

Expected returns and expected dividend growth

Martin Lettau^a, Sydney C. Ludvigson^{a*}

^a*Department of Finance, Stern School of Business, New York University 44 West Fourth Street, New York, New York 10012, USA*

^{a*}*Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, New York, 10003, USA*

(Received 12 August 2003; received in revised form 15 March, 2004; accepted 24 May 2004)

Abstract

We investigate a consumption-based present-value relation that is a function of future dividend growth and find that changing forecasts of dividend growth are an important feature of the post-war U.S. stock market, despite the failure of the dividend-price ratio to uncover such variation. In addition, dividend forecasts are found to *covary* with changing forecasts of excess stock returns over business cycle frequencies. This covariation is important because positively correlated fluctuations in expected dividend growth and expected returns have offsetting effects on the log dividend-price ratio. The market risk premium and expected dividend growth thus vary considerably more than is apparent using the log dividend-price ratio alone as a predictive variable.

JEL classification: G12, G10

Keywords: Risk premia, dividend growth, cash-flow predictability, return predictability.

Lettau acknowledges financial support from the National Science Foundation. Ludvigson acknowledges financial support from the Alfred P. Sloan Foundation, from the National Science Foundation and the CV Starr Center for Applied Economics (NYU). This material is based upon work supported by the National Science Foundation under Grant No. 0224944. We thank Jushan Bai, John Y. Campbell, Kenneth French, Mark Gertler, Jonathan Lewellen, Anthony Lynch, Lucrezia Reichlin, Peter Schotman, William Schwert (the editor), Charles Steindel, Ross Valkanov, Charles Whiteman, an anonymous referee, and seminar participants at Duke University, INSEAD, London Business School, London School of Economics, Ohio State University, the New School for Social Research, New York University, SUNY Albany, the University of Iowa, the University of Maryland, the University of Montreal, Yale, the SITE 2001 summer conference, the CEPR Summer 2002 Finance Symposium and the 2003 American Finance Association meetings for helpful comments, and Nathan Barczy for excellent research assistance. Any errors or omissions are the responsibility of the authors, and do not necessarily reflect the views of the National Science Foundation.

*Corresponding author contact information: E-mail address: sydney.ludvigson@nyu.edu

0304-405X/02/ \$ see front matter © 2005 Published by Elsevier Science B.V. All rights reserved

1. Introduction

It is not necessary to delve far into recent surveys of the asset pricing literature to uncover a few key empirical results that have come to represent stylized facts, part of the “standard view” of U.S. aggregate stock market behavior:

1. Large predictable movements in dividends are not apparent in U.S. stock market data. In particular, the log dividend-price ratio does not have important long-horizon forecasting power for the growth in dividend payments.¹
2. Returns on aggregate stock market indexes in excess of a short-term interest rate are highly forecastable over long horizons. The log dividend-price ratio is extremely persistent and forecasts excess returns over horizons of many years.²
3. Variance decompositions of dividend-price ratios show that changing forecasts of future excess returns comprise almost all of the variation in dividend-price ratios. These findings form the basis for the conclusion that expected dividend growth is approximately constant.³

Empirical evidence on the behavior of the dividend-price ratio has transformed the way financial economists perceive asset markets. It has replaced the age-old view that expected returns are approximately constant, with the modern-day view that time-variation in expected returns constitutes an important part of aggregate stock market variability. Even the extraordinary behavior of stock prices during the late 1990s has not unraveled this transformation. Academic researchers have responded to this episode by emphasizing that—in contrast to stock market dividends—movements in aggregate stock prices have always played an important role historically in restoring the dividend-price ratio to its mean, even though the persistence of the dividend-price ratio implies that such restorations can sometimes take many years to materialize (Heaton and Lucas, 1999; Campbell and Shiller, 2001; Cochrane, 2001, Ch. 20; Fama and French, 2002; Campbell, 2003).

Yet there are noticeable cracks in the standard academic paradigm of predictability based on the dividend-price ratio. On the one hand, several researchers, focusing primarily on forecasting horizons of less than a few years, have argued that careful statistical analysis provides

¹A large literature documents the poor predictability of dividend growth by the dividend-price ratio over long horizons; see for example, Campbell (1991, 2003) and Cochrane (1991, 1994, 1997, 2001). Ang and Bekaert (2001) find somewhat stronger predictability; we discuss their results further below.

²See Fama and French (1988), Campbell and Shiller (1988), Hodrick (1992), Campbell et al. (1997), Cochrane (1997), Cochrane (2001, Ch. 20), and Campbell (2003).

³See Campbell (1991); Cochrane (1991); Hodrick (1992); Campbell et al. (1997), Ch. 7; Cochrane (2001), Ch. 20; Campbell (2003).

little evidence that the log dividend-price ratio forecasts returns (for example, Nelson and Kim, 1993; Stambaugh, 1999; Ang and Bekaert, 2001; Valkanov, 2003; Goyal and Welch, 2003). These findings have led some to conclude that changing forecasts of excess returns make a negligible contribution to fluctuations in the aggregate stock market.

On the other hand, researchers have employed predictive variables other than the dividend-price ratio and found evidence that excess returns on the aggregate stock market are strongly forecastable at horizons as short as a few quarters. Lettau and Ludvigson (2001a) find that excess stock returns are forecastable at horizons over which the dividend-price ratio has comparably weak forecasting power, namely at “business cycle” frequencies, those ranging from a few quarters to several years. Such predictable variation in returns is revealed not by the slow-moving dividend-price ratio, but instead by an empirical proxy for the log consumption-wealth ratio, denoted cay_t , a variable that captures deviations from the common trend in consumption, asset (nonhuman) wealth, and labor income. The consumption-wealth variable cay_t is less persistent than the dividend-price ratio, consistent with the finding that the former forecasts returns over shorter horizons than the latter.

Taken together, these empirical findings raise a series of puzzling questions. If some statistical analyses imply that the dividend-price ratio doesn’t forecast returns, does it follow that expected excess returns are approximately constant? If so, then why does cay_t have predictive power for excess returns at horizons ranging from a few quarters to several years? Moreover, if results using cay_t suggest that business cycle variation in expected returns is present, why does the dividend-price ratio have difficulty capturing this variation?

This paper argues that it is possible to make sense of these seemingly contradictory findings and in the process provide empirical answers to the questions posed above. We study a consumption-based present-value relation that is a function of future dividend growth. The evidence we present has two key elements. First, using data on aggregate consumption and dividend payments from aggregate (human and nonhuman) wealth, we show that changing forecasts of stock market dividend growth *do* make an important contribution to fluctuations in the post-War U.S. stock market, despite the failure of the dividend-price ratio to uncover such variation. Although U.S. dividend growth is known to have some short-run forecastability arising from the seasonality of dividend payments, to our knowledge this study is one of the few to find important predictability in direct long-horizon regressions, and in particular at horizons over which excess stock returns have been found to be forecastable. The second key element of our empirical findings is that the dividend forecasts we uncover are found to positively covary with changing forecasts of excess stock returns. Menzly et al. (2004) argue that such covariation can arise naturally in general equilibrium models that include a forecastable component in dividend growth. Brennan and Xia (2001) present a general equilibrium model that incorporates time-variation in the conditional mean of dividend growth.

Our findings help resolve the puzzles discussed above, for two reasons. First, the results help explain why the log dividend-price ratio has been found to be a relatively weak predictor of U.S. dividend growth, despite the evidence presented here that dividend growth is highly forecastable. Movements in expected dividend growth that are positively correlated with movements in expected returns have offsetting effects on the log dividend-price ratio. Second, the results help explain why business cycle variation in expected excess returns is uncovered by cay_t , but is not well captured by the dividend-price ratio. Movements in expected returns that are positively correlated with movements in expected dividend growth will have offsetting effects on the log dividend-price ratio, but not necessarily on the log consumption-wealth ratio.

We emphasize two implications of our findings. First, expected dividend growth is not constant, but instead varies over horizons ranging from one to six years. Thus, the variation in expected dividend growth that we uncover occurs at business cycle frequencies, not the ultra-low frequencies that dominate the sampling variability of the log dividend-price ratio. Second, common variation in expected returns and expected dividend growth will make it more difficult for the log dividend-price ratio to display statistically significant predictive power for future returns, consistent with evidence reported in Nelson and Kim (1993), Stambaugh (1999), Ang and Bekaert (2001), Valkanov (2003), and Goyal and Welch (2003), as well as for future dividend growth. Such findings should not be interpreted as an indication that expected returns are constant, however. On the contrary, the results in this paper suggest that the log dividend-price ratio will have difficulty revealing business cycle variation in the equity risk-premium precisely *because* expected returns fluctuate at those frequencies, and covary with changing forecasts of dividend growth. The findings presented here therefore suggest not only that expected returns vary, but that they vary by far more (over shorter horizons) than what can be revealed using the log dividend-price ratio alone as a predictive variable.

The rest of the paper is organized as follows. In the next section, we present an expression linking aggregate consumption and dividend payments from aggregate wealth, to the expected future growth rates of income flows from aggregate wealth. This delivers a present-value relation for future dividend growth in terms of observable variables. We then move on in Section 3 to discuss the data, and present results from estimating the common trend in log consumption and measures of the dividend payments from aggregate wealth. For the main part of our analysis, we focus on findings using the growth in dividends on the Center for Research in Security Prices (CRSP) value-weighted stock market index. This is done in order to make our results directly comparable with those from the existing asset pricing literature. Nevertheless, in Section 5.3 we show that our main conclusions are not altered by including aggregate share repurchases in the measure of dividends. In Section

4 we present the outcome of long-horizon forecasting regressions for dividend growth and excess returns on the U.S. stock market. Section 5 discusses one possible explanation for why expected dividend growth could be positively correlated with expected returns. Section 6 concludes.

2. A consumption-based present-value relation for dividend growth

We consider a representative agent economy in which all wealth, including human capital, is tradable. Let W_t be beginning-of-period aggregate wealth (defined as the sum of human capital, H_t , and nonhuman or asset wealth, A_t) in period t and let $R_{w,t+1}$ be the net return on aggregate wealth. For expositional convenience, we consider a simple accumulation equation for aggregate wealth, written

$$W_{t+1} = (1 + R_{w,t+1})(W_t - C_t). \quad (1)$$

Labor income Y_t does not appear explicitly in this equation because of the assumption that the market value of tradable human capital is included in aggregate wealth.⁴ Throughout this paper we use lower case letters to denote log variables, e.g., $c_t \equiv \log(C_t)$.

Defining $r \equiv \log(1 + R)$, Campbell and Mankiw (1989) derive an expression for the log consumption-aggregate wealth ratio by taking a first-order Taylor expansion of Eq. (1), solving the resulting difference equation for log wealth forward, and imposing a transversality condition. Ignoring a constant, the resulting expression holds to a first-order approximation:

$$c_t - w_t = E_t \sum_{i=1}^{\infty} \rho_w^i (r_{w,t+i} - \Delta c_{t+i}), \quad (2)$$

where $\rho_w \equiv 1 - \exp(\overline{c - w})$. This expression says that the log consumption-wealth ratio embodies rational forecasts of returns and consumption growth. We omit unimportant linearization constants in the equations throughout the paper.

Equation (2) is of little use in empirical work because aggregate wealth, w_t , includes human capital, which is not observable. Lettau and Ludvigson (2001a) address this problem by reformulating the bivariate cointegration relation between c_t and w_t as a trivariate cointegration relation involving three observable variables, namely c_t , a_t , and y_t , where a_t is the log of nonhuman or asset wealth, and y_t is log labor income. The resulting empirical “proxy” for

⁴None of the derivations below are dependent on this assumption. In particular, Eq. (3), below, can be derived from the analogous budget constraint in which human capital is nontradeable: $A_{t+1} = (1 + R_{a,t+1})(A_t + Y_t - C_t)$, where, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^j (1 + R_{a,t+i})^{-i} Y_{t+j}$.

the log consumption-aggregate wealth ratio is a consumption-based present-value relation involving future returns to asset wealth:⁵

$$cay_t \equiv c_t - \omega a_t - (1 - \omega) y_t = E_t \sum_{i=1}^{\infty} \rho_w^i (\omega r_{a,t+i} - \Delta c_{t+i} + (1 - \omega) \Delta y_{t+1+i}), \quad (3)$$

where ω is the average share of asset wealth, A_t , in aggregate wealth, W_t , and $r_{a,t}$ is the log return to asset wealth, A_t . Under the maintained hypothesis that asset returns, consumption growth and labor income growth are covariance-stationary, Eq. (3) says that consumption, asset wealth, and labor income are cointegrated, and that deviations from the common trend in c_t , a_t , and y_t summarize expectations of returns to either asset wealth, consumption growth, or labor income growth, or some combination of all three. The wealth shares ω and $(1 - \omega)$ are cointegration coefficients. We discuss their estimation below. The cointegration residual on the left-hand side of Eq. (3) is denoted cay_t for short. The cointegration framework says that, if risk premia vary over time (for any reason), cay_t is a likely candidate for predicting future excess returns. Both Eq. (2) and Eq. (3) are consumption-based present-value relations involving future returns to wealth.

In this paper we employ the same accounting framework to construct a consumption-based present-value relation involving future dividend growth from aggregate wealth. We can move from the consumption-based present-value relation involving future returns, Eq. (3), to one involving future dividend growth, by expressing the market value of assets in terms of expected future returns and expected future income flows. The general derivation is given in Campbell and Mankiw (1989), and the application to our setting is given in Appendix A. This derivation delivers a present-value relation involving the log of consumption, c_t , the log of dividends from asset wealth, d_t , and the log of dividends from human wealth, y_t :

$$cdy_t \equiv c_t - \nu d_t - (1 - \nu) y_t = E_t \sum_{i=1}^{\infty} \rho_w^i (\nu \Delta d_{t+i} + (1 - \nu) \Delta y_{t+i} - \Delta c_{t+i}). \quad (4)$$

Eq. (4) is a consumption-based present-value relation involving future dividend growth. Under the maintained hypothesis that Δd_t , Δy_t , and Δc_t are covariance-stationary, Eq. (4) says that consumption, dividends from asset wealth, and dividends from human capital (labor income) are cointegrated, and that deviations from their common trend—given by the left-hand side of Eq.(4)—provide a rational forecast of either dividend growth, labor income

⁵We will often refer loosely to cay_t as a proxy for the log consumption-aggregate wealth ratio, $c_t - w_t$. More precisely, Lettau and Ludvigson (2001a) find that cay_t is a proxy for the important predictive components of $c_t - w_t$ for future returns to asset wealth. Nevertheless, the left-hand side of (3) will be proportional to $c_t - w_t$ under the following conditions: first, expected labor income growth and consumption growth are constant and, second, the conditional expected return to human capital is proportional to the return to nonhuman capital.

growth, or consumption growth, or some combination of all three. The income shares ν and $(1 - \nu)$ are cointegration coefficients. We discuss their estimation further below. The cointegration residual on the left-hand side of Eq. (4) is denoted cdy_t , for short. The model implies that the consumption-aggregate dividend ratio is stationary and mean-reverting. Two recent theoretical models that specify the consumption-aggregate dividend ratio as a mean-reverting process are Santos and Veronesi (2000) and Longstaff and Piazzesi (2004).

