NBER WORKING PAPER SERIES

EXPECTED RETURNS AND EXPECTED DIVIDEND GROWTH

Martin Lettau
Sydney C. Ludvigson

Working Paper **9605**
http://www.nber.org/papers/w9605

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2003

Lettau acknowledges financial support from the National Science Foundation. Ludvigson acknowledges financial support from the Alfred P.Sloan Foundation and from the National Science Foundation. We thank Jushan Bai, John Y. Campbell, Kenneth French, Mark Gertler, Rick Green, Anthony Lynch, Lucrezia Reichlin, Peter Schotman, Charles Steindel, 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 Barczi for excellent research assistance. Any errors or omissions are the responsibility of the authors. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Expected Returns and Expected Dividend Growth
Martin Lettau and Sydney C. Ludvigson
NBER Working Paper No. 9605
April 2003
JEL No. G12, G10

ABSTRACT

We investigate a consumption-based present value relation that is a function of future dividend growth. Using data on aggregate consumption and measures of the dividend payments from aggregate wealth, we show that changing forecasts of dividend growth make an important contribution to fluctuations in the U.S. stock market, despite the failure of the dividend-price ratio to uncover such variation. In addition, these dividend forecasts are found to covary with changing forecasts of excess stock returns. The variation in expected dividend growth we uncover is positively correlated with changing forecasts of excess returns and occurs at business cycle frequencies, those ranging from one to six years. Because positively correlated fluctuations in expected dividend growth and expected returns have offsetting affects on the log dividend-price ratio, the results imply that both the market risk-premium and expected dividend growth vary considerably more than what can be revealed using the log dividend-price ratio alone as a predictive variable.

Martin Lettau
Department of Finance
Stern School of Business
New York University
44 West Fourth Street
New York, NY 10012-1126
and NBER
mlettau@stern.nyu.edu

Sydney C. Ludvigson
Department of Economics
New York University
269 Mercer Street, 7th Floor
New York, NY 10001
and NBER
sydney.ludvigson@nyu.edu

1 Introduction

One does not have 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; Lewellen (2001); Campbell (2002); Fama and French (2002)). These researchers maintain that, despite the market's unusual behavior recently, changing forecasts of excess returns make important contributions to fluctuations in the aggregate stock market.

[1] A large literature documents the poor predictability of dividend growth by the dividend-price ratio over long horizons, for example, Campbell (1991); Cochrane (1991); Cochrane (1994); Cochrane (1997); Cochrane (2001); Campbell (2002). Ang and Bekaert (2001) find somewhat stronger predictability; we discuss their results further below.

[2] See Fama and French (1988), Campbell and Shiller (1988); Hodrick (1992); Campbell, Lo, and MacKinlay (1997); Cochrane (1997); Cochrane (2001), Ch. 20; Campbell (2002).

[3] See Campbell (1991); Cochrane (1991); Hodrick (1992); Campbell, Lo, and MacKinlay (1997), Ch. 7; Cochrane (2001), Ch. 20; Campbell (2002).

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 less than a few years, have argued that careful statistical analysis provides little evidence that the log dividend-price ratio forecasts returns (for example, Nelson and Kim (1993); Stambaugh (1999); Ang and Bekaert (2001); Valkanov (2001)). 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, other researchers have found that excess returns on the aggregate stock market are strongly forecastable at horizons far shorter than those over which the persistent dividend-price ratio predominantly varies. 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 latter.

Taken together, these empirical findings raise a series of puzzling questions. Do the intermediate horizon statistical analyses using the dividend-price ratio imply 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 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 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. Second, these dividend forecasts are found to positively covary with changing forecasts of excess stock returns.

These 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 US 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, they can explain why business cycle variation in expected excess returns is captured by cay_t, but 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 affects 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 significant predictive power for future returns as well as future dividend growth, consistent with evidence reported in Nelson and Kim (1993), Stambaugh (1999), Ang and Bekaert (2001) and Valkanov (2001)). Such a failure is not an indication that expected returns are constant, however. On the contrary, 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. These findings 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 this 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 paid from the CRSP value-weighted stock market index, 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 US stock market. Section 5 discusses one possible explanation for why expected dividend growth might vary over time, and be positively correlated with expected returns, despite the fact that firms may have an incentive to smooth dividend payments if shareholders desire smooth consumption paths. Section 6 concludes.

2 A Consumption-Based Present Value Relation for Dividend Growth

This section presents a consumption-based present value relation for future 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; $R_{w,t+1}$ is the net return on aggregate wealth. For expositional convenience, we consider a simple accumulation equation for aggregate wealth, written

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

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 (1), solving the resulting difference equation for log wealth forward, and imposing a transversality condition.[5] The resulting expression holds to a first-order approximation:[6]

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

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.

Equation(2) is of little use in empirical work because aggregate wealth includes human capital, which is not observable. Lettau and Ludvigson (2001a) address this problem by reformulating the bivariate cointegrating relation between c_t and w_t as a trivariate cointegrating 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 the log consumption-aggregate wealth ratio is a consumption-based present value relation

[4]None of the derivations below are dependent on this assumption. In particular, equation (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, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^{j} (1 + R_{a,t+i})^{-i} Y_{t+j}$.

[5]This transversality condition rules out rational bubbles.

[6]We omit unimportant linearization constants in the equations throughout the paper.

involving future returns to asset wealth[7]

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

where ω is the average share of asset wealth, A_t, in aggregate wealth, W_t, $r_{a,t}$ 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, (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 asset wealth, consumption growth, labor income growth, or some combination of all three. The wealth shares ω and $(1-\omega)$ are cointegrating coefficients. We discuss their estimation further below. The cointegrating residual on the left-hand-side of (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 (2) and (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, (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 and the logs of dividends from asset wealth, d_t, and human wealth, y_t, which takes the form

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

Equation (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, equation (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 (4)) provide a rational forecast of dividend growth, labor income growth, consumption growth, or some combination of all three. The income shares ν and $(1-\nu)$ are cointegrating coefficients. We discuss their estimation further below. The cointegrating residual on the left-hand-side of (4) is denoted cdy_t, for short.

[7]We 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.

It is instructive to compare equation (4) with the proxy for the consumption-aggregate wealth ratio, (3), studied in Lettau and Ludvigson (2001a). Equation (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, equation (4) says that if investors want to maintain flat consumption paths, fluctuations in cdy_t summarize expectations of the growth in future dividends to aggregate wealth. 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 may forecast stock market dividend growth, if in fact expected dividend growth varies over time. In this case, (4) says that cdy_t is a state variable that summarizes changing forecasts of dividend growth to asset wealth.

It is also instructive to compare (4) and (3) with the linearized expression for the log dividend-price ratio. Following Campbell and Shiller (1988) the log dividend-price ratio may be written (up to a first-order approximation) as

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

where p_t be the log price of stock market wealth, which pays the dividend, d_t, $\rho \equiv \frac{1}{1+\exp(\overline{d-p})}$, and r_t is the log return to stock market wealth.[8] This equation says that if the log dividend-price ratio is high, agents must be expecting high future returns on stock market wealth, or low dividend growth rates. Many studies, cited in the introduction, have documented that $d_t - p_t$ explains little of the variability in future dividend growth; as a consequence, expected dividend growth is often modelled as constant.

Equation (5) can be simplified if we assume that expected stock returns follow a first-order autoregressive process, $E_t r_{t+1} \equiv x_t = \phi x_{t-1} + \xi_t$. With this specification for expected stock returns, and if expected dividend growth is constant, the log dividend-price ratio takes the form

$$d_t - p_t = E_t \sum_{i=0}^{\infty} \rho^i (r_{t+1+i} - \Delta d_{t+1+i}) = \frac{x_t}{1 - \rho\phi}. \tag{6}$$

When expected dividend growth is constant, the log dividend-price ratio does not forecast dividend growth at any horizon but instead forecasts long-horizon stock returns, because it

[8]Like those above for cay_t and cdy_t, this expression ignores inconsequential linearization constants.

captures time-varying expected returns, x_t. Equation (6) shows that, under the standard view that expected dividend growth is approximately constant, any and all variation in expected returns must be revealed by variation in the dividend-price ratio.

