

Expected Returns and Expected Dividend Growth*

Martin Lettau
New York University and CEPR

Sydney C. Ludvigson
New York University

PRELIMINARY

Comments Welcome

First draft: July 24, 2001

This draft: September 3, 2002

*Lettau: Department of Finance, Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012-1126; Email: mlettau@stern.nyu.edu, Tel: (212) 998-0378; Fax: (212) 995-4233; <http://www.stern.nyu.edu/~mlettau>. Ludvigson: Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003; Email: sydney.ludvigson@nyu.edu, Tel: (212) 998-8927; Fax: (212) 995-4186; <http://www.econ.nyu.edu/user/ludvigsons/>. 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 John Y. Campbell, Mark Gertler, Anthony Lynch, Lucrezia Reichlin, Peter Schotman and seminar participants at Duke University, New York University, the University of Maryland, the SITE 2001 conference, and the CEPR Summer 2002 Finance Symposium for helpful comments, and Nathan Barczy for excellent research assistance. Any errors or omissions are the responsibility of the authors.

Expected Returns and Expected Dividend Growth

Abstract

We develop 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 “business cycle” variation in expected returns, and the results suggest that a substantial fraction of the variation in expected dividend growth is common to variation in expected excess returns. Movements in expected dividend growth that are entirely common to movements in expected returns have no effect on the log dividend-price ratio.

An implication of these findings is that the log dividend-price ratio will have difficulty predicting both dividend growth and excess returns at business cycle frequencies. Such a failure of predictive power is not an indication that risk-premia are constant, however. On the contrary, the results presented here imply that the log dividend-price ratio will have difficulty revealing business cycle variation in both the equity risk-premium and expected dividend growth precisely *because* expected returns fluctuate at those frequencies, and covary with changing forecasts of dividend growth. The findings 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.

JEL: G12, G10.

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; Fama and French (2002); Campbell (2001); Lewellen (2001)).

¹A 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); Campbell (2001); Cochrane (2001).

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

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

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.

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 provide 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 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 quarters to several years? Moreover, if business cycle variation in expected returns is apparent, 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, the evidence presented here is the first, to our knowledge, to suggest that there is important predictability at longer horizons, and at horizons over which excess stock returns have been found to be forecastable. Second, 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 business cycle variation in expected returns, the latter captured by movements in cay_t , and our results suggest that a substantial fraction of its variation is common to variation in expected excess stock returns.

These findings help resolve the puzzles discussed above, for several reasons. First, they can explain why business cycle variation in expected excess returns that is well captured by cay_t , is not well captured by the dividend-price ratio. Movements in expected dividend growth that are *common* to movements in expected returns have offsetting effects on the log dividend-price ratio. Second, the results help explain why the dividend-price ratio has consistently been found to be a relatively weak predictor of US dividend growth, despite the evidence presented here that dividend growth is highly forecastable. Again the reason is that movements in expected dividend growth that are common to movements in expected returns have offsetting effects on the log dividend-price ratio. Third, although common movements in expected returns and expected dividend growth have offsetting effects on the dividend-price ratio, such movements will not have offsetting effects on the log consumption-wealth ratio, or on cay_t . It follows that cay_t should be affected by both business cycle and very long horizon variation in expected returns, a phenomenon that would make cay_t less persistent than the dividend-price ratio, consistent with the data.

To understand intuitively why dividend growth might be forecastable, it is useful to recall the interpretation offered in Lettau and Ludvigson (2001a) for why cay_t might forecast returns. According to that interpretation, forward-looking investors who want to maintain flat consumption paths over time will set consumption so as to “smooth out” transitory variation in wealth arising from time-variation in expected returns, implying that cay_t is likely to reveal changing forecasts of future returns to aggregate wealth. The same logic suggests that aggregate consumption may also contain information about future dividend payments from aggregate wealth. We refer to the total dividend payments from aggregate wealth as *aggregate dividends*. Consumption and components of aggregate dividends (including stock market dividends) are likely to share a common trend, and deviations from this common trend should be a function of expected future dividend growth, just as deviations from

the common trend in consumption and aggregate wealth should be a function of expected future returns to aggregate wealth. Consistent with this hypothesis, we find that log consumption is cointegrated with empirical measures of dividend payments from aggregate wealth, and that deviations from their common stochastic trend reveal changing forecasts of dividend growth to the stock market component of aggregate wealth. This result is directly analogous to the finding that cay_t reveals changing forecasts of future returns to the stock market component of aggregate wealth (Lettau and Ludvigson (2001a)).

Our approach represents a departure from the standard one of studying models of expected dividend variation, without considering how those expectations are determined in relation to aggregate consumption. In a classic paper, Miller and Modigliani (1961) provided assumptions under which the value of the firm is independent of dividend payout policy, with the striking conclusion that finance theory provides no prediction as to the behavior of dividends paid on equity. Here we argue that, although finance theory provides no prediction about dividends, consumer theory provides a prediction about optimal consumption, which is inextricably tied to aggregate dividends in the long-run. Thus, the long-run behavior of aggregate dividends is pinned down by theory, and may be inferred from observable consumption behavior.

We emphasize four implications of our findings. First, the log dividend-price ratio is likely to fail statistical tests of return predictability at anything but extremely long horizons, consistent with the 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. Once we entertain the possibility that there is substantial comovement in expected returns and expected dividend growth at some frequencies, the logic that the dividend-price ratio should reveal time-variation in expected returns over those horizons is stood on its head. Movements in expected returns that are entirely common to movements in expected dividend growth have no effect on the log dividend-price ratio. 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.

Second, 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 invest-

ment opportunity set that occur over business cycle horizons. This implication is especially relevant for the growing body of literature on strategic asset allocation, in which the log dividend-price ratio is used as a proxy for time-variation in the investment opportunity set, and as an input into the optimal asset allocation decision of a long-horizon investor.⁴

Third, the results presented here imply that modeling stocks as if dividends pay a multiple $\lambda > 1$ of log aggregate consumption is unlikely, by itself, to provide a complete resolution of the equity premium puzzle. Several researchers have noted that introducing such autonomous variation in dividend growth in equilibrium asset pricing models may offer one way of rationalizing the observed equity premium, since those models can in principle make the dividend claim far more risky than the consumption claim. Campbell (1986), Abel (1999), Bansal and Yaron (2000) and Campbell (2001) develop models of this form. If these models are modified so that consumption and aggregate dividends are cointegrated, however, the dividend claim cannot be more risky than the consumption claim in the long-run, no matter how volatile dividends are relative to consumption in the short-run. We illustrate this point using a simple example, which shows that if the adjustment parameter governing the error correction representation for consumption and dividends takes on empirically plausible values, the dividend claim will be only slightly more risky than the consumption claim.

Finally, our findings imply that the log dividend-price ratio and the log consumption-wealth ratio may, at times, give very different signals about the future path of stock prices. The most recent episode in history provides an example. The level of aggregate stock market valuation at the end of 2000 would still require a 75 percent decline in stock prices to restore the dividend-price ratio to its historical mean; by contrast, the empirical measure of the consumption-wealth ratio was largely restored to its sample mean after the broad market declines in late 2000 and early 2001 (Lettau and Ludvigson (2001c)).

The rest of this paper is organized as follows. In the next section, we lay out the theoretical framework linking aggregate consumption and dividend payments from aggregate wealth, to the expected future growth rates of dividends, and show how we express 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. Here we emphasize our use of annual data to insure that any forecastability of dividend growth we uncover is not attributable to the seasonality of dividend payments. Section 4 presents the outcome of forecasting regressions for dividend growth on the US stock market. Sec-

⁴For a lucid summary of this literature, see Campbell and Viceira (2001).

tion 5 aims to quantify the common variation in expected dividend growth and expected returns by modeling expected returns and expected dividend growth in a simple principal components framework, and by employing the frequency-domain measures of comovement developed in Croux, Forni, and Reichlin (2002). These findings reinforce the conclusion that persistent variation in the log dividend-price ratio is better described as low frequency variation in forecasts of excess stock market returns than in forecasts of dividend growth, consistent with the arguments in Heaton and Lucas (1999), Campbell and Shiller (2001), Cochrane (2001), Fama and French (2002), Campbell (2001), and Lewellen (2001). In Section 6 we discuss some implications of our findings for the equity premium puzzle, mentioned above. Finally, Section 7 discusses one possible explanation for why expected returns might be positively correlated with expected dividend growth on US stock markets, even though firms may have an incentive to smooth dividend payments if shareholders themselves desire smooth consumption paths. Section 8 concludes.

2 A Consumption-Based Present Value Relation for Dividend Growth

This section develops 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

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

Labor income Y_t does not appear explicitly in this equation because of the assumption that the market value of tradable human capital is included in aggregate wealth.⁵ 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

⁵None 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, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^j (1 + R_{a,t+i})^{-i} Y_{t+j}$.

condition.⁶ The resulting expression is:⁷

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

where $\rho_w \equiv 1 - \exp(\overline{c - w})$. This expression says that the log consumption-wealth ratio embodies rational forecasts of returns and consumption growth.

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 involving future returns to asset wealth⁸

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

where ω is the average share of asset wealth, A_t , in aggregate wealth, W_t , $r_{a,t}$ is the log return to asset wealth, A_t , and α_a and α_y are parameters to be estimated, discussed further below. 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 cointegrating residual on the left-hand-side of (3) is denoted cay_t for short. Both (2) and (3) are consumption-based present-value relations involving future returns to wealth.