It is instructive to compare cdy_t with the proxy for the consumption-aggregate wealth ratio cay_t in Eq. (3), studied in Lettau and Ludvigson (2001a). Eq. (3) says that if investors want to maintain flat consumption paths (i.e., expected consumption growth is approximately constant), fluctuations in cay_t reveal expectations of future returns to asset wealth, provided that expected labor income growth is not too volatile. This implication was studied in Lettau and Ludvigson (2001a). Analogously, Eq. (4) says that if investors want to maintain flat consumption paths, fluctuations in cdy_t summarize expectations of future dividend growth. This implication of the framework is studied here. Investors with flat consumption paths will appear to smooth out transitory fluctuations in dividend income stemming from time-variation in expected dividend growth. Consumption should be high relative to its long-run trend relation with d_t and y_t when dividend growth is expected to be higher in the future, and low relative to its long-run trend with d_t and y_t when dividend growth is expected to fall. Moreover, if expected consumption growth and expected labor income growth do not vary much, cdy_t should display relatively little predictive power for future consumption and labor income growth, but could forecast stock market dividend growth, if in fact expected dividend growth varies over time. In this case, Eq. (4) says that cdy_t is a state variable that summarizes changing forecasts of dividend growth to asset wealth.

The framework presented above does not predict which variables on the right-hand side of Eqs. (3) and (4) have time-varying conditional means and could therefore be forecastable. It states merely that, if there is predictable variation (for any reason) in the variables on the right-hand side of Eqs. (3) and (4), cay_t and cdy_t are good candidates for uncovering that variation. Of course, by imposing additional theoretical structure (structure that specifies preferences and technology), predictions could be developed about which variables ought to be forecastable according to specific theories. Importantly, however, the empirical investigation of this paper does not require imposing such additional structure, and we view this as one its main virtues. The results we obtain exploit only cointegration, a phenomenon that can be motivated by the logic of a simple budget constraint identity, applicable to a wide variety of theoretical structures.

It is instructive to compare the consumption-based present-value relations Eqs. (3) and (4) with the linearized expression for the log dividend-price ratio. Following Campbell and Shiller (1988), the log dividend-price ratio can be written (up to a first-order approximation

and ignoring constants) as

$$d_t - p_t = E_t \sum_{i=0}^{\infty} \rho^i (r_{t+1+i} - \Delta d_{t+1+i}), \quad (5)$$

where p_t be the log price of stock market wealth, d_t is the log dividend, $\rho \equiv \frac{1}{1+\exp(\bar{d}-p)}$, and r_t is the log return to stock market wealth. This equation says that if the log dividend-price ratio is high, agents must be expecting either high future returns on stock market wealth or low dividend growth rates.

Like the consumption-based expressions developed above, this expression does not predict which variables on the right-hand side should be forecastable; it suggests only that, if there is predictable variation in returns or dividend growth, $d_t - p_t$ is a good candidate for uncovering that variation. Many studies, cited in the introduction, have documented that $d_t - p_t$ forecasts returns over long horizons but explains little of the variability in future dividend growth. Other studies find that the forecasting power of $d_t - p_t$ for future excess returns over shorter horizons is statistically tenuous. As a consequence, expected dividend growth is often modeled as constant, and questions remain about the predictive power of $d_t - p_t$ for future returns. We argue here that such findings are consistent with an entirely different explanation of the data: both expected dividend growth and expected returns vary, but the two contain a common component and are therefore positively correlated. This makes forecastable movements in dividends and returns hard to detect statistically using the dividend-price ratio.

This be seen by simplifying the notation of Eq. (5) to

$$d_t - p_t = \eta_{r,t} - \eta_{d,t}, \quad (6)$$

where $\eta_{r,t} \equiv E_t \sum_{i=0}^{\infty} \rho^i r_{t+1+i}$ and $\eta_{d,t} \equiv E_t \sum_{i=0}^{\infty} \rho^i \Delta d_{t+1+i}$. Suppose $\eta_{r,t}$ and $\eta_{d,t}$ are each equal to the sum of an independent component and a common component:

$$\begin{aligned} \eta_{d,t} &= x_{d,t} + v_t \\ \eta_{r,t} &= x_{r,t} + v_t, \end{aligned}$$

where $x_{d,t}$ is the independent component of expected dividend growth, $x_{r,t}$ is the independent component of expected return volatility, and v_t is the component common to both expected returns and expected dividend growth. By assumption, $x_{d,t}$ and $x_{r,t}$ are uncorrelated, so that $\text{Var}(d_t - p_t) = \text{Var}(x_{d,t}) + \text{Var}(x_{r,t})$. Thus, it is clear that the common component, v_t , will have no influence on the dividend-price ratio, making a potentially important forecastable component of dividends and returns hard to detect using the dividend-price ratio. But the expressions in Eqs. (3) and (4) show that the forecasting power of cay_t for returns and

cdy_t for dividend growth is not compromised by such common variation in expected returns and expected dividend growth, because these variables are not (offsetting) functions of *both* expected returns and expected dividend growth. Thus, unlike $d_t - p_t$, cay_t captures variation in both components of expected returns, $x_{r,t}$ and v_t , while cdy_t captures variation in both components of expected dividend growth, $x_{d,t}$ and v_t . These considerations motivate the use of the consumption-based ratios developed above to uncover possible time-variation in expected returns and expected dividend growth that might be missed by the dividend price ratio.

Eq. (6) also implies that, if dividend growth rates are forecastable by some variable other than $d_t - p_t$, and if $d_t - p_t$ does not forecast dividend growth, it must be that the forecastable component in dividend growth is offset by a simultaneous movement in expected returns. The evidence we present below using consumption-based ratios is consistent with such a scenario.

It is important to emphasize that forecastable variation in dividend growth is entirely consistent with the “excess volatility” phenomenon of LeRoy and Porter (1981) and Shiller (1981). Cochrane (1991) illustrates this phenomenon by providing a variance decomposition of Eq. (5):

$$\text{Var}(p_t - d_t) = \text{Cov}\left(p_t - d_t, \sum_{j=1}^{\infty} \rho^{j-1} \Delta d_{t+j}\right) - \text{Cov}\left(p_t - d_t, \sum_{j=1}^{\infty} \rho^{j-1} r_{t+j}\right). \quad (7)$$

If discount rates were constant, returns would be unforecastable and the last term on the right-hand side would be zero. In this case, all variation in price-dividend ratios arises from covariation between $p_t - d_t$ and future dividend growth, i.e., from changing dividend forecasts where $d_t - p_t$ is the forecasting variable. Cutting the sums off at a 15-year horizon, Cochrane shows empirically that almost all of the variation in price-dividend ratios is attributable to changing return forecasts (the last covariance term on the right-hand side); price-dividend ratios are therefore said to be excessively volatile. This phenomenon remains a robust feature of the stock market even today; see Cochrane (2001) for an update of this variance decomposition using more recent data. Notice, however, that prices can be excessively volatile at the same time that dividend growth is highly forecastable, as long as most of the variation in expected dividend growth is attributable to the common component, v_t , in Eq. (6). This component has no influence on the dividend-price ratio and therefore no influence on the variance decomposition Eq. (7). It follows that substantial predictable variation in dividend growth can be entirely consistent with the excess volatility phenomenon documented by LeRoy and Porter (1981), Shiller (1981), and Cochrane (1991).

The approximate consumption identities developed above serve to motivate and interpret an investigation of whether consumption-based present-value relations might be informative

about the future path of dividend growth, asset returns, labor income growth or consumption growth. Nothing in the empirical investigation itself, however, is dependent on these approximations. But we can assess the implications of the framework by investigating whether empirical measures of these present-value relations contain information about the future path of consumption growth, labor income growth, or dividend growth from the aggregate stock market. We do so next.

3. The common trend in consumption, dividends and labor income

3.1. Data and preliminary analysis

Before we can estimate a common trend between consumption and measures of aggregate dividends, we need to address two data issues that arise from the use of aggregate consumption and dividend/income data. First, the log of asset wealth, a_t , is a measure of real, per capita household net worth, which includes all financial wealth, housing wealth, and consumer durables. Durable goods are properly accounted for as part of nonhuman wealth, A_t , a component of aggregate wealth, W_t , and so should not be accounted for as part of consumption or treated purely as an expenditure. Treating durables purchases purely as an expenditure (by, e.g., removing them from A_t and including them in C_t) is improper because it ignores the evolution of the asset over time, which must be accounted for by multiplying the stock by a gross return. In the case of many durable goods, this gross return would be less than one and would consist primarily of depreciation. The budget constraint Eq. (1) thus applies to the *flow* of consumption, C_t ; durables expenditures are excluded from C_t because they represent replacements and additions to a capital stock (investment), rather than a service flow from the existing stock. Total flow consumption is unobservable because we lack observations on the service flow from much of the durables stock. We therefore follow Blinder and Deaton (1985) and Campbell (1987) and use the log of real per capita expenditures on nondurables and services (excluding shoes and clothing), as a measure of c_t . This measure of consumption typically comprises over 85 % of total personal consumption expenditures. An internally consistent cointegration relation can then be obtained if we assume that the log of (unobservable) real total flow consumption is cointegrated with the log of real nondurables and services expenditures.

Second, we have experimented with constructing various empirical measures of nonstock dividends by adding forms of non-equity income from private capital, the largest component of which is interest income, to stock market dividends. In our sample, however, we find the strongest evidence of a common trend between log consumption, log stock market dividends,

and log labor income. A likely explanation is that the inflationary component of nominal interest income, along with the explicit inflation tax on interest income inherent in the U.S. tax code, makes real interest income difficult to measure. Such tax treatments can create peculiar trends in interest income that have nothing in particular to do with the evolution of permanent real interest income. These problems are especially evident in our sample during the 1970s and 1980s when nominal interest income grew rapidly because of inflation. In addition, we do not directly observe dividend payments from some forms of nonhuman, nonfinancial household net worth, primarily physical capital.⁶

The theory calls for the real component of nominal interest income, which is not directly observable. The nominal value of interest-bearing assets is observable, and can be put in real terms by deflating by a price level to get the component that should be associated with real consumption. But to obtain real interest income, one needs to also multiply this value by a real interest rate for the asset. (This is important because real interest income was negative for periods of the 1970s when inflation was high.) To construct the appropriate real interest rate for each asset class requires obtaining data on the nominal interest rate for that asset and subtracting off inflation. The difficulty is that the asset data are in the Flow of Funds accounts while nominal interest data are in the National Income and Product Accounts, and the asset classes are not the same in the two data bases.

Some researchers have documented a significant decline in the percentage of firms paying tax-inefficient dividends in data since 1978 (e.g., Fama and French, 2001). It might seem that such a phenomenon would create inflation-induced trends in aggregate stock market dividend income similar to those for nominal interest income. An inspection of the dividend data from the CRSP value-weighted index, however, reveals that—with the exception of the unusually large one-year decline in dividends in 2000, discussed below—the average annual growth rate of dividends has not declined precipitously over the period since 1978, or over the full sample. The average annual growth rate of real, per capita dividends is in fact higher, 5.6%, from 1978 through 1999, than the growth rate for the period 1948 to 1978. The annual growth rate for the whole sample (1948-2001) is 4.2%.

Fortunately, it is not necessary to include every dividend component from aggregate wealth in the expression Eq. (4) to obtain a consumption-based present-value relation that is a function of future stock market dividend growth, the object of interest in this study. As long as the excluded forms of dividend payments are cointegrated with the included forms

⁶One possibility is to use the product side of the national income accounts to estimate income flows as the residual from GDP less reported dividend and labor income. This approach creates its own problems, however, because it requires the income and product sides of the national accounts to be combined, and there is no way to know how much of the statistical discrepancy between the two is attributable to underestimates of income versus overestimates of output.

(as models with balanced growth would suggest), the framework above implies that the included dividend measures can be combined with consumption to obtain a valid cointegration relation. In this study, we use stock market dividends as a measure of dividend payments from nonhuman (asset) wealth, and we use d_t to denote stock market dividends from now on. If nonstock/nonlabor forms of dividend income are cointegrated with the dividend payments from stock market wealth, d_t , and/or human capital, y_t , the framework above implies a cointegration relation among c_t , stock market dividends d_t , and labor income y_t , and the resulting cointegration residual should reveal expectations over long horizons of either future Δd_t , Δy_t , or Δc_t , or some combination of all three.

These data considerations have two implications for our analysis. First, as discussed in Lettau and Ludvigson (2001a), the cointegration parameters in Eqs. (3) and (4) are likely to sum to a number less than one if, over the period covered by our data sample, there has been a secular decline in the share of nondurables and services in total consumption. In this case the log of total consumption can be expected to be proportional (up to a stationary error term) to the log of nondurables and services expenditures, with the constant of proportionality greater than unity (see Lettau and Ludvigson, 2001, for further discussion). Second, since some components of aggregate dividends are omitted in Eq. (4), the sum of the cointegration parameters in Eq. (3) is unlikely to be identical to the sum of the cointegration coefficients in Eq. (4). Both of these conclusions are based on a Monte Carlo analysis. For these reasons, we denote the wealth shares ω and $(1 - \omega)$ generically as cointegration coefficients α_a and α_y , respectively, and likewise denote the shares ν and $(1 - \nu)$ generically as cointegration coefficients β_d and β_y , respectively. The parameters $\hat{\alpha}_a$, $\hat{\alpha}_y$, $\hat{\beta}_d$, and $\hat{\beta}_y$ can be estimated using either single-equation or system cointegration methods. The estimated values of the cointegration residuals cay_t and cdy_t are denoted \widehat{cay}_t and \widehat{cdy}_t , respectively. Throughout this paper, we use “hats” to denote the estimated values of parameters.

