

NBER WORKING PAPER SERIES

THE PREDICTABILITY OF RETURNS WITH REGIME SHIFTS IN CONSUMPTION
AND DIVIDEND GROWTH

Anisha Ghosh
George M. Constantinides

Working Paper 16183
<http://www.nber.org/papers/w16183>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2010

We thank John Cochrane, Rick Green, Lars Hansen, John Heaton, Burton Hollifield, Christian Julliard, Oliver Linton, Bryan Routledge, Duane Seppi, Pietro Veronesi and seminar participants at Carnegie Mellon University, the University of Chicago, and the Sixth Annual Early Career Women in Finance Mini-Conference for helpful comments. We remain responsible for errors and omissions. Ghosh acknowledges financial support from Carnegie Mellon University. Constantinides acknowledges financial support from the Center for Research in Security Prices of the University of Chicago Booth School of Business. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Anisha Ghosh and George M. Constantinides. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Predictability of Returns with Regime Shifts in Consumption and Dividend Growth
Anisha Ghosh and George M. Constantinides
NBER Working Paper No. 16183
July 2010
JEL No. E44,G12

ABSTRACT

The predictability of the market return and dividend growth is addressed in an equilibrium model with two regimes. A state variable that drives the conditional means of the aggregate consumption and dividend growth rates follows different time-series processes in the two regimes. In linear predictive regressions over 1930-2009, the market return is predictable by the price-dividend ratio with R^2 11.7% if the probability of being in the first regime exceeds 50%; and dividend growth is predictable by the price-dividend ratio with R^2 28.3% if the probability of being in the second regime exceeds 50%. The model-implied state variables perform significantly better at predicting the equity, size, and value premia, the aggregate consumption and dividend growth rates, and the variance of the market return than linear regressions with the market price-dividend ratio and risk free rate as predictive variables.

Anisha Ghosh
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213
anishagh@andrew.cmu.edu

George M. Constantinides
The University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
gmc@ChicagoBooth.edu

1 Introduction

Stock return predictability has for long been the subject of both theoretical and empirical research in financial economics. Attempts to predict the aggregate stock market return have a long history in finance going back to as early as 1920 when Dow (1920) explored the role of dividend ratios in predicting the market return. Over the last three decades, the academic literature has explored numerous financial variables as potential predictors of the market return and equity premium. The price-dividend ratio has received extensive scrutiny as a predictive variable because, as a mathematical identity, all variation in the price-dividend ratio must be accounted for by changing expectations on future returns and/or future dividend growth (Campbell and Shiller (1988)). Welch and Goyal (2008) review this literature and undertake a comprehensive study of the in-sample and out-of-sample performance of these variables in predicting the equity premium. They conclude that "by and large, these models have predicted poorly both in-sample (IS) and out-of-sample (OOS) for 30 years now; these models seem unstable, as diagnosed by their out-of-sample predictions and other statistics; and these models would not have helped an investor with access only to available information to profitably time the market." These conclusions are controversial.

Campbell and Thompson (2008) show that, when restrictions are imposed on the theoretically expected sign of the regression coefficient and the fitted value of the equity premium, the out-of-sample R^2 improves but is still small. However, they argue that even a small value of R^2 is economically meaningful for mean-variance investors. Cochrane (2008) provides a defense of return predictability by arguing that return and dividend growth predictability are intimately related and that the absence of dividend growth predictability gives stronger evidence against the null that returns are not predictable than does the presence of return predictability in the historical data.

In this paper, we shed light on this debate by arguing that there exist (at least) two economic regimes. The market return is predictable by the price-dividend ratio in the first regime; and the market dividend growth is predictable by the price-dividend ratio in the second regime. We identify the regimes in the context of the dynamic equilibrium asset pricing model with two regimes, proposed in Constantinides and Ghosh (2009b). The probability that the economy is in the first regime is obtained as a non-linear function of the market price-dividend ratio and risk free rate, with parameters estimated from the Euler equations of the market return, the risk free rate, and the cross-section of size and book-to-market equity-sorted portfolio returns plus unconditional moments of the consumption and dividend processes. Furthermore, this non-linearity cannot be captured by simple nonlinear functions like a quadratic function of the market price-dividend ratio and risk free rate.

Over the period 1930 – 2009, in all years when the probability of being in the first regime exceeds 50%, in-sample linear predictive regressions of the realized one-year real market return and realized real dividend growth on the lagged log price-dividend ratio

have (adjusted) \bar{R}^2 11.7% for the market return and negative \bar{R}^2 for dividend growth. By contrast, in the second regime, the \bar{R}^2 for the market return is only 0.2% and for dividend growth is 28.3%. We also find that the equity premium and the returns on portfolios of “Small”, “Large”, “Growth”, and “Value” stocks are predictable by the market-wide price-dividend ratio in the first regime (when dividend growth is not predictable) with statistically significant coefficients and \bar{R}^2 varying from 3.1% for the “Small” portfolio to 12.5% for the “Large” portfolio. The price-dividend ratio performs poorly at predicting returns (except the return on “Small” stocks) in the second regime with the \bar{R}^2 varying from -3.1% for the “Large” portfolio to 1.5% for the “Value” portfolio.

In the model, a state variable x_t that simultaneously drives the conditional means of the aggregate consumption and dividend growth rates reverts to its unconditional mean with a process that differs across two regimes. Based on his information set, the consumer observes x_t and also calculates the posterior probability, p_t , that the economy is in the first regime. The conditional means of the aggregate consumption and dividend growth rates are affine functions of the two state variables (x_t, p_t) . The market-wide log price-dividend ratio and risk free rate are approximately affine functions of (x_t, p_t) and their product, thereby rendering the (potentially latent) state variables and the expected return of each asset class known nonlinear functions of the price-dividend ratio and risk free rate. The model parameters are estimated from the Euler equations of the market return, the risk free rate, and the cross-section of size and book-to-market-equity sorted portfolio returns plus unconditional moments of the consumption and dividend processes.

We show that the model has superior in-sample forecasting performance for the equity premium and its variance relative to a linear forecasting model with the market-wide price-dividend ratio and risk free rate as predictive variables. Moreover, unlike linear forecasting regressions with the price-dividend ratio and risk free rate as predictive variables, the model-implied state variables have robust forecasting performance across subperiods.

While most of the predictability literature focuses on predicting the aggregate US stock market return and equity premium, the literature on the time series forecastability of the cross-section of size and book-to-market-equity sorted portfolio returns is scant. Forecastability of the cross-section of returns is important for at least two reasons. First, the historical size premium (9.4%) and value premium (7.3%) are of the same order of magnitude as the equity premium (7.9%), based on arithmetic annual returns. Therefore, the predictability of these premia is important in active portfolio management. Second, it is also important in providing an alternative channel to examine the empirical plausibility of a given set of state variables that purport to explain the cross-section of returns. We show that the model has superior forecasting performance for the size and value premia relative to the linear forecasting model that is robust

across subperiods.

Finally, we demonstrate that our model retains its predictive power out-of-sample. We examine the out-of-sample predictive performance of the model in two ways. First, we show that the model-implied state variables give an out-of-sample R^2 of 5.1% for the equity premium and 10.1% for the market return over the period 1976 – 2009. Second, we use the central insight of the model of changing economic regimes that makes the equity premium and market return predictable by the market-wide price-dividend ratio in the first regime but not in the second regime. This procedure gives out-of-sample R^2 of 7.3% and 9.1%, respectively, for the equity premium and market return. When used as predictive variables in a linear predictive model, the price-dividend ratio and risk free rate have poor predictive performance giving out-of-sample R^2 of -1.6% and -2.5% , respectively, for the equity premium and market return.

Our paper is related to equilibrium models by Bansal and Shaliastovich (2010), Bansal and Yaron (2004), Drechsler (2009), Hansen, Heaton and Li (2008), Lettau and Ludvigson (2001), and Menzly, Santos, and Veronesi (2004) with implications on forecasting the market return and dividend growth.

Our paper is also related to Brandt and Kang (2004), Koijen and Van Binsbergen (2009), Pastor and Stambaugh (2009), and Rytchkov (2007), who focus on return predictability using filtering techniques. While these are reduced form models, we rely on an equilibrium model and avoid using filtering techniques by arguing that, under the model assumptions, the (potentially latent) state variables and the expected return of each asset class are known nonlinear functions of observable financial variables like the price-dividend ratio and risk free rate.

Finally, our work is related to Lettau and Van Nieuwerburgh (2008), Pastor and Stambaugh (2001), and Paye and Timmermann (2006) who find evidence of structural breaks and argue that allowing for these breaks has important implications for return predictability.

The paper is organized as follows. In Section 2, we present the regime shifts model. We express the price-dividend ratio, risk free rate, and expected equity premium as functions of the state variables (x_t, p_t) . The annual data over the period 1930 – 2009 are discussed in Section 3. In Section 4, we estimate the model parameters by *GMM* from the set of the Euler equations for the market return, the risk free rate, and portfolios of "Small", "Large", "Growth" and "Value" stocks, and the unconditional moments of the consumption and dividend growth processes. Using the point estimates of the model parameters, we invert the expressions for the price-dividend ratio and risk free rate as functions of the state variables and express the state variables as functions of the price-dividend ratio and risk free rate.

Armed with the time series of the state variables, we address the questions raised in this paper. Section 5 presents empirical evidence that the predictability of returns and dividend growth differ significantly in the two-regimes. In Section 6, we present evidence on the in-sample and out-of-sample predictability of the equity, size, and value

premium. In Section 7, we present evidence on the predictability of the variance of the market return. Section 8 concludes.

2 The Model and Implications for Predictability

We consider the regime shift model proposed in Constantinides and Ghosh (2009b). Here we provide a brief discussion of the model and its implications for the predictability of the equity premium, size premium, value premium, consumption growth, and dividend growth (see Constantinides and Ghosh (2009b) for further details).

2.1 Model

The model stipulates that the state variable, x_t , that simultaneously drives the conditional means of the aggregate consumption and dividend growth rates reverts to its unconditional mean with a process that differs across two regimes:

$$x_{t+1} = \rho_{s_{t+1}} x_t + \varphi_e \sigma_{s_{t+1}} e_{t+1}, \quad (1)$$

$$\Delta c_{t+1} = \mu + x_t + \sigma_{s_{t+1}} \eta_{t+1}, \quad (2)$$

$$\Delta d_{t+1} = \mu_d + \phi x_t + \varphi_d \sigma_{s_{t+1}} u_{t+1}, \quad (3)$$

where c_{t+1} is the logarithm of the aggregate consumption level; d_{t+1} is the logarithm of the aggregate stock market dividends; and $s_t = 0, 1$ is a second state variable that denotes the economic regime. The persistence parameter, ρ_{s_t} , of the state variable x_t and the level of its volatility, σ_{s_t} , are generally different in the two regimes. The shocks e_{t+1} , η_{t+1} , and u_{t+1} are assumed to be distributed with mean 0 and variance 1 and independent of the past.

