

What Information Drives Asset Prices?

Anisha Ghosh

McGill University

anisha.ghosh@mcgill.ca

George M. Constantinides

University of Chicago and NBER

gmc@chicagobooth.edu

Abstract

The market price-dividend ratio is highly correlated with inflation and labor market variables but not with aggregate consumption and GDP. We build a model with learning from inflation, or earnings, or a combination thereof. The estimated model rationalizes the moments of consumption and dividend growth, market return, price-dividend ratio, and real and nominal term structures and the low predictive power of the price-dividend ratio for consumption and dividend growth 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 the signal.

October 19, 2017

Keywords: macroeconomic news; learning models; CPI signal; earnings per hour signal; asset pricing

JEL classification: E3, G12, G14

We thank Rui Albuquerque, Daniel Andrei, Alexander David, Bruce Grundy, Lars Hansen, Zhiguo He, John Heaton, Christopher Hrdlicka, 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 University of Washington, the London School of Economics, McGill University, Oxford University, and the 2017 WFA meeting for their helpful advice and feedback. We thank Yiran Fan for outstanding research assistance and Chuck Burke for technical assistance. Constantinides received financial support from the Center for Research in Security Prices, the University of Chicago, as trustee/director of the DFA group of funds and SW7 Holdings and as member of the advisory board of FTSE Russell and the Cook County Illinois Investment Policy Committee.

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 from the history of aggregate consumption and GDP alone. These models fare poorly in explaining several features of stock market data, including the high equity premium, low risk free rate, high variability of the price-dividend ratio, and low predictability of consumption and dividend growth by the price-dividend ratio.

In this paper we show that expanding the information set of investors to include macroeconomic variables in addition to consumption and GDP plays a central role in improving the performance of this class of models. Reliable identification of the information set of investors is important for a host of applications, in addition to the application considered in this paper. 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, e.g., Hansen and Richard (1987)). Therefore, if at all possible, summarizing the investors' information set with a small-dimensional set of variables is useful. Also recent research highlights how investors' subjective beliefs, or belief distortions relative to the rational expectations benchmark, can be extracted from observed asset prices via the Euler equations of consumption (see, e.g., 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 investors' information set. Our paper suggests that just one macroeconomic variable—a linear combination of over one hundred macroeconomic variables that loads heavily on inflation and labor market variables (the second principal component, or $2^{nd} PC$), along with consumption growth, goes a long way towards proxying for investors' relevant information set.

We find that the broad category of macroeconomic information that is most highly correlated with the market-wide price-dividend ratio consists of inflation variables, including the rate of change in the Consumer Price Index for all Urban Consumers (CPI-U). Labor market variables, including the change in average hourly earnings and average hours of production in private non-farm payrolls in different sectors, constitute the second category of macroeconomic variables most correlated with the price-dividend ratio. These are also the two classes of macroeconomic variables that, according to FactSet, Bloomberg users pay the most attention to (see, Ai and Bansal (2016)). The high correlation between the price-dividend ratio and the above two classes of macroeconomic variables is not just a feature of U.S. data but is, in fact, also observed in most of the G7 countries, as we discuss in Section 5.

However there is some variability in the correlations between the price-dividend ratio and these macroeconomic variables across subperiods, with the correlations flipping signs in certain subperiods. This is because CPI-U growth and growth in earnings per hour are unreliable *univariate* signals regarding expectations of economic growth and the resulting price-dividend ratio. Expectations of high economic growth lead to wage pressure and inflation depending on productivity growth, tightness of the labor supply, and the ability of firms to automate production, push employees to work longer hours and harder, or outsource work to foreign contractors.

To the contrary we find that the 2^{nd} PC extracted from a broad cross section of over one hundred macroeconomic variables (that loads most heavily on inflation and labor market variables) has consistently high positive correlation with the price-dividend ratio in all subsamples. Moreover the high correlation also obtains out-of-sample. It appears that the 2^{nd} PC blends signals from inflation, the labor market, and other macroeconomic variables in a way that captures the economic conditions under which expectations of high economic growth lead to wage pressure and inflation.

On the other hand the price-dividend ratio has negligible correlation with the contemporaneous consumption and GDP growth as well as moving averages of current and lagged growth rates. This is contrary to the implications of learning from consumption and GDP histories alone where investors' beliefs and, therefore, equilibrium asset prices are functions of the consumption and GDP histories alone. Once again this feature of the data is observed not only in the United States but also in most G7 countries.

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 consumption and GDP histories alone. Therefore, in this paper, we consider a learning process where investors learn not only from the history of consumption but also from an additional signal. To shed light on the potential sources of the signal, we set the signal equal to judiciously chosen macroeconomic variables: a combination of macroeconomic variables (2nd PC) highly correlated with the price-dividend ratio, or CPI-U growth, or change in earnings per hour.

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 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 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 additional signal with innovations orthogonal to those of aggregate consumption and dividend growth. 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 and volatility of the market return and 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 and dividend growth, the additional macro signal, 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 consumption history alone fails along a number of dimensions. It implies essentially zero volatility of the price-dividend ratio, grossly at odds with its sample counterpart of 0.40. It, therefore, equates the volatility of the market return with that of the dividend growth thereby failing to explain the excess volatility puzzle. The model-implied autocorrelation of the price-dividend ratio is only 0.16, at odds with the autocorrelation of 0.89 observed in the data. Thus the model fails to generate the high persistence in the market-wide price-dividend ratio—one of the most robust features observed in the data. The model implies much higher correlation between the consumption growth rate and the innovation in the risk-free rate than that observed in the data (0.47 versus -0.01). Finally, since the model-implied price-dividend ratio is essentially constant, the model counterfactually implies a pro-

cyclical expected market return (driven by pro-cyclical expected dividend growth) and flat conditional volatility of the market return.

Not only does the main 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 0.77 with the historical time series when the signal is set equal to the 2nd PC of the macro variables and correlation 0.73 with the historical time series when the signal is set equal to the innovation in the CPI-U. 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.

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. The question, therefore, arises as to what drives the superior performance of the main model. The crux of the intuition lies in the contrast between the properties of the state and beliefs processes in the two models. In the main model the state and the beliefs processes have high persistence and low variance while in the alternative model these processes have lower persistence and higher 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. The current beliefs are also 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. Furthermore, the high persistence in the beliefs process, combined with preference for early resolution of uncertainty, yields a high equity premium and low risk free rate, consistent with the data.

In the alternative model, on the other hand, beliefs are less persistent and, therefore, current beliefs generally 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.

Also 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 heteroscedasticity of the consumption growth rate or the additional macro signal (the volatilities of consumption growth and the signal are set 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. In contrast the long-run risks model of Bansal and Yaron (2004) critically relies on the heteroscedasticity of consumption growth in order to generate time-variation in the equity premium. Moreover, Constantinides and Ghosh (2011) argue that, contrary to the model's implications, the conditional variance of expected consumption growth is fairly flat and fails to capture the large time variation in the equity premium observed in the data.

We further assess the empirical plausibility of the main model 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 and dividend growth rates, consistent with the data. This also contrasts with the long-run risks model that implies orders of magnitude higher R^2 in forecasting regressions of the short and long horizon consumption and dividend growth rates on the price-dividend ratio. The superior performance of the main model is rendered possible because it does not rely on the persistence of consumption and dividend growth to generate key asset pricing results, relying instead on the persistence of other macro variables in the information set of investors.

Finally we compare the average nominal and real yields on Treasury bonds with their model-implied counterparts. This is a challenging test because real yields beyond one year and nominal yields of any maturity are not targeted in the estimation of the model parameters. For all maturities the model-implied nominal yields closely match the data. Also for all maturities the model-implied real yields closely match the data, a feat that eludes many alternative models.