The data used in this study are annual per capita variables, measured in 1996 dollars, and span the period 1948-2001. Following Cochrane (1994), we use annual data in order to insure that any forecastability of dividend growth we uncover is not attributable to the seasonality of dividend payments. Stock market dividends are measured as dividends on the CRSP value-weighted index and are scaled to match the units of consumption and labor income. Appendix B provides a detailed description of the sources and definitions of the data used in this study.

Table 1 first presents summary statistics for the log of real per capita consumption growth, labor income growth, dividend growth, the change in the log of the CRSP price index, Δp_t , and the change in the log of household net worth, Δa_t , all in annual data. Real dividend growth is considerably more volatile than consumption and labor income, having a standard deviation of 12% compared to 1.1% and 1.8% for consumption and labor income growth,

respectively. It is somewhat less volatile than the log difference in the CRSP value-weighted price index, which has a standard deviation of 16%, but is still more volatile than the log difference in net worth, which has a standard deviation of 4%. Consumption growth and labor income growth are strongly positively correlated, as are Δp_t and Δa_t . Annual real consumption growth and real dividend growth have a weak correlation of -0.16.

We begin our analysis by testing for both the presence and number of cointegration relations in the system of variables $\mathbf{x}'_t \equiv [c_t, d_t, y_t]'$. Such tests have already been performed for the system $\mathbf{v}'_t = [c_t, a_t, y_t]'$ on quarterly data in Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2004). The results using annual data are contained in Appendix C. We assume that all of the variables contained in \mathbf{x}_t and \mathbf{v}_t are first-order integrated, or $I(1)$, an assumption verified by unit root tests. Test results presented in Appendix C suggest the presence of a single cointegration relation for both vector time series. Normalizing each cointegration vector by setting the coefficient on consumption to unity, we denote the single cointegration relation for $\mathbf{v}'_t = [c_t, a_t, y_t]'$ as $\boldsymbol{\alpha}' = (1, -\alpha_d, -\alpha_y)'$, and for $\mathbf{x}'_t = [c_t, d_t, y_t]'$ as $\boldsymbol{\beta}' = (1, -\beta_d, -\beta_y)'$.

The evidence strongly suggests the presence of a single, *trivariate*, cointegration relation, i.e., one that involves all three variables in each cointegrated system. This of course does not imply that there is no bivariate cointegration in population, as cointegration tests are known to have low power in finite samples when the relations involved are characterized by slowly mean-reverting deviations from trend. Regardless of the number of linearly independent cointegration relations present, it is always the case that econometric procedures can only identify the space spanned by these cointegration relations; a normalization is required to obtain linearly independent relations.

The cointegration parameters α_d , α_y and β_d , β_y must be estimated. We use a dynamic least squares procedure that delivers asymptotically optimal estimates (Stock and Watson, 1993). Two leads and lags of the first differences of Δy_t and Δd_t are used in the dynamic least squares regression. This procedure estimates $\hat{\boldsymbol{\beta}}' = (1, -0.13, -0.68)'$. The Newey-West corrected t -statistics (Newey and West, 1987) for these estimates are -10.49 and -34.82, respectively. We denote the estimated cointegration residual $\hat{\boldsymbol{\beta}}' \mathbf{x}_t$ as \widehat{cdy}_t . The estimated cointegration vector for $\mathbf{v}'_t = [c_t, a_t, y_t]'$ is $\hat{\boldsymbol{\alpha}}' = (1, -0.33, -0.57)'$, very close to that obtained in Lettau and Ludvigson (2001a) using quarterly data. The Newey-West corrected t -statistics for these estimates are -14.32 and -30.48, respectively. The cointegration residual $\hat{\boldsymbol{\alpha}}' \mathbf{x}_t$ is denoted \widehat{cay}_t .

Table 2 displays autocorrelation coefficients for $d_t - p_t$, \widehat{cay}_t , and \widehat{cdy}_t . It is well known that the dividend-price ratio is very persistent. In annual data from 1948 to 2001 it has a first-order autocorrelation 0.87, a second-order autocorrelation of 0.72. and a third-order autocorrelation of 0.59. The autocorrelations of \widehat{cdy}_t and \widehat{cay}_t are much lower and die out

more quickly. Their first-order autocorrelation coefficients are 0.46 and 0.46, respectively; their second order autocorrelation coefficients are 0.14 and 0.20, respectively.

In Fig. 1 we plot the de-measured values of \widehat{cdy}_t and \widehat{cay}_t over the period 1948-2001. The sample correlation between \widehat{cdy}_t and \widehat{cay}_t is 0.41. The figure shows that the two consumption-based present-value relations tend to move together over time, although there are some notable episodes in which they diverge. One such episode is the year 2000, when an extraordinary 30% decline in recorded dividends (an extreme outlier in our sample) pushed \widehat{cdy}_t well above its historical mean.

To better understand the time-series properties of \widehat{cay}_t and \widehat{cdy}_t , it is useful to examine estimates of vector error-correction representations, or cointegrated vector autoregressions (VARs) for $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$. The vector error-correction representation for $(c_t, a_t, y_t)'$ takes the form

$$\Delta \mathbf{x}_t = \mathbf{v} + \gamma \widehat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1} + \mathbf{\Gamma}(L) \Delta \mathbf{x}_{t-1} + \mathbf{e}_t, \quad (8)$$

where $\Delta \mathbf{x}_t$ is the (3×1) vector of log first differences for $(\Delta c_t, \Delta a_t, \Delta y_t)'$, \mathbf{v} and γ are (3×1) vectors, $\mathbf{\Gamma}(L)$ is a finite order distributed lag operator, and $\widehat{\boldsymbol{\alpha}}$ is the (3×1) vector of previously estimated cointegration coefficients. A directly analogous set of equations holds for $\Delta \mathbf{x}_t = (\Delta c_t, \Delta d_t, \Delta y_t)'$ where $\widehat{\boldsymbol{\alpha}}$ is replaced by $\widehat{\boldsymbol{\beta}}$. Note that in the error-correction representation, $\widehat{\boldsymbol{\alpha}}' \mathbf{x}_{t-1}$ is just the cointegration residual, \widehat{cay}_t . Thus, in the error-correction representation, the growth rate in each variable is regressed on lags of the growth rates of all the other variables in the system, and on one lag of the estimated cointegration residual, \widehat{cay}_t in the case of the system involving $(\Delta c_t, \Delta a_t, \Delta y_t)'$, and \widehat{cdy}_t in the case of the system involving $(\Delta c_t, \Delta d_t, \Delta y_t)'$.

Table 3 presents the results of estimating first-order cointegrated VARs for c_t , a_t and y_t , and for c_t , d_t , and y_t .⁷ The table reveals several noteworthy properties of the data on consumption, household wealth, stock market dividends, and labor income. First, Panel A shows that the estimated cointegration residual \widehat{cay}_{t-1} is a strong predictor of wealth growth. Both consumption and labor income growth are somewhat predictable by lags of either consumption growth and/or wealth growth, as noted elsewhere (Flavin, 1981; Campbell and Mankiw, 1989), but the adjusted R^2 statistics (especially for the labor income equation) are lower than those for the asset regression. More importantly, the cointegration residual \widehat{cay}_{t-1} is an economically and statistically significant determinant of next period's asset growth, but not next period's consumption or labor income growth. This finding implies that asset wealth is mean-reverting, and adjusts over long horizons to match the smoothness of consumption

⁷The VAR lag lengths were chosen in accordance with findings from Akaike and Schwartz tests. The second system is also studied in Ludvigson and Steindel (1999).

and labor income. These results are consistent with those in Lettau and Ludvigson (2001a) using quarterly data.

Second, Panel B shows that the estimated cointegration residual \widehat{cdy}_{t-1} is a strong predictor of dividend growth. These results come from estimates from a cointegrated VAR for c_t , d_t , and y_t . The findings are analogous to those for the cointegrated VAR involving c_t , a_t , and y_t . Consumption and labor income are predictable by lagged consumption and wealth growth, but not by the cointegration residual \widehat{cdy}_{t-1} . What is strongly predictable by the cointegration residual is stock market dividend growth: \widehat{cdy}_{t-1} is both a statistically significant and economically important predictor of next year's dividend growth, Δd_t . These findings imply that when log dividends deviate from their habitual ratio with log labor income and log consumption, it is not consumption or labor income but dividends that are forecast to slowly adjust until the cointegration equilibrium is restored. As for asset wealth, dividends are mean-reverting and adapt over long horizons to match the smoothness in consumption and labor income.

4. Long-horizon forecasting regressions

A more direct way to understand mean reversion is to investigate regressions of long-horizon returns and dividend growth on the consumption ratios \widehat{cdy}_{t-1} and \widehat{cay}_{t-1} . The theory behind Eqs. (3) and (4) makes clear that both ratios should track longer-term tendencies in asset markets, rather than provide accurate short-term forecasts of booms or crashes. To investigate such a possibility, Table 4 presents the results of univariate regressions of the return on the CRSP value-weighted stock market index in excess of the three-month Treasury bill rate, at horizons ranging from one to six years. In each regression, the dependent variable is the H -period log excess return, $r_{t+1} - r_{f,t+1} + \dots + r_{t+H} - r_{f,t+H}$, where $r_{f,t}$ is used to denote the Treasury bill rate, or “risk-free” rate. The independent variable is either $d_t - p_t$, \widehat{cay}_t , or \widehat{cdy}_t . The table reports the estimated regression coefficient, the adjusted R^2 statistic in square brackets, and a heteroskedasticity and autocorrelation-consistent t -statistic for the hypothesis that the regression coefficient is zero (in parentheses). The table also reports, in curly brackets, the rescaled t -statistic recommended by Valkanov (2003) for the hypothesis that the regression coefficient is zero. We discuss this rescaled t -statistic below. Table 5 presents the same output for predicting long-horizon CRSP dividend growth, $\Delta d_{t+1} + \dots + \Delta d_{t+H}$.

The first row of Table 4 shows that the log dividend-price ratio has little power for forecast aggregate stock market returns from one to six years in this sample. These results differ from those reported elsewhere, primarily because we have included the last few years of stock market data in the sample. Variation in the log dividend-price ratio is too persistent to

display statistical evidence of stationarity in this sample (see Appendix C). The extraordinary increase in stock prices in the late 1990s substantially weakens the statistical evidence for predictability by $d_t - p_t$ that had been a feature of previous samples. If we end the sample in 1998, the log dividend-price ratio displays forecasting power for excess returns, but its strongest forecasting power is exhibited over horizons that are far longer than that exhibited by the consumption-wealth ratio proxy, \widehat{cay}_t (see Lettau and Ludvigson, 2001).⁸ It is possible that future samples will restore the statistical forecasting power of $d_t - p_t$ for future returns in these types of regressions. Lettau et al. (2003), however, find evidence of a structural break in the mean price-dividend ratio in the mid-1990s, raising the possibility that the historical relation between $d_t - p_t$ and future returns has changed.

In contrast to the dividend-price ratio, the second row of Table 4 shows that \widehat{cay}_t has statistically significant forecasting power for future excess returns at horizons ranging from one to six years. This evidence is consistent with that reported in Lettau and Ludvigson (2001a) using quarterly data. Using this single variable alone achieves an \overline{R}^2 of 0.25 for excess returns at a one-year horizon, 0.44 for excess returns over a two-year horizon, and 0.50 for excess returns over a six-year horizon.

The remaining row of Table 4 gives an indication of the forecasting power of \widehat{cdy}_t for long-horizon excess returns. At a one-year horizon, \widehat{cdy}_t displays little statistical forecasting power for future returns in this sample. For returns over all longer horizons, however, this present-value relation for dividend growth displays forecasting power for future returns. In addition, the coefficients from these predictive regressions are positive, so high values of \widehat{cdy}_t forecast high excess returns as do high values of \widehat{cay}_t . The t -statistics are above four for all horizons in excess of one year, and the \overline{R}^2 statistic rises from 0.19 at a two-year horizon to 0.36 at a six-year horizon. Because both \widehat{cay}_t and \widehat{cdy}_t are positively related to future excess returns, the results imply that both capture some component of time-varying expected returns.

We now turn to forecasts of long-horizon dividend growth. Table 5 displays results from the same forecasting exercise for long-horizon dividend growth, presented in Table 4 for excess returns. Previous studies (cited in the introduction) have found little forecastability in dividend growth using the dividend-price ratio as a predictor. Table 5, by contrast, suggests that $d_t - p_t$ seems to forecast dividend growth, but with the wrong (positive) sign. Such

⁸Some researchers using different statistical approaches find that the dividend-price ratio remains a predictor of excess stock returns even in samples that include recent data. Lewellen (2003) notes that when the dividend-price ratio is very persistent but nonetheless stationary, episodes in which the dividend yield deviates from its long-run mean for an extended period of time do not necessarily constitute evidence against predictability. Similar results are reported in recent work by Campbell and Yogo (2002), who find evidence of return predictability by financial ratios if one is willing to rule out an explosive root in the ratios.

predictability is at odds with the theory: high prices (relative to dividends) should forecast higher real dividend growth, not lower. This finding on the sign has been reported elsewhere (e.g., Cochrane, 1991, Campbell, 2003), but the statistical significance of the coefficient displayed in Table 5 is new. Table 5 also shows, however, that the statistical significance is entirely driven by the last two observations in our sample, which include the highly unusual year-2000 observation for dividends. Using conventional t -statistics, there is no evidence of a statistically reliable relation between the dividend-price ratio and future dividend growth from one-to-three-year horizons in the 52 years of data over the period 1948-1999. Using the Valkanov corrected t -statistics, there is no evidence at any horizon in either sample. We are therefore reluctant to place emphasis on this result.

