

Interpreting Risk Premia Across Size, Value, and Industry Portfolios

Ravi Bansal*

Fuqua School of Business,
Duke University

Robert F. Dittmar

Kelley School of Business,
Indiana University

Christian T. Lundblad

Kelley School of Business,
Indiana University

First Draft: July 2002

This Draft: December 2002

*Corresponding author. Author contact information: Bansal (email: ravi.bansal@duke.edu, tel: (919) 660-7758), Dittmar (email: rdittmar@indiana.edu, tel: (812) 855-2694), and Lundblad (email: clundbla@indiana.edu, tel: (812) 855-9714). The usual disclaimer applies.

Abstract

In this paper, we model dividend and consumption growth rates as a vector-autoregression (VAR), from which we measure the long-run response of dividend growth rates to consumption shocks. We find that this long-run cash flow beta can justify well over 50% of the difference in risk premia across size, book-to-market, and industry sorted portfolios. Interestingly, the long-run cash flow betas explain about 50% of the dispersion in the standard CAPM-based portfolio betas for these assets. Our economic model highlights the reasons for the failure of the market beta to justify the cross-section of risk premia. The market beta is itself a weighted combination of cash flow betas and additional priced sources of risk. Each risk source's beta may be significant; however, a weighted combination of the betas may not be significant in explaining the cross-section of risk premia, as each source of risk carries a distinct price. Our results indicate that differences in exposure of cash-flows to aggregate economic fluctuations as captured by aggregate consumption movements contain very valuable information regarding differences in risk premia. In all, our results indicate that the size, book-to-market and industry spreads are not puzzling from the perspective of economic models.

1 Introduction

The focus of this paper is to characterize the systematic sources of priced risks in the cross-section of returns from the perspective of general equilibrium models. The empirical work of Hansen and Singleton (1982, 1983) underscores the importance of consumption risks in understanding risk premia. A consistent implication of these consumption based models is that the link between cash flows and aggregate consumption is a key input in determining an asset's exposure to and compensation for risk. Our approach emphasizes the long-run links between cash flows and consumption, and shows that this relation is empirically important for interpreting risk premia. In addition, we also focus on understanding why the commonly used market based CAPM has difficulties in justifying the cross-section of risk premia across assets.

We concentrate on characterizing the sources of risk inherent in size, book-to-market, and industry sorted portfolios. These portfolios have been at the center of the asset pricing literature over the past two decades. These sorts produce economically meaningful risk premia; from 1949 through 2001, size sorted decile portfolios generate premia of 4.14% per annum, book-to-market sorted portfolios generate premia of 6.07% per annum, and industry groupings produce a spread of 3.62% per annum. As the empirical literature has shown, the return premia of these dimensions pose a considerable challenge to economic models.

We explore the sources of these differences in average returns by examining the implications of a general economic model. In this model, returns are assumed to be generated by realized shocks to current and expected future cash flow growth. Further, asset cash flows are explicitly linked to the dynamics of aggregate consumption. In this setting, we show that differences in the long-run response of cash flows to a unit consumption shock (i.e., the cash flow beta) should explain cross-sectional variation in risk premia. When we additionally allow risk premia to fluctuate, the cash flow beta can be augmented by the traditional market beta in the determination of variation in risk premia across assets. This version of our economic model also highlights some of the reasons why the usual market beta of an asset may fail to capture differences in risk premia across assets.

A key dimension of this paper is the measurement of the cash flow beta. We model the consumption and dividend growth rate dynamics as a vector autoregression (VAR). The long-run cash flow beta for a given asset can be obtained from this VAR as the long-run

response of cash flow growth to a unit shock in consumption. The first paper to focus on the empirical measurement of cash flow betas, Bansal, Dittmar, and Lundblad (2001), argues that covariation between dividend growth rates and consumption at long lags provides sharp information regarding risk premia on assets. In contrast to their paper, we provide the joint transition dynamics of cash flows and consumption in order to measure the cash flow betas. Our setup permits the explicit analytical expression of the cash flow beta, which improves the precision of the associated parameter estimates. Additionally, we examine the links between cash flow betas and market betas, and analyze the reasons for the failure of standard market betas to capture risk premia across assets. Finally, we incorporate industry portfolios in our analysis, which pose their own unique empirical challenges as documented in Fama and French (1997).

As predicted by the theory, we find that the price of risk associated with the cash flow betas is highly significant and positive. To confirm our statistical inference, we conduct Monte Carlo experiments to examine the finite sample distribution of the price of risk and the cross-sectional R^2 . This finite sample distribution accounts for estimation error in the VAR dynamics of consumption and dividend growth. The point estimate in annual data for the price of cash flow beta risk is 0.357 (S.E. 0.177), with an adjusted cross-sectional R^2 of 57%. Our cash-flow beta's do well on all three sorts; their correlation with the mean returns is 0.79 for book-to-market, 0.81 for size, and 0.57 for industry portfolios. Evidence from quarterly data corroborates our annual results, with the quarterly price of risk and explanatory power comparable to that found in the annual data. To conduct statistical inference we also provide finite sample distributions for the cross-sectional risk prices and the cross-sectional R^2 for the quarterly data.

We further present a model based on Epstein and Zin (1989) preferences, similar to that developed in Bansal and Yaron (2002). This model highlights the conditions under which the cash-flow betas will explain the cross-section of risk premia. Further, it also provides insights into the failure of the market betas to capture cross-sectional risk premia. In this model, asset returns are driven both by cash flow news and changing risk premia; the risk premium fluctuates due to changes in aggregate economic uncertainty (i.e., consumption volatility). The result is that the cross-section of risk premia is determined both by an asset's cash flow beta and its beta with respect to news about aggregate risk premia. The standard market beta is a weighted combination of these different betas, where each of these sources of risk bears a different price. Consequently, the market beta may fail to explain the cross-section of

risk premia. As predicted by this argument, there is a meaningful relation between the cash flow beta and the market beta. The correlation between these two quantities is 0.65. The message implied by this evidence is that the cash flow beta is an important source of risk in isolation, and explains a considerable degree of the cross-sectional variation in observed risk premia.

In all, our empirical evidence indicates that the exposure of dividends to movements in the aggregate economy, as measured by consumption, contains very valuable information regarding differences in risk premia across assets. Dividend streams that have larger exposure to aggregate consumption news also offer higher risk premia. The work of Lettau and Ludvigson (2001) and Jagannathan and Wang (1996) highlight alternative channels for explaining differences in risk premia across assets. Our work augments the understanding of the determinants of risk premia by focusing on the links between cash flows and consumption.