Note that much of the existing asset pricing literature relies on latent and hard-to-measure state variables (slow moving expected consumption growth in long run risks models, consumption

habits in habit formation models, and expected size and frequency of disasters in rare-disaster models). Our paper represents an advance in this regard in that it proposes that investors learn about the state of the economy from information proxied by observable macroeconomic variables and demonstrates that the model performance fares favorably relative to these alternative paradigms.

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), Lettau, Ludvigson, and Wachter (2008), 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 in the presence of money illusion and address the changing sign of the correlation 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 the discount rate.

The paper also draws on the empirical evidence in Albuquerque, Eichenbaum, and Rebelo (2016), 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 states, 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 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.

The remainder of the paper is organized as follows. In Section 2 we present the model and derive its pricing implications. We discuss the data in Section 3. In Section 4 we launch a systematic investigation of the sources of macroeconomic information that drive aggregate prices. The empirical methodology and estimation results are presented in Section 5, along with a comparison between the main model and the alternative one. In Section 6 we discuss the economic intuition underlying the results. In Section 7 we present the model implications on the real and nominal term structures. We conclude in Section 8.

2 The Model and Solution

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

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

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

where c_t and d_t are the logarithms of aggregate consumption and stock market dividend, respectively, in period t . The volatilities of consumption growth, σ_c , and dividend growth, σ_d ,

are assumed to be constant across the two regimes. Our modeling choice of constant volatility of consumption and dividend growth in the two states is primarily made to highlight that our key results obtain even in the absence of such heteroscedasticity for which there is limited empirical evidence. 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

$$\begin{bmatrix} \pi_1 & 1-\pi_2 \\ 1-\pi_1 & \pi_2 \end{bmatrix}, \quad (3)$$

where $\pi_i \equiv \text{prob}(s_t = i | s_{t-1} = i)$ and $0 < \pi_i < 1$ for $i=1,2$. The unconditional probability of $s_t = 1$ is $(1-\pi_2)/(2-\pi_1-\pi_2)$ and its expected duration is $(1-\pi_1)^{-1}$ 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 slow economic growth.

We assume that investors do not observe the regime at time t but learn from a history of signals, $\mathbb{F}(t)$. We assume that the investors' history of signals is $\mathbb{F}(t) = \{c_\tau, x_\tau\}_{\tau=-\infty, \dots, t}$, where x_t is a scalar reflecting additional variables over $[t-1, t]$ 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, $\mathbb{F}(t) = \{c_\tau\}_{\tau=-\infty, \dots, t}$. 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 history alone, $\mathbb{F}(t) = \{c_\tau\}_{\tau=-\infty, \dots, t}$, leads to poor results for two reasons. First consumption growth is not very persistent. As we shall see in Sections 5 and 6 it is important that the signal be persistent. Second consumption growth is poorly correlated with the market-wide price-dividend ratio—a regression of the price-dividend ratio on the contemporaneous consumption growth produces a statistically

insignificant slope coefficient and $R^2 = 1.7\%$. Moreover the R^2 remains small when the contemporaneous consumption growth is replaced by a moving average of current and lagged growth rates—the inclusion of five and ten lags produce R^2 of 0.72% and 0.05%, respectively. As we show in Section 4 these are substantially smaller than those obtained with some other macro variables such as inflation variables.

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

$$x_{t+1} = \mu_{x,s_{t+1}} + \sigma_x \varepsilon_{x,t+1}, \quad (4)$$

where $\varepsilon_{x,t+1}$ is *i.i.d.* standard normal and orthogonal to $\varepsilon_{c,t+1}$ and $\varepsilon_{d,t+1}$.

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

$$\mu_i = \begin{bmatrix} \mu_{c,i} \\ \mu_{x,i} \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_c^2 & 0 \\ 0 & \sigma_x^2 \end{bmatrix}. \quad (5)$$

Investors assign probability p_t that the economy is in the first regime at date t :

$$p_t \equiv \text{prob}(s_t = 1 | \mathbb{F}(t)). \quad (6)$$

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

$$\begin{aligned} & g(u_{t+1} | \mathbb{F}(t)) \\ &= \frac{f(p_t)}{2\pi |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(u_{t+1} - \mu_1)^T \Sigma^{-1} (u_{t+1} - \mu_1)\right) + \frac{1-f(p_t)}{2\pi |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(u_{t+1} - \mu_2)^T \Sigma^{-1} (u_{t+1} - \mu_2)\right), \end{aligned} \quad (7)$$

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

$$\begin{aligned} & f(p_t) \equiv \text{prob}(s_{t+1} = 1 | \mathbb{F}(t)) \\ &= \text{prob}(s_{t+1} = 1 | s_t = 1) \times \text{prob}(s_t = 1 | \mathbb{F}(t)) + \text{prob}(s_{t+1} = 1 | s_t = 2) \times \text{prob}(s_t = 2 | \mathbb{F}(t)) \\ &= \pi_1 p_t + (1 - \pi_2)(1 - p_t) \\ &= 1 - \pi_2 + (\pi_1 + \pi_2 - 1) p_t. \end{aligned} \quad (8)$$

Upon observing $(\Delta c_{t+1}, x_{t+1})$ at time $t+1$, the value of p_{t+1} is updated with Bayes' rule as

$$p_{t+1} | (\Delta c_{t+1}, x_{t+1}, \mathbb{F}(t)) = \frac{f(p_t) e^{-\frac{(\Delta c_{t+1} - \mu_{c,1})^2}{2\sigma_c^2} - \frac{(x_{t+1} - \mu_{x,1})^2}{2\sigma_x^2}}}{2\pi\sigma_c\sigma_x g(\Delta c_{t+1}, x_{t+1} | \mathbb{F}(t))}. \quad (9)$$

Therefore the conditional expectation of p_{t+1} is $f(p_t)$:

$$\begin{aligned} & E[p_{t+1} | \mathbb{F}(t)] \\ &= \iint \{p_{t+1} | (\Delta c_{t+1}, x_{t+1}, \mathbb{F}(t))\} g(\Delta c_{t+1}, x_{t+1} | \mathbb{F}(t)) d\Delta c_{t+1} dx_{t+1} \\ &= \iint \frac{f(p_t) e^{-\frac{(\Delta c_{t+1} - \mu_{c,1})^2}{2\sigma_c^2} - \frac{(x_{t+1} - \mu_{x,1})^2}{2\sigma_x^2}}}{2\pi\sigma_c\sigma_x g(\Delta c_{t+1}, x_{t+1} | \mathbb{F}(t))} g(\Delta c_{t+1}, x_{t+1} | \mathbb{F}(t)) d\Delta c_{t+1} dx_{t+1} \\ &= f(p_t) \iint \frac{e^{-\frac{(\Delta c_{t+1} - \mu_{c,1})^2}{2\sigma_c^2} - \frac{(x_{t+1} - \mu_{x,1})^2}{2\sigma_x^2}}}{2\pi\sigma_c\sigma_x} d\Delta c_{t+1} dx_{t+1} \\ &= f(p_t) \end{aligned} \quad (10)$$

and the unconditional mean of p_t is $\bar{p} = \frac{1 - \pi_2}{2 - \pi_1 - \pi_2}$.

Our model nests the version typically assumed in the literature where investors learn from the consumption history alone: setting $\mu_{x,1} = \mu_{x,2}$ 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 consumption history alone.

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

$$U_t = \left\{ (1 - \delta)(C_t)^{1-1/\psi} + \delta \left(E_t \left[(U_{t+1})^{1-\gamma} \right] \right)^{\frac{1-1/\psi}{1-\gamma}} \right\}^{1/(1-1/\psi)}, \quad (11)$$

where δ is the subjective discount factor, γ is the RRA coefficient, ψ is the EIS, and

$\theta \equiv \frac{1-\gamma}{1-1/\psi}$. As shown in Epstein and Zin (1989), the stochastic discount factor (SDF) is

$$SDF_{t+1} = \delta^\theta \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} R_{c,t+1}^{\theta-1}, \quad (12)$$

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

We solve the model numerically with value function iteration. A detailed description of the numerical solution procedure is available in the Online Appendix A1. The obtained highly nonlinear solutions for the equilibrium wealth-consumption and price-dividend ratios and the conditional mean and volatility of the market return highlight the importance of eschewing the Campbell and Shiller (1988) log-linearization or any other form of approximation in solving the model.