Given his information set, $F(t)$, the representative consumer observes x_t and calculates his subjective probability, p_t , at time t of being in regime $s_t = 0$:

$$p_t \equiv \text{Prob}(s_t = 0 | F(t)) \quad (4)$$

We do not take a stand on the content of the information set, $F(t)$. In one extreme case, it may be limited to the history of consumption, dividends, and past realizations of x . In the other extreme case, it may include all publicly available information. Furthermore, we do not take a stand on the optimality of the filter that the consumer applies to form his belief, p_t . The econometrician does not directly observe the state variables, p_t and x_t , and, hence, they are latent.

We assume that s_t follows a Markov process with the following transition probability matrix:

$$\Pi = \begin{pmatrix} \pi_0 & 1 - \pi_1 \\ 1 - \pi_0 & \pi_1 \end{pmatrix}, \quad (5)$$

where $0 < \pi_i < 1$ for $i = 0, 1$. Thus, the consumer's probability of being in regime $s_{t+1} = 0$ at time $t + 1$, given his information set, $F(t)$, is

$$Prob(s_{t+1} = 0|F(t)) = \pi_0 p_t + (1 - \pi_1)(1 - p_t) \equiv f(p_t). \quad (6)$$

Note that $0 < f(p_t) < 1$ for all p_t , $0 \leq p_t \leq 1$.

Once the consumer updates his information set at time $t + 1$, his probability of being in regime $s_{t+1} = 0$ at time $t + 1$ is $p_{t+1} \equiv Prob(s_{t+1} = 0|F(t + 1))$. We assume that the consumer's expectations are unbiased in that

$$p_{t+1} = f(p_t) + \varepsilon_{t+1}, \quad (7)$$

where $E[\varepsilon_{t+1}|F(t)] = 0$.

We make the following assumptions regarding the shocks η_{t+1} , u_{t+1} , e_{t+1} , and ε_{t+1} :

$$E[y_{t+1}|F(t), s_{t+1} = 0] = E[y_{t+1}|s_{t+1} = 0] \equiv y(0), \text{ a constant}, \quad (8)$$

where $y = \eta$, u , e , and ε ;

$$E[y_{t+1}w_{t+1}|F(t), s_{t+1} = 0] = E[y_{t+1}w_{t+1}|F(t)] \equiv \sigma_{y,w}, \text{ a constant}, \quad (9)$$

where y , $w = \eta$, u , e , ε , and $y \neq w$; and

$$E[y_{t+1}^2|F(t), s_{t+1} = 0] = E[y_{t+1}^2] = 1, \quad (10)$$

where $y = \eta$, u , and e .

Equation (8) recognizes that the means of the residuals η_{t+1} , u_{t+1} , e_{t+1} , and ε_{t+1} , conditional on the regime at time $t + 1$, may differ from their unconditional value of zero. Equation (9) recognizes that the residuals η_{t+1} , u_{t+1} , e_{t+1} , and ε_{t+1} may be correlated. Finally, equation (10) limits the number of parameters to be estimated by setting the second moments of the residuals η_{t+1} , u_{t+1} , and e_{t+1} , conditional on the regime at time $t + 1$, equal to their unconditional value of one.

We assume that the consumer has the version of Kreps and Porteus (1978) preferences adopted by Epstein and Zin (1989) and Weil (1989). These preferences allow for a separation between the coefficient of risk aversion and the elasticity of intertemporal substitution. The utility function is defined recursively as

$$V_t = \left[(1 - \delta)C_t^{\frac{1-\gamma}{\theta}} + \delta (E[V_{t+1}^{1-\gamma}|F(t)])^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \quad (11)$$

where δ denotes the subjective discount factor, $\gamma > 0$ is the coefficient of risk aversion, $\theta \equiv \frac{1-\gamma}{1-\frac{1}{\psi}}$, and $\psi > 0$ is the elasticity of intertemporal substitution. Note that the sign of θ depends on the relative magnitudes of γ and ψ . The standard time-separable power utility is obtained as a special case when $\theta = 1$, i.e. $\gamma = \frac{1}{\psi}$.

For this specification of preferences, Epstein and Zin (1989) and Weil (1989) show that, for any asset j , the first-order conditions of the consumer's utility maximization yield the following Euler equations,

$$E[\exp(m_{t+1} + r_{j,t+1})|F(t)] = 1, \quad (12)$$

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1)r_{c,t+1}, \quad (13)$$

where m_{t+1} is the natural logarithm of the intertemporal marginal rate of substitution, $r_{j,t+1}$ is the continuously compounded return on asset j , and $r_{c,t+1}$ is the unobservable continuously compounded return on an asset that delivers aggregate consumption as its dividend each period.

We rely on log-linear approximations for the log return on the consumption claim, $r_{c,t+1}$, and that on the market portfolio (the observable return on the aggregate dividend claim), $r_{m,t+1}$, as in Campbell and Shiller (1988),

$$r_{c,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + \Delta c_{t+1}, \quad (14)$$

$$r_{m,t+1} = \kappa_{0,m} + \kappa_{1,m} z_{m,t+1} - z_{m,t} + \Delta d_{t+1}, \quad (15)$$

where z_t is the log price-consumption ratio and $z_{m,t}$ the log price-dividend ratio. In equation (14), $\kappa_1 = \frac{e^{\bar{z}}}{1+e^{\bar{z}}}$ and $\kappa_0 = \log(1 + e^{\bar{z}}) - \kappa_1 \bar{z}$ are log-linearization constants, where \bar{z} denotes the long run mean of the log price-consumption ratio. Similarly, in equation (15), $\kappa_{1,m} = \frac{e^{\bar{z}_m}}{1+e^{\bar{z}_m}}$ and $\kappa_{0,m} = \log(1 + e^{\bar{z}_m}) - \kappa_{1,m} \bar{z}_m$, where \bar{z}_m denotes the long run mean of the log price-dividend ratio.

Note that the current model specification involves two state variables, x_t and p_t . We conjecture and verify the following approximate expressions for the log price-consumption ratio and log price-dividend ratio at date t , respectively, (see Appendices A.1 and A.2 in Constantinides and Ghosh (2009b) for derivations, expressions, and intuition for the parameters $A_0(0)$, $A_1(0)$, $A_0(1)$, $A_1(1)$, $A_{0,m}(0)$, $A_{1,m}(0)$, $A_{0,m}(1)$, and $A_{1,m}(1)$):

$$z_t = p_t [A_0(0) + A_1(0)x_t] + (1 - p_t) [A_0(1) + A_1(1)x_t], \quad (16)$$

$$z_{m,t} = p_t [A_{0,m}(0) + A_{1,m}(0)x_t] + (1 - p_t) [A_{0,m}(1) + A_{1,m}(1)x_t]. \quad (17)$$

The continuously compounded risk free rate, $r_{f,t}$, between periods t and $t + 1$, is a function of the two latent state variables and their product (see Appendix A.3 in Constantinides and Ghosh (2009b) for derivation and expressions for the parameters $A_{0,f}$, $A_{1,f}$, $A_{2,f}$, $A_{3,f}$),

$$r_{f,t} = A_{0,f} + A_{1,f}x_t + A_{2,f}p_t + A_{3,f}p_t x_t. \quad (18)$$

2.2 Predictive Implications for Returns and Growth Rates

Equations (15), (17), and (3) imply that the expected market return is given by:

$$E[r_{m,t+1}|F(t)] = B_0 + B_1x_t + B_2p_t + B_3p_tx_t. \quad (19)$$

Hence, from Equations (19) and (18), the expected equity premium is given by:

$$\begin{aligned} E[(r_{m,t+1} - r_{f,t})|F(t)] &= E_0 + E_1x_t + E_2p_t + E_3p_tx_t, \\ E_i &= B_i - A_{i,f}, \quad i = 0, 1, \dots, 3. \end{aligned} \quad (20)$$

The model generates time-varying expected returns and equity premium. The coefficients $\{B_i, E_i\}_{i=0}^3$ are known functions of the underlying time-series and preference parameters. Under the assumption that the dividend growth processes of the "Small", "Large", "Growth" and "Value" portfolios are similar to that for the market, the expected returns on these portfolios can also be shown to be affine functions of the state variables, x and p , and their product.

The regime shifts model also has implications for the predictability of the aggregate consumption and dividend growth rates. The time series specification of the model implies that the expected consumption growth rate is given by

$$\begin{aligned} E(\Delta c_{t+1}|F(t)) &= \mu + x_t + E[\sigma_{s_{t+1}}\eta_{t+1}|F(t)] \\ &= \mu + x_t + (\sigma_0 - \sigma_1)\eta(0)f(p_t), \end{aligned} \quad (21)$$

and the expected dividend growth rate is given by

$$\begin{aligned} E(\Delta d_{t+1}|F(t)) &= \mu_d + \phi x_t + \varphi_d E[\sigma_{s_{t+1}}u_{t+1}|F(t)] \\ &= \mu_d + \phi x_t + \varphi_d (\sigma_0 - \sigma_1)u(0)f(p_t), \end{aligned} \quad (22)$$

both linear functions of the state variables, x_t and p_t .

3 Data

We consider the predictive performance of the model at the annual frequency, using annual data over the entire available sample period 1930 – 2009. The asset menu consists of the market return, risk free rate, and portfolios of "Value", "Growth", "Small" capitalization, and "Large" capitalization stocks. Our market proxy is the Centre for Research in Security Prices (CRSP) value-weighted index of all stocks on the NYSE, AMEX, and NASDAQ. The proxy for the annual real risk free rate is the inflation-adjusted rolled-over return of one-month Treasury Bills from Ibbotson Associates. The

equity premium is the difference in average returns on the market and the risk free rate. The construction of the size and book-to-market portfolios is as in Fama and French (1993). In particular, for the size sort, all NYSE, AMEX, and NASDAQ stocks are allocated across 10 portfolios according to their market capitalization at the end of June of each year. Value-weighted returns on these portfolios are then computed over the following twelve months. NYSE breakpoints are used in the sort. "Small" and "Large" denote the bottom and top market capitalization deciles, respectively. The size premium is the difference in average returns between the "Small" and "Large" portfolios. Similarly, value-weighted returns are computed for portfolios formed on the basis of BE/ME at the end of June of each year using NYSE breakpoints. The BE used in June of year t is the book equity for the last fiscal year end in $t - 1$ and ME is the price times shares outstanding at the end of December of $t - 1$. "Growth" and "Value" denote the bottom and top BE/ME deciles, respectively. The value premium is the difference in average returns between the "Value" and "Growth" portfolios. Annual returns for the "Small", "Large", "Growth", and "Value" portfolios are computed by compounding monthly returns within each year. The premia are computed as the difference in the average annual returns.