The remainder of this paper is organized as follows. In section 2, we discuss the model for cash flow betas when discount rates are constant. Our strategy for estimating these betas is discussed in section 3. Section 4 discusses the empirical evidence. We analyze the economic implications of our framework in section 5. Section 6 provides concluding remarks.

2 Cash flow Betas

In this section, we provide the arguments that motivate our cash flow beta. For any asset i , consider the Campbell and Shiller (1988) linear approximation for the log return:

$$r_{i,t} = \kappa_{i,0} + g_{i,t} + \kappa_{i,1}pd_{i,t} - pd_{i,t-1} \quad (1)$$

where $pd_{i,t}$ is the log price dividend ratio, $g_{i,t}$ the log dividend growth rate, and $r_{i,t}$ the log return ($\kappa_{i,0}$ and $\kappa_{i,1}$ are parameters in the linearization). Under this approximation (1), one can derive the following present value implication for the log price dividend ratio

$$pd_{i,t} = \frac{\kappa_{i,0}}{(1 - \kappa_{i,1})} + E_t\left[\sum_{j=1}^{\infty} \kappa_{i,1}^j g_{i,t+j} - \sum_{j=1}^{\infty} \kappa_{i,1}^j r_{i,t+j}\right] \quad (2)$$

Further, if we assume that expected returns are constant through time, the return innovation can be expressed as follows:

$$r_{i,t} - E_{t-1}[r_{i,t}] \equiv e_{r_{i,t}} = g_{i,t} - E_{t-1}[g_{i,t}] + E_t\left[\sum_{j=1}^{\infty} \kappa_{i,1}^j g_{i,t+j}\right] - E_{t-1}\left[\sum_{j=1}^{\infty} \kappa_{i,1}^j g_{i,t+j}\right] \quad (3)$$

Note that the case where both expected returns and expected growth rates for cash flow can vary is considered in section 5.

2.1 Cash Flow Dynamics

To determine the long-run dividend exposures to consumption shocks, we first must characterize the dynamic processes for consumption and dividends. Log consumption growth, $g_{c,t}$, is assumed to follow an $AR(J)$ process

$$g_{c,t} = \mu_c + \sum_{j=1}^J \rho_{c,j} g_{c,t-j} + \eta_{c,t}, \quad (4)$$

and (log) dividend growth rates follow

$$\begin{aligned} g_{i,t} &= \mu_i + \sum_{k=1}^{k=K} \gamma_{i,k} g_{t-k} + u_{i,t} \\ u_{i,t} &= \sum_{j=1}^L \rho_{j,i} u_{i,t-j} + b_i \eta_{c,t} + \zeta_{i,t} \end{aligned} \quad (5)$$

where $\zeta_{i,t}$ is uncorrelated with consumption innovations as stated above. Without loss of generality assume that $K \geq J$.

To characterize the evolution of the system, let $1 + (K + L) = q$. The $q \times 1$ vector z_t is

$$z'_t = [g_{i,t} \ u_{i,t} \cdots u_{i,t-(L-1)} \ g_{c,t} \cdots g_{c,t-(K-1)}] \quad (6)$$

The dynamics of consumption and dividend growth can then be expressed as

$$z_t = \mu + Az_{t-1} + Gu_t \quad (7)$$

where A and G are $q \times q$ matrices. Note that consumption feeds into the future dynamics of dividends, but dividends do not feed back into consumption. The $q \times 1$ vector u_t has its first elements as $\zeta_{i,t}$ and its last element as $\eta_{c,t}$; all other elements of u_t are zero.

To allow for the effect of κ_1 , we define the matrix A_κ as $\kappa_1 A$. From equation (3), it follows that $e_{r_i,t}$ is the first element of the matrix

$$[I + \sum_{j=1}^{\infty} A_\kappa^j] G u_t = [I - A_\kappa]^{-1} G u_t \quad (8)$$

The cash flow beta, $\beta_{i,t}$, equals the first element of $[I - A_\kappa]^{-1} G \iota$, where ι has an element one corresponding to the consumption innovation and zero elsewhere. Note that the return innovation is

$$e_{r_i,t} = \beta_{i,d} \eta_{c,t} + \zeta_{i,t};$$

where $\beta_{i,d} \eta_{c,t}$ is the return response to aggregate consumption news and $\zeta_{i,t}$ represents the cash flow news specific to the asset. Note also that $\zeta_{i,t}$ and $\eta_{c,t}$ are uncorrelated. $\beta_{i,d}$ is determined by the reaction of the infinite sum of dividend growth rates to consumption news; that is, the accumulated impulse response of dividend growth rates to a unit consumption shock. We call $\beta_{i,d}$ the long-run cash flow beta. In other words, this beta provides the response of the present value of future dividend growth to a unit consumption shock.

To gain some intuition into what this risk measure captures, note that the long-run cash flow consumption beta with $K = L = J = 1$ is

$$\beta_{i,d} = \frac{\kappa_{i,1} \gamma_{i,1}}{1 - \kappa_{i,1} \rho_{c,1}} + \frac{b_i}{1 - \kappa_{i,1} \rho_{i,1}} \quad (9)$$

which reflects both the contemporaneous correlation between dividend and consumption shocks, b_i , and the effect of current consumption growth on future dividends, γ_i . In general, the cash flow beta for asset i will be

$$\beta_{i,d} = \frac{\sum_k \kappa_{i,1}^k \gamma_{i,k}}{1 - \sum_j \kappa_{i,1}^j \rho_{c,j}} + \frac{b_i}{1 - \sum_l \kappa_{i,1}^l \rho_{i,l}} \quad (10)$$

When equality is imposed ($\gamma_{i,k} = \gamma_i$), then $\sum_k \kappa_{i,1}^k \gamma_{i,k} = \gamma_i \sum_k \kappa_{i,1}^k$. This expression measures the average covariance between dividend growth and the lagged, K -period growth rate of consumption.