3 Description of the Data

For our main empirical results we use US annual data over the entire available sample period from 1964 to 2013. The starting date of 1964 is dictated by the availability of data on a broad cross section of macroeconomic variables that serve as signals in our main model. 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 CPI 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 CPI 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 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.

We obtain annual yields on Treasury bonds from the Federal Reserve Bank of St. Louis website. Nominal yields on one-, three-, five-, and ten-year nominal bonds are available from 1962 to 2013; on two-year bonds from 1976 to 2013; on seven-year bonds from 1969 to 2013; and on 30-year bonds from 1977 to 2013, with intermission from 2003 to 2005. As the empirical counterpart to real risk free bonds, yields on five-, seven-, and ten-year Treasury Inflation Protected Securities (TIPS) are available from 2003 to 2013, and on 20-year TIPS from 2004 to 2013.

The consumption data consists of 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 CPI.

We obtain panel data 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 (2011) argues that the market understands the subsequent revisions but revisions still matter.

4 What Are Potential Signals?

We shed light on potential sources of the signal in our main model over and above aggregate consumption history. In particular we show that a substantial fraction of the variation in the price-dividend ratio is explained by variations in macroeconomic variables but not by the one variable commonly assumed in the literature, namely aggregate consumption growth. Moreover, we show that this feature of the data holds not only in the United States, but also in most G7 countries.

We relate the US market-wide price-dividend ratio to a broad cross-section of publicly available macroeconomic variables. The 106 macroeconomic variables we consider may be broadly classified into the following six categories: (1) output and income, which includes personal

income, industrial production index, and capacity utilization measures; (2) labor markets, which includes the unemployment rate, unemployment insurance claims, average hourly earnings, average hours of production, 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 and inventories, and new orders for different types of 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 PPI and CPI for different goods and services, and the commodity prices index. Data on all these variables are available from 1964 onwards.

The price-dividend ratio is highly persistent but its first-order autocorrelation coefficient, at 0.89, is less than 0.9 from 1964 to 2013. We difference some of the macroeconomic variables to make them stationary. We refer the reader to Sydney Ludvigson's website for a detailed description of the macroeconomic variables and the transformations applied to make them stationary. Each of the transformed variable has first-order autocorrelation coefficient less than 0.9.

We relate the price-dividend ratio to the macroeconomic variables by running univariate regressions of the log price-dividend ratio on the level (if stationary) or the first difference of the 106 macroeconomic variables. In figure 1 we present the R^2 of these regressions. The figure reveals that inflation variables and, to a lesser extent, labor market variables 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. In light of the concern that regressions of one persistent variable (the price-dividend ratio) on other persistent variables (the macroeconomic variables) may produce spurious results, in figure 2 we present the R^2 from univariate regressions of the first difference of the log price-dividend ratio on the first difference of the 106 transformed macroeconomic variables. The R^2 s are lower than in figure 1 but the rankings are the same as in figure 1.

[Figures 1 and 2 here]

The inflation variables include the CPI-U, the PPI, and the implicit price deflator for personal consumption expenditures. The regression of the price-dividend ratio on the CPI-U growth 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 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 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 in 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.

In contrast a regression of the 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. Similar results obtain for the dividend growth rate—contemporaneous dividend growth produces an $R^2 = 0.7\%$, and a moving average of dividend growth with five and ten lags produce $R^2 = 3.0\%$ and 4.8%, respectively.

Next we show that the high correlations between the price-dividend ratio and inflation and labor market variables, and the low correlation between the price-dividend ratio and consumption growth are not artifacts of US data. Over the 1964-2016 period the correlation between the price-

dividend ratio and the growth in the CPI is large and negative in six out of the seven G7 countries – -79.7% in France, -69.7% in the UK, -66.8% in the US, -64.5% in Canada, -55.5% in Germany, and -42.1% in Japan.¹ Similarly the high correlation between the price-dividend ratio and labor market variables is observed in most of the G7 countries. In particular the correlation between the price-dividend ratio and average hourly earnings growth is -68.5% in France, -67.0% in the UK, -63.7% in Canada, -58.4% in Germany, and -54.4% in the US. Finally the correlation between the price-dividend ratio and real per capita consumption growth is small in most G7 countries—13.9% in Japan, 8.4% in the US, 0.9% in the UK, -0.8% in Canada, and -11.9% in France.

Turning back to the US data, the first, second, and third principal components of the 106 transformed macroeconomic variables explain 40.6%, 20.0%, and 8.8%, respectively, of the variation of these macroeconomic variables. We relate the price-dividend ratio to each of the principal components by running univariate regressions of the log price-dividend ratio on each of the first six principal components and obtain R^2 of 0%, 49%, 0%, 12%, 6%, and 1%, respectively. Essentially the second principal component of changes in the macroeconomic variables ($2^{nd} PC$) is highly correlated with the market-wide price-dividend ratio while the other principal components are not. Figure 3 presents the R^2 from univariate regressions of the $2^{nd} PC$ on each of the 106 transformed macroeconomic variables. Not surprisingly the $2^{nd} PC$ loads most heavily on inflation and labor market variables. In figure 4 we present the R^2 from univariate regressions at the quarterly frequency. The R^2 s are very similar to those at the annual frequency.

[Figures 3 and 4 here]

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 5 displays the time series of the log price-dividend ratio, consumption and GDP growth, CPI growth, and hourly earnings growth; and figure 6 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 consumption growth, GDP growth, and the first and third principal components of the

¹ Data for the G7 countries are obtained from the Global Financial Database.

covariance matrix are neither strongly correlated with the price-dividend ratio nor serve as informative signals in the context of our model, as we show in Section 5.

[Figures 5 and 6 here]

The above results suggest that inflation and labor market variables and the 2nd PC (that heavily loads on inflation and labor market variables) play an important role in explaining variation in asset prices. In Table 1 we report regressions of the log price-dividend ratio on CPI-U growth, earnings per hour growth, and the 2nd PC over the full period from 1964 to 2013 (panel A) as well as two non-overlapping subperiods from 1964 to 1988 (panel B) and from 1989 to 2013 (panel C). We also report these regressions in first differences. Over the full sample period as well as both subperiods, these variables have statistically significant explanatory power for the log price-dividend ratio. The last row of panel C presents out-of-sample R^2 s over the latter subperiod from 1989 to 2013.² The out-of-sample R^2 s are large for the CPI-U growth and 2nd PC and, more importantly, similar in magnitude to the in-sample values obtained over the same subperiod.

[Table 1 here]

These results should be interpreted with caution. The CPI-U growth and earnings per hour growth are unreliable *univariate* signals regarding expectations of economic growth and the resulting price-dividend ratio. Expectations of high economic growth lead to wage pressure and inflation depending on productivity growth, tightness of the labor supply, and the ability of firms to automate production, push employees to work longer hours and harder, or outsource work to foreign contractors, as manifested by the flipping of correlation signs in certain subperiods. The correlation between CPI-U growth and the price-dividend ratio is -62%, -50%, and -33% in the subperiods from 1966 to 1985, from 1986 to 2005, and from 2006 to 2011, respectively; and the correlation between change in earnings per hour and the price-dividend ratio is -48%, 15%, and -53% in the subperiods from 1966 to 1985, from 1986 to 2005, and from 2006 to 2011, respectively. While the correlation between CPI-U growth and the price-dividend ratio seems consistently negative, the correlation is close to zero at -6% since the turn of the century. By contrast the correlation between the 2nd PC and the price-dividend ratio is consistently positive in any subperiod – the correlations are 62%, 49%, and

² The out-of-sample R^2 is computed by estimating the regression parameters over the first subperiod from 1964 to 1988 and then using the estimated parameters to obtain the fitted values of the regression in the second subperiod from 1989 to 2013.