It is useful to consider the behavior of the log dividend-price ratio in a simple example for which expected dividend growth is not constant. Suppose that expected dividend growth varies according to a first-order autoregressive process,

$$E_t \Delta d_{t+1} \equiv g_t = \psi g_{t-1} + \zeta_t. \tag{7}$$

As is evident from (5), the effect of such variation on the log dividend-price ratio depends on the correlation between expected dividend growth and expected returns. For example, if the two are positively correlated, expected returns may be modeled as having two components, one component common to expected dividend growth, and another component independent of expected dividend growth. In this case we may write $E_t r_{t+1} = \beta g_t + x_t$, where $\beta > 0$ is the loading on expected dividend growth that generates a positive correlation between $E_t r_{t+1}$ and $E_t \Delta d_{t+1}$, and x_t is a component of expected returns that is independent of expected dividend growth.[9] Note that when $\beta = 1$, all of the variation in expected dividend growth is common to variation in expected returns.

Combining $E_t r_{t+1} = \beta g_t + x_t$ with (5), the log dividend-price ratio becomes

$$d_t - p_t = E_t \sum_{j=0}^{\infty} \rho (r_{t+1+i} - \Delta d_{t+1+i}) \tag{8}$$

$$= \frac{1}{1-\rho\phi} x_t - \frac{1-\beta}{1-\rho\psi} g_t. \tag{9}$$

Equation 9 shows that, when β is greater than zero, the relationship between $d_t - p_t$ and both expected dividend growth and expected returns will be obfuscated. When all of the variation in expected dividend growth is common to variation in expected returns, $\beta = 1$ and the expression is precisely the same as (6) for the case in which expected dividend growth is constant. In this instance, the log dividend-price ratio will have no forecasting power for future dividend growth even though, by construction, expected dividend growth varies over time. This is because positively correlated fluctuations in expected dividend growth and expected returns have offsetting affects on the log dividend-price ratio. The log dividend-price ratio will also have no forecasting power for one component of expected returns, namely g_t, because that component is completely offset by variation in expected dividend growth. When $0 < \beta < 1$, $d_t - p_t$ will still have difficulty revealing changing forecasts of stock market dividend growth, because it only captures a portion, $(1 - \beta)$, of time-variation n expected dividend growth; the remaining portion is not revealed because it is common to time-varying expected returns. It will also only capture a portion, x_t, of

[9] The loading on x_t is normalized to unity. This normalization is without loss of generality, since the specification above can always be redefined as $E_t r_{t+1} = \beta g_t + \gamma \tilde{x}_t$ as $E_t r_{t+1} = \beta g_t + x_t$ where $x_t = \gamma \tilde{x}_t$.

time-varying expected returns, because the remaining portion, βg_t, is more than offset by variation in expected dividend growth, $-g_t$. Notice that these problems do not affect the two consumption-based ratios discussed above, because they are not simultaneous functions of expected returns and expected dividend growth. These considerations motivate the use of the consumption-based ratios developed above to uncover possible time-variation in expected returns and expected dividend growth.

The framework developed above, with its approximate consumption identities, serves merely 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. The empirical investigation itself, discussed in the next section, is not dependent on these approximations. Nevertheless, we may assess the implications of framework presented above by investigating whether such present-value relations are informative 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, we use nondurables and services expenditure as a measure of aggregate consumption. This measure is a subset of total consumption, which is unobservable because we don't have a measure of the service flow from the stock of durable goods. Note that it would be inappropriate to use total personal consumption expenditures as a measure of consumption in this framework. This series includes expenditures on durable goods, which represent replacements and additions to the capital stock (investment), rather than the service flow from the existing stock. Durables expenditures are properly accounted for as part of nonhuman wealth, A_t, a component of aggregate wealth, W_t.[10]

[10]Treating durables purchases purely as an expenditure removing them from A_t and including them in C_t is also improper because doing so 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 consist primarily of depreciation.) What should be used in this budget constraint for C_t is not total expenditures but total consumption, of which the service flow from the durables stock is one part. But the service flow is unobservable, and is not the same as the investment expenditures on consumption goods. An assumption of some sort is necessary, and we follow Lettau and Ludvigson (2001a) by assuming that the log of unobservable real total consumption, c_t^T, is a multiple, $\lambda > 1$ of the log of real nondurables and services expenditure, c_t, plus a stationary random component, ϵ_t. Under this assumption, the observable

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, and creates 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.[11,12] In addition, we do not directly observe dividend payments from some forms of nonhuman, nonfinancial household net worth, primarily physical capital.[13]

Fortunately, it is not necessary to include every dividend component from aggregate wealth in the expression (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 (as models with balanced growth would suggest), the framework above implies that the included dividend measures may be combined with consumption to obtain a valid cointegrating relation. In this study, we use stock market dividends as a measure of dividend payments from nonhuman (asset) wealth, and use d_t to denote stock market dividends from now on. If nonstock/nonlabor forms of dividend income are cointegrated with the dividend payments

---

log of real nondurables and services expenditures, c_t, appears in the cointegrating relation (3).

[11]The real component of nominal interest income is not directly observable. Nominal interest income can be put in real terms by deflating by a price level to get the component which should be associated with real consumption, but one would still need to subtract the product of some inflation rate and the stock of financial assets from this amount. Measurement is complicated because the stock data are in the flow of funds while the nominal interest data are in the National Income and Product Accounts, and the components do not match precisely.

[12]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 problem with trends in stock market dividend income similar to those for 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 total dollar value of CRSP value-weighted dividends (in real, per capita terms) 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 5.6% from 1978 through 1999, greater than the growth rate for the period 1948 to 1978. The annual growth rate for the whole sample (1948-2001) is 4.2%.

[13]One response to this point is to use the product side of the national income accounts to estimate income flows of such components of wealth 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.

from stock market wealth, d_t, and/or human capital, y_t, the framework above implies a cointegrating relation among c_t, stock market dividends, d_t, and labor income y_t, and the resulting cointegrating 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. First, imply that the cointegrating coefficients in (3) and (4) should not sum to one. As discussed in Lettau and Ludvigson (2001a), the cointegrating parameters in (3) and (4) are likely to sum to a number less than one because only a fraction of total consumption based on nondurables and services expenditure is observable (see Lettau and Ludvigson (2001a)). Second, they have implications for the sums of the cointegrating coefficients in (3) and (4). Denote the shares wealth shares ω and $(1-\omega)$ generically as cointegrating coefficients α_a and α_y, respectively. Likewise, denote the shares ν and $(1-\nu)$ generically as cointegrating coefficients β_d and β_y, respectively. Since some components of aggregate dividends are omitted in (4), the sums $\widehat{\alpha}_a + \widehat{\alpha}_y$ and $\widehat{\beta}_d + \widehat{\beta}_y$, (where "hats" denote estimated values), are unlikely to be identical in finite samples.[14] The parameters $\widehat{\alpha}_a$, $\widehat{\alpha}_y$, $\widehat{\beta}_d$, and $\widehat{\beta}_y$ may be estimated using either single equation or system methods. The estimated values of the cointegrating residuals cay_t and cdy_t are denoted $\widehat{cay}_t$ and $\widehat{cdy}_t$, respectively.

The data used in this study are annual, per capita variables, measured in 1996 dollars, and span the period 1948 to 2001. 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. Annual data is also used because the outcome of both tests for, and estimation of, cointegrating relations can be highly sensitive to seasonal adjustments. 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 all the data used in this study.

Table 1 first presents summary statistics for 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 percent 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 percent, but still more volatile than the log difference in networth, which has a standard deviation of 4 percent. 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 by testing for both the presence and number of cointegrating 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

[14] These conclusions are based on our own Monte Carlo analysis.