In this paper we use the same accounting framework to construct a consumption-based present value relation involving future dividend growth. The objective of this paper is to

⁶This transversality condition rules out rational bubbles.

⁷We omit unimportant linearization constants in the equations throughout the paper.

⁸We 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 two conditions: first, expected labor income is constant, and second, the return to human capital is either constant or proportional to the return to nonhuman capital. Although we do not observe the return to human capital, Lettau and Ludvigson (2001c) find that expected future labor income growth does not appear to vary much in aggregate data.

study the behavior of a particular component of aggregate dividends, namely dividends to stock market wealth, which we denote d_t . 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. A complete derivation is given in Appendix A. This delivers a present-value relation involving the log of consumption and the logs of dividends from stock market wealth and nonstock dividends including primarily labor income, the dividend from human capital. Rather than creating additional notation, we denote nonstock dividends as y_t , since estimates of national income shares suggest that labor income is by far the most important component of nonstock income produced by private factors of production. The resulting present value relation takes the form

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

where ν is the average share of stock market wealth in aggregate wealth, and β_a and β_y are parameters to be estimated, discussed further below.

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, stock market dividends, and dividends from other forms of aggregate wealth (primarily human capital) should be cointegrated, and that deviations from their common trend (given by the left-hand-side of (4)) provide a rational forecast of either dividend growth, labor income growth, consumption growth, or some combination of all three. We denote the cointegrating residual on the left-hand-side of (4) as cdy_t , for short.

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 the cay_t summarize expectations of future returns to assets wealth, provided that expected labor income growth is not too volatile. This implication was studied in Lettau and Ludvigson (2001a). Those results indicate that cay_t has little predictive power for consumption and labor income growth, but instead forecasts excess returns on the aggregate stock market. Notice that if cay_t forecasts only asset returns, (3) says that it is a state variable that summarizes changing forecasts of future returns to asset wealth. 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. Thus, consumption should be high relative to its long-run trend relation with d_t and y_t in anticipation of high dividend growth in the future, and low in anticipation of low future dividend growth. Moreover, if expected consumption growth and expected labor income growth do not vary much, as previous research suggests, 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 stock market 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, a subset of total consumption which includes the unobservable service flow from the stock of durable goods.

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 labor income. 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, creates peculiar trends in interest income that have nothing in particular to do with the evolution of permanent income. These problems are especially evident in our sample during the 1970s and 1980s when nominal interest income grew rapidly because of inflation.⁹ In addition, we do not directly observe at least one component of the income flow from aggregate wealth, namely the dividend payments from some forms of nonhuman, nonfinancial household net worth, primarily physical capital.¹⁰ Nevertheless, if nonstock/nonlabor forms of dividend income are cointegrated with the dividend payments d_t and y_t (as models with balanced growth would suggest), the framework above implies a cointegrating relation among c_t , d_t and labor income y_t , and the resulting cointegrating residual will reveal expectations over long-horizons of either future Δd_t , Δy_t or Δc_t , or some combination of all three. For these reasons we focus in this paper on results based on using consumption, c_t , stock market dividends, d_t , and labor income, y_t , to form an estimate of a cointegrating residual cdy_t .

Taken together, these data issues imply that the cointegrating coefficients in both (3) and (4) should not sum to one. As discussed in Lettau and Ludvigson (2001a), the cointegrating parameters α_a and α_y in (3) are, in principle, equal to the shares ω and $(1 - \omega)$; in practice, the estimated values of these parameters 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)). The same issues apply to the cointegrating parameters β_d and β_y , which are in principle equal to the shares ν and $1 - \nu$. In addition, the sums of estimated coefficients (where “hats” denote estimated values), $\widehat{\alpha}_a + \widehat{\alpha}_y$ and $\widehat{\beta}_d + \widehat{\beta}_y$, are unlikely to be identical, since a component of aggregate dividends is omitted in (4).

⁹Some 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. In fact, 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%.

¹⁰One 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 requires that the income and product sides of the national accounts be combined, however, a procedure that creates its own measurement difficulties since 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.

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.

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 $\mathbf{v}'_t = [c_t, a_t, y_t]'$ in Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2001c), and those tests suggest the presence of a single cointegrating relation among those variables. We refer the reader to those papers for details and simply note here that there is strong evidence of cointegration among c_t , a_t , and y_t .¹¹ The results of cointegration tests for c_t , d_t , and y_t are contained in Appendix C of this paper. We assume all of the variables contained in \mathbf{x}_t are first order integrated, or $I(1)$, an assumption verified by unit root tests. In addition, the findings presented in the Appendix C suggest the presence of a single cointegrating vector for the three variables in \mathbf{x}_t . The cointegrating coefficient on consumption is normalized to one, and we denote the single cointegrating relation for $\mathbf{x}'_t = [c_t, d_t, y_t]'$ as $\boldsymbol{\beta}' = (1, -\beta_d, -\beta_y)'$.

The cointegrating parameters α_d and α_y must be estimated. We use a dynamic least squares procedure which generates super-consistent estimates of β_d and β_y (Stock and Watson (1993)).¹² 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.

¹¹Cointegration tests in Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2001c) were based on quarterly data. The outcome of these tests is not altered by using the annual data in this study.

¹²Two leads and lags of the first differences of Δy_t and Δd_t are used in the dynamic least squares regression.

3.2 An Illustrative Example

In the next section we discuss empirical results which suggest that expected dividend growth and expected excess returns contain a common component. Before considering those findings, it is useful to consider a simple example that illustrates how common variation in expected dividend growth and expected returns affects the log dividend-price ratio. Let p_t be the log price of stock market wealth, which pays the dividend, d_t . Following Campbell and Shiller (1988) the log dividend-price ratio may be written up to a first-order approximation as

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

where $\rho = 1 / (1 + \exp(\overline{d - p}))$, and r_t is the log return to stock market wealth. This equation says that if the log dividend-price ratio is high, agents must be expecting 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 virtually none of the variability of dividend growth over long-horizons, and as a consequence, expected dividend growth is often modelled as constant. Instead, what has been found to be forecastable over long-horizons is excess returns, suggesting that risk premia vary slowly over time. Notice that the consumption-based present value relation for future dividend growth, (4), is an alternative valuation metric for capturing possible time-variation in expected dividend growth.

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

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

Clearly the log price dividend ratio does not forecast dividend growth at any horizon but instead forecasts long-horizon stock returns. Forecasting regressions of long-horizon returns on the log dividend-price ratio will display R^2 statistics that are hump-shaped in the horizon, depending on the value of ϕ .¹³ This example 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 follows that long-horizon forecasting regressions of excess returns provide one way of assessing whether risk-premia vary.

¹³Campbell, Lo, and MacKinlay (1997), Chapter 7, provides an illustration of this point.

Now suppose, in contrast to the standard view, that expected dividend growth is not constant but instead varies according to the first-order autoregressive process,

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

The effect of such variation on the dividend-price ratio depends on the correlation between expected dividend growth and expected returns. For example, if variation in expected dividend growth is common to variation in expected returns, expected returns may contain two components, $E_t r_{t+1} = g_t + x_t$, so that the log dividend-price ratio becomes

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

$$= \left(\frac{g_t}{1 - \rho\psi} + \frac{x_t}{1 - \rho\phi} \right) - \frac{g_t}{1 - \rho\psi} = \frac{x_t}{1 - \rho\phi}. \quad (9)$$

This expression is precisely the same as (6) for the case in which expected dividend growth is constant. The log dividend-price ratio still forecasts returns because it captures one component of excess returns, x_t , that is independent of expected dividend growth, but it does not forecast dividend growth even though, by construction, expected dividend growth varies over time. Moreover, variation in $d_t - p_t$ does not capture one component of time-varying expected returns, g_t , because that component is common to time-varying expected dividend growth. Variation in expected dividend growth is correlated with changing forecasts of future returns and has completely offsetting effects on the log dividend-price ratio. It follows that fluctuations in the log dividend-price ratio will be driven only by one of the two components of expected returns, x_t . If we relax the assumption that expected dividend growth is approximately constant, it no longer follows that any and all variation in expected returns must be revealed by variation in the dividend-price ratio.

By contrast, the ability of cay_t to reveal variation in expected returns and the ability of cdy_t to reveal variation in expected dividend growth is not affected by whether expected returns and expected dividend growth are correlated with one another. To see this, suppose for expositional clarity that consumption growth is i.i.d., i.e., $\Delta c_{t+1} = \epsilon_{t+1}$, where ϵ_{t+1} is an unforecastable error, and consider a simple example in which aggregate wealth is equal to stock market wealth and aggregate dividends are equal to stock market dividends. In this case, the share of stock market wealth in aggregate wealth is unity, implying that cay_t in (3) collapses to the log consumption-wealth ratio, $c_t - w_t = c_t - p_t$, and cdy_t in (4) collapses to the log consumption-dividend ratio $c_t - d_t$. Combining (4) and (7), it is straightforward to show that the log consumption-dividend ratio forecasts dividend growth and the common

component of expected returns:¹⁴

$$c_t - d_t = E_t \sum_{i=1}^{\infty} \rho^i (\Delta d_{t+i} - \Delta c_{t+i}) = \frac{g_t}{1 - \rho\psi}, \quad (10)$$

while, using (3), the log consumption-wealth ratio captures both components of expected returns:

$$c_t - w_t = E_t \sum_{i=1}^{\infty} \rho^i (r_{t+i} - \Delta c_{t+i}) = \frac{g_t}{1 - \rho\psi} + \frac{x_t}{1 - \rho\phi}. \quad (11)$$

Equation (11) says that the log consumption-wealth ratio is a better predictor of returns than $d_t - p_t$ because it captures both components of time-varying expected returns, g_t and x_t . If the former is less persistent than the latter, $c_t - w_t$ will also be less persistent than $d_t - p_t$ and will forecast returns over shorter horizons than $d_t - p_t$.

