilon_{\Delta CPI,t+1}$ is *i.i.d.* standard normal and orthogonal to $\varepsilon_{c,t+1}$ and $\varepsilon_{d,t+1}$. Equation (19) has three parameters, $\mu_{\Delta CPI,1}, \mu_{\Delta CPI,2}$, and $\sigma_{\Delta CPI}$, which is two more parameters than that in equation (4) in which we set $\mu_{x,1} = 0$ and $\sigma = 1$. However the CPI dynamics now have to match four additional moment restrictions, the mean, variance, and first-order autocorrelation of inflation growth and the correlation between the price-dividend ratio and inflation growth.

The results are presented in Table 3 and cover the period from 1964 to 2013. The starting date of 1964 is chosen to coincide with the date when data on the broad cross-section of macro variables become available. The moments estimated from the data over this period are somewhat different from the corresponding moments estimated from the data from 1929 to 2013, as reported in Table 1. In the first two panels of Table 3 we display the sample moments and the model-generated moments of the consumption and dividend growth rates, market return, risk free rate, and market-wide price-dividend ratio. In the "Data" row we report the moments computed from the data along with standard errors (Newey-West (1987) corrected using two lags) in parentheses. In the "Model" row we present the model-generated moments along with the 95% confidence intervals in square brackets. The model-generated moments are calculated analytically whenever analytical solutions are available and from a single long simulation of length one million otherwise. The 95% confidence intervals are obtained from 10,000 simulations of eighty-four years each, the same size as the historical sample.

[Table 3 here]

Even though the signal is forced to be the CPI growth instead of being left unspecified and there are two additional overidentifying restrictions in the estimation, the model does well in matching the data. In the first panel of Table 3 we see that the model explains the mean and volatility of consumption and dividend growth and the auto-correlation of dividend growth but does not match the auto-correlation of consumption growth. Recall that the model with unspecified signal does not match the latter moment either (Table 1).

In the second panel of Table 3 the model rationalizes the mean and auto-correlation of the risk free rate, the mean, volatility, and auto-correlation of the market return and price-dividend ratio but implies a lower volatility of the risk free rate than the volatility observed in the data. Recall that the model with unspecified signal does not match the latter moment either (Table 1). In the third panel the model matches well the correlation between CPI growth and the price dividend ratio but lesser so the mean, volatility, and the autocorrelation of CPI growth.

In the fourth panel, the model-implied RRA is similar to the model-implied RRA in the model with unspecified signal but the IES is 1.9 which is higher than in the model-implied IES in the model with unspecified signal (Table 1). The other parameter estimates are reasonable. In particular, the two regimes are very persistent with persistence comparable to the persistence of the regimes when the signal is left unspecified (Tables 1). Overall, the model retains its good empirical fit when the signal is constrained to equal CPI growth.[4]

In addition to assessing the model fit for the unconditional moments of consumption and dividend growth rates and returns, when the signal is constrained to be an observable variable, such as the change in the CPI, we can compute the correlation between the historically observed price-dividend ratio and that implied by the model. This is because the model-implied price-dividend ratio is a known function of the investors' beliefs, which in turn is a known function of the history of the consumption growth and the change in the CPI. Therefore the time series of the consumption growth and the change in the CPI can be used to obtain the time series of the model-implied price-dividend ratio. Using the parameter estimates in Table 3, the model-implied time series of the price-dividend ratio has a correlation of 0.69 with the price-dividend ratio observed in the data. The alternative model, on the other hand, produces a correlation of zero between the model-implied and historical price-dividend ratios. This is because the model-implied price-dividend ratio is an almost flat function of the beliefs in the alternative model.

Also note that the correlation between the rate of change of the CPI and the price-dividend ratio is consistently high over the period 1964-2013. In particular the correlation is -67.5% from 1964 to 2013, and remains high at -70.1% and -42.7% over the non-overlapping sub-periods 1964-1989 and 1990-2013,

[4] The 10-year nominal bond yield does just as well as the rate of inflation as a signal.

respectively. Therefore the high correlation is not driven by the high inflation in the 1970s followed by the stagflation in the 1980s. Moreover, the high correlation between the change in the CPI and price-dividend ratio is also observed over the longer period since the end of the Great Depression—the correlation is -39.6% over the post war period from 1947 to 2013 and -33.3% from 1935 to 2013.

We also provide another evidence in favor of the risk channel highlighted in our model. The rate of inflation is chosen as a potential signal because of its high correlation with the market-wide price-dividend ratio. If innovations to the inflation process drive variations in the price-dividend ratio due to the mechanism highlighted in the model, then the innovations to the inflation process should also forecast future cash flow growth. We confirm that this is indeed the case. In particular the rate of inflation forecasts one-year ahead dividend growth with a statistically significant slope coefficient of -0.76 and $R^2 = 7.8\%$ in the historical sample. The corresponding model-implied values of the slope coefficient and R^2, obtained from a single long simulation of length one million, are close to their data counterparts at -0.70 and 11.7%, respectively. The simulated 95% confidence intervals of the slope coefficient and R^2, obtained from 10,000 simulations of length equal to the historical sample, contain the respective sample values. Moreover, not only does the rate of inflation forecast the one-year ahead dividend growth, but it forecasts longer horizon growth rates as well: it forecasts the two-year ahead dividend growth with a statistically significant slope of -1.29 and $R^2 = 8.8\%$, and the three-year ahead dividend growth with a statistically significant slope of -1.52 and $R^2 = 8.1\%$. As with the one-year ahead dividend growth, the model-implied values of these statistics—slope of -1.35 and $R^2 = 17.3\%$ for the two-year dividend growth, and slope of -1.99 and $R^2 = 20.1\%$ for the three-year dividend growth—are close to their data counterparts. This lends further support to the central insight highlighted in the model.

In Section 4 we contrasted the model fit when investors learn from consumption history and a latent signal (Table 1) with the results when investors learn from consumption history alone (Table 2) over the period 1929-2013 and concluded that the latter model fails along several dimensions. We extend this comparison over the subperiod 1964 to 2013 and reach a similar conclusion. We re-estimate the model when investors learn from consumption history alone over the subperiod 1964 to 2013 and present the results in Table 4. The model fails along several dimensions. The model-implied volatility of the price-dividend ratio is essentially zero, the expected market return is too low and the correlation of consumption growth with the price-dividend ratio is far more negative than in the data. These results demonstrate that the alternative model does not fit the data.

[Table 4 here]

The change in average hourly earnings of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing sector (ΔEAR) and the second principal component of changes in the macroeconomic variables $(MACRO)$ are highly correlated with the market-wide price-dividend ratio while the first and third principal components are not. (The first, second, and third principal components explain 40.6%, 20.0%, and 8.8%, respectively, of the variation in the macro variables.) Therefore, as a second example, we set the signal equal to the change in the average hourly earnings (Table 4) and, as a third example, we set the signal equal to the second principal component (Table 5). The results are strikingly similar to the results in Table 3.

[Tables 5 and 6 here]

Recall that the target of our paper is to highlight the pivotal role of expanding the information set of the investors beyond the consumption history in addressing several seemingly puzzling aspects of asset market data. In this section we show that CPI growth, average hourly earnings of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing sector, and the second principal component of changes in macroeconomic variables are major signals that affect the level of market prices and returns. Whereas we cannot claim that these are the particular signals that investors focus on and regardless of what information investors are considering, we find empirical evidence that their focus manifests itself in price-level and labor-market variables.

7 Discussion of the Results and Interpretation of the Economic Regimes

In Sections 4 and 6 we showed that the main model fits the data well while the alternative model fails along several dimensions. The success of the main model is a particularly strong result because it derives solely from expanding the information set of investors to include either an unspecified signal or one of a number of macro variables over and above consumption growth. The question then arises as to what drives the superior performance of the main model. We illustrate the intuition using the parameter estimates in Table 3 when the signal is set equal to the rate of change in the CPI.

The crux of the intuition lies in the response of the price-dividend ratio to changes in beliefs, p_t, and, therefore, to changing expectations about future returns and dividend growth. In the main model, the belief process has high persistence and low variance, while in the alternative model the belief process has low persistence and high variance. In the main model the current value of p_t is highly informative about the future and, therefore, the price-dividend ratio responds sharply to changes in p_t. This makes the price-dividend ratio highly volatile in this model. By contrast, in the alternative model, the current value of p_t provides little information about the future and, therefore, the price-dividend ratio is essentially constant. Below, we elaborate on this intuition.

7.1 The Beliefs Process

The estimates of the transition probabilities for both the main and alternative models are high. This implies that there is high uncertainty about the future when p_t is around 0.5 but not when p_t is near its boundaries of 0 or 1. Investors have a strong precautionary savings motive around $p_t = 0.5$ that declines as p_t approaches its boundaries. This generates non-linearities in the price-dividend ratio, expected market return, and conditional variance of the market return.

In the main model the investors' beliefs are a function of consumption growth and inflation growth. The estimated first-order auto-correlation coefficient of inflation growth is 0.78 and the model-implied autocorrelation is 0.43 while the estimated first-order auto-correlation coefficient of consumption growth is 0.51 and the model-implied autocorrelation is 0.02. Therefore p_t has higher persistence in the main model than in the alternative model. The above features of the learning process explain why the price-dividend ratio is highly volatile in the main model. The price-dividend ratio moves in response to changing expectations about future returns and future dividend growth. In the main model the current realization of p_t is informative about the future and, therefore, the price-dividend ratio is highly responsive to changes in p_t. In figure 4 we illustrate the conditional variance of the beliefs, $\mathrm{var}\left(p_{t+1} \mid p_t\right)$, as a function of p_t in the main model (red line).and the alternative model (blue line).

[Figure 4 here]

7.2. The Risk Free Rate

In figure 5 we plot the risk free rate as a function of the probability of the expansion state in the main model (red line) and in the alternative model (blue line). In the alternative model the risk free rate is (almost) linearly increasing in the probability of the expansion state p_t . Similarly, in the main model, the risk free rate is also low during recessions, consistent with the empirical evidence, and mostly increasing in the probability of the expansion state. However the risk free rate is a nonlinear function of the probability in this model and, for very high values of p_t, is decreasing in p_t. The latter feature of the model obtains because, when the probability is very high, the substitution effect for the dividend claim is very strong (because the mean of the dividend growth rate is very different in the two regimes and the transition probabilities are very high). Therefore, investors buy this asset rather than borrowing against the future to smooth consumption. This does not happen in the alternative model because the income effect always dominates the substitution effect in this model (because the mean of dividend growth is less different between the two regimes and the transition probabilities are lower than those in the main model).

