

Reconciling the Return Predictability Evidence In-Sample Forecasts, Out-of-Sample Forecasts, and Parameter Instability *

Martin Lettau

New York University Stern School of Business, CPER and NBER

Stijn Van Nieuwerburgh

New York University Stern School of Business

November 7, 2005

Abstract

Evidence of stock return predictability by financial ratios is still controversial, as documented by inconsistent results for in-sample and out-of-sample regressions and by substantial parameter instability. This paper shows that these seemingly incompatible results can be reconciled if the assumption of a fixed steady-state mean of the economy is relaxed. We find strong empirical evidence in support of shifts in the steady-state and propose simple methods to adjust financial ratios for such shifts. The forecasting relationship of adjusted price ratios and future returns is statistically significant, stable over time, and present in out-of-sample tests. We also show that shifts in the steady-state are responsible for the parameter instability and poor out-of-sample performance of unadjusted price ratios that are found in the data. Our conclusions hold for a variety of financial ratios and are robust to changes in the econometric technique used to estimate shifts in the steady-state.

*Lettau: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; mlettau@stern.nyu.edu; Tel: (212) 998-0378; <http://www.stern.nyu.edu/~mlettau>. Van Nieuwerburgh: Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012; svnieuwe@stern.nyu.edu; Tel: (212) 998-0673; <http://www.stern.nyu.edu/~svnieuwe>. We thank Yakov Amihud, John Campbell, Sydney Ludvigson, Eli Ofek, Matthew Richardson, Robert Whitelaw and the seminar participants at NYU and McGill University for comments.

1 Introduction

The question of whether stock returns are predictable has received an enormous amount of attention. This is not surprising because the existence of return predictability is not only of interest to practitioners but also has important implications for financial models of risk and return. One branch of the literature asserts that expected returns contain a time-varying component that implies predictability of future returns. Due to its persistence, the predictive component is stronger over longer horizons than over short horizons. Classic predictive variables are financial ratios, such as the dividend-price ratio, the earnings-price ratio, and the book-to-market ratio (Rozeff (1984), Fama and French (1988), Campbell and Shiller (1988), Cochrane (1991), Goetzman and Jorion (1993), Hodrick (1992), Lewellen (2004), and many others), but other variables have also been found to be powerful predictors of long-horizon returns (e.g., Lettau and Ludvigson (2001), Lustig and Van Nieuwerburgh (2005a), Menzly, Santos, and Veronesi (2004), Piazzesi, Schneider, and Tuzel (2004)). Moreover, these studies conclude that growth rates of fundamentals, such as dividends or earnings, are much less forecastable than returns, suggesting that most of the variation of financial ratios is due to variations in expected returns.

These conclusions are controversial because the forecasting relationship of financial ratios and future stock returns exhibits a number of disconcerting features. First, correct inference is problematic because financial ratios are extremely persistent; in fact, standard tests leave the possibility of unit roots open. Nelson and Kim (1993), Stambaugh (1999), Ang and Bekaert (2001), Ferson, Sarkissian, and Simin (2003), and Valkanov (2003) conclude that the statistical evidence of forecastability is weaker once tests are adjusted for high persistence. Second, financial ratios have poor out-of-sample forecasting power, as shown in Bossaerts and Hillion (1999) and Goyal and Welch (2003, 2004). Third, and related to the poor out-of-sample evidence, the forecasting relationship of returns and financial ratios exhibits significant instability over time. For example, in rolling 30-year regressions of annual Center for Research in Security Prices (CRSP) value-weighted returns on lagged log dividend-price ratios, the ordinary least squares (OLS) regression coefficient varies between zero and 0.5 and the associated R^2 ranges from close to zero to 30% depending on the subsample. Not surprisingly, the hypothesis of a constant regression coefficient is routinely rejected (Viceira (1996), Paye and Timmermann (2003)).

In addition to concerns that return forecastability might be spurious, the benchmark model of time-varying expected returns faces additional challenges. The extreme persistence of price ratios implies that expected returns have to be extremely persistent as well. But if shocks to expected returns have a half-life of many years or even decades, as implied by the high persistence of financial ratios, they are unlikely to be linked to many plausible economic risk

factors, such as those linked to business cycles. Instead, researchers have to identify slow-moving factors that are primary determinants of equity risk. In addition, the extraordinary valuation ratios in the late 1990s represent a significant challenge for the benchmark model. Given the historical record of returns, fundamentals, and prices, it is exceedingly unlikely that persistent stationary shocks to expected returns are capable of explaining price multiples like those seen in 1999 or 2000.

How can models with time-varying expected returns be reconciled with these serious challenges? In this paper, we propose an explanation for the puzzling empirical patterns; namely, the possibility that changes in the steady-state mean of financial ratios are caused, for example, by changes of the steady-state growth rate of economic fundamentals and/or expected return of equity. Why might we expect the steady-state to change over time? Some possibilities include permanent technological innovations that change the long-term growth rate of the economy or improved risk sharing, changes in stock market participation, changes in the tax code, or lower macroeconomic volatility that decrease the long-term expected return of equity.

Changes in the steady-state have dramatic effects on the relationship of returns and price ratios because they cause the mean of the price ratios to change permanently. In particular, predictability regressions are effected because price ratio regressors would be non-stationary. However, as we show below, deviations of the price ratios from their steady-state values are stationary. Thus, the appropriate return forecasting specifications do not include non-stationary price ratios themselves but only stationary deviations from steady states. Our empirical results conclude that such “adjusted” price ratios have favorable properties compared to unadjusted price ratios. For example, forecasting relationships of returns with lagged adjusted price ratios are much more stable over time and have superior out-of-sample forecasting power relative to unadjusted price ratios and to “naive” models. Adjusted price ratios are also less persistent and less volatile than unadjusted price ratios, suggesting that changes in the steady-state account for a fair portion of their overall volatility.

These results show that seemingly incompatible views of the forecasting relationship of stock returns and price ratios can be reconciled if the assumption of a fixed steady-state of price ratios is relaxed. In this sense, the branch of the literature that finds supportive evidence for return predictability and the branch that emphasizes the instability and poor out-of-sample forecastability are both correct. We find that returns are indeed forecastable, but low-frequency shifts in the mean of the price ratios cause the forecasting relationship to be unstable and reduce the out-of-sample forecasting power if financial ratios are not adjusted for the presence of permanent shifts in its mean.

Several papers have explored the impact of structural breaks on return predictability. For example, Viceira (1996) and Paye and Timmermann (2003) reported evidence in favor of breaks in the OLS coefficient in the forecasting regression of returns on the lagged dividend-price ratio.

Our focus is instead on shifts in the mean of financial ratios, which in turn render the forecasting relationship unstable if such shifts are not taken into account. In other words, in contrast to Viceira (1996) and Paye and Timmermann (2003), we focus on the behavior of the mean of price ratios instead of the behavior of the slope coefficient. Pastor and Stambaugh (2001) use a Bayesian framework to estimate breaks in the equity premium. They found several shifts in the equity premium since 1834 and identified the sharpest drop in the 1990s, which is consistent with the timing of the shift in price ratios identified in this paper.

Our paper is also related to the recent literature on inference in forecasting regressions with persistent regressors (see e.g., Amihud and Hurwicz (2004), Ang and Bekaert (2001), Campbell and Yogo (2002), Lewellen (2004), and Torous, Volkanov, and Yan (2004)). In these papers, asymptotic distributions for OLS regressions were derived under the assumption that the forecasting variable is a close-to unit, yet stationary, root process. In contrast, we allow for the presence of a small but econometrically important non-stationary component in forecasting variables.

The rest of the paper is organized as follows. In Section 2 we establish that the standard dividend-price ratio does not significantly forecast stock returns or dividend growth. In contrast, we find much stronger evidence for return predictability in various subsamples. The slope coefficient in the return equation is much smaller in the full sample than in any of the constituent subsamples, which confirms the instability of the forecasting relationship over time. In Section 3.1, we show how changes in the steady-state affect the dividend-price ratio. Filtering out this non-stationary component yields adjusted price ratios that have strong in-sample and out-of-sample return predictability (Section 3). The relationship between returns and adjusted price ratios is stable over time. We consider various methods to estimate the adjusted price ratios. Our first adjustment allows for structural breaks in the mean of financial ratios. For the dividend-price ratio, we show evidence for either one break in the early 1990s or two breaks around 1954 and 1994. In Section 4, we find similar break dates for other valuation ratios such as the earnings-price ratio and the book-to-market value ratio. Our predictability findings continue to hold for these valuation ratios. In Section 5 we use two alternative adjustments. We first estimate a Hamilton (1989) regime-switching model and use real-time estimates of the probabilities of regime switches. This leads to slightly later break dates, but the predictability results are the same. We then use a recursive procedure that estimates both the break dates and the means of the regimes in real-time. Again we find strong predictability in-sample, but the out-of-sample predictability is somewhat weaker because there is substantial estimation uncertainty about the mean of the new regime. In Section 6 we consider a vector error correction model that includes the return and dividend growth predictability equations and imposes a joint present value restriction on the slope parameters from both equations. We find that this restriction is satisfied when we use the adjusted dividend-price ratio as an independent variable, but

not when we use the unadjusted series. Finally, in Section 7, we find that our simple model serves as a plausible data generating process. It is able to replicate both the findings of no predictability when the unadjusted dividend-price ratio is used and the findings of in-sample and out-of-sample predictability when the adjusted series is used.

2 Instability of Forecasting Relationships

In this section we document the instability of the forecasting relationship between returns, dividend growth, and the lagged dividend-price ratio. The forecasting relationship of returns and other financial ratios (such as the earnings-price ratio and the book-to-market ratio) and alternative measures of dividends (such as accounting for repurchases or considering only dividend-paying firms) are similar and will be presented later. The data are based on annual CRSP value-weighted returns from 1927 to 2004 and are described in detail in Appendix A. The top panel of Figure 1 shows the estimation results for the forecasting regression of demeaned returns on the demeaned lagged dividend-price ratio using 30-year rolling windows:

$$r_{t+1} - \bar{r} = \kappa_r(dp_t - \overline{dp}) + \tau_{t+1}^r, \quad (1)$$

where r_t denotes the log return, dp_t denotes the log dividend-price ratio $d_t - p_t$, and \bar{r} and \overline{dp} denote the sample means of returns and the log dividend-price ratio in each of the subsamples, respectively. The top panel plots the slope coefficient κ_r along with two standard error bands. The instability of the forecasting relationship is strikingly illustrated by the variation of the return predictability coefficient over time. The estimates of κ_r are around 0.5 in the subsamples ending in the late 1950s and in the samples ending in the early 1980s to the mid 1990s. In contrast, κ_r is much smaller for the samples ending in the mid 1960s and is close to zero and statistically insignificant in samples ending in the late 1990s and early 2000s. Similarly, the R^2 of the forecasting regression displays instability with values ranging from 34% in 1982 to 0% at the end of the 1990s (not shown). This evidence has led some researchers to conclude that the dividend-price ratio does not forecast stock returns, or at least not robustly so. Not surprisingly, the hypothesis of a constant regression coefficient is routinely rejected.

We also estimate a predictability regression for demeaned dividend growth rates:

$$\Delta d_{t+1} - \bar{d} = \kappa_d(dp_t - \overline{dp}) + \tau_{t+1}^d, \quad (2)$$

where d_t denotes log dividends and \bar{d} denotes the sample mean of dividend growth. Dividend growth rates are even less forecastable than returns. For most of the sample, the point estimate is not statistically significantly different from zero, and the regression R^2 never exceeds 16%

(not shown). Interestingly, the dividend-price ratio at the end of the 1990s seems to forecast neither stock returns nor dividend growth. This is a conundrum from the perspective of any present value model (see Section 3.1).

[Figure 1 about here.]

The left two columns of Table 1, denoted “No Break,” report the coefficients κ_r and κ_d from equations (1) and (2) and their asymptotic standard errors for the entire 1927-2004 sample, as well as various subsamples. The first row shows that the dividend-price ratio marginally predicts stock returns (first column); the coefficient is significant at the 5% level if asymptotic standard errors are used for inference. However, small sample standard errors computed from a bootstrap simulation suggest that the coefficient κ_r is not statistically different from zero for the entire sample.¹ The dividend-price ratio does not forecast dividend growth (third column) at conventional significance levels. Thus, we cannot reject the hypothesis that the dividend-price ratio forecasts neither dividend growth nor returns.

Rows 2 and 3 report the results for two non-overlapping samples that span the entire period: 1927-1991 and 1992-2004. We will justify the particular choice of subsamples in Section 3. The estimates of κ_r display a remarkable pattern across subsamples: In both subsamples κ_r is much larger than its estimate in the whole sample. In fact, the estimates are almost identical in the two subsamples: .2353 in the 1927-1991 subsample compared to .2351 in the later 1992-2004 subsample. Yet, when we join the two subsamples, the point estimate drops to .094. In addition, κ_r is strongly statistically significant in both subsamples but only marginally significant in the whole sample. Confirming the instability of κ_r estimates, row 4 reports the results of a Chow test, which rejects the null hypothesis of no structural break in 1991 at the 4% level. Finally, the dividend growth forecasting relationship displays less instability, and the coefficient remains insignificant in both subsamples.

The pattern of κ_r is not unique to the specific subsamples chosen. We obtain very similar results when we use three non-overlapping subsamples: 1927-1954, 1955-1994, and 1995-2004 (bottom half of Table 1). Again, we find that the return predictability coefficient κ_r is estimated to be much higher in each of the three subsamples than in the entire subsample. In row 5, the predictability coefficient is .09, whereas it is .51, .38, and .53 in rows 6, 7, and 8 respectively.

¹Asymptotic standard errors may be a poor indicator of the estimation uncertainty in small samples, and the p -values for the null of no predictability may be inaccurate. The asymptotic corrections advocated by Hansen and Hodrick (1980) have poor small sample properties. Ang and Bekaert (2001) find that use of those standard errors leads to over-rejection of the no-predictability null. The bootstrap exercise imposes the null of no predictability and asks how likely it is to observe the κ_r coefficients estimated reported in the first column of Table 1. We find that the small sample p -value for κ_r is 6.8% compared to an asymptotic p -value of 4.1%. We also conduct a second bootstrap exercise to find the small-sample bias in the return coefficient. Consistent with Stambaugh (1999), we find an upward bias. If the true value is .094, the bootstrap exercise estimates a coefficient of .115. Detailed results are available upon request.

Moreover, it is statistically significant in each subsample. Row 9 shows that we strongly reject the joint null hypothesis of parameter stability in 1954 and 1994. For dividend growth, the evidence is more mixed. We fail to reject the same null hypothesis of no breaks in 1954 and 1994, but the κ_d coefficient is marginally statistically different from zero in rows 7 and 8.

[Table 1 about here.]

We conclude that the forecasting relationship between returns and the dividend-price ratio is unstable over time. Coefficient estimates of κ_r are almost identical in non-overlapping subsamples, but the point estimate for the whole sample is much lower than it is in each of the subsamples.

Next, we investigate what might explain this intriguing pattern of the regression coefficient that links returns to past dividend-price ratios.

3 Steady-State Shifts and Forecasting

3.1 Changes in the Mean of Price Ratios

The standard specification of stock returns and forecasting variables assumes that all processes are stationary around a constant mean. For example, Stambaugh (1986, 1999), Mankiw and Shapiro (1986), Nelson and Kim (1993), and Lewellen (1999) considered the following model:

$$r_{t+1} = \bar{r} + \kappa_r y_t + \tau_{t+1}^r \quad (3)$$

$$y_t = \bar{y} + v_t. \quad (4)$$

The mean of the forecasting variable y_t , \bar{y} , is constant and the stochastic component v_t is assumed to be stationary, often specified as an AR(1) process. Means of financial ratios are determined by properties of the steady-state of the economy. For example, the mean of the log dividend-price ratio \bar{dp} is a function of the growth rate \bar{d} of log dividends and expected log return \bar{r} in steady-state:

$$\bar{dp} = \log(\exp(\bar{r}) - \exp(\bar{d})) - \bar{d}, \quad (5)$$

whereas the stochastic component depends on expected future deviations of returns and dividend growth from their steady-state values (Campbell and Shiller (1988)):

$$dp_t = \bar{dp} + E_t \sum_{j=1}^{\infty} \rho^{j-1} [(r_{t+j} - \bar{r}) - (\Delta d_{t+j} - \bar{d})], \quad (6)$$

where $\rho = (1 + \exp(\bar{dp}))^{-1}$ is a constant. Similar equations can be derived for other financial ratios (e.g., Vuolteenaho (2000)). Berk, Green, and Naik (1999) show how stock returns and

book-to-market ratios are related in a general equilibrium model.