Next, we explore the ability of the estimated long-run cash flow beta to explain the cross-sectional variation in observed average returns for market capitalization, book-to-market ratio, and industry sorted portfolios (30 portfolios in all). In section 5, we provide detailed economic motivation for why the long-run cash flow beta should explain the cross-sectional differences in risk premia. This motivation leads to the specification

$$R_{i,t} = \lambda_0 + \lambda_c \beta_{i,d} + v_{i,t} \quad (11)$$

In equation (11), $R_{i,t}$ are the observed returns for asset i . The cross-sectional price of risk parameters λ_0 and λ_c , as shown in section 5, are determined by preference parameters. The above equation imposes the restriction that the differences in average returns across assets reflect differences only in $\beta_{i,d}$. Also, we will subsequently explore the pricing implications of the long-run beta in a model that facilitates a more general preference specification and a time-varying cost of capital.

3 Estimation

To explore the long-run relationship between consumption and dividend growth, we first estimate the dynamic processes described for consumption and dividend growth rates. Note that in estimation we remove the unconditional mean from all the dividend growth rate and consumption growth rate series and use these demeaned series in estimating the dynamics of consumption and dividend growth rates. We use GMM for estimation, and consider the following set of moment conditions for estimation. First, the consumption dynamics can be estimated using the moment conditions:

$$E[\mathbf{g}_{0,t}] = E[\eta_{c,t} g_{c,t-j}] = 0 \quad (12)$$

for $j = 1 \cdots J$. This expression gives us J moment conditions associated with estimating the consumption dynamics. We estimate the dividend growth dynamics with the following moment restrictions:

$$E[\mathbf{g}_{1i,t}] = \begin{pmatrix} E[u_{i,t} g_{c,t-k}] \\ E[u_{i,t-l} \zeta_{i,t}] \\ E[\eta_{c,t} \zeta_{i,t}] \end{pmatrix} = \mathbf{0} \quad (13)$$

for $k = 1 \cdots K$, and $l = 1 \cdots L$. The last moment condition estimates b_i . This expression yields $(K + L + 1)$ moment conditions for each dividend growth under consideration, and J moment conditions associated with estimating the consumption growth dynamics. For N assets we consequently have $J + N(K + L + 1)$ moment conditions and the same number of parameters. For the annual data we will set $J = K = 1$ and $L = 2$. In addition, we also consider the cross-sectional restrictions

$$E[\mathbf{g}_{2,t}] = \begin{pmatrix} \sum_i E[R_{i,t} - (\lambda_0 + \lambda_c \beta_{i,d})] \\ \sum_i E[(R_{i,t} - (\lambda_0 + \lambda_c \beta_{i,d})) \beta_{i,d}] \end{pmatrix} = \mathbf{0} \quad (14)$$

The final two moment conditions ensure an exactly-identified system where the GMM based estimates for the relevant risk prices, λ_0 and λ_c , are equivalent to those obtained under ordinary least squares. Taken together, this yields $2 + J + N(K + L + 1)$ parameters, and the same number of orthogonality conditions.

With 30 assets and 4 parameters to characterize the dividend growth rates, the dimension of the optimal GMM weight matrix would be at least 120×120 , which is impossible to estimate given the number of time-series observations. In practice, since the joint optimal GMM weighting matrix becomes too large, we utilize the following weighting matrix for the calculation of standard errors:

$$\mathbf{W}^{-1} = \begin{pmatrix} E[\mathbf{g}_{0,t} \mathbf{g}'_{0,t}] & 0 & \cdots & \cdots & 0 \\ \mathbf{0} & (E[\mathbf{g}_{1i,t} \mathbf{g}'_{1i,t}]) & \cdots & \cdots & \mathbf{0} \\ \vdots & \cdots & \ddots & \cdots & \vdots \\ \mathbf{0} & \cdots & \cdots & (E[\mathbf{g}_{1N,t} \mathbf{g}'_{1N,t}]) & \mathbf{0} \\ \mathbf{0} & \cdots & \cdots & \mathbf{0} & E[\mathbf{g}_{2,t} \mathbf{g}'_{2,t}] \end{pmatrix} \quad (15)$$

That is, the weighting matrix is a block-diagonal matrix of the covariance of the moment conditions. The resulting weighting matrix is HAC-adjusted following Newey and West (1987). It is important to note that the standard errors on the time-series parameters for a given (univariate) dividend growth rate utilize the full GMM weight matrix—and hence are quite reasonable. The system associated with the estimating the risk prices is exactly identified; that is, the point estimates correspond to the OLS estimates. However, the standard errors for the risk prices, that is λ_0 and λ_c , do not take account of the error in estimating the time-series parameters that go into the construction of the cash-flow betas. For this reason we also report the monte carlo-based finite sample distribution for the risk

prices and the cross-sectional R^2 that takes account of the estimation error of all the time series and cross-sectional parameters for all assets at the same time. The details of this monte carlo are provided in section 4.

3.1 Data

3.1.1 Aggregate Cash Flows and Factors

Our empirical exercise is conducted on data sampled at both the annual and quarterly frequency. We collect (at both frequencies) seasonally adjusted real per capita consumption of nondurables plus services data from the NIPA tables available from the Bureau of Economic Analysis. Also, to convert returns and other nominal quantities, we also take the associated personal consumption expenditures (PCE) deflator from the NIPA tables. The mean of the annual real consumption growth rate series over the period spanning 1949 through 2001 is 0.0212 with standard deviation of 0.0114, and the mean of the inflation series is 0.0354 per annum with a standard deviation of 0.0254. Quarterly figures are comparable. For subsequent analysis, we also measure the aggregate market portfolio return as the return on the CRSP value-weighted index of stocks.

3.1.2 Portfolio Menu

We consider portfolios formed on firms' market value, book-to-market ratio, and industry classification. Our rationale for examining portfolios sorted on these characteristics is that size and book-to-market based sorts are the basis for the factor model examined in Fama and French (1993). Additionally, industry sorted portfolios have posed a particularly challenging feature from the perspective of systematic risk measurement (see Fama and French (1997)). We focus on one-dimensional sorts on these characteristics as this procedure typically results in over 150 firms in each decile portfolio which facilitates a more accurate measurement of the consumption exposure of dividends; it is important to limit the portfolio specific variation in dividend growth rates, and a larger number of firms in a given portfolio helps achieve this.

Market Capitalization Portfolios

We form a set of value-weighted portfolios on the basis of market capitalization. The

set of all firms covered by CRSP are ranked on the basis of their market capitalization at the end of June of each year using NYSE capitalization breakpoints. In Table 1, we present means and standard deviations of market value-weighted returns for size quintile portfolios. The table displays a significant size premium over the post-war sample period; the mean real return on the lowest quintile firms is 13.59% per annum, contrasted with a return of 9.45% per annum for the highest quintile. The means and standard deviations of these portfolios are similar to those reported in previous work.