[Figure 5 here]

7.3. The Price-Dividend Ratio

In figure 6 we display the price-dividend ratio as a function of the probability of the economic expansion state, p_t, in the main model (red line) and the alternative model (blue line). In the main model the price-dividend ratio is sharply increasing and convex in the probability of the expansion state (red line). This nonlinearity justifies our approach of eschewing the Campbell and Shiller (1988) log-linearization or any other form of approximation in solving the model. By contrast in the alternative model the price-dividend ratio is (almost) flat in the probability of the expansion state (blue line) because the current beliefs are not very informative about the future. The latter model is, therefore, unable to generate the observed volatility of the price-dividend ratio. To help improve its performance Johannes, Lochstoer, and Mou (2016) introduce model and parameter uncertainty. Even though the investors learn from consumption alone, model and parameter uncertainty slows down the learning process and imparts higher persistence in p_t compared to the persistence of consumption growth and this improves the model fit to the data.

[Figure 6 here]

7.4. The Market Return

In figure 7 we display the expected market return as a function of the probability of the expansion state. The alternative model implies that the expected market return is an increasing function of p_t (blue line), leading to the counterfactual prediction of a procyclical expected market return. This occurs because in the alternative model the price-dividend ratio is almost constant and the expected dividend growth is an increasing function of p_t.

[Figure 7 here]

In the main model, on the other hand, when $p_t \approx 0.5$ there is maximum uncertainty about the current regime. The estimates of the preference parameters suggest a strong preference for early resolution of uncertainty. Therefore the expected market return is at its highest around this point (red line). As p_t increases from 0.5 towards 1 there is less uncertainty that the economy is in the good regime. Therefore the expected market return is decreasing in p_t over this range. A more formal way of stating this is that, over this range of p_t, an increase in p_t causes a sharp increase in the price-dividend ratio and this has the effect of decreasing the expected return that dominates the increase in the expected return arising from the increase in the expected dividend growth and the increase in expected future price-dividend ratio due to the rise in p_t.

The situation is more complicated for p_t between 0 and 0.5. As p_t increases over this range, there are two competing forces. On the one hand there is more uncertainty about the regime that causes the expected market return to increase; and, on the other hand, there is a lower probability of being in the second regime and this causes the expected market return to decrease. The former effect dominates because of the strong preference for early resolution of uncertainty, causing the expected market return to increases with an increase in p_t over this range. A more formal way of stating this is that, over this range of p_t, an increase in p_t causes a less sharp increase in the price-dividend ratio (the function is less convex over this range) and this has the effect of decreasing the expected return that is dominated by the increase in the expected return arising from the increase in the expected dividend growth and the increase in expected price-dividend ratio due to the rise in p_t.

In figure 8 we display the expected market return as a function of the price-dividend ratio. In the main model the expected market return is strongly non-linear in the price-dividend ratio and peaks in the middle value of the price-dividend ratio when uncertainty is at its highest (red line). This highly non-linear pattern is unlike the common practice of predicting the market return with the price-dividend ratio with linear regressions. In the alternative model the plot of the expected market return as a function of the price-dividend ratio does not make sense because the price-dividend ratio is insensitive to changes in p_t (blue line).

[Figure 8 here]

In figure 9 we display the conditional variance of the market return in both models. The conditional variance of the market return depends on the conditional variance of the price-dividend ratio and dividend growth. The conditional variance of both of these variables depend on the conditional variance of p_{t+1}. Now the conditional variance of p_{t+1} is an inverted *U*-shaped function of p_t—being the highest when $p_t \approx 0.5$

and declining when p_t approaches its boundaries of 0 or 1. In the alternative model the price-dividend ratio is almost constant and, therefore, the conditional variance of the market return is driven only by the conditional variance of dividend growth. The latter is small to be consistent with the data and therefore the conditional variance of the market return is low and almost flat (blue line). In the main model uncertainty about the current regime peaks at around $p_t \approx 0.5$ and, therefore, the conditional variance of the market return peaks at around $p_t \approx 0.5$ (red line). To summarize, the main model, unlike the alternative model, generates strong time-variation in the conditional mean and variance of the stock market return. And it does so without relying on heteroscedasticity of the aggregate consumption and dividend growth rates—a phenomenon for which there is limited empirical evidence.

[Figure 9 here]

8 Predictive Regressions

We further assess the empirical plausibility of the main and the alternative models with predictive regressions for the consumption growth, dividend growth and market return. Both models imply that the expected consumption and dividend growth rates are linearly increasing in the beliefs. However the models differ in their implications on the expected market return: the main model implies that the expected market return is an inverted *U*-shaped function of the beliefs while the alternative model implies that the expected market return is linearly increasing in the beliefs.

The time series of beliefs are extracted from the observed price-dividend ratio and risk free rate. In figure 10 we present the time series of the beliefs, p_t for the main model (red line) as well as for the alternative model (blue line).[5] Also shown in the figure are the NBER-dated recession periods as shaded rectangles and the major stock market downturns in the twentieth century identified in Barro and Ursua (2012) as black-dashed vertical lines. The extracted beliefs significantly differ in the main and alternative models and we further address the issue with predictive regressions.

[Figure 10 here]

[5] Note that, in each year *t*, we have a system of two equations (one for the price-dividend ratio and one for the risk free rate) in one unknown, p_t. We extract p_t by minimizing a weighted non-linear least squares criterion function. The weighting matrix used is a diagonal matrix with the diagonal entries equal to the inverse of the variances of the price-dividend ratio and the risk free rate.

In Table 7 we present the results of predicting consumption growth, dividend growth, and market return with beliefs. In the data panels beliefs are extracted from the observed price-dividend ratio and risk free rate while in the model panels beliefs are obtained from a long simulation of the model of length one million years.

[Table 7 here]

In panel A we present the results for the main model. In predictive regressions of consumption growth the coefficient of p_t and the R^2 are close to zero in both cases where p_t stands for the data-implied and the model-implied beliefs. In predictive regressions of dividend growth the data-implied coefficient of p_t is 0.11 but statistically insignificant and the R^2 is 4.2%. The model-implied coefficient of p_t is 0.14 and statistically significant and the R^2 is 19.3%, larger than in the data. We conclude that the data-implied and the model-implied predictive regressions of consumption and dividend growth are in agreement for the main model.

As shown in figure 5 the model implies that the expected market return is first increasing and then decreasing in p_t. Therefore we include p_t^2 as an additional regressor in the predictive regression for the market return. In both the data driven and the model-driven regressions Consistent with the model, the coefficient of p_t is positive and the coefficient of p_t^2 is negative but the order of magnitude of the coefficients are different. The R^2 is small at 4.4% in the data, consistent with its model-implied value of 3.3%.

In results available from the authors upon request we also show that the model generates the low predictive power of the price-dividend ratio for the one-year ahead as well as long horizon consumption growth rates, consistent with the data. This stands in stark contrast with the long-run risks model of Bansal and Yaron (2004) that implies an order of magnitude higher R^2 in forecasting regressions of the one-year and three-year ahead consumption growth rates and two orders of magnitude higher R^2 in a forecasting regression of the five-year ahead consumption growth rate on the price-dividend ratio (see, e.g., Beeler and Campbell (2012)). The superior performance of our main model is rendered possible because it does not rely on the persistence of consumption growth to generate key asset pricing results, relying instead on the persistence of the additional macro variables in the information set.

In panels C and D of Table 7 we present the results corresponding to panels A and B, respectively, but for the alternative model. For the consumption growth the regression coefficient is 0.006 and the R^2 is 1.8% in the data-driven forecasting regressions and 0.27 and 36.4%, respectively, in the model-driven forecasting regressions. For the dividend growth the regression coefficient is -0.41 and the R^2 is 3.2% in the data-driven forecasting regressions and 0.61 and 3.9%, respectively in the model-driven forecasting regressions. In the forecasting regressions of the market return all regression coefficients have the wrong sign. These findings further demonstrate that the alternative model does not fit the data.

9 Concluding Remarks

The market-wide price-dividend ratio is strongly correlated with two groups of macroeconomic variables, price variables including the rate of change in the CPI and PPI and labor market variables including the average hourly earnings, average hours, and number of employees in different sectors. However the price-dividend ratio has very small correlation with the aggregate consumption growth, the one variable that investors are assumed to learn from in an extensive literature on learning in financial markets. This suggests that the poor empirical performance of these learning models may potentially be explained by the stringent information set imposed on investors, namely the assumption that they learn from the history of consumption alone. In this paper we explore the role of expanding the information set of investors in explaining several stylized facts of stock market data.

We present a model of a real exchange economy with learning about the state of the economy from the consumption history and a latent signal. The model offers an explanation of the equity premium and risk free rate puzzles. It also rationalizes the mean and, more importantly, the volatility of the market-wide price-dividend ratio, thereby accounting for the excess volatility puzzle. Finally it is also consistent with the low predictive power of the price-dividend ratio for future consumption growth. The signal allows the investors to update their beliefs regarding the probability on whether the economy is in an expansion or recession state. The model demonstrates that the market rationally processes macroeconomic information in forming beliefs and setting prices.

We re-estimate a nested version of the model in which the investors learn from the consumption history alone, a specification commonly used in the literature. The model fails to address the equity premium puzzle and the observed high volatility of the market return relative to dividend growth. It significantly understates the volatility of the price-dividend ratio. It also performs poorly in predictive regressions.

Our analysis highlights the pivotal role of expanding the information set of the investors in addressing several seemingly puzzling aspects of asset market data. In particular we show that CPI growth, average hourly earnings of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing sector, and the second principal component of changes in macroeconomic variables are proxies of the set of signals that affect the level of market prices and returns.

References

Ai, H., 2010, "Information Quality and Long-Run Risk: Asset Pricing Implications," *Journal of Finance* 65: 1333-1367.

Ai, H. and R. Bansal, 2016, "Risk Preferences and the Macro Announcement Premium," working paper, Duke University.

Albuquerque, R., M. Eichenbaum, and S. Rebelo, 2012, "Valuation Risk and Asset Pricing," *Journal of Finance*, forthcoming.

Andrei, D., M. Hasler, and A. Jeanneretz, 2016, "Asset Pricing Implications of Learning about Long-Run Risk," working paper, UCLA.