A crucial assumption is that the steady-state of the economy is constant over time: The average long-run growth rate of the economy as well as the average long-run return of equity are fixed and not allowed to change. However, if either the steady-state growth rate or expected return were to change, the effects on financial ratios and their stochastic relationships with returns would be profound. Even relatively small changes in long-run growth and/or expected return have large effects on the mean of the dividend-price ratio, as can be seen from (5). The effects of steady-state shifts on other valuation ratios, such as the earnings-price ratio and the book-to-market ratio, are similar. In this paper, we entertain the possibility that the steady change of the U.S. economy has indeed changed since 1926, and we study the effect of these changes on the forecasting relationship of returns and price ratios. Why might the steady-state change over time? Possible explanations include technological innovations that raise the long-term growth rate of the economy (\bar{d}), and improved risk sharing, changes in stock market participation, changes in the tax code, or lower macroeconomic volatility that might lower the long-run required return of equity (\bar{r}). Either effect would lower \bar{dp} .

A steady-state is characterized by long-run growth and expected return. Any short-term deviation from steady-state is expected to be only temporary and the economy is expected to return to its steady-state eventually. Thus, steady-state growth and expected return must be constant in expectations, but the steady-state might shift unexpectedly. Correspondingly, we assume that $E_t \bar{r}_{t+j} = \bar{r}_t$, $E_t \bar{d}_{t+j} = \bar{d}_t$, $E_t \bar{dp}_{t+j} = \bar{dp}_t$.²

The framework introduced above can be easily adapted to allow for shifts in the steady-state. The details of the derivation are given in the appendix. Just as in the case with constant steady-state (6), the log dividend-price ratio is the sum of the steady-state dividend-price ratio and the discounted sum of expected returns minus expected dividend growth in excess of steady-state growth and returns:

$$dp_t = \bar{dp}_t + E_t \sum_{j=1}^{\infty} \rho_t^{j-1} [(r_{t+j} - \bar{r}_t) - (\Delta d_{t+j} - \bar{d}_t)], \quad (7)$$

where $\rho_t = (1 + \exp(\bar{dp}_t))^{-1}$. The important difference of (7) compared to (6) is that the mean of the log dividend-price ratio is no longer constant. In fact, it not only varies over time but it is non-stationary. If, for example, the steady-state growth rate increases permanently, the steady-state dividend-price ratio decreases and the current log dividend-price ratio declines permanently. While the log dividend-price ratio contains a non-stationary component it is important to note that *deviations* of dp_t from steady states are stationary as long as deviations

²Although the log dividend-price ratio is a nonlinear function of steady-state returns and growth, we assume that the steady-state log dividend-price ratio is also (approximately) a martingale: $E_t \bar{dp}_{t+j} = \bar{dp}_t$. This assumption is justified for the specific processes for steady-state returns and growth that we will consider below.

of dividend growth and returns from their respective steady states are stationary, an assumption we maintain throughout the paper.³ In other words, the dividend-price ratio dp_t itself contains a non-stationary component \overline{dp}_t but the appropriately demeaned dividend-price ratio $dp_t - \overline{dp}_t$ is stationary. The implications for forecasting regressions with the dividend-price ratio are immediate. First, in the presence of steady-state shifts, a non-stationary dividend-price ratio is not a well-defined predictor and this non-stationarity could be the cause for the empirical patterns described in the previous section. Second, the dividend-price ratio must be adjusted to remove the non-stationary component \overline{dp}_t to render a stationary process.

While we emphasized the effect of steady-state shifts on the dividend-price ratio, the intuition carries through to other financial ratios. Changes in the steady-state have similar effects on the earnings-price ratio and the book-to-market ratio. However, other permanent changes in the economy, such as changes in payout policies, could affect different ratios differently. In the following section, we provide evidence that steady-state shifts have occurred in our sample and propose simple methods to adjust financial ratios for such shifts.

3.2 Steady State Shifts in the Dividend-Price Ratio

Has the steady-state relationship of growth rates and expected returns shifted since the beginning of our sample in 1926? If so, have these shifts affected the stochastic relationship between returns and price ratios? In this section we use econometric techniques that exploit the entire sample to detect changes in the steady-state. In the next section, we study how investors in real-time might have assessed the possibility of shifts in the steady states without the benefit of the whole sample. In both cases, there is strong empirical evidence in favor of changes in the steady-state and we find that such changes have dramatic effects on the forecasting relationship of returns and price ratios. We suggest a simple, yet effective, adjustment to the dividend-price ratio and revisit the forecasting equations from Section 2. We first study shifts in the dividend-price ratios in detail and consider alternative ratios later.

Our econometric specification is directly motivated by the framework that allows for changes in the steady-state laid out in the previous section. Equation (7) implies that the log dividend-price ratio is the sum of a non-stationary component and a stationary component. In this

³Of course, in a finite sample it is impossible to conclusively distinguish a truly permanent change from an extremely persistent one. Thus, our insistence of non-stationarity might seem misguided. However, the important insight is that the dividend-price ratio is not only a function of (less) persistent changes in expected growth rates and expected returns that could potentially have cyclical sources but is also affected by either extremely persistent or permanent structural changes in the economy. This distinction turns out to be very useful, as we will show in the remainder of the paper. In this sense our assumption of true non-stationarity can be regarded to include “extremely persistent but stationary.” In a finite sample, the conclusions will be the same in either setting. The distinction of “permanent” versus “extremely persistent” is important, however, for structural asset pricing models because permanent shocks might have much larger impact on prices than very persistent ones.

section, we model the non-stationary component as a constant that is subject to rare structural breaks as in Perron (1989).

The full line in each of the panels of Figure 3 shows the log dividend-price ratio from 1927 to 2004. Visually, the series displays evidence non-stationarity. Especially the bull market of the 1990s seems hard to reconcile with a stationary model. The dividend-price ratio has gone up since, but at the end of our sample in 2004, prices would have to fall an additional 46% for the dividend-price ratio to return to its historical mean. A first explanation we entertain is that the bull market of the 1990s represents a sequence of extreme realizations from a stationary distribution.

The solid line in Figure 2 shows the smoothed empirical distribution of the log dividend-price ratio dp_t . This distribution has a fat left tail, mainly due to the observations in the last 15 years. To investigate whether this is a typical plot from a stationary distribution, we conduct two exercises. Following Campbell, Lo, and MacKinlay (1997), Stambaugh (1999), Campbell and Yogo (2002), Ang and Bekaert (2001), and many others we estimate an AR(1) process for the log dividend-price ratio. First, in a bootstrap exercise, we draw from the empirical distribution with replacement. The smoothed bootstrap distribution is the dash-dotted line in the figure. Second, we compute the density of dp_t using Monte Carlo simulations from an estimated AR(1) model with normal innovation. This density is plotted as the dashed line. The graph shows that neither the bootstrap nor the Monte Carlo can replicate the fat left tail that we observe in the data. Interestingly, the stationary model also cannot generate the right tail of the empirical distribution. In summary, it is unlikely that the dp_t data sample from 1927 to 2004 was generated by a stationary distribution.

[Figure 2 about here.]

An alternative explanation is that the long-run mean of the log dividend-price ratio is subject to structural breaks. To investigate this possibility, we test the null hypothesis of no break against the alternative hypotheses of one or two breaks with unknown break dates. We use the Perron sup- F test statistic, which is reported in the first row of Table 2. The null hypothesis of no break is strongly rejected (the p -value is less than 1%) in favor of a break in 1991 or two breaks in 1954 and 1994. The test statistic is slightly higher for the two-break case. The null of one break is rejected in favor of the alternative of two breaks at the 5% level but not at the 2.5% level. The null of two breaks against the alternative of three breaks is not rejected (not shown). The data seem to strongly favor one or two breaks, rather than zero or three, but the relative evidence for one or two breaks is not as strong and only slightly in favor of two breaks.

[Table 2 about here.]

In the last column we report the estimated change in the log dividend-price ratio before and after the break. In the one-break case, the change in \overline{dp} is -.86, whereas in the two-break case, the first change in 1954 is -.37 and the second change is -.78. The two plots in the left column of Figure 3 overlay the long-run mean \overline{dp} on the raw dp series. For now, we are agnostic as to whether the break(s) is (are) due to a change in the long-run mean of dividend growth or expected returns, or a combination of the two. We return to this question later. It is worth emphasizing, however, that the date of the shift in the dividend-price ratio is consistent with the breaks in the equity premium identified by Pastor and Stambaugh (2001).

[Figure 3 about here.]

This result motivates us to construct two adjusted dividend-price series, one for the one-break case and one for the two-break case. For each, we simply subtract the mean in the relative subsample(s). In the one-break case with break date τ , the adjusted $dp = d - p$ ratio, \widetilde{dp} is:

$$\widetilde{dp}_t = \begin{cases} dp_t - \overline{dp}_1 & \text{for } t = 1, \dots, \tau \\ dp_t - \overline{dp}_2 & \text{for } t = \tau, \dots, T, \end{cases} \quad (8)$$

where \overline{dp}_1 is the sample mean for 1927-1991 and \overline{dp}_2 is the sample mean for 1992-2004. The adjusted dp ratio in the two-break case is defined analogously. The right column of Figure 3 illustrates this procedure graphically.

The bottom half of Table 2 compares the autocorrelation properties of the unadjusted and adjusted dp series. As is well known, the raw dp series is very persistent. The first and second order autocorrelations are .91 and .81. The null hypothesis of a unit root cannot be rejected, according to an Augmented Dickey Fuller (ADF) test (third column). In contrast, the two adjusted \widetilde{dp} series are much less persistent; the first order autocorrelation drops to .77 and .61, respectively. The null of a unit root in the adjusted series is rejected at the 4% and 1% levels. Interestingly, the volatility of the adjusted series is only half as large as for the unadjusted series (last column). This substantially alleviates the burden on standard asset pricing models to match the volatility of the price-dividend ratio, once the non-stationary nature of the mean dp ratio has been taken into account.

3.3 Forecasting with the Adjusted Dividend-Price Ratio

We now revisit the return and dividend growth predictability equations (1) and (2), but use the adjusted dividend-price ratios instead of the raw series as predictor variable.

The second and fourth columns of Table 1 show the estimation results of the return and dividend growth predictability regressions using \widetilde{dp} , respectively. Rows 1-4 are for the one-break case; rows 5-9 are for the two-break case. Because the adjusted dividend-price ratio is

the same as the raw series with each subsample, the results in rows 2 and 3 are unchanged. But now in row 1, we find that the adjusted dividend-price ratio significantly predicts stock returns. The coefficient for the entire sample is .235, just as in the two subsamples! It is now clear that the low point estimate for κ_r in the left column was due to averaging across regimes. Not taking the non-stationarity of the dp ratio into account severely biases the point estimate for κ_r downwards. Furthermore, row 4 shows that the evidence for a break in the forecasting relationship between returns and the dividend-price ratio has disappeared. The null hypothesis of no structural break can no longer be rejected when using \widetilde{dp} . The full sample regression R^2 is 10%, more than twice the value of the first column. The results for dividend growth predictability remain largely unchanged. This is not surprising given that we did not detect much instability in the relationship between Δd_{t+1} and dp_t to begin with.

The rolling window estimates confirm this result.⁴ The middle panel of Figure 1 shows that the coefficient κ_r is much more stable in the one-break case than in the no-break case (top panel). In particular, its value in the 1990s hovers around .3, compared to 0 without the adjustment. Likewise, the regression R^2 is also more stable and does not drop off in the 1990s. The same exercise shows that the dividend growth relationship is stable and that κ_d never moves far from zero (not shown). The evidence for dividend growth predictability is weak at best.⁵

The bottom panel of Table 1 uses \widetilde{dp} , adjusted for breaks in 1954 and 1994. The full sample estimate for κ_r is now .455 (row 5). The coefficient is highly significant.⁶ The full sample regression R^2 is 22%. Dividend growth is not predictable. The Chow test in row 9 finds no evidence for instability in either forecasting equation. The bottom panel of Figure 1 shows that the rolling estimates for κ_r are very stable when we use \widetilde{dp} adjusted for two breaks. The point estimate hovers around .4. The return regression R^2 goes up as high as 40%.

We conclude that taking changes in the long-run mean of the dividend-price ratio into account is crucial for forecasts of stock returns. In Section 3.1 we extended the model to allow for non-stationarity in \overline{dp} . In this section we examined a simple form of non-stationarity, a structural break. Appropriately adjusting the dividend-price ratio for the structural break strengthens the evidence for return predictability, but not dividend growth predictability. Fi-

⁴In the rolling window estimation we assume that the break in \overline{dp} is caused by a break in mean expected returns \bar{r} . The alternative assumption that the break is in the long-run growth rate of the economy \bar{g} gives identical results.

⁵The lack of predictive power of the dividend-price ratio for dividend growth does not imply that dividend growth is not forecastable because any correlated movement in expected returns and expected dividend growth cancels in $d - p$, as shown in Lettau and Ludvigson (2005).

⁶A bootstrap analysis confirms that the small sample p -value (asymptotic p -value) is 1.11% (0.00%) in the one-break case 0.00% (0.00%) in the two-break case. A second bootstrap exercise shows that the small sample bias in the coefficients is small relative to their magnitude. In the one-break case, the bias is .019 (we estimate .254 when the true coefficient is .235). In the two-break case, the bias is .013 (we estimate .468 when the true value is .455). Detailed results are available upon request.

nally, the in-sample return predictability evidence stands up to the usual problem of persistent regressor bias (Nelson and Kim (1993), Stambaugh (1999), Ang and Bekaert (2001) and Valkanov (2003)) because the adjusted dividend-price ratio is much less persistent.

3.4 Out-Of-Sample Predictability

The previous statements pertained to in-sample predictability. In this section we compare the out-of-sample forecasting properties of the adjusted dividend-price ratio to the unadjusted series and the random walk model.

We follow the approach taken by Goyal and Welch (2003) and predict one-year ahead returns with this year's price-dividend ratio. The first forecasting regression uses 20 years of data, so that the first forecasted return is the one in 1946. For all future years we use expanding windows. We compare four different forecasters. The first is the current sample mean return implied by the 'naive' random walk model. The second is the standard unadjusted dividend-price ratio dp . The third and fourth predictors are the dividend-price ratios $\tilde{d}p$, adjusted for one break in 1991 or two breaks in 1954 and 1994. Panel 1 of Table 3 reports the mean absolute forecast error and the root mean-squared forecasting error. Comparing the second to the first row, we find the result of Goyal and Welch (2003): the random walk model has superior out-of-sample properties compared to the standard dividend-price ratio specification. The latter's prediction errors are almost 1% per year higher. However, using the adjusted dividend-price ratios substantially reduces the out-of-sample forecasting error. The mean absolute error (MAE) and root mean squared error (RMSE) of the dividend-price ratio adjusted for a single break are lower than those for the random walk model. Even more impressively, the out-of-sample forecasting power of the dividend-price ratio adjusted for two breaks is dramatically improved compared to the unadjusted ratio and to the random walk model. The RMSE and MAE are reduced by 12-15% compared to the random walk model.

[Table 3 about here.]

The results reported in this section show that forecasting with the unadjusted dividend-price ratio series results in coefficient instability in the forecasting regression, unreliable inference (insignificance in small samples, and results depending on the subsample), and poor out-of-sample predictability. These disconcerting properties are due to a non-stationary component that shifts the mean of the dividend-price ratio. Adjusting for such changes in the mean of the dividend-price ratio improves the forecasting relationship between returns and the dividend-price ratio in all three dimensions. The predictability coefficient is stable over time, least squares coefficient estimates are highly significant, and the adjusted dividend-price ratio has forecasting power out-of-sample.

4 Other Financial Ratios

While the dividend-price ratio has been the classic prediction variable at least in the academic literature, it is useful to investigate to what extent our results are robust to a different measure of payouts. Lamont (1998) finds that the log earnings-price ratio ep forecasts returns. We find very much the same patterns for the earnings-price ratio as for the dividend-price ratio. The earnings data start in 1946 and are described in Appendix A. The book-to-market ratio is computed from the same earnings and dividend data using the clean-surplus method.

[Table 4 about here.]

Table 4 shows that the null hypothesis of no structural break in the ep ratio is strongly rejected in favor of one or two breaks (first row). The Perron test estimates a 1990 break date in the one-break case and 1953 and 1994 break dates in the two-break case. These line up almost perfectly with the dp break dates in Table 2. One other often used valuation ratio, the log book-to-market ratio (bm) also displays strong evidence of two breaks with similar break dates in 1953 and 1990. Clearly, there is evidence for a permanent or strongly persistent component in all valuation ratios.