37%, respectively, over subperiods 1966 to 1985, 1986 to 2005, and 2006 to 2011; and 19% since 2001 to the present. The same applies to correlations between their growth rates. It appears that the 2^{nd} PC blends signals from inflation, the labor market, and other macroeconomic variables in a way that captures the economic conditions under which expectations of high economic growth lead to wage pressure and inflation. The changing economic conditions are also highlighted in David and Veronesi (2013) who study learning from inflation shocks and address the changing sign of the correlation between stock and bond returns.

We present predictive regressions of 1-, 2-, 3-, and 5-year ahead consumption and dividend growth and market return by the 2^{nd} PC (Table 2, left panels). The rationale for these predictive regressions is as follows. If certain macroeconomic news signals drive the variation in the price-dividend ratio then these signals should forecast future cash flow growth. Table 2 confirms that this is indeed the case. The 2^{nd} PC has statistically significant predictive power for dividend growth one, two, and five years ahead (the coefficient is statistically insignificant for three years ahead dividend growth); and it has statistically insignificant predictive power for consumption growth and market return. The right panel of Table 2 presents the model-implied predictive regressions that produce results very similar to those obtained in the data and are discussed in Section 5. The corresponding predictive regressions using the CPI-U growth and earnings per hour growth as predictor variables are available in the Online Appendix A2 and produce results very similar to those obtained with the 2^{nd} PC .

[Table 2 here]

5 The Information that Drives Asset Prices

5.1 Estimation Methodology

We show that consumption growth is an incomplete signal and that investors' beliefs are driven by some signal(s) in addition to consumption growth. In particular, we demonstrate that the model fits the data quite well when the signal is set equal to the 2^{nd} PC , CPI-U growth, or earnings per hour growth. The results highlight the informational role of macroeconomic variables and suggest

that just one macroeconomic variable 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 $2^{nd} PC$. Thus equation (4) that describes the dynamics of the signal becomes:

$$\left(2^{nd} PC\right)_{t+1} = \mu_{2^{nd} PC, s_{t+1}} + \sigma_{2^{nd} PC} \varepsilon_{2^{nd} PC, t+1}, \quad (13)$$

where $\varepsilon_{2^{nd} PC, t+1}$ is *i.i.d.* standard normal and orthogonal to $\varepsilon_{c, t+1}$ and $\varepsilon_{d, t+1}$. Thus, the model has fifteen parameters: two parameters of the regime transition matrix (π_1 and π_2); three parameter of the signal distribution ($\mu_{2^{nd} PC, 1}$, $\mu_{2^{nd} PC, 2}$, $\sigma_{2^{nd} PC}$); seven parameters of the time-series processes of aggregate consumption and dividend growth ($\mu_{c, 1}$, $\mu_{c, 2}$, $\mu_{d, 1}$, $\mu_{d, 2}$, σ_c , σ_d , ρ); and three preference parameters (δ , γ and ψ). We estimate the parameters using GMM to match the following twenty-two sample moments: the unconditional mean, variance, and first-order autocorrelation of consumption growth, dividend growth, $2^{nd} PC$, market return, market-wide price-dividend ratio, and risk free rate; the correlation between consumption and dividend growth rates; the correlation between the consumption growth rate and the price-dividend ratio; the correlation between the dividend growth rate and the price-dividend ratio; and the correlation between the $2^{nd} PC$ and the price-dividend ratio. Therefore, we have an over-identified system of twenty-two moment restrictions and fifteen parameters.

5.2 Estimation Results

The estimation results are presented in Table 3. In the first three panels, we display the sample moments and the model-generated moments of the consumption and dividend growth rates, $2^{nd} PC$, 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 fifty years each, the same size as the historical sample.

[Table 3 here]

The model does reasonably well in matching the data. In the first panel of Table 3 we see that the model matches the mean (0.014 versus 0.019) and volatility (0.015 versus 0.013) of consumption growth but not its autocorrelation (0.022 versus 0.45). However the model-implied second-, third- and fifth-order autocorrelations are 0.021, 0.020, and 0.018, respectively, in closer agreement with the corresponding data-implied autocorrelations at 0.188, 0.065, and 0.046, respectively. Models that match precisely the observed first-order autocorrelation of consumption growth typically substantially overestimate the higher-order autocorrelations. Moreover it has been argued in the literature that the high first-order autocorrelation is potentially an artifact 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 sensitive to the precise sample period and the measure of consumption used. It takes the value 0.45 over the 1964 to 2013 sample period when nondurables and services consumption is used as the measure of consumption expenditures, while it is close to zero (-0.06) over the longer period from 1890 to 2009 where total consumption is the measure of consumption expenditures. Second if we were to take at face value the first-order autocorrelation of consumption growth in the 1964 to 2013 sample (0.45), we would conclude that consumption growth is an informative signal. In Section 5.3 we re-estimate the model when the investors learn from the consumption history alone and find that this alternative model performs poorly.

The model does less well in matching the mean (0.025 versus 0.010), volatility (0.142 versus 0.067) and autocorrelation (0.038 versus 0.270) of dividend growth, the latter two possibly due to dividend smoothing. The model matches the correlation of consumption and dividend growth (0.345 versus 0.323) and generates the low correlations between consumption growth and the price-dividend ratio and between dividend growth and the price-dividend ratio observed in the data.

In the second panel of Table 3, the model-implied mean risk free rate is 0.018, consistent with the sample value of 0.015. The model also generates the persistence in the risk free rate observed in the data.

The model-implied mean equity return is 0.040, close to the sample value of 0.046. The model-implied volatility of the market return is 0.20, close to the sample value of 0.18.

The autocorrelation of the market return is low both in the data and the model. The model-implied volatility and autocorrelation of the price-dividend ratio are 0.45 and 0.95, respectively, in agreement with their sample values of 0.42 and 0.90, respectively.

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 the volatility of the market-wide price-dividend ratio, thereby partly accounting for the excess volatility puzzle.

In the third panel the model matches well the moments of the 2nd PC and the high correlation between the 2nd PC and the price-dividend ratio. Not reported in the table, the model-implied second-, third-, and fifth-order autocorrelations of the 2nd PC are 0.62, 0.59, and 0.53, respectively, in reasonable agreement with the corresponding data-implied autocorrelations at 0.49, 0.40, and 0.46, respectively.

In the fourth panel we display the point estimates of the parameters along with the associated standard errors in parentheses. The point estimates of the coefficient of RRA, 14.6, and the EIS, 1.6, suggest preference for early resolution of uncertainty. 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. It is tempting to refer to the first regime as the regime of economic expansions and the second regime as the regime which encompasses economic contractions and recoveries but in Section 6 we provide caveats to such an interpretation.

In the fourth panel the conditional means of the 2nd PC are 0.75 and -2.66 in the two states, respectively, and the volatility is 0.97, rendering the signal highly informative, as it is drawn from two very different distributions in the two regimes. The signal is also highly persistent with a first-order autocorrelation coefficient of 0.65. At first sight it may seem puzzling that the signal has persistence that is an order of magnitude higher than those of consumption and dividend growth (0.02 and 0.04, respectively) despite all of these variables being driven by the same underlying state. However a 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, $(\mu_{y,1} - \mu_{y,2})^2 / \sigma_y^2$. For the consumption and dividend

growth rates this ratio is 0.13 and 0.24, respectively, while for the 2nd PC this ratio is two orders of magnitude higher at 12.29. This explains the substantially higher persistence of the signal compared to that of the consumption and dividend growth rates. We argue in Section 6 that this feature of the signal plays a key role in generating the asset pricing results. Moreover it does so without relying on excessive predictability of consumption and dividend growth rates that are difficult to measure in the data.