Also used in the empirical analysis are the price-dividend ratio and dividend growth rate of the market portfolio. These two time series are computed using the monthly returns with and without dividends on the market portfolio obtained from the CRSP files. The monthly dividend payments within a year are added to obtain the annual aggregate dividend, i.e. we do not reinvest dividends either in T-Bills or in the aggregate stock market.¹ 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.

Finally, the consumption data consists of the per capita personal consumption expenditure on nondurable goods obtained from the Bureau of Economic Analysis. All nominal quantities are converted to real, using an $ARMA(1, 1)$ forecast of the annual inflation.

4 Parameter Estimation

The parameters are estimated from the Euler equations of the market return, the risk free rate, and the "Small", "Large", "Growth", and "Value" portfolio returns

¹Koijen and Van Binsbergen (2010) consider the predictability of returns and dividend growth in the context of a reduced-form present value model for two alternative reinvestment strategies for the dividends: TBill-invested dividends that give a dividend growth series very similar to that used in this paper and market-invested dividends that give rise to a highly volatile dividend growth series having volatility about twice that for TBill-invested dividends. While the results presented in the paper are for cash-invested dividends, very similar results are obtained for market-invested dividends and are available from the authors on request.

plus unconditional moments of the consumption and dividend processes, using the GMM approach. The lagged log price-dividend ratio of the market and the risk free rate are used as instruments. The Euler equations for the six assets along with the chosen instruments give 18 moment restrictions. To this set of pricing restrictions, we add 5 moment restrictions implied by the time-series specification of the model. These moments correspond to the unconditional means and variances of aggregate consumption and dividend growth rates and the covariance between consumption and dividend growth rates. Thus, we have a total of 23 moment conditions. The total number of parameters to be estimated is 21: the 3 preference parameters (γ, ψ, δ); the 16 time-series parameters ($\mu, \mu_d, \phi, \varphi_d, \rho_0, \rho_1, \sigma_0, \sigma_1, \pi_0, \pi_1, \varphi_e, e(0), \eta(0), u(0), \varepsilon(0), \sigma_{\varepsilon,e}$); and 2 combinations of all the parameters that appear in the Euler equations.

Note that the pricing kernel is a function of the aggregate consumption growth rate and the two latent (from the point of view of the econometrician) state variables, x_t and p_t . Our estimation methodology involves inversion of two non-linear equations (17) and (18) to express the latent state variables, x_t and p_t , as functions of the observables, $z_{m,t}$ and $r_{f,t}$. This procedure yields quadratic equations for x_t and p_t , with coefficients that depend on $z_{m,t}$ and $r_{f,t}$, and the time-series and preference parameters. Solving the equations gives two pairs of solutions for x_t and p_t . We report results obtained using the bigger root of the quadratic equations as this choice minimizes the value of the GMM criterion function.

The estimation results are reported in Table 1. The first row reports the point estimates of the parameters along with the associated standard errors in parentheses. The persistence parameter of the state variable, x , in the two regimes takes values 0.60 and 0.94, respectively. This suggests that in the first regime, consumption and dividend dynamics are driven by a high frequency component that has a half-life just over 1 year. In the second regime, x has a half-life of just over 11 years. The volatility of x takes values 3.5% and 0.5%, respectively, in the two regimes. These findings suggest the presence of two regimes, one in which consumption and dividend growth rates are more persistent and less volatile and the other during which the growth rates are much less persistent and have higher volatility. The point estimates of the transition probabilities, π_0 and π_1 , suggest that the duration of the regimes are 10 and 3.3 years, respectively.

The point estimates of the subjective discount factor (0.956) and the risk aversion coefficient (12) are economically sensible. The point estimate of the IES is 0.9 and is smaller than one. However, the standard error is 0.18 and we cannot reject values of the IES slightly greater than one.

The table also reports the model-implied and the historical values of the equity premium, risk free rate, size premium, and value premium. The historically observed average level of the risk free rate is 0.9% with standard error of 0.6%. The model generates an average risk free rate of 0.4%. The model generates an equity premium of 10.3%, which is within the one standard error interval of the 7.9% value in the data.

The model also generates a size premium of 4.7%, that is within the one standard error interval of the 9.4% value in the data. Finally, the model produces a value premium of 3.8%, which is within the 95% confidence interval of the value 7.3% in the data

Note that the GMM estimation procedure examines the ability of the model to simultaneously explain the pricing restrictions given by the Euler equations and the restrictions on the unconditional moments of aggregate consumption and dividend growth rates implied by the time-series specification of the model. Therefore, the estimates of the time-series and preference parameters in Table 1 are also consistent with the time-series specification of the model. The unconditional means of consumption and dividend growth rates are 1.5% and 1.7%, respectively, in the data. The model-implied values of these moments are 3.4% and 6.4%, respectively. The unconditional variances of consumption and dividend growth rates are 0.06% and 1.4%, respectively, in the data while the model implies values of 0.18% and 3.9%, respectively. Finally, consumption and dividend growth have a correlation of 0.59 in the historical sample while the corresponding value in the model is 0.46.

Note that the market-wide log price-dividend ratio and risk free rate are approximately affine functions of x_t , p_t , and, $x_t p_t$ (equations (17) and (18), respectively). The coefficients $\{A_{i,m}(j)\}_{i,j=0}^1$ and $\{A_{i,f}\}_{i=0}^3$ are known functions of the underlying time-series and preference parameters. Therefore, using the point estimates of the parameters and the time series of the price-dividend ratio and risk free rate, we extract the time series of the state variables x_t and p_t and use them in the forecasting regressions for returns and growth rates. This gives the following expressions for the state variables x_t and p_t in terms of the market-wide log price-dividend ratio and risk free rate:

$$p_t = \frac{-b_t \pm \sqrt{b_t^2 - 4ac_t}}{2a}, \quad (23)$$

where

$$\begin{aligned} a &= -2.20 \times 10^{-12}, \\ b_t &= -78.8r_{f,t} - (3.55 \times 10^{-15}) z_{m,t} + 3.42, \\ c_t &= 6.36r_{f,t} - 1.11z_{m,t} + 0.94, \end{aligned}$$

and

$$x_t = \frac{r_{f,t} - 0.052 + (2.79 \times 10^{-14}) p_t}{1.11 + (3.55 \times 10^{-15}) p_t}. \quad (24)$$

We choose the bigger root of p_t (and the corresponding value of x_t) as this choice minimizes the GMM criterion function. In years when the bigger root of p_t exceeds one, we set $p_t = 0.99$ and in years when the bigger root of p_t is negative, we set $p_t = 0.01$.

In Figure 1, we plot the time-series of the probability that the economy is in the first regime over the period 1930 – 2008. The shaded areas mark years with at least one quarter in recession, as defined by the NBER. The vertical dashed lines mark major stock market crashes, as defined in Mishkin and White (2002). The probability is a nonlinear function of the market price-dividend ratio and risk free rate.

Figure 2 depicts the nonlinear relationship between the probability and the market-wide price-dividend ratio and risk free rate. It is a 3-dimensional plot of the probability of being in the first regime (z-axis) as a function of the price-dividend ratio (x-axis) and risk free rate (y-axis) over the period 1930-2009.

5 Economic Interpretation of the Two Regimes

The point estimates of the model parameters in Table 1 imply that the first regime has expected duration 10 years and the consumption and dividend growth dynamics are driven by a high-frequency state variable that has half-life just over one year; while the second regime has expected duration 3.3 years and the consumption and dividend growth dynamics are driven by a low-frequency state variable that has half-life just over 11 years. These properties suggest that the regimes capture features of the economy other than the business cycle.

In Table 2, we present the mean, variance, and annual autocorrelation of the dividend, consumption, and GDP growth, the risk free rate, the market-wide price-dividend ratio, the market return, and the equity, size, and value premia.² In Panel *A*, we present these summary statistics for the 54 years over the period 1931 – 2008 in which the posterior probability that the economy is in the first regime equals or exceeds 50%. In Panel *B*, we present these summary statistics for the 24 years over the period 1931 – 2008 in which the posterior probability that the economy is in the first regime is below 50%. Given the small size of these subsamples, the standard errors are large and differences in the point estimates across the two regimes are often statistically insignificant. However, the differences across several macroeconomic and financial variables make a compelling case that the two regimes are different. The average aggregate consumption growth rate is 1.7% (s.e. 0.5%) in the first regime whereas it is much smaller at 0.9% (s.e. 0.8%) in the second regime. The GDP growth is also higher (4.0% with s.e. 1.1%) in the first regime than in the second regime (2.7% with s.e. 1.7%). Similarly, the aggregate dividend growth rate is much higher at 3.8% (s.e. 1.7%) in the first regime compared to –3.1% (s.e. 3.0%) in the second regime. The average risk free rate is 1.6% (s.e. 0.6%) in the first regime while it is negative 0.7% (s.e. 2.1%) in the second

²Let t_j, t_k, \dots denote consecutive (but not necessarily adjacent) years in which the posterior probability that the economy is in the first regime equals or exceeds 50%. The first order autocorrelation of dividend growth in Panel *A* is calculated as the correlation of Δd_{t_j} with $\Delta d_{t_{j+1}}$ and not as the correlation of Δd_{t_j} with Δd_{t_k} . The other autocorrelations reported in the table are calculated accordingly.

regime. The mean of the market-wide price-dividend ratio is slightly higher (3.50) in the first regime compared to the second (3.08). The average market return is 9.9% (s.e. 3.1%) in the first regime and -1.5% (s.e. 3.7%) in the second one. The variance of the market return is higher in the second regime than the first one. The equity premium is also much higher at 8.3% (s.e. 2.9%) in the first regime relative to a negative 0.7% (s.e. 4.7%) in the second regime. Most of the size premium occurs in the first regime while the value premium is similar in the two regimes.