Some researchers have argued that there were persistent changes in firms' payout policies in the 1990s and have argued to adjust dividend-price ratios for repurchases (Fama and French (2001), Grullon and Michaely (2002), and Boudoukh, Michaely, Richardson, and Roberts (2004)). First, we find no evidence for a break in the payout ratio $de = d - e$ at the 10% level (row 3 of Table 2). This is consistent with the view that both \overline{dp} and \overline{ep} contain structural breaks. Second, even if there was a break in the \overline{de} ratio and we took the point estimate for \overline{de} in the subsamples, we would find that more than three-fourths of the change in the dividend-price ratio comes from a change in the earnings-price ratio and less than one-fourth from a change in \overline{dp} : $\overline{dp} = \overline{ep} + \overline{de}$. In particular, for our S&P 500 sample from 1946-2004 with a 1991 break date, we find a change in \overline{dp} of -.81, a change in \overline{ep} of -.63, and a change in \overline{de} of -.18.

To further investigate the role of repurchases and the role of a changing composition in CRSP, we consider two additional valuation ratios. First, we consider the CRSP universe without NASDAQ stocks. Arguably, removing NASDAQ stocks goes a long way towards eliminating new economy and non-dividend paying companies that became more prevalent in the 1990s.⁷ We compute the dividend-price ratio, dp^{nas} , and the dividend growth for this group. This time

⁷Fama and French (2001) document that the fraction of non-financial, non-utility firms that paid dividends declined by almost 45% between 1978 and 1999. However, most of that decline is attributable to new firms and to small firms. They write: "The characteristics of dividend payers (large profitable firms) do not change much after 1978." We take this group to be the value-weighted CRSP index without NASDAQ stocks. This series starts to deviate from the full sample series in 1973. We verified that the dividend growth rate of this set of firms did not change in the 1990s. Average dividend growth from 1927-1991 was 5.45%. Average dividend growth from 1992-2004 was 5.46%.

series has properties very similar to those of the series with the NASDAQ. The Perron tests in row 5 of Table 2 show a break of 75% in 1992, close to the 86% change in the full sample series in 1991. For the two-break case, the break dates and magnitudes are also very similar: 1954 and 1995 and -35% and -70%.

Second, we use the Boudoukh, Michaely, Richardson, and Roberts (2004) repurchase yield data, available from 1971 onwards, and construct a corrected dividend-price ratio and dividend growth rate series. We label this repurchase-adjusted dividend price series dp^{rep} .⁸ The case favored by the data is a three-break case with break dates in 1957, 1973, and 1990 (see row 6 of Table 2). We show below that these two adjustments do not materially affect our predictability results with the standard dividend-price ratio presented earlier. This leads us to conclude that structural changes in payout policies and/or the composition of firms in the 1990s can only explain a small part of the change in the dividend-price ratio.

We first turn to the in-sample forecasting regressions with the earnings-price ratio. When we use the earnings-price ratio as a return predictor, we obtain similar results to what we reported for the dividend-price ratio in Table 1. The first row of Table 5 shows that when the unadjusted earnings-price ratio is the independent variable, the slope coefficient is .119. Just like with the dividend-price ratio regressions, this coefficient displays parameter instability among subsamples: The full sample point estimate is lower than the estimates in all subsamples, and the Chow test of no break has a p -value of only .15 (not shown). The next two columns show that this bias is due to averaging over subsamples. Once we use the adjusted \tilde{ep} ratio, the instability disappears and the full sample point estimate increases to .215 in both the one-break and two-break case. These coefficients are twice the size of the ones obtained with the unadjusted ep ratio and are measured precisely. The regression R^2 almost doubles. The slope coefficient is very similar to the one we found in the first panel of Table 1: $\kappa_r = .235$. One difference from the results in Table 1 is that the adjusted earnings-price ratio also significantly forecasts earnings growth, with a negative sign (not shown).

[Table 5 about here.]

For the unadjusted lagged log book-to-market ratio $bm = b - m$, we find that the predictability coefficient κ_r is only marginally significant. The point estimate is .07, lower than the point estimates in the subsamples 1927-1952 (.26), 1953-1990 (.44), and 1991-2004 (.72), all of which are strongly significant. Again, this downward bias is due to averaging over the break(s). The full sample point estimate increases to .255 with the one-break adjusted \widetilde{bm} series as regressor and to .308 with the two-break adjusted \widetilde{bm} series. The regression R^2 increases from 3% in the first column to 19% in the last column.

⁸We note that this is just one possible adjustment. The correct adjustment depends on the investor under consideration (e.g., an investor's cash flows are adjusted for aggregate repurchases, but not for seasoned equity offerings nor initial public offerings).

The return predictability findings for dp^{rep} and dp^{nas} are similar to the benchmark dp results. First, using dp^{nas} without break adjustment, we find a point estimate for κ_r of .11. This point estimate is lower than in either subsample (.24 for 1927-1992 and .30 in 1993-2004). Once we use the break-adjusted series, the point estimate more than doubles to .250. Just as for the standard dp ratio, the downward bias comes from averaging over the break. The break-adjusted point estimate is close to the .235 we found for the sample that includes the NASDAQ. We obtain further increases in the point estimate and the R^2 in the two-break case. Second, using dp^{rep} , the full sample return predictability coefficient is .19, higher than the .09 for the standard dp series, but again lower than in either subsample (.25 in 1927-1990 and .53 in 1991-2004). Clearly, adjusting for repurchases improves the forecasting power of the dp ratio. However, adjusting for the breaks is important and further strengthens the case for predictability. In the preferred case of three breaks, the predictability coefficient is .58, three times its unadjusted value. The regression R^2 is also three times higher.

Finally, we note that the out-of-sample predictability results for the earnings-price ratio are in line with the results with the dividend-price ratio reported in Table 3. When we use the unadjusted earnings-price ratio to forecast one-year ahead returns, the mean absolute forecasting error is 13.44% on average, higher than for the random walk model (13.31%). However, when we use the break-adjusted earnings-price ratio, the forecast error is only 11.8% per year (in both the one-break and two-break cases).

We conclude that the other financial ratios indicate a predictability pattern similar to that of the dividend-price ratio. Without the adjustment for a change in its long-run mean, the relationship between one-year ahead returns and the financial ratios is unstable over time, the evidence for forecastability is weaker, and the out-of-sample forecasts do not beat the simple random walk model. However, once we filter out the non-stationary component, we find a stable forecasting relationship, a large predictability coefficient, and out-of-sample forecasting power that improves significantly on the random walk model. The fact that the results are so similar for earnings and dividend data suggests that an explanation that exclusively rests on changing payout policies misses the most important structural changes in the economy: changes in long-run growth rate or long-run expected returns.

5 Real-Time Dividend-Price Adjustments

Implicit in the previous exercise is the assumption that the investor is endowed with the knowledge of the break dates. Both in-sample and out-of-sample predictability results used an ex-post break adjustment with the knowledge of the entire data sample. In this section we relax this assumption and consider real-time (ex-ante) estimates of the adjusted dividend-price ratio only. We restrict the time- t information set of investors to include only past values of dividends,

prices and returns. This is a rather restrictive assumption because other variables might help investors estimate changes in the steady-state with higher accuracy. In this sense, the results in this section represent a conservative case to gauge the real-time forecasting power of adjusted dividend-price ratios.

5.1 Estimation

We start by arguing that an investor who only had access to real-time data would infer similar break dates as the ones from Section 3. The only difference is that the break dates are estimated slightly later than in the ex-post break tests. This would be the natural outcome in a model in which the investor has to learn the break date(s). We entertain two approaches: a recursive mean estimation and a regime-switching model.

Before showing the estimation results, we first confirm our earlier conclusion (based on ex-post data, see Figure 2) that it is extremely unlikely that the dividend-price ratio sample is drawn from a stationary distribution, based on real-time data only. Figure 4 shows a recursive (i.e., real-time) estimation of the empirical distribution of the log dividend-price ratio. In each year, the investor estimates an AR(1) model for dp , using data up to the current year. she then bootstraps from the available sample to compute the empirical distribution of the log dividend-price ratio. Each year she recomputes the 2.5, 5, 95, and 97.5 percentiles of the bootstrapped distribution (dashed and dotted lines in the figure). The figure also plots the realized dividend-price ratio, in deviation from its recursive sample mean (full line). In 1958, the investor is quite confident that the realized dividend-price ratio is far below the mean; it hits the 2.5 percentile of the empirical distribution. Likewise, in 1994, the observed dividend-price ratio falls in the 2.5% tail of the distribution. Between 1995 and 1999, the investor is almost certain that the observed dividend-price ratio has *not* been drawn from a stationary distribution. Interestingly, these ‘crossing’ dates are almost identical to the break dates estimated in the previous section. This shows that the permanent changes in the dividend-price ratio that were identified by the ex-post break tests do not rely on having the benefit of the entire sample through 2004. Even an investor in real-time would have concluded that extreme observations of dividend-price ratios are unlikely to be generated by a stationary process.

[Figure 4 about here.]

We repeat the forecasting procedures with two alternative real-time adjustments of the dividend-price ratio. First, we consider a real-time adjusted dividend-price ratio series \widehat{dp}^P estimated from the Perron model. In each year t , the investor estimates the Perron structural break test using all data available up to year t . she uses the estimated mean dp ratio associated with the last regime, \overline{dp}_t^P , as the estimate for the time- t long-run mean of the dividend-price

ratio. The adjusted dividend-price series is built up recursively, where the year t observation is given by

$$\widehat{dp}_t^P = dp_t - \overline{dp}_t^P. \quad (9)$$

The initial sample is 1927-1951, so the first adjustment pertains to 1951. The left panel of Figure 5 plots the adjusted series \widehat{dp}_t^P alongside the raw $d - p$ series. Comparing Figure 5 with Figure 3 shows that the recursively adjusted dividend-price ratio closely resembles the behavior of the dividend-price series adjusted using the two ex-post break dates.

[Figure 5 about here.]

Alternatively, the investor estimates a Hamilton regime-switching model in real-time (Hamilton (1989)). The top panel of Figure 6 shows the real-time estimate of the probability that the dividend-price ratio is drawn from the low regime, when two regimes are considered. Starting in 1991, the investor puts non-zero probability on a shift to the low $d - p$ regime. By 1995, she is more than 50% certain that the shift occurred. The mean estimates for the two regimes coincide with our ex-post estimates. When she considers three regimes instead, the investor increases the probability of a switch from the high to the middle dp regime in 1954. By 1960 she is more than 50% certain that the first shift occurred (see middle panel of Figure 6). In 1990, she starts to attribute probability mass to the low- p regime, and by 1996, she is more than 50% confident that the economy left the middle- dp regime for the low- dp regime.

[Figure 6 about here.]

We define a second, real-time adjusted log dividend-price ratio series \widehat{dp}^H for the Hamilton model:

$$\widehat{dp}_t^H = dp_t - \overline{dp}_t^j, \quad (10)$$

where \overline{dp}_t^j is the mean associated with the regime that has the highest posterior likelihood given data up to period t . So, the adjusted dividend-price series from the regime-shift model, \widehat{dp}^H , are constructed similarly to the ex-post series \widetilde{dp} . The only difference is the timing of the breaks: 1995 instead of 1991 in the two-regime case and 1960 and 1996 instead of 1954 and 1994 in the three-regime case. The right panel of Figure 5 shows the adjusted dividend-price ratio series, \widehat{dp}^H for the two-regime and three-regime case and compares them to the raw dp series.

As was the case with the two ex-post adjusted series (\widetilde{dp}), the real-time adjusted series (\widehat{dp}) are much less persistent than the raw dp series. The null of a unit root is strongly rejected.⁹

⁹Results available upon request.

5.2 Forecasting

We repeat the in-sample estimation of the forecasting equations (1) and (2) with the real-time adjusted dividend-price ratio series \widehat{dp}^P and \widehat{dp}^H . Table 6 compares the results of using these adjusted series to those using the unadjusted series. Columns 2 and 3 report the results for \widehat{dp}^P . The point estimate for κ_r is .30 and is measured precisely, indicating that returns are forecastable by this real-time adjusted dividend-price ratio. This value for κ_r is between the values based on the ex-post adjustment with one break ($\kappa_r = .23$) and two breaks ($\kappa_r = .45$) from Section 3.3. This makes sense, because the Perron break model is re-estimated every period. On the other hand, the dividend growth predictability coefficient κ_d is not statistically different from zero. Using the Chow test, we also verified that the relationship between returns and dividend growth on the one hand and the Perron-adjusted dividend-price ratio on the other hand is stable across these break dates. We now turn to the forecasting regressions with the Hamilton adjustment.

Columns 4-7 show the forecasting regressions when we use \widehat{dp}^H as the predictor. The results are similar to those using the ex-post adjustment in Table 1. In particular, in the two-regime (one-break) model, the point estimate for κ_r is .218 with a small standard error. Having the (real-time) break in 1995 instead of in 1991 (ex-post break date) leaves the predictability coefficient largely unchanged; it is .235 in the first panel of Table 1. In the three-regime (two-break) case, the point estimate is .353 (row 6), somewhat lower than the .455 estimate with ex-post break adjustment, but still strongly significant. The R^2 is 14%, 8% lower than in the ex-post adjustment, but still 10% higher than in the no-adjustment case. Thus, even though the estimation of the break date entails some uncertainty in small samples, Table 6 confirms that our results are robust to a plausible change in the break date.

[Table 6 about here.]

Figure 7 plots estimates of the slope coefficient κ_r in 30-year rolling regressions using the three different real-time adjustments. The return forecasting regressions using the real-time dividend-price ratio is significantly more stable than those for the unadjusted dividend-price ratio plotted in Figure 1.

[Figure 7 about here.]

Comparing the coefficient estimates for κ_r highlights one important distinction between the Perron and Hamilton adjustments. In the Hamilton adjustment, the investor estimates the probability of each regime in real time but has access to the entire time series to estimate the mean dp ratio in the different regimes. In the Perron procedure, the investor not only estimates the break dates in real-time, but also the long-run mean in the current regime. Clearly,

estimating that new long-run mean based on a few data points incurs a lot of measurement error. For example, if the Perron investor in 2000 detects a break in 1997, she only has three data points to estimate the new \overline{dp}^P . The Hamilton investor knows what the new mean \overline{dp} is. A satisfactory resolution that avoids any look-ahead bias in the Hamilton model would be to conduct a full Bayesian analysis in which the investor estimates the number of regimes, the regime switching dates, and their associated long-run means based on real-time available data and prior information. Pettenuzzo and Timmermann (2004) have worked out such an estimation in the context of an asset allocation problem.

5.3 Out-Of-Sample Predictability

Finally, we return to the out-of-sample forecastability question. The results with the Hamilton procedure are similar to the out-of-sample predictions with ex-post break adjustment. Rows 5 and 6 in Panel 1 of Table 3 show that the RMSE using \widehat{dp}^H in the three-regime (two-regime) case is .1498 (.1548), lower than the random walk model (.1605) and the model with the unadjusted dp ratio (.1685). The same is true for the MAE.

The last row of the same panel reports the MAE and RMSE using the recursive Perron dividend-price ratio \widehat{dp}^P . Both out-of-sample forecasting errors are smaller than those obtained with the unadjusted dividend-price ratio. The MAE is marginally smaller than it is for the random walk model, but the RMSE is slightly larger. As expected, the performance of the real-time adjusted out-of-sample prediction is somewhat worse than the one reported in Section 3.4.

As we pointed out, this goes back to the difficulty in estimating the mean of a new regime, after a structural break occurs. In real life the investor has access to data other than the dividend-price ratio to make inference about the new long-run mean of dividend growth or expected returns. Indeed, Ludvigson and Ng (2005) found that two new factors, a volatility and a risk-premium factor, summarize the information in a large number of financial series and have strong forecasting power for future returns that is largely orthogonal to commonly used valuation ratios such as the dp ratio. This is the sense in which the Perron method provides a worst-case scenario for the investor.

6 Long-Horizon Predictability

An important component of the empirical work on return predictability uses long-horizon regressions. In this section we provide a framework for analyzing long-horizon predictability. We use the ex-post dividend-price ratio as the predictor and derive theoretical restrictions that link return and dividend growth to lagged dividend-price ratios at different horizons. Rather than

estimating regressions for various horizons separately, the advantage of this approach is that fewer parameters have to be estimated. Moreover, the restrictions impose that the estimates across horizons are consistent with each other.