Rows 2 and 3 present the results of predictive regressions using \widehat{cay}_t and \widehat{cdy}_t . The consumption-based present-value relation for future dividend growth, \widehat{cdy}_t , has strong forecasting power for future dividend growth at horizons ranging from one to six years. The individual coefficients are highly statistically significant, and the regression results suggest that the variable explains between 20% and 40% of future dividend growth, depending on the horizon. Lettau and Ludvigson (2001a) find that \widehat{cay}_t has predictive power for future returns; Row 2 shows that it also has statistically significant predictive power for dividend growth rates in our sample, with high values of \widehat{cay}_t predicting high dividend growth rates. The forecasting power of \widehat{cay}_t is, however, weaker than that displayed by \widehat{cdy}_t at every horizon in excess of one year (row 3). For example, at a four-year horizon, \widehat{cdy}_t explains about 21% of the variation in dividend growth, while \widehat{cay}_t explains 14%. At a six-year horizon, \widehat{cdy}_t explains 38% percent of the variation in dividend growth, while \widehat{cay}_t explains 18%. Still, just as for excess returns, the results suggest that both \widehat{cay}_t and \widehat{cdy}_t capture some component of time-varying expected dividend growth.

Table 6 shows that neither \widehat{cay}_t or \widehat{cdy}_t have long-horizon forecasting power for consumption growth or labor income growth. The first panel suggests that consumption growth is slightly predictable by these variables at a one-year horizon, but this is only because the variables proxy for serial correlation in consumption growth: their predictive power vanishes when lagged consumption growth is included as a regressor (Panel 2). Thus, we find predictable variation in dividend growth by \widehat{cdy}_t that is independent of predictable variation in consumption growth. This relates to how consumption and dividends are typically modeled theoretically. Dividends are often modeled in the macro-asset pricing literature as a simple levered claim to aggregate consumption, a specification that allows dividends to be more volatile than aggregate consumption but makes the two perfectly correlated. This makes some sense because stock market cash flows are in fact far more volatile than aggregate consumption (Table 1). The evidence presented here, however, suggests that dividends and consumption are not perfectly correlated: the consumption-aggregate dividend

ratio varies over time and predicts future dividend growth, but not future consumption growth. This evidence is supportive of more sophisticated modeling approaches to specifying the relation between consumption and dividends such as those found in recent work by Santos and Veronesi (2000) and Longstaff and Piazzesi (2004). These models specify that the consumption-dividend ratio is mean-reverting and can generate forecastable variation in dividend growth not evident in consumption growth, consistent with the empirical evidence presented here. The authors show that this type of modeling strategy can generate a substantial equity premium and greater volatility in stock returns compared to models in which dividends are a simple deterministic function of aggregate consumption.

The results in Tables 4 and 5 suggest that there is common variation in expected returns and expected dividend growth. The consumption-wealth ratio proxy, \widehat{cay}_t , which is a strong predictor of excess stock market returns, is also a predictor of stock market dividend growth. Conversely, \widehat{cdy}_t , a strong predictor of stock market dividend growth, is also a predictor of excess stock market returns. A natural question is whether either variable has independent information for future excess returns and dividend growth. To address this question, Table 7 presents the results of multivariate regressions of long-horizon excess returns (upper panel) and dividend growth (lower panel) using \widehat{cay}_t and \widehat{cdy}_t as regressors. For comparison, we also include $d_t - p_t$ in the multivariate regressions. Table 7 shows that, in forecasting long-horizon excess returns, \widehat{cdy}_t contains little information about future returns that is independent of that contained in \widehat{cay}_t : at most forecasting horizons, \widehat{cay}_t drives out \widehat{cdy}_t ; at the horizons for which this is not true, \widehat{cdy}_t adds only marginally to the overall fit. But note that even though both variables convey information about future returns and future dividend growth, \widehat{cay}_t contains some information about future returns that is independent of that contained in \widehat{cdy}_t . This suggests the presence of an independent component in expected excess returns, corresponding to the component $x_{r,t}$ above. The log dividend-price ratio has some independent forecasting power for returns at five- and six-year horizons, but none at earlier horizons. This suggests the presence of a long-horizon component in expected returns that is independent of the business-cycle component captured by \widehat{cay}_t .

The second panel of Table 7 shows that much the opposite pattern is borne out in long-horizon forecasting regressions for dividend growth: although both \widehat{cay}_t and \widehat{cdy}_t seem to have some independent information about future dividend growth over one or two year horizons, in a bivariate regression \widehat{cdy}_t drives out \widehat{cay}_t in forecasting future dividend growth at all forecasting horizons greater than three years. For a forecasting horizon of three years, the information contained in \widehat{cay}_t and \widehat{cdy}_t is apparently sufficiently similar that the regression has difficulty distinguishing their independent effects (although \widehat{cdy}_t is statistically significant at the 6% percent level). Accordingly, \widehat{cay}_t and \widehat{cdy}_t are not marginally significant predictors of dividend growth over a three year horizon, but they are strongly jointly significant (the

p -value for the F -test is less than 0.000001). In the trivariate regression including $d_t - p_t$, we find that, at horizons greater than one-year, the dividend-price ratio is a marginally significant predictor, but again high prices forecast low real dividend growth in this sample.

The findings in Table 7 suggest that much of the variation in expected dividend growth is common to variation in expected returns, at least for two- and three-year horizons. The findings also suggest that there is a component of expected returns, as captured by cay_t , that moves independently of expected dividend growth. This is analogous to the example presented in Section 2. If much of the variation in expected dividend growth is common to variation in expected returns, we would not expect innovations in expected dividend growth to have an important effect on the log dividend-price ratio, for the reasons discussed in Section 2. By contrast, if there were a component of expected returns that is independent of expected dividend growth, we would expect innovations in expected returns to have a positive effect on the log dividend-price ratio, again for the reasons discussed in Section 2.

One way to evaluate these possibilities is to estimate elasticities of the dividend-price ratio with respect to innovations in expected dividend growth and expected returns, as captured by \widehat{cdy}_t and \widehat{cay}_t , respectively. Such estimates can be accomplished by running regressions of $d_t - p_t$ on innovations in \widehat{cdy}_t and \widehat{cay}_t . The output below is generated by regressing $d_t - p_t$ on the residuals, $\varepsilon_{cdy,t}$ and $\varepsilon_{cay,t}$, from first-order autoregressions for \widehat{cdy}_t and \widehat{cay}_t , respectively. The lagged log dividend-ratio is also included as a regressor to control for the substantial persistence in $d_t - p_t$. citeferson/sarkissian/simin:03 point out that, when the predictable component of the regressand is large (as it is for the extremely persistent price-dividend ratio), spurious regression is a concern if the regressor is also sufficiently persistent, even if it is stationary. The potential for this problem can be eliminated by including the lagged dependent variable in the regression. The estimation output from our regressions using data from 1948 to 2001, with t -statistics in parentheses, is

$$\begin{aligned}
 d_t - p_t &= \underset{(-1.45)}{-0.06} + \underset{(18.89)}{0.96} (d_{t-1} - p_{t-1}) - \underset{(-1.0)}{1.31} \varepsilon_{cdy,t} \\
 d_t - p_t &= \underset{(-1.41)}{-0.05} + \underset{(22.02)}{0.97} (d_{t-1} - p_{t-1}) + \underset{(2.73)}{4.24} \varepsilon_{cay,t}.
 \end{aligned}$$

These results confirm the intuition suggested by the long-horizon forecasting regressions presented above. Innovations in expected dividend growth, as proxied by $\varepsilon_{cdy,t}$, have no statistically significant effect on $d_t - p_t$, consistent with the finding that much of the variation in expected dividend growth is common to variation in expected returns. By contrast, innovations in expected returns, as proxied by $\varepsilon_{cay,t}$, are statistically significant at conventional significance levels, consistent with the presence of a component of expected returns that moves independently of expected dividend growth and is therefore correlated with $d_t - p_t$. Of course, $d_t - p_t$ could still have trouble displaying statistically reliable forecasting power

for future returns, because unlike cay_t , it has no way of isolating such an independent component.

It is natural to ask whether other variables typically used in the business cycle and stock return forecasting literatures have forecasting power for dividend growth. Table 8 provides an answer from univariate long-horizon regressions of dividend growth on a host of other forecasting variables. One regression (row 1) uses just one-period lagged dividend growth as a predictor. We also report the results of using the term spread, TRM (ten-year Treasury bond minus the one-year Treasury bond, denoted), the default spread, DEF (the BAA minus AAA corporate bond rates), a short-term interest rate, $RREL$ (the annualized three-month Treasury bill rate minus its four-quarter moving average), and the Cochrane and Piazzesi (2002) bond risk factor (a linear combination of forward rates) as predictive variables for dividend growth. We use the largest sample period over which each of these variables is available (see the notes to Table 8). Lagged dividend growth has modest predictive power at one- and three-year horizons, but the R^2 statistics are small. Other than lagged dividend growth, most of these variables have little forecasting power. One exception is Cochrane and Piazzesi's bond risk factor, which has statistically significant forecasting power for one-year-ahead dividend growth, with a modest R^2 of 3%. (Since this variable is a monthly variable, we used the end-of-year value of the variable as our annual observation. This had greater predictive power than the average over the year.) In general, however, the results suggest that the consumption-based ratios studied here are superior as predictors of stock market dividend growth to most other variables used in the stock return and business cycle forecasting literatures. Finally, we note that while even a naive measure of detrended dividends created from a one-sided band-pass filter has some univariate forecasting power for future dividend growth, results (not reported) show that its forecasting power is driven out in a multivariate regression by \widehat{cdy}_t .

We can also ask whether \widehat{cdy}_t and \widehat{cay}_t have forecasting power for other measures that might be related to aggregate consumption and/or dividend growth, such as aggregate output and investment. Results from forecasting real per capita GDP and investment growth, by including \widehat{cdy}_t or \widehat{cay}_t in multivariate forecasting regressions that also use $RREL$, TRM , and DEF , show that \widehat{cdy}_t and \widehat{cay}_t have strongly significant marginal forecasting power over a one-year horizon, but generally do not display statistically significant forecasting power at longer horizons. Lettau and Ludvigson (2002) considered the forecasting power of \widehat{cay}_t for future investment growth in quarterly data, and report coefficients with the same negative sign as here; see Lettau and Ludvigson (2002) for an economic interpretation. The results reported in this study—using annual data and a different sample—are qualitatively similar, but seem to imply slightly less forecasting power at longer horizons than the previous study showed. These results are available upon request.

4.1. Related empirical findings

The evidence presented above suggests that there is important predictability of dividend growth over long horizons, and that predictable variation in dividend growth is correlated with that in excess returns. To our knowledge, such evidence is largely new. Other researchers (cited in the introduction) have found that dividend growth predictability—if evident at all in long-horizon regressions—occurs at relatively short horizons and is not highly correlated with predictable variation in excess returns. More recently, Kothari et al. (2004) study the stock market’s reaction to aggregate earnings news and find indirect evidence that expected cash flows and discount rates move together over time, consistent with the findings reported in this paper. Goyal and Welch (2003) show that the dividend-price ratio forecasts itself more than it forecasts returns or dividend growth. This finding, in conjunction with an approximate identity defining the log stock return, implies that there must be some offsetting movements in expected returns and expected dividend growth. Ang (2002) investigates the forecastability of long-horizon dividend growth for the aggregate stock market using annual data from 1927-2000. Although Ang concludes that there may be some long-horizon forecastability of dividend growth based on results from rolling forward a first-order vector autoregression for dividend yields, dividend growth rates and returns, he finds little evidence of predictability in long-horizon dividend growth from direct long-horizon regressions. These findings are consistent with those of the earlier papers cited in the introduction which use the log dividend-price ratio as a predictive variable, and our own results using $d_t - p_t$, reported above.

One recent study that finds predictability of dividend growth in direct long-horizon regressions is Ang and Bekaert (2001), who report results based on observations from the fourth quarter of 1952 to the fourth quarter of 1999 on the S&P 500 stock market index. Like Ang (2002), they confirm the earlier findings of Campbell (1991) and Cochrane (1991) that dividend growth is largely unpredictable by the dividend-price ratio in univariate long-horizon forecasting regressions. As Campbell (1991) and Cochrane (1991) emphasize, such findings imply that changing forecasts of future returns must comprise most of the variation in the dividend-price ratio. But Ang and Bekaert (2001) find that the dividend-price ratio has marginal predictive power for future dividend growth in a multivariate regression once the earnings yield is included as a regressor. (The earnings yield also has marginal predictive power.) There are two main differences between our predictability results and those in Ang and Bekaert (2001). First, the joint forecasting power of the dividend yield and the earnings yield for dividend growth is concentrated at shorter horizons than in regressions using \widehat{cdy}_t and \widehat{cay}_t . Second, the R^2 s for the regressions using the former variables are substantially lower than those using the latter. For example, in the sample used in Ang and Bekaert

(2001), the dividend yield and the earnings yield jointly explain about 21% of dividend growth one year ahead, and about 13% over a five-year horizon. The comparable numbers using \widehat{cdy}_t alone as a predictive variable are 31% and 34%. These numbers are higher than those reported in Table 4 because we use the slightly shorter sample employed by Ang and Bekaert (2001) in order to make the results directly comparable to theirs.

4.2. Additional statistical tests

4.2.1. Multivariate long-horizon forecasting regressions⁹

The cointegration coefficients in \widehat{cay}_t and \widehat{cdy}_t are estimated using the full sample. This estimation strategy is appropriate for testing the forecasting implications of the theoretical framework above, because sufficiently large samples of data are necessary to recover the true cointegration coefficients, and there is no implication (either from the theoretical framework or from statistical theory) that \widehat{cay}_t and \widehat{cdy}_t should forecast the right-hand side variables in Eqs. (3) and (4) unless the cointegration coefficients have converged to their true values. Fortunately, cointegration coefficients are “superconsistent,” converging to their true values at a rate proportional to the sample size T , and can therefore be treated as known in second-stage forecasting regressions. It follows that a valid test of the predictability implied by the theoretical cointegration framework in Eqs. (3) and (4) requires the use of the full sample to estimate the cointegration coefficients in \widehat{cay}_t and \widehat{cdy}_t . This issue is discussed in more detail in Lettau and Ludvigson (2001b).

A separate issue concerns not whether \widehat{cay}_t and \widehat{cdy}_t have forecasting power, but whether a practitioner, operating in the early part of our sample and without access to the whole sample to estimate cointegration coefficients, could have exploited the forecasting power of \widehat{cay}_t and \widehat{cdy}_t . Out-of-sample or subsample analysis is often used to assess questions of this nature. A difficulty with these procedures, however, is that the subsample analysis inherent in out-of-sample forecasting tests entails a loss of information, and as such makes out-of-sample tests substantially less powerful than in-sample forecasting tests (Inoue and Kilian, 2002). This loss of power means that out-of-sample (and subsample) analyses can fail to reveal true forecasting power that even a practitioner could have had in real time.