$\mathbf{v}_t' = [c_t, a_t, y_t]'$ in Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2002). The results are contained in Appendix C of this paper. We assume 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 the Appendix C suggest the presence of a single cointegrating relation for both vector time series. We denote the single cointegrating 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 cointegrating parameters α_d, α_y and β_d, β_y must be estimated. We use a dynamic least squares procedure which generates asymptotically optimal estimates (Stock and Watson (1993)).[15] This procedure estimates $\widehat{\boldsymbol{\beta}}' = (1, -0.13, -0.68)'$. The Newey-West corrected t-statistics (Newey and West (1987)) for these estimates are -10.49 and -34.82, respectively. We denote the estimated cointegrating residual $\widehat{\boldsymbol{\beta}}'\mathbf{x_t}$ as $\widehat{cdy}_t$. The estimated cointegrating vector for $\mathbf{v}_t' = [c_t, a_t, y_t]'$ is $\widehat{\boldsymbol{\alpha}}' = (1, -0.29, -0.60)'$, very similar 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.

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 2000 it has a first order autocorrelation 0.88, a second order autocorrelation of 0.72 and a third order autocorrelation of 0.60. The autocorrelations of $\widehat{cdy}_t$ and $\widehat{cay}_t$ are much lower and die out more quickly. Their first order autocorrelation coefficients are 0.48 and 0.55, respectively; their second order autocorrelation coefficients are 0.13 and 0.22 respectively.

In Figure 1 we plot the demeaned values of $\widehat{cdy}_t$ and $\widehat{cay}_t$ over the period 1948 to 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 $d_t - p_t$, $\widehat{cay}_t$, and $\widehat{cdy}_t$, it is useful to examine estimates of error-correction representations for $(d_t, p_t)'$, $(c_t, a_t, y_t)'$ and $(c_t, d_t, y_t)'$. Table 3 presents the results of estimating first-order cointegrated vector autoregressions (VARs) for d_t and p_t, for c_t, a_t and y_t, and for c_t, d_t, and y_t.[16] For dividends and prices, the theoretical cointegrating vector $(1, -1)'$ is imposed; for the other two systems, the cointegrating vectors are estimated as discussed above. 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 log dividend-price ratio has little ability to forecast future dividend growth or price growth in the cointegrated VAR. Variation in the log dividend-price

[15] Two leads and lags of the first differences of Δy_t and Δd_t are used in the dynamic least squares regression.

[16] The 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).

ratio is too persistent to display statistical evidence of cointegration in this sample, a result made apparent by the absence of a statistically significant error-correction representation in Panel A (although see the discussion below of the findings in Lewellen (2001) and Campbell and Yogo (2002)). Second, Panel B shows that estimation of the cointegrating 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 cointegrating 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 and labor income. These results are consistent with those in Lettau and Ludvigson (2001a) using quarterly data.

Panel C displays estimates from a cointegrated VAR for c_t, d_t, and y_t. The results 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 cointegrating residual $\widehat{cdy}_{t-1}$. What is strongly predictable by the cointegrating 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 dividends, rather than consumption or labor income, that is forecast to slowly adjust until the cointegrating 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 onto the consumption ratios $\widehat{cdy}_{t-1}$ and $\widehat{cay}_{t-1}$. The theory behind (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. We focus in this paper on explaining the historical behavior of forecastable components of stock market dividend growth, and their relation to forecastable components of excess stock market returns. 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 6 years. In each regression, the dependent variable is the H-period log excess return, $r_{t+1} - r_{f,t+1} + ... + r_{t+H} - r_{f,t+H}$, 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 (2001) 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, $\Delta d_{t+1} + ... + \Delta d_{t+H}$. As hinted at by the results reported in Table 3, neither $\widehat{cay}_t$, or $\widehat{cdy}_t$ has any important long-horizon forecasting power for consumption or labor income growth; to conserve space, we do not report those results here.

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 6 years in this sample. Again, these results differ from those reported elsewhere, primarily because we have included the last few years of stock market data in the sample. 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 (2001a)).[17] By contrast, 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 0.27 for excess returns at a one-year horizon, 0.49 for excess returns over a two year horizon, and 0.52 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, indicating that a high $\widehat{cdy}_t$ forecasts high excess returns just as a high $\widehat{cay}_t$ forecasts high excess returns. The t-statistics are above four for all horizons in excess of one year, and the $\overline{R}^2$ statistic rises from .20 at a three year horizon to .32 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

[17] Other statistical approaches find that the dividend-price ratio remains a strong predictor of excess stock returns even in samples that include recent data. Lewellen (2001) notes that when the dividend-price ratio is very persistent but nonetheless stationary, episodes in which the dividend yield remains deviated from its long-run mean for an extended period of time will 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.

the same forecasting exercise for long horizon dividend growth as presented above for long-horizon excess returns. In this sample, which includes data in the last half of the 1990s, the log dividend-price ratio displays some forecasting power for future dividend growth (row 1), but has the wrong sign (positive), consistent with evidence in Campbell (2002) who also uses data that include the second half of the 1990s. 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 percent of future dividend growth, depending on the horizon.

Lettau and Ludvigson (2001a) found that $\widehat{cay}_t$ had 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 20 percent of the variation in dividend growth, while $\widehat{cay}_t$ explains 9 percent. At a five year horizon, $\widehat{cdy}_t$ explains about 28 percent of the variation in dividend growth, while $\widehat{cay}_t$ explains 10 percent. 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.

The results in Tables 5 and 6 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 predictive power for excess returns and dividend growth. To address this question, Table 6 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. The table shows that, in forecasting long horizon excess returns, $\widehat{cdy}_t$ contains no information about future returns that is independent of that contained in $\widehat{cay}_t$: at all forecasting horizons, $\widehat{cay}_t$ drives out $\widehat{cdy}_t$. 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_t in the discussion above.

The second panel of Table 6 shows that much the opposite pattern is borne out in long-horizon forecasting regressions for dividend growth: $\widehat{cdy}_t$ drives out $\widehat{cay}_t$ in forecasting future dividend growth at all forecasting horizons greater than three years. But for forecasting horizons between 2 and 3 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 2 and 3 year horizons, but they are strongly jointly significant (the p-value for the F-test is less than 0.000001).

This latter finding suggests that much of the variation in expected dividend growth may be common to variation in expected returns, at least for two and three year horizons. The findings also suggest that there may be a component of expected returns that moves independently of expected dividend growth. Note that 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.

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. 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 $d_t - p_t$. The estimation output from these 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. These findings reinforce the conclusion that persistent variation in the log dividend-price ratio is better described as capturing some low frequency component of expected excess returns than variation in expected dividend growth, consistent with the arguments in Heaton and Lucas (1999), Campbell and Shiller (2001), Cochrane (2001), Fama and French (2002), and Lewellen (2001); Campbell (2002).

4.1 Related Empirical Findings

In summary, 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 of important predictability in dividend growth, correlated with important forecastable movements in excess returns, 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, 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 does find predictability of dividend growth in direct long-horizon regressions is Ang and Bekaert (2001), who report results based on observations from 1952:Q4 to 1999:Q4 on the S&P 500 stock market index. Like Ang (2002), they also 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 dividend growth must comprise little of the variation in the dividend-price ratio. But, Ang and Bekaert (2001) do find that the dividend-price ratio has *marginal* predictive power for future dividend growth in a multivariate regression once the earnings yield is also 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-squares 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 percent of dividend growth one year ahead, and about 13 percent a five year horizon. The comparable numbers using $\widehat{cdy}_t$ alone as a predictive variable are 31 percent and 34 percent.[18]

4.2 Additional Statistical Tests

4.2.1 Multivariate Long-Horizon Forecasting Regressions[19]

The cointegrating coefficients in $\widehat{cay}_t$ and $\widehat{cdy}_t$ are estimated using the full sample. This estimation strategy is appropriate for testing the theoretical framework above, because suf-

[18]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.