Recall the return and dividend growth predictability equations (1) and (2)

$$\tilde{r}_{t+1} = \kappa_r \widetilde{dp}_t + \tau_{t+1}^r \quad (11)$$

$$\Delta \tilde{d}_{t+1} = \kappa_d \widetilde{dp}_t + \tau_{t+1}^d, \quad (12)$$

where variables with a tilde are appropriately demeaned and stationary. Subtracting (12) from (11) and using the log-linear approximation for log returns $\tilde{r}_{t+1} = \widetilde{dp}_t - \rho \widetilde{dp}_{t+1} + \Delta \tilde{d}_{t+1}$ yields the implied AR(1) process for the dividend-price ratio in (13):

$$\widetilde{dp}_{t+1} = \phi \widetilde{dp}_t + \tau_{t+1}^{dp}, \text{ where} \quad (13)$$

$$1 - \rho\phi = \kappa_r - \kappa_d \quad (14)$$

$$\rho\tau_{t+1}^{dp} = \tau_{t+1}^d - \tau_{t+1}^r, \quad (15)$$

where the innovations are linked by $\rho\tau_{t+1}^{dp} = \tau_{t+1}^d - \tau_{t+1}^r$. The model imposes a non-linear present value restriction (14) on the predictability coefficients κ_r and κ_d . Because $\rho < 1$ and stationarity implies $|\phi| < 1$, $\kappa_r - \kappa_d$ must be positive. This is another way of saying that either returns ($\kappa_r \neq 0$) or dividend growth ($\kappa_d \neq 0$) have to be forecastable (or both). Most researchers work with equations (11) and (13); we work with (11) and (12) instead because long-horizon restrictions are more easily derived in this case. Iterating forward on equations (11) and (12), we obtain the H -period dividend growth and return forecasting equations

$$\frac{1}{H} \sum_{j=1}^H \tilde{r}_{t+j} = \kappa_r(H) \widetilde{dp}_t + \tau_{t,t+H}^r \quad (16)$$

$$\frac{1}{H} \sum_{j=1}^H \Delta \tilde{d}_{t+j} = \kappa_d(H) \widetilde{dp}_t + \tau_{t,t+H}^d, \quad (17)$$

where

$$\kappa_d(H) = \kappa_d \frac{1}{H} \left(\frac{1 - \phi^H}{1 - \phi} \right) \quad (18)$$

$$\kappa_r(H) = \kappa_r \frac{1}{H} \left(\frac{1 - \phi^H}{1 - \phi} \right). \quad (19)$$

Let N be the number of horizons $H > 1$. Then the joint system of one-year ahead and H -year ahead predictability regressions for returns and dividend growth contains $2 + 2N$ equations but

only two free parameters (κ_d, κ_r) . The parameter ϕ is implied by the present value constraint (14). The typical approach in the literature is to estimate univariate long-horizon predictability equations without imposing these restrictions. Instead, we estimate the entire system of long-horizon regressions jointly. This estimation procedure takes the high correlation of regression coefficients at different horizons, pointed out by Boudoukh, Richardson, and Whitelaw (2005), explicitly into account.

The estimation routine we describe below finds parameters to match the entire ‘term structure’ of univariate predictability coefficients. Figure 8 illustrates this procedure for $H = \{1, 3, 5, 7, 10\}$. The top panel plots the univariate predictability coefficients obtained from standard OLS regressions of H -year ahead returns and dividend growth on the unadjusted dividend-price ratio dp , as well the predictability coefficients $(\kappa_d(H), \kappa_r(H))$ implied by the model. To match the pattern of the ten OLS predictability coefficients, the optimization routine chooses a value for κ_r above the value from the one-year ahead univariate regression.

[Figure 8 about here.]

We start by estimating the one-period ahead equations for returns and dividend growth (11-12). These are the same equations we estimated in Section 3.3, but the model provides two new insights. First, the estimation delivers a value for the autocorrelation coefficient of the dividend-price ratio ϕ because we impose the present value constraint (14). Second, we use the break adjusted series \widetilde{dp} from Section 3 in the estimation.

Because return and dividend growth series were demeaned by their sample averages, the previous sections implicitly assumed a break in \overline{dp} without associated break in \bar{r} or \bar{d} . The model tells us that a break in \overline{dp} must be associated with a break in either \bar{r} or \bar{d} , or both. First, we assume that \bar{d} is constant and focus on changes in the expected returns, consistent with the evidence on breaks in the equity premium in Pastor and Stambaugh (2001). The implied change in \bar{r} with the change in \overline{dp} can be inferred from $\bar{r}_t = (1 + \bar{d}) \exp(\overline{dp}_t) + \bar{d}$. The top panel of Table 7 shows how large the change in \bar{r} is (row 2) corresponding to the change in the mean *dividend-price* ratio in the data (row 1). The left panel is for the one-break case, the right panel for the two-break case. The observed change in \overline{dp}_t implies a decline in mean expected returns of 2.6% in 1991 or a dual decline of 1.7% in 1954 and 2% in 1994, assuming long-run dividend growth did not change. Alternatively, it can stem from an increase in long-run dividend growth of 2.5% in 1991 or a dual increase of 1.6% in 1954 and 2% in 1994, when mean expected returns are held constant. In the results reported below, we choose to correct \bar{r}_t , but this choice turns out to be unimportant for the point estimates.¹⁰ The second panel

¹⁰The reason is that returns and dividend growth are very volatile, compared to the change in their mean implied by the change in \overline{dp} . For the same reason, the results reported below are virtually identical if we assume that the break takes place in \bar{d} instead. This validates the results in Section 3.3.

reports the change in the mean earnings-price ratio in the various subsamples, as well as the implied change in the long-run mean return or long-run mean dividend growth rate. Appendix A describes how the latter two are computed. We find that the change in the mean return of -2.47% that accounts for the change in the earnings-price ratio \bar{ep} (bottom panel) is very similar to the -2.60% change that accounted for the change in the dividend-price ratio \bar{dp} (top panel). The same is true when the long-run growth rate does all the adjustment.

[Table 7 about here.]

Panel A of Table 8 reports the estimation of the one-year ahead system. Row 1 uses the unadjusted dp series, whereas rows 2 and 3 use the adjusted series \widetilde{dp} for the one-break and two-break case respectively. The GMM estimation uses the OLS normal conditions to estimate κ_d and κ_r . Therefore, the point estimates are identical to the ones reported in Table 1. Three differences are worth pointing out. First, the adjustment in \bar{r} delivers slightly lower standard errors for κ_r in rows 2 and 3. Second, as foreshadowed by Table 2, the point estimates for ϕ are substantially lower when we use the adjusted dividend-price ratio: .81 and .69 in rows 2 and 3 compared to .95 in row 1. Third, the first number in the last column reports the violation of the present value constraint (14) by the univariate (OLS) coefficient estimates $\kappa_{d,ols}$ and $\kappa_{r,ols}$, expressed in the same units as κ_d and κ_r . Row 1 shows that using the unadjusted dp ratio leads to violations of the present value constraint. They are half the size of the estimated κ_r . Yet, when we use the adjusted dp ratio, constraint (14) is satisfied.

[Table 8 about here.]

Next, we estimate the one-period ahead equations for returns and dividend growth (11-12) jointly with the long-horizon regressions (16-17). We select a small number ($N = 2$) of long-horizon moments, corresponding to $H = \{1, 3, 5\}$. The joint system of one-year, three-year, and five-year ahead predictability regressions for returns and dividend growth contains $2 + 2N = 6$ equations and $2N = 4$ restrictions. Panel B of Table 8 reports the results.

The point estimates for (κ_d, κ_r) are similar to those obtained from the one-year ahead system in panel A. Row 4, which uses the unadjusted dp ratio, fails to find evidence for return predictability or dividend growth predictability at the 5% level. The point estimate falls from .094 in panel A to .068 in panel B. Using long-horizon information makes the case for return predictability weaker when the unadjusted dp series is used. Furthermore, the estimate for $\phi = .99$ and its standard error, implied by the estimates for κ_d and κ_r through (14) and computed using the delta method, indicate that we cannot reject the null hypothesis of a unit root in the dividend-price ratio.

The results in rows 5 and 6 look very different. Once we use the adjusted ratio \widetilde{dp} , we find strong evidence for return predictability. The point estimates remain large: .210 in the

one-break case and .409 in the two-break case. Moreover, the asymptotic standard errors on κ_r are reduced. The reason is that we use restrictions of the term structure of predictability coefficients that cannot be uncovered by estimating the long-horizon moments in isolation; imposing these constraints improves the inference on κ_r and κ_d . Put differently, the univariate OLS long-horizon coefficients violate the present value constraint. The first number in the last column reports the average violation across the three constraints (the RMSE); it is 19% in row 1. This violation is lower in rows 5 and 6. The second number in the last column reports the average moment violation (RMSE) and measures the degree to which the imposed restrictions are satisfied. In row 5, the average moment violation is only 8.5%, less than half as big as in row 4.

The middle and bottom panels of Figure 8 graphically illustrate the long-horizon estimation results with the adjusted dividend-price ratio for the larger $\{1, 3, 5, 7, 10\}$ -year system. First, the one-period coefficients are very similar to the ones we reported for the $\{1, 3, 5\}$ -year system. Second, the two panels show that when the adjusted dividend-price ratio is used, the pattern of GMM long-horizon regression slopes generated by our model lines up almost perfectly with the univariate OLS regression slopes. Put differently, the normal conditions for the long-horizon moments are satisfied at the model-implied long-horizon predictability coefficients. This contrasts with the first panel, which is for the unadjusted dp ratio. There, the OLS coefficients are quite different from the GMM estimates. The correction supports both our specification and our main argument.

The results for the $\{1, 3, 5\}$ -year and $\{1, 3, 5, 10\}$ -year systems are representative of the results we found for different numbers of long-horizon moments and choices of horizons. Imposing long-horizon information confirms the results of the earlier sections: Returns are predictable by the dividend-price ratio, once its non-stationary component is removed.

7 Monte Carlo Simulations

In this section, we provide further evidence that the model in (11-13) captures the moments of the data well. We use a Monte Carlo exercise to show *(i)* that this model replicates the failures that are found in the in-sample and out-of-sample predictability literature using the unadjusted dp ratio, *(ii)* that it matches the moments once the dividend-price ratio is properly adjusted for and after taking into account small sample inference issues.

For the Monte Carlo exercise, we specify a structural model for the joint behavior of expected returns, expected dividend growth, and dividend growth innovations. In Appendix C, we derive the regression residuals $\tau = (\tau^d, \tau^r, \tau^{dp})$ as functions of the structural innovations and show how to identify the structural parameters from the parameters of the vector error correction model (VECM) in equations (11-13). We back out the structural parameters from estimates

of the VECM parameters and simulate the structural model generating 10,000 time series for returns, dividend growth and the dividend-price ratio of length $T = 78$. We then run univariate predictability equations on the model-generated data and compare the parameter estimates to the true predictability coefficients and to the data. The Monte Carlo exercise also serves as a way to investigate the small sample properties of the regression coefficients. Appendix D describes the algorithm in detail.

We start by studying the properties of univariate return predictability regressions. The first row of Table 9 reports the ‘true’ predictability coefficients for the one-year, three-year, and five-year horizon equations, as well as the theoretical regression R^2 .¹¹ In the top panel, the true parameters come from an estimation of the VECM that assumes a break in dp in 1991; in the bottom panel we specify two breaks in 1954 and 1994.

The second and fourth rows report the same coefficients and R^2 estimated on simulated data. Row 2 uses the unadjusted dp ratio, and row 4 the adjusted \widetilde{dp} to predict returns. The second row shows that the model with the unadjusted dp ratio fails to detect the return predictability that is present in the data. The simulation-based estimate of κ_r and the regression R^2 are close to zero and do not increase enough with the horizon. The one-year ahead coefficient is not significant when we use the small-sample standard error. Moreover, this failure of the model matches the failure in the data. Row 3 shows the results of the same univariate regressions in the data using the unadjusted dp ratio. The slope coefficients, standard errors, and R^2 line up closely with the results from the Monte Carlo simulation without break adjustment. In row 4, we adjust each Monte Carlo series for a break in 1991 (top panel) or two breaks in 1954 and 1994 (bottom panel). The predictability coefficients at all horizons now line up closely with their true values. The model with adjusted dp ratio recovers the predictability pattern of row 1. Moreover, the predictability coefficients, their standard errors, and the regression R^2 from the model in row 4 match the ones from regressions of observed returns on observed adjusted dividend-price ratios (row 5). Results for dividend growth regressions are not reported but the simulations recover the lack of predictability in the true κ_d coefficients implied by the VECM.

[Table 9 about here.]

Comparing row 4 to row 1, we notice that there is some small sample bias. In line with the findings of Stambaugh (1999), the estimate for κ_r is upward biased. In the first panel, the true value of the slope coefficient in the one-year ahead return regression is .222 versus .264 in simulation, a bias of .042. At the five-year horizon, the upward bias is only .016. Likewise, the R^2 of the regression is slightly upward biased in the simulation: 10.5% versus 8.9% at one-year horizon and 28% versus 25% at the five-year horizon. In the two-break case reported in the

¹¹As Campbell, Lo, and MacKinlay (1997) point out, there is no closed form solution for the long-horizon R^2 . We approximate it by simulating the structural model for 100,000 periods.

second panel, the bias is smaller. The one-year ahead coefficient estimate is .483 versus the true value of .456. The upward bias disappears at the five-year horizon. Overall, the bias is small relative to the magnitude of the coefficients, and therefore does not affect our conclusions. Finally, the small sample standard errors, averaged across Monte Carlo simulations of the same length of the data, are very similar to the asymptotic standard errors from the data (middle columns).

Panels 2 and 3 of Table 3 report the out-of-sample prediction errors from a second Monte Carlo exercise. We simulate the structural model under the null hypothesis that the data generating process has one break in 1991 (panel 2) or has two breaks in 1954 and 1994 (panel 3). We compare the same three out-of-sample forecasting exercises as in the data (panel 1). In addition, we look at the forecast errors when we only correct for the second break and not the first one. The ‘true’ parameters in panels 2 and 3 are the same and were obtained from the VECM parameters estimated under the assumption of two breaks in 1954 and 1994 for the period 1947-2004, the same forecasting period as in panel 1. The details are shown in Appendix D. The Monte Carlo exercise regenerates the pattern we found in the data. When the unadjusted dividend-price ratio is used as a forecasting variable, the out-of-sample prediction errors are large and close to the random walk errors. On the other hand, when we implement the one-break or two-break adjustment in the model, the simulated data generate substantially lower prediction errors, mimicking the improvement in the data.

The simple model (11-13) replicates the patterns of univariate one-year ahead and long-horizon regression results in-sample as well as one-year ahead out-of-sample prediction errors found in the data. In particular, it regenerates (i) the failures of using the unadjusted dp ratio as a predictor, and (ii) the successes of using the adjusted \widetilde{dp} ratio.

8 Conclusion

The macroeconomics literature has recently turned to models with persistent changes in fundamentals to explain the dramatic change in valuation ratios in the bull market of the 1990s. Most such models imply a persistent decline in expected returns or in the equity premium. Lettau, Ludvigson, and Wachter (2004) argued that a persistent decline in the volatility of aggregate consumption growth leads to a decline in the equity premium. Another class of models argues for persistent improvements in the degree of risk-sharing among households or regions, either due to developments in the market for unsecured debt or the market for housing-collateralized debt (Krueger and Perri (2005) and Lustig and Van Nieuwerburgh (2005b)). In the model of Lustig and Van Nieuwerburgh (2005c), the improvement in risk sharing implies a persistent decline in the equity premium. McGrattan and Prescott (2005) argued that persistent changes in the tax code can explain the persistent decline in the equity premium. Lastly, models of

limited stock market participation argue that the gradual entry of new participants has persistently depressed equity premia (Vissing-Jorgensen (2002), Calvet, Gonzalez-Eiras, and Sadini (2003), and Guvenen (2003)). Other models argue that there was a persistent increase in the long-run growth rate of the economy in the 1990s (Quadrini and Jermann (2003) and Jovanovic and Rousseau (2003)).

In this paper we argue that such a persistent downward change in expected returns and/or a persistent upward change in the long-run growth rate of the economy leads to a persistent decline in the mean of financial ratios. Such breaks in the mean of valuation ratios have important effects on estimation and inference of return forecasting regressions. We consider various econometric techniques to detect shifts in the mean of price ratios and suggest a simple procedure to extract that stationary component. The adjusted price ratios robustly forecast returns. The predictability is also present out-of-sample. At the same time, we show that shifts in the steady-state expected returns and growth rate of fundamentals are responsible for the instability of the return forecasting relation and lack of out-of-sample return predictability.