With these considerations in mind, we now provide an alternative approach to assessing the forecasting power of \widehat{cay}_t and \widehat{cdy}_t . The approach we propose eliminates the need to estimate cointegration parameters using the full sample in a first-stage regression, but at the same time avoids the power problems inherent in out-of-sample and subsample analyses. We

⁹We are grateful to Jushan Bai for pointing out the possibility of using the methodology used in this subsection.

consider single-equation, multivariate regressions taking the form

$$z_{t+h} = a + b_1 c_t + b_2 a_t + b_3 y_t + u_t, \quad (9)$$

where a , b_1 , b_2 , and b_3 are regression coefficients to be estimated, and the dependent variable z_{t+h} is either the h -period excess return on the CRSP value-weighted index or the h -period dividend growth rate on the CRSP value-weighted index. Rather than estimating the cointegration relation among c_t , a_t , and y_t in a first-stage regression and then using the cointegration residual as the single right-hand side variable, the regression Eq. (9) uses the multiple variables involved in the cointegration relation as regressors directly. If there is a relation between the left-hand side variable to be forecast, and some stationary linear combination of the regressors c_t , a_t , and y_t , the regression can freely estimate the non-zero coefficients b_1 , b_2 , and b_3 that generate such a relation. For this exercise, we maintain the hypothesis that the left-hand-side variable is stationary, while the right-hand side variables are $I(1)$. Then, under the null hypothesis that $(c_t, a_t, y_t)'$ has a single cointegration relation, it is straightforward to show that the limiting distributions for b_1 , b_2 , and b_3 will be standard, implying that the forecasting regression Eq. (9) will produce valid R^2 and t -statistics. Because this procedure does not require any first-stage estimation of cointegration parameters, it is clear that the forecasting regression Eq. (9), in particular its coefficients and R^2 statistics, cannot be influenced by such a prior analysis.

Table 9 reports long-horizon regression results for excess returns and dividend growth, from an estimation of Eq. (9) and a directly analogous multivariate regression in which c_t , d_t , and y_t are the regressors. The table reports the coefficient estimates at the top of each cell, heteroskedasticity and serial correlation robust t -statistics in parentheses, and adjusted R^2 statistics in square brackets. Inference on b_1 , b_2 , and b_3 can be accomplished by re-writing Eq. (9) so that the hypotheses to be tested are written as a restrictions on $I(0)$ variables (Sims et al. (1990)). For example, the hypothesis $b_1 = 0$ can be tested by rewriting Eq. (9) as

$$\begin{aligned} z_{t+h} &= a + b_1 [c_t - \omega a_t - (1 - \omega) y_t] + [b_2 + b_1 \omega] a_t + [b_3 + b_1 (1 - \omega)] y_t + u_t \\ &= a + b_1 [cay_t] + [b_2 + b_1 \omega] a_t + [b_3 + b_1 (1 - \omega)] y_t + u_t. \end{aligned}$$

It follows that the ordinary least squares estimate of b_1 has a limiting distribution given by

$$\sqrt{T} (\hat{b}_1 - b_1) \longrightarrow N \left(0, \frac{\sigma_u^2}{T \sum_{t=1}^T (cay_t - \overline{cay})^2} \right),$$

where σ_u^2 denotes the variance of u_t , and \overline{cay} is the sample mean of cay_t . These may be evaluated by using the full sample estimates, \widehat{cay}_t . A similar rearrangement can be used to

test hypotheses about b_2 and b_3 . Note that the full sample estimates of the cointegration coefficients are only required for inference about the forecasting exercise—they do not affect the forecasting exercise itself.

The results in Table 9 are consistent with those obtained using \widehat{cay}_t and \widehat{cdy}_t as forecasting variables. In almost every case, the individual coefficients on each regressor are strongly statistically significant as predictive variables for excess returns and dividend growth, and the R^2 statistics indicate that the regressors jointly explain about the same fraction of variation in future returns and future dividend growth explained by the individual regressors \widehat{cay}_t and \widehat{cdy}_t . For example, the multivariate regression with c_t , a_t , and y_t explains 26% of one-year-ahead excess returns, whereas \widehat{cay}_t explains 25%. The multivariate regression with c_t , d_t , and y_t explains about 24% of the variation in one-year-ahead dividend growth, whereas \widehat{cdy}_t explains 21%. The adjusted R^2 statistics are slightly lower than in the regressions that include only \widehat{cdy}_t or \widehat{cay}_t because there are three regressors instead of one. These results do not support the conclusion that \widehat{cay}_t and \widehat{cdy}_t have forecasting power merely because they are estimated in a first stage, using data from the whole sample period.

4.2.2. Small sample robustness

There are at least two potential econometric hazards with interpreting the long-horizon regression results using \widehat{cdy}_t and \widehat{cay}_t , presented above. One is that the use of overlapping data in long-horizon regressions can skew statistical inference in finite samples. Valkanov (2003) shows that, in finite samples where the forecasting horizon is a nontrivial fraction of the sample size, the t -statistics of long-horizon regression coefficients do not converge to a well-defined distribution, and the finite-sample distributions of R^2 statistics in long-horizon regressions do not converge to their population values. A second possible econometric hazard with interpreting the long-horizon regression results presented in the previous section occurs because (like most long-horizon forecasting variables) \widehat{cdy}_t and \widehat{cay}_t are persistent variables, which, although predetermined, are not exogenous. This lack of exogeneity can create a small-sample bias in the regression coefficient that works in the direction of indicating predictability even where none is present (Nelson and Kim, 1993 and Stambaugh, 1999).

To address these potential inference problems, we perform three robustness checks. The first is to compute the rescaled t/\sqrt{T} statistic (where T is the sample size), recommended by Valkanov (2003). The second is to use vector autoregressions to impute long-horizon R^2 statistics, rather than estimating them directly from long-horizon regressions. The third is to perform both bootstrapped estimates of the empirical distribution of the predictive regression coefficients and adjusted R^2 statistics under the null of no predictability.

We begin by discussing the rescaled t/\sqrt{T} statistic, which is due to Valkanov (2003). Valkanov shows that, when there is nontrivial overlap in the residuals of long-horizon regressions, the t -statistic for whether the predictive variable is statistically different from zero diverges at rate $T^{1/2}$. Thus, Valkanov proposes testing for statistical significance by using a rescaled t/\sqrt{T} statistic, which has a well-defined limiting distribution. The distribution of this rescaled statistic is nonstandard, however, and depends on two nuisance parameters, δ and c . The parameter δ measures the covariance between innovations in the variable to be forecast, and innovations in some forecasting variable, call it X_t . The parameter c measures deviations from unity in the highest autoregressive root for X_t , in a decreasing (at rate T) neighborhood of 1. Both of these parameters can be consistently estimated using the methodology described in Valkanov (2003). With these estimates in hand, the rescaled t -statistic, t/\sqrt{T} , can be compared against critical values provided in Valkanov (2003).

The rescaled t -statistics for our application are reported in curly brackets in Table 4, for univariate predictive regressions of excess returns on \widehat{cay}_t and \widehat{cdy}_t , and in Table 5 for univariate predictive regressions of dividend growth on \widehat{cay}_t and \widehat{cdy}_t . The tables report both the statistic itself and whether its value implies that the predictive coefficient in each regression is statistically significant at the 5% 2.5% and 1% levels. According to this rescaled t -statistic, \widehat{cay}_t is a powerful forecaster of excess returns (statistically significant at the 1% level) at every horizon ranging from one to six years, as is \widehat{cdy}_t at all but the one-year horizon (Table 4). For future dividend growth (Table 5), the rescaled t -statistic implies that \widehat{cdy}_t is a statistically significant predictor at the 1% percent level at every horizon from one to six years, whereas \widehat{cay}_t is a statistically significant predictor of dividend growth at the 1% level at every horizon ranging from one to four years. According to these statistics, the forecasting power of \widehat{cay}_t and \widehat{cdy}_t for long-horizon excess stock market returns and stock market dividend growth is robust to accounting for biases arising from the use of overlapping data in finite samples.

Finite sample problems with overlapping data can also be avoided by using vector autoregression to impute implied long-horizon R^2 statistics, rather than estimating them directly from long-horizon returns. The methodology for measuring long-horizon statistics by estimating a VAR has been covered by Campbell (1991), Hodrick (1992), and Kandel and Stambaugh (1989), and we refer the reader to those articles for further details. We present the results from using this methodology in Table 10, which investigates the long-horizon predictive power of \widehat{cay}_t and \widehat{cdy}_t in two ways. First, we can use the estimated cointegrated VARs (or error-correction representations) reported in Table 3 to impute long-horizon R^2 statistics for dividend growth, Δd_t , and wealth growth, Δa_t . Following Cochrane 1994, Appendix D, this is done by writing the error-correction representation in AR(1) notation (see Cochrane (1994), Appendix D), and then applying the standard VAR methodology

for imputing long-horizon statistics to the resulting AR(1) representation. Note that wealth growth is highly correlated with the log return on the CRSP value-weighted index and therefore proxies for r_t (the correlation is over 0.89 in our sample). The cointegrated VAR for c_t , d_t , and y_t in Panel C of Table 3 can be used to impute long-horizon R^2 statistics for forecasting Δd_t with \widehat{cdy}_t , while the cointegrated VAR for c_t , a_t , and y_t in Panel B of Table 3 can be used to impute long-horizon R^2 statistics for forecasting Δa_t with \widehat{cay}_t . Second, we can use standard bivariate, first-order VARs for dividend growth Δd_t and either \widehat{cay}_t or \widehat{cdy}_t , and for stock returns r_t and either \widehat{cay}_t or \widehat{cdy}_t , to compute the implied long-horizon R^2 statistics. This latter approach is more common. In either approach, we calculate an implied R^2 statistic using the coefficient estimates of the VAR and the estimated covariance matrix of the VAR residuals.

Table 10 shows that the pattern of the implied R^2 statistics from the vector autoregressions is similar to those produced from the direct long-horizon regressions. Moreover, the results from using the cointegrated VARs (estimates are presented in Table 3) are very similar to those using the bivariate VARs. For example, the implied adjusted R^2 statistics for forecasting dividend growth with \widehat{cdy}_t in the bivariate VAR (row 4) peaks at 0.2 for a three-year horizon, whereas the comparable R^2 statistic using the cointegrated VAR is 0.17 over this same horizon (row 2). A similar pattern holds for the implied R^2 statistics for forecasting with \widehat{cay}_t : the implied R^2 statistic for forecasting excess returns with \widehat{cay}_t in the bivariate VAR is as high as 49% at a three-year horizon (row 5), whereas the implied R^2 statistic for forecasting Δa_t with \widehat{cay}_t in the cointegrated VAR is 36% (row 1). We conclude that the evidence favoring predictability of dividend growth and excess stock returns by \widehat{cdy}_t and \widehat{cay}_t is robust to the VAR methodology.

Our third method for addressing potential finite sample biases is to estimate the empirical distribution of regression coefficients and adjusted R^2 statistics from predictive regressions in which \widehat{cay}_t and \widehat{cdy}_t are used as forecasting variables. Table 11 presents results based on a bootstrap simulation conducted under the null hypothesis of no predictability (i.e., residuals for the dependent variable are generated by regressions on a constant). We use first-order autoregressive specifications as our reduced form models for \widehat{cay}_t and \widehat{cdy}_t . The standard bootstrap is not consistent if the data series have a near-unit root. However, \widehat{cay}_t and \widehat{cdy}_t do not appear well-characterized as near-unit root processes, since—unlike the log dividend-price ratio—standard cointegration tests strongly reject the hypothesis that they are $I(1)$ random variables. We generate artificial sequences of excess returns and dividend growth are generated by drawing randomly (with replacement) from the sample residuals, under the null of no predictability. The simulations are repeated 10,000 times. Nelson and Kim (1993) also perform randomization, which differs from bootstrapping only in that sampling is without replacement; our results are very similar using randomization. To avoid difficulties caused

by the use of overlapping data, we focus here on one-year-ahead regressions.

Table 11 summarizes the estimated sampling distribution for the slope coefficient and the R^2 statistic in univariate forecasting regressions of annual excess returns and annual dividend growth. The results of each simulation are nearly identical. In almost every case, the estimated predictability coefficient and R^2 statistic lie outside of the 95% confidence interval based on the empirical distribution under the null of no predictability. In most cases they lie outside of the 99% confidence interval. The one exception is for the case in which excess returns are regressed on the one-year lagged value of \widehat{cdy} ; in this case, we cannot reject the hypothesis that one-step-ahead forecasting power of \widehat{cdy}_t is not statistically distinguishable from zero. This is not surprising, since even the standard asymptotic statistics suggest that \widehat{cdy}_t only has significant predictive power for returns at horizons longer than one year. For all of the other regressions and forecasting horizons, we find that our estimated slope coefficients and R^2 statistics are large relative to their sampling distributions under the null of no predictability. A Monte Carlo simulation of 10,000 artificial time-series generated from the Normal distribution under the null of no predictability generated almost identical results. We conclude that the predictability findings documented here are robust to correcting for small sample biases in the regression coefficients and R^2 statistics.

4.3. Including Share Repurchases

So far we have focused on measuring dividends as the actual cash paid to shareholders of the CRSP value-weighted index. We do this in order to make our results directly comparable with the existing literature, which has focused on forecasting the growth rate in this particular measure of dividends. This measure is of interest because it represents the predominant form of payout to shareholders over much of the post-War period. Moreover, as noted by Campbell and Shiller (2001), traditional dividends are an appealing indicator of fundamental value for long-term shareholders, because the end-of-period share price becomes trivially small when discounted from the end to the beginning of a long holding period.

Nonetheless, there is a growing view that changing corporate finance policy has led many firms, in recent years, to compensate shareholders through repurchase programs rather than through dividends (Fama and French, 2001; Grullon and Michaely, 2002). Still, large firms with high earnings have continued to increase traditional dividend payouts over time (DeAngelo, DeAngelo and Skinner, 2002). The impact on aggregate dividends is therefore unclear. In this section we show that our main conclusions are not altered by adjusting dividends to account for share repurchase activity.