Bansal, R. and I. Shaliastovich, 2011, "Learning and Asset Price Jumps," *Review of Financial Studies* 24: 2738-2780.

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

Barro, R. J. and J. F. Ursùa, 2012, "Rare Macroeconomic Disasters," *Annual Review of Economics* 4: 83-109.

Basak, S., 2005, "Asset Pricing with Heterogeneous Beliefs," *Journal of Banking and Finance* 29: 2849-2881.

Beeler, J. and J. Y. Campbell, 2012, "The Long-Run Risks Model and Aggregate Asset Prices: an Empirical Assessment," *Critical Finance Review* 1: 141-182.

Bidder, R. and I. Dew-Becker, 2016, "Long-Run Risk Is the Worst-Case Scenario, *American Economic Review* 106: 2496-2527.

Campbell, J., Y. C. Pflueger, and L. M. Viceira, 2015, "Monetary Policy Drivers of Bond and Equity Risks," working paper, Harvard University.

Campbell, J. Y. and R. J. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1: 195-227.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer, 2016, "Parameter Learning in General Equilibrium: The Asset Pricing Implications," *American Economic Review* 106: 664-698.

Constantinides, G. M. and A. Ghosh, 2011, "Asset Pricing Tests with Long Run Risks in Consumption Growth," *Review of Asset Pricing Studies* 1: 96-136.

Croce, M. M., M. Lettau, and S. C. Ludvigson, 2015, "Investor Information, Long-Run Risk, and the Term Structure of Equity," *Review of Financial Studies* 28: 706-742.

D'Acundo, F., D. Hoang, and M. Weber, 2016, "The Effect of Unconventional Fiscal Policy on Consumption Expenditure," working paper, University of Chicago.

David, A. and P. Veronesi, 2013, "What Ties Return Volatilities to Fundamentals and Price Valuations?" *Journal of Political Economy* 121: 682-746.

Drechsler, I., 2013, "Uncertainty, Time-Varying Fear, and Asset Prices," *Journal of Finance* 68: 1843-1889.

Duffee, G. R., 2005, "Time Variation in the Covariance between Stock Returns and Consumption Growth," *Journal of Finance* 60: 1673-1712.

Epstein, L. G. and M. Schneider, 2003, "Recursive Multiple-Priors," *Journal of Economic Theory* 113: 1–31.

Epstein, L. R. and S. E. Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* 57: 937-969.

Gilbert, T., 2016, Information Aggregation around Macroeconomic Announcements: Revisions Matter," *Journal of Financial Economics*, forthcoming.

Greenwald, D. I., M. Lettau, and S. C. Ludvigson, 2015, "Origins of Stock Market Fluctuations," working paper, New York University.

Gurkaynak, R. S., B. Sack, and E. Swanson, 2005, "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models," *American Economic Review* 95: 425-436.

Hall, R. E., 2014, "High Discounts and High Unemployment," Working paper, Stanford University/

Hansen, L. P., J. Heaton, and N. Li, 2008, "Consumption Strikes back: Measuring Long-Run Risk," *Journal of Political Economy* 116: 260-302.

Hansen, L.P., P. Hansen, P., and P.A Mykland, 2016, "Measuring Belief Distortions: An Information-Theoretic Approach," Working paper, University of Chicago.

Hansen, L. P. and T. J. Sargent, 2001, "Robust Control and Model Uncertainty," *American Economic Review* 91: 60–66.

Johannes, M., L. Lochstoer, and Y. Mou, 2016, "Learning about Consumption Dynamics," *Journal of Finance* 71: 551-600.

Klibanoff, P., M. Marinacci, and S. Mukerji, 2005, “A Smooth Model of Decision Making under Ambiguity,” *Econometrica* 73: 1849–92.

Lettau, M. and S. Ludvigson, 2001, “Consumption, Aggregate Wealth, and Expected Stock Returns,” *Journal of Finance* 56: 815-849.

Lettau, M. and S. Ludvigson, 2013, “Shocks and Crashes,” *National Bureau of Economic Macroeconomics Annual*, edited by J. Parker and M. Woodford, MIT Press.

Li, G., 2005, “Information Quality, Learning and Stock Market Returns,” *Journal of Financial and Quantitative Analysis* 40: 595-620.

Lucas, Jr., R. E., 1978, “Asset Prices in an Exchange Economy,” *Econometrica* 46: 1429-1445.

Maccheroni, F., M. Marinacci, and A. Rustichini, 2006. “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica* 74: 1447–98.

Nakamura, E., J. Steinsson, R. J. Barro, and J. F. Ursùa, 2013, “Crises and Recoveries in an Empirical Model of Consumption Disasters,” *American Economic Journal: Macroeconomics* 5: 35-74.

Newey, W. K. and K. D. West, 1987, “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica* 55: 702-708.

Pastor, L. and P. Veronesi, 2009, “Learning in Financial Markets,” *Annual Review of Financial Economics* 1: 361-381.

Van Nieuwerburgh, S. and L. Veldkamp, 2006, “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics* 53: 753-772.

Veronesi, P., 2000, “How Does Information Quality Affect Stock Returns?” *Journal of Finance* 55: 807-837.

Weil, P., 1990, “Nonexpected Utility in Macroeconomics,” *Quarterly Journal of Economics* 105: 29–42.

Table 1: Model Fit and Parameter Estimates with Learning from both Consumption and an Unspecified Signal, 1929-2013

The table reports estimation results and model fit for the main model with learning from both consumption growth and an unspecified signal, using annual data over the entire available sample period 1929-2013. The parameters are estimated using GMM. Seventeen moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, and the covariance between consumption growth and the price-dividend ratio. The number of parameters to be estimated is twelve. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c, \Delta d)$	$corr(\Delta c, p/d)$
Data	.019	.022	.475	.012	.113	.183	.567	.135
	(.003)	(.004)	(.149)	(.013)	(.019)	(.208)	(.133)	(.156)
Model	.005	.030	.025	.001	.207	.172	.489	.168
	[-.002 ,.013]	[.026, .035]	[-.205 ,.219]	[-.089 ,.089]	[.171, .235]	[-.119 ,.348]	[.303, .640]	[-.078, .358]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.005	.030	.672	.061	.200	.009	3.393	.455	.862
	(.005)	(.005)	(.121)	(.019)	(.018)	(.217)	(.080)	(.049)	(.054)
Model	.005	.009	.902	.050	.287	.028	3.057	.434	.902
	[-.003, .012]	[.003, .009]	[.474, .957]	[-.026, .129]	[.208, .368]	[-.205, .219]	[2.681, 3.433]	[.133, .436]	[.474, .957]

Parameter Estimates						
γ	ψ	δ	π_1	π_2	μ	$\rho_{c,d}$
13.940	.558	.990	.952	.950	-6.730	.469
(.0000)	(.0000)	(.0011)	(.0004)	(.0001)	(.0000)	(.0000)
$\mu_{c,1}$	$\mu_{c,2}$	$\mu_{d,1}$	$\mu_{d,2}$	σ_c	σ_d	
.010	.000	.090	-.091	.030	.186	
(.0021)	(.0021)	(.0208)	(.0200)	(.0023)	(.0000)	

Table 2: Model Fit and Parameter Estimates with Learning only from Consumption, 1929-2013

The table reports estimation results and model fit with consumption growth as the only signal using annual data over the entire available sample period 1929-2013. The parameters are estimated using GMM. Seventeen moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, and the covariance between consumption growth and the price-dividend ratio. The number of parameters to be estimated is eleven. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c, \Delta d)$	$corr(\Delta c, p/d)$
Data	.019	.022	.475	.012	.113	.183	.567	.135
	(.003)	(.004)	(.149)	(.013)	(.019)	(.208)	(.133)	(.156)
Model	-.001	.033	.317	.008	.170	.044	.493	-.756
	[-.010, .011]	[.027, .034]	[.121, .507]	[-.049, .017]	[.119, .162]	[-.178, .250]	[.067, .452]	[-.921, -.842]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.005	.030	.672	.061	.200	.009	3.393	.455	.862
	(.005)	(.005)	(.121)	(.019)	(.018)	(.217)	(.080)	(.049)	(.054)
Model	.008	.016	.420	.040	.170	.046	3.423	.002	.213
	[-.001, .013]	[.017, .020]	[.284, .655]	[-.023, .043]	[.119, .162]	[-.177, .252]	[3.619, 3.621]	[.002, .002]	[.255, .635]

Parameter Estimates					
γ	ψ	δ	π_1	π_2	$\rho_{c,d}$
14.226	.629	.973	.731	.690	.451
(.0005)	(.0148)	(.2177)	(.0359)	(.2445)	(.0035)
$\mu_{c,1}$	$\mu_{c,2}$	$\mu_{d,1}$	$\mu_{d,2}$	σ_c	σ_d
.026	-.032	.060	-.051	.016	.161
(.0918)	(.0967)	(.0775)	(.0943)	(.0455)	(.0424)

Table 3: Model Fit and Parameter Estimates with Learning from both Consumption and the Change in CPI, 1964-2013

The table reports estimation results and model fit for the main model with consumption growth and the change in the CPI for all urban consumers (CPI-U) as signals, using annual data over the entire available sample period 1964-2013. The parameters are estimated using GMM. Twenty-one moment restrictions are used in the GMM, namely the mean, variance, and first-order autocovariance of the consumption and dividend growth rates, the rate of inflation, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, the covariance between consumption growth and the price-dividend ratio, and the covariance between the rate of inflation and the price-dividend ratio. The number of parameters to be estimated is fourteen. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c, \Delta d)$	$corr(\Delta c, p/d)$
Data	.020	.013	.513	.016	.072	.269	.253	-.046
	(.002)	(.002)	(.161)	(.012)	(.010)	(.187)	(.171)	(.147)
Model	.011	.016	.019	.011	.100	.259	.418	.155
	[.006 ,.016]	[.013, .019]	[-.276 ,.266]	[-.083 ,.053]	[.070, .125]	[-.254 ,.519]	[.159, .625]	[-.164, .409]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.014	.019	.680	.056	.179	-.013	3.617	.401	.891
	(.004)	(.002)	(.136)	(.023)	(.022)	(.271)	(.095)	(.047)	(.064)
Model	.014	.002	.594	.041	.205	-.003	3.582	.439	.919
	[.012, .015]	[.000, .002]	[-.087, .814]	[-.030, .081]	[.086, .366]	[-.437, .363]	[2.823, 3.771]	[.010, .637]	[-.082, .942]