References

- AMIHUD, Y., AND C. M. HURWICH (2004): “Predictive Regressions: A Reduced-Bias Estimation Method,” *Financial and Quantitative Analysis*, 39, 813–841.
- ANG, A., AND G. BEKAERT (2001): “Stock Return Predictability: Is it There?,” NBER Working Paper No. 8207.
- BERK, J. B., R. C. GREEN, AND V. NAIK (1999): “Optimal Investment, Growth Options and Security Returns,” *Journal of Finance*, 54, 1153–1607.
- BOSSAERTS, P., AND P. HILLION (1999): “Implementing Statistical Criteria to Select Return Forecasting Models: What Do We Learn?,” *Review of Financial Studies*, 12, 405–428.
- BOUDOUGH, J., R. MICHAELY, M. RICHARDSON, AND M. ROBERTS (2004): “On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing,” NBER Working Paper No. 10651.
- BOUDOUGH, J., M. RICHARDSON, AND R. F. WHITELAW (2005): “The Myth of Long-Horizon Predictability,” New York University Working Paper.
- CALVET, L., M. GONZALEZ-EIRAS, AND P. SADINI (2003): “Financial Innovation, Market Participation and Asset Prices,” NBER Working Paper No. 9840.
- CAMPBELL, J. Y. (1991): “A Variance Decomposition for Stock Returns,” *Economic Journal*, 101, 157–179.
- CAMPBELL, J. Y., A. W. LO, AND C. MACKINLAY (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.
- 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.
- CAMPBELL, J. Y., AND M. YOGO (2002): “Efficient Tests of Stock Return Predictability,” Unpublished paper, Harvard University.
- COCHRANE, J. H. (1991): “Explaining the Variance of Price-Dividend Ratios,” *Review of Financial Studies*, 5(2), 243–280.
- FAMA, E. F., AND K. R. FRENCH (1988): “Dividend Yields and Expected Stock Returns,” *Journal of Financial Economics*, 22, 3–27.
- (2001): “Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay,” *Journal of Financial Economics*, 60(1), 3–43.
- FERSON, W. E., S. SARKISSIAN, AND T. T. SIMIN (2003): “Spurious regressions in financial economics?,” *Journal of Finance*, 58(4), 1393–1413.

- GOETZMAN, W. N., AND P. JORION (1993): "Testing the Predictive Power of Dividend Yields," *Journal of Finance*, 48, 663–679.
- GOYAL, A., AND I. WELCH (2003): "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 49(5), 639–654.
- (2004): "A Comprehensive Look at the Empirical Performance of the Equity Premium Prediction," Working Paper Brown University.
- GRULLON, G., AND R. MICHAELY (2002): "Dividends, Share Repurchases, and the Substitution Hypothesis," *Journal of Finance*, 57(4), 1649–1684.
- GUVENEN, M. F. (2003): "A Parsimonious Macroeconomic Model for Asset Pricing: Habit Formation or Cross-sectional Heterogeneity?," Rochester Center for Economic Research Working Paper No. 499.
- HAMILTON, J. D. (1989): "A New Approach to the Analysis of Nonstationary Returns and the Business Cycle," *Econometrica*, 57, 357–384.
- HANSEN, L. P., AND R. J. HODRICK (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88(5), 829–853.
- HODRICK, R. (1992): "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5, 357–386.
- JOVANOVIC, B., AND P. L. ROUSSEAU (2003): "Specific Capital and the Division of Rents," Unpublished paper, New York University.
- KRUEGER, D., AND F. PERRI (2005): "Does Income Inequality Lead To Consumption Inequality? Some Empirical Findings and a Theoretical Explanation," *Review of Economic Studies*, forthcoming.
- LAMONT, O. (1998): "Earnings and Expected Returns," *Journal of Finance*, 53, 1563–87.
- LETTAU, M., AND S. C. LUDVIGSON (2001): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.
- (2005): "Expected Returns and Expected Dividend Growth," *Journal of Financial Economics*, 76, 583–626.
- LETTAU, M., S. C. LUDVIGSON, AND J. A. WACHTER (2004): "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?," NBER Working Paper No. 10270.
- LEWELLEN, J. W. (1999): "The Time Series Relations Among Expected Return, Risk, and book-to-market," *Journal of Financial Economics*, 54, 5–53.

- (2004): “Predicting Returns With Financial Ratios,” *Journal of Financial Economics*, 74(2), 209–235.
- LUDVIGSON, S. C., AND S. NG (2005): “The Empirical Risk-Return Relation: A Factor Analysis Approach,” Unpublished manuscript.
- LUSTIG, H., AND S. VAN NIEUWERBURGH (2005a): “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 60(3), 1167–1219.
- (2005b): “How Much Does Household Collateral Constrain Regional Risk Sharing,” Unpublished manuscript, University of California at Los Angeles and New York University.
- (2005c): “Quantitative Asset Pricing Implications of Housing Collateral Constraints,” Unpublished manuscript, University of California at Los Angeles and New York University.
- MANKIW, N. G., AND M. SHAPIRO (1986): “Risk and Return: Consumption Beta versus Market Beta,” *Review of Economics and Statistics*, 68, 452–59.
- MCGRATTAN, E. R., AND E. C. PRESCOTT (2005): “Taxes, Regulations, and the Value of U.S. and U.K. Corporations,” *Review of Economic Studies*, 72.
- MENZLY, L., T. SANTOS, AND P. VERONESI (2004): “Understanding Predictability,” *Journal of Political Economy*, 112(1), 1–47.
- NELSON, C. C., AND M. J. KIM (1993): “Predictable Stock Returns: The Role of Small Sample Bias,” *Journal of Finance*, 43, 641–661.
- PASTOR, L., AND R. F. STAMBAUGH (2001): “The Equity Premium and Structural Breaks,” *Journal of Finance*, 56(4), 1207–1239.
- PAYE, B. S., AND A. TIMMERMANN (2003): “Instability of Return Prediction Models,” Unpublished manuscript, University of California San Diego.
- PERRON, P. (1989): “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis,” *Econometrica*, 57(6), 1361–1401.
- PETTENUZZO, D., AND A. TIMMERMANN (2004): “Optimal Asset Allocation under Structural Breaks,” Unpublished paper, University of California San Diego.
- PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2004): “Housing, Consumption, and Asset Pricing,” Unpublished paper, University of Chicago.
- QUADRINI, V., AND U. JERMANN (2003): “Stock Market Boom and the Productivity Gains of the 1990s,” Unpublished paper, University of Southern California.
- ROZEFF, M. S. (1984): “Dividend Yields are Equity Risk Premia,” *Journal of Portfolio Management*, 49, 141–160.

- STAMBAUGH, R. F. (1986): "Bias in Regressions with Lagged Stochastic Regressors," Working Paper University of Chicago.
- (1999): "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.
- TOROUS, W., R. VOLKANOV, AND S. YAN (2004): "On Predicting Returns with nearly Integrated Explanatory Variables," *Journal of Business*, 77, 937–966.
- VALKANOV, R. (2003): "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics*, 68, 201–232.
- VICEIRA, L. (1996): "Testing For Structural Change in the Predictability of Asset Returns," Unpublished manuscript, Harvard University.
- VISSING-JORGENSEN, A. (2002): "Limited Asset Market Participation and Intertemporal Substitution," *Journal of Political Economy*, 110(4), 825–853.
- VUOLTEENAHO, T. (2000): "Understanding the Aggregate Book-Market Ratio and its Implications to Current Equity-Premium Expectations," Unpublished paper, Harvard University.

A Data Description

Data used in sections 2, 3, 5, and 6 We use annual end-of-year data from 1926-2004 from CRSP on the value weighted market return (NASDAQ, NYSE, AMEX), with and without dividend capitalization. We denote the capitalized net return series by R , $1 + R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t}$, and the uncapitalized series by R^{ex} , $1 + R_{t+1}^{ex} = \frac{P_{t+1}}{P_t}$. We construct

1. the dividend yield $\frac{D_{t+1}}{P_t}$ as the difference of capitalized and uncapitalized series,
2. the dividend level as the product of the dividend yield $\frac{D_{t+1}}{P_t}$ and the uncapitalized price P_t ,
3. the dividend-price ratio $\frac{D_{t+1}}{P_{t+1}}$ as the ratio of the dividend level D_{t+1} and the uncapitalized price at time $t + 1$, P_{t+1} ,
4. dividend growth as the change in the dividend level $\frac{D_{t+1} - D_t}{D_t}$
5. we check that capitalized returns satisfy $1 + R_{t+1} = \left(1 + \frac{D_{t+1} - D_t}{D_t}\right) \frac{D_t}{P_t} \left(1 + \frac{1}{\frac{D_{t+1}}{P_{t+1}}}\right)$

We then define log returns $r_{t+1} = \log(1 + R_t)$, log dividend growth $\Delta d_{t+1} = \log\left(1 + \frac{D_{t+1} - D_t}{D_t}\right)$, and the lagged log dividend-price ratio $dp_t \equiv d_t - p_t = \log\left(\frac{D_t}{P_t}\right)$. The latter is our independent variable in the analysis. We note that this is not equal to the log dividend yield, $d_{t+1} - p_t$ which some other researchers have used.

Data used in Section 4 We obtained monthly S&P 500 dividend-price and price-earnings ratios and end-of-month composite price index data from Haver. The data are from January 1946-December 2004. This allows us to construct annual end-of-year series for dividends, earnings and prices. This data set also contains returns and dividends. This is the data we use in Section 7 of the main text. Book-to-market ratios are from Vuolteenaho (2000) for 1927-1999. The observations for 1999-2004 are constructed using the clean-surplus ratio method with same Haver earnings and dividend data.

The second panel of Table 7 reports the subsample and full sample means of the log earnings-price ratio (\overline{ep}), dividend-earnings ratio (payout rate), and the implied mean changes in long-run expected returns \bar{r} or expected dividend growth \bar{d} . The implied change in \bar{r} , holding $\bar{d} = 0.0709$ fixed, is computed as:

$$\bar{r} = \bar{d} + (1 + \bar{d}) \frac{\overline{D}}{\overline{E}} \frac{\overline{E}}{\overline{P}}.$$

Likewise, the implied change in \bar{d} , holding $\bar{r} = .1098$ fixed, is computed as:

$$\bar{d} = \frac{\bar{r} - \frac{\overline{D}}{\overline{E}} \frac{\overline{E}}{\overline{P}}}{1 + \frac{\overline{D}}{\overline{E}} \frac{\overline{E}}{\overline{P}}}.$$

B A Modified Log-Approximation

In this section we extend the Campbell-Shiller log-linear approximation to allow for time-varying steady-state growth rates and returns. This framework is a useful organizing principle for the empirical analysis in the main text, but most of our specifications do not impose the approximation. Our results do not depend on the accuracy of the approximation, but the framework helps to understand the intuition and implications.

The gross return of an asset is defined as

$$R_{t+1} \equiv \frac{P_{t+1} + D_{t+1}}{P_t} \quad (20)$$

$$= \frac{D_{t+1}}{D_t} \frac{1 + P_{t+1}/D_{t+1}}{P_t/D_t} \quad (21)$$

As of period t , the steady-state (gross) growth rate of dividends is \bar{D}_t and steady-state expected (gross) returns are \bar{R}_t , implying a steady-state level of the price-dividend ratio $\bar{P}D_t$:

$$\bar{P}D_t = \frac{\bar{D}_t}{\bar{R}_t - \bar{D}_t}. \quad (22)$$

Taking logs and rearranging yields the steady-state log price dividend ratio

$$\bar{p}d_t = \bar{d}_t - \log(\exp(\bar{r}_t) - \exp(\bar{d}_t)). \quad (23)$$

Instead of presuming that the steady-state growth rates and expected returns are constant, we allow for the possibility that the steady-state may change over time. The only requirement that we impose on the steady-state log returns and log growth rates is that they are martingales:

Assumption 1. $E_t[\bar{r}_{t+j}] = \bar{r}_t$ and $E_t[\bar{d}_{t+j}] = \bar{d}_t$.

In other words, the steady-state is constant in expectation only. Although the log price dividend ratio is a nonlinear function of steady-state returns and growth, we assume for simplicity that the steady-state log P/D ratio is also (approximately) a martingale:

Assumption 2. $E_t[\bar{p}d_{t+j}] = \bar{p}d_t$.

This assumption follows from assumption 1 for the specific processes for steady-state returns and growth that we consider in the main text: an unforecastable break in \bar{r} and/or \bar{d} .

Log-linearizing (21) around the steady-state in $t+1$ and expressing the variables in deviations from steady-state yields

$$\begin{aligned} pd_t - \bar{p}d_t &= (\Delta d_{t+1} - \bar{d}_t) - (r_{t+1} - \bar{r}_t) + \rho_{t+1}(pd_{t+1} - \bar{p}d_{t+1}) \\ &\quad + \Delta \bar{p}d_{t+1} + \Delta \bar{r}_{t+1} - \Delta \bar{d}_{t+1} \end{aligned} \quad (24)$$

where $pd_t = p_t - d_t$ and $\rho_{t+1} = \frac{\exp(\bar{p}d_{t+1})}{1 + \exp(\bar{p}d_{t+1})}$. The last two assumptions state that ρ_t is a martingale, and that deviations from the mean price-dividend ratio are uncorrelated with ρ :

Assumption 3. $E_t[\rho_{t+j}] = \rho_t$ and $E_t[\rho_{t+j}(pd_{t+j} - \bar{p}d_t)] = 0$.

Under assumptions 1, 2, and 3, we can take conditional expectations and solve the expectational difference equation for pd_t :

$$pd_t - \bar{p}d_t = E_t[\Delta d_{t+1} - \bar{d}_t] - E_t[r_{t+1} - \bar{r}_t] + \rho_t E_t[pd_{t+1} - \bar{p}d_{t+1}] \quad (25)$$

$$= \sum_{j=1}^{\infty} \rho_t^{j-1} E_t[\Delta d_{t+j} - \bar{d}_t] - E_t[r_{t+j} - \bar{r}_t] \quad (26)$$

The log price dividend ratio is the sum of the steady-state price dividend ratio and the discounted sum of expected dividend growth minus expected returns in excess of steady-state growth and returns:

$$pd_t = \overline{pd}_t + \sum_{j=1}^{\infty} \rho_t^{j-1} E_t[\Delta \tilde{d}_{t+j}] - E_t[\tilde{r}_{t+j}], \quad (27)$$

where $\Delta \tilde{d}_{t+j} = \Delta d_{t+j} - \bar{d}_t$ and $\tilde{r}_{t+j} = r_{t+j} - \bar{r}_t$. The expression in the main text for dp_t follows from (27) and $dp_t = -pd_t$.

C A Structural Model

Setup We propose a model for log dividend growth Δd and log returns r , where expected dividend growth z and expected returns x follow an AR(1) with autoregressive coefficient ϕ :

$$\Delta d_{t+1} - \bar{d} = z_t + \epsilon_{t+1} \quad z_{t+1} = \phi z_t + \zeta_{t+1} \quad (28)$$

$$r_{t+1} - \bar{r} = x_t + \eta_{t+1} \quad x_{t+1} = \phi x_t + \xi_{t+1} \quad (29)$$

where \bar{d} is the long-run mean log dividend growth and \bar{r} is the long-run mean return. The model has three fundamental shocks: a dividend innovation ϵ_{t+1} , an innovation in expected dividends ζ_{t+1} , and an innovation in expected returns ξ_{t+1} . Campbell (1991)'s return decomposition implies that

$$\eta_{t+1} = -\frac{\rho}{1-\rho\phi} \xi_{t+1} + \frac{\rho}{1-\rho\phi} \zeta_{t+1} + \epsilon_{t+1}. \quad (30)$$

We assume that all three errors are serially uncorrelated and have zero cross-covariance at all leads and lags: $Cov(\epsilon_{t+1}, \zeta_{t+j}) = 0$, $\forall j \neq 1$, and $Cov(\epsilon_{t+1}, \xi_{t+j}) = 0$, $\forall j \neq 1$, except $Cov(\zeta_t, \xi_t) = \chi$ and $Cov(\zeta_t, \epsilon_t) = \lambda$.

The Dividend-Price Ratio In steady-state, the log dividend-price ratio is $\overline{dp} = \log\left(\frac{\bar{r}-\bar{d}}{1+\bar{d}}\right)$, hence $\overline{DP} \approx \frac{\bar{r}-\bar{d}}{1+\bar{d}}$ and $\rho \approx \frac{1+\bar{d}}{1+\bar{r}}$. The log dividend-price ratio can be written as

$$dp_t = d_t - p_t = \overline{dp} + \frac{x_t - z_t}{1-\rho\phi}. \quad (31)$$

This equation clearly shows that the demeaned dividend-price ratio is an imperfect forecaster of returns. Returns are predicted by x_t (29) which not only contains the demeaned dp ratio, but also expected dividend growth z_t : $x_t = (1-\rho\phi)(dp_t - \overline{dp}) + z_t$.