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

ficiently large samples of data are necessary to recover the true cointegrating 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 (3) and (4) unless the cointegrating coefficients have converged to their true values. Fortunately, cointegrating 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 subsequent estimation. It follows that a valid test of the theoretical cointegration framework in (3) and (4) requires the use of the full sample to estimate the cointegrating coefficients.[20]

A separate issue concerns not whether the theoretical framework is correct, but whether a practitioner, operating in early part of our sample and without access to the whole sample to estimate cointegrating 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 can lead such tests to be substantially less powerful than in-sample forecasting tests (Inoue and Kilian (2002)). This means that out-of-sample (and subsample) tests can fail to reveal true forecasting power that even a practitioner could have had in real time. This pattern that would be exacerbated in any investigation of long-horizon forecasting power.

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 cointegrating 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. To do so, we consider single equation, multivariate regressions taking the form

$$z_{t+h} = a + b_1 c_t + b_2 a_t + b_3 y_t + u_t, \tag{10}$$

where 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 cointegrating relation among c_t, a_t, and y_t in a first stage regression and then using the cointegrating residual as the single right-hand-side variable, the regression (10) uses the multiple variables involved in the cointegrating 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 which generate such a relation. For this excercise, we maintain the hypothesis that the left-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 cointegrating 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 (10) will produce valid adjusted

[20] This issue is discussed in more detail in Lettau and Ludvigson (2001b).

R^2 and t-statistics. Because this procedure does not require any first-stage estimation of cointegrating parameters, it is clear that the forecasting regression (10), its coefficients and R^2 statistics, cannot be influenced by such a prior analysis.

Table 6 reports long-horizon regression results for excess returns and dividend growth, from an estimation of (10) 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.[21] The results are broadly 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 as do the individual regressors $\widehat{cay}_t$ and $\widehat{cdy}_t$. For example, the multivariate regression with c_t, a_t, and y_t as regressors explains about 26 percent of one year ahead excess returns, whereas $\widehat{cay}_t$ explains 27 percent. Similarly, the multivariate regression with c_t, d_t, and y_t explains about 24 percent of the variation in one-year-ahead dividend growth, whereas $\widehat{cdy}_t$ explains 20 percent. 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 $\widehat{cay}$, presented above. One is that the use of overlapping data in long-horizon regressions can skew statistical inference in finite samples. Valkanov (2001) shows that, in finite samples where the forecasting horizon is a nontrivial fraction of the sample size, (i) the t-statistics of long-horizon regression coefficients do not converge to

[21] Inference on b_1, b_2, and b_3 can be accomplished by re-writing (10) so that the hypotheses to be tested are written as a restrictions on $I(0)$ variables (Sims, Stock, and Watson (1990)). For example, the hypothesis $b_1 = 0$ can be tested by rewriting (10) 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

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

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 cointegrating coefficients are only required to do inference about the forecasting excercise–they do not affect the forecasting excercise itself.

a well defined distribution, and (ii) 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 $\widehat{cay}$ 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 (2001). Second, we use vector autoregessions to impute long-horizon R^2 statistics, rather than estimating them directly from long-horizon regressions. Third, we perform both bootstrapped and Monte Carlo 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. Valkanov (2001) 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 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 may be consistently estimated using the methodology described in Valkanov (2001). With these estimates in hand, the rescaled t-statistic, $t/\sqrt{T}$, may be compared against critical values provided in Valkanov (2001).

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 table reports 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 percent 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. These results do not support the conclusion that the forecasting power of $\widehat{cay}_t$ and $\widehat{cdy}_t$ for long-horizon excess stock market returns and stock market dividend growth can be attributed to biases arising from the use of overlapping data in finite samples.

Finite sample problems with overlapping data in long-horizon regressions may also be avoided by using vector autoregression to impute implied long-horizon R^2 statistics for univariate forecasting regressions, 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 8, which investigates the long-horizon predictive power of $\widehat{cay}_t$ and $\widehat{cdy}_t$ for future returns and future dividend growth using bivariate, first-order VARs. For each forecasting horizon we consider, we calculate an implied R^2 statistic using the coefficient estimates of the VAR and the estimated covariance matrix of the VAR residuals.

Table 8 shows that the pattern of the implied R^2 statistics from the vector autoregressions is very similar to those from the produced from the single equation long-horizon regressions. The implied adjusted R^2 statistics for forecasting dividend growth with $\widehat{cdy}_t$ (row 3) peaks at 0.2 for a three year horizon. This forecasting power is consistently greater than that obtained from a simple autoregression for dividend growth (row 1). 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 $\widehat{cay}_t$ is as high as 49% at a three year horizon; for forecasting dividend growth with $\widehat{cay}_t$, it reaches 24% at a three year horizon. Thus, the evidence favoring predictability of dividend growth and excess stock returns using $\widehat{cdy}_t$ and $\widehat{cay}_t$ is robust to the VAR methodology, implying that the size of the long-horizon R^2 statistics cannot be readily attributed to inference problems with the use of overlapping data in finite samples.

An alternate 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 9 presents results based on two methodologies which yield very similar results: a bootstrap and a Monte Carlo simulation, both conducted under the null hypothesis of no predictability (i.e., residuals for the dependent variable are generated by regressions on a constant). For both simulations, we use first-order autoregressive specifications our reduced form models for $\widehat{cay}_t$ and $\widehat{cdy}_t$.[22] For the bootstrap, 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.[23] The simulations were repeated 10,000 times. For the Monte Carlo simulation, 10,000 artificial time-series equal to the size of our data set were generated under the null of no predictability by taking random draws from a normal distribution; the notes to Table

[22]It is known that 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.

[23]Nelson and Kim (1993) also perform randomization, which differs from bootstrapping only in that sampling is without replacement. We also performed the simulations using randomization and found that the results were not affected by this change.

9 provides details. To avoid difficulties caused by the use of overlapping data, we focus here on the one-year ahead regressions presented in Tables 5 and 6.

Table 9 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. Panel A presents the bootstrap results; Panel B, the Monte Carlo results. The results of each simulation are nearly identical. In almost every case, the estimated predictability coefficient and R^2 statistic lies outside of the 95 percent confidence interval based on the empirical distribution under the null of no predictability. In most cases they lie outside of the 99 percent 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 indistinguishable 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. In summary, these results, like those discussed above using the rescaled t-statistic and VAR-imputed R^2 statistics, do not support the conclusion that the predictability of excess returns and dividend growth documented here is can be attributed to small sample biases in the regression coefficients or 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)), even if large firms with high earnings have continued to increase traditional dividend payouts over time (DeAngelo, DeAngelo, and Skinner (2002)). 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 share repurchase expenditures for the Industrial Compustat firms reported in Grullon and Michaely (2002) to our measure of dividends.[24] 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, repurchases 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 10 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. 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 $\widehat{cdy}$ are both strong predictors of the long-horizon growth rates in this series, with t-statistics that begin above 4 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 may be possible over long horizons (as revealed by the dividend-price ratio), it may be more difficult over horizons corresponding more closely to the business cycle. Several researchers have presented evidence that is suggestive of this hypothesis. Gertler and Hubbard (1993) study firm-level data from Compustat and find that firm dividend payouts are lower during a slow-down in economic growth and higher during periods of economic expansion. Bernanke and Gertler (1989) and Bernanke, Gertler, and Gilchrist (1996) present theoretical and empirical evidence of countercyclical variation in the 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 counter-cyclical: it rises during an

[24] We add gross repurchases to our measure of dividends, since those data are readily available from the published work of Grullon and Michaely (2002). This procedure is conservative for our purpose, since the sum of share repurchases and traditional dividend payments would only be closer to our original series if we instead added net repurchases (net of new issues).

economic slow-down and falls during periods of economic growth (Fama and French (1989); Ferson and Harvey (1991); Lettau and Ludvigson (2001a)). 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. This suggests that consumers may be better able to smooth transitory fluctuations in their dividend income than managers are able to smooth fluctuations in earnings. 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 11 presents some evidence that is supportive of this hypothesis using earnings data for NYSE firms. The earnings data are from Lewellen (2001) and are operating earnings before depreciation to market value. Unfortunately, only a short sample is available that is limited by when Compustat data are available: 1964-2000.[25] Table 11 reports that earnings growth is predictably higher when predictable dividend growth, as captured by $\widehat{cdy}_t$, is 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 four for one to three year forecasting horizons, and in excess of three for a four year horizon. The univariate forecasting regression explains about 14 percent of the variation in earnings growth 4 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 forecast to grow more quickly in the future. These results are consistent with the hypothesis manager's dividend smoothing ability is imperfect over business cycles.