Book-to-Market Portfolios

Book values are constructed from Compustat data. The book-to-market ratio at year t is computed as the ratio of book value at fiscal year end $t - 1$ to CRSP market value of equity at calendar year $t - 1$.¹ All firms with Compustat book values covered in CRSP are ranked on the basis of their book-to-market ratios at the end of June of each year using NYSE book-to-market breakpoints. Sample statistics for these data are also presented in Table 1. The highest book-to-market firms earn average real returns of 15.14% per annum, whereas the lowest book-to-market firms average 9.07% per annum.

Industry Portfolios

Value-weighted industry portfolios are formed by sorting NYSE, AMEX, and NASDAQ firms by their CRSP SIC Code at the beginning of each month. Industry definitions follow those in Fama and French (1997). We specifically utilize definitions for ten industries: i1, Nondurable Goods, i2, Durable Goods, i3, Manufacturing, i4, Energy, i5, Chemicals, i6, Telecommunications, i7, Utilities, i8, Wholesale, Retail, and Services, i9, Financial, and i10, Other.² Sample statistics for these data are also presented in Table 1. The mean real returns range from 8.74% for the Financial industry to 12.36% for Nondurables.

3.2 Portfolio Cash Dividends

To measure the long-run cash flow beta, we also need to extract the cash dividend payments associated with each portfolio discussed in the previous section. Our construction of the

¹We thank Ken French for providing us the value-weighted book-to-market-sorted portfolio data. For a detailed discussion of the formation of the book-to-market variable, refer to Fama and French (1993).

²Industry definitions follow those provided by Kenneth French at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

dividend series is the same as that in Campbell (2000). Let the total return per dollar invested be

$$R_{t+1} = h_{t+1} + y_{t+1}$$

where h_{t+1} is the price appreciation and y_{t+1} the dividend yield (i.e., dividends at date $t + 1$ per dollar invested at date t). More clearly stated, h_{t+1} represents the ratio of the per dollar value of the portfolio at time $t + 1$ to time t , $\frac{V_{t+1}}{V_t}$, and y_{t+1} represents the per dollar dividends paid by the portfolio at time $t + 1$ divided by per dollar value at time t , $\frac{D_{t+1}}{V_t}$. We directly observe both R_{t+1} and the price gain series h_{t+1} for each portfolio; hence, we construct the dividend yield as $y_{t+1} = R_{t+1} - h_{t+1}$.³ The level of the cash dividends we employ in the paper is extracted as follows

$$D_{t+1} = y_{t+1}V_t$$

where

$$V_{t+1} = h_{t+1}V_t$$

with $V_0 = 100$. Hence, the dividend series that we use, D_t , corresponds to the total dividends given out by a mutual fund at t that extracts the dividends and reinvest the capital gains. The ex-dividend value of the mutual fund is V_t and the per dollar total return for the investors in the mutual fund is

$$R_{t+1} = \frac{V_{t+1} + D_{t+1}}{V_t} = h_{t+1} + y_{t+1}$$

which is precisely the CRSP total return for each portfolio.

We construct the level of cash dividends, D_t , for the size, book-to-market, and industry portfolios on a monthly basis. From this series, we construct the annual and quarterly levels of dividends by summing the dividends *within* the period under consideration. As the dividend yields still have strong seasonalities at the quarterly frequency, we also employ a trailing four quarter average of the quarterly dividends to construct the deseasonalized quarterly dividend series. This procedure is consistent with the approach in Hodrick (1992), Heaton (1993), and Bollerslev and Hodrick (1995). The annual cash dividends, however, are accumulated over the full twelve months, and are then free from this seasonality issue thus requiring no further adjustments. These series are converted to real by the personal consumption deflator. Log growth rates are constructed by taking the log first difference of the cash dividend series. Statistics for the annual dividend growth rates of the portfolios

³We thank Ken French for providing both the total return and price appreciation series for the book-to-market portfolios.

under consideration are presented in Table 1.

4 Empirical Evidence

For the purposes of estimation, we assume (for the annual data spanning the postwar 1949-2001 period) that the log consumption growth rate, $g_{c,t}$, follows an AR(1) process, and accordingly, the log cash dividend growth rate, $g_{i,t}$, depends only upon the lagged consumption growth rate. That is, we assume $K = 1$. Further, we assume that shocks to the dividend growth rate, $u_{i,t}$, follow an AR(2) process ($L = 2$). Taken together, the dynamic process for the *demeaned* annual consumption and dividend growth rate data that we consider:

$$\begin{aligned}
g_{c,t+1} &= \rho_c g_{c,t} + \eta_{t+1} \\
g_{i,t+1} &= \gamma_i g_{c,t} + u_{i,t+1} \\
u_{i,t+1} &= b_i \eta_{t+1} + \rho_{1,i} u_{i,t} + \rho_{2,i} u_{i,t-1} + \zeta_{i,t+1} \\
\beta_{i,d} &= \frac{\kappa_{i,1} \gamma_i}{1 - \kappa_{i,1} \rho_c} + \frac{b_i}{1 - \kappa_{i,1} \rho_{1,i} - \kappa_{i,1}^2 \rho_{2,i}}
\end{aligned} \tag{16}$$

In this case, the long-run cash flow beta, $\beta_{i,d}$, is determined both by the contemporaneous covariance between the dividend and consumption shock, b_i , and the effect the consumption growth rate has upon future dividends, embodied in the coefficient γ_i ; in both cases, the autoregressive nature of the processes magnify the effects accordingly.⁴ Note, our results appear to be qualitatively robust to alternative choices for K and L .

The parameter estimates for this model are presented in Table 2 for the annual frequency. Estimates of γ_i for the characteristic-sorted portfolios are presented in Table 2 along with HAC-adjusted standard errors. As shown in the table, a clear pattern emerges in the projection of cash dividend growth rates on the lagged consumption growth rate. Sorting on market capitalization produces a pattern in γ_i . For example, the small firm portfolio displays a sensitivity to lagged consumption growth of 0.23 (S.E. 1.77) compared to -0.01 (S.E. 0.63) for the large firm portfolio. The pattern is most pronounced *within* the decile sort. Also, the book-to-market sorted portfolios produce large spreads in γ_i ; the high book-to-market

⁴Note, that $\kappa_{i,1}$ is estimated to be equivalent to $1/(1 + \exp(\overline{d-p}))$, where $(\overline{d-p})$ is the average log dividend price ratio. $\kappa_{i,1}$ is, on average, 0.964 for annual data, and 0.988 for quarterly data. Incorporating $\kappa_{i,1}$ in the calculation of the long-run beta does not materially impact our results. For example, if we assume $\kappa_{i,1}=1$ for all assets, our results are materially unchanged.

firms' sensitivity to lagged consumption growth is 4.54 (S.E. 2.59) compared to 0.48 (S.E. 2.26) for the low book-to-market firms. The pattern among industry-sorted portfolios is less identifiable. Despite strong cross-sectional significance documented below, the estimates of γ_i are associated with large standard errors.