Deriving the Reduced Form Model This structural model implies a reduced form model that recovers the two predictability equations from Section (2):

$$(\Delta d_{t+1} - \bar{d}) = \kappa_d(dp_t - \overline{dp}) + \tau_{t+1}^d \quad (32)$$

$$(r_{t+1} - \bar{r}) = \kappa_r(dp_t - \overline{dp}) + \tau_{t+1}^r \quad (33)$$

$$(dp_{t+1} - \overline{dp}) = \phi(dp_t - \overline{dp}) + \tau_{t+1}^{dp}. \quad (34)$$

The third equation of this Vector Error Correction Model is an AR(1) process for the dividend-price ratio. Because of equation (6), the dividend-price ratio is the difference of two AR(1) processes with the same root ϕ ,

which is again an AR(1) process. We also considered a model where expected dividend growth has a different autoregressive coefficient $\psi \neq \phi$: $z_{t+1} = \psi z_t + \zeta_{t+1}$. In that case the dividend-price ratio is an ARMA(1,1) with roots $\phi + \psi$ and $-\phi\psi$. Since the dp ratio is well described by an AR(1) model in the data, we set $\phi = \psi$.

The slope coefficients are related to the structural parameters:

$$\kappa_d = \frac{Cov(\Delta d_{t+1}, d_t - p_t)}{Var(d_t - p_t)} = \frac{-(1 - \rho\phi)(\sigma_\zeta^2 - \chi)}{\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi} \quad (35)$$

$$\kappa_r = \frac{Cov(r_{t+1}, d_t - p_t)}{Var(d_t - p_t)} = \frac{(1 - \rho\phi)(\sigma_\xi^2 - \chi)}{\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi} \quad (36)$$

The innovations to the VECM, $\tau = (\tau^d, \tau^r, \tau^{dp})$, are given by:

$$\tau_{t+1}^d = \epsilon_{t+1} + x_t \begin{pmatrix} -\kappa_d \\ 1 - \rho\phi \end{pmatrix} + z_t \begin{pmatrix} \kappa_r \\ 1 - \rho\phi \end{pmatrix} \quad (37)$$

$$\tau_{t+1}^r = \epsilon_{t+1} + x_t \begin{pmatrix} -\kappa_d \\ 1 - \rho\phi \end{pmatrix} + z_t \begin{pmatrix} \kappa_r \\ 1 - \rho\phi \end{pmatrix} - \rho \frac{\xi_{t+1} - \zeta_{t+1}}{1 - \rho\phi} \quad (38)$$

$$\tau_{t+1}^{dp} = \frac{\xi_{t+1} - \zeta_{t+1}}{1 - \rho\phi}. \quad (39)$$

The structural model imposes a restriction on the innovation vector: $\rho\tau_{t+1}^{dp} = (\tau_{t+1}^d - \tau_{t+1}^r)$. Another way to write this restriction is as a restriction on a weighted sum of κ_r and κ_d :

$$\kappa_r - \kappa_d = 1 - \rho\phi.$$

We call this restriction the *present value constraint*.

Identification Leaving aside the mean parameters $(\bar{d}, \bar{r}, \bar{dp})$ which play no role in the demeaned system, the structural parameter vector is $\Theta = (\phi, \sigma_\zeta, \sigma_\xi, \sigma_\epsilon, \chi, \lambda)$. The variance σ_η^2 is implied by (30).

The vector of parameters from the reduced form model is $b = (\kappa_d, \kappa_r, \phi, \Sigma_\tau)$, where Σ_τ is the variance-covariance matrix of τ . There are six unique elements in this covariance matrix and the present value constraint imposes three restrictions on these six elements. In addition, the present value constraint imposes a restriction on one of the elements in $(\kappa_r, \kappa_d, \phi)$. Hence, there are five unique elements in $b = (\kappa_r, \kappa_d, \sigma_{\tau^r}, \sigma_{\tau^d}, \sigma_{\tau^r, \tau^d})$, where σ_{τ^r} and σ_{τ^d} are the standard deviations of τ^r and τ^d , and σ_{τ^r, τ^d} is the covariance between the two. When we estimate the non-singular system of equations (32) and (33), we can use these five coefficients to identify five out of six structural parameters in Θ . Therefore, whenever we back out structural parameters from the reduced form estimates, we find it convenient to tabulate results for a range of values for V_z , which measures the contribution of expected dividend growth to the variance of the dividend-price ratio:

$$V_z = \frac{\sigma_z^2}{(1 - \rho\phi)^2 \sigma_{dp}^2} = \frac{\sigma_\zeta^2}{(\sigma_\xi^2 + \sigma_\zeta^2 - 2\chi)}. \quad (40)$$

D Monte Carlo Simulations

In-sample The Monte Carlo exercise simulates the model under the null hypothesis of 1 break or 2 breaks in the dividend-price ratio. I then asks what the univariate long-horizon slope coefficients and R^2 are in a small sample (of the same length as the data) when we use the unadjusted dp as regressor versus the adjusted \widetilde{dp} . For the one-break case reported in the top panel of Table 9, we use the following algorithm:

step 1 To find the true parameters, we estimate the 1-,3-,5-year VECM under the assumption of 1 break in 1991 for the full sample. This delivers reduced form estimates:

$$b = (.0121, .2098, .1410, .1868, .0176).$$

Throughout, $\rho = .9616$, the value implied by the mean price-dividend ratio in the sample.

step 2 We invert these reduced form parameters to obtain structural parameters based on the identification scheme described above. We need to take a stance on the fraction of the variance in the dividend-price ratio attributable to expected dividend growth (equation 40). We set $V_z = 0.3$, but the results are not sensitive to this choice. The implied structural parameters are:

$$\Theta = (.8344, .0182, .0355, .1371, .0004, -.0005).$$

step 3 In each Monte Carlo iteration, we draw a new 178×3 vector of i.i.d. standard normal variables. The structural shocks $(\varepsilon_t, \zeta_t, \xi_t)$ are these standard normal variables pre-multiplied by $\Sigma_u^{\frac{1}{2}}$, where Σ_u is the covariance matrix of the structural innovations:

$$\Sigma_u = \begin{bmatrix} \sigma_\varepsilon^2 & \lambda & 0 \\ \lambda & \sigma_\zeta^2 & \chi \\ 0 & \chi & \sigma_\xi^2 \end{bmatrix}$$

step 4 We recursively build up time-series for x , $r - \bar{r}$, z , and $\Delta d - \bar{d}$ according to (29) and (28). We form a time series for the demeaned dividend-price ratio from (31).

step 5a Under the null hypothesis of 1 break in 1991, the *break-adjusted* series for returns, dividend growth and the dividend-price ratio are obtained by adding *constant* long-run means \bar{r} , \bar{d} , and \bar{dp} to the demeaned series. I.e. the adjusted dp series is the demeaned series plus the mean over the *entire sample* 1927-2004, which is $\bar{dp} = -3.272$. Likewise, the adjusted dividend growth series Δd is formed by adding to the demeaned series $\Delta d - \bar{d}$ the sample mean $\bar{d} = .0432$. We then back out $\bar{r} = \bar{d} + (1 + \bar{d}) \exp(\bar{dp}) = .0827$, and form the adjusted return series as the demeaned series plus this \bar{r} .

step 5b On the other hand, the *unadjusted* series still displays a break in 1991. To obtain the unadjusted series we need to add in a *different mean* before and after 1991. As before, we assume \bar{d} did not change, so that the change in \bar{dp} entirely comes from a change in \bar{r} . The unadjusted dp series is obtained by adding in the 1927-1991 mean (-3.133) before 1991 and the 1992-2004 mean (-3.968) after 1991. Likewise for returns we add in .0886 before 1991 and .0629 after 1991. These are the subsample means \bar{dp} and \bar{r} that were reported in lines 1 and 2 of Table 7 (left panel).

step 6 We form annualized, cumulative long-horizon returns and dividend growth rates $\frac{1}{N} \sum_{j=1}^N \Delta r_{t+j}$ and $\frac{1}{N} \sum_{j=1}^N \Delta d_{t+j}$ in the same way. There is one set of unadjusted series and one set of adjusted series.

step 7 After the formation, we discard the first 100 observations (burin-in), and are left with the same number of observations as in the data: 78- longest horizon+1. For the 1-, 3-, and 5-year system we report on, the longest horizon is 5, so all statistics are computed with 74 observations.

step 8 We then run univariate 1-, 3-, and 5-year ahead predictability regressions of *adjusted* returns and dividend growth on both *unadjusted* and *adjusted* dividend-price ratios. We keep track of the predictability coefficients and regression R^2 .

step 9 We repeat this procedure 50,000 times and report the average coefficients and R^2 across Monte Carlo iterations.

The Monte Carlo exercise in the two-break case is exactly analogous (bottom panel of Table 9). In step 1, we estimate the 1-, 3-, and 5-year ahead VECM system under the null of 2-breaks in 1954 and 1994. The VECM coefficient vector is

$$b = (0.0795, 0.4089, 0.1391, 0.1743, 0.0157).$$

The implied structural parameter vector for $V_z = 0.3$ is

$$\Theta = (0.6974, 0.0293, 0.0635, 0.1339, 0.0014, -0.0012).$$

To construct the unadjusted dp and return series we use the means reported in the right panel of Table 7.

Out-of-sample We conduct a separate Monte Carlo exercise for the out-of-sample predictability (panels 2 and 3 in Table 3). We simulate the model under the null hypothesis of two breaks in 1954 and 1994. For simplicity we use the same true parameters in both panels 2 and 3. They are obtained from estimating the 1-year ahead VECM under the assumption of two breaks in 1954 and 1994, but for the period 1947-2004, which is the same sample period as we use in the data (panel 1 of Table 3). We then construct adjusted and unadjusted dp ratios as in the above algorithm, and run the same out-of-sample predictions as in the data.

Table 1: Forecasting Returns and Dividend Growth with the Dividend-Price Ratio.

This table reports estimation results for the equations $\Delta d_{t+1} - \bar{d} = \kappa_d(dp_t - \bar{dp}) + \tau_{t+1}^d$ and $r_{t+1} - \bar{r} = \kappa_r(dp_t - \bar{dp}) + \tau_{t+1}^r$. The table reports point estimates and standard errors in parentheses of κ_r and κ_d , as well as regression R^2 in square brackets. The parameters $(\bar{r}, \bar{d}, \bar{dp})$ are the sample means of log returns r , log dividend growth Δd and the log dividend-price ratio dp . The top panel compares the case of no break in the log dividend-price ratio (\bar{dp} is fixed) with the case where there is a break in the log dividend-price ratio: \bar{dp}_1 is the sample mean log dividend-price ratio for 1927-1991 and \bar{dp}_2 is the mean for 1992-2004. The estimation is by GMM, where the moments are the OLS normal conditions. Standard errors are by Newey-West with four lags. Row 1 reports results for the full sample; rows 2 and 3 report results for two subsamples. Row 4 reports the F-statistic and associated p -value from a Chow test with null hypothesis of no structural break in 1991 in the forecasting equations. The bottom panel compares the case of no break in the log dividend-price ratio (\bar{dp} is fixed) with the case where there are two breaks in the log dividend-price ratio: \bar{dp}_1 is the sample mean log dividend-price ratio for 1927-1954 (row 6), \bar{dp}_2 is the mean for 1955-1994 (row 7), and \bar{dp}_3 is the mean for 1995-2004 (row 8). Row 9 panel reports the F-statistic and associated p -value from a Chow test with null hypothesis of no structural breaks in 1954 and 1994 in the forecasting equations.

Sample	Returns		Dividend Growth	
	No Break	1 Break ('91)	No Break	1 Break ('91)
1927-2004	.094 (.046) [.038]	.235 (.058) [.100]	.005 (.037) [.000]	.019 (.047) [.001]
1927-1991	.235 (.065) [.087]	.235 (.065) [.087]	.014 (.053) [.001]	0.014 (.053) [.001]
1992-2004	.235 (.134) [.199]	.235 (.134) [.199]	.035 (.103) [.006]	.035 (.103) [.006]
Chow F -stat	3.408	.134	.114	.024
p -val	[.038]	[.875]	[.892]	[.977]
Sample	No Break	2 Breaks ('54,'94)	No Break	2 Breaks ('54,'94)
1927-2004	.094 (.046) [.038]	.455 (.081) [.223]	.005 (.037) [.000]	.124 (.073) [.032]
1927-1954	.510 (.175) [.163]	.510 (.175) [.163]	.037 (.182) [.002]	.037 (.182) [.002]
1955-1994	.383 (.106) [.240]	.383 (.106) [.240]	.142 (.077) [.064]	.142 (.077) [.064]
1995-2004	.532 (.129) [.546]	.532 (.129) [.546]	.226 (.097) [.126]	.226 (.097) [.126]
Chow F -stat	4.390	.235	.998	.500
p -val	[.003]	[.918]	[.414]	[.736]

Table 2: Tests for Change in Mean of Log Dividend-Price Ratio

The top half of the table reports the test-statistic of a supF Perron structural break tests of the null hypothesis of no break against the alternative hypothesis of one (first row) or two (second row) breaks with unknown break date. It reports the p -value of the test statistic, as well as the resulting break date. The last column reports the estimated change in means before and after the break(s). The bottom half of the table reports first and second order autocorrelation coefficients, an Augmented Dickey Fuller test, testing the null hypothesis of a unit root (and associated p -value), and the time-series standard deviation. It does so for the unadjusted log dividend-price ratio (row 1), the log price ratio adjusted for a change in its mean in 1991 (row 2), and the log dividend-price ratio adjusted for a change in its mean in 1954 and 1994 (row 3).

Structural Break Tests of dp						
	H_0	H_1	$supF$ -Test	p -value	Date(s)	$\Delta\bar{dp}$
dp	0 break	1 break	13.7	< 1%	1991	-.86
	0 break	2 breaks	28.9	< 1%	1954, 1994	-.37, -.78
Persistence Properties of Adjusted Dividend-Price Ratio						
	AC(1)	AC(2)	ADF Test	p -val	s.d.	
dp , unadjusted	.91	.81	-1.383	.586	.42	
\widetilde{dp} , adjusted, 1 break	.77	.55	-3.016	.038	.26	
\widetilde{dp} , adjusted, 2 breaks	.61	.23	-4.731	.010	.20	

Table 3: Out-of-Sample Predictability.

The table reports one-period-ahead return forecast errors based on the Random Walk model (row 1), based on the forecasting equation $r_{t+1} - \bar{r} = \kappa_r(dp_t - \bar{dp}) + \tau_{t+1}^r$ with fixed \bar{dp} (row 2), and based on the same forecasting equation but with a change in \bar{dp} in 1991 (row 3), and two changes in the mean \bar{dp} in 1954 and 1994 (row 4). In rows 5 and 6, we use \widehat{dp}^H as forecaster, the dividend-price ratio adjusted with the Hamilton procedure, using two regimes or three regimes. In row 6, we use the recursively adjusted dividend-price ratio \widehat{dp}^P , using the Perron procedure. All numbers denote returns per annum. The second and third panels report results from a Monte Carlo exercise. We simulate the structural model under the null hypothesis that the data generating process has one break in 1991 (panel 2) or has two breaks in 1954 and 1994 (panel 3). We compare the same three out-of sample forecasting exercises as in the data (panel 1). Except in the last panel, we also look at the forecast errors when we only correct for the second break and not the first one. The structural parameters in panels 2 and 3 are the same and were obtained from the Vector Error Correction Model (VECM) parameters estimated under the assumption of two breaks in 1954 and 1994 for the period 1947-2004, the same forecasting period as in panel 1.

	Mean absolute error	Root mean squared error
Panel 1: Data		
Random Walk	.1338	.1605
Unadjusted dp	.1411	.1685
\widetilde{dp} - 1 break	.1309	.1558
\widetilde{dp} - 2 breaks	.1158	.1421
\widehat{dp}^H - recursive Hamilton, 1 break	.1309	.1548
\widehat{dp}^H - recursive Hamilton, 2 breaks	.1158	.1498
\widehat{dp}^P - recursive Perron	.1336	.1637
Panel 2: Monte Carlo - 1 break		
Random Walk	.1212	.1525
Unadjusted dp	.1210	.1555
\widetilde{dp} - ex-post, 1 break	.1058	.1334
Panel 3: Monte Carlo - 2 breaks		
Random Walk	.1222	.1533
Unadjusted dp	.1203	.1516
\widetilde{dp} - ex-post, 1 break	.1165	.1493
\widetilde{dp} - ex-post, 2 breaks	.1064	.1336

Table 4: Tests for Change in Mean of Financial Ratios

The top half of the table reports the test-statistic of a supF Perron structural break tests of the null hypothesis of no break against the alternative hypothesis of one (first row) or two (second row) breaks with unknown break date. It reports the p -value of the test statistic, as well as the resulting break date. The last column reports the estimated change in means before and after the break(s). These tests are performed for the log earnings-price ratio $ep = e - p$, the log book value -to-market value of equity ratio $bm = b - m$, the log dividend-earnings ratio $de = d - e$, the log dividend-price ratio adjusted for repurchases dp^{rep} , and the log dividend-price ratio of the universe of CRSP firms that *excludes* the NASDAQ firms dp^{nas} .