Inflation Growth (ΔCPI)				
	$E[\Delta CPI]$	$\sigma(\Delta CPI)$	AC1(ΔCPI)	corr(ΔCPI , p/d)
Data	.041	.027	.778	-.675
	(.006)	(.005)	(.115)	(.064)
Model	.023	.049	.425	-.726
	[.001,.083]	[.030,.065]	[-.248,.662]	[-.860,-.416]

Parameter Estimates							
γ	ψ	δ	π_1	π_2	$\mu_{c,1}$	$\mu_{c,2}$	$\rho_{c,d}$
14.834	1.999	.990	.989	.933	.012	.006	.407
(.0000)	(.0000)	(.0028)	(.0035)	(.0023)	(.0020)	(.0038)	(.0000)
$\mu_{d,1}$	$\mu_{d,2}$	$\mu_{\Delta CPI,1}$	$\mu_{\Delta CPI,2}$	σ_c	σ_d	$\sigma_{\Delta CPI}$	
.032	-.122	.010	.107	.016	.085	.036	
(.0263)	(.0077)	(.0050)	(.0045)	(.0095)	(.0002)	(.0040)	

Table 4: Model Fit and Parameter Estimates with Learning only from Consumption, 1964-2013

The table reports estimation results and model fit for the alternative model using annual data over the entire available sample period 1964-2013. The parameters are estimated using GMM. Seventeen moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, and the covariance between consumption growth and the price-dividend ratio. The number of parameters to be estimated is eleven. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c, \Delta d)$	$corr(\Delta c, p/d)$
Data	.020	.013	.513	.016	.072	.269	.253	-.046
	(.002)	(.002)	(.161)	(.012)	(.010)	(.187)	(.171)	(.147)
Model	-.003	.017	.481	.002	.122	.047	.459	-.824
	[-.010, .011]	[.027, .034]	[.121, .507]	[-.049, .017]	[.119, .162]	[-.178, .250]	[.067, .452]	[-.921, -.842]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.014	.019	.680	.056	.179	-.013	3.617	.401	.891
	(.004)	(.002)	(.136)	(.023)	(.022)	(.271)	(.095)	(.047)	(.064)
Model	.016	.017	.801	.024	.121	.050	3.829	.007	.782
	[-.001, .013]	[.017, .020]	[.284, .655]	[-.023, .043]	[.119, .162]	[-.177, .252]	[3.619, 3.621]	[.002, .002]	[.255, .635]

Parameter Estimates					
γ	ψ	δ	π_1	π_2	$\rho_{c,d}$
14.837	.542	.976	.843	.958	.444
(.0002)	(.0166)	(.1622)	(.0269)	(.0727)	(.0059)
$\mu_{c,1}$	$\mu_{c,2}$	$\mu_{d,1}$	$\mu_{d,2}$	σ_c	σ_d
.022	-.010	.059	-.013	.011	.118
(.4060)	(.1012)	(.2300)	(.0668)	(.1959)	(.0842)

Table 5: Model Fit and Parameter Estimates with Learning from both Consumption and the Change in Hourly Earnings, 1964-2013

The table reports estimation results and model fit for the main model, with consumption growth and the change in average hourly earnings of production and nonsupervisory employees on private nonfarm payrolls in the manufacturing sector (ΔEAR) as signals, using annual data over the entire available sample period 1964-2013. The parameters are estimated using GMM. Twenty-one moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, the change in the average hourly earnings, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, the covariance between consumption growth and the price-dividend ratio, and the covariance between the change in the average hourly earnings and the price-dividend ratio. The number of parameters to be estimated is fourteen. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c, \Delta d)$	$corr(\Delta c, p/d)$
Data	.020	.013	.513	.016	.072	.269	.253	-.046
	(.002)	(.002)	(.161)	(.012)	(.010)	(.187)	(.171)	(.147)
Model	.011	.016	.019	.015	.096	.241	.410	.157
	[.005 ,.016]	[.013, .019]	[-.278 ,.265]	[-.075 ,.055]	[.068, .118]	[-.255 ,.492]	[.148, .618]	[-.152, .410]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.014	.019	.680	.056	.179	-.013	3.617	.401	.891
	(.004)	(.002)	(.136)	(.023)	(.022)	(.271)	(.095)	(.047)	(.064)
Model	.014	.002	.602	.042	.200	.001	3.692	.443	.924
	[.012, .015]	[.000, .002]	[-.074, .790]	[-.029, .082]	[.090, .347]	[-.433, .376]	[2.901, 3.893]	[.019, .631]	[-.070, .945]

Hourly Earnings Growth (ΔEAR)				
	$E[\Delta EAR]$	$\sigma(\Delta EAR)$	AC1(ΔEAR)	corr(ΔEAR , p/d)
Data	.042	.024	.902	-.642
	(.006)	(.005)	(.115)	(.064)
Model	.033	.027	.358	-.685
	[.021,.064]	[.017,.034]	[-.251,.605]	[-.824,-.440]

Parameter Estimates							
γ	ψ	Δ	π_1	π_2	$\mu_{c,1}$	$\mu_{c,2}$	$\rho_{c,d}$
14.988	1.966	.989	.990	.940	.012	.006	.396
(.0000)	(.0000)	(.0029)	(.0053)	(.0061)	(.0083)	(.0075)	(.0001)
$\mu_{d,1}$	$\mu_{d,2}$	$\mu_{\Delta EAR,1}$	$\mu_{\Delta EAR,2}$	σ_c	σ_d	$\sigma_{\Delta EAR}$	
.035	-.104	.026	.073	.016	.082	.021	
(.0250)	(.0154)	(.0065)	(.0161)	(.0137)	(.0024)	(.0060)	

Table 6: Model Fit and Parameter Estimates with Learning from both Consumption and the 2nd Principal Component of Changes in Macroeconomic Variables, 1964-2013

The table reports estimation results and model fit for the main model, using consumption growth and the 2nd principal component of changes in macroeconomic variables $(MACRO)$ as signals over the period 1966-2011. The parameters are estimated using GMM. Twenty-one moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, the change in the average hourly earnings, market return, risk free rate and the market-wide price-dividend ratio, the covariance between consumption and dividend growth, the covariance between consumption growth and the price-dividend ratio, and the covariance between the change in the average hourly earnings and the price-dividend ratio. The number of parameters to be estimated is fourteen. Panel C presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A and B present the sample moments (with standard errors in parentheses below) and the corresponding model-implied moments (with simulated 95% confidence intervals in square brackets below) for the consumption and dividend growth rates (Panel A) and asset prices and returns (Panel B). The confidence intervals are obtained as the 5th and 95th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends								
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$corr(\Delta c,\Delta d)$	$corr(\Delta c, p/d)$
Data	.019	.013	.450	.010	.067	.270	.323	-.021
	(.002)	(.002)	(.181)	(.011)	(.010)	(.197)	(.139)	(.156)
Model	.008	.018	.020	.010	.103	.205	.425	.148
	[.002 ,.014]	[.014, .021]	[-.288 ,.261]	[-.056 ,.070]	[.075, .119]	[-.252 ,.401]	[.147, .638]	[-.230, .389]

Prices									
	$E[r_f]$	$\sigma(r_f)$	$AC1(r_f)$	$E[r_m]$	$\sigma(r_m)$	$AC1(r_m)$	$E[p/d]$	$\sigma(p/d)$	$AC1(p/d)$
Data	.015	.018	.639	.046	.183	-.018	3.610	.415	.896
	(.004)	(.002)	(.150)	(.025)	(.023)	(.281)	(.102)	(.048)	(.064)
Model	.023	.003	.952	.045	.171	.020	3.475	.441	.952
	[.020, .026]	[.000, .003]	[-.065, .935]	[-.014, .102]	[.095, .284]	[-.291, .211]	[2.975, 3.865]	[.000, .449]	[-.064, .935]

Second Principal Component (*MACRO*)				
	$E[MACRO]$	$\sigma(MACRO)$	AC1(*MACRO*)	corr(*MACRO*, p/d)
Data	.000	1.011	.758	.703
	(.231)	(.153)	(.120)	(.065)
Model	-.037	1.216	.849	.944
	[-1.291,1.140]	[.342,1.289]	[-.229,.859]	[.343,.960]

Parameter Estimates							
γ	ψ	Δ	π_1	π_2	$\mu_{c,1}$	$\mu_{c,2}$	$\rho_{c,d}$
14.789	0.847	.990	.978	.974	.010	.005	.408
(.0001)	(.0008)	(.0494)	(.0196)	(.0285)	(.0215)	(.0240)	(.0212)
$\mu_{d,1}$	$\mu_{d,2}$	$\mu_{MACRO,1}$	$\mu_{MACRO,2}$	σ_c	σ_d	σ_{MACRO}	
.053	-.042	1.071	-1.234	.017	.091	.405	
(.0439)	(.0260)	(.0706)	(.1237)	(.0084)	(.0015)	(.0748)	

Table 7: Predictive Regressions, Annual Data 1964-2013

The table presents model-implied predictive regressions for the aggregate consumption growth, dividend growth and market return. The model implies that the expected aggregate consumption and dividend growth rates are linear functions of the probability of being in the first regime p_t, while the expected market return is a highly nonlinear function of p_t. Panels A and B present results for the main model, while Panels C and D present results for the alternative model.