One way to adjust dividends for such shifts in corporate financial policy is to add dollars spent on repurchases to dividends. We do so here by adding aggregate gross share repurchase

expenditures for the Industrial Compustat firms reported in Grullon and Michaely (2002) to our measure of dividends. Gross repurchases are readily available from the published work of Grullon and Michaely (2002), and this procedure is conservative for our purpose because the sum of share repurchases net of new issues and traditional dividend payments would be that much closer to our original series. These data cover the period 1972 to 2000 and are added to the CRSP dividends after being scaled to match the units of our original dividend series. As Grullon and Michaely (2002) note, repurchase activity prior to 1972 represented a tiny fraction of shareholder compensation for U.S. corporations; thus, the traditional dividend series should provide an accurate measure of actual payouts in data prior to 1972.

Table 12 presents the results of univariate long-horizon forecasting regressions for the growth in dividends plus repurchase activity, using \widehat{cay}_t and \widehat{cdy}_t as forecasting variables in separate regressions.¹⁰ The results should be compared with those in Table 5, which presents the analogous findings using CRSP value-weighted dividends. Comparing the output from the two tables, it is immediately evident that the inclusion of share repurchases does not alter the main conclusions obtained from using traditional dividends: \widehat{cay}_t and \widehat{cdy}_t are both strong predictors of the long-horizon growth rates in this series, with t -statistics that begin above 4.0 for horizons at one year and increase, and R -squared statistics that are in line with those in Table 4. We conclude that adjusting dividends for repurchases does not alter the main finding in this paper, namely that the growth in compensation to shareholders is forecastable in post-War data, and over horizons previously associated exclusively with return forecastability.

5. Why might expected dividend growth covary with expected returns?

If investors themselves desire smooth consumption paths, why don't managers perfectly smooth dividend payments? Some authors have noted that dividends are smoother than earnings, consistent with the hypothesis that managers do some dividend smoothing (Cochrane, 1994; Lamont, 1998). One possibility is that although dividend-smoothing is possible over long horizons (as revealed by the slow-moving dividend-price ratio), it is more difficult over horizons corresponding to the business cycle. Several researchers have presented evidence that is suggestive. Gertler and Hubbard (1993) study firm-level data from Compustat and find that dividend payouts are lower during a slowdown in economic growth than they are during periods of economic expansion. Bernanke and Gertler (1989) and Bernanke et al. (1996) present theoretical and empirical evidence of countercyclical variation in the

¹⁰The cointegration relation cdy_t is reestimated using the new dividend series inclusive of repurchases.

external finance premium, suggesting that managers who need to finance long-term projects have a greater need to retain earnings in recessions than in expansions. The equity risk premium also appears countercyclical: it rises during an economic slowdown and falls during periods of economic growth (Fama and French, 1989; Ferson and Harvey, 1991; Lettau and Ludvigson, 2001). Taken together, these findings suggest that high risk premia occur in periods of economic recession and coincide with a temporarily low stock price, temporarily low earnings, *and* temporarily low dividends. In such a world, consumers might be better able to smooth consumption than managers are able to smooth dividends. If this is true, earnings growth should be predictably higher when, according to \widehat{cdy}_t , dividend growth and excess stock returns are predictably higher.

Table 13 presents some evidence that is supportive of this hypothesis using earnings data for NYSE firms. The earnings data are from Lewellen (2003) and are the ratio of operating earnings before depreciation to market value. Unfortunately, the sample is limited by the availability of Compustat data to the period 1964-2000. We use Lewellen's data because earnings per share is contaminated by variability in share issuance. Table 13 reports that earnings growth is predictably higher when dividend growth is predictably higher. The regressor \widehat{cdy}_t is strongly statistically significant as a predictor of earnings growth at business cycle frequencies, with t -statistics in excess of 4.0 for one-to-three-year forecasting horizons, and in excess of 3.0 for a four-year horizon. The univariate forecasting regression explains about 17% of the variation in earnings growth four years hence. Thus, when consumption is high relative to its common trend with d_t and y_t , both dividends and earnings are temporarily low and are forecast to grow more quickly in the future. These results are consistent with the hypothesis that manager's dividend-smoothing ability is imperfect over business cycles.

6. Conclusion

This paper presents evidence that changing forecasts of dividend growth are an important feature of the post-War U.S. stock market, despite the failure of the dividend-price ratio to uncover such variation. Although these findings contradict the common conclusion that expected dividend growth is roughly constant, they reinforce the textbook conclusion that expected returns are time-varying and make an important contribution to aggregate stock market fluctuations. Dividend forecasts covary with changing forecasts of excess stock returns, and are positively correlated with business cycle variation in expected returns. Such fluctuations in expected returns and expected dividend growth have offsetting effects on the dividend-price ratio. The findings provide at least a partial explanation for why the consumption-wealth ratio has been found superior to the log dividend-price ratio as a predictor of excess stock market returns over medium-term horizons.

Our results suggest that an important component of time-varying expected returns and time-varying expected dividend growth is not captured by the log dividend-price ratio, or likely by other aggregate financial ratios. The results also imply that time-varying investment opportunities will be poorly captured by variation in the log dividend-price ratio, because it fails to reveal significant movements in the investment opportunity set that occur over business cycle horizons. This implication should be of special relevance to the growing body of literature on strategic asset allocation, in which the log dividend-price ratio is often used as a proxy for time-variation in the investment opportunity set, and as an input into the optimal asset allocation decision of a long-horizon investor. For a lucid summary of this literature, see Campbell and Viceira (2001).

The results in this paper hold for post-War data. There is little evidence of forecastability in dividend growth by \widehat{cdy}_t once pre-War (available from 1929) data are included in the sample. A likely explanation is the relatively poor quality of pre-War consumption and labor income data. For example, data-collection methodologies changed discretely at the beginning of the post-War period, and Romer (1989) finds that pre-War GDP estimates significantly exaggerate the size of cyclical fluctuations in the pre-War era. The upshot for our empirical application is that there is no evidence of a cointegration relation among c_t , d_t , and y_t in samples that include pre-War data. Since the forecasting power of \widehat{cdy}_t is predicated on cointegration, it is not difficult to understand why we find little predictive power in such samples.

While our findings help to explain some puzzling features of post-War U.S. stock market behavior, we caution that the results provide but one piece of a larger puzzle concerning the behavior of the dividend-price ratio, especially in more recent years. For example, the aggregate price-dividend ratio soared to unprecedented levels in the 1990s and even today remains above historical norms. Lettau et al. (2003) provide a partial explanation for this surge based on evidence of an unprecedented decline in macroeconomic risk. There is also a growing view that shifts in corporate financial policy may have created persistent changes in dividend growth rates (e.g., Fama and French 2001). It is too soon to tell whether such shifts in corporate financial policy will be sustained. But trend-like shifts in price-earnings ratios and other valuation measures suggest that factors other than changes in corporate payout policy must have contributed to the behavior of aggregate financial ratios in recent data. Whatever the reason for these changes, the results presented here suggest that some of the differences at cyclical frequencies between the log dividend-price ratio and the log consumption-wealth ratio have been attributable historically to changing forecasts of long-horizon dividend growth.

References

- Ang, A., 2002. Characterizing the ability of dividend yields to predict future dividends in log-linear present value models, unpublished paper, Graduate School of Business, Columbia University.
- Ang, A., Bekaert, G., 2001. Stock return predictability: Is it there?, nBER Working Paper No. 8207.
- Bernanke, B. S., Gertler, M., March 1989. Agency costs, net worth and business cycle fluctuations. *American Economic Review* 79, 14–31.
- Bernanke, B. S., Gertler, M., Gilchrist, S., 1996. The financial accelerator and the flight to quality. *The Review of Economics and Statistics* 78 (1), 1–15.
- Blinder, A. S., Deaton, A., 1985. The time series consumption function revisited. *Brookings Papers on Economic Activity* 2, 465–511.
- Brennan, M. J., Xia, Y., 2001. Stock price volatility and the equity premium. *Journal of Monetary Economics* 47, 249–283.
- Campbell, J. Y., 1987. Does saving anticipate declining future labor income? an alternate test of the permanent income hypothesis. *Econometrica* 55, 1249–73.
- Campbell, J. Y., 1991. A variance decomposition for stock returns. *Economic Journal* 101, 157–179.
- Campbell, J. Y., 2003. Consumption-based asset pricing. In: Constantinides, G., Harris, M., Stulz, R. (Eds.), *Handbook of the Economics of Finance* forthcoming. North-Holland, Amsterdam.
- Campbell, J. Y., Lo, A. W., MacKinlay, C., 1997. *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.
- Campbell, J. Y., Mankiw, G., 1989. Consumption, income, and interest rates: Reinterpreting the time series evidence. In: Blanchard, O., Fischer, S. (Eds.), *NBER Macroeconomics Annual: 1989*. MIT Press, Cambridge, MA, pp. 185–216.
- Campbell, J. Y., Shiller, R. J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1, 195–227.
- Campbell, J. Y., Shiller, R. J., 2001. Valuation ratios and the long-run stock market outlook: An update, nBER working paper No. 8221.
- Campbell, J. Y., Viceira, L., 2001. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, London, UK.
- Campbell, J. Y., Yogo, M., 2002. Efficient tests of stock return predictability, unpublished paper, Harvard University.

- Cochrane, J. H., June 1991. Explaining the variance of price-dividend ratios. *Review of Financial Studies* 5 (2), 243–280.
- Cochrane, J. H., 1994. Permanent and transitory components of gdp and stock prices. *Quarterly Journal of Economics* 109, 241–265.
- Cochrane, J. H., 1997. Where is the market going? uncertain facts and novel theories. *Economic Perspectives*, Federal Reserve Bank of Chicago 11, 1–37.
- Cochrane, J. H., 2001. *Asset Pricing*. Princeton University Press, Princeton, NJ.
- Cochrane, J. H., Piazzesi, M., 2002. Bond risk premia, unpublished paper, University of Chicago Graduate School of Business.
- DeAngelo, H., DeAngelo, L., Skinner, D. J., 2002. Are dividends disappearing? dividend concentration and the consolidation of earnings, unpublished paper, University of Michigan.
- Fama, E. F., French, K. R., 1988. Dividend yields and expected stock returns. *Journal of Financial Economics* 22, 3–27.
- Fama, E. F., French, K. R., 1989. Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics* 25, 23–49.
- Fama, E. F., French, K. R., 2001. Disappearing dividends: Changing firm characteristics or lower propensity to pay. *Journal of Financial Economics* 60 (1), 3–43.
- Fama, E. F., French, K. R., April 2002. The equity premium. *Journal of Finance* 57 (2), 637–659.
- Person, W. E., Harvey, C. R., 1991. The variation of economic risk premiums. *Journal of Political Economy* 99, 385–415.
- Flavin, M., 1981. The adjustment of consumption to changing expectations about future income. *Journal of Political Economy* 89 (5), 974–1009.
- Gertler, M., Hubbard, R. G., 1993. Corporate financial policy, taxation, and macroeconomic risk. *Rand Journal of Economics* 24 (2), 286–303.
- Goyal, A., Welch, I., May 2003. Predicting the equity premium with dividend ratios. *Management Science* 49 (5), 639–654.
- Grullon, G., Michaely, R., August 2002. Dividends, share repurchases, and the substitution hypothesis. *Journal of Finance* 57 (4), 1649–1684.
- Heaton, J., Lucas, D., 1999. Stock prices and fundamentals. In: Blanchard, O. J., Fischer, S. (Eds.), *NBER Macroeconomics Annual: 1999*. MIT Press, Cambridge, MA.

- Hodrick, R., 1992. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies* 5, 357–386.
- Inoue, A., Kilian, L., 2002. In-sample or out-of-sample tests of predictability: Which one should we use?, unpublished paper, Department of Economics, University of Michigan.
- Kandel, S., Stambaugh, R. F., 1989. Modeling expected stock returns for long and short horizons, unpublished Paper, Rodney L. White Center, Wharton School, University of Pennsylvania.
- Kothari, S., Lewellen, J., Warner, J., 2004. Stock returns, aggregate earnings surprises and behavioral finance, unpublished Paper, Sloan School of Management, MIT.
- Lamont, O., 1998. Earnings and expected returns. *Journal of Finance* 53, 1563–87.
- LeRoy, S., Porter, R., 1981. The present value relation: Tests based on variance bounds. *Econometrica* 49, 555–557.
- Lettau, M., Ludvigson, S. C., June 2001a. Consumption, aggregate wealth and expected stock returns. *Journal of Finance* 56 (3), 815–849.
- Lettau, M., Ludvigson, S. C., 2001b. Measuring and modeling variation in the risk-return tradeoff, unpublished paper, New York University.
- Lettau, M., Ludvigson, S. C., 2002. Time-varying risk premia and the cost of capital: An alternative implication of the q theory of investment. *Journal of Monetary Economics* 49, 31–66.
- Lettau, M., Ludvigson, S. C., March 2004. Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption. *American Economic Review* 94 (1), 276–299.
- Lettau, M., Ludvigson, S. C., Wachter, J. A., 2003. The declining equity premium: What role does macroeconomic risk play?, unpublished paper, New York University.
- Lewellen, J. W., 2003. Predicting returns with financial ratios. *Journal of Financial Economics*, forthcoming.
- Longstaff, F. A., Piazzesi, M., 2004. Corporate earnings and the equity premium. *Journal of Financial Economics* forthcoming.
- Menzly, L., Santos, T., Veronesi, P., February 2004. Understanding predictability. *Journal of Political Economy* 112 (1), 1–47.
- Nelson, C. C., Kim, M. J., 1993. Predictable stock returns: The role of small sample bias. *Journal of Finance* 43, 641–661.
- Newey, W. K., West, K. D., 1987. A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.

- Romer, C. D., 1989. The pre-war business cycle reconsidered: New estimates of gross national product 1869-1908. *Journal of Political Economy* 97 (1), 1–37.
- Santos, J., Veronesi, P., 2000. Labor income and predictable stock returns, unpublished paper, University of Chicago.
- Shiller, R. J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.
- Sims, C. A., Stock, J. H., Watson, M. W., January 1990. Inference in linear time-series models with some unit roots. *Econometrica* 58 (1), 113–144.
- Stambaugh, R. F., 1999. Predictive regressions. *Journal of Financial Economics* 54, 375–421.
- Stock, J. H., Watson, M. W., 1993. A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica* 61, 783–820.
- Valkanov, R., 2003. Long-horizon regressions: Theoretical results and applications. *Journal of Financial Economics* 68, 201–232.