We also present the contemporaneous covariance between the consumption and cash flow growth rate shocks, b_i , in Table 2. This parameter measures the immediate response of each asset's cash flow growth rate to an aggregate shock. As can be seen, sorting on market capitalization produces a pronounced pattern in the contemporaneous relationship between consumption and dividend shocks. For example, the small firm portfolio's estimated b_i is 5.03 (S.E. 1.87) compared to 1.56 (S.E. 0.42) for the large firm portfolio. Similarly, the book-to-market sorted portfolios produce large spreads in b_i . The estimated b_i for high book-to-market firms is 7.03 (S.E. 2.81) compared to 2.98 (S.E. 1.50) for the low book-to-market firms. This pattern is even more pronounced across all ten book-to-market sorted portfolios, with several of the low book-to-market portfolios displaying negative contemporaneous covariance between their dividend growth and aggregate shocks. Finally, here, the pattern among industry-sorted portfolios is somewhat more identifiable, with shocks to the dividend growth rate on the durable goods industry displaying the most pronounced covariance with aggregate shocks, with an estimate of 7.14 (S.E. 1.41). Note, the durable goods industry also displays the largest average return of the post-war period. Finally, unlike the projection coefficients, γ_i , the contemporaneous covariances are generally estimated with time-series precision (the standard errors presented are HAC-adjusted).

In Table 2, we also document the sum of the autoregressive coefficients for the portfolio-specific dividend growth rate shocks. Many of these coefficients are small in size and not significant, but there are some exceptions (see the first two size-sorted portfolios, for example). Additionally, the first order autocorrelation coefficient in consumption growth is estimated to be 0.33 (S.E. 0.11). Our estimates of the long-run cash flow beta (see equation (16)) will facilitate this serial correlation.

Finally, we also present the implications of the previously estimated parameters for the long-run cash flow beta, $\beta_{i,d}$, for each of the 30 portfolios. This is the key parameter of interest, as it describes each portfolio's long-run dividend response to an aggregate consumption shock. Further, according to the model presented above, this parameter is the sole measure of exposure to systematic risk which determines risk premia in the cross-section. Accordingly, we will explore the ability of the long-run cash flow beta to explain cross-sectional variation

in average returns across the 30 portfolios. As can be seen in equation (16), the long-run cash flow beta is essentially the sum of the projection coefficient describing the response of dividend growth to lagged consumption, γ_i , and the contemporaneous covariance between shocks to dividend and aggregate consumption growth, b_i , adjusted for serial correlation in each series. Empirically, the estimated long-run cash flow betas differ dramatically across the portfolios, generally in line with their observed average returns. For example, we document a large long-run cash flow beta spread in market capitalization portfolios; the $\beta_{i,d}$ for the small firm portfolio is 7.81 (S.E. 3.71), whereas the same for the large firm portfolio is only 1.59 (S.E. 1.22). The same pattern emerges for the book-to-market sorted portfolios; the estimated $\beta_{i,d}$'s for the low and high book-to-market portfolio are 4.28 (S.E. 5.24) and 15.13 (S.E. 5.65), respectively, in line with the large observed dispersion in average returns across high and low book-to-market portfolios. Finally, a less pronounced pattern emerges with the industry sorted portfolios, with the durable goods industry displaying the largest, by far, estimated long-run cash flow beta at 7.68 (S.E. 3.81). The lowest long-run cash flow beta among the industry-sorted portfolio is associated with the telecommunications industry, which displays the third lowest average return in Table 1. HAC-adjusted standard errors, computed using the delta method, demonstrate that the long-run cash flow betas are generally estimated with precision in the time-series.

4.1 The Long-run Cash Flow Beta and the Cross-section of Returns

In this section, we examine the ability of the long-run cash flow beta, $\beta_{i,d}$, to explain the cross-sectional variation of observed equity risk premia. Effectively, we perform standard cross-sectional regressions using the 30 decile portfolios (10 size, 10 book-to-market, and 10 industry). The estimated cross-sectional risk premia restriction is stated in equation (14), with λ_0 and λ_c as the cross-sectional parameters of interest, given the estimated long-run cash flow beta. Table 3 (Panel A) documents the ability of the estimated long-run cash flow betas to explain the cross-section of average returns. Our results demonstrate that the estimated price of consumption risk is both positive and significant; the OLS estimate of λ_c is 0.36, with a HAC-adjusted t -statistic of 2.02. The GMM based standard errors account for the time-series variation in measured returns. Further, the adjusted R^2 is 57%. Within portfolios sorts, this relationship holds as well; for example, the correlations between average returns

and the long-run cash flow betas are 0.79, 0.81, and 0.57 for the size, book-to-market and industry portfolio, respectively. Consistent with the large cross-sectional R^2 , the estimated long-run cash flow beta can explain a considerable portion of the cross-sectional variation in measured risk premia associated with this set of portfolios.

We also explore the cross-sectional regression for quarterly sampled data from the second quarter of 1949 through the fourth quarter of 2001. We estimate the time-series parameters as above, but the lag lengths are modified to represent the temporal dependence at the quarterly frequency; i.e. $K = 4$ and $L = 8$. Despite the alternative measurement frequency, the time-series estimates, not reported but available upon request, are qualitatively similar to those presented in Table 2 for the annual data. In Panel B of Table 3, we present the associated cross-sectional regressions for the data measured at the quarterly frequency. Generally speaking, the results are very similar to the implications of the annual data presented in Panel A. As above, the estimated price of consumption risk is significantly positive; the OLS estimate of λ_c is 0.078 (note that average returns are measured at the quarterly frequency as well), with a HAC-adjusted t -statistic of 2.22. Further, the adjusted R^2 is 50%, suggesting that the cross-sectional explanatory power of the long-run cash flow beta is not specific to measurement frequency. This stands in sharp contrast to standard factor based models, for which measurement frequency is known to be an issue.