Main Model						
	Panel A: Data			Panel B: Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Intercept	.017 (.013)	-.082 (.069)	.173 (.955)	.006*** [-.113, .150]	-.109*** [-.698. .777]	.028*** [-24.12, 23.95]
Lag p	.002 (.015)	.111 (.078)	.055 (2.440)	.006*** [-.139, .127]	.140*** [-.749, .731]	1.063*** [-48.97, 54.06]
Lag p^2			-.213 (1.518)			-1.059*** [-29.78, 24.95]
R^2	.000	.042	.044	.014 [.000,.135]	.193 [.000,.442]	.033 [.006,.529]

Alternative Model						
	Panel C: Data			Panel D: Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Intercept	.018*** (.003)	.028* (.015)	.084** (.039)	-.009*** [-.012, -.003]	-.010*** [-.046, .032]	.010*** [-.028, .053]
Lag p	.006 (.006)	-.041 (.033)	-.452* (.253)	.027*** [-.024, .035]	.061*** [-.239, .227]	.116*** [-.932, 1.032]
Lag p2			.574* (.305)			-.053** [-2.746, 2.652]
R^2	.018	.032	.071	.364 [.001,.567]	.039 [.000,.178]	.044 [.002,.217]

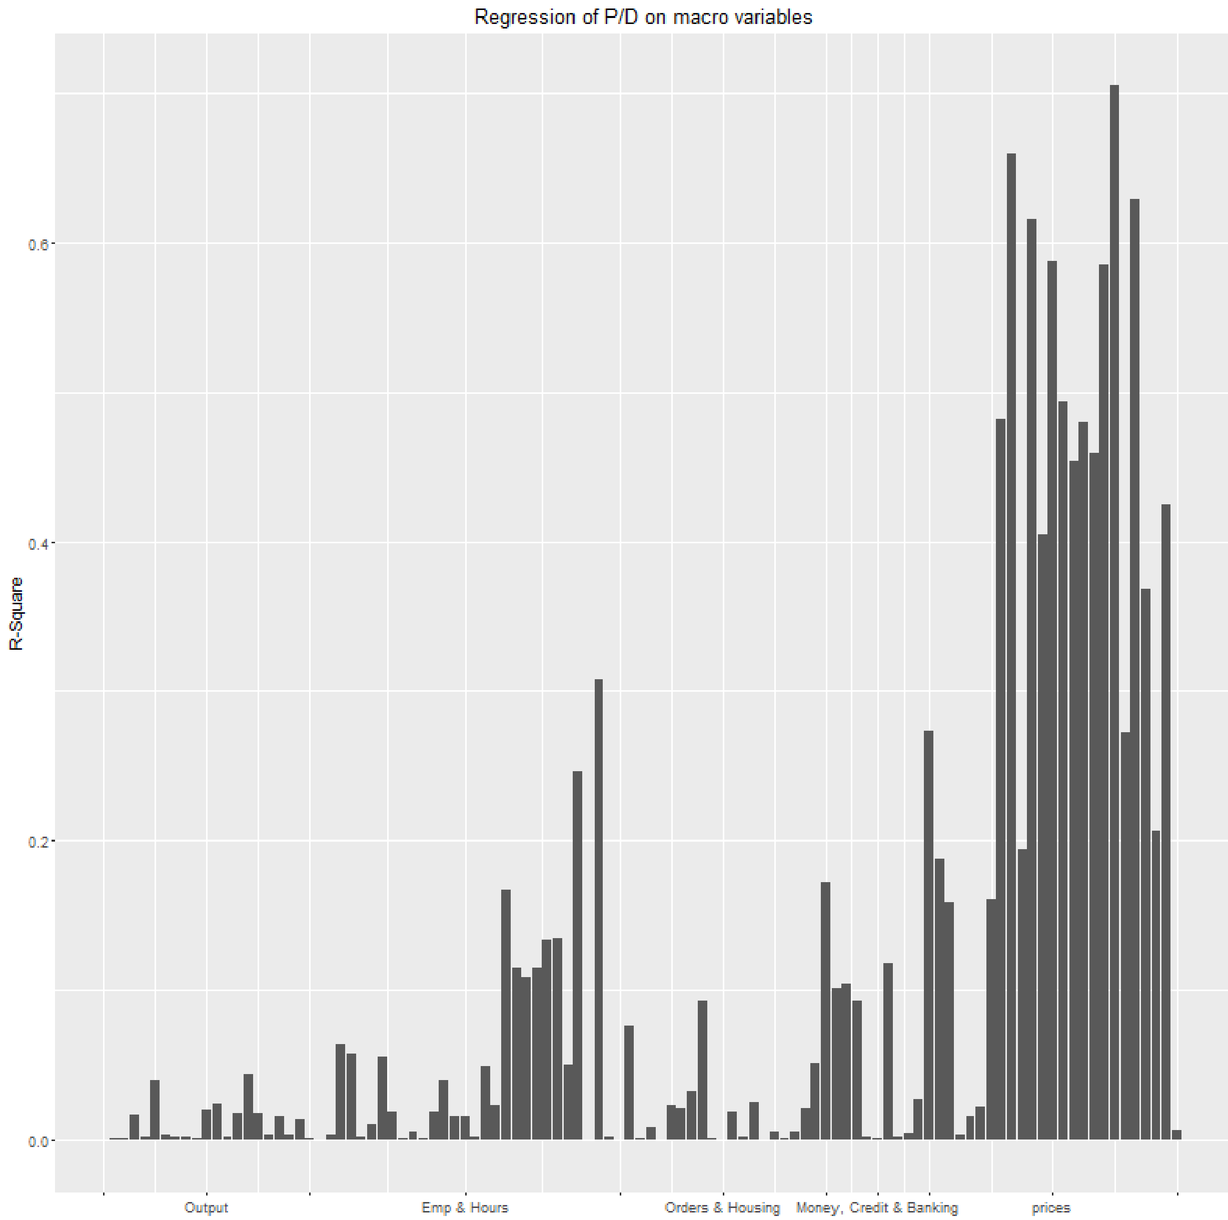

Figure 1: The R^2 in univariate regressions of the time series of the price-dividend ratio on each of the 106 macroeconomic variables over the 1964-2011 sample period.

Figure 2: The time series of macroeconomic variables over the 1964-2011 sample period.

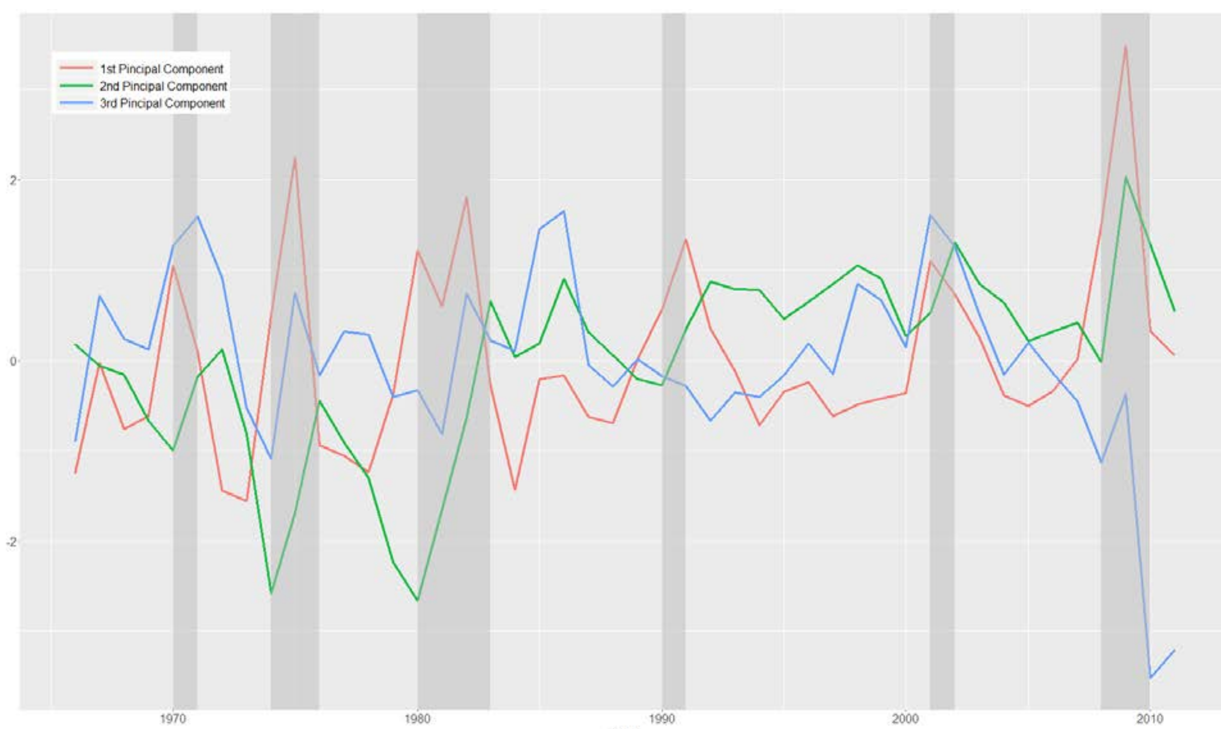

Figure 3: The time series of the first three principal components of the 106 macroeconomic variables over the 1964-2011 sample period.