The model-implied price-dividend ratio is a known function of investors' beliefs which in turn is a known function of the history of consumption growth and the 2nd PC. We use the time series of consumption growth and the 2nd PC and obtain the time series of the model-implied price-dividend ratio. Using the parameter estimates in Table 3 the correlation between the model-implied price-dividend ratio and the price-dividend ratio observed in the data is 0.77. The alternative model, on the other hand, produces correlation 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.

In the Online Appendix we present robustness checks. First, in the Online Appendix A3, we reestimate the model at the quarterly frequency and verify that the results are very similar to those in Table 3. Second, in the Online Appendix A4, we constrain the signal to equal CPI-U growth or the earnings per hour growth and obtain strikingly similar results. Finally, in unreported results, we set the signal equal to GDP growth or the first or third principal component of the macroeconomic variables and obtain results similar to the learning from consumption alone model, thereby demonstrating that these variables do not serve as informative signals in the context of our model.

5.3 Estimation Results for the Alternative Model (Learning only from consumption history)

We re-estimate the model when investors learn from consumption history alone. The results are reported in Table 4. The main shortcoming of this model is that it implies essentially zero volatility of the price-dividend ratio, grossly at odds with its sample counterpart of 0.40. The flat price-dividend ratio causes the volatility of the market return to equal the volatility of the dividend growth rate,

contrary to the observed higher volatility of the market return relative to dividend growth. Thus the model fails to explain the excess volatility puzzle.

Furthermore, the model-implied autocorrelation of the price-dividend ratio is only 0.16, at odds with the autocorrelation of 0.89 observed in the data. Thus the model fails to generate the high persistence in the market-wide price-dividend ratio—one of the most robust features observed in the data. Finally, as shown in Section 6, since the model-implied price-dividend ratio is essentially constant, the model counterfactually implies a pro-cyclical expected market return (driven by pro-cyclical expected dividend growth) and flat conditional volatility of the market return. Some shortcomings of this model have earlier been highlighted in Johannes, Lochstoer, and Mou (2016) and our results confirm and extend their findings.

[Table 4 here]

5.4 Model-Implied Predictive Regressions

In Section 4 we motivated the choice of the 2nd PC (and CPI-U growth and earnings per hour growth) as signals by presenting evidence in the left panels of Table 2 that this variable predicts dividend growth at various horizons. In the right panels of this table we present the corresponding model-implied predictive regressions and compare them to the data-implied ones.

We use the point estimates in Table 3 and generate a simulated time series of the consumption and dividend growth, market return, and the 2nd PC of length one million years. In the right panels of Table 2 we present the coefficient estimates and R^2 from the predictive regressions on the simulated data. We estimate the 95% confidence intervals of the coefficients and R^2 by performing these predictive regressions in 10,000 simulated samples of length 50 years each which is the same length as the historical sample.

Consistent with the data, the 2nd PC has little forecasting power for consumption growth at any horizon. The 2nd PC forecasts one-year ahead consumption growth with a coefficient of 0.001 and R^2 1.3% in the data. The corresponding model-implied values are remarkably similar at 0.001 and 1.8%, respectively. Similar results obtain for the longer horizons. Likewise, consistent with the data, the 2nd PC does not forecast the market return at any horizon. The 2nd PC mildly forecasts one-year ahead dividend growth with a coefficient of 0.021 and R^2 9.8% in the data. The

corresponding model-implied values are 0.014 and 2.9%, respectively. Similar results obtain for the longer horizons. Finally similar results, reported in the Online Appendix A2, obtain for the CPI growth regressions and the earnings per hour growth regressions. Overall the predictive regressions provide strong support for the economic mechanism highlighted in the model as well as the ability of the model to quantitatively match predictive patterns observed in the data.

6 Discussion of the Results and Interpretation of the Economic Regimes

In Section 5 we showed that the main model fits the data reasonably 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 one of a number of macroeconomic 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 where the signal is set equal to the 2nd PC.

The crux of the intuition lies in the response of the price-dividend ratio to changes in beliefs, ρ_t , and, therefore, to changing expectations about future dividend growth. In the main model the belief process has high persistence (first-order autocorrelation coefficient of 0.95) and low variance while in the alternative model the belief process has lower persistence (first-order autocorrelation coefficient of 0.75) and high variance. In the main model the current value of ρ_t is highly informative about the future and, therefore, the price-dividend ratio sharply responds to changes in ρ_t . This makes the price-dividend ratio highly volatile in this model. By contrast in the alternative model the current value of ρ_t typically provides little information about the future and, therefore, the price-dividend ratio is essentially constant. Below we elaborate on this intuition.

6.1 Interpretation of the Economic Regimes

Whereas it is tempting to interpret the first regime as the regime of economic expansions and the second regime as the regime which encompasses economic recessions and recoveries there are caveats to such an interpretation. We earlier pointed out in figures 5

and 6 that most macroeconomic variables are either pro-cyclical or counter-cyclical but 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. Consumption growth, GDP growth, and the first and third principal components of the covariance matrix of the macroeconomic variables are neither strongly correlated with the price-dividend ratio nor serve as informative signals.

The correlations of recessions with the 2nd PC, CPI-U growth, and earnings growth are higher at -41%, 47%, and 33%, respectively, but lower than 50%. We extract the time series of the beliefs process from the observed price-dividend ratio and risk free using the parameter estimates in Table 3 for the main model and in Table 4 for the alternative model. For the main model the correlation of beliefs with the business cycle is -24% and for the alternative model the correlation is 0.6%.³ Figure 7 plots the time series of the beliefs in the main model from 1964 to 2013. The results suggest that the regimes are correlated, albeit imperfectly, with the business cycle.

[Figure 7 here]

6.2 The Beliefs Process

In figure 8 we illustrate the conditional variance of the beliefs, $\text{var}(p_{t+1} | p_t)$, as a function of p_t in the main model (red line). The estimates of the transition probabilities for the main model are 0.99 and 0.96, respectively, compared to 0.98 and 0.80, respectively, for the alternative model, i.e. the transition probability for the second regime is much higher in the main model than in the alternative model. If $p_t = 0$, we have $E[p_{t+1} | p_t = 0] = 0.04$, $\text{Var}[p_{t+1} | p_t = 0] = 0.03$, and, therefore, p_{t+1} is likely to remain close to 0 with a fair degree of certainty. In other words, the investor is able to reliably forecast the regime in the next period given the current period realizations of consumption growth and the rate of inflation. A similar argument holds when $p_t = 1$; in fact, in this case $E[p_{t+1} | p_t = 1] = 0.99$ and the conditional variance of p_{t+1} is even closer to zero ($\text{Var}[p_{t+1} | p_t = 1] = 0.006$) than when $p_t = 0$ and, therefore, p_{t+1} is likely to be

³ Correlation with recessions is computed as the correlation with a dummy variable that takes the value 1 in a year if at least one of the quarters in that year is in an NBER-designated recession period and takes the value 0 otherwise.

close to 1 with an even greater degree of certainty. In contrast when $p_t = 0.5$ we have $Var[p_{t+1}|p_t = 0.5] = 0.21$ and it is difficult for the investor to forecast the regime next period. To summarize, there is a lot of uncertainty about the future when p_t is around 0.5 but not when p_t is near its boundaries of 0 or 1. This implies that the investor has a strong precautionary savings motive around $p_t = 0.5$ that declines as p_t approaches its boundaries. This generates non-linearities in the equilibrium solutions for the price-dividend ratio, the expected market return, and the conditional variance of the market return.

[Figure 8 here]

The above results stand in contrast to the alternative model. In figure 6 we also illustrate the conditional variance of the beliefs, $var(p_{t+1} | p_t)$, as a function of p_t in the alternative model (blue line). The conditional variance is an inverted U -shaped function of p_t as in the main model but, unlike the main model, it has a value higher than 0.13 for all values of p_t between 0 and 0.76. When the variance is 0.13 the conditional volatility is 0.36 and the 95% confidence interval for the distribution of $(p_{t+1}|p_t)$ covers almost the entire permissible region. This means that the investor observes the realization of C_t and accurately infers the regime at date t because $\mu_{c,1} = 0.016$ is very different from $\mu_{c,2} = -0.034$ and the volatility of consumption growth is low at 0.019. However, this typically provides little information as to which regime the economy will be in next period (except when the current p_t is very high). Thus there is a lot of uncertainty about the regime next period although there is low uncertainty about the current regime. Once C_{t+1} is realized the investor accurately updates her beliefs regarding the regime. This implies that p_t has lower persistence and higher variance.