Structural Break Tests						
	H_0	H_1	$supF$ -Test	p -value	Date(s)	$\Delta mean$
ep	0 break	1 break	15.5	< 1%	1990	-.67
	0 break	2 breaks	18.0	< 1%	1953, 1994	-.50, -.62
bm	0 break	1 break	9.3	< 10%	1953	-.80
	0 break	2 breaks	17.9	< 1%	1953, 1990	-.71, -.33
de	0 break	1 break	5.6	> 10%	1993	-.24
	0 break	2 breaks	5.3	> 10%	1990, 1993	+.27, -.46
dp^{nas}	0 break	1 break	10.2	< 5%	1992	-.75
	0 break	2 breaks	18.1	< 1%	1954, 1995	-.35, -.70
dp^{rep}	0 break	1 break	3.7	> 10%	1990	-.43
	0 break	2 breaks	4.5	> 10%	1954, 1991	-.23, -.34
	0 break	3 breaks	20.1	< 1%	1957, 1973, 1990	-.46, +.50, -.57

Table 5: Forecasting Returns With Other Financial Ratios.

This table reports estimation results for the equation $r_{t+1} - \bar{r} = \kappa_r(y_t - \bar{y}) + \tau_{t+1}^r$, where y is the log earnings-price ratio $ep = e - p$ in the first row, the log book-to-market value ratio $bm = b - m$ in the second row, the dividend-price ratio without the NASDAQ firms dp^{nas} in the third row, and the repurchase adjusted dividend-price ratio dp^{rep} in the fourth row. The table reports point estimates and standard errors in parentheses of κ_r , and the regression R^2 in brackets. The regressor is the unadjusted valuation ratio in the first column, the one-break adjusted valuation ratio in the second column, and the two-break adjusted valuation ratio in the third column. For the predictor dp^{rep} , we report the three-break case as well. The break dates for all regressors are reported in Table 4. The sample is 1946-2004 in row 1, and 1927-2004 in all other rows.

Predictor y	No Break	1 Break	2 Breaks	3 Breaks
ep	.119 (.030) [.104]	.214 (.039) [.190]	.216 (.045) [.185]	
bm	.070 (.036) [.030]	.255 (.063) [.154]	.308 (.064) [.188]	
dp^{nas}	.110 (.048) [.043]	.250 (.056) [.105]	.417 (.090) [.182]	
dp^{rep}	.191 (.054) [.079]	.282 (.065) [.126]	.361 (.084) [.161]	.576 (.097) [.250]

Table 6: Forecasting with the Dividend-Price Ratio - Real-Time Adjustments

This table redoes the estimation in table 1, but estimates break dates in real-time. For the one-break case, a regime switching model is estimated for dp . The probability of switching from the high to the low dp regime exceeds 0.50 for the first time in 1995. Likewise, for the two-break case, a regime switching model with three regimes is estimated. The 0.50-probability switch from the high to the medium dp regime occurs in 1960; the switch from the medium to the low dp regime occurs in 1996 at the 50% level.

	\widehat{dp}^P - Rec. Perron		\widehat{dp}^H - 2 Regimes		\widehat{dp}^H - 3 Regimes	
	Returns	Div. growth	Returns	Div. growth	Returns	Div. growth
κ	.299	.153	.218	.027	.353	.109
(s.e.)	(.092)	(.102)	(.066)	(.047)	(.094)	(.069)
$[R^2]$	[.073]	[.037]	[.086]	[.003]	[.139]	[.026]

Table 7: Implied Changes in Steady-state Expected Returns and Dividend Growth

The top panel (bottom panel) of the table reports the mean log dividend-price ratio (log earnings-price ratio) in the subsamples, as well as the difference between the two. On the second row is reports the mean return \bar{r} in the subsamples if all of the changes in \bar{dp} (\bar{ep}) in the first row were attributable to changes in mean returns. The third row reports mean dividend growth rates \bar{d} in the subsamples if all of the changes in \bar{dp} (\bar{ep}) in the first row were attributable to changes in mean dividend growth. The left panel reports the case of one break; the right panel the case of two breaks. For the log dividend-price ratio the break date is estimated to be 1991 for the one break case (1954 and 1994 for the two break case). For the log earnings-price ratio, the break is estimated to be 1990 for the one break case (1953 and 1994 for the two break case). In the bottom panel, we also report the change in the payout rate, the log dividend-earnings ratio de . The sample for the top panel is 1927-2004, the sample for the bottom panel is 1946-2004 (see data appendix A).

Dividend-Price Ratio								
	27-91	92-04	Change	27-54	55-94	95-04	Changes	
\bar{dp}	-3.133	-3.968	-0.835	-2.940	-3.301	-4.086	-.362, -.785	
$\bar{r} \downarrow, \bar{d} \text{ constant}$	8.86%	6.29%	-2.60%	9.83%	8.16%	6.07%	-1.67%, -2.09%	
$\bar{d} \uparrow, \bar{r} \text{ constant}$	5.00%	7.54%	2.54%	4.07%	5.68%	7.76%	1.61%, 2.08%	
Earnings-Price Ratio								
	46-90	91-04	Change	46-53	54-94	95-04		
\bar{ep}	-2.540	-3.202	-0.662	-2.217	-2.652	-3.267	-.435, -.616	
Payout rate \bar{de}	-.665	-.779	-.114	-.606	-.661	-.889	-.056, -.227	
$\bar{r} \downarrow, \bar{d} \text{ constant}$	11.72%	9.24%	-2.47%	13.70%	11.25%	8.87%	-2.46%, -2.38%	
$\bar{d} \uparrow, \bar{r} \text{ constant}$	6.38%	8.79%	2.41%	4.52%	6.83%	9.16%	2.31%, 2.33%	

Table 8: Estimation with Long-Horizon Moments

This table reports GMM estimates for the parameters $(\kappa_d, \kappa_r, \phi)$, their asymptotic standard errors and p -values. The results in panel A are for the system with one-year ahead equations for dividend growth and returns ($H = 1, N = 0$). The results in panel B are for the system with one-year, three-year and five-year ahead equations for dividend growth and returns ($H = \{1, 3, 5\}, N = 2$). The estimation is by GMM. The first stage weighting matrix is the identity matrix. The asymptotic standard errors and p -values are computed using the Newey-West HAC procedure (second stage weighting matrix) with four lags in panel A and $H = 5$ lags in panel B. The first number in the last column denotes the present-value constraint violation of the univariate OLS slope estimators $(1 - \rho\phi^{ols})^{-1}(\kappa_r^{ols} - \kappa_d^{ols})$. In panel B this number is the average violation of the 3 constraints, one at each horizon. It is expressed in the same units as κ_d and κ_r . The second number in the last column reports the average moment violation. In panel A that number is not available (N/A) because the system is exactly identified. The dividend-price ratio in rows 1 and 4 is the unadjusted one. In rows 2 and 5, the dividend-price ratio is adjusted for one break in 1991 (see equation 8), and in rows 3 and 6, it is the series adjusted for two breaks in 1954 and 1994. All estimation results are for the full sample 1927-2004.

	κ_d	κ_r	ϕ	PV violation	moment violation
Panel A: No Long-Horizon Moments $H = \{1\}$					
No Break	.005 (.037)	.094 (.046)	.945 (.052)	-.046	N/A
1 Break ('91)	.019 (.047)	.2353 (.0554)	.813 (.052)	.004	N/A
2 Breaks ('54, '94)	.124 (.073)	.4553 (.0792)	.694 (.070)	-.001	N/A
Panel B: Long-Horizon Moments $H = \{1, 3, 5\}$					
No Break	.021 (.018)	.068 (.038)	.990 (.032)	.189	.205
1 Break ('91)	.012 (.019)	.210 (.043)	.834 (.042)	.076	.085
2 Breaks ('54, '94)	.080 (.065)	.409 (.078)	.697 (.060)	.100	.144

Table 9: Long-Horizon Predictability in Data and Monte Carlo Exercise.

The table reports results from univariate regressions of cumulative long-horizon returns on the log dividend-price ratio. The left columns denote slope coefficients for one-year, three-year, and five-year horizon regressions; the middle columns report standard errors for the slope coefficients; the right columns report the corresponding regression R^2 . The rows labeled ‘Data, dp unadj.’ denote regressions run with real data using the unadjusted log dividend-price ratio as independent variable. The rows labeled ‘Data, dp adj.’ denote regressions run with real data using the log dividend-price ratio, adjusted for a break in the mean \bar{dp} . In the top panel there is one break in this mean in 1991; in the bottom panel there are two breaks in 1954 and 1994. The results for the data are contrasted with the results from a Monte Carlo exercise. The return and dividend growth system is estimated until 1991 in the top panel (1954 in the bottom panel) on real data. The estimated parameters imply ‘true’ structural parameters. The theoretical long-horizon slope coefficients and regression R^2 are reported in the row with label ‘True Values in Sim.’ For these parameters, the structural model is then simulated 10,000 times for 78 periods. The row ‘Sim., dp unadj.’ denotes the Monte Carlo average slope coefficient and R^2 statistic using the unadjusted log dividend-price ratio as independent variable. The row ‘Sim., dp adj.’ also reports regression coefficients and statistics of regressions on artificial data, but now the independent variable comes from a model where the mean \bar{r} is adjusted to equal the change in \bar{dp} in the data.

Horizon	κ_r			$s.e.$			R^2 (%)		
	1	3	5	1	3	5	1	3	5
One Break in 1991									
True Values in Sim.	.222	.187	.159	N/A	N/A	N/A	8.9	20.0	25.2
Sim., dp unadj.	.115	.092	.075	.064	.034	.025	5.1	11.5	14.6
Data, dp unadj.	.087	.095	.071	.065	.037	.026	2.4	8.1	9.2
Sim., dp adj.	.264	.214	.175	.093	.048	.034	10.5	22.7	28.0
Data, dp adj.	.220	.200	.150	.084	.047	.033	8.6	20.0	22.0
Two Breaks in 1954 and 1994									
True Values in Sim.	.456	.330	.249	N/A	N/A	N/A	20.7	37.0	38.8
Sim., dp unadj.	.128	.089	.064	.063	.032	.023	6.2	11.5	12.8
Data, dp unadj.	.087	.095	.071	.065	.037	.026	2.4	8.1	9.2
Sim., dp adj.	.483	.340	.249	.110	.052	.037	21.4	37.5	38.6
Data, dp adj.	.442	.355	.231	.102	.054	.041	20.7	37.4	30.8

Figure 1: Forecasting Returns - Rolling Regressions

The top panel plots estimation results for the equation $r_{t+1} - \bar{r} = \kappa_r (dp_t - \bar{dp}) + \tau_{t+1}^r$. It shows the estimates for κ_r using 30-year rolling windows. The dashed line in the left panels denote the point estimate plus or minus one standard deviation. The parameters \bar{r} and \bar{dp} are the sample means of log returns r and the log dividend-price ratio dp . The data are annual for 1927-2004. The middle panel gives the slope coefficient κ_r from a regression where the right-hand side variable is \widetilde{dp} , adjusted for 1 break in 1991 (see Section 3.3). The bottom panel gives the slope coefficient κ_r from a regression where the right-hand side variable is \widetilde{dp} , adjusted for 2 breaks in 1954 and 1994 (see Section 3.3).

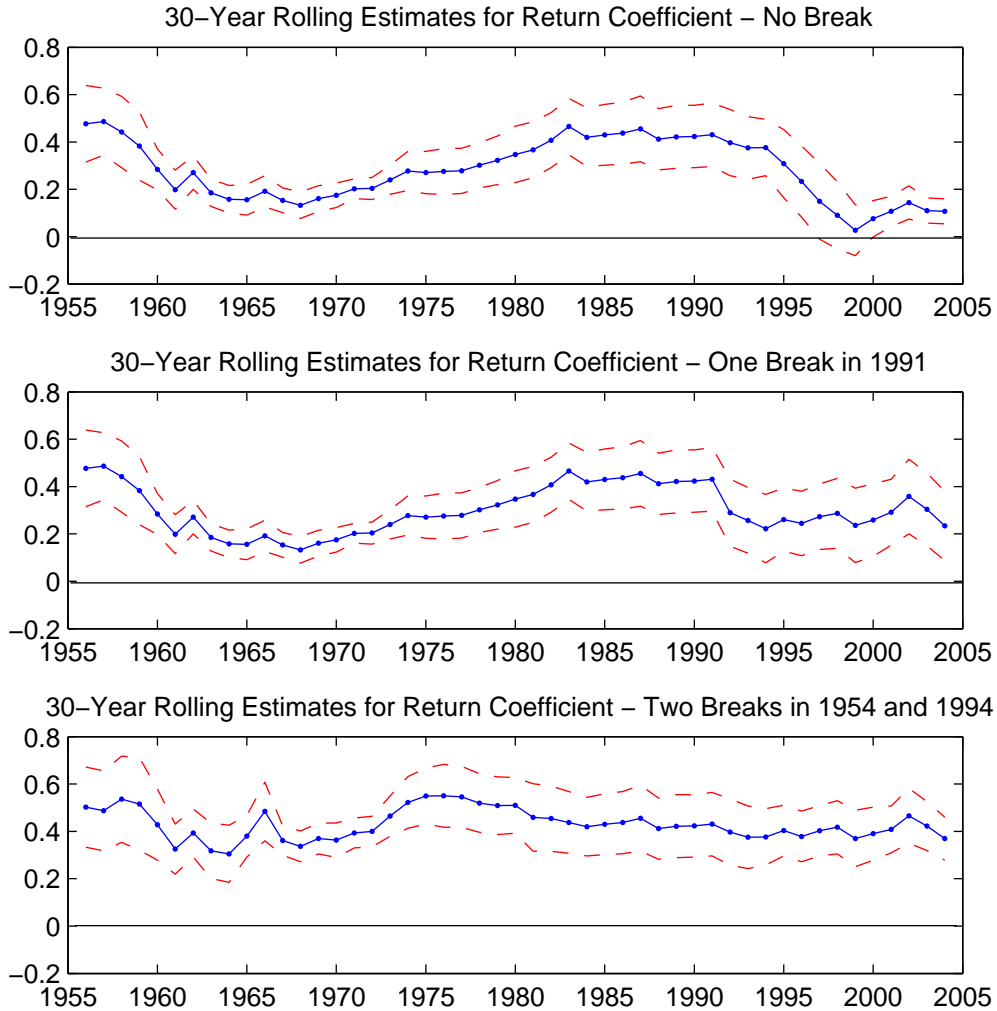


Figure 2: The Empirical Distribution of the Dividend-Price Ratio.

The figure plots the smoothed empirical distribution of the log dividend-price ratio dp (solid line), alongside the smoothed density obtained from drawing from the empirical distribution with replacement (bootstrap, dash-dotted line), and the smoothed density from a Monte-Carlo exercise (dashed line).

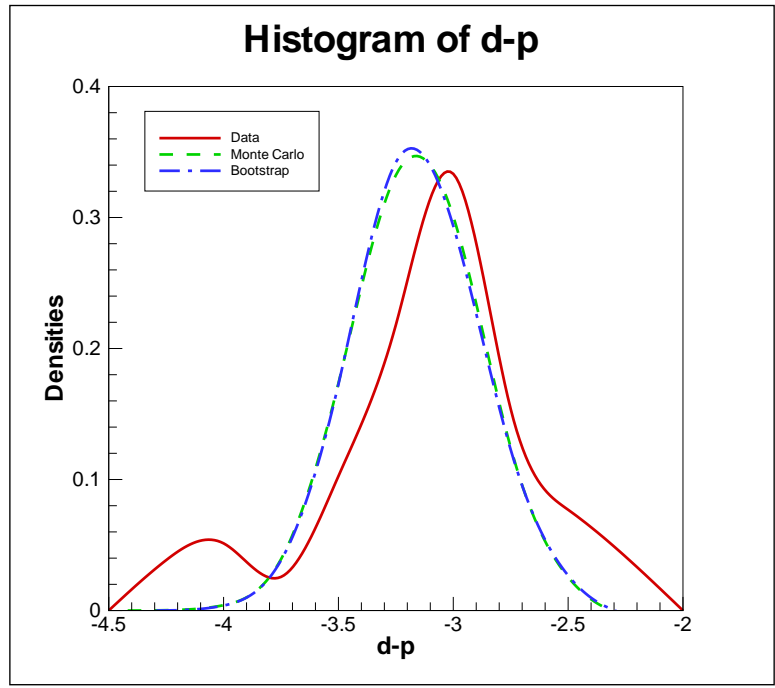


Figure 3: Change in the Mean of the Dividend-Price Ratio.