4 Long-Horizon Forecasting Regressions

We now turn to forecasting results using \widehat{cay}_t and \widehat{cdy}_t as predictive variables. Table 1 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.

How does the persistence of \widehat{cdy}_t compare to $d_t - p_t$ and \widehat{cay}_t ? Table 2 displays autocorrelation coefficients. 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

¹⁴In this illustrative example, $\rho = \rho_w$ since all wealth is invested in the stock market. More generally, we assume that the steady state stock market dividend-price ratio and consumption-aggregate wealth ratio are sufficiently similar so that $\rho \approx \rho_w$ and this simple example captures the essence of the problem.

coefficients are 0.48 and 0.55, respectively; their second order autocorrelation coefficients are 0.13 and 0.22 respectively.

To better understand the time-series properties of $d_t - p_t$, \widehat{cay}_t , and \widehat{cdy}_t in our annual data set, 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 vectorautoregressions (VAR) for d_t and p_t , for c_t , a_t and y_t , and for c_t , d_t , and y_t .¹⁵ 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 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)). 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, 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).

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 dividend growth: \widehat{cdy}_{t-1} is both a statistically significant and economically important predictor of next year's dividend growth, Δd_t . This finding—in conjunction with the evidence that the cointegrating residual \widehat{cdy}_{t-1} is a persistent variable—implies that although there is

¹⁵The 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).

some short-run predictability in the growth of consumption and labor income (as exhibited by the dependence of these variables on lagged growth rates), it is dividend growth that exhibits error-correction behavior and therefore predictability over long horizons. Indeed, regressions of long-horizon growth rates on lagged growth rates and \widehat{cdy}_{t-1} (not shown) demonstrate that dividend growth is substantially more forecastable than consumption or labor income growth as the horizon over which these variables are measured increases. Thus, the significant relation between \widehat{cdy}_{t-1} and future dividend growth displayed in Table 3 means 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. Dividends are mean reverting and adapt over long-horizons to match the smoothness in consumption and labor income.

Figure 1 plots the demeaned values of \widehat{cdy}_t and \widehat{cay}_t over the period spanning 1948 to 2001. The figure shows that the two 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. According to the results just presented, the current levels of \widehat{cdy}_t are forecasting a increase in future dividend growth rates. By contrast, \widehat{cay}_t had been forecasting a sharp decrease in stock market returns as of 1999, but has now been largely restored to its long-run mean after the broad market declines of 2000 and 2001.

A more direct way to understand mean reversion in dividend growth is to investigate regressions of long-horizon dividend growth onto the cointegrating residual \widehat{cdy}_{t-1} . The theory behind (3) and (4) makes clear that both the log consumption-wealth ratio and the log consumption-dividend ratio should track longer-term tendencies in asset markets rather than provide accurate short-term forecasts of booms or crashes.¹⁶ We focus in this paper on explaining the historical behavior of excess stock market returns and dividend growth. Table

¹⁶The forecasting tests in this paper are intentionally in-sample tests. This makes our results directly comparable with existing literature which has investigated the predictability of long-horizon dividend growth using in-sample regressions. In addition, both the theoretical framework presented above, and the hypothesis of cointegration imply long-horizon predictability, but not stable predictive power over short subsamples of a finite data set. Inoue and Kilian (2002) show that the subsample analysis inherent in out-of-sample forecasting tests causes them to be significantly less powerful than in-sample forecasting tests, a pattern that would be exacerbated in an investigation of long-horizon forecasting power. See Lettau and Ludvigson (2001b) for further discussion on the relative merits of in-sample and out-of-sample forecasting tests using consumption-based present value relations.

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, $r_{t+1} - r_{f,t+1} + \dots + r_{t+H} - r_{f,t+H}$, where $r_{f,t}$ is used to denote the Treasury bill rate, or “risk-free” rate. The independent variable is either $d_t - p_t$, \widehat{cay}_t , or \widehat{cdy}_t . The table reports the estimated regression coefficient, the adjusted R^2 statistic in square brackets, and a heteroskedasticity and autocorrelation-consistent t -statistic for the hypothesis that the regression coefficient is zero in parentheses. The table also reports, in curly brackets, the rescaled t -statistic recommended by Valkanov (2001) for the hypothesis that the regression coefficient is zero. We discuss this rescaled statistic below. Table 5 presents the same output for predicting long-horizon CRSP dividend growth, $\Delta d_{t+1} + \dots + \Delta d_{t+H}$.

The first row of Table 4 shows that the log dividend-price ratio has little power for forecast aggregate stock market returns from one to 6 years in this sample, according to conventional statistical analysis. 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 It should be noted, however, that although recent data has weakened the statistical evidence in favor of predictability by the dividend-price ratio according to conventional statistical analyses, Lewellen (2001) finds that the dividend-price ratio remains a strong predictor of excess stock returns even in samples that include recent data. Noting that the dividend-price ratio is very persistent, Lewellen incorporates the information conveyed by the sample autocorrela-

¹⁷Campbell (2001) reports that $d_t - p_t$ explains about 40 percent of the variation in excess stock returns at a 16 quarter horizon in data from 1947 to 1998. His results differ from those reported in Table 4 for two reasons. First, the unusual movements in the dividend-price ratio over just the last three years of our sample have had a dramatic impact on the statistical evidence on predictability by the log dividend-price ratio. Second, Campbell performs long-horizon regressions using quarterly data, rather than annual as we use here. Quarterly and annual regressions of long-horizon returns differ in the number of overlapping residuals relative to the number of observations. Valkanov (2001) shows that the finite-sample distributions of R^2 statistics in long-horizon regressions do not converge to their population values when there are overlapping residuals, with the degree of divergence dependent on the amount of overlap. This finding implies that different data sets (e.g., quarterly versus annual) are likely to generate different distributions for this statistic. We address this difficulty below by using a vector autoregressive approach to impute long-horizon R^2 statistics.

tion of the dividend-price ratio when assessing its predictive power. Episodes in which the dividend-yield remains deviated from its long-run mean for an extended period of time do not necessarily constitute evidence against predictability over very long-horizons when the forecasting variable is extremely persistent. Similar results are reported in recent work by Campbell and Goto (2002), who find evidence of return predictability by financial ratios if one is willing to rule out an explosive root in the ratios. These results corroborate previous findings that the dividend yield captures very slow-moving forecasts of excess returns, but has considerably more difficulty predicting medium-horizon returns, over business cycle frequencies.

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 \bar{R}^2 of 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 the estimated present value relation, \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, the 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 \bar{R}^2 statistic rises from .20 at a three year horizon to .32 at a six year horizon. Thus, the results suggest that both \widehat{cay}_t and \widehat{cdy}_t capture some component of time-varying expected returns.

Table 5 displays results from the same forecasting exercise for long horizon dividend growth. 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 (2001) 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 . Previous research 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 evidence presented above is the first, to our knowledge, to suggest that there is important predictability of dividend growth over long horizons in direct long-horizon regressions. Researchers, cited in the introduction, have found that dividend growth predictability—if evident at all in these regressions—occurs at relatively short horizons. More recently, Ang (2002) performs direct long-horizon regressions using annual data from 1927-2000 and finds some evidence that the dividend-price ratio predicts dividend growth at a one-year horizon, but no evidence that it does so at longer horizons. These findings are consistent with those of the earlier papers cited in the introduction.

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 , which is a strong predictor of stock market dividend growth, is also a predictor of excess stock market returns. Does either variable have independent predictive power for excess returns and dividend growth? To address this question, Table 6 presents the results of multivariate regressions of long-horizon returns and dividend growth using both \widehat{cay}_t and \widehat{cdy}_t as regressors. The first panel of Table 6 shows that, in forecasting long horizon excess returns, \widehat{cay}_t drives out \widehat{cdy}_t , implying that the latter contains no information about future returns that is independent of that contained in \widehat{cay}_t . \widehat{cay}_t does contain some information for future returns that is independent of that contained in \widehat{cdy}_t , suggesting the existence a component of time-varying expected returns that is not associated with rational forecasts of dividend growth.

The second panel of Table 6 shows that 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. For horizons between 2 and 3 years, the information contained in \widehat{cay}_t and \widehat{cdy}_t is sufficiently similar that the regression has difficulty distinguishing their independent effects (although \widehat{cdy}_t is statistically significant at the 6 percent level). The two variables are strongly jointly significant as a predictor of long-horizon dividend growth (the p -value for the F -test is less than 0.000001), a finding that is consistent with results discussed below which suggest that much of the variation in expected dividend growth is common to variation in expected returns. In addition, for every

forecasting horizon in excess of one year, the regression results imply that \widehat{cay}_t contains no information about future dividend growth that is independent of that already contained in \widehat{cdy}_t .