To explore the small-sample features of our estimator, we conduct a simulation-based Monte Carlo analysis. The small sample distribution may be particularly important since the long-run cash flow beta is not always precisely measured in the time-series. For most of the portfolios, $\beta_{i,d}$ is significantly different from zero, but the projection of dividend growth on lagged consumption growth, γ_i , is generally not. Despite this issue, the cross-sectional price of consumption risk, λ_c , does appear to be estimated precisely with more than 50% of the cross-sectional dispersion in risk premia explained. Collectively, this requires more careful consideration, and in consequence, we consider an additional simulation based experiment to ensure that our results reflect the economic content of our model rather than random chance.

We conduct the following Monte Carlo experiment, in which we simulate 10,000 samples of both annual and quarterly measured aggregate consumption growth of the same size as is available in our sample (1949-2001). This experiment simulates under the alternative hypothesis that our model is incorrect. That is, we effectively assume that the price of consumption risk and the long-run cash flow beta, $\beta_{i,d}$, are zero. For the annual data,

demeaned consumption is simulated from an AR(1) process

$$\hat{g}_{c,t+1} = \hat{\rho}_c \hat{g}_{c,t} + \hat{\eta}_{c,t+1} \quad (17)$$

where $\hat{\rho}_c$ is the autoregressive parameter for consumption estimated in the data, and $\hat{\eta}_{c,t+1}$ is simulated from a normal distribution with standard deviation equal to σ_η , which corresponds to the standard deviation of the consumption growth residual in the data. The simulated consumption growth rates and demeaned *observed* dividend growth rates are used to estimate the time-series parameters in equation (16). That is, we re-estimate the long-run cash flow beta for each iteration as follows:

$$\begin{aligned} \hat{g}_{c,t+1} &= \rho_c \hat{g}_{c,t} + \hat{\eta}_{t+1} \\ g_{i,t+1} &= \gamma_i \hat{g}_{c,t} + u_{i,t+1} \\ u_{i,t+1} &= b_i \hat{\eta}_{t+1} + \rho_{1,i} u_{i,t} + \rho_{2,i} u_{i,t-1} + \zeta_{i,t+1} \\ \beta_{i,d} &= \frac{\kappa_{i,1} \gamma_i}{1 - \kappa_{i,1} \rho_c} + \frac{b_i}{1 - \kappa_{i,1} \rho_{1,i} - \kappa_{i,1}^2 \rho_{2,i}} \end{aligned} \quad (18)$$

where each portfolio's demeaned dividend growth rate, $g_{i,t}$, is the actual observed quantity for each portfolio. For each iteration, we then run the standard cross-sectional regression:

$$R_{i,t} = \lambda_0 + \lambda_c \beta_{i,d} + v_{i,t} \quad (19)$$

where $R_{i,t}$ is the *observed* real return for each portfolio. As the simulated consumption growth is independent of all the dividend growth rates, the population values of the long-run cash flow betas, $\beta_{i,d}$, are zero, and therefore the population value of λ_c is also zero. This Monte Carlo experiment provides finite sample empirical distributions for λ_c and the adjusted R^2 for the cross-sectional projection. For each iteration, we store the HAC-adjusted t -statistic and the $\bar{R}^2$. The simulation for the quarterly frequency is conducted analogously, adjusting the dynamic process we assume as in the empirical section presented above.

The results of this experiment are presented in Table 3. The distribution for the HAC-adjusted t -statistic on the estimated price of risk, λ_c , and the cross-sectional adjusted R^2 are presented in Panel A for the annual data and Panel B for the quarterly observed data. The t -statistic distribution is essentially centered at zero (the population value) for both frequencies. This evidence suggests that our point estimates for λ_c are statistically significant,

as our estimated t -statistic of 2.02 (for the annual data) and 2.22 (for the quarterly data) are in the far right hand tail of the empirical distribution. These t -statistics are at or near the 95% quantile, which is consistent with a rejection of the null hypothesis that no *positive* cross-sectional relationship exists at the 5% confidence level. As additional evidence in favor of the relationship between the measured average returns and the long-run cash flow beta, $\bar{R}^2$'s of 57% (for the annual data) and 50% (for the quarterly data) are in the far right tail of the adjusted R^2 empirical distribution for these data; in both cases, the empirical $\bar{R}^2$ exceeds the 97.5% critical value. Collectively, this experiment suggests that our empirical results reflect the true economic content of the estimated long-run cash flow beta rather than random chance. In an economy in which asset returns and dividend growth are independent of consumption growth, the probability of observing these estimated magnitudes of λ_c and the cross-sectional $\bar{R}^2$ is extremely low.

5 Economic Motivations

In this section, we explore the implications of extending the equilibrium model to facilitate more general preference specifications. In particular, for the time-nonseparable preferences developed in Epstein and Zin (1989) (EZ), the Intertemporal Marginal Rate of Substitution (IMRS) is

$$M_{t+1} = \exp \left\{ \theta \ln \delta - \frac{\theta}{\psi} g_{c,t+1} - (1 - \theta) r_{c,t+1} \right\}. \quad (20)$$

$g_{c,t+1}$ is the growth rate (in logs) of consumption and $r_{c,t+1}$ is the return (in logs) on an asset that pays off aggregate consumption each period. Further, δ is a time preference parameter and ψ the intertemporal elasticity of substitution. The parameter, $\theta \equiv \frac{1-\alpha}{1-\frac{1}{\psi}}$, wherein α represents the coefficient of relative risk aversion. Under this parameterization, the innovation in the (log) IMRS in this model is determined by

$$\eta_{m,t} = -\frac{1-\alpha}{\psi-1} \eta_{c,t} - \frac{\psi\alpha-1}{\psi-1} \eta_{r_{c,t}} \quad (21)$$

where the innovation in the return on the consumption asset is $\eta_{r_{c,t}}$ and $\eta_{c,t}$ is the innovation in consumption growth. It is well recognized that the innovation to the return to the consumption stream, $\eta_{r_{c,t}}$, is endogenous to the model. For example, when consumption growth is assumed to be an AR(1) process with Gaussian innovations, equation (21) leads

to a single-factor risk premium specification—we refer to this as Model 1. In Model 1, $\eta_{r_c,t}$ is a scalar multiple of the consumption innovation (i.e, perfectly correlated with consumption innovations) and as shown below, the cash flow’s long-run consumption beta is sufficient to characterize risk premia across assets. An alternative model (called Model 2) that leads to a two-factor specification follows where consumption dynamics are also characterized by stochastic volatility. In this case, expected returns may be time varying. Further, we have a two-factor model and the average return on assets is also determined by the covariation with $\eta_{r_c,t}$. This model captures the intuition that the risk premia on different assets are determined by the risk associated with the long-run cash flow consumption beta’s and the exposure of assets to a factor that determines time-varying risks.