The top left panel plots the log dividend-price ratio $dp_t = d_t - p_t$ as well as its sample means \bar{dp}_1 in the subsample 1927-1991 and \bar{dp}_2 in the subsample 1992-2004. The bottom left panel overlays the subsample means \bar{dp}_1 in 1927-1954, \bar{dp}_2 in 1955-1994, and \bar{dp}_3 in 1995-2004. The top right panel plots the adjusted dividend-price ratio $\tilde{dp}_t = dp_t - \bar{dp}_1, t = 1, \dots, \tau$ and $dp_t - \bar{dp}_2, t = \tau, \dots, T$. The bottom right panel plots the adjusted dividend-price ratio in the two-break case. In the two bottom panels, the adjusted series is rescaled so that it coincides with the adjusted series for the first subsample.

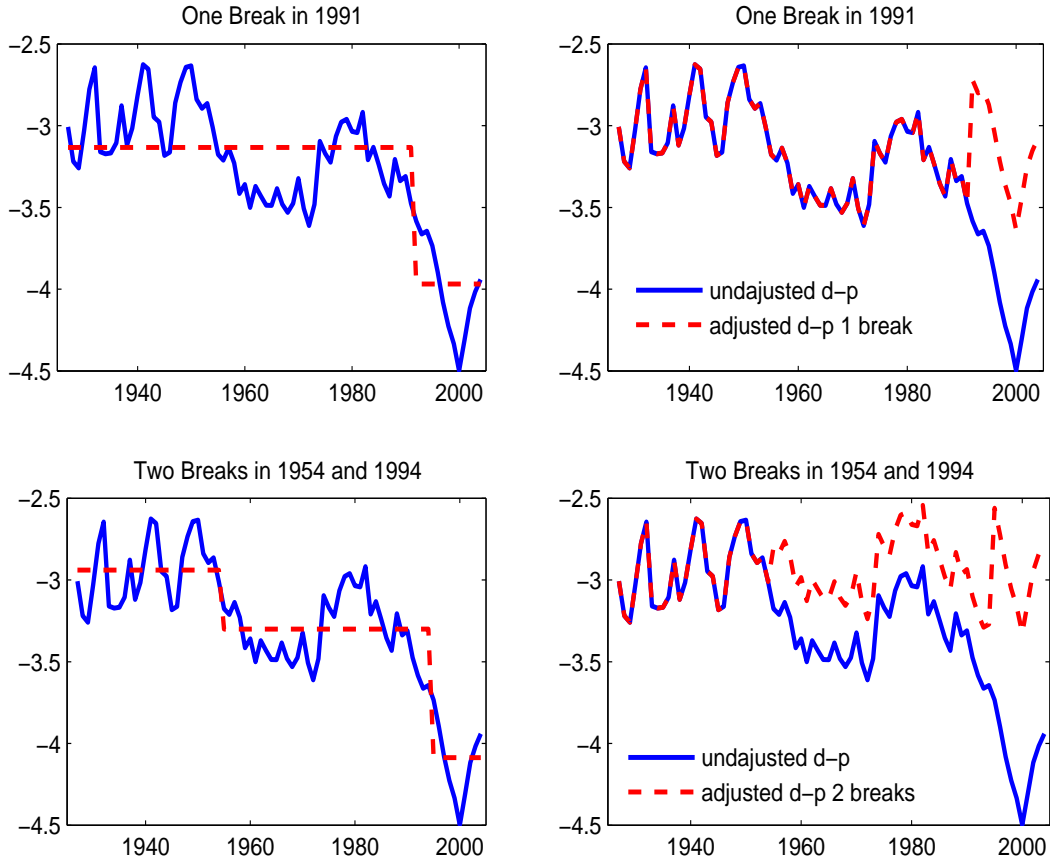


Figure 4: Recursive Estimation of the Empirical Distribution of the Dividend-Price Ratio.

We recursively estimate an AR(1) for the log dividend-price ratio dp , $dp_{t+1} = c + \phi dp_t + \tau_{t+1}^{dp}$ using data up to time $t + 1$, and bootstrap percentiles of the empirical distribution by drawing with replacement from the residuals $\{\tau_1^{dp}, \dots, \tau_{t+1}^{dp}\}$. The dashed lines represent the 2.5, 5, 95, and 97.5 percentiles of the bootstrapped distribution. The initial sample is 1927-1951. Each successive exercise adds one year of data. The solid line represents the observed log-dividend-price ratio in deviation from its sample mean.

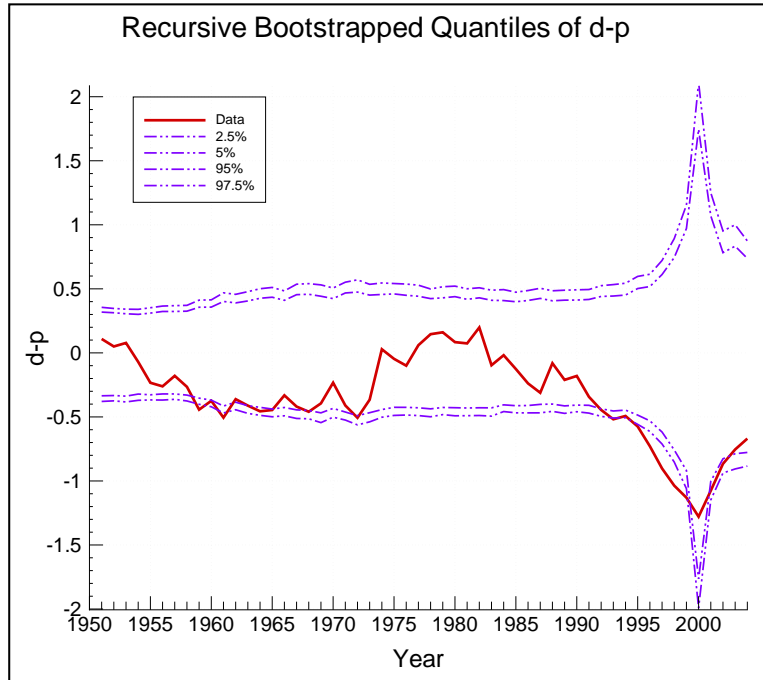


Figure 5: The Recursively Adjusted Dividend-Price Ratio.

The left panel plots the adjusted dividend-price ratio \widehat{dp}_t^P for the Perron model with recursive break estimation. The right panel plots the adjusted dividend-price ratio, \widehat{dp}_t^H , for Hamilton regime switching model, both with two and three regimes. In both figures, the adjusted series is rescaled so that it coincides with the adjusted series for the first subsample.

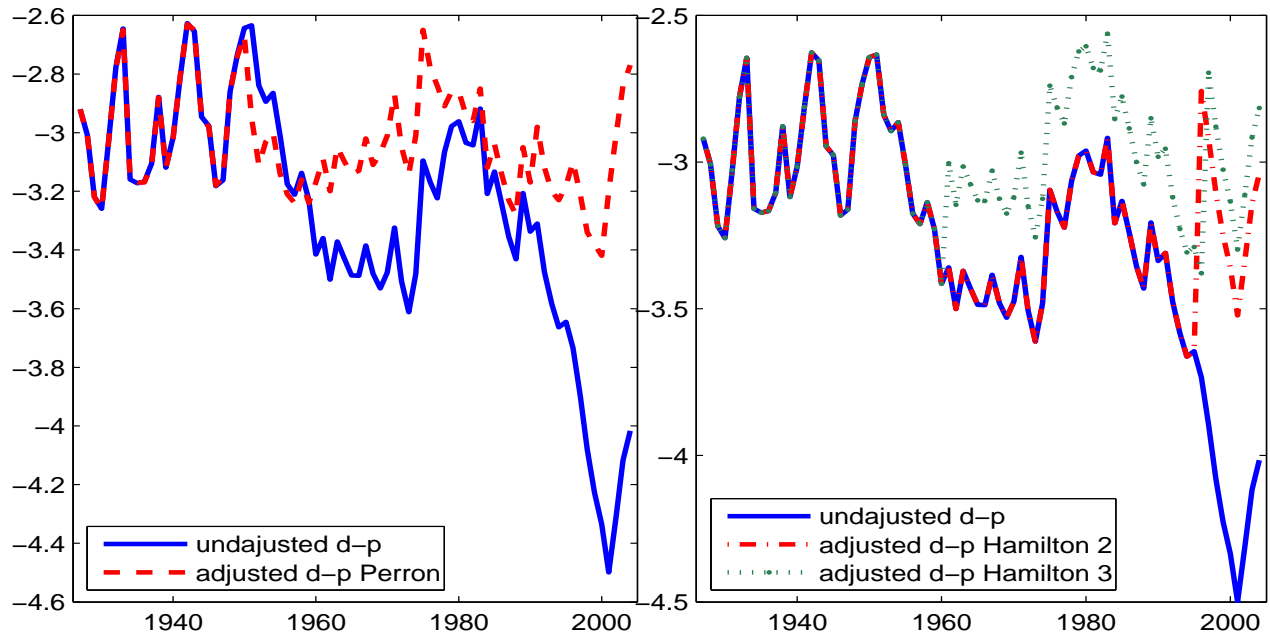


Figure 6: Regime Switching Model for the Log Dividend-Price Ratio - Two Regimes

this figure reports estimation results from a Hamilton regime-switching model for the log dividend-price ratio $d - p$.

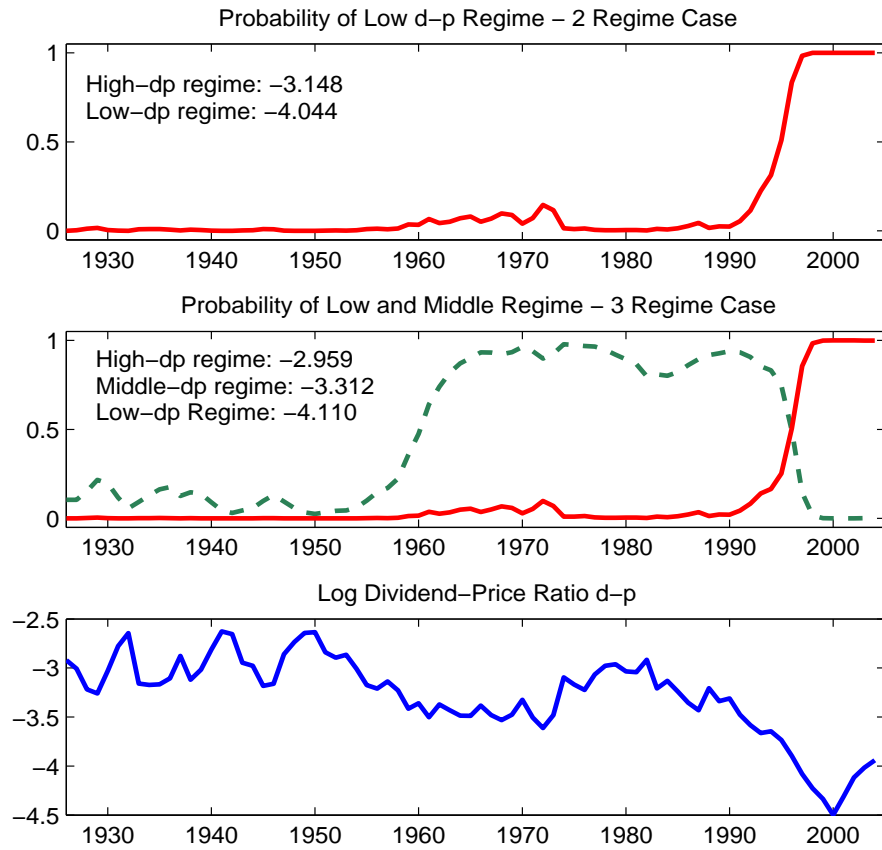


Figure 7: Forecasting Returns - Rolling Estimates, Real Time Adjustments

The figure plots κ_r estimates for 30-year rolling regressions of $r_{t+1} - \bar{r} = \kappa_r \hat{d}p_t + \tau_{t+1}^r$. The dividend-price ratio $\hat{d}p_t$ is recursively adjusted as described in three different ways: (i) recursive-Perron in the top panel, (ii) a 2-regime Hamilton model in the middle panel, and (iii) a 3-Regime Hamilton model in the bottom panel. The data are annual for 1927-2004.

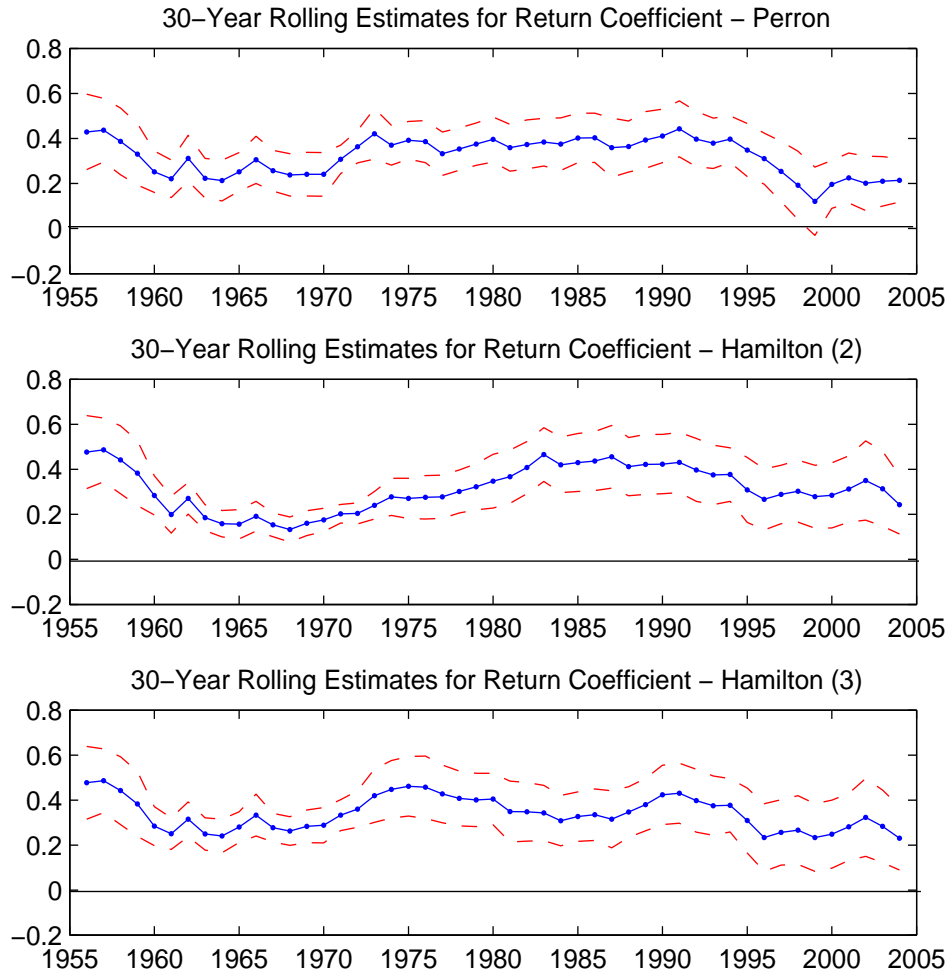


Figure 8: Estimation with Long-Horizon Moments

This figure compares the univariate OLS long-horizon regression coefficients, $\kappa_d(H) - OLS$ and $\kappa_r(H) - OLS$ to the GMM estimates that impose the present-value restriction (14), $\kappa_d(H) - GMM$ and $\kappa_r(H) - GMM$. The system contains 10 equations, 5 return and 5 dividend growth equations. The horizons (in years) are $H \in \{1, 3, 5, 7, 10\}$. The top panel uses the unadjusted \tilde{dp} ratio as predictor, the middle panel uses the \tilde{dp} ratio adjusted for one break in 1991, and the bottom panel uses the \tilde{dp} ratio adjusted for two breaks in 1954 and 1994.