The above feature of the learning process explains why the price-dividend ratio is flat in the alternative model. The price-dividend ratio moves in response to changing expectations about future returns and/or future dividend growth. The current realization of p_t provides little information about the future for over three quarters of possible values of p_t and, therefore, the price-dividend ratio is unresponsive to changes in p_t . To help improve the performance of the model, Johannes, Lochstoer, and Mou (2016) introduce model and parameter uncertainty. Even

though the investor learns 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.

6.3 The Risk Free Rate

In figure 9 we plot the risk free rate as a function of the probability of the first state in the main model (red line) and in the alternative model (blue line). In the main model the risk free rate is U-shaped. We earlier cautioned that in the main model the second state should not be interpreted as the recession state. Indeed the risk free rate is lowest not when the probability of being in the second state is the highest ($\rho_t = 0$) but around $\rho_t \approx 0.75$ when the conditional variance of beliefs and, therefore, uncertainty is high.

In the alternative model the risk free rate is monotonically increasing in the probability of the first state and is lowest when the probability of being in the second state is the highest ($\rho_t = 0$).

[Figure 9 here]

6.4 The Price-Dividend Ratio

In figure 10 we display the price-dividend ratio as a function of the probability of the expansion state, ρ_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 being in the first 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 (blue line) because, as explained in Section 7.2, 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.

[Figure 10 here]

6.5 The Expected Market Return

In figure 11 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 ρ_t (blue line), leading to the counterfactual prediction of 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 ρ_t .

[Figure 11 here]

In the main model, on the other hand, when $p_t = 0.5$ 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 ρ_t increases between 0 and 0.5 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 higher probability of being in the first 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 increase with an increase in ρ_t over this range. As ρ_t increases from 0.5 towards 1 uncertainty is decreasing that the economy is in the good regime. Therefore the expected market return is decreasing in ρ_t over this range.

In figure 12 we display the expected market return as a function of the price-dividend ratio. In the main model the expected market return is strongly concave in the price-dividend ratio. This highly non-linear pattern is unlike the common practice of predicting the market return with the price-dividend ratio with a linear regression. 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 ρ_t (blue line).

[Figure 12 here]

In figure 13 we display the conditional variance of the market return. 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 ρ_{t+1} . Now the conditional variance of ρ_{t+1} is an inverted U -shaped function of ρ_t —being the highest when $\rho_t \approx 0.5$ and declining when ρ_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 around $p_t = 0.5$ and, therefore, the conditional variance of the market return peaks around $p_t = 0.5$ (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 13 here]

Note that the main model's implications that the expected market return and the conditional variance of the market return are the highest and the risk free rate lowest around $p_t = 0.5$ may seem contrary to conventional wisdom that these should be the highest and lowest, respectively, when $p_t \approx 0$, i.e. when we are certain that the economy is in a recession. To help interpret this implication of the model, consider the 2008 financial crisis and its aftermath. This was not a macroeconomic disaster episode like the Great Depression. However, there was considerable uncertainty at the time about the future of the economy. The period 2008-2013 can, therefore, be characterized as a period of low growth (negative growth for the years 2008 and 2009) and low inflation. Thus, in the context of our model, it can be characterized as a period with significant uncertainty, i.e. p_t lies away from its boundaries. Note that expected returns and conditional variance of returns were arguably exceptionally high during this period, consistent with the model – the volatility of the market return over this 6-year period was 26.4% compared to a volatility of 17.9% over the entire period 1964-2013 in Table 6. Note, however, that consumption growth was not more volatile over this 6-year period compared to the whole period (1.1% versus 1.3%). This is, once again, consistent with the model.

7 The Nominal and Real Term Structures

The recent macro finance literature has increasingly emphasized the importance of a unified framework for pricing different classes of financial assets. In this section we assess the performance of our main model in explaining the average term structure of interest rates. This is a challenging test of the model because neither nominal nor real yields of any maturity (other than the short term real interest rate) are targeted in the estimation of the model parameters.

We first consider the nominal term structure. Note that, since we argue that the rate of inflation is one of the major components of the signal over and above consumption growth and model the dynamics of the inflation process, our main model has implications for the nominal term structure. Table 5, panel A, presents the means and volatilities of yields on nominal bonds with maturities spanning one to 30 years, over the entire available sample period from 1962 to 2013. The table shows that the nominal term structure is strongly upward sloping, with the mean yields varying from 5.5% for one-year bonds to 7.3% for 30-year bonds. The volatilities of the yields, on the other hand, are slightly declining with maturity, varying from 3.2% for one-year bonds to 2.8% for 30-year bonds. We use the parameter estimates, reported in the Online Appendix A4, for the model where the signals are consumption and CPI growth to obtain the corresponding model-implied averages and volatilities of yields. These are reported in panel B. Consistent with the data, the model generates an upward sloping term structure. Moreover, the mean yields vary from 5.8% for one-year bonds to 6.9% for 30-year bonds, in close agreement with their data counterparts. Because, as we discuss below, the model-implied real term structure is flat, the main model, unlike the long run risks model, does not rely on a high inflation risk premium to match the nominal term structure. Also, as in the data, the volatilities of the yields decline with maturity.

[Table 5 here]

We next consider the term structure of real bonds. Table 5, panel C, reports the means and volatilities of Treasury Inflation Protected Securities (TIPS) with maturities five to 20 years, over the entire available sample period from 2003 to 2013. The data suggest a mild positive slope of the term structure with the mean yields varying from 0.8% for five-year bonds to 1.7% for 20-year bonds that should be interpreted with caution owing to the small sample size. As shown in panel D of Table 5, the model essentially implies a flat real term structure at 1.5%. This represents an

improvement over long run risks models that counterfactually imply a steep downward sloping real term structure. Note that our results obtain although real yields with maturities beyond one year are not targets in the estimation. The reason for this difference between our framework and the long run risks model is that the former, unlike the latter, does not rely on high persistence of the aggregate consumption growth. In the long run risks model consumption growth is strongly positively autocorrelated and this makes long terms bonds a hedge against consumption risk, causing the term premium to be negative. In our main model, on the other hand, the autocorrelation of consumption growth is close to zero rendering the real term structure essentially flat.

Overall the main model matches well the nominal and real term structures of interest rates, although none of the model parameters is estimated to match these moments. This lends further support to the economic mechanism highlighted in the model.

8 Concluding Remarks

The market-wide price-dividend ratio is strongly correlated with certain categories of macroeconomic variables, particularly inflation, labor market variables, and the 2nd PC of the macroeconomic variables that heavily loads on inflation and labor market variables. By contrast the price-dividend ratio has very small correlation with aggregate consumption and GDP growth, the variables 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 and GDP 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 consumption history and an additional signal set equal to an observed macroeconomic variable, for example, the 2nd PC, CPI-U growth, or earnings per hour growth. The model offers an explanation of the equity premium and risk free rate puzzles. It rationalizes the mean and, more importantly, the volatility of the market-wide price-dividend ratio, thereby accounting for the excess volatility puzzle. It matches the average real and nominal yields of

Treasury bonds, both at the short and long ends of the term structure. Finally, it is also consistent with the low predictive power of the price-dividend ratio for future consumption and dividend growth. Our findings suggest 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 consumption history alone. The model fails to address the excess volatility puzzle, i.e. the observed high volatility of the market return relative to dividend growth. It implies essentially zero volatility of the price-dividend ratio and significantly understates its autocorrelation. It also implies a procyclical expected market return and near constant conditional volatility of returns, in sharp contrast to the strong countercyclical patterns in these moments observed in the data. 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 future research one should endogenize the dependence of asset prices on certain macroeconomics variables and explain why inflation captures expectations for future economic growth in some periods but not in others, possibly depending on the tightness of the labor supply.