5.1 Model 1: Constant Risk Premia

In the first model, we derive implications for risk premia in an economy where the risk premia on all assets are constant. As in equation (4), consumption growth follows an AR process with lag length 1:

$$g_{c,t+1} = \mu_c + \rho_c g_t + \eta_{c,t+1}$$

In this case, the innovation to the consumption portfolio is

$$\eta_{r_c,t} = (1 + b_c)\eta_{c,t},$$

where $b_c = \frac{1-\psi^{-1}}{1-\kappa\rho_c}$. Substituting this expression into equation (21) implies that

$$\eta_{m,t} = \lambda_1 \eta_{c,t},$$

where $\lambda_1 = \alpha + b_c \frac{\psi\alpha-1}{\psi-1}$. When consumption is assumed to be *i.i.d*, then $b_c = 0$ and λ_1 equals the risk aversion parameter α . Note that the return innovation to the consumption stream, $\eta_{r_c,t}$, is perfectly correlated with the innovations in consumption. These assumptions lead to a single-factor model with constant risk premia. The single factor prices risks associated with consumption news.

With constant cost of capital, it is also straightforward to show that for any asset i , the return innovation is:

$$r_{i,t} - E_{t-1}[r_{i,t}] \equiv e_{r_i,t} = \beta_{i,d}\eta_{c,t} + \zeta_{i,t} \quad (22)$$

where, as before, the term $\beta_{i,d}\eta_{c,t} + \zeta_{i,t}$ represents news about cash flows. The arithmetic risk premium on any asset, approximately, is determined by $cov_t(\lambda_1\eta_{c,t+1}, e_{r_i,t+1})$, hence the risk premium is

$$E_t[R_{i,t+1} - R_{f,t+1}] = \beta_{i,d}\lambda_1 Var(\eta_{c,t}) \quad (23)$$

In this model, the cross-sectional differences in risk premia are determined by the differences in the long-run exposure of cash flows to consumption news.

The cross-sectional implications in equation (23) motivate the set of cross-sectional restrictions that we have explored in the previous section. Under these assumptions, the long-run exposure of cash flows to consumption news may be modeled using a VAR for cash flow and consumption growth. The resulting restriction on cross-sectional risk premia is

$$E[R_{i,t+1}] = \lambda_0 + \beta_{i,d}\lambda_c \quad (24)$$

which follows from equation (23), with the average risk free rate being λ_0 and the price of risk for consumption $\lambda_1 Var(\eta_c)$ being λ_c . Consequently, under this set of assumptions, risk premia in the cross section are driven only by the price of risk associated with risk inherent in aggregate consumption growth.

5.2 Model 2: Risk and Return with Time-Varying Risk Premia

To allow for the possibility that risk premia are time-varying, we begin by assuming that consumption growth innovations are characterized by stochastic volatility. In this setting, the innovation in the return to the consumption stream is given by

$$\eta_{r_c,t} = b_1\eta_{c,t} + \sum_K b_K\eta_{K,t}$$

The terms $\eta_{K,t}$ correspond to the innovations in risk premia—for simplicity we assume that all the risk sources are uncorrelated. This can be motivated by a model that captures fluctuating consumption volatility as in Bansal and Yaron (2002) or can simply be viewed as a version of the ICAPM in Merton (1973). The innovations to the returns on the consumption stream are not perfectly correlated with consumption innovations as in Model 1.

Consider the innovation in the return to any asset i , where the risk premium on the asset

varies:

$$r_{i,t} - E_{t-1}[r_{i,t}] \equiv e_{r_{i,t}} = \beta_{i,d}\eta_{c,t} + \zeta_{i,t} - \sum_k \beta_{i,k}\eta_{k,t} \quad (25)$$

In addition to cash flow news, changes in expected returns also affect the return innovations. Consider the covariation in return innovations with innovations in the pricing kernel, $cov_t(\frac{1-\alpha}{\psi-1}\eta_{c,t} + \frac{\psi\alpha-1}{\psi-1}\eta_{r_{c,t}}, e_{r_{i,t}})$. Given our assumptions above, this covariation implies that the risk premium can be (approximately) stated as

$$E_t[R_{i,t+1} - R_{f,t+1}] = \beta_{i,d}\lambda_1 Var_t(\eta_{c,t}) + \sum_k \beta_{i,k}\lambda_k Var_t(\eta_{k,t}) \quad (26)$$

Economically, equation (26) captures the intuition that risk premia on assets are determined by cash-flow beta's and by variables that may influence the expected returns.

This model also allows us to interpret the links between market betas and expected returns. Note that the market return innovation will also satisfy equation (25). Hence, the covariance between return innovations to an asset i and the market is

$$\beta_{i,d}\beta_{mk,d}var(\eta_c) + \sum_k \beta_{i,k}\beta_{mk,k}var(\eta_k) \quad (27)$$

Note that across assets this covariance is a weighted average reflecting all risks: the cash flow risk, $\beta_{i,d}$, and risks associated with expected return news. The market beta of an asset will also reflect a weighted average of these two individual betas. However, while each individual beta may be important (and significantly priced), a weighted average of the two betas may fail to appear to be a priced risk source, as each beta carries different prices of risks (see equation (26)). This is one potential reason why market betas may fail to explain the cross-section of average returns.

One way to evaluate this proposition is to consider a regression of the market betas on the cash flow betas. This regression provides a sense of how much of the market beta is driven by the cash flow beta. Since the residual from this projection would only approximately identify the weighted average term $\sum_L \beta_{i,L}var(\eta_L)$, it may not be useful in explaining risk premia across assets. That is, the portion of market beta that is orthogonal to cash flow beta may be insufficient to capture risk attributable to aggregate economic uncertainty. However, as we are able to identify $\beta_{i,d}$ separately, we can still infer the percentage of the cross-section of market betas that are driven by the cash flow betas.