In Table 3, we present the results of linear regressions of the consumption growth rate, the dividend growth rate and returns with the lagged market-wide log price-dividend ratio as predictive variable in the two regimes. The first regime captures periods of consumption and dividend growth unpredictability and return predictability; the second regime captures periods of consumption and dividend growth predictability and return unpredictability. In Panel A, the aggregate consumption and dividend growth rates are not predictable by the price-dividend ratio, having statistically insignificant coefficients and small \bar{R}^2 of 1.0% and -1.9% , respectively.³ However, returns are strongly predictable by the price-dividend ratio. The market return and equity premium have statistically significant slope coefficients and \bar{R}^2 11.7% and 12.1%, respectively. The price-dividend ratio also has superior predictive ability for the cross-section of size and book-to-market-equity sorted portfolio returns with \bar{R}^2 3.1%, 12.5%, 11.1%, and 5.4%, for the "Small", "Large", "Growth", and "Value" portfolios, respectively.

The second regime captures periods of consumption and dividend growth predictability and return unpredictability. The price-dividend ratio strongly forecasts the aggregate consumption and dividend growth rates. The slope coefficients in the predictive regressions are significantly positive and the \bar{R}^2 rises from 1.0% in Panel A to 10.5% in Panel B for the consumption growth rate and from -1.9% in Panel A to 28.3% in Panel B for the dividend growth rate. The price-dividend ratio performs poorly in predicting the market return and equity premium in this regime. The regressions have statistically insignificant slope coefficients and \bar{R}^2 two and one orders of magnitude smaller, respectively, compared to their counterparts in Panel A. The price-dividend ratio also performs poorly in predicting the cross-section of returns in this regime, with the exception of the portfolio of "Small" capitalization stocks.

Taken as a whole, the results in Tables 2 and 3 suggest that the economy exhibits different characteristics across regimes. The differences in predictability across regimes shed light on why the empirical evidence on predictability which does not explicitly account for regime shifts is not robust in subperiods and its interpretation is controversial; and why recognition of structural breaks has important implications for return predictability (Lettau and Van Nieuwerburgh (2008), Pastor and Stambaugh (2001), and Paye and Timmermann (2006)).

³Throughout the paper, we use \bar{R}^2 to denote the adjusted R^2 .

6 Forecasting the Equity, Size, and Value Premia

We examine the ability of the regime shifts model to forecast the equity, size, and value premia with regressions on the model state variables, x and p , and their product. We compare the results with corresponding linear regressions on the market-wide price-dividend ratio and risk free rate. In Section 6.1, we estimate the model parameters over the period 1931 – 2009, extract the time series of the state variables, and perform in-sample forecasting regressions over the period 1931 – 2009. In Section 6.2, we estimate the model parameters over the subperiod 1931 – 1975, extract the time series of the state variables, and perform in-sample forecasting regressions over the non-overlapping subperiod 1976 – 2009. In Section 6.3, we estimate the model parameters over the subperiod 1931 – 1975, extract the time series of the state variables, and perform out-of-sample predictive regressions over the subperiod 1976 – 2009. In all cases, the model-implied regressions outperform the regressions based on the price-dividend ratio and risk free rate.

6.1 In-Sample Forecasting: 1931-2009

The expected equity premium implied by the model is an affine function of the two state variables and their product (equation (20)). We estimate the model parameters over the period 1931 – 2009 and extract the time series of the state variables. We perform an in-sample forecasting regression of the realized equity premium on the state variables and their product. The regression coefficients are marginally significant and the \bar{R}^2 is 5.0% (Table 4, Panel A). The two state variables, x and p , are highly non-linear functions of the aggregate log price-dividend ratio and risk free rate (see equations (17) and (18)). Therefore, the expected equity premium is a highly nonlinear function of the price-dividend ratio and risk free rate. We investigate whether this nonlinearity is important by performing linear forecasting regressions of the realized equity premium on the aggregate log price-dividend ratio (Row 2) and the log price-dividend ratio and risk free rate (Row 3). The regression coefficients are marginally significant and the \bar{R}^2 is 2.8% and 3.8%, respectively. This indicates that linear forecasting regressions do not capture the highly nonlinear dependence of the expected equity premium on the log price-dividend ratio and risk free rate.

The superior predictive performance of the model is also revealed in Figure 3, Panel A that plots the realized equity premium (black solid line) along with its predicted value from the forecasting regression implied by the regime shift model (green dotted line) and a linear forecasting regression using the log market-wide price-dividend ratio as a predictor variable (red dashed line). Note that the time series of the equity premium predicted by the model lines up more closely with the actual realized time series compared to the time series predicted by the price-dividend ratio. In particular, the price-dividend ratio, unlike the state variables of the regime shift model, fails to

account for the sharp movements in the equity premium in the historical data including the Great Depression of the early 30s followed by a very quick recovery and the huge run-up in asset prices in the mid-90s. To further illustrate these observations, Figure 4, Panel *B* plots the cumulative squared demeaned equity premium minus the cumulative squared regression residual from the alternative forecasting regression specifications : the predictive regression implied by the model (black solid line), and a linear predictive regression with the log price-dividend ratio as a predictor variable (red dashed line). An increase in a line indicates better performance of the named model relative to the equity premium mean while a decrease in a line indicates better performance of the equity premium mean. The figure reveals the superior predictive performance of the regime shifts model relative to the other predictor variables that is particularly pronounced during the Great depression, World War II, and the run-up in the 90s.

Equations (23) and (24) imply that the expected equity premium in equation (20) and the expected return of each asset class are highly nonlinear functions of the price-dividend ratio and risk free rate. Moreover, the nonlinearity cannot be captured by including as additional predictor variables (in addition to the price-dividend ratio and risk free rate), the square of the price-dividend ratio, the square of the risk free rate, or interaction terms of the price-dividend ratio and risk free rate in linear forecasting regressions. Figure 4 plots the expected equity premium as a function of the price-dividend ratio and risk free rate. The figure reveals a highly nonlinear relationship that cannot be captured by simple nonlinear functions of the price-dividend ratio and risk free rate.

Moreover, the superior forecasting performance of the model is robust across sub-periods. The model-implied state variables give an \bar{R}^2 of 8.7%, when attention is restricted to the subperiod 1976 – 2009 - the period over which the most widely used predictive variables perform poorly as noted in Welch and Goyal (2008). The price-dividend ratio gives an \bar{R}^2 of only 2.2% over this subperiod and the inclusion of the risk free rate further lowers the \bar{R}^2 to -1.0% .

The historical size premium (9.4%) and value premium (7.3%) are of the same order of magnitude as the equity premium (7.9%), based on arithmetic annual returns. The predictability of these premia is important in active portfolio management. It is also important in providing an alternative channel to examine the empirical plausibility of a given set of state variables that purport to explain the cross-section of returns. The results of predictive regressions for the full sample period 1931 – 2009 are presented in Table 4.

Panel *B* displays results for the size premium. The first row displays results of a regression with x , p , and their product as predictive variables. The \bar{R}^2 of the regression is 9.0%. The second row displays results of a linear regression with the market-wide log price-dividend ratio as predictive variable. The \bar{R}^2 is only 1.0% - nine times smaller than that obtained from the model-implied regression in Row 1. Row 3 displays results

from a linear regression with the risk free rate as an additional predictive variable. The \bar{R}^2 increases marginally to 2.8%, but is still more than three times smaller than that obtained from the regression in Row 1.

In Panel *C*, the results on predicting the value premium are similar to those in Panel *B*. The predictive regression with x , p , and their product as predictive variables has \bar{R}^2 3.3%. The linear regression with the lagged market-wide log price-dividend ratio as regressor has \bar{R}^2 -0.7%. The inclusion of the risk free rate further lowers the \bar{R}^2 to -2.0%.

Figure 5, Panel *A* (*C*) plots the realized size (value) premium (black solid line) along with its predicted value from the forecasting regression implied by the regime shift model (green dotted line) and a linear forecasting regression using the log market-wide price-dividend ratio as a predictor variable (red dashed line). Panel *B* (*D*) plots the cumulative squared demeaned size (value) premium minus the cumulative squared regression residual from the alternative forecasting regression specifications: the predictive regression implied by the model (black solid line), and a linear predictive regression with the log price-dividend ratio as a predictor variable (green dashed line). An increase in a line indicates better performance of the named model relative to the portfolio mean return while a decrease in a line indicates better performance of the mean return. The figure reveals the substantially superior predictive performance of the regime shifts model relative to the mean return that is particularly pronounced during the Great depression and the run-up in the 90s.

Note that the results of the predictive regression of the realized equity premium on the aggregate log price-dividend ratio and risk free rate do not support the implication of the single-regime Bansal and Yaron (2004) model that the equity premium is an affine function of the aggregate log price-dividend ratio and risk free rate.⁴

6.2 In-Sample Forecasting: 1976-2009

We reexamine the ability of the regime shifts model to forecast the equity, size, and value premia over the subperiod 1976 – 2009 for two reasons. First, it facilitates comparison with the extant literature that documents poor in-sample (and out-of-sample) performance of predictive models over this subperiod. Second, it allows us to estimate the model parameters over the first subperiod 1931 – 1975 and examine the forecasting performance of the model over the non-overlapping second subperiod

⁴This implication of the Bansal and Yaron (2004) model follows from two observations. First, the aggregate log price-dividend ratio and interest rate are affine functions of the two state variables - the conditional mean of consumption growth rate and the conditional variance of its innovation. Therefore, the two state variables are affine functions of the the aggregate log price-dividend ratio and interest rate. Second, the equity premium is an affine function of the conditional variance of the innovation of the consumption growth rate. Hence, the model predicts that the equity premium is an affine function of the price-dividend ratio and interest rate.

1976 – 2009. The forecasting performance of the model is even stronger over the subperiod compared to linear forecasting regressions with the price-dividend ratio and risk free rate as predictive variables. This demonstrates that the superior forecasting performance of the model over the full sample period 1931 – 2009 is not due to the potential look-ahead bias introduced by estimating the model parameters over the same period over which we forecast the premia. The results are reported in Table 5.

In Panel *A*, we report results for the equity premium. The first row displays results of a regression with x , p , and their product as predictive variables. The \bar{R}^2 of the regression is 7.0%. The second row displays results of a linear regression with the market-wide log price-dividend ratio as predictive variable. The \bar{R}^2 is only 2.2% - a third of that obtained from the model-implied regression in Row 1. The inclusion of the risk free rate lowers the \bar{R}^2 further to -1.0% . The poor forecasting performance of the price-dividend ratio and risk free rate over the last three decades is consistent with the findings reported in Welch and Goyal (2008).