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, 2016, "Valuation Risk and Asset Pricing," *Journal of Finance* 71: 2861-2904.
- 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.
- 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.
- 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., 2011, Information Aggregation around Macroeconomic Announcements: Revisions Matter," *Journal of Financial Economics* 101: 114-131.
- 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, 2013, "Shocks and Crashes," *National Bureau of Economic Macroeconomics Annual*, edited by J. Parker and M. Woodford, MIT Press.
- Lettau, M., S. Ludvigson, and J.A. Wachter, 2008, "The Declining Equity Premium: What Role does Macroeconomic Risk Play?" *Review of Financial Studies* 21: 1653-1687.
- 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.
- 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: Regressions of the Log Price-Dividend Ratio on CPI Growth, Earnings per Hour Growth, and 2nd PC over the Full Sample Period and Subperiods

The table reports regressions of the log price-dividend ratio and its first difference on the CPI-U growth, earnings per hour growth, and 2nd PC (and their first difference) over the full period from 1964 to 2013 (panel A) and two subperiods from 1964 to 1988 (panel B) and from 1989 to 2013 (panel C). The last row of panel C presents out-of-sample R^2 s over the subperiod from 1989 to 2013.

Panel A: 1964-2013						
	y(t) = Δ CPI		y(t) = Δ EARNINGS		y(t) = 2 nd PC	
	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)
y(t)	-10.229 (1.613)		-10.795 (1.862)		.289 (.044)	
Δ y(t)		-5.006 (1.391)		-4.246 (2.744)		.132 (.034)
Intercept	4.032 (.078)	.009 (.023)	4.072 (.090)	.007 (.026)	3.610 (.044)	.009 (.024)
R ²	45.60%	21.59%	41.17%	4.85%	49.39%	25.43%
Panel B: 1964-1988						
	y(t) = Δ CPI		y(t) = Δ EARNINGS		y(t) = 2 nd PC	
	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)
y(t)	-4.745 (1.007)		-4.024 (1.596)		.120 (.036)	
Δ y(t)		-4.334 (1.436)		-3.719 (2.841)		.128 (.038)
Intercept	3.569 (.062)	.001 (.029)	3.547 (.010)	-.005 (.033)	3.368 (.042)	.004 (.030)
R ²	49.14%	28.36%	21.66%	6.93%	34.55%	35.28%
Panel C: 1989-2013						
	y(t) = Δ CPI		y(t) = Δ EARNINGS		y(t) = 2 nd PC	
	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)	p/d(t)	Δ p/d(t)
y(t)	-13.478 (4.896)		-1.306 (8.687)		.274 (.124)	
Δ y(t)		-6.918 (3.163)		-6.046 (6.755)		.138 (.066)
Intercept	4.286 (.143)	.013 (.038)	3.954 (.235)	.018 (.041)	3.771 (.100)	.014 (.040)
R ²	24.79%	17.86%	0.10%	3.51%	16.00%	17.73%
OOS R ²	14.38%	15.37%	-0.33%	2.99%	11.73%	17.64%

Table 2: Predictive Regressions of 1-, 2-, 3-, and 5-year ahead Consumption and Dividend Growth and Market Return by the 2nd PC

The left panels of the table report predictive regressions of the 1-, 2-, 3-, and 5-year ahead consumption and dividend growth and market return by the 2nd PC from 1966 to 2011. The right panels present the corresponding model-implied predictive regressions.

Panel A: 1 year						
	Data			Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Lagged 2 nd PC	.001 (.002)	.021 (.010)	.002 (.027)	.001 [-.004,.004]	.014 [-.037,.041]	-.001 [-.084,.033]
Intercept	.019 (.002)	.009 (.010)	.050* (.028)	.014 [.008,.020]	.024 [-.046,.085]	.038 [-.081,.124]
R ²	.013	.098	.000	.018 [.000,.140]	.029 [.000,.153]	.000 [.000,.169]
Panel B: cumulative 2 years						
	Data			Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Lagged 2 nd PC	.001 (.003)	.033 (.016)	-.003 (.040)	.002 [-.006,.006]	.026 [-.053,.064]	-.003 [-.118,.050]
Intercept	.038 (.003)	.019 (.016)	.097 (.040)	.028 [.016,.039]	.048 [-.079,.161]	.075 [-.139,.217]
R ²	.004	.094	.000	.032 [.000,.196]	.049 [.000,.227]	.000 [.000,.200]
Panel C: cumulative 3 years						
	Data			Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Lagged 2 nd PC	.000 (.005)	.034 (.021)	.005 (.047)	.003 [-.007,.009]	.039 [-.066,.087]	-.004 [-.150,.064]
Intercept	.057 (.004)	.027 (.020)	.142 (.046)	.043 [.024,.058]	.073 [-.123,.236]	.113 [-.200,.308]
R ²	.000	.059	.000	.043 [.000,.252]	.068 [.000,.287]	.000 [.000,.247]
Panel D: cumulative 5 years						
	Data			Model		
	Δc	Δd	r_m	Δc	Δd	r_m
Lagged 2 nd PC	-.005 (.005)	.048 (.026)	-.008 (.061)	.005 [-.009,.013]	.061 [-.085,.129]	-.007 [-.209,.088]
Intercept	.097 (.005)	.057 (.026)	.246 (.060)	.071 [.040,.097]	.121 [-.207,.389]	.188 [-.302,.493]
R ²	.021	.080	.001	.059 [.000,.339]	.095 [.000,.365]	.001 [.000,.332]

Table 3: Model Fit and Parameter Estimates with Learning from both Consumption Growth and 2nd Principal Component, 1964-2011

The table reports estimation results and model fit for the main model, using consumption growth and the 2nd PC as signals over the period 1966-2011. The parameters are estimated using GMM. Twenty-two moment restrictions are used in the GMM, namely the mean, variance, and first-order auto-covariance of the consumption and dividend growth rates, the second principal component of changes in macro variables, 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, the covariance between dividend growth and the price-dividend ratio, and the covariance between the second principal component and the price-dividend ratio. The number of parameters to be estimated is fifteen. Panel D presents the parameter estimates along with asymptotic standard errors in parentheses. The standard errors are Newey-West (1987) corrected using two lags. Panels A, B and C 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), asset prices and returns (Panel B), and second principal component (Panel C). The confidence intervals are obtained as the 2.5th and 97.5th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends									
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\rho_{\Delta c, \Delta d}$	$\rho_{\Delta c, p/d}$	$\rho_{\Delta d, p/d}$
Data	.019 (.002)	.013 (.002)	.450 (.181)	.010 (.011)	.067 (.010)	.270 (.197)	.323 (.139)	-.021 (.156)	.108 (.154)
Model	.014 [.008, .019]	.015 [.012, .018]	.022 [-.293, .269]	.025 [-.040, .075]	.142 [.111, .170]	.038 [-.286, .276]	.345 [.054, .572]	.153 [-.214, .398]	.194 [-.236, .425]
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 (.004)	.018 (.002)	.639 (.150)	.046 (.025)	.183 (.023)	-.018 (.281)	3.610 (.102)	.415 (.048)	.896 (.064)
Model	.018 [.016, .018]	.001 [.000, .002]	.610 [-.103, .894]	.040 [-.022, .095]	.202 [.132, .331]	-.017 [-.353, .258]	4.397 [3.534, 4.639]	.450 [.003, .553]	.950 [-.096, .943]
2 nd PC									
	$E[2^{nd} PC]$		$\sigma(2^{nd} PC)$		$AC1(2^{nd} PC)$		$\rho_{2^{nd} PC, p/d}$		
Data	.000 (.231)		1.011 (.153)		.758 (.120)		.703 (.065)		
Model	-.001 [-2.635, .977]		1.712 [.800, 2.089]		.646 [-.266, .750]		.836 [.367, .905]		
Parameter Estimates									
γ	ψ	δ	π_1	π_2	$\mu_{c,1}$	$\mu_{c,2}$	ρ		
14.570 (.0000)	1.642 (.0002)	.990 (.1465)	.990 (.0317)	.964 (.1028)	.015 (.0044)	.010 (.0058)	.325 (.0008)		
$\mu_{d,1}$	$\mu_{d,2}$	$\mu_{2^{nd} PC,1}$	$\mu_{2^{nd} PC,2}$	σ_c	σ_d	$\sigma_{2^{nd} PC}$			
.040 (.0266)	-.029 (.0881)	.746 (.0895)	-2.661 (.0613)	.014 (.0261)	.140 (.0349)	.972 (.0958)			