Comparing adjusted R^2 statistics in rows 2 and 3 of Table 4, it is evident that fluctuations in \widehat{cdy}_t account for about half of the variation in future excess returns that is captured by \widehat{cay}_t at a three year horizon, and more than half at a six year horizon. Similarly, fluctuations in \widehat{cay}_t explain about half of variation in future dividend growth that is captured by \widehat{cdy}_t at four and five year horizons. Thus, these statistics are also consistent with the hypothesis that expected dividend growth and expected returns are positively correlated, and that a significant fraction of their movement is common. Section 5, below, pursues two ways of quantifying this common movement.

4.1 Additional Statistical Tests

There are at least two potential econometric hazards with interpreting the long-horizon regression results just presented. 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 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}_t are persistent variables, which, although predetermined, are not exogenous. This lack of exogeneity can create a small sample bias in the regression coefficient that works in the direction of indicating predictability even where none is present (Nelson and Kim (1993) and Stambaugh (1999)).

To address these potential inference problems, we perform three robustness checks. The first is to compute the rescaled t/\sqrt{T} statistic (where T is the sample size), recommended by Valkanov (2001). In contrast to the standard t -statistic, this rescaled t has a well defined distribution in finite samples that is straightforward to simulate, once two nuisance parameters have been estimated. Second, we use vector autoregressions to impute long-horizon R^2 statistics rather than estimating them directly from long-horizon regressions. Third, we perform bootstrapped estimates of the empirical distribution of the predictive regression coefficients and adjusted R^2 statistics under the null of no predictability; these bootstrapped estimates

are reported in Table 8 and discussed below. Both the Valkanov rescaled t -statistic and the bootstrapping methodology provide distinct approximations to the true finite sample distribution of the relevant statistics. Since the two approximations are unlikely to be the identical in any given sample, we look at the complete picture and perform both exercises.

The rescaled t -statistics 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 . After obtaining estimates of the nuisance parameters for our application, we compute the rescaled t -statistic and use Valkanov's critical values to determine statistical significance; the table reports the rescaled t -statistic 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. The analysis suggests that the forecasting power of \widehat{cay}_t and \widehat{cdy}_t for excess stock market returns and stock market dividend growth cannot 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 autoregressions 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. Hodrick (1992) also compares the asymptotic distributions of these statistics to their empirical distributions under the null and under alternatives, and finds that they have good finite sample properties. We present the results from using this methodology in Table 7, which investigates the long-horizon predictive power of both \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 7 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 R^2 statistics for forecasting dividend growth with \widehat{cdy}_t (row 3) peaks around three years, according to this metric, equal to 0.20. In addition, the implied R^2 statistics are considerably larger using lagged \widehat{cdy}_t as a predictive variable for both excess returns and dividend growth (rows 3 and 6) than when a simple autoregressive specification for returns or dividend growth is employed (rows 1 and 4). 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. We present the results of doing so in Table 8. The methodology is based on bootstrap simulation under the null hypothesis that expected excess returns and dividend growth are constant (i.e., the coefficient on \widehat{cay}_t and \widehat{cdy}_t are restricted to be zero in the excess return and dividend growth equations), and a first-order autoregressive model for the predictive variables \widehat{cay}_t and \widehat{cdy}_t .¹⁸ Artificial sequences of excess returns and dividend growth are generated by drawing randomly (with replacement) from the sample residual pairs.¹⁹ The simulations were repeated 10,000 times. 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 8 summarizes the estimated sampling distribution for the slope coefficient and the R^2 statistic in univariate regressions of annual excess returns for the CRSP index on lagged \widehat{cay}_t and \widehat{cdy}_t (top panel), as well as univariate regressions of annual dividend growth for the CRSP index on lagged \widehat{cay}_t and \widehat{cdy}_t (bottom panel). In almost every case, the estimated regression coefficient and R^2 statistic lies outside of the 95 percent confidence interval based on the empirical distribution. In most cases they lie outside of the 99 percent confidence

¹⁸It is known that the standard bootstrap is not consistent if the data series have a near-unit root. Although \widehat{cay}_t and \widehat{cdy}_t are relatively persistent variables (see Table 1), they 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.

¹⁹Nelson 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.

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 strong predictive power for returns at horizons longer than one year. For all of the other regressions and forecasting horizons, however, we find that the slope coefficients and R^2 statistics are large relative to their sampling distributions. For both the dividend growth and excess return forecasting regressions, the estimated regression coefficients and R^2 statistics are large compared with their empirical values generated under the null of no predictability. These results, like those discussed above using the rescaled t -statistic and VAR-imputed R^2 statistics, imply that the predictability of excess returns and dividend growth documented here is unlikely to be attributable to small sample biases in the regression coefficients or R^2 statistics.

4.2 Interpretation

The empirical results presented above indicate that fluctuations in the log consumption-aggregate dividend ratio are informative about the future path of dividend growth. One theoretical framework that is consistent with these findings is the rational expectations framework presented above and in Lettau and Ludvigson (2001a). If investors want relatively flat consumption paths over time, forward looking agents will attempt to smooth out transitory fluctuations in both wealth and total income from aggregate wealth, or aggregate dividends. The present-value relation cay_t will be high when agents are expecting high future returns, and low when they are expecting low future returns. Similarly, the present value relation cdy_t will be high when agents are expecting high future dividend growth, and low when agents are expecting low dividend growth. The finding that returns and dividend growth are forecastable does not mean that agents have perfect foresight, or even that investors can foresee a large fraction of the variation in future returns and future dividend growth. But they are consistent with some ability on the part of stockholders to gauge the future path of returns and dividend growth. Accordingly, the findings suggest a somewhat different perspective on investor behavior than that found in recent theoretical treatments which explain the variability in the dividend-price ratio by assuming that investors make systematic errors in forecasting dividend growth (e.g., Barsky and De Long (1993)).

Although rational, forward looking behavior is consistent with the findings presented

above, explanations for these phenomena which involve some form of irrational behavior or irrational expectations may also be consistent. One possibility, raised by the work of Campbell and Mankiw (1989), is that stockholders follow a simple rule-of-thumb in which they consume their total income every period. If income itself is close to a random walk, such a rule could imply that consumption and income could define the trend in wealth (implying that the consumption-wealth ratio forecasts returns) without investors possessing any real foresight. The evidence presented here, however, exposes the potential limitations of such an explanation for stockholder behavior. Although *labor* income, the dividend to human capital, may be close to a random walk, dividends from stock market wealth have an important transitory component, yet consumption tracks just the long-term components in equity dividends. Any explanation of stockholders' behavior based on a simple rule of thumb of this type could only be consistent with these findings by sheer coincidence, since such a rule must consistently make no important systematic errors in forecasting dividend growth. This observation does not, of course, rule out an explanation based on heterogeneity in stock market participation, a possibility raised by the work of Mankiw and Zeldes (1991), Vissing-Jørgensen (1999), Vissing-Jørgensen (2002), and others. For example, nonstockholders may follow a simple Campbell-Mankiw rule of thumb while stockholders are forward-looking and make an optimal consumption choice. Such an explanation is in fact very close in spirit to the original Campbell and Mankiw (1989) framework. Lacking a long time-series on the consumption of stockholders, however, such a hypothesis may be difficult to test empirically.

5 The Comovement in Expected Returns and Expected Dividend Growth

This section pursues two ways of quantifying the comovement between expected excess returns and expected dividend growth, as measured by the state variables, or driving processes, cay_t and cdy_t , respectively. We begin by investigating a simple principal components framework for the bivariate system $(\widehat{cay}_t, \widehat{cdy}_t)'$. We use a standard eigenvalue decomposition to identify principal components in this bivariate system (see Theil (1971)). Figure 2 plots the first principal component, along with our estimate of cdy_t . The first principal component is almost identical to \widehat{cdy}_t itself. The behavior of these two variables is well-described by the first principal component, which explains 80.2 percent of the variation in the system $(\widehat{cay}_t, \widehat{cdy}_t)'$. Thus, a large fraction of the variation in expected returns and expected divi-

dividend growth is indeed common, and the common component is very well captured by \widehat{cdy}_t . Moreover, the finding that the first principal component is almost identical to \widehat{cdy}_t suggests that virtually all of the variation in expected dividend growth is common to variation in expected returns. The first principal component is less highly correlated with \widehat{cay}_t , suggesting that some variation in expected returns is independent of that in expected dividend growth. In short, most of the variation in rational forecasts of long-horizon dividend growth is common to variation in rational forecasts of long-horizon returns, but not the other way around. Given that common variation in expected returns and expected dividend growth has offsetting effects on the log dividend price ratio, the findings are consistent with the well-documented result that the log dividend-price ratio contains significantly more information about future long-horizon returns than it does about long horizon dividend growth.