Table 1: Summary statistics

	Δc_t	Δy_t	Δd_t	Δp_t	Δa_t
Univariate summary statistics					
Mean (in %)	2.01	2.30	4.01	6.88	2.45
Standard Deviation (in %)	1.14	1.83	12.24	16.13	4.05
Correlation matrix					
Δc_t	1.00	0.78	-0.13	-0.00	0.32
Δy_t		1.00	-0.10	-0.10	0.18
Δd_t			1.00	0.64	0.52
Δp_t				1.00	0.83
Δa_t					1.00

Notes: This table reports summary statistics for annual growth of real per capita consumption Δc_t , labor income Δy_t , dividends on the CRSP value-weighted index Δd_t , the CRSP value-weighted price index Δp_t , and asset wealth Δa_t . The sample spans the period 1948-2001.

Table 2: Autocorrelations of ratios

Ratio	ρ_1	ρ_2	ρ_3	ρ_4
$d - p$	0.87	0.72	0.59	0.47
$c - 0.33 a - 0.57 y$	0.46	0.14	0.05	0.00
$c - 0.13 d - 0.68 y$	0.46	0.20	0.24	0.16

Notes: This table reports autocorrelations of ratios involving consumption c_t , labor income y_t , dividends on the CRSP value-weighted index d_t , the CRSP value-weighted price index p_t , and asset wealth a_t . ρ_i denotes the autocorrelation of order i (in years). The cointegrating coefficients in the last two rows are estimates using dynamic least squares with two leads and lags. The sample is annual and spans the period 1948-2001.

Table 3: Estimates from cointegrated VARs

Panel A: (c, a, y)			
Dependent variable	Equation		
	Δc_t	Δy_t	Δa_t
$\frac{\Delta c_{t-1}}{(t\text{-stat})}$	0.267 (1.279)	0.449 (1.220)	-0.523 (-0.696)
$\frac{\Delta y_{t-1}}{(t\text{-stat})}$	-0.039 (-0.294)	-0.148 (-0.641)	0.433 (0.916)
$\frac{\Delta a_{t-1}}{(t\text{-stat})}$	0.112 (2.777)	0.128 (1.794)	0.392 (2.702)
$\frac{\widehat{cay}_{t-1}}{(t\text{-stat})}$	-0.007 (-0.053)	0.102 (0.457)	1.726 (3.803)
\bar{R}^2	0.199	0.050	0.207

Panel B: (c, d, y)			
Dependent variable	Equation		
	Δc_t	Δy_t	Δd_t
$\frac{\Delta c_{t-1}}{(t\text{-stat})}$	0.469 (2.284)	0.652 (1.869)	-0.136 (-0.066)
$\frac{\Delta y_{t-1}}{(t\text{-stat})}$	-0.074 (-0.572)	-0.156 (-0.709)	-0.252 (-0.176)
$\frac{\Delta d_{t-1}}{(t\text{-stat})}$	0.029 (2.311)	0.052 (2.389)	-0.129 (-0.917)
$\frac{\widehat{cdy}_{t-1}}{(t\text{-stat})}$	-0.038 (-0.408)	0.219 (1.377)	2.400 (2.314)
\bar{R}^2	0.179	0.098	0.104

Notes: The table reports estimated coefficients from cointegrated first-order vector autoregressions of the column variable on the row variable; c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. t -statistics are reported in parentheses. Estimated coefficients that are significant at the 5% level are highlighted in boldface. The sample is annual and spans the period 1948-2001.

Table 4: Univariate long-horizon regressions – excess stock returns

h -period regression: $\sum_{i=1}^h (r_{t+i} - r_{f,t+i}) = k + \gamma z_t + \epsilon_{t,t+h}$						
Horizon h (in years)						
$z_t =$	1	2	3	4	5	6
$d_t - p_t$	0.14 (1.90) {0.26} [0.08]	0.24 (1.40) {0.19} [0.10]	0.27 (1.21) {0.16} [0.10]	0.34 (0.73) {0.10} [0.10]	0.52 (0.84) {0.11} [0.16]	0.73 (1.12) {0.15} [0.23]
\widehat{cay}_t	5.87 (4.15) {0.51***} [0.25]	10.50 (5.58) {0.75***} [0.44]	11.93 (6.86) {0.90***} [0.40]	12.54 (6.72) {0.78***} [0.32]	16.31 (7.18) {0.96***} [0.36]	21.66 (7.82) {0.81***} [0.50]
\widehat{cdy}_t	1.50 (1.45) {0.26} [0.01]	5.54 (7.54) {1.10***} [0.19]	6.36 (4.89) {0.59***} [0.23]	6.90 (5.07) {0.61***} [0.22]	8.30 (4.69) {0.74***} [0.22]	11.81 (4.96) {0.69***} [0.36]

Notes: This table reports the results of h -period regressions of returns on the CRSP value-weighted index in excess of a 3-month Treasury bill rate, $r_{r,t}$, on the variable listed in the first column: $\sum_{i=1}^h (r_{t+i} - r_{f,t+i}) = k + \gamma z_t + \epsilon_{t,t+h}$, where the z_t are the cointegration residuals listed in the first column, c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. \widehat{cay}_t and \widehat{cdy}_t are estimated cointegrating residuals for the systems $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$, respectively. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), the t/\sqrt{T} test suggested in Valkanov (2001) in curly brackets, and adjusted R^2 statistics in square brackets. Significant coefficients using the standard t -test at the 5% level are highlighted in boldface. Significance at the 5%, 2.5%, and 1% level of the t/\sqrt{T} test using Valkanov's (2001) critical values is indicated by *, **, and ***, respectively. The sample is annual and spans the period 1948-2001.

Table 5: Univariate Long-horizon Regressions – Dividend Growth

h -period regression: $d_{t+h} - d_t = k + \gamma z_t + \epsilon_{t,t+h}$						
$z_t =$	Horizon h (in years)					
	1	2	3	4	5	6
$d_t - p_t$	0.09 (2.94) {0.40} [0.07]	0.18 (2.11) {0.29} [0.15]	0.19 (2.70) {0.37} [0.13]	0.23 (2.27) {0.31} [0.14]	0.29 (2.70) {0.37} [0.15]	0.34 (2.41) {0.33} [0.19]
$d_t - p_t$ (1948-1999)	0.07 (1.57) {0.30} [0.02]	0.11 (1.63) {0.21} [0.03]	0.12 (1.60) {0.22} [0.03]	0.17 (1.98) {0.19} [0.06]	0.23 (2.05) {0.23} [0.08]	0.30 (2.44) {0.27} [0.14]
\widehat{cay}_t	4.50 (6.89) {1.00***} [0.31]	5.87 (7.15) {0.91***} [0.35]	5.45 (4.63) {0.56***} [0.22]	5.26 (3.88) {0.51***} [0.14]	6.30 (3.81) {0.40*} [0.16]	6.77 (2.31) {0.25} [0.18]
\widehat{cdy}_t	2.76 (4.74) {0.73***} [0.21]	3.97 (6.02) {0.91***} [0.26]	3.68 (5.18) {0.65***} [0.21]	4.02 (4.32) {0.54***} [0.21]	5.24 (5.69) {0.75***} [0.29]	6.10 (4.16) {0.50***} [0.38]

Notes: This table reports results from h -period regression of dividend growth on the CRSP value-weighted index: $d_{t+h} - d_t = k + \gamma z_t + \epsilon_{t,t+h}$, where the z_t are the cointegration residuals listed in the first column, c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. \widehat{cay}_t and \widehat{cdy}_t are estimated cointegrating residuals for the systems $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$, respectively. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), the t/\sqrt{T} test suggested in Valkanov (2001) in curly brackets, and adjusted R^2 statistics in square brackets. Significant coefficients using the standard t -test at the 5% level are highlighted in boldface. Significance at the 5%, 2.5%, and 1% level of the t/\sqrt{T} test using Valkanov's (2001) critical values is indicated by *, **, and ***, respectively. The sample is annual and spans the period 1948-2001 with the exception of the second panel, which use the 1948-1999 sample.

Table 6: Long horizon regressions – consumption and labor income growth

Variables	Horizon h (in years)					
	1	2	3	4	5	6
<i>h</i> -period regression: consumption growth						
\widehat{cay}_t	-0.21 (-2.42) [0.06]	-0.17 (-1.03) [0.00]	-0.26 (-0.98) [0.00]	-0.40 (-1.33) [0.01]	-0.54 (-1.53) [0.02]	-0.52 (-1.59) [0.01]
\widehat{cdy}_t	-0.15 (-2.16) [0.06]	-0.15 (-1.18) [0.01]	-0.16 (-1.09) [0.00]	-0.18 (-1.09) [-0.01]	-0.27 (-1.45) [0.01]	-0.33 (-1.76) [0.01]
multivariate <i>h</i> -period regression: consumption growth						
Δc_{t-1}	0.31 (2.47)	0.39 (1.62)	0.35 (1.19)	0.22 (0.52)	0.00 (-0.01)	-0.06 (-0.13)
\widehat{cay}_t	-0.11 (-1.27) [0.13]	-0.05 (-0.36) [0.03]	-0.13 (-0.55) [0.00]	-0.31 (-1.14) [0.00]	-0.54 (-1.66) [0.00]	-0.54 (-1.89) [-0.01]
Δc_{t-1}	0.32 (2.52)	0.37 (1.32)	0.36 (1.22)	0.31 (0.76)	0.13 (0.30)	0.04 (0.10)
\widehat{cdy}_t	-0.11 (-1.73) [0.15]	-0.09 (-0.76) [0.04]	-0.10 (-0.75) [0.01]	-0.13 (-0.95) [-0.01]	-0.24 (-1.65) [-0.01]	-0.32 (-1.98) [-0.01]
<i>h</i> -period regression: labor income growth						
\widehat{cay}_t	-0.15 (-1.00) [0.00]	-0.05 (-0.21) [-0.02]	-0.20 (-0.49) [-0.01]	-0.33 (-0.59) [-0.01]	-0.57 (-0.91) [0.00]	-0.78 (-1.28) [0.01]
\widehat{cdy}_t	-0.02 (-0.15) [-0.02]	0.12 (0.58) [-0.01]	0.09 (0.33) [-0.02]	0.16 (0.46) [-0.02]	0.12 (0.33) [-0.02]	-0.09 (-0.22) [-0.02]

Notes: This table reports results from univariate *h*-period regression of consumption and labor income growth: c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. \widehat{cay}_t and \widehat{cdy}_t are estimated cointegrating residuals for the systems $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$, respectively. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected *t*-statistics (in parentheses), the t/\sqrt{T} test suggested in Valkanov (2001) in curly brackets, and adjusted R^2 statistics in square brackets. Significant coefficients using the standard *t*-test at the 5% level are highlighted in boldface. The sample is annual and spans the period 1948-2001.

Table 7: Multivariate long-horizon regressions

Variables	Horizon h (in years)					
	1	2	3	4	5	6
<i>h</i> -period regression: excess stock returns						
$d_t - p_t$	0.02 (0.30)	0.02 (0.14)	0.09 (0.59)	0.18 (0.88)	0.35 (1.36)	0.46 (1.86)
\widehat{cay}_t	5.59 (3.10)	9.01 (3.59)	8.96 (2.78)	8.61 (1.89)	12.76 (2.57)	15.68 (3.27)
\widehat{cdy}_t	0.03 (0.03) [0.22]	2.14 (2.08) [0.44]	2.58 (2.33) [0.41]	2.72 (1.95) [0.35]	1.60 (1.09) [0.42]	3.30 (2.54) [0.60]
\widehat{cay}_t	5.90 (3.50)	9.23 (4.29)	9.88 (4.18)	9.65 (2.84)	13.61 (3.66)	16.68 (4.18)
\widehat{cdy}_t	-0.07 (-0.10) [0.24]	2.10 (2.01) [0.45]	2.50 (2.03) [0.41]	2.95 (1.73) [0.33]	2.60 (1.31) [0.36]	4.81 (2.50) [0.53]
<i>h</i> -period regression: dividend growth						
$d_t - p_t$	0.03 (0.67)	0.08 (1.33)	0.12 (1.87)	0.17 (2.19)	0.19 (2.28)	0.21 (2.13)
\widehat{cay}_t	3.23 (2.94)	3.49 (2.43)	2.39 (1.47)	1.26 (0.69)	1.07 (0.52)	0.37 (0.12)
\widehat{cdy}_t	1.91 (3.34) [0.37]	2.46 (2.62) [0.43]	2.38 (2.46) [0.30]	2.91 (2.17) [0.28]	4.06 (3.44) [0.34]	5.06 (2.80) [0.43]
\widehat{cay}_t	3.63 (3.88)	4.47 (2.83)	3.58 (1.67)	2.21 (1.07)	1.53 (0.75)	0.83 (0.25)
\widehat{cdy}_t	1.79 (3.50) [0.38]	2.30 (2.37) [0.41]	2.28 (1.90) [0.26]	3.11 (2.47) [0.21]	4.60 (4.34) [0.28]	5.75 (3.11) [0.37]

Notes: This tables reports results from h -period regression of returns on the CRSP value-weighted index in excess of a three-month Treasury bill rate (top panel) and dividend growth (bottom panel). c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. \widehat{cay}_t and \widehat{cdy}_t are estimated cointegrating residuals for the systems $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$, respectively. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), and adjusted R^2 statistics in square brackets. Significant coefficients at the 5% level are highlighted in boldface. The sample is annual and spans the period 1948-2001.