Panel *B* displays results for the size premium. The first row shows that the regression with x , p , and their product as predictive variables yields statistically significant coefficients of x and p and an \bar{R}^2 of 25.9%. The second row shows that the coefficient of the price-dividend ratio is statistically insignificant and the \bar{R}^2 is -3.0% . Row 3 displays results from a linear regression with the risk free rate as an additional predictive variable. The coefficient of the risk free rate is statistically significant and the \bar{R}^2 rises to 15.5% but is still much smaller than that obtained from the model-implied regression in Row 1.

In Panel *C*, we report results on forecasting the value premium. The forecasting regression with x , p , and their product as predictive variables has a statistically significant coefficient of xp , and \bar{R}^2 2.3%. The linear regression with the lagged market-wide log price-dividend ratio as regressor has \bar{R}^2 -2.2% . The inclusion of the risk free rate further lowers the \bar{R}^2 to -3.2% .

Figure 6, Panel *A* plots the realized equity premium (black solid line) along with its predicted value from the forecasting regression implied by the regime shift model (green dotted line) and a linear forecasting regression using the log market-wide price-dividend ratio as a predictor variable (red dashed line). Figure 6, Panel *B* plots the cumulative squared demeaned equity premium minus the cumulative squared regression residual from the alternative forecasting regression specifications : the predictive regression implied by the model (black solid line), and a linear predictive regression with the log price-dividend ratio as a predictor variable (red dashed line). Figure 7 reports analogous plots for the market return and Figure 8 for the size and value premia. Note that the time series of the premia predicted by the model line up much more closely with the actual realized time series compared to the time series predicted by the price-dividend ratio.

6.3 Out-of-Sample Prediction: 1976-2009

Whereas many models that forecast the equity premium and/or market return in-sample in certain subperiods spectacularly fail to predict out-of-sample, we demonstrate that our model retains its predictive power out-of-sample. We examine the out-of-sample performance of our model forecasts in two ways.

Our first approach to examining the out-of-sample performance of our model forecasts relies on the observation that the two state variables, x and p , and their product should predict the equity premium and market return. At each year t , starting from 1975, we forecast the equity premium and market return in the year $t + 1$ as follows. First, we estimate the model parameters over the period 1931 – 1975 and extract the time series of the state variables. This approach is conservative because we do not use all the information in the history from 1931 to time t in estimating the model parameters. Second, we estimate the coefficients of x , p , and xp from a regression over the period 1931 to time t . Campbell and Thompson (2008) point out that the rolling out-of-sample predictive regressions are estimated over short sample periods, particularly at the beginning of the forecast evaluation period, and can, therefore, easily generate perverse results, such as a negative coefficient when theory suggests that the coefficient should be positive. In our out-of-sample forecast evaluation, we impose two restrictions suggested in Campbell and Thompson (2008): *a*) we set the regression coefficients to zero whenever they have the “wrong” sign (different from the theoretically expected sign estimated over the full sample), and *b*) we assume that investors rule out a negative equity premium or market return, and set the forecast to the zero whenever it is negative.

The out-of-sample performance of these forecasts is evaluated using an out-of-sample R^2 statistic as in Campbell and Thompson (2008) and Welch and Goyal (2008):

$$R_{OS}^2 = 1 - \frac{MSE_A}{MSE_N}, \quad (25)$$

where MSE_A denotes the mean-squared prediction error from the predictive regression implied by the model and MSE_N denotes the mean-squared prediction error of the historical average return. If the out-of-sample R^2 is positive, then the predictive regression has lower mean-squared prediction error than the historical average return.

The results are reported in Table 6. The first row of Panel *A* reports the out-of-sample results for the model-implied predictive regression for the equity premium. The out-of-sample R^2 is economically large at 5.1%. The first row of Panel *B* shows that the model-implied predictive regression for the market return has an even higher out-of-sample R^2 of 10.1%.

Our second approach to examining the out-of-sample performance of our model forecasts uses the central insight of the model of changing economic regimes that makes the equity premium and market return predictable by the market-wide price-dividend ratio in the first regime but not in the second regime (see Table 3). We estimate the

model parameters over the period 1931 – 1975 and extract the time series of the state variable, p_t . At each year $t + 1$, starting from 1975, we perform the following regression using data for all prior years:

$$r_{m,t+1} - r_{f,t} = \alpha_0 + \alpha_1 I_{\{p_t > 0.5\}} z_{m,t} + v_{t+1}. \quad (26)$$

We use the coefficient estimates to predict the equity premium for the subsequent year. Equation (26) implies that for those time periods in which the probability of being in the first regime, p_t , is bigger than 0.5, the price-dividend ratio is used to predict the equity premium, whereas in the time periods when $p_t < 0.5$ the equity premium is assumed not to be forecastable. We perform a similar regression for the market return. As for the first approach in Row 1, we impose restrictions on the signs of the slope coefficients and on the signs of the forecasts.

The results are reported in Table 6. Row 2 of Panel *A* shows that the predictive regression (26) for the equity premium gives an out-of-sample R^2 of 7.3%. The results in Panel *B* for the market return provide even stronger evidence in favour of the two-regime model. The model-implied predictive regression (26) for the market return gives a large out-of-sample R^2 of 9.1%.

We compare the predictive performance of the model-implied regressions in Rows 1 and 2 to a specification that ignores the presence of regimes and performs a linear predictive regression of the realized equity premium and market return on the lagged log price-dividend ratio. As for the model-implied regressions in Rows 1 and 2, we impose restrictions on the signs of the slope coefficients and on the signs of the forecasts. Row 3 of Panel *A* shows that the linear regression model gives out-of-sample R^2 of 0.2% for the equity premium - an order of magnitude smaller than those obtained from the model-implied regressions in Rows 1 and 2. Row 3 of Panel *B* shows that the log price-dividend ratio gives an out-of-sample R^2 of -4.6% for the market return - more than two orders of magnitude smaller than that obtained from the model-implied predictive regression in Row 1 and more than one order of magnitude lower than that obtained from the model-implied regression in Row 2.

Row 4 shows that addition of the risk free rate to the linear regression model worsens its out-of-sample predictive performance and gives negative out-of-sample R^2 of -1.6% and -2.5%, respectively, for the equity premium and market return. The poor out-of-sample predictive performance of the price-dividend ratio and risk free rate over the last thirty years has also been reported in Welch and Goyal (2008).

Figure 9, Panel *A* plots the realized equity premium (black solid line) along with its predicted value from the predictive regression (26) implied by the regime shift model (green dotted line) and a linear predictive regression using the log market-wide price-dividend ratio as a predictor variable (red dashed line). Figure 9, Panel *B* plots the cumulative squared demeaned equity premium minus the cumulative squared regression residual from the alternative forecasting regression specifications : the predictive regression implied by the model (black solid line), and a linear predictive regression

with the log price-dividend ratio as a predictor variable (red dashed line). Figure 10 reports analogous plots for the market return.

7 Forecasting the Variance of Market Return

We estimate the conditional variance of the annual market return as the sum of squares of the twelve monthly log returns. In Table 7, we report the results of predictive regressions of this conditional variance over 1931–2009 on x , p , and their product (Row 1), the lagged aggregate log price-dividend ratio (Row 2), and the lagged aggregate log price-dividend ratio and risk free rate (Row 3). In Row 1, the regression coefficient on xp is statistically significant and the \bar{R}^2 of the regression is economically large at 6.5%. In Rows 2 and 3, the regression coefficients on the log price-dividend ratio and risk free rate are not statistically significant and the values of \bar{R}^2 are much smaller than that in the regression of Row 1.

The superior performance of the model in forecasting the conditional variance of the annual market return is illustrated in Figure 11 that plots the realized variance (black solid line) along with its predicted value from the forecasting regression implied by the regime shift model (green dotted line) and a linear forecasting regression using the log market-wide price-dividend ratio as a predictor variable (red dashed line). Note that the time series of the variance predicted by the model lines up more closely with the actual realized time series compared to the time series predicted by the price-dividend ratio. In particular, the price-dividend ratio, unlike the state variables of the regime shift model, fails to account for the sharp movements in the variance in the historical data including the Great Depression of the early 30s, the Oil shock in the mid 70s, and the 1987 crash.

8 Concluding Remarks

We address the predictability of returns and of consumption and dividend growth in an equilibrium model with two regimes. The novel state variable is the probability that the economy is in the first regime. The economy exhibits different characteristics across regimes. The first regime captures periods of dividend growth unpredictability and return predictability; while the second regime captures periods of dividend growth predictability and return unpredictability. The differences in predictability across regimes shed light on the controversial interpretation of the extant empirical evidence on predictability and, in particular, the lack of robustness across subperiods. We show that the model-implied state variables perform significantly better at predicting the equity, size, and value premia and the variance of the market return over 1931–2009 and 1976–2009 than linear regressions with the market log price-dividend ratio and log risk free rate as predictive variables.

The economy exhibits other differences across regimes as well. The aggregate consumption and dividend growth rates are much higher in the first regime compared to the second. The average market return and equity premium are substantially higher in the first regime than in the second one. The variance of the market return and premium are higher in the second regime than in the first one. The size premium is higher in the first regime than in the second one. The value premium exhibits reversal in the first regime and momentum in the second one.

The first regime has expected duration 10 years and the consumption and dividend growth dynamics are driven by a high-frequency state variable that has half-life just over a year; while the second regime has expected duration 3.3 years and the consumption and dividend growth dynamics are driven by a low-frequency state variable that has half-life just over 11 years. These properties suggest that the regimes capture features of the economy other than the business cycle. High on our agenda is an understanding of the economic forces that differentiate the regimes.

A related goal is the investigation on the number of regimes that are needed to adequately describe the economy since there is no a priori reason that there should be only two economic regimes. The challenge is the judicious increase of the number of regimes in a model that retains computational and empirical tractability.

High also on our agenda is a unified theoretical framework that explains both the historically observed levels of returns of different classes of assets as well as their time series predictability. The current paper focuses on equities but the methodology is general and applicable to bonds, derivatives, and other asset classes.