The second way we quantify the common variation in \widehat{cay}_t and \widehat{cdy}_t is to employ the frequency domain measure of comovement developed by Croux, Forni, and Reichlin (2002). Croux, Forni and Reichlin call this measure of comovement “dynamic correlation” when applied to two series. We refer the reader to that paper for details and merely outline the approach here. Dynamic correlation measures how important are cycles of different frequencies in accounting for the comovement between series. Here we use dynamic correlation to quantify the comovement between our empirical state variables for expected returns and expected dividend growth. The dynamic correlation statistic is defined

$$\rho_{xy}(\lambda) = \frac{C_{xy}(\lambda)}{\sqrt{S_x(\lambda)S_y(\lambda)}},$$

where λ is the frequency $-\pi \leq \lambda \leq \pi$, C_{xy} is the co-spectrum, and S_x and S_y are the spectral density functions of x and y respectively.²⁰

Figure 2 displays four plots. The first three are the individual spectra for the dividend-price ratio, \widehat{cay}_t and \widehat{cdy}_t . Recall that a white noise series has a flat spectrum. Thus the very steep spectrum for the dividend-price ratio indicates that the low frequency components of $d_t - p_t$ are by far the most important determinants of its sample variance, so much so that the variable appears nonstationary in this sample. By contrast, the low frequency components of \widehat{cay}_t and \widehat{cdy}_t are far less important; the spectra for these two variables are much flatter and they take on much lower values at low frequencies than does the spectrum for the log dividend-price ratio. These results are not surprising given the autocorrelation coefficients reported in Table 2.

²⁰The co-spectrum is the real part of the cross-spectral density between two series.

The fourth panel of Figure 2 displays the dynamic correlations between each series. The panel shows that \widehat{cay}_t and \widehat{cdy}_t are most highly correlated at periods ranging from 2 to 4 years; they have a correlation of almost 70 percent at frequencies corresponding to 2 years and a correlation in excess of 50 percent at horizons less than 4 years. By contrast, \widehat{cay}_t is most highly correlated with $d_t - p_t$ at low frequencies, when it is *least* correlated with \widehat{cdy}_t . In addition, although \widehat{cdy}_t is not highly correlated with $d_t - p_t$ at any frequency, it is least correlated with $d_t - p_t$ at low frequencies, where most of the sampling variation in the dividend-price ratio occurs.

Taken together, these results imply that the low-frequency variation in $d_t - p_t$ that so dominates its sampling variation is better described as capturing slow-moving forecasts of excess returns than time-varying expected dividend growth. Variation in expected dividend growth is substantial but occurs at higher frequencies than those over which the $d_t - p_t$ ratio predominately varies. This is consistent with the hypothesis that the common variation in expected dividend growth and expected returns occurs at business cycle frequencies, those corresponding to periods of a few years. The dividend-price ratio is most highly correlated with the consumption-wealth ratio at low frequencies, when the comovement between expected returns and expected dividend growth is least pronounced. The common components analysis suggests that variation in expected dividend growth is almost all common to variation in expected returns, and therefore has no effect on the dividend-price ratio.

6 Implications for the Equity Premium Puzzle

There has been a growing interest in equilibrium asset pricing models that explicitly connect consumption to dividends.²¹ One reason for this interest is the recognition that autonomous variation in dividends and consumption offer one way of rationalizing the observed equity premium, since those models can in principal generate volatile stock returns without substantial variation in consumption or real interest rates. Bansal and Yaron (2000) and Campbell (2001) consider models of this form. Neither of these models, however, allow aggregate consumption and dividends to be cointegrated. In this section we consider how implications for the equity premium in models with autonomous dividend variability are affected by allowing consumption and measures of aggregate dividends to be cointegrated.

The framework above implies that aggregate consumption can be thought of as the divi-

²¹See, for example, Abel (1999), Bansal and Yaron (2000), Bansal, Dittmar, and Lundblad (2001), Campbell (1986), and Campbell (2001).

dividend paid on all invested wealth, $d_{w,t}$. Many authors have assumed that the aggregate stock market is a good proxy for aggregate wealth and thus set consumption equal to the stock market dividend, d_t (e.g., Grossman and Shiller (1981), Lucas (1978); Mehra and Prescott (1985)). Alternatively, Campbell (1986) and Abel (1999) assume that dividends on equity equal aggregate consumption raised to a power λ , so that, in logs, $d_t = \lambda c_t$. Abel (1999) shows that λ can be thought of as a measure of leverage, with levered equity represented by values of $\lambda \geq 1$. Values for λ greater than one make dividend growth more volatile than consumption growth and therefore help resolve the equity premium puzzle. We follow Abel (1999) and assume that $\lambda \geq 1$. Campbell (2001) generalizes this set-up to allow for separate variation in dividends and consumption by specifying a simple bivariate model taking the form

$$\begin{aligned}\Delta c_{t+1} &= \epsilon_{t+1} \\ \Delta d_{t+1} &= \lambda \epsilon_{t+1} + \eta_{t+1},\end{aligned}\tag{12}$$

where η_{t+1} is an autonomous component in dividends, assumed to be uncorrelated with ϵ_{t+1} . In this model, there is no cointegration between c_t and d_t .

The purpose of the calculations in this section is to examine the role of cointegration in affecting the equity premium. Such a specification makes sense if there is a saving technology so that agents are not forced to consume their endowment (dividend) every period. In this case, consumption and aggregate dividends will be cointegrated and move together over the long term, but will not necessarily be equal in every period.

To build intuition, it is useful to begin by considering a simple example in which stock market wealth is equal to aggregate wealth, implying $w_t = p_t$, and aggregate dividends, $d_{w,t}$, are equal to stock market dividends, d_t . In this case, the share of stock market wealth in aggregate wealth is unity, implying that cay_t in (3) collapses to $c_t - w_t = c_t - p_t$, and cdy_t in (4) collapses to $c_t - d_t$. Below we relax this simplifying assumption and consider a case in which c_t is cointegrated with d_t and y_t , consistent with empirical evidence reported above.

We may allow for cointegration between c_t and d_t by modifying (12) so that consumption and stock market dividends have the following error-correction representation

$$\begin{aligned}\Delta c_{t+1} &= \epsilon_{t+1} \\ \Delta d_{t+1} &= \chi(c_t - d_t) + \lambda \epsilon_{t+1} + \eta_{t+1}.\end{aligned}\tag{13}$$

Equation (13) generalizes equation (12). The last two terms in (13) are the same as in (12) and account for the sensitivity of dividends to innovations in consumption, allowing

for some independent variation between consumption and dividends. The first term arises because consumption and dividends are cointegrated; χ , which we will refer to as the *adjustment* parameter, measures the effect on dividend growth of last period's cointegrating error, $(c_t - d_t)$. Notice that the budget constraint analysis discussed above implies dividends can no longer be a levered version of the *level* of consumption, $d_t = \lambda c_t$, $\lambda > 1$, otherwise consumption would eventually become a negligible fraction of dividends. Instead, dividends may respond to a consumption *innovations*, ϵ_{t+1} , in a levered fashion, but cointegration requires the volatility of dividend growth to be identical to that of consumption growth when measured over sufficiently long horizons.

In this simple bivariate model, χ tells us how quickly dividends adjust to restore the consumption-aggregate dividend ratio to its long-run mean, subsequent to an equilibrium-distorting shock, or "cointegration gap." Higher values of χ correspond to swifter adjustments and less persistent variation in $(c_t - d_t)$. Equation (13) collapses to (12) in the special case where dividends and consumption are not cointegrated, which in this simple example occurs as $\chi \rightarrow 0$. This example assumes that consumption growth is unforecastable, so that deviations from the common trend in consumption and dividends are eventually eliminated by a subsequent movement in dividend growth, consistent with the empirical findings presented above.

How does cointegration between consumption and dividends affect risk premia in this simple example? To address this question, we follow classic papers on the equity risk premium and assume there is a representative agent that maximizes a time-separable constant relative risk aversion (CRRA) utility function defined over aggregate consumption

$$U(C_t) = \frac{C_t^{1-\gamma} - 1}{1-\gamma},$$

where γ is the coefficient of relative risk aversion. For expositional convenience, we assume that expected asset returns are constant and aggregate consumption and asset returns are jointly lognormal and homoskedastic. The expected log equity risk premium (adjusted for a Jensen's inequality term) is equal to

$$E_t r_{t+1} - r_{f,t+1} + \sigma_r^2/2 = \gamma \text{Cov}(r_{t+1}, \Delta c_{t+1}). \quad (14)$$

In addition, Campbell (1991) shows that asset return surprises are given by the following expression, up to a first-order approximation

$$r_{t+1} - E_t r_{t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j}$$

If expected returns are constant, return surprises are entirely attributable to revisions in expectations about future dividend growth. In this case, the risk-premium of the consumption claim is

$$E_t r_{t+1} - r_{f,t+1} + \sigma_r^2/2 = \gamma \sigma_\epsilon^2. \quad (15)$$

To derive the risk-premium of the dividend claim, use the cointegrated system (13). The unexpected return of the dividend claim reduces to

$$r_{t+1} - E_t r_{t+1} = \frac{\lambda(1-\rho) + \rho\chi}{1-\rho + \rho\chi} \epsilon_{t+1} + \frac{1-\rho}{1-\rho + \rho\chi} \eta_{t+1}.$$

Defining $\Gamma \equiv \frac{\lambda(1-\rho) + \rho\chi}{1-\rho + \rho\chi}$, equation (14) implies that the risk premium on the dividend claim is

$$E_t r_{t+1} - r_{f,t+1} + \sigma_r^2/2 = \gamma \Gamma \sigma_\epsilon^2. \quad (16)$$

We may use (16) to illustrate how the equity premium is affected by cointegration. Notice that Γ is increasing in the leverage parameter, λ , but decreasing in the adjustment parameter χ . In the model of equation (13), only dividends participate in the adjustment needed to restore the consumption-dividend ratio to its long-run mean, subsequent to an innovation in $c_t - d_t$. The higher is χ , the more quickly this adjustment takes place and the less persistent is $c_t - d_t$. If dividends are not cointegrated with consumption, the adjustment never takes place, and $\chi = 0$. In this case, $\Gamma = \lambda$ and the model collapses to the simple levered models of the type considered by Campbell (1986) and Abel (1999). By contrast, if consumption and dividends are cointegrated and dividends fully adjust in the period immediately following an innovation in the consumption-dividend ratio, $\chi = 1$ and $\Gamma = \lambda(1-\rho) + \rho \approx 1$, since ρ is close to one. In this case the equity risk-premium of the dividend claim is almost identical to that of the consumption claim, and the presence of leverage does not help resolve the equity premium puzzle. In summary, relative to the risk premium of the consumption claim, the risk premium of the dividend claim is decreasing in χ , and is therefore largest when there is no cointegration ($\chi = 0$).