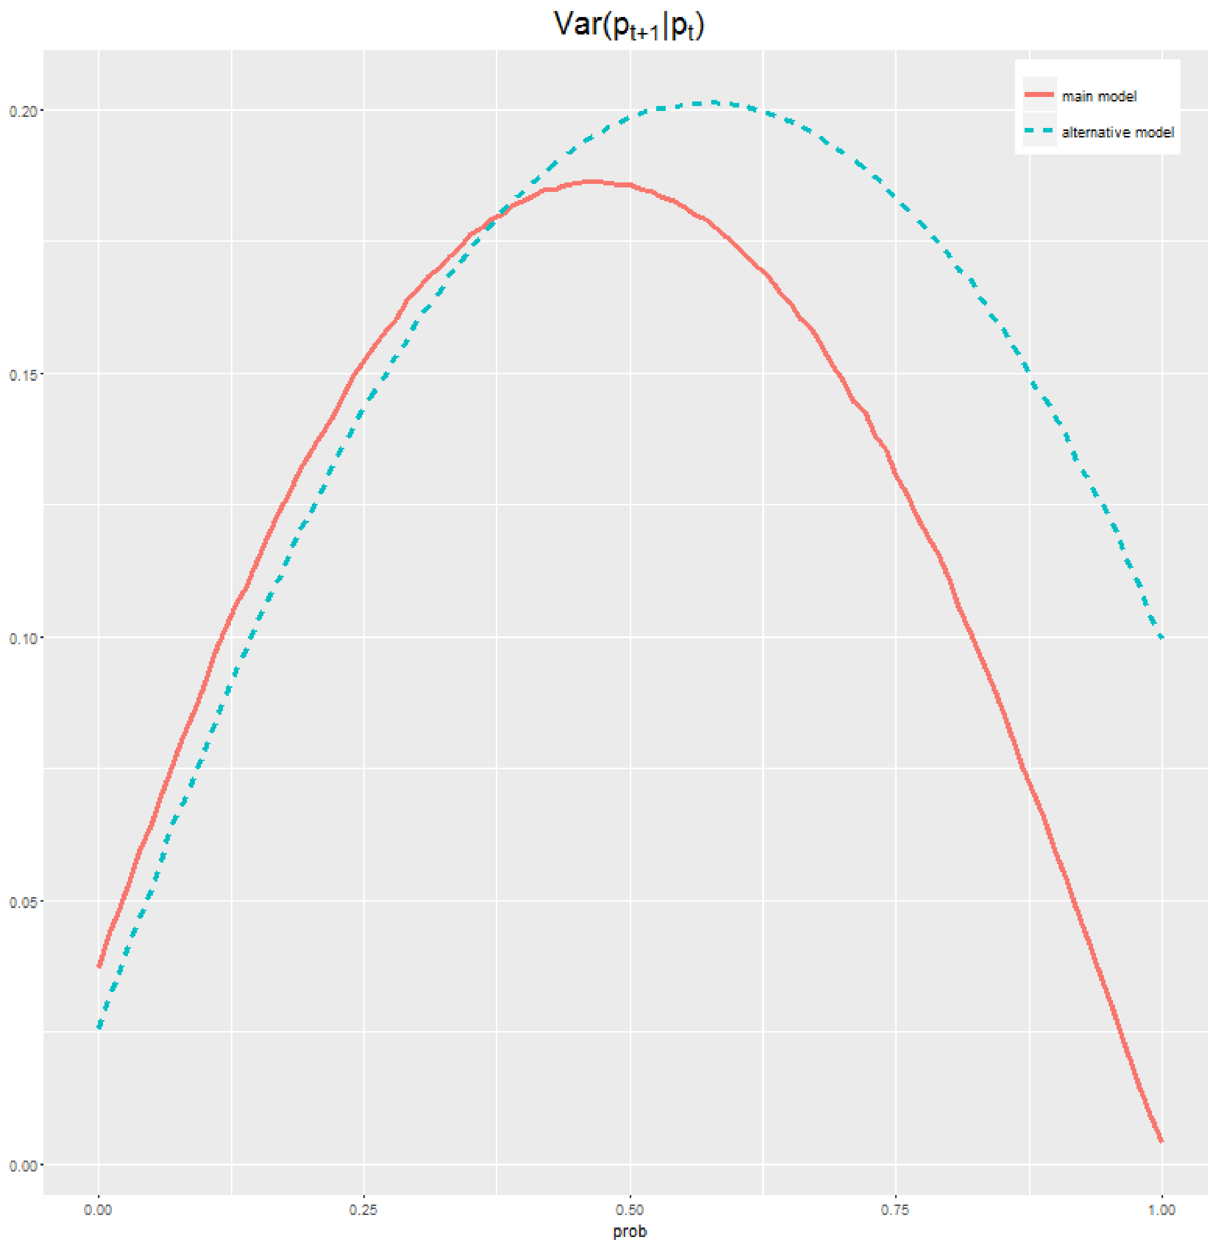

Figure 4: The conditional variance of the beliefs, $\operatorname{var}\left(p_{t+1} \mid p_t\right)$, as a function of the beliefs p_t in the main model (red line) and the alternative model (blue line). The conditional variance is computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1964-2013.

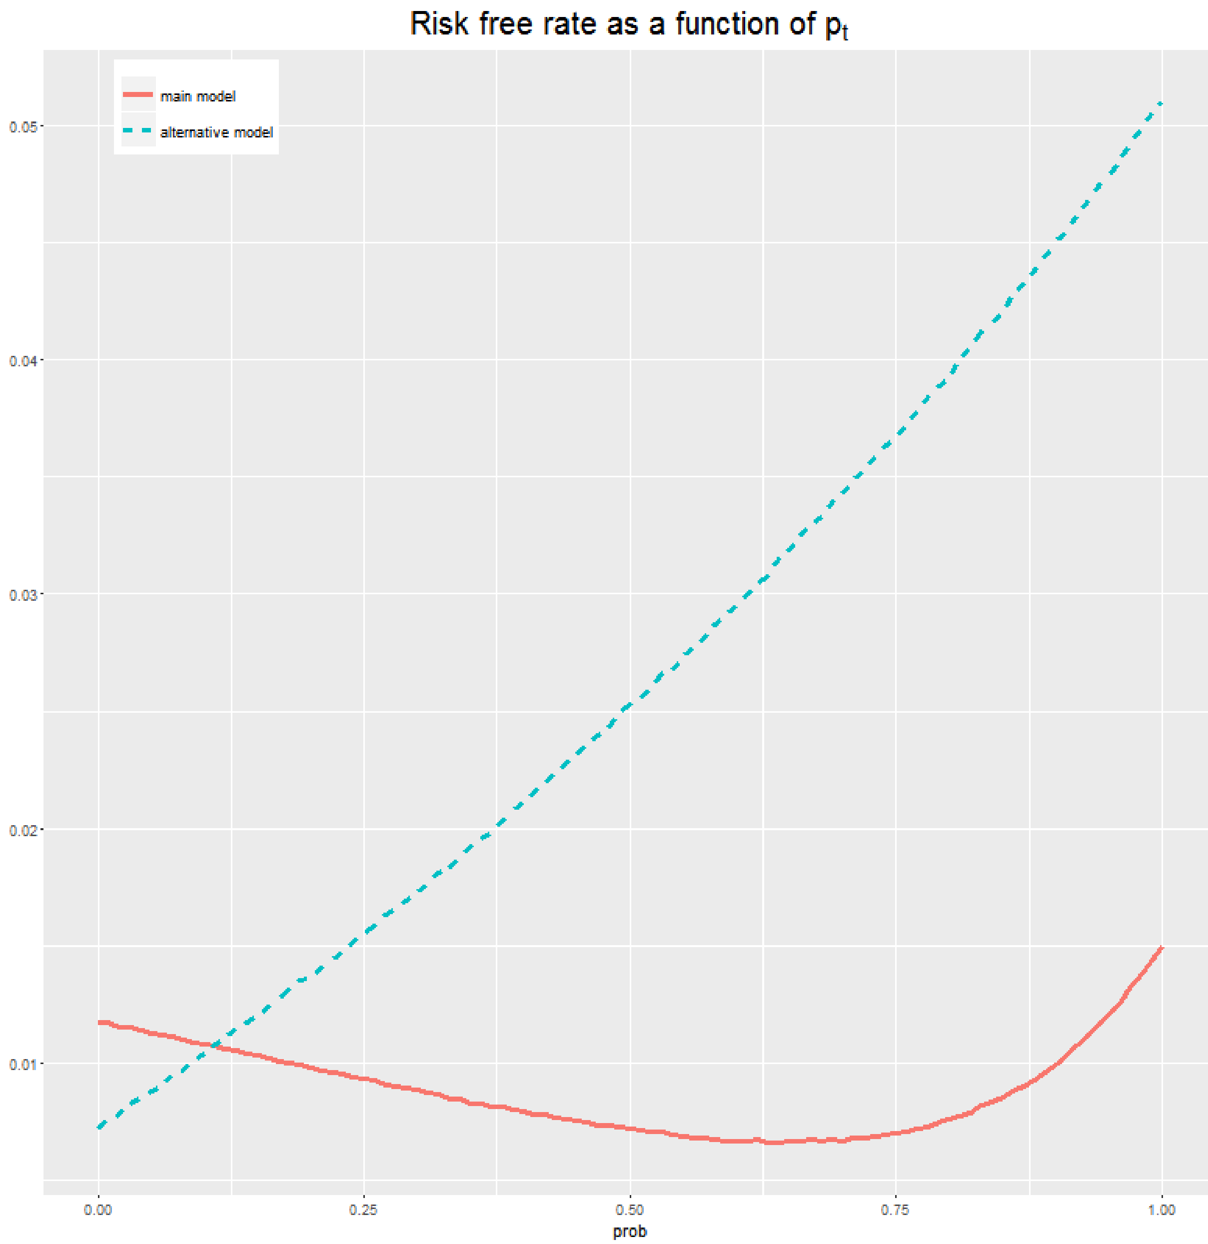

Figure 5: The risk free rate as a function of the beliefs in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1929-2013.