References

- [1] Bansal, R., and Yaron, A., (2004), "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 1481-1509.
- [2] Bansal, R., and Shaliastovich, I., (2008), "Learning and Asset Price Jumps," *Working paper*.
- [3] Brandt, M. W., and Kang, Q., (2004), "On the Relation Between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach," *Journal of Financial Economics*, 72, 217-257.
- [4] Campbell, J. Y., and Shiller, R.J., (1988), "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195-227.
- [5] Campbell, J. Y., and Thompson, S., (2008), "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?," *Review of Financial Studies*, 21(4), 1509-1531.
- [6] Cochrane, J. H., (2008), "The Dog That Did Not Bark: A Defense of Return Predictability," *Review of Financial Studies*, 21(4). 1533-1575.
- [7] Constantinides, G., and Ghosh, A., (2009a), "Asset Pricing Tests with Long Run Risks in Consumption Growth," *Working paper*.
- [8] Constantinides, G., and Ghosh, A., (2009b), "The Cross-Section of Returns with Regime Shifts in Consumption and Dividend Growth," *Working paper*.
- [9] Dow, C. H., (1920), *Scientific Stock Speculation*. The Magazine of Wall Street.
- [10] Drechsler, I., (2009), "Uncertainty, Time-Varying Fear, and Asset Prices," *Working paper*.
- [11] Epstein, L. R., and Zin, S. E., (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-969.
- [12] Hansen, L.P., Heaton, J., and Li, N., (2008), "Consumption Strikes Back," *Journal of Political Economy*, 116(2), 260-302.
- [13] Koijen, R. S. J., and Van Binsbergen, J. H., (2009), "Predictive Regressions: A Present-Value Approach," *Forthcoming in Journal of Finance*.
- [14] Kreps, D. M., and Porteus, E. L., (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46(1), 185-200.

- [15] Lettau, M., and Ludvigson, S., (2001), "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 56, 815-849.
- [16] Lettau, M., and Van Nieuwerburgh, S., (2008), "Reconciling the Return Predictability Evidence," *Review of Financial Studies*, 21(4), 1607-1652.
- [17] Menzly, L., Santos, T., and Veronesi, P., (2004), "Understanding Predictability," *Journal of Political Economy*, 112(1), 1-47.
- [18] Pastor, L., and Stambaugh, R. F., (2001), "The Equity Premium and Structural Breaks," *Journal of Finance*, 56, 1207-1239.
- [19] Pastor, L., and Stambaugh, R. F., (2009), "Predictive Systems: Living with Imperfect Predictors," *Journal of Finance*, 64, 1583-1628.
- [20] Paye, B., and Timmerman, A., (2006), "Instability of Return Prediction Models," *Journal of Empirical Finance*, 13 (3), 274-315.
- [21] Rytchkov, O., (2007), "Filtering Out Expected Dividends and Expected Returns," *Unpublished paper*, Texas A&M.
- [22] Weil, P., (1989), "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24(3), 401-421.
- [23] Welch, I., and Goyal, A., (2008), "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455-1508.

Table 2: Summary Statistics in the two Regimes

<i>Panel A: Regime 1</i>			
	$E(.)$	$Var(.)$	$AC(1)$
<i>Dividend Growth</i>	0.038 (0.017)	0.008 (0.003)	0.22 (0.08)
<i>Consumption Growth</i>	0.017 (0.005)	0.0003 (0.0002)	0.19 (0.09)
<i>GDP Growth</i>	0.040 (0.011)	0.001 (0.0003)	0.26 (0.12)
<i>Risk Free Rate</i>	0.016 (0.006)	0.001 (0.0007)	0.43 (0.05)
<i>Price Dividend Ratio</i>	3.50 (0.82)	0.17 (0.06)	0.90 (0.34)
<i>Market</i>	0.099 (0.031)	0.029 (0.009)	0.06 (0.10)
<i>Equity Premium</i>	0.083 (0.029)	0.029 (0.008)	0.06 (0.10)
<i>Size Premium</i>	0.058 (0.038)	0.047 (0.014)	0.36 (0.11)
<i>Value Premium</i>	0.045 (0.027)	0.042 (0.012)	-0.13 (0.13)
<i>Panel B: Regime 2</i>			
	$E(.)$	$Var(.)$	$AC(1)$
<i>Dividend Growth</i>	-0.031 (0.030)	0.023 (0.012)	-0.05 (0.49)
<i>Consumption Growth</i>	0.009 (0.008)	0.001 (0.001)	0.37 (0.48)
<i>GDP Growth</i>	0.027 (0.017)	0.004 (0.001)	0.56 (0.55)
<i>Risk Free Rate</i>	-0.007 (0.021)	0.005 (0.002)	0.74 (0.85)
<i>Price Dividend Ratio</i>	3.08 (1.07)	0.17 (0.09)	0.84 (0.39)
<i>Market</i>	-0.015 (0.037)	0.054 (0.021)	-0.06 (0.14)
<i>Equity Premium</i>	-0.007 (0.047)	0.060 (0.024)	-0.12 (0.23)
<i>Size Premium</i>	0.016 (0.068)	0.070 (0.033)	0.55 (0.38)
<i>Value Premium</i>	0.042 (0.060)	0.049 (0.021)	0.47 (0.28)

The table reports the mean, variance, and first order autocorrelation of the aggregate dividend, consumption, and GDP growth rates, the risk free rate, the market-wide price-dividend ratio, the market return, and the equity, size, and value premia in the two regimes over the period 1930-2009.

Table 3: Forecastability in the two Regimes

<i>Panel A: Regime 1</i>			
	<i>const.</i>	<i>log (P/D)</i>	<i>Adj-R²</i>
<i>Consumption Growth</i>	-0.01 (0.02)	0.01 (0.01)	0.010
<i>Dividend Growth</i>	0.02 (0.10)	0.004 (0.03)	-0.019
<i>Market</i>	0.61 (0.18)	-0.15 (0.05)	0.117
<i>Equity Premium</i>	0.60 (0.18)	-0.15 (0.05)	0.121
<i>Small</i>	0.70 (0.34)	-0.16 (0.10)	0.031
<i>Large</i>	0.60 (0.18)	-0.15 (0.05)	0.125
<i>Growth</i>	0.65 (0.21)	-0.16 (0.06)	0.111
<i>Value</i>	0.71 (0.29)	-0.17 (0.08)	0.054

<i>Panel B: Regime 2</i>			
	<i>const.</i>	<i>log (P/D)</i>	<i>Adj-R²</i>
<i>Consumption Growth</i>	-0.09 (0.05)	0.032 (0.017)	0.105
<i>Dividend Growth</i>	-0.64 (0.20)	0.20 (0.06)	0.283
<i>Market</i>	0.34 (0.35)	-0.11 (0.11)	0.002
<i>Equity Premium</i>	0.40 (0.37)	-0.13 (0.12)	0.011
<i>Small</i>	1.11 (0.53)	-0.36 (0.17)	0.132
<i>Large</i>	0.17 (0.34)	-0.06 (0.11)	-0.031
<i>Growth</i>	0.19 (0.35)	-0.07 (0.11)	-0.028
<i>Value</i>	0.63 (0.53)	-0.20 (0.17)	0.015

The table reports the predictive performance of the market-wide log price-dividend ratio in each of the two regimes for the aggregate consumption growth, aggregate dividend growth, market return, equity premium, and portfolios of "Small", "Large", "Growth", and "Value" stocks over the period 1930 – 2008.

Table 4: Forecasting Equity, Size, Value Premia 1931-2009

<i>Panel A: Equity Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
0.07 (0.06)	-1.70 (1.10)	-0.03 (0.10)	2.38 (3.22)			0.050
0.36 (0.17)				-0.09 (0.05)		0.028
0.33 (0.17)				-0.08 (0.05)	-0.60 (0.45)	0.038
<i>Panel B: Size Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
-0.01 (0.07)	-1.02 (1.25)	-0.12 (0.11)	-3.37 (3.65)			0.090
0.28 (0.20)				-0.07 (0.06)		0.010
0.25 (0.19)				-0.06 (0.06)	-0.86 (0.52)	0.028
<i>Panel C: Value Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
0.03 (0.06)	-1.48 (1.17)	0.03 (0.11)	2.75 (3.44)			0.033
0.17 (0.18)				-0.04 (0.05)		-0.007
0.16 (0.18)				-0.03 (0.05)	-0.06 (0.48)	-0.020

The table reports results from forecasting regressions for the equity, size, and value premia over the period 1931-2009. Panels A, B, and C report results for the equity, size, and value premia, respectively. The first row of each panel reports the regression coefficients along with the associated standard errors, and the adjusted- R^2 from the forecasting regression of the realized premium on x , p , and xp . The second row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression of the realized premium on the lagged aggregate log price-dividend ratio. The third row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression of the realized premium on the lagged aggregate log price-dividend ratio and the log risk free rate.

Table 5: Forecasting Equity, Size, Value Premia 1976-2009

<i>Panel A: Equity Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
0.08 (0.04)	-2.64 (2.03)	-0.004 (0.09)	7.00 (4.74)			0.070
0.36 (0.23)				-0.08 (0.06)		0.022
0.35 (0.25)				-0.08 (0.07)	0.06 (1.10)	-0.010
<i>Panel B: Size Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
-0.09 (0.04)	-4.32 (2.15)	0.26 (0.10)	6.64 (5.04)			0.259
-0.03 (0.29)				0.01 (0.08)		-0.030
0.19 (0.27)				-0.02 (0.07)	-3.39 (1.20)	0.155
<i>Panel C: Value Premium</i>						
<i>const.</i>	<i>x</i>	<i>p</i>	<i>xp</i>	<i>log (P/D)</i>	<i>r_f</i>	<i>Adj-R²</i>
0.03 (0.05)	-3.23 (2.41)	0.18 (0.11)	9.23 (5.63)			0.023
0.22 (0.28)				-0.04 (0.08)		-0.022
0.29 (0.29)				-0.05 (0.08)	-1.05 (1.29)	-0.032

The table reports results from forecasting regressions for the equity, size, and value premia over the subperiod 1976-2009. Panels A, B, and C report results for the equity, size, and value premia, respectively. The first row of each panel reports the regression coefficients along with the associated standard errors, and the adjusted- R^2 from the forecasting regression of the realized premium on x , p , and xp . The second row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression of the realized premium on the lagged aggregate log price-dividend ratio. The third row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression of the realized premium on the lagged aggregate log price-dividend ratio and the log risk free rate.

Table 6: Out-of-Sample Prediction, 1976-2009

<i>Panel A: Equity Premium</i>	
	R^2
x, p, xp	0.051
$\log(P/D) \times I_{\{p_t > 0.5\}}$	0.073
$\log(P/D)$	0.002
$\log(P/D), r_f$	-0.016
<i>Panel B: Market Return</i>	
	R^2
x, p, xp	0.101
$\log(P/D) \times I_{\{p_t > 0.5\}}$	0.091
$\log(P/D)$	-0.046
$\log(P/D), r_f$	-0.025


The table reports results from out-of-sample predictive regressions for the equity premium and market return over the period 1976-2009. Panels A and B report results for the equity premium and market return, respectively. The first row of Panel A (B) reports the out-of-sample R^2 from the model-implied predictive regression of the realized equity premium (market return). The second row of Panel A (B) reports the out-of-sample R^2 from the predictive regression of the realized equity premium (market return) on the lagged aggregate log price-dividend ratio. The third row of Panel A (B) reports the out-of-sample R^2 from the predictive regression of the realized equity premium (market return) on the lagged aggregate log price-dividend ratio and the log risk free rate.