To get a feel for the magnitude of this effect on the equity premium, consider the following parameterization: $\gamma = 50$, $\sigma_\epsilon = 1.1\%$ (from Table 1), $\lambda = 10$ and $\rho = 0.95$. The risk premium of a dividend claim that is not cointegrated with consumption ($\chi = 0$) is 6.05%, in line with post-war U.S. data. If dividends adjust very quickly to close the cointegration gap, the risk premium on the dividend claim decreases dramatically. For $\chi = 1$ one obtains a risk premium of only 0.605%. For more interesting intermediate values, the effect of cointegration is still important even if the adjustment parameter is substantially smaller; for example, if $\chi = 0.5$, the risk premium is only 1.2%.

The value for risk-aversion required to match the empirical equity premium in this parameterization is considerably larger than that required in Abel (1999) to match the equity premium, as is the number for λ , which he estimates is consistent with the relative volatilities of consumption and dividends when it takes the value $\lambda = 2.74$. The reason for these differences is that Abel calibrates his model to data used in Mehra and Prescott (1985), which includes the significantly more volatile pre-war consumption data and smaller equity risk premium. Those values, as reported by Mehra and Prescott (1985) in data from 1889-1978, are $\sigma_\epsilon = 3.6\%$ and $E_t r_{t+1} - r_{f,t+1} = 4.48\%$. Using this figure for σ_ϵ , Abel obtains a risk-premium of 3.9% for the log-linear approximate version of his model by setting $\lambda = 2.74$, and relative risk aversion equal to 11.048. The same value for the equity risk-premium (3.92%) may be obtained in the model of this section by setting $\chi = 0$ (so that the framework collapses to the simple levered case without cointegration), and by plugging the values $\lambda = 2.74$, $\gamma = 11.048$, and $\sigma_\epsilon = 3.6\%$ into (16).

So far we have considered a simple example in which consumption is cointegrated with stock market dividends. The analysis above, however, suggests not that consumption and stock market dividends are cointegrated, but that consumption and measures of aggregate dividends, or c_t , d_t , and labor income y_t , are cointegrated. It is straightforward to generalize the analysis further to allow for cointegration among consumption, dividends and labor income, as in (4). For example, we might modify (13) so that c_t , a_t , and y_t have the following error-correction specification

$$\begin{aligned}\Delta c_{t+1} &= \epsilon_{t+1} \\ \Delta y_{t+1} &= \alpha \epsilon_{t+1} + u_{t+1} \\ \Delta d_{t+1} &= \chi \widehat{cdy}_t + \lambda \epsilon_{t+1} + \pi u_{t+1} + \eta_{t+1}.\end{aligned}\tag{17}$$

The set of equations in (17) generalizes those in (13) and is directly analogous. Requiring $\alpha > 1$ is analogous to requiring $\lambda > 1$, so that α can be thought of as a “leverage” parameter for the dividend to human capital, y_t . Similarly, the autonomous innovation u_{t+1} is analogous to the autonomous innovation η_{t+1} . The cointegrating residual is now $\widehat{cdy}_t \equiv c_t - \widehat{\beta}_d d_t - \widehat{\beta}_y y_t$, rather than $c_t - d_t$. The parameter χ tells us how quickly dividends adjust to restore \widehat{cdy}_t to its long-run mean; higher values of χ again correspond to swifter adjustments and less persistent variation in \widehat{cdy}_t . Moreover, this case is comparable to the previous example in that stock market dividends and human capital dividends may behave in a levered fashion in response to innovations in consumption growth, but the relative volatility of the growth rates in c_t , d_t , and y_t is governed in the long-run by their common trends, which insure that the precise

linear combination of variables given by \widehat{cdy}_t is stationary. Finally, this example models log consumption and log labor income as random walks, thereby capturing the empirical finding documented above that only stock market dividends, Δd_{t+1} , are significantly related to \widehat{cdy}_t . Although the data suggest that log consumption and log labor income are not exactly random walks, the empirical findings presented above suggest that (17) is likely to be a reasonable approximation of the vector time-series properties for these variables.

Under the assumptions about preferences and the stochastic properties of returns and consumption made above, the risk-premium on the consumption claim is again given by (15). To derive the risk-premium of the dividend claim, use the cointegrated system (17). In this case, the unexpected return of the dividend claim reduces to

$$r_{t+1} - E_t r_{t+1} = \frac{\lambda(1-\rho) + \rho\chi(1-\widehat{\beta}_y\alpha)}{1-\rho + \rho\widehat{\beta}_d\chi} \epsilon_{t+1} + \frac{\pi(1-\rho) - \rho\chi\beta_y}{1-\rho + \rho\widehat{\beta}_d\chi} u_{t+1} + \frac{1-\rho}{1-\rho + \rho\widehat{\beta}_d\chi} \eta_{t+1}.$$

Defining $\Psi \equiv \frac{\lambda(1-\rho) + \rho\chi(1-\widehat{\beta}_y\alpha)}{1-\rho + \rho\widehat{\beta}_d\chi}$, equation (14) implies that the risk premium on the dividend claim is

$$E_t r_{t+1} - r_{f,t+1} + \sigma_r^2/2 = \gamma\Psi\sigma_\epsilon^2.$$

The parameter Ψ collapses to Γ when the cointegrating coefficient on labor income, $\widehat{\beta}_y$, is zero, and the cointegrating coefficient on stock market dividends, $\widehat{\beta}_d$, is one.

We may again gain an understanding of the magnitude of this effect on the equity premium by considering the parameterization used in Abel (1999): $\gamma = 11.048$, $\sigma_\epsilon = 3.6\%$, $\lambda = 2.74$ and $\rho = 0.95$. We also set $\widehat{\beta}_y = 0.68$ and $\widehat{\beta}_d = 0.13$ to match the empirical estimates reported above. The parameter α may be calibrated by matching the covariance between Δc_t and Δy_t implied by (17) with that from the US data displayed in Table 1. The covariance implied by (17) is

$$\text{Cov}(\Delta c_t, \Delta y_t) = \alpha\sigma_\epsilon^2,$$

where σ_ϵ^2 is the variance of consumption growth. Given the estimates of the standard deviations and correlation of Δc_t and Δy_t , a value of $\alpha = 1.26$ is obtained.

Under these parameter values, the risk-premium of the consumption claim is 1.43%. When there is no cointegration ($\chi \rightarrow 0$), the risk-premium of the dividend claim again collapses to that implied by Abel's specification and the presence of leverage generates the much larger 3.92% premium. By contrast, when cointegration among c_t , d_t , and y_t is present and $\chi = 2.4$ (the estimated value reported in Table 3), the risk-premium of the dividend claim is only slightly larger than that of the consumption claim, equal to 1.91%. This occurs because an innovation in consumption growth causes \widehat{cdy}_t to deviate from its mean,

and therefore necessitates a subsequent adjustment in dividends in order to restore \widehat{cdy}_t to its mean. Such an adjustment reduces the covariance between news about *future* dividend growth and today's consumption innovation, thereby reducing the equity premium. To obtain an equity premium of 3.92% in this setting (when $\chi = 2.4$), an unrealistically large value for the stock market leverage parameter is required, equal to $\lambda = 13$.

Of course, the model above could account for a higher equity risk-premium without leverage if it were modified to allow for catching up with the Joneses preferences, as in Abel (1999). Abel emphasizes, however, that one cannot rely on catching up with the Joneses preferences alone to match both the equity premium and the volatility of the short-term riskless rate. To rationalize both phenomena, leverage must be sufficiently high, since only leverage increases the equity premium implied by the model without increasing the volatility of the short-term riskless rate. The calculations in this section suggest that the presence of cointegration among consumption, stock market dividends and labor income, once calibrated to match the estimates of the cointegrating relation we document here, can significantly limit the scope for leverage to increase the equity risk-premium implied by the model. Even if deviations from the common trend in consumption, dividends and labor income are quite persistent and dividends are considerably more volatile than consumption in the short-run, the dividend stream may not be significantly more risky than the consumption stream if all three variables are tied together in the long-run. In the simple asset pricing model studied above, even large values of leverage are insufficient to explain the equity risk premium when dividends, consumption and labor income are cointegrated.