6 Conclusion

This paper presents evidence that changing forecasts of dividend growth make an important contribution to fluctuations in the 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 affects 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.

The findings suggest that an important component of time-varying expected returns and

[25] We use Lewellen's data and not earnings per share since that measure is contaminated by variability in share issuance.

time-varying expected dividend growth is not captured by the log dividend-price ratio, or likely by other aggregate financial ratios. This stacks the deck against such financial ratios in statistical tests of return or dividend growth predictability. 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 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.[26]

We caution that the findings presented here provide but one piece of a larger puzzle concerning the behavior of the dividend-price ratio, especially that more recently. There is a growing view that shifts in corporate financial policy may have created persistent changes in dividend growth rates. For example, firms have been distributing an increasing fraction of total cash paid to shareholders in the form of stock repurchases (e.g., Fama and French (2001)). It is too soon to tell whether such shifts in corporate financial policy will be sustained. At the same time, stock prices relative to earnings and other measures of economic fundamentals have followed patterns similar to that of the dividend-price ratio (Campbell and Shiller (2001)), while the consumption-based valuation ratio for dividend growth studied here has been less affected. These observations suggest that factors other than changes in corporate payout policy may be partly responsible for 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 between the log dividend-price ratio and the log consumption-wealth ratio have been attributable historically to changing forecasts of long-horizon dividend growth.

[26]For a lucid summary of this literature, see Campbell and Viceira (2001).

Appendix A: Derivation of cdy_t

Equation (4) is based on the derivation in Campbell and Mankiw (1989) for the relation between log consumption and the log of total income flows from aggregate wealth. Campbell and Mankiw move from the consumption-based present value relation involving future returns, (the consumption-wealth ratio), to one involving future income flows. A derivation is given in Campbell and Mankiw (1989) and here.

W_t is total wealth, which consists of N_t shares at time t, each of which have an ex-dividend price, P_t, and dividend payment, I_t:

$$W_t = N_t(P_t + I_t). \tag{11}$$

The return on aggregate wealth is defined

$$R_{t+1} = \frac{P_{t+1} + I_{t+1}}{P_t}. \tag{12}$$

Combining (11) and (12),

$$\frac{W_{t+1}}{N_{t+1}} = R_{t+1}(\frac{W_t}{N_t} - I_t). \tag{13}$$

Equation (13) can be written

$$W_{t+1} = R_{t+1}\left(N_t + \Delta N_{t+1}\right)\left(\frac{W_t}{N_t} - I_t\right) ==>$$

$$W_{t+1} = R_{t+1}\left(W_t - I_t N_t + (W_t - I_t N_t)\frac{\Delta N_{t+1}}{N_t}\right)$$

Note that from (11), $(W_t - I_t N_t) = N_t P_t$. Thus,

$$W_{t+1} = R_{t+1}(W_t - I_t N_t + P_t \Delta N_{t+1}).$$

The term $P_t \Delta N_{t+1}$ is net-new investment, i.e. the net issuance of new shares, ΔN_{t+1}, valued at the ex-dividend price P_t. Investors save by reinvesting a portion of their dividend income in the asset markets.

Equation (13) is of the same form as the accumulation equation for total wealth, $W_{t+1} = R_{t+1}\left(W_t - C_t\right)$, and can be linearized in the same way. Campbell and Mankiw do so and derive

$$i_t - w_t = -n_t + E_t \sum_{i=1}^{\infty} \rho^i (r_{t+i} - \Delta i_{t+i}) + \text{constant}, \tag{14}$$

where lower case letters denote log variables. Note that i_t in (14) is the log per share dividend. To obtain total dividends, I_t must be multiplied by the number of shares N_t; or in logs we need $i_t + n_t$. Adding n_t on both sides of (14) delivers a present value relation relating log total dividends to log total wealth:

$$i_t^T - w_t = +E_t \sum_{i=1}^{\infty} \rho^i (r_{t+i} - \Delta i_{t+i}^T + \Delta n_{t+i}) + \text{constant},$$

where i_t^T denotes total (rather than per share) income from aggregate wealth, $i_t^T \equiv i_t + n_t$.

Combining (14) with the log-linearized expression for the log consumption wealth ratio

$$c_t - w_t = E_t \sum_{i=1}^{\infty} \rho^i (r_{t+i} - \Delta c_{t+i}), \tag{15}$$

yields

$$c_t - i_t^T = E_t \sum_{i=1}^{\infty} \rho^i (\Delta i_{t+i}^T - \Delta n_{t+i} - \Delta c_{t+i}) + \text{constant}. \tag{16}$$

This equation is a more general version of (3.7) in Campbell and Mankiw:

$$c_t - i_t^T = E_t \sum_{i=1}^{\infty} \rho^i (\Delta i_{t+i}^T - \Delta c_{t+i}) + \text{constant}. \tag{17}$$

Campbell and Mankiw arrive at (17) by normalizing (in the last step), N_t, the number of shares in each period, to equal one. Although equation (16) implies that $c_t - i_t^T$ may related to future changes in the log of the number of shares of asset wealth, this implication is not interesting because the pure number of shares is continuously renormalized by stock splits and reverse splits. Note also that the notation in Campbell and Mankiw (1989) is unfortunate, because in their text (and in their equation (3.7)), y_t is used to denote log total income (what we denote i^T here), whereas in their appendix (where they derive equation (17)), y_t denotes the log of income per share, i_t.

Equation (4) is based Campbell and Mankiw's (16), but differs in two respects. First, Campbell and Mankiw assume a particular functional form for investor preferences, and therefore set $E_t \Delta c_{t+i} = \mu + \sigma E_t r_{t+i}$. Second, equation (16) is expressed in terms of the total income flow from aggregate wealth, i_t^T, whereas as in (4), this total is decomposed into its asset wealth and human wealth components using the relation $i_t^T \approx \nu d_t + (1 - \nu)\, y_t$, where ν is the steady state share of income from asset wealth in total income. Together these assumptions yield the expression

$$cdy_t \equiv c_t - \nu d_t - (1 - \nu)\, y_t = E_t \sum_{i=1}^{\infty} \rho_w^i (\nu \Delta d_{t+i} + (1 - \nu) \Delta y_{t+i} - \Delta c_{t+i} - \nu \Delta n_{t+i}). \tag{18}$$

For this simple framework, we have assumed that the number of "shares" of human capital are constant, since human wealth is not traded on a stock market. This assumption is inconsequential for the substance of the derivation, since it merely determines whether Δn_{t+i} in (18) is multiplied by the constant ν. Finally, we follow Campbell and Mankiw (1989) and avoid carrying the term $\nu \Delta n_{t+i}$ around by making an arbitrary normalization that the number shares is always unity. This delivers equation (4) in the text.

The steady state income shares ν and $(1 - \nu)$ can be related to the steady state wealth shares ω and $(1 - \omega)$. To see this, assume that the steady state of the economy is characterized by balanced growth at some gross rate $1 + g$, and a constant return on aggregate wealth, $R_{w,t} \equiv R$. These assumptions are standard features of equilibrium growth models. Equation (15) implies the steady state value of beginning-of-period aggregate wealth is given by

$$
\begin{aligned}
W_t &= \sum_{i=0}^{\infty} (1+R)^{-i}\, C_{t+i} \\
&= \sum_{i=0}^{\infty} (1+R)^{-i} \left(I_{t+i} N_{t+i} + P_t \Delta N_{t+1+i} \right) \\
&= \sum_{i=0}^{\infty} (1+R)^{-i} \left(D_{t+i} + P_t \Delta N^A_{t+1+i} + Y_{t+i} \right),
\end{aligned}
$$

where N^A_{t+1} denotes the change in the number of shares of asset wealth at time $t+1$. Using the expression above, and noting that the steady state ratio of aggregate wealth to consumption is given by $(1+R)\,/\,(R-g)$, it is straightforward to show that the steady share of asset wealth in aggregate wealth, ω, is given by

$$
\omega = \frac{D_t + (\pi - 1)\, P_t N^A_t}{D_t + (\pi - 1)\, P_t N^A_t + Y_t},
$$

where, D_t, $P_t N^A_t$, and Y_t all grow deterministically at rate $1+g$, $\pi \equiv (1+g)\,/\,(1+r-kr+kg)$, and where $k \equiv I^T_t / C_t \geq 1$, is the steady state ratio of total income to total consumption. Notice that when $k = 1$ (there is no saving in steady state), we have $\pi = 1$ and

$$
\omega = \frac{D_t}{D_t + Y_t} = \nu,
$$

and income shares equal wealth shares.