Table 7: Forecasting Market Return Variance, 1931-2009

<i>const.</i>	x	p	xp	$\log(P/D)$	r_f	$Adj-R^2$
0.004 (0.001)	0.015 (0.015)	-0.003 (0.001)	-0.027 (0.045)			0.065
0.005 (0.002)				-0.001 (0.001)		0.010
0.005 (0.002)				-0.001 (0.001)	0.003 (0.006)	0.001

The table reports results from the predictive regression for the market return variance over the period 1931-2009. The sum of squares of the monthly returns is used as a proxy for the conditional variance. The first row of reports the regression coefficients along with the associated standard errors, and the adjusted- R^2 from the forecasting regression on the two state variables, x and p , and their product. The second row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression on the lagged aggregate log price-dividend ratio. The third row reports the regression coefficients along with the associated standard errors and the adjusted- R^2 from the forecasting regression on the lagged aggregate log price-dividend ratio and the log risk free rate.

Figure 1: Posterior Probability of First Regime



The figure plots the probability of being in the first regime over 1930 – 2008 along with the NBER recessions (shaded columns) and the major stock market crashes (vertical dashed lines) identified in Mishkin and White (2002). The aggregate log price-dividend ratio and the risk free rate are affine functions only of the two state variables and their product. Hence, we can invert the system to extract the two state variables as known functions of the observable aggregate log price-dividend ratio and risk free rate. This procedure of inverting two non-linear equations yields quadratic equations for x_t and p_t , with coefficients that depend on $z_{m,t}$ and $r_{f,t}$, and the time-series and preference parameters. Solving the equations gives two pairs of solutions for x_t and p_t as functions of the observables, $z_{m,t}$ and $r_{f,t}$. We choose the bigger root as this minimizes the GMM criterion function.





Figure 2: The figure depicts the nonlinear relationship between the probability of being in the first regime and the market-wide price-dividend ratio and risk free rate. It plots the probability of the z-axis

Panel A: Realized and Predicted Equity Premium



Panel B: Cumulative SSE Difference




Figure 3: Panel A plots the realized equity premium along with its model forecasted value over 1931-2009. Panel B shows the annual in-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the equity premium. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned equity premium minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned equity premium minus the cumulative squared LM residual.

Expected Equity Premium as a Linear Function of P/D and σ^2




Figure 4: The figure depicts the nonlinear relationship between the expected equity premium and the market-wide price-dividend ratio and risk free rate.







Figure 5: Panel A (C) plots the realized size (value) premium along with its model forecasted value over 1931-2009. The model implies that the expected size and value premia are affine functions of the two state variables, x and p , and their product, xp . Panel B (D) shows the annual in-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the equity premium. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned size (value) premium minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned size (value) premium minus the cumulative squared LM residual.

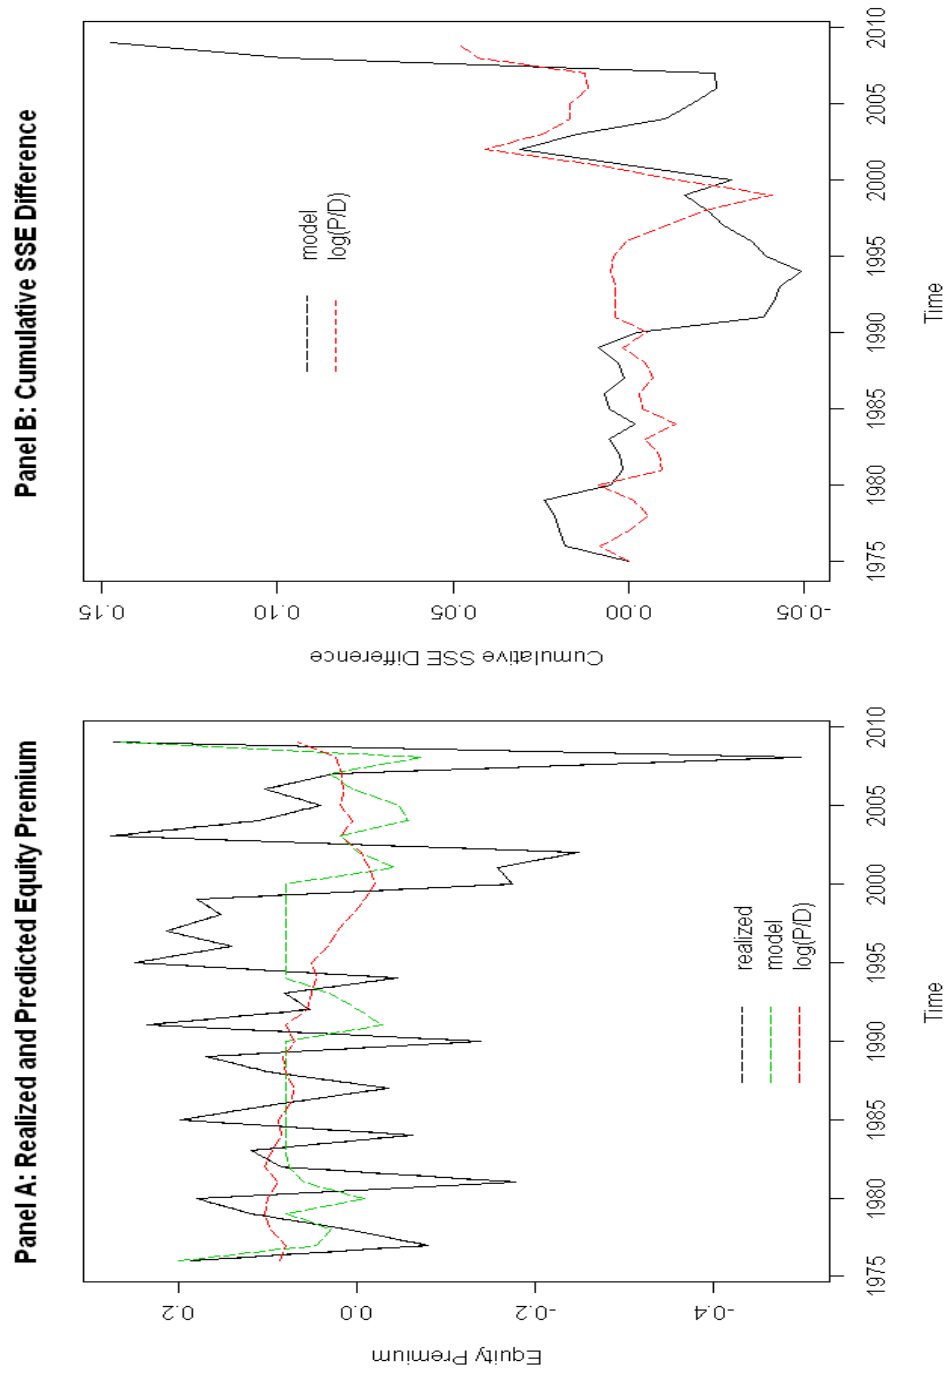


Figure 6: Panel A plots the realized equity premium along with its model forecasted value over 1976-2009. Panel B shows the annual in-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the equity premium. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned equity premium minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned equity premium minus the cumulative squared LM residual.

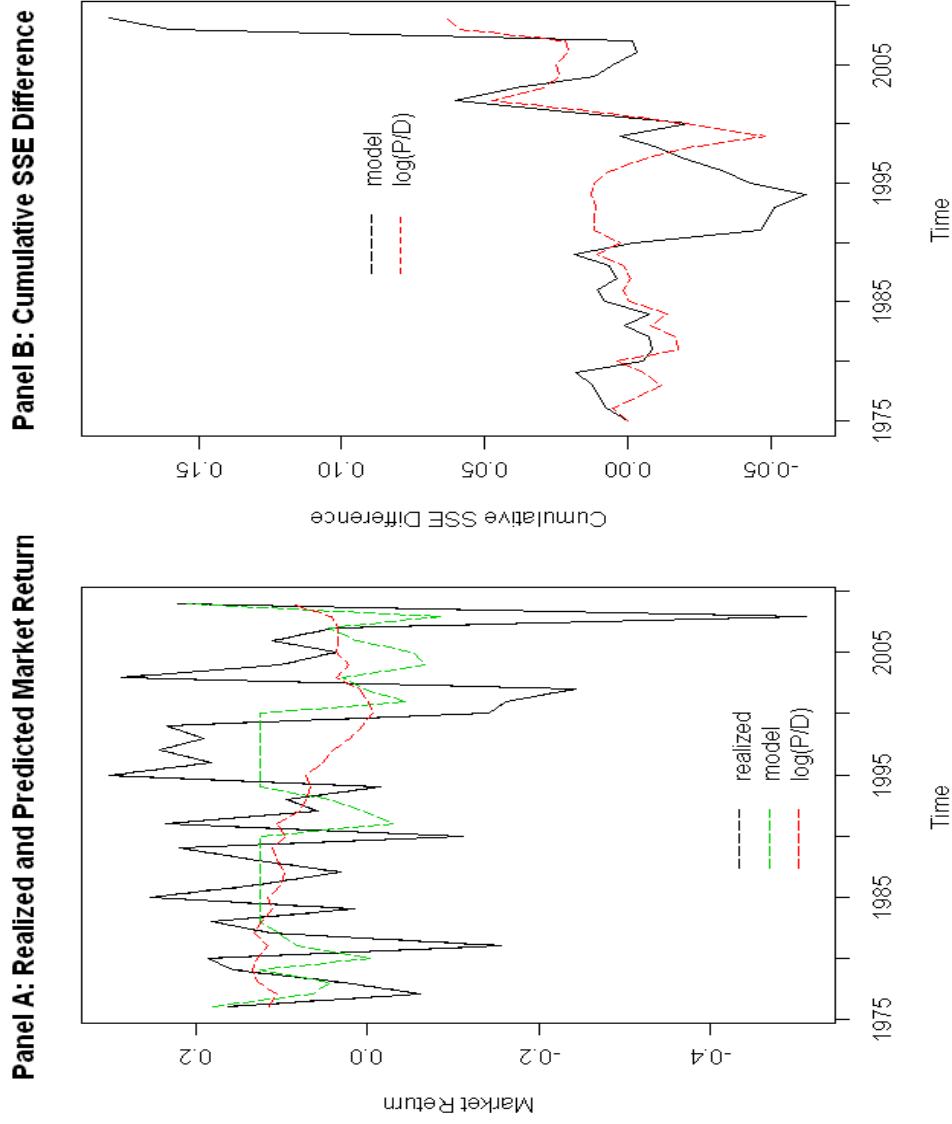


Figure 7: Panel A plots the realized market return along with its model forecasted value over 1976-2009. Panel B shows the annual in-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the market return. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned market return minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned market return minus the cumulative squared LM residual.