7 Why Might Expected Dividend Growth Covary with Expected Returns?

Why might expected dividend growth be positively correlated with expected Returns? Some authors have noted that the present value model for dividends and prices, (5), together with the hypothesis that managers smooth dividends, can be used to interpret the empirical behavior of the dividend-price ratio. Such explanations imply that expected dividend growth is roughly constant. For example, Cochrane (1994) argues that if managers perfectly smooth dividends by setting them equal to the discounted value of future earnings (discounted at the risk-free rate), dividends will follow a random walk, explaining why dividend growth is close to unforecastable by the log dividend-price ratio. The investigation of this paper finds that

dividend growth is not unforecastable, but is instead predictable over horizons of several years, horizons shorter than those corresponding to frequencies over which the persistent dividend-price ratio primarily varies. Of course, dividends are smoother than earnings, and the dividend payout ratio has some forecasting power for future dividends (Lamont (1998)), consistent with the hypothesis that managers do some dividend smoothing. But the results presented suggest that dividend smoothing by managers may be imperfect.

If investors themselves desire smooth consumption paths, why don't managers perfectly smooth dividend payments? One possibility is that although dividend-smoothing may be possible over long horizons (implying that dividend growth may be unforecastable by variables such as the dividend-price ratio whose sampling variation is dominated by low-frequency fluctuation), it may be more difficult over horizons corresponding to business cycle fluctuations. 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 vary with macroeconomic conditions: they are lower during a slow-down in economic growth and higher during periods of economic expansion. Others have documented that the relative cost of obtaining external finance rises during an economic slow-down, suggesting that managers who need to finance long-term projects have a greater need to retain earnings in recessions than in expansions. Bernanke and Gertler (1989) and Bernanke, Gertler, and Gilchrist (1996) present theoretical and empirical evidence of countercyclical variation in the external finance premium. The equity risk-premium also appears counter-cyclical: it rises during an 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. According to this hypothesis, consumers may be better able to smooth transitory fluctuations in earnings than managers, implying that earnings growth should be predictably higher when, according to \widehat{cdy}_t , dividend growth and excess stock returns are predictably higher.

Table 9 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: 1964-2000.²² Table 9 presents long-horizon

²²We use Lewellen's data and not earnings per share since that measure is contaminated by variability in share issuance.

forecasts of earnings growth using \widehat{cdy}_t as a predictive variable. Earnings growth is predictably higher when predictable dividend growth, as captured by \widehat{cdy}_t , is higher. The present-value relation, \widehat{cdy}_t , is strongly statistically significant as a predictor of earnings growth, 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. The results suggest that when consumption is high relative to its common trend with stock market dividends and labor income, both dividends and earnings are temporarily low and forecast to grow more quickly in the future. These findings are consistent with the hypothesis that consumers are better able to smooth consumption than managers are able to smooth dividends through business cycles.

8 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. We find that aggregate consumption, stock market dividends and labor income are cointegrated, and that deviations from their common trend are typically restored by a forecastable movement in stock market dividend growth. 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. 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 have several implications for asset pricing. First, they imply that the log dividend-price ratio will have difficulty predicting both dividend growth and excess returns over anything but extremely long horizons precisely because expected excess returns fluctuate and covary with expected dividend growth at those horizons. Thus, the results suggest that an important component of time-varying expected returns and time-varying expected dividend growth is not captured by the log dividend-price ratio, or likely by other aggregate financial ratios. This stacks the deck against such financial ratios in statistical tests of return or dividend growth predictability. Second, the findings 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. Third, autonomous variation in dividends as a levered version of consumption is unlikely to provide a complete resolution of the equity premium puzzle. If consumption, stock market dividends and labor income are cointegrated, as the evidence presented here suggests, the dividend claim may not be significantly more risky than the consumption claim even for large values of leverage.

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, however, 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.

Appendix A: Derivation of Consumption-Based Loglinear Present Value Expressions

This appendix derives the consumption-based present value expression for dividend growth, cdy_t . First, note that cay_t is derived by expressing the market value of human capital in terms of future expected returns to human wealth and future dividend flows from human wealth, namely labor income (Lettau and Ludvigson (2001a)). This allows a reformulation of 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 y_t denotes log labor income. Denote the net return to non-human capital $R_{a,t}$ and the net return to human capital $R_{h,t}$, and assume that human capital, H_t , takes the form, $H_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^j (1 + R_{h,t+i})^{-i} Y_{t+j}$, so that labor income is the dividend to human wealth. A log-linear approximation of H_t yields

$$h_t = \kappa_H + y_t + E_t \sum_{j=1}^{\infty} \rho_h^j (\Delta y_{t+j} - r_{h,t+j}) \quad (18)$$

where κ_H is a constant and $\rho_h \equiv 1/(1 + \exp(\overline{y} - \overline{h}))$. We follow Campbell and Mankiw (1989) and assume that $\rho_h = \rho_w \equiv 1/(1 + \exp(\overline{c} - \overline{w}))$. (The equations below can easily be extended to relax this assumption but nothing substantive is gained by doing so since the generalization would merely yield more complicated of future expected returns and future expected income flows from wealth.) An approximation for log aggregate wealth, W_t , as a function of its component elements, A_t and H_t is given by

$$w_t \approx \omega a_t + (1 - \omega) h_t. \quad (19)$$

Combining (2), (18) and (19) furnishes an approximate equation for the log consumption-aggregate wealth ratio given in equation (3).

To derive cdy_t , we express the market value of stock market wealth in terms of future expected returns to stock market wealth and future dividend flows from stock market wealth. We do the same for all nonstock wealth. The present value relation for stock market wealth, S_t ,

$$S_t = E_t \sum_{j=0}^{\infty} \prod_{i=0}^j (1 + R_{t+i})^{-i} D_{t+j} \quad (20)$$

is analogous to the more familiar present value relation for prices and dividends:

$$P_{s,t} = E_t \sum_{j=1}^{\infty} \prod_{i=1}^j (1 + R_{s,t+i})^{-i} D_{s,t+j}, \quad (21)$$

where $P_{s,t}$ and $D_{s,t}$ in (21) denote price and dividends per share, respectively. However, there are two differences. The first is that (20) uses total market value, S_t , and total dividends, rather price

per share, P_t , and dividends per share, as in (21). The second is that we use a different timing convention in order to be consistent with the timing convention for aggregate wealth in (1). Since wealth is measured at the beginning of the period, asset values are cum-dividend values, rather than the typically assumed ex-dividend value associated with measuring price at the end of the period. To obtain a dividend-based consumption ratio, we suppose total wealth consists of N_t shares, and let $P_t \equiv P_{s,t}N_t$ be the ex-dividend value of an asset making a dividend payment D_t :

$$\begin{aligned} S_t &= (P_t + D_t) \\ &= N_t(P_{s,t} + D_{s,t}). \end{aligned}$$

Then the return on wealth can be written

$$R_{t+1} = (P_{t+1} + D_{t+1}) / P_t,$$

and $S_t = P_t + D_t$ is the cum-dividend stock market value at time t . Thus the expression for the return can be written in terms of cum-dividend total value:

$$S_{t+1} = R_{t+1} (S_t - D_t),$$

which is in the same form as (1) and can be linearized in the same way. We obtain a loglinear approximate expression

$$s_t = \kappa_s + d_t + E_t \sum_{j=1}^{\infty} \rho_w^j (\Delta d_{t+j} - r_{t+j}), \quad (22)$$

where κ_s is a constant. The same derivation can be applied to all nonstock forms of aggregate wealth (primarily human capital) denoted \widetilde{W}_t , where \widetilde{W}_t includes human capital and nonstock forms of net worth. In the text we conserve on notation by using Y_t to denote both the income flows from human capital and the income flows from all nonstock forms of net worth, since the latter is primarily human capital. To be precise, we now use \widetilde{Y}_{t+1} for the latter, implying that

$$\widetilde{W}_{t+1} = \widetilde{R}_{t+1} (\widetilde{W}_t - Y_t),$$

and a loglinear approximate expression for \widetilde{W}_t takes the form

$$\widetilde{w}_t = \kappa_{\widetilde{w}} + \widetilde{y}_t + E_t \sum_{j=1}^{\infty} \rho_w^j (\Delta \widetilde{y}_{t+j} - \widetilde{r}_{t+j}), \quad (23)$$

where \widetilde{r}_t is the log return to \widetilde{W}_t . Implicitly we are assuming that the mean dividend-price ratio for each form of aggregate wealth equals the mean consumption-wealth ratio, since the same parameter

ρ_w appears in (18), (22) and (23). An approximation for log aggregate wealth, W_t , as a function of its component elements, S_t and \widetilde{W}_t is given by

$$w_t \approx \nu s_t + (1 - \nu) \widetilde{w}_t, \quad (24)$$

where ν is the average share of stock market wealth in aggregate wealth. Combining (23), (22), (24) and (2), we obtain (4).