Table 8: Other forecasting variables – dividend growth, univariate regressions

Variables	Horizon h (in years)					
	1	2	3	4	5	6
Δd_{t-1}	-0.26 (-2.34) [0.05]	-0.24 (-1.87) [0.02]	-0.32 (-2.30) [0.03]	-0.17 (-1.48) [-0.01]	-0.42 (-1.92) [0.04]	-0.34 (-1.55) [0.02]
TRM_{t-1}	-0.66 (-0.59) [-0.02]	0.03 (0.02) [-0.02]	2.13 (0.97) [-0.01]	3.49 (1.35) [0.01]	4.10 (1.30) [0.02]	2.85 (0.98) [0.00]
DEF_{t-1}	-0.02 (-0.53) [-0.02]	-0.01 (-0.19) [-0.02]	0.02 (0.46) [-0.02]	0.06 (1.09) [0.00]	0.07 (1.13) [0.00]	0.11 (1.97) [0.04]
$RREL_{t-1}$	-0.85 (-0.97) [-0.02]	0.77 (0.71) [-0.02]	0.79 (0.48) [-0.02]	-1.63 (-1.97) [-0.02]	-0.39 (-0.26) [-0.02]	-2.74 (-1.46) [-0.01]
CP_{t-1}	-0.01 (-5.91) [0.03]	-0.01 (-1.13) [-0.01]	0.00 (-0.46) [-0.03]	0.00 (0.53) [-0.03]	0.00 (0.38) [-0.03]	0.02 (1.19) [0.01]

Notes: This table reports results from h -period regression of dividend growth. TRM_t is the term spread, or the difference between the ten-year Treasury bond yield and the three-month Treasury bond yield; DEF_t is the BAA corporate Bond rate minus the AAA corporate bond rate; $RREL_t$ is the relative bill rate and CP is the Cochrane-Piazzesi (2003) bond factor. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), and adjusted R^2 statistics in square brackets. Significant coefficients at the 5% level are highlighted in boldface. The sample is annual and spans the period 1948-2001 for the regression including lagged dividend growth, 1953-2001 for the regressions including TRM, DEF, and RREL, and 1965-2001 for the regressions including CP.

Table 9: Multivariate long-horizon regressions

Variables	Horizon h (in years)					
	1	2	3	4	5	6
<i>h</i> -period regression: excess stock returns on c , a , and y						
c_t	6.50 (4.76)	11.96 (6.18)	13.48 (6.37)	13.92 (6.48)	17.34 (6.91)	22.53 (6.92)
a_t	-1.83 (-4.57)	-3.29 (-5.82)	-3.60 (-5.83)	-3.67 (-5.85)	-5.06 (-6.90)	-6.76 (-7.11)
y_t	-4.01 (-4.91) [0.26]	-7.43 (-6.43) [0.49]	-8.50 (-6.73) [0.48]	-8.85 (-6.89) [0.40]	-10.56 (-7.05) [0.41]	-13.52 (-6.96) [0.55]
<i>h</i> -period regression: excess stock returns on c , d , and y						
c_t	1.53 (1.59)	5.44 (6.56)	7.12 (4.97)	8.38 (5.52)	10.36 (5.69)	14.70 (5.62)
d_t	-0.07 (-0.60)	-0.54 (-5.10)	-0.55 (-3.00)	-0.58 (-3.03)	-0.74 (-3.20)	-1.16 (-3.49)
y_t	-1.30 (-1.98) [0.00]	-4.03 (-7.09) [0.15]	-5.48 (-5.58) [0.23]	-6.50 (-6.26) [0.25]	-7.96 (-6.39) [0.24]	-11.05 (-6.18) [0.37]
<i>h</i> -period regression: dividend growth on c , a , and y						
c_t	4.31 (6.26)	5.04 (6.60)	4.52 (4.64)	4.70 (4.16)	6.18 (4.16)	6.61 (3.34)
a_t	-1.44 (-7.16)	-1.91 (-8.56)	-2.05 (-7.22)	-2.49 (-7.56)	-3.46 (-7.99)	-3.70 (-6.41)
y_t	-2.44 (-5.93) [0.29]	-2.67 (-5.86) [0.37]	-2.10 (-3.60) [0.29]	-1.88 (-2.78) [0.29]	-12.31 (-2.61) [0.38]	-22.50 (-2.12) [0.40]
<i>h</i> -period regression: dividend growth on c , d , and y						
c_t	2.14 (3.92)	2.88 (4.77)	2.41 (3.59)	2.39 (2.74)	3.28 (3.86)	4.08 (2.99)
d_t	-0.45 (-6.49)	-0.64 (-8.28)	-0.62 (-7.23)	-0.69 (-6.20)	-0.88 (-8.15)	-0.98 (-5.66)
y_t	-1.22 (-3.28) [0.24]	-1.60 (-3.86) [0.33]	-1.21 (-2.64) [0.29]	-1.09 (-1.82) [0.31]	-1.57 (-2.70) [0.40]	-2.13 (-2.28) [0.48]

Notes: See next page.

Notes: This table reports results from h -period regression of returns on the CRSP value-weighted index in excess of a three-month Treasury bill rate, and dividend growth. c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), and d_t is log stock market dividends. For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), and adjusted R^2 statistics (in square brackets). Significant coefficients at the 5% level are highlighted in boldface. The distribution of the coefficient estimates is as follows. Consider a regression $z_t = \mu + \beta_1 x_{1t} + \beta_2 x_{2t} + \beta_3 x_{3t} + \epsilon_t$, where x_1, x_2 , and x_3 are cointegrated and the cointegrating vector is $(1, -\theta_2, -\theta_3)$. Let $\eta_t = x_{1t} - \theta_2 x_{2t} - \theta_3 x_{3t}$. Then the OLS estimate of β_1 has a limiting distribution of $\sqrt{T}(\hat{\beta}_1 - \beta_1) \rightarrow N(0, \sigma^2(1/T \sum_{t=1}^T (\eta_t - \bar{\eta}_t)^2)^{-1})$ where $\bar{\eta}_t$ is the mean of η_t and σ^2 is the variance of ϵ . Note that η_t depends on the cointegrating vector. The standard error is Newey-West corrected. The sample is annual and spans the period 1948-2001.

Table 10: Implied long-horizon R^2 from VARs

Variable	Implied R^2 for forecast horizon H					
	1	2	3	4	5	6
			VECM (c, a, y)			
Δa_t	0.31	0.39	0.40	0.36	0.32	0.28
			VECM (c, d, y)			
Δd_t	0.17	0.17	0.16	0.18	0.18	0.19
			VAR ($\Delta d, \widehat{cay}_t$)			
Δd_t	0.37	0.36	0.36	0.39	0.47	0.48
			VAR ($\Delta d, \widehat{cdy}_t$)			
Δd_t	0.19	0.17	0.18	0.21	0.29	0.34
			VAR (r, \widehat{cay}_t)			
r_t	0.25	0.32	0.33	0.49	0.52	0.62
			VAR (r, \widehat{cdy}_t)			
r_t	0.03	0.22	0.25	0.33	0.40	0.38

Note: The table reports implied (unadjusted) H -period R^2 statistics estimated from vector error correction models (VECM) and vector autoregressions (VAR). The sample is annual and spans the period 1948-2001.

Table 11: Small sample inference of slope and R^2

x_t	$\hat{\beta}$	95% CI	99% CI	R^2	95% CI	99% CI
excess returns						
\widehat{cay}_t	5.87	(-3.12, 3.10)	(-4.56, 4.48)	0.25	(-0.02, 0.06)	(-0.02, 0.11)
\widehat{cdy}_t	1.50	(-2.45, 2.48)	(-3.53, 3.66)	0.01	(-0.02, 0.05)	(-0.02, 0.10)
dividend growth						
\widehat{cay}_t	4.50	(-1.79, 1.79)	(-2.62, 2.59)	0.31	(-0.02, 0.06)	(-0.02, 0.11)
\widehat{cdy}_t	2.76	(-1.54, 1.56)	(-2.29, 2.21)	0.21	(-0.02, 0.06)	(-0.02, 0.11)

Notes: This table reports confidence intervals from a bootstrap procedure. Ten thousand artificial time series of the size of our data set are generated under the null hypothesis of no predictability. The data-generating process is $z_{t+1} = \alpha + e_{t+1}$, $x_{t+1} = \mu + \phi x_{t-1} + v_{t+1}$ where z_t is either excess returns or dividend growth and x_t is either *cay* or *cdy*. The parameters in the data-generating process are set to sample estimates for both the bootstrap and the Monte Carlo. We then run OLS regressions $z_{t+1} = \alpha + \beta x_t + u_{t+1}$ and study the distributions of $\hat{\beta}$ and the R^2 . We draw (with replacement) from the residuals of the system estimated under the null hypothesis. The columns denoted $\hat{\beta}$ and R^2 report our empirical estimates using annual data from 1948 to 2001.

Table 12: Univariate long-horizon regressions – including share repurchases

h -period regression: $d_{t+h} - d_t = k + \gamma z_t + \epsilon_{t,t+h}$						
$z_t =$	Horizon h (in years)					
	1	2	3	4	5	6
$d_t - p_t$	0.09 (1.76) [0.01]	0.10 (1.12) [0.01]	0.10 (0.81) [0.00]	0.11 (0.70) [0.00]	0.15 (0.81) [0.01]	0.19 (0.88) [0.02]
\widehat{cay}_t	4.41 (4.52) [0.24]	6.31 (5.02) [0.26]	6.85 (4.21) [0.21]	7.05 (3.64) [0.18]	8.67 (5.07) [0.21]	9.93 (4.68) [0.26]
\widehat{cdy}_t	4.25 (5.77) [0.20]	5.05 (5.23) [0.19]	4.58 (2.99) [0.12]	4.78 (2.37) [0.11]	6.49 (3.28) [0.16]	8.27 (4.04) [0.26]

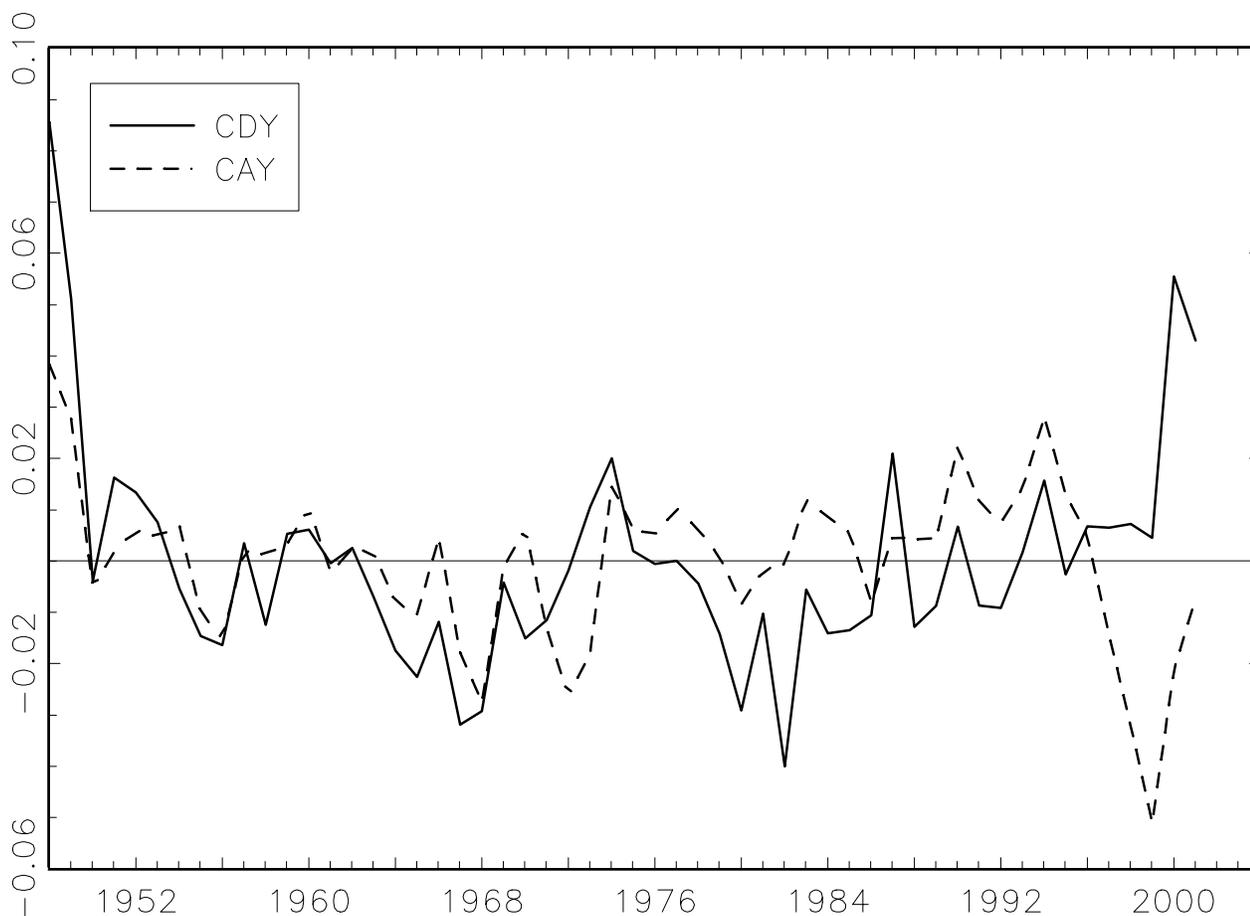
Notes: This table reports results from h -period regression of dividend growth on the CRSP value-weighted index: $d_{t+h} - d_t = k + \gamma z_t + \epsilon_{t,t+h}$, where dividends d are adjusted to include share repurchases using the estimates in Grullon and Michaely (2002). For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), and adjusted R^2 statistics in square brackets. Significant coefficients using the standard t -test at the 5% level are highlighted in boldface. The sample is annual and spans the period 1948-2000, since the repurchases data from Grullon and Michaely are only available through 2000.

Table 13: Long-horizon regression – earnings growth

Variables	Horizon h (in years)					
	1	2	3	4	5	6
\widehat{cdy}_t	2.35 (3.79) [0.09]	3.80 (3.01) [0.08]	5.11 (3.28) [0.09]	7.04 (3.54) [0.17]	6.98 (2.56) [0.15]	7.22 (2.21) [0.15]

Notes: This table reports results from h -period regression of earnings growth: $e_{t+h} - e_t = k + \beta \widehat{cdy}_t + \epsilon_{t,t+h}$. The earnings data are from Lewellen (2001). For each regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t -statistics (in parentheses), and adjusted R^2 statistics in square brackets. Significant coefficients at the 5% level are highlighted in bold face. The sample is annual and spans the period 1964-2000.

Figure 1: CDY and CAY



Notes: \widehat{cay}_t and \widehat{cdy}_t are estimated cointegrating residuals for the systems $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$, respectively. c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. The sample is annual and spans the period 1948-2001.