Appendix B: Data Description

The sources and description of each data series we use are listed below.

CONSUMPTION

Consumption is measured as either total personal consumption expenditure or expenditure on nondurables and services, excluding shoes and clothing. The quarterly data are seasonally adjusted at annual rates, in billions of chain- weighted 1996 dollars. The components are chain-weighted together, and this series is scaled up so that the sample mean matches the sample mean of total personal consumption expenditures. Our source is the U.S. Department of Commerce, Bureau of Economic Analysis.

AFTER-TAX LABOR INCOME

Labor income is defined as wages and salaries + transfer payments + other labor income - personal contributions for social insurance - taxes. Taxes are defined as [wages and salaries/(wages and salaries + proprietors' income with IVA and Ccadj + rental income + personal dividends + personal interest income)] times personal tax and nontax payments, where IVA is inventory valuation and Ccadj is capital consumption adjustments. The annual data are in current dollars. Our source is the Bureau of Economic Analysis.

WEALTH

Total wealth is household net wealth in billions of current dollars, measured at the end of the period. Stock market wealth includes direct household holdings, mutual fund holdings, holdings of private and public pension plans, personal trusts, and insurance companies. Nonstock wealth is the residual of total wealth minus stock market wealth, and includes ownership of privately traded companies in noncorporate equity. Our source is the Board of Governors of the Federal Reserve System.

DIVIDENDS

Dividends are constructed from the CRSP index returns. The CRSP dividends, $D_{c,t}$, are scaled by the average ratio of stock market wealth, S_t to the price of the value-weighted CRSP index, $P_{c,t}$ to reflect dollar values, i.e., $D_t \equiv E(S_t/P_{c,t})D_{c,t}$.

POPULATION

A measure of population is created by dividing real total disposable income by real per capita disposable income. Consumption, wealth, labor income, and dividends are in per capita terms. Our source is the Bureau of Economic Analysis.

PRICE DEFLATOR

The nominal after-tax labor income, stock market dividend and wealth data are deflated by the personal consumption expenditure chain-type deflator (1996=100), seasonally adjusted. (Source: Bureau of Economic Analysis.)

Appendix C: Cointegration Tests

This appendix presents the results of cointegration tests. Dickey-Fuller tests for the presence of a unit root in c, y, a, d, and p (not reported) are consistent with the hypothesis of a unit root in those series.

Table C-I reports test statistics corresponding to two cointegration tests. Reported in the far right column are Phillips and Ouliaris (1990) residual based cointegration test statistics. The table shows both the Dickey-Fuller t-statistic and the relevant five and 10 percent critical values. The test is carried out without a deterministic trend in the static regression. We applied the data dependent procedure suggested in Campbell and Perron (1991) for choosing the appropriate lag length in an augmented Dickey-Fuller test. This procedure suggested that the appropriate lag length was one for both the $(c, a, y)'$ system and the $(c, d, y)'$ system. The tests reject the null of no cointegration both systems at the five percent level. The persistent dividend-price ratio displays no evidence favoring cointegration in this sample.

Table C-I also reports the outcome of testing procedures suggested by Johansen (1988) and Johansen (1991) that allow the researcher to estimate the number of cointegrating relationships. This procedure presumes a p-dimensional vector autoregressive model with k lags, where p corresponds to the number of stochastic variables among which the investigator wishes to test for cointegration. For our application, $p = 3$. The Johansen procedure provides two tests for cointegration: under the null hypothesis, H_0, that there are exactly r cointegrating relations, the 'Trace' statistic supplies a likelihood ratio test of H_0 against the alternative, H_A, that there are p cointegrating relations, where p is the total number of variables in the model. A second approach uses the 'L-max' statistic to test the null hypothesis of r cointegrating relations against the alternative of $r + 1$ cointegrating relations.

The critical values obtained using the Johansen approach also depend on the trend characteristics of the data. We present results allowing for linear trends in data, but assuming that the cointegrating relation has only a constant. See the articles by Johansen for a more detailed discussion of these trend assumptions. In choosing the appropriate trend model for our data, we are guided by both theoretical considerations and statistical criteria. Theoretical considerations imply that the long-run equilibrium relationship between consumption, labor income and wealth do not have deterministic trends, although each individual data series may have deterministic trends. The Table also reports the 90 percent critical values for these statistics.

Both the L-max and Trace test results establish evidence of a cointegrating relation among log consumption, log labor income, and the log of household wealth, and among log consumption, log dividends and the log of labor income. Table C-I shows that we may reject the null of no cointegration against the alternative of one cointegrating vector. In addition, we cannot reject the null hypothesis of one cointegrating relationship against the alternative of two or three.

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., and G. Bekaert (2001): "Stock Return Predictability: Is it There?," NBER Working Paper No. 8207.

Bernanke, B. S., and M. Gertler (1989): "Agency Costs, Net Worth and Business Cycle Flutuations," *American Economic Review*, 79, 14–31.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1996): "The Financial Accelerator and the Flight to Quality," *The Review of Economics and Statistics*, 78(1), 1–15.

Campbell, J. Y. (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (2002): "Consumption-Based Asset Pricing," in *Handbook of the Economics of Finance* forthcoming, ed. by G. Constantinides, M. Harris, and R. Stulz. North-Holland, Amsterdam.

Campbell, J. Y., A. W. Lo, and C. MacKinlay (1997): *The Econometrics of Financial Markets*. Princeton University Press, Princeton, NJ.

Campbell, J. Y., and G. Mankiw (1989): "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in *NBER Macroeconomics Annual: 1989*, ed. by O. Blanchard, and S. Fischer, pp. 185–216. MIT Press, Cambridge, MA.

Campbell, J. Y., and P. Perron (1991): "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots," in *NBER Macroeconomics Annual: 1991*, ed. by O. J. Blanchard, and S. Fischer. MIT Press, Cambridge, MA.

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

——— (2001): "Valuation Ratios and the Long-Run Stock Market Outlook: An Update," NBER working paper No. 8221.

Campbell, J. Y., and L. Viceira (2001): *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford Univeristy Press, London, UK.

Campbell, J. Y., and M. Yogo (2002): "Efficient Tests of Stock Return Predictability," Unpublished paper, Harvard University.

Cochrane, J. H. (1991): "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5(2), 243–280.

——— (1994): "Permanent and Transitory Components of GDP and Stock Prices," *Quarterly Journal of Economics*, 109, 241–265.

——— (1997): "Where is the Market Going? Uncertain Facts and Novel Theories," *Economic Perspectives, Federal Reserve Bank of Chicago*, 11, 1–37.

——— (2001): *Asset Pricing*. Princeton University Press, Princeton, NJ.

DeAngelo, H., L. DeAngelo, and D. J. Skinner (2002): "Are Dividends Disappearing? Dividend Concentration and the Consolidation of Earnings," Unpublished paper, University of Michigan.

Fama, E. F., and K. R. French (1988): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–27.

——— (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

——— (2001): "Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay," *Journal of Financial Economics*, 60(1), 3–43.

——— (2002): "The Equity Premium," *Journal of Finance*, 57(2), 637–659.

Ferson, W. E., and C. R. Harvey (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., and R. G. Hubbard (1993): "Corporate Financial Policy, Taxation, and Macroeconomic Risk," *Rand Journal of Economics*, 24(2), 286–303.