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 single 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

- ABEL, A. B. (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33.
- 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.
- BANSAL, R., R. F. DITTMAR, AND C. T. LUNDBLAD (2001): "Consumption, Dividends, and the Cross-Section of Equity Returns," Unpublished paper, Fuqua School of Business, Duke University.
- BANSAL, R., AND A. YARON (2000): "Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles," Unpublished paper, Fuqua School, Duke University.
- BARSKY, R. B., AND B. DE LONG (1993): "Why Does the Stock Market Fluctuate?," *Quarterly Journal of Economics*, 108(2), 291–311.
- BERNANKE, B. S., AND M. GERTLER (1989): "Agency Costs, Net Worth and Business Cycle Fluctuations," *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. (1986): "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics*, 101, 785–804.
- (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.
- (2001): "Consumption-Based Asset Pricing," Unpublished paper, Harvard University.
- CAMPBELL, J. Y., AND M. GOTO (2002): "Efficient Tests of Stock Return Predictability," Unpublished paper, Harvard University.
- 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 University Press, London, UK.
- 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.
- CROUX, C., M. FORNI, AND L. REICHLIN (2002): "A Measure of Comovement for Economic Variables: Theory and Empirics," *The Review of Economics and Statistics*, forthcoming.
- 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.
- GROSSMAN, S. J., AND R. J. SHILLER (1981): "The Determinants of the Variability of Stock Market Prices," *American Economic Review*, 71, 222–227.
- 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.
- (2001c): "Understanding Trend and Cycle in Asset Values: Bulls, Bears, and the Wealth Effect on Consumption," Unpublished Paper, New York University.
- LEWELLEN, J. W. (2001): "Predicting Returns With Financial Ratios," Unpublished paper, MIT Sloan School of Management.
- LUCAS, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1446.
- MANKIW, N. G., AND S. P. ZELDES (1991): "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, 29(1), 97–112.

- MEHRA, R., AND E. PRESCOTT (1985): "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- MILLER, M., AND F. MODIGLIANI (1961): "Dividend Policy, Growth, and the Valuation of Shares," *Journal of Business*, 36(4), 411–433.
- NELSON, C. C., AND M. J. KIM (1993): "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance*, 43, 641–661.
- NEWKEY, 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.
- 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.
- THEIL, H. (1971): *Principles of Econometrics*. John Wiley and Sons, New York, NY.
- VALKANOV, R. (2001): "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics*, forthcoming.
- VISSING-JØRGENSEN, A. (1999): "Limited Stock Market Participation and the Equity Premium Puzzle," Unpublished manuscript, University of Chicago.
- (2002): "Limited Asset Market Participation and Intertemporal Substitution," *Journal of Political Economy*, 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} (<i>t</i> -stat)	-0.194 (-1.059)		0.364 (1.352)
Δp_{t-1} (<i>t</i> -stat)	-0.192 (-1.441)		-0.210 (-1.079)
$d_{t-1} - p_{t-1}$ (<i>t</i> -stat)	0.103 (2.205)		0.070 (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} (<i>t</i> -stat)	0.267 (1.279)	0.449 (1.220)	-0.523 (-0.696)
Δy_{t-1} (<i>t</i> -stat)	-0.039 (-0.294)	-0.148 (-0.641)	0.433 (0.916)
Δa_{t-1} (<i>t</i> -stat)	0.112 (2.777)	0.128 (1.794)	0.392 (2.702)
\widehat{cay}_{t-1} (<i>t</i> -stat)	-0.007 (-0.053)	0.102 (0.457)	1.726 (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} (<i>t</i> -stat)	0.469 (2.284)	0.652 (1.869)	-0.136 (-0.060)
Δy_{t-1} (<i>t</i> -stat)	-0.074 (-0.572)	-0.156 (-0.709)	-0.252 (-0.176)
Δd_{t-1} (<i>t</i> -stat)	0.029 (2.311)	0.052 (2.389)	-0.129 (-0.917)
\widehat{cdy}_{t-1} (<i>t</i> -stat)	-0.038 (-0.408)	0.219 (1.377)	2.400 (2.314)
\bar{R}^2	0.179	0.098	0.104

Notes: The table reports estimated coefficients from cointegrated first-order vector autoregressions of the column variable on the row variable; c_t is log consumption, y_t is log labor income, a_t is log asset wealth (net worth), d_t is log stock market dividends, and p_t is the log CRSP value-weighted price index. *t*-statistics are reported in parentheses. Estimated coefficients that are significant at the 5% level are highlighted in 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: $\sum_{i=1}^h r_{t+i} = k + \gamma z_t + \epsilon_{t,t+h}$						
Horizon h (in years)						
$z_t =$	1	2	3	4	5	6
$d_t - p_t$	0.14 (1.90) [0.08]	0.24 (1.40) [0.10]	0.27 (1.21) [0.10]	0.34 (0.73) [0.10]	0.52 (0.84) [0.16]	0.73 (1.12) [0.23]
\widehat{cay}_t	6.48 (4.19) {0.57***} [0.27]	11.78 (5.42) {0.74***} [0.49]	13.23 (5.42) {0.74***} [0.46]	13.62 (5.27) {0.72***} [0.37]	16.81 (7.07) {0.96***} [0.39]	21.94 (5.46) {0.74***} [0.52]
\widehat{cdy}_t	1.32 (1.47) {0.20} [0.01]	5.21 (7.38) {1.00***} [0.16]	6.11 (4.13) {0.56***} [0.20]	6.77 (4.28) {0.58***} [0.20]	18.09 (4.92) {0.67***} [0.20]	11.40 (4.45) {0.61***} [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: $d_{t+h} - d_t = k + \gamma z_t + \epsilon_{t,t+h}$						
	Horizon h (in years)					
$z_t =$	1	2	3	4	5	6
$d_t - p_t$	0.09 (2.94) [0.07]	0.18 (2.11) [0.15]	0.19 (2.70) [0.13]	0.23 (2.27) [0.14]	0.29 (2.70) [0.15]	0.34 (2.41) [0.19]
\widehat{cay}_t	4.74 (6.26) {0.85***} [0.29]	5.89 (4.86) {0.66***} [0.30]	4.90 (3.33) {0.45***} [0.16]	4.30 (2.80) {0.38***} [0.09]	5.13 (2.17) {0.30*} [0.10]	5.72 (1.50) {0.20} [0.12]
\widehat{cdy}_t	2.74 (4.06) {0.55***} [0.20]	3.95 (5.84) {0.79***} [0.24]	3.65 (4.13) {0.56***} [0.20]	3.99 (3.60) {0.49***} [0.20]	5.24 (5.38) {0.73***} [0.28]	6.13 (3.65) {0.50***} [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

Variables	Horizon h (in years)					
	1	2	3	4	5	6
<i>h</i> -period regression: excess stock returns						
\widehat{cay}_t	6.94 (3.27)	11.15 (3.73)	12.22 (3.82)	11.95 (3.43)	15.33 (3.86)	18.47 (4.12)
\widehat{cdy}_t	-0.74 (-0.81)	0.89 (0.69)	1.14 (0.80)	1.70 (1.02)	1.46 (0.76)	3.44 (1.81)
	[0.27]	[0.48]	[0.45]	[0.36]	[0.38]	[0.53]
<i>h</i> -period regression: dividend growth						
\widehat{cay}_t	3.71 (3.08)	4.27 (1.95)	2.62 (1.09)	0.62 (0.31)	-0.30 (-0.14)	-0.82 (-0.24)
\widehat{cdy}_t	1.64 (2.57)	2.29 (1.86)	2.58 (1.87)	3.72 (2.74)	5.37 (4.78)	6.48 (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: 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 8: Bootstrap Confidence Intervals of Slope and R^2

x_t	$\widehat{\beta}$	95% CI	99% CI	R^2	95% CI	99% CI
$r_{t+1} = \alpha + \beta x_t + u_{t+1}; \quad x_{t+1} = \mu + \phi x_{t-1} + v_{t+1}$						
\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)
$\Delta d_{t+1} = \alpha + \beta x_t + u_{t+1}; \quad x_{t+1} = \mu + \phi x_{t-1} + v_{t+1}$						
\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)

Notes: This tables reports confidence intervals from a bootstrap procedure. 10,000 artificial time series are generated under the null hypothesis $\beta = 0$ by drawing (with replacement) from the residuals of the system estimated under the null hypothesis. The columns denoted $\widehat{\beta}$ and R^2 report data estimates using annual data from 1948 to 2001.

Table 9: 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 table reports results from h -period regression of earnings growth: $e_{t+h} - e_t = k + \beta \widehat{cdy}_t + \epsilon_{t,t+h}$. The earnings data are from Lewellen (2001). For each regression, the table reports OLS estimates of the regressors, Newey 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

Variables	L-max Test			Trace Test			<i>t</i> -Test		
	$H_0 : r =$	0	1	2	$H_0 : r =$	0	1	2	$H_0 : \text{no cointegration}$
<i>10% Critical Values</i>		10.60	2.71		13.31	2.71			-2.60
<i>5% Critical Values</i>		14.07	3.76		15.41	3.76			-2.93
<i>d, p</i>		6.06	4.56		10.62	4.56			-0.47
<i>10% Critical Values</i>		13.39	10.60	2.71	26.70	13.31	2.71		-3.52
<i>5% Critical Values</i>		20.97	14.07	3.76	29.68	15.41	3.76		-3.80
<i>c, a, y</i>		25.34	6.57	0.07	31.98	6.64	0.07		-4.13
<i>c, d, y</i>		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





Figure 2

First Principal Component and cdy



First Principal Component and cay






Figure 3


A: Spectrum of cdy



B: Spectrum of cay



C: Spectrum of d-p



D: Dynamic Correlation