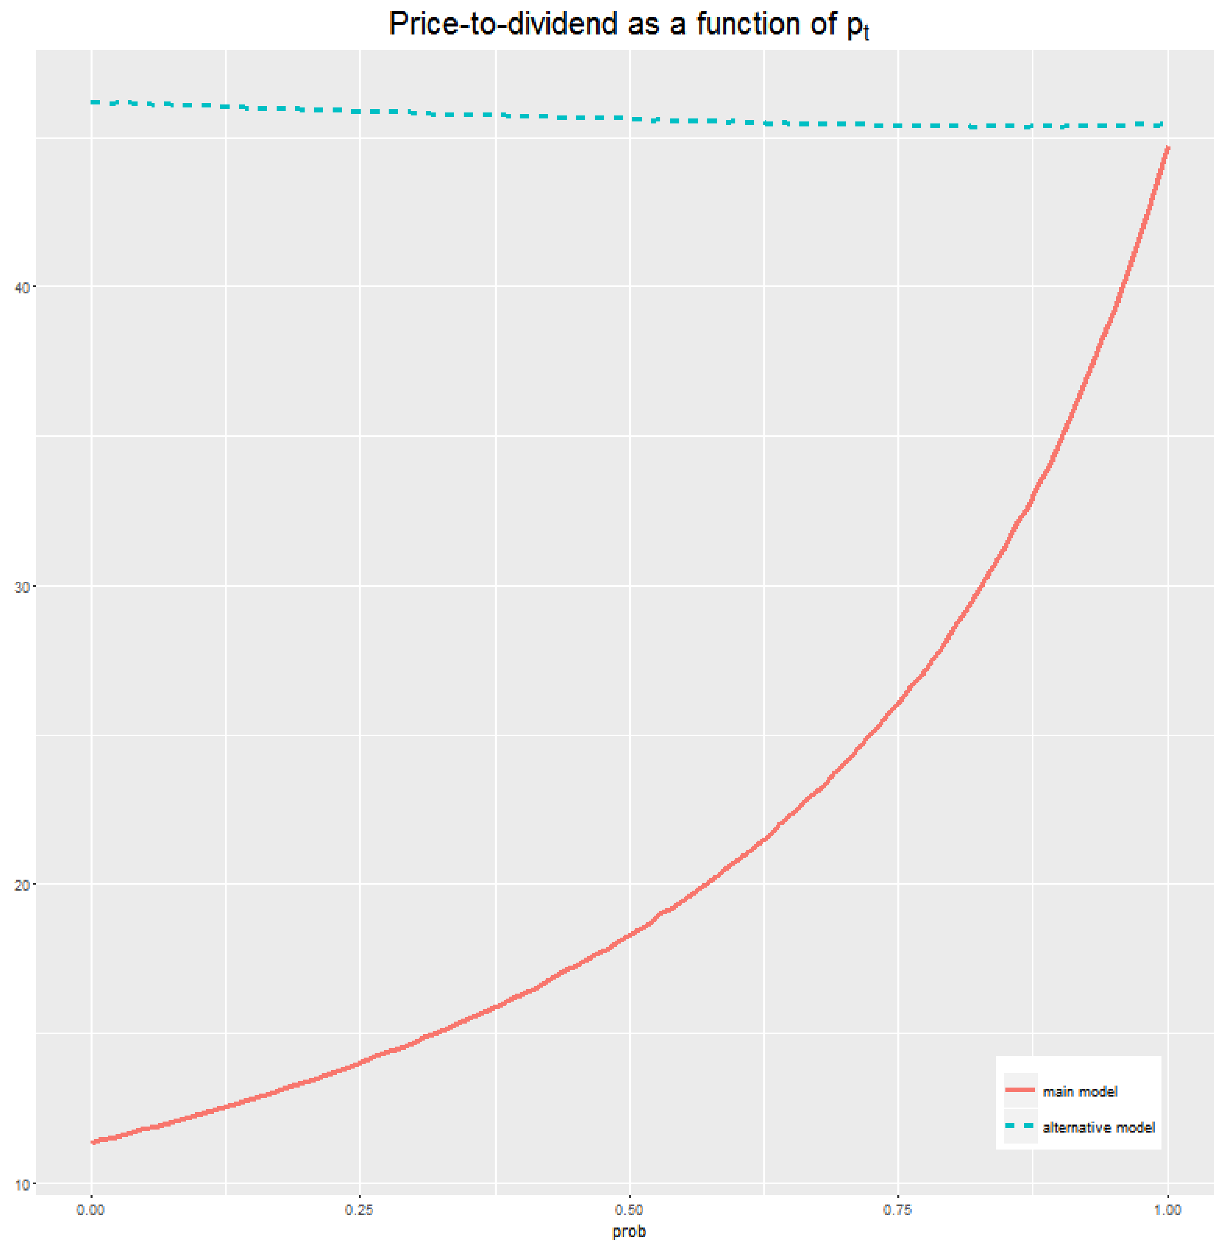

Figure 6: The price-dividend ratio as a function of the beliefs in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1964-2013.

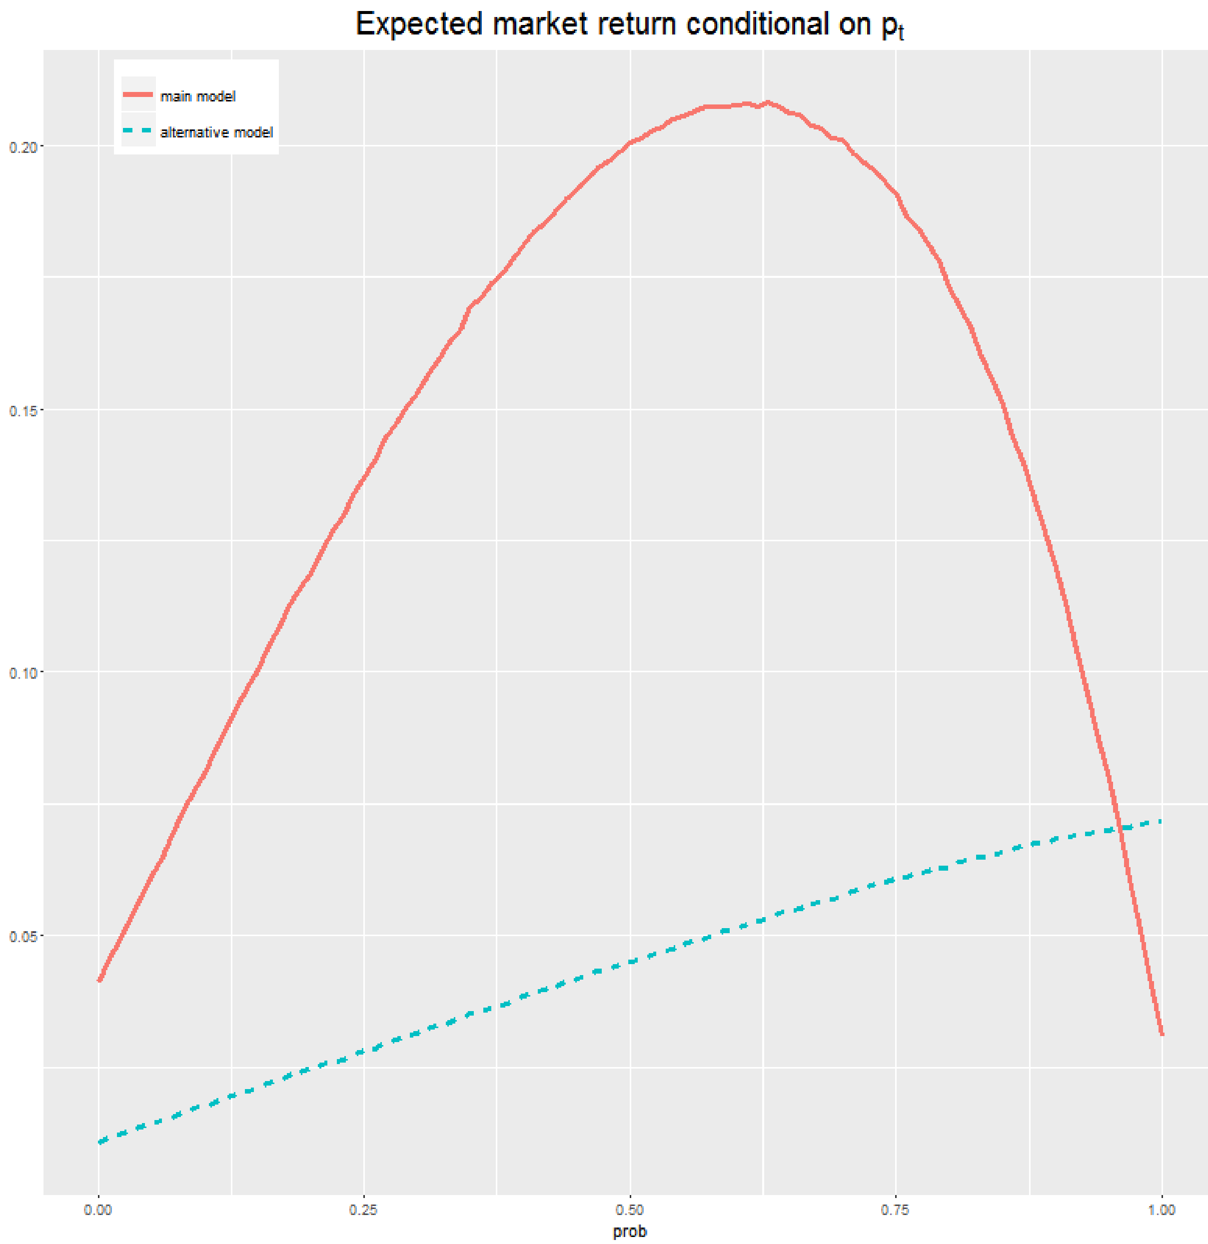

Figure 7: The expected market return as a function of the beliefs in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1964-2013.