Grullon, G., and R. Michaely (2002): "Dividends, Share Repurchases, and the Substitution Hypothesis," *Journal of Finance*, 57(4), 1649–1684.

Heaton, J., and D. Lucas (1999): "Stock Prices and Fundamentals," in *NBER Macroeconomics Annual: 1999*, ed. by O. J. Blanchard, and S. Fischer. 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., and L. Kilian (2002): "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?," Unpublished paper, Department of Economics, University of Michigan.

Johansen, S. (1988): "Statistical Analysis of Cointegrating Vectors," *Journal of Economic Dynamics and Control*, 12, 231–254.

——— (1991): "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models," *Econometrica*, 56, 1551–1580.

Kandel, S., and R. Stambaugh (1989): "Modeling Expected Stock Returns for Long and Short Horizons," Unpublished Paper , Rodney L. White Center, Wharton School, University of Pennsylvania.

Lamont, O. (1998): "Earnings and Expected Returns," *Journal of Finance*, 53, 1563–87.

Lettau, M., and S. C. Ludvigson (2001a): "Consumption, Aggregate Wealth and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

——— (2001b): "Measuring and Modeling Variation in the Risk-Return Tradeoff," Unpublished paper, New York University.

——— (2002): "Understanding Trend and Cycle in Asset Values," Unpublished Paper, New York University.

Lewellen, J. W. (2001): "Predicting Returns With Financial Ratios," Unpublished paper, MIT Sloan School of Management.

Nelson, C. C., and M. J. Kim (1993): "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance*, 43, 641–661.

Newey, W. K., and K. D. West (1987): "A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Phillips, P., and S. Ouliaris (1990): "Asymptotic Properties of Residual Based Tests for Cointegration," *Econometrica*, 58, 165–193.

Sims, C. A., J. H. Stock, and M. W. Watson (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., and M. Watson (1993): "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems," *Econometrica*, 61, 783–820.

Valkanov, R. (2001): "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics,* forthcoming.

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, CRSP-VW dividends Δd_t, CRSP-VW price Δp_t and asset wealth Δa_t. The sample spans the period 1948 to 2001.

Table 2: Autocorrelations of Ratios

Ratio	ρ_1	ρ_2	ρ_3	ρ_4
$d - p$	0.875	0.724	0.596	0.473
$c - 0.29\, a - 0.60\, y$	0.551	0.130	0.085	0.051
$c - 0.13\, d - 0.68\, y$	0.475	0.217	0.258	0.171

Notes: This table reports autocorrelations of ratios involving consumption c_t, labor income y_t, CRSP-VW dividends d_t, CRSP-VW price 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 2 leads and lags. The sample is annual and spans the period 1948 to 2001.

Table 3: Estimates From Cointegrated VARs

Panel A: (d, p)

Dependent Variable	Equation Δd_t	Δp_t
Δd_{t-1}	-0.194	0.364
(t-stat)	(-1.059)	(1.352)
Δp_{t-1}	-0.192	-0.210
(t-stat)	(-1.441)	(-1.079)
$d_{t-1} - p_{t-1}$	0.103	0.070
(t-stat)	(2.205)	(1.021)
$\bar{R}^2$	0.183	0.046

Panel B: (c, a, y)

Dependent variable	Equation Δc_t	Δy_t	Δa_t
Δc_{t-1}	0.267	0.449	-0.523
(t-stat)	(1.279)	(1.220)	(-0.696)
Δy_{t-1}	-0.039	-0.148	0.433
(t-stat)	(-0.294)	(-0.641)	(0.916)
Δa_{t-1}	**0.112**	0.128	**0.392**
(t-stat)	(2.777)	(1.794)	(2.702)
$\widehat{cay}_{t-1}$	-0.007	0.102	**1.726**
(t-stat)	(-0.053)	(0.457)	(3.803)
$\bar{R}^2$	0.199	0.050	0.207

Panel C: (c, d, y)

Dependent variable	Equation Δc_t	Δy_t	Δd_t
Δc_{t-1}	**0.469**	0.652	-0.136
(t-stat)	(2.284)	(1.869)	(-0.060)
Δy_{t-1}	-0.074	-0.156	-0.252
(t-stat)	(-0.572)	(-0.709)	(-0.176)
Δd_{t-1}	**0.029**	**0.052**	-0.129
(t-stat)	(2.311)	(2.389)	(-0.917)
$\widehat{cdy}_{t-1}$	-0.038	0.219	**2.400**
(t-stat)	(-0.408)	(1.377)	(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 bold face. The sample is annual and spans the period 1948 to 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	0.24	0.27	0.34	0.52	0.73
	(1.90)	(1.40)	(1.21)	(0.73)	(0.84)	(1.12)
	[0.08]	[0.10]	[0.10]	[0.10]	[0.16]	[0.23]
$\widehat{cay}_t$	**6.48**	**11.78**	**13.23**	**13.62**	**16.81**	**21.94**
	(4.19)	(5.42)	(5.42)	(5.27)	(7.07)	(5.46)
	{0.57***}	{0.74***}	{0.74***}	{0.72***}	{0.96***}	{0.74***}
	[0.27]	[0.49]	[0.46]	[0.37]	[0.39]	[0.52]
$\widehat{cdy}_t$	1.32	**5.21**	**6.11**	**6.77**	**18.09**	**11.40**
	(1.47)	(7.38)	(4.13)	(4.28)	(4.92)	(4.45)
	{0.20}	{1.00***}	{0.56***}	{0.58***}	{0.67***}	{0.61***}
	[0.01]	[0.16]	[0.20]	[0.20]	[0.20]	[0.32]

Notes: This tables reports the results of h-period regressions of CRSP-VW returns 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 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 and 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 bold face. 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 to 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**	**0.18**	**0.19**	**0.23**	**0.29**	**0.34**
	(2.94)	(2.11)	(2.70)	(2.27)	(2.70)	(2.41)
	[0.07]	[0.15]	[0.13]	[0.14]	[0.15]	[0.19]
$\widehat{cay}_t$	**4.74**	**5.89**	**4.90**	**4.30**	**5.13**	**5.72**
	(6.26)	(4.86)	(3.33)	(2.80)	(2.17)	(1.50)
	{0.85***}	{0.66***}	{0.45***}	{0.38***}	{0.30*}	{0.20}
	[0.29]	[0.30]	[0.16]	[0.09]	[0.10]	[0.12]
$\widehat{cdy}_t$	**2.74**	**3.95**	**3.65**	**3.99**	**5.24**	**6.13**
	(4.06)	(5.84)	(4.13)	(3.60)	(5.38)	(3.65)
	{0.55***}	{0.79***}	{0.56***}	{0.49***}	{0.73***}	{0.50***}
	[0.20]	[0.24]	[0.20]	[0.20]	[0.28]	[0.37]

Notes: This tables reports results from h-period regression of CRSP-VW dividend growth: $d_{t+h} - d_t = k + \gamma\ z_t + \epsilon_{t,t+h}$, where 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 and 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 bold face. 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 to 2001.

Table 6: Multivariate Long-horizon Regressions

	Horizon h (in years)					
Variables	1	2	3	4	5	6
	h-period regression: excess stock returns					
$\widehat{cay}_t$	**6.94**	**11.15**	**12.22**	**11.95**	**15.33**	**18.47**
	(3.27)	(3.73)	(3.82)	(3.43)	(3.86)	(4.12)
$\widehat{cdy}_t$	-0.74	0.89	1.14	1.70	1.46	3.44
	(-0.81)	(0.69)	(0.80)	(1.02)	(0.76)	(1.81)
	[0.27]	[0.48]	[0.45]	[0.36]	[0.38]	[0.53]
	h-period regression: dividend growth					
$\widehat{cay}_t$	**3.71**	4.27	2.62	0.62	-0.30	-0.82
	(3.08)	(1.95)	(1.09)	(0.31)	(-0.14)	(-0.24)
$\widehat{cdy}_t$	**1.64**	2.29	2.58	**3.72**	**5.37**	**6.48**
	(2.57)	(1.86)	(1.87)	(2.74)	(4.78)	(3.32)
	[0.34]	[0.35]	[0.21]	[0.18]	[0.26]	[0.36]

Notes: This tables reports results from h-period regression of CRSP-VW returns in excess of a 3-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 and 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 1948 to 2001.