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

The table reports estimation results and model fit for the alternative model using annual data over the sample period 1964-2013. The parameters are estimated using GMM. Eighteen 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, the covariance between consumption growth and the price-dividend ratio, and the covariance between dividend 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 2.5th and 97.5th percentiles from 10,000 simulations of the same length as the historical sample.

Consumption and Dividends									
	$E[\Delta c]$	$\sigma(\Delta c)$	$AC1(\Delta c)$	$E[\Delta d]$	$\sigma(\Delta d)$	$AC1(\Delta d)$	$\rho_{\Delta c, \Delta d}$	$\rho_{\Delta c, p/d}$	$\rho_{\Delta d, p/d}$
Data	.020	.013	.513	.016	.072	.269	.253	-.046	.102
	(.002)	(.002)	(.161)	(.012)	(.010)	(.187)	(.171)	(.147)	(.137)
Model	.014	.019	.300	.018	.113	.018	.349	.129	.089
	[.000, .020]	[.013, .028]	[-.262, .645]	[-.016, .051]	[.091, .136]	[-.279, .269]	[.075, .576]	[-.553, .717]	[-.265, .417]
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	.018	.011	.745	.046	.114	.014	3.567	.003	.161
	[.006, .021]	[.001, .021]	[-.102, .892]	[.012, .079]	[.092, .137]	[-.284, .265]	[3.566, 3.569]	[.001, .005]	[-.192, .558]
Parameter Estimates									
	γ	ψ	δ	π_1	π_2	ρ			
	13.536	.849	.990	.988	.801	.329			
	(.0021)	(.0234)	(.0560)	(.0440)	(.0721)	(.0041)			
	$\mu_{c,1}$	$\mu_{c,2}$	$\mu_{d,1}$	$\mu_{d,2}$	σ_c	σ_d			
	.016	-.034	.022	-.052	.015	.112			
	(.0153)	(.2217)	(.0263)	(.3073)	(.0888)	(.0331)			

Table 5: Nominal and Real Yields, 1964-2013

The table reports the average nominal and real yields on Government bonds with maturity up to 30 years and their model-implied counterparts using the point estimates reported in Table 9 for the model where the signals are consumption growth and CPI growth.

Panel A: Nominal Yields in the Data							
	1-year	2-year	3-year	5-year	7-year	10-year	30-year
Mean	0.055	0.059	0.059	0.061	0.066	0.065	0.073
Volatility	0.032	0.035	0.030	0.029	0.029	0.027	0.028
Panel B: Model-Implied Nominal Yields							
	1-year	2-year	3-year	5-year	7-year	10-year	30-year
Mean	0.058	0.058	0.059	0.060	0.062	0.063	0.069
Volatility	0.022	0.021	0.021	0.020	0.019	0.018	0.013
Panel C: Real Yields in the Data							
	5-year	7-year	10-year	20-year	30-year		
Mean	0.008	0.011	0.014	0.017	-		
Volatility	0.012	0.010	0.009	0.007	-		
Panel D: Model-Implied Real Yields							
	5-year	7-year	10-year	20-year	30-year		
Mean	0.015	0.015	0.015	0.015	0.015		
Volatility	0.001	0.001	0.001	0.001	0.000		




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




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




Figure 3: The R^2 in univariate regressions at the annual frequency of the 2nd PC on each of the 106 transformed macroeconomic variables over the sample period from 1964 to 2011.




Figure 4: The R^2 in univariate regressions at the quarterly frequency of the 2nd PC on each of the 106 transformed macroeconomic variables over the sample period from 1964 to 2011.




Figure 5: The time series of macroeconomic variables and the price-dividend ratio over the sample period from 1964 to 2011.




Figure 6: The time series of the first three principal components of the 106 macroeconomic variables over the sample period from 1964 to 2011.




Figure 7: The time series of beliefs in the main model extracted from the price-dividend ratio and risk free rate using the parameter estimates in Table 7. The sample period is from 1966 to 2011.

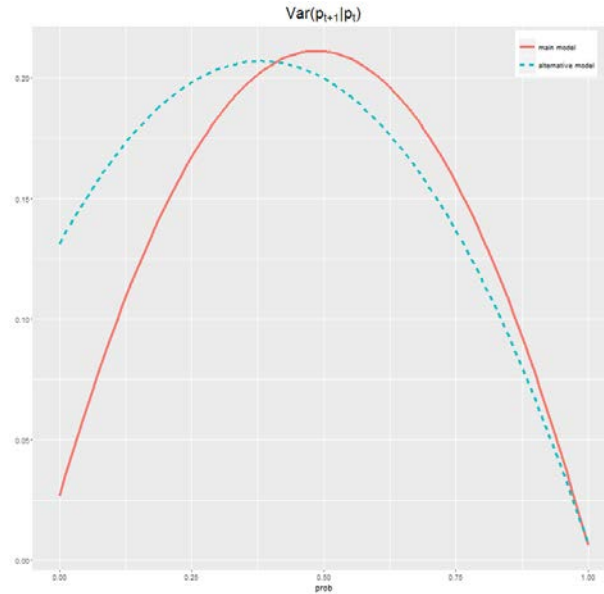


Figure 8: The conditional variance of the beliefs, $\text{var}(p_{t+1} | p_t)$, as a function of the beliefs p_t in the main model (red line) and the alternative model (blue line). The conditional variance is computed using the point estimates of the model parameters in Table 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.

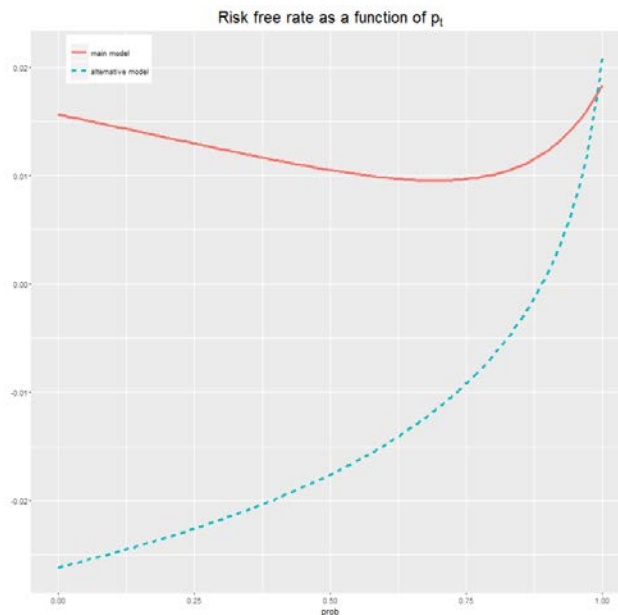


Figure 9: The risk free rate as a function of the beliefs in the main model (red line) and the alternative model (blue line), computed using the point estimates of the model parameters in Table 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.




Figure 10: 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 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.




Figure 11: 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 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.




Figure 12: 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 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.




Figure 13: 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 7 for the main model and in Table 8 for the alternative model. The sample period is from 1966 to 2011.