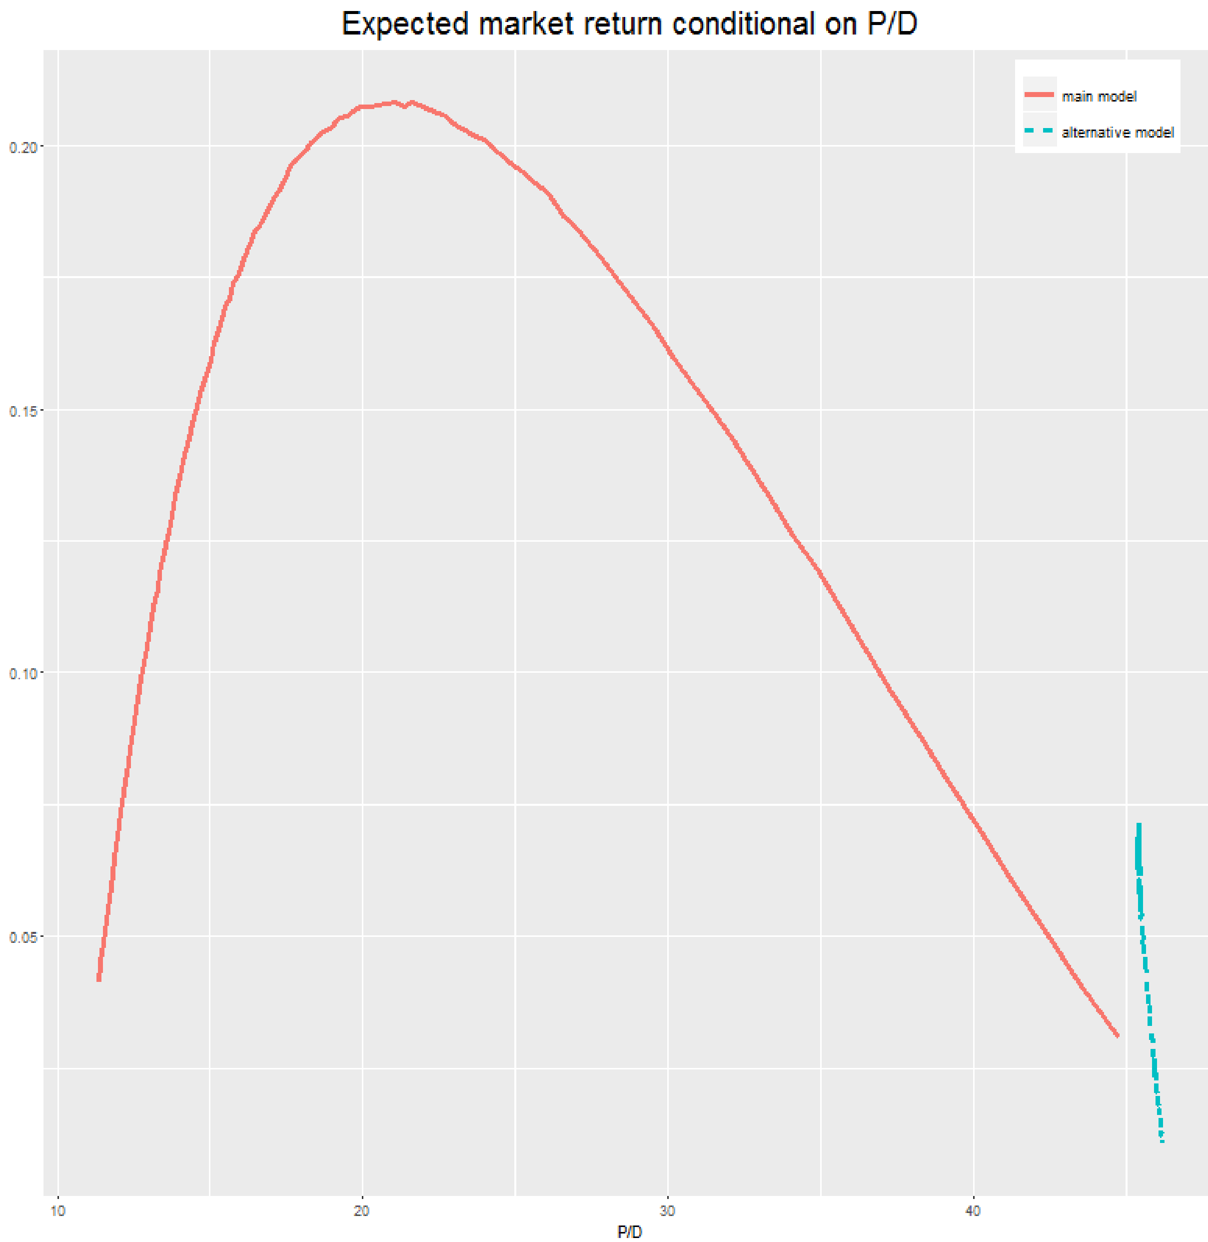

Figure 8: The expected market return as a function of the price-dividend ratio in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1964-2013.

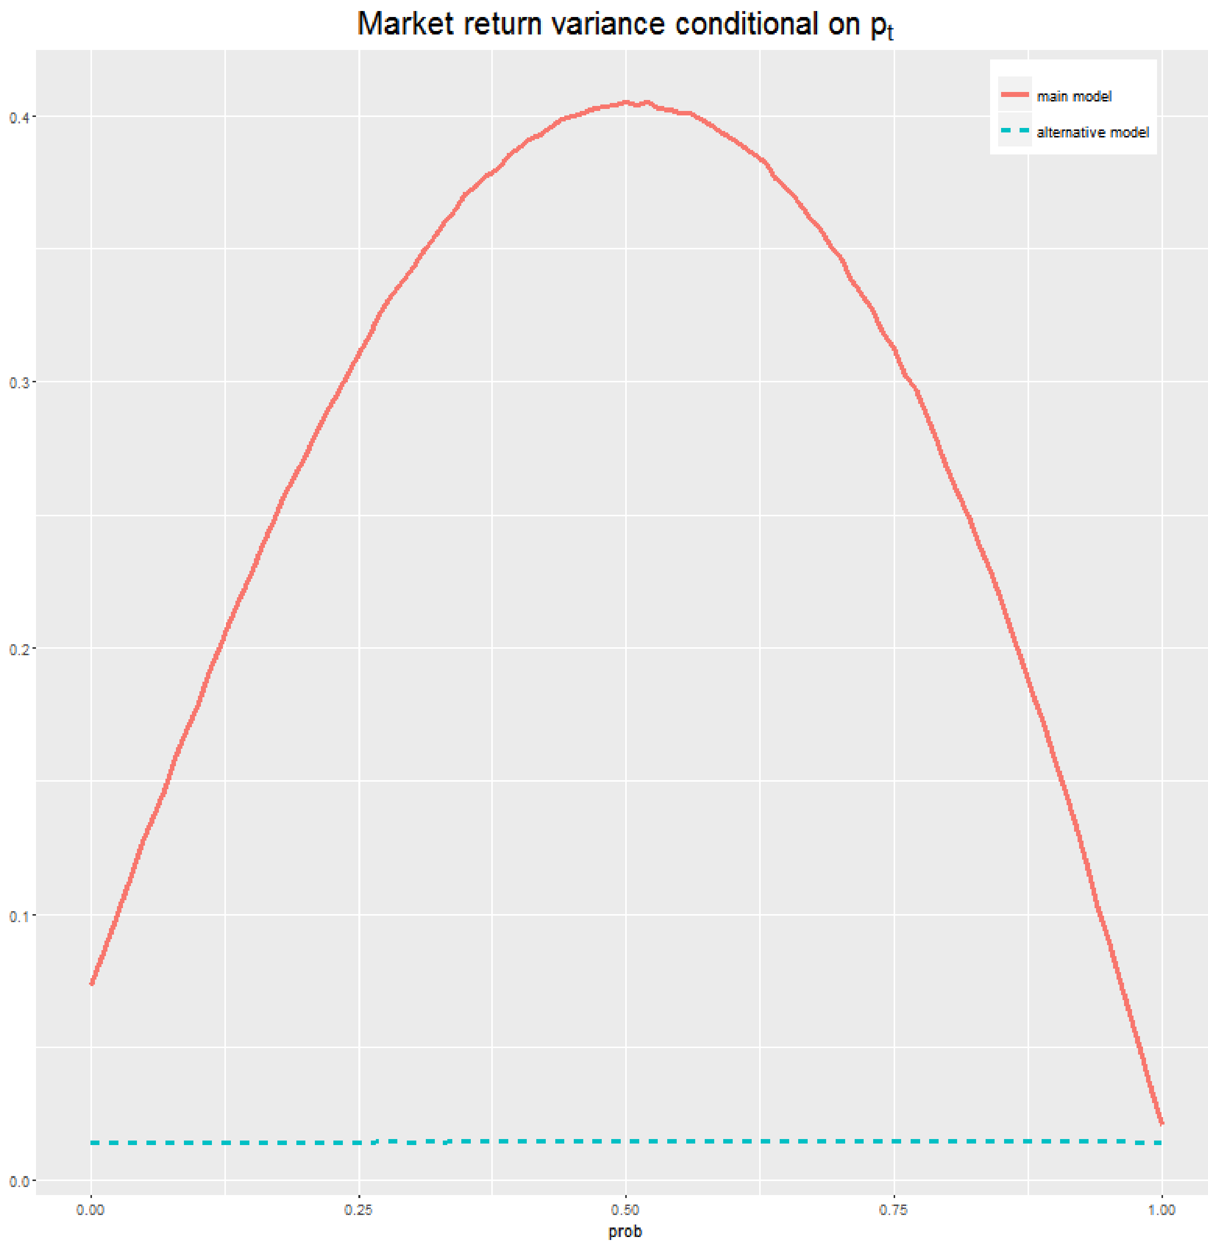

Figure 9: The conditional variance of the market return as a function of the beliefs in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 3 for the main model and in Table 4 for the alternative model. The sample is annual over the sample period 1929-2013.

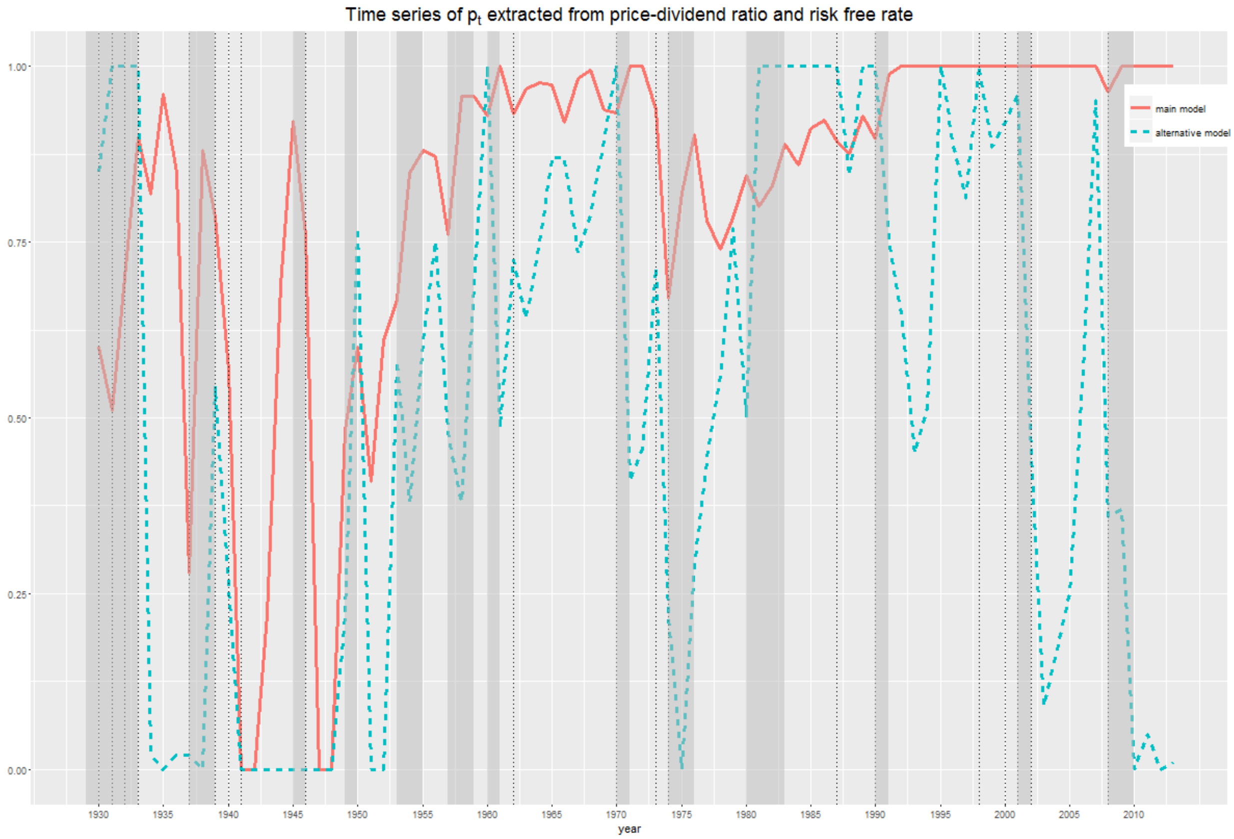

Figure 10: The time series of the beliefs, p_t , extracted from the observed price-dividend ratio and risk free rate, using the point estimates of the model parameters in Table 3 for the main model (red line) and in Table 4 for the alternative model (blue line). The sample is annual over the entire available sample period 1929-2013.