Table 7: Multivariate Long-horizon Regressions

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

Notes: See next page.

Notes: This tables reports results from h-period regression of CRSP-VW returns in excess of a 3-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 and 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 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 $\sqrt{T}(\widehat{\beta}_1 - \beta_1) \rightarrow N(0, \sigma^2(1/T\sum_{t=1}^{T}(\eta_t - \bar{\eta}_t)^2)^{-1})$ 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 to 2001.

Table 8: Implied Long-Horizon R^2 from VARs

row	Variables	Implied R^2 for Forecast Horizon H					
		1	2	3	4	5	6
1	Δd_t	0.12	0.09	0.05	0.05	0.04	0.03
2	$\Delta d_t,\ \widehat{cay}_t$	0.34	0.31	0.24	0.20	0.17	0.14
3	$\Delta d_t,\ \widehat{cdy}_t$	0.17	0.19	0.20	0.19	0.19	0.19
4	r_t	0.08	0.09	0.05	0.03	0.03	0.03
5	$r_t,\ \widehat{cay}_t$	0.36	0.52	0.49	0.42	0.39	0.36
6	$r_t,\ \widehat{cdy}_t$	0.20	0.26	0.26	0.28	0.31	0.32

Note: The table reports implied R^2 statistics for H-year dividend growth and excess returns obtained from second-order vector autoregressions. The column denoted "Variables" lists the variables included in the VAR. The implied (unadjusted) R^2 statistics for dividend growth in rows 1, 2 and 3 and excess returns in rows 4, 5 and 6 for horizon H are calculated from the estimated parameters of the VAR and the estimated covariance matrix of VAR residuals. Row 1 gives the implied R^2 statistic for forecasting dividend growth with lagged dividend growth, row 2 with lagged $\widehat{cay}_t$ and row 3 with lagged $\widehat{cdy}_t$. Row 4 gives the implied R^2 statistic for forecasting excess stock market returns with lagged returns, row 5 with lagged $\widehat{cay}_t$, and row 6 with lagged $\widehat{cdy}_t$. The sample is annual and spans the period 1948 to 2001.

Table 9: Small Sample Inference of Slope and R^2

x_t	$\widehat{\beta}$	95% CI	99% CI	R^2	95% CI	99% CI
			Panel A: Bootstrap			
			excess returns			
$\widehat{cay}_t$	6.48	(-2.78, 3.86)	(-4.13, 5.44)	0.27	(-0.01, 0.06)	(-0.02, 0.11)
$\widehat{cdy}_t$	1.32	(-2.09, 2.96)	(-3.10, 4.15)	0.01	(-0.01, 0.07)	(-0.02, 0.10)
			dividend growth			
$\widehat{cay}_t$	4.74	(-2.12, 2.61)	(-3.10, 3.78)	0.29	(-0.02, 0.06)	(-0.02, 0.11)
$\widehat{cdy}_t$	2.74	(-1.32, 2.19)	(-1.94, 3.09)	0.20	(-0.02, 0.06)	(-0.02, 0.11)
			Panel B: Monte Carlo			
			excess returns			
$\widehat{cay}_t$	6.48	(-3.31, 3.17)	(-4.73, 4.60)	0.27	(-0.02, 0.05)	(-0.02, 0.10)
$\widehat{cdy}_t$	1.32	(-2.45, 2.44)	(-3.54, 3.52)	0.01	(-0.02, 0.06)	(-0.02, 0.11)
			dividend growth			
$\widehat{cay}_t$	4.74	(-1.96, 1.98)	(-2.97, 2.96)	0.29	(-0.02, 0.06)	(-0.02, 0.11)
$\widehat{cdy}_t$	2.74	(-1.58, 1.57)	(-2.13, 2.29)	0.20	(-0.02, 0.06)	(-0.02, 0.11)

Notes: This tables reports confidence intervals from a bootstrap procedure (Panel A) and a Monte Carlo simulation (Panel B). 10,000 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 dividends 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 $\widehat{\beta}$ and the R^2. In the bootstrap we draw (with replacement) from the residuals of the system estimated under the null hypothesis. For the Monte Carlo analysis the residuals e and v are drawn from a normal distribution. The columns denoted $\widehat{\beta}$ and R^2 report our empirical estimates using annual data from 1948 to 2001.

Table 10: 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	0.10	0.10	0.11	0.15	0.19
	(1.76)	(1.12)	(0.81)	(0.70)	(0.81)	(0.88)
	[0.01]	[0.01]	[0.00]	[0.00]	[0.01]	[0.02]
$\widehat{cay}_t$	**4.66**	**6.36**	**6.52**	**6.51**	**7.97**	**9.24**
	(4.58)	(4.52)	(3.44)	(2.97)	(4.15)	(3.95)
	[0.24]	[0.25]	[0.19]	[0.15]	[0.18]	[0.22]
$\widehat{cdy}_t$	**4.28**	**5.10**	**4.59**	**4.77**	**6.47**	**8.29**
	(5.67)	(5.05)	(2.91)	(2.32)	(3.23)	(3.98)
	[0.20]	[0.19]	[0.12]	[0.10]	[0.16]	[0.24]

Notes: This tables reports results from h-period regression of CRSP-VW dividend growth: $d_{t+h} - d_t = k + \gamma\ z_t + \epsilon_{t,t+h}$, 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 and 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 bold face. The sample is annual and spans the period 1948 to 2000, since the repurchases data from Grullon and Michaely are only available through 2000.

Table 11: Long-horizon Regression – Earnings Growth

Variables	Horizon h (in years) 1	2	3	4	5	6
$\widehat{cdy}_t$	**2.16**	**3.46**	**4.73**	**6.68**	**6.75**	**7.05**
	(4.88)	(6.50)	(4.51)	(3.56)	(2.20)	(2.01)
	[0.07]	[0.06]	[0.07]	[0.14]	[0.13]	[0.13]

Notes: This tables 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 and 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 to 2000.

Table C-I: Cointegration Tests

	L-max Test				Trace Test				t-Test
Variables	$H_0 : r =$	0	1	2	$H_0 : r =$	0	1	2	H_0 : no cointegration
10% Critical Values		12.10	2.82			13.31	2.71		-2.60
5% Critical Values		14.04	3.96			15.41	3.76		-2.93
d, p		6.06	**4.56**			10.62	**4.56**		-0.47
10% Critical Values		18.70	12.10	2.82		26.70	13.31	2.71	-3.52
5% Critical Values		20.78	14.04	3.96		29.68	15.41	3.76	-3.80
c, a, y		**25.34**	6.57	0.07		**31.98**	6.64	0.07	**-4.13**
c, d, y		**27.58**	5.36	1.08		**34.01**	6.43	1.08	**-3.77**

Notes: The first two columns report the L-max and Trace test statistics described in Johansen (1988) and Johansen (1991). The former tests the null hypothesis that there are r cointegrating relations against the alternative of r+1; the latter tests the null of r cointegrating relations against the alternative of p, where p is the number of variables in the cointegrated system. The last column reports the Dickey-Fuller test for $d_t - p_t$ and the Phillips-Ouliaris (1990) cointegration test for (c, a, y) and (c, d, y). The critical values for the Phillips-Ouliaris tests allow for trends in the data while the Dickey-Fuller regression does not include a trend, since according to the theory, there should be no trend in d-p. One lag was used for all tests. The variables are consumption c_t, labor income y_t, CRSP-VW dividends d_t, CRSP-VW price p_t and asset wealth a_t. The null hypothesis is no cointegration; significant statistics at the 10% level are highlighted in bold face. The sample is annual and spans the period 1948 to 2001.

Figure 1: CDY and CAY