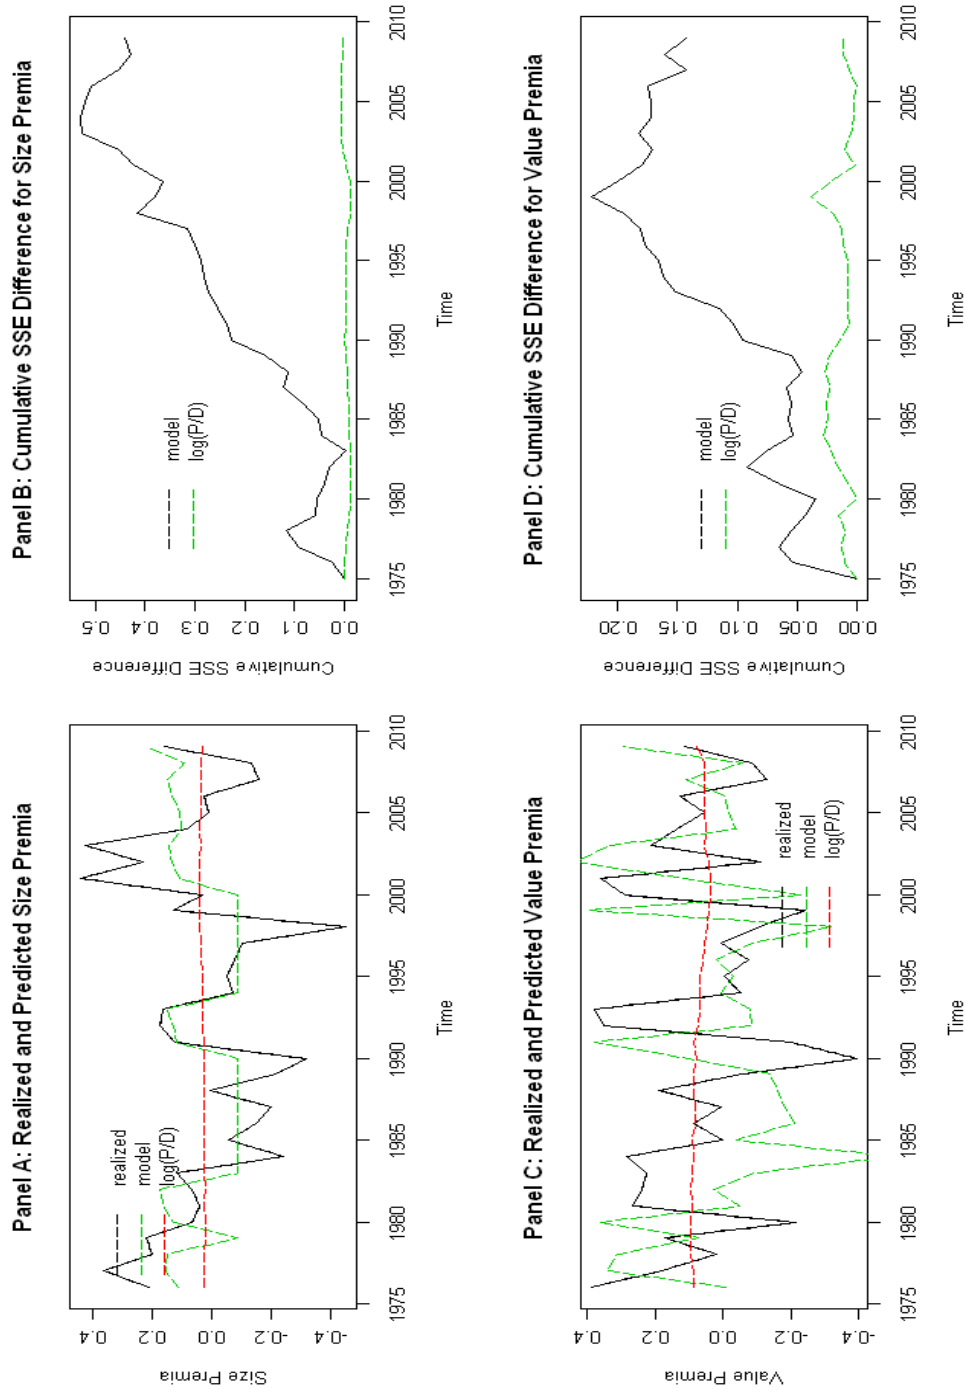



Figure 8 shows the annual in-sample performance of the model-implied predictive regression for the size premium and value premium over the subperiod 1976-2009. The figure plots the "Cumulative SSE Difference" is the cumulative squared demeaned portfolio return minus the cumulative squared model-implied regression residual. In other words, it is the cumulative squared prediction error of the historical mean return minus the cumulative squared prediction error of the regime shift model.

Panel A: Equity Premium: Out-of-Sample



Panel B: Cumulative SSE Difference




Figure 9: Panel A plots the realized equity premium along with its out-of-sample predicted value over 1976-2009. Panel B shows the annual out-of-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the equity premium. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned equity premium minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned equity premium minus the cumulative squared LM residual.

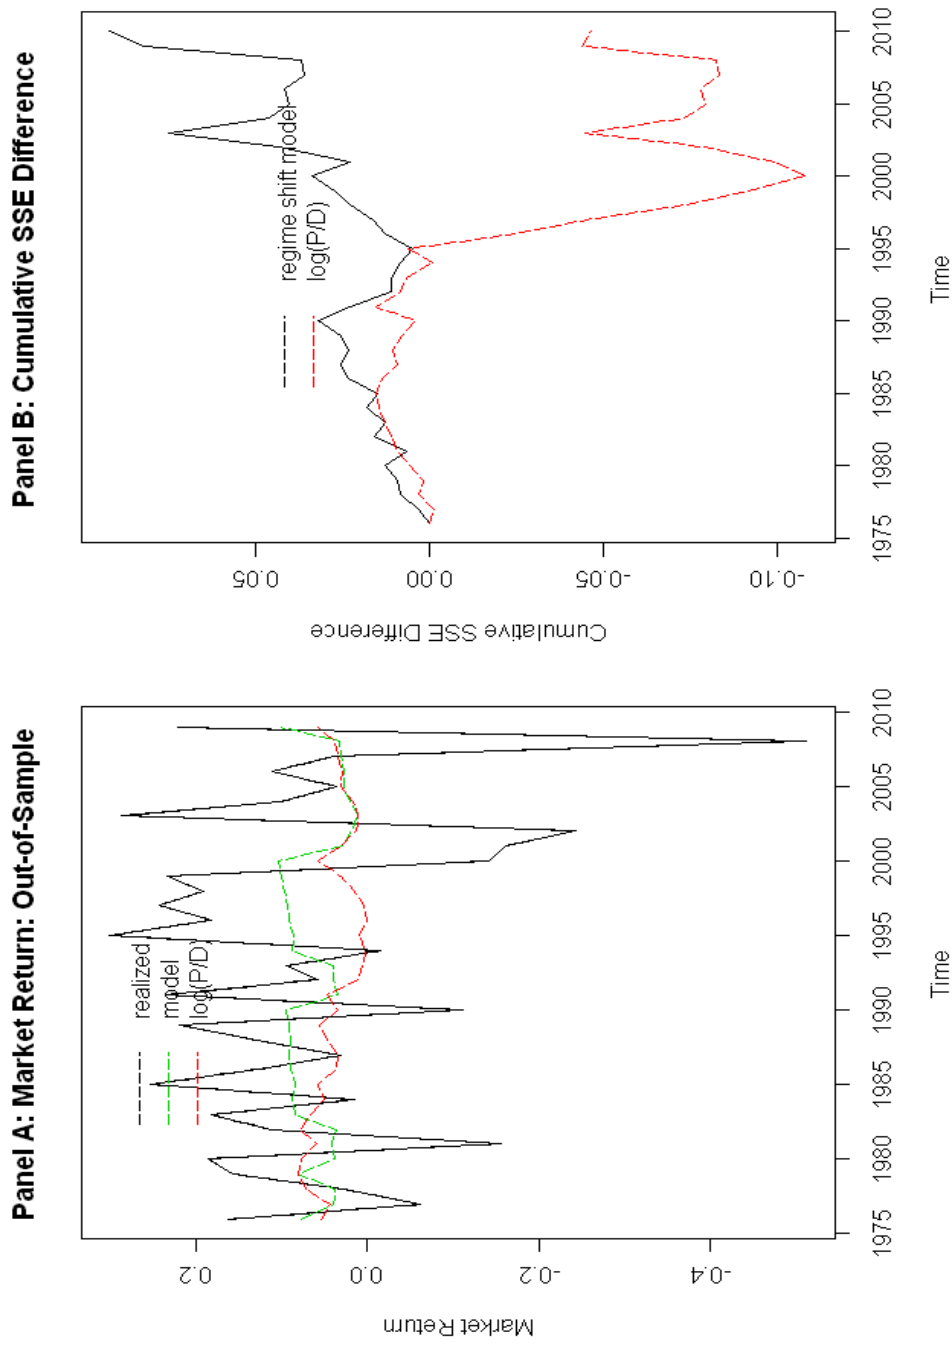


Figure 10: Panel A plots the realized market return along with its out-of-sample predicted value over 1976-2009. Panel B shows the annual out-of-sample performance of the model-implied predictive regression (RSM) as well as a linear forecasting regression that uses the lagged log price-dividend ratio as a predictor variable (LM) for the equity premium. The figure plots the "Cumulative SSE Difference" defined as the cumulative squared demeaned market return minus the cumulative squared RSM residual. It also plots the cumulative squared demeaned market return minus the cumulative squared LM residual.




Figure 11: The figure plots the realized market variance along with its model-predicted value over 1931 – 2009. The black solid line plots the annual realized variance that is calculated as the sum of squares of the monthly log market returns. The green dotted line plots the model-predicted variance and is obtained as the fitted value from a regression of the realized variance on the two state variables and their product. The red dashed line shows the predictive power of the log price-dividend ratio for the future variance by plotting the fitted value from a regression of the realized variance on the lagged log price-dividend ratio.