5.3 Multi Factor Model: Estimation and Empirical Results

To consider the implications of the addition of the market beta to the cross-sectional explanatory power of the long-run beta, we first must obtain a time-series estimate of the usual market beta. To do this, we augment the set of orthogonality conditions to include

$$E[(r_{i,t} - \bar{r}_{i,t})(r_{m,t} - \bar{r}_m)] = 0 \quad (28)$$

where, to maintain parsimony in estimation as above, we also demean all returns and market betas are estimated asset-by-asset. In addition, we also consider the modified cross-sectional restrictions

$$R_{i,t} = \lambda_0 + \lambda_c \beta_{i,d} + \lambda_m \beta_{i,m} + e_{i,t} \quad (29)$$

In reporting the standard errors for the price of risk parameters—all standard errors, as before are HAC-adjusted.

In Table 4 (Panel A), we report the time-series estimates of the market betas (see equation (28)) for the annual data. As can be seen the market betas, $\beta_{i,m}$, are estimated in the time-series with precision. Additionally, it appears that the market betas across the market capitalization sorted portfolios exhibit a strong pattern in accordance with their observed average returns (a well-known result), but this is not true for the book-to-market sorted portfolios; the high and low book-to-market portfolios have very similar estimates of the market beta. The industry sorted portfolios display a less pronounced pattern, but the durable goods industry is associated with the largest market beta and does display the largest average return. Quarterly estimates of the market betas (not reported) are broadly similar, and are available upon request. In accordance with the observed patterns, there is some evidence that the market beta *alone* can explain some of the cross-sectional variation in observed risk premia as implied by the standard Capital Asset Pricing Model. For example, for the annual data, the estimated market price of risk, λ_m , is both positive and significant, with an associated adjusted R^2 of 25%; however, at the quarterly frequency, λ_m is not significant, and the adjusted R^2 falls to 10%. In sum, this result suggests the possibility at least that there might be some incremental improvement in the more general two-factor model that facilitates risk prices for both the long-run cash flow beta and the market beta, as discussed above.

In Table 4, the results for the two-factor specification are presented for both the annual

(Panel B) and quarterly (Panel C) data. As can be seen, the inclusion of the market beta into the cross-sectional regressions, as in equation (29), does not dramatically affect the explanatory power of the long-run cash flow beta. Estimated risk-prices for the market beta are not statistically different from zero in either case. For example, for the annual frequency, the estimated risk price associated with the market portfolio covariance, λ_m , is 0.29, but with a HAC adjusted t -statistics of only 0.07. Conversely, the estimated risk price for the long-run cash flow beta with respect to consumption, λ_c , is 0.35, with a HAC adjusted standard error of 2.49. Further, the adjusted r-squared is 0.56 for the annual frequency, nearly identical to the r-square associated with the one-factor specification—using only the long run cash flow betas. Taken together, this evidence suggests that the cross-sectional explanatory presented above for the market beta, albeit quite weak, is entirely subsumed by ability of the long-run cash flow beta to capture observed risk premia. Finally, these results are broadly similar the quarterly frequency, suggesting that this evidence is robust to the frequency measurement.

The implication of the above result is that market beta does not capture a significant portion of the cross-sectional variation in risk premia relative to cash flow beta. However, as discussed above, the market beta does have some power for explaining the cross-section of returns. As suggested above, we regress the portfolio market betas on the cash flow betas. These untabulated results suggest that the long run cash flow betas explain a considerable portion of the market betas; the projection of the market beta's on the long run cash flow betas yields a point estimate is 0.349 and the regression adjusted R^2 is 0.42. A significant part of the explanatory power of market betas is attributable to their ability to proxy for cash flow beta and, as predicted, the relationship is positive.

We then regress average returns on cash flow betas and the residual from the projection of market betas on cash flow betas described above. This residual has no incremental explanatory power for average returns; the associated point estimate of 0.327 has a t -statistic of 0.077. The fact that the market beta has a considerable link to cash flow beta, and that the residual market beta has no explanatory power, is consistent with the economic arguments present in Section 5.

6 Conclusion

This paper documents a striking empirical observation. Cash flow betas, described as the long run response of dividends from a unit shock to aggregate consumption, can explain more than 50% of the cross-sectional variation in observed average returns across a challenging collection of size, book-to-market, and industry sorted portfolios. The cash flow betas correlation with mean returns is 0.79 for book-to-market, 0.81 for size, and 0.57 for industry portfolios. Across both the annual and quarterly frequencies, we measure the cash flow betas by estimating the joint time-series dynamics for both aggregate consumption and portfolio specific dividend growth rates using a VAR.

Our measures of the cash flow beta also explain about 50% of the dispersion in standard CAPM betas, despite the fact that the market betas alone do not explain much of the average return variation. We describe a model which allows for time-variation in both expected consumption growth and aggregate uncertainty, for which each risk source will require a distinct price. Under this specification, it is not surprising that the usual market beta, as a weighted average of an asset's exposure to both of these potential sources of risk, does not do a particularly good job explaining the cross-section of observed returns. From the perspective of our cash flow beta's model, the pronounced cross-sectional variation in average returns across these portfolios does not appear to be particularly puzzling. Our model captures the economic intuition that cash-flows of different assets have different long run exposures to fluctuations in the consumption in the economy, and that this exposure has considerable capacity to explain differences in mean returns across assets.

References

- Bansal, R., R. F. Dittmar, and C. T. Lundblad, 2001, “Consumption, Dividends, and the Cross-Section of Equity Returns,” working paper, Duke University and Indiana University.
- Bansal, R., and A. Yaron, 2002, “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” working paper, Duke University.
- Bollerslev, T., and R. Hodrick, 1995, “Financial Market Efficiency Tests,” in *Handbook of Applied Econometrics*, ed. by M. H. Pesaran, and M. R. Wickens. Blackwell, Oxford and Malden, MA, vol. 1 of *Macroeconomics*, chap. Chapter 9.
- Campbell, J., 2000, “Asset Pricing at the Millineum,” *Journal of Finance*, 55, 1515–1567.
- Campbell, J., and R. Shiller, 1988, “Stock Prices, Earnings, and Expected Dividends,” *Journal of Finance*, 43, 661–676.
- Epstein, L. G., and S. E. Zin, 1989, “Substitution, Risk Aversion, and the Intertemporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969.
- Fama, E., and K. French, 1993, “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics*, 33, 3–56.
- , 1997, “Industry Costs of Equity,” *Journal of Financial Economics*, 43, 153–193.
- Hansen, L. P., and K. Singleton, 1982, “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, 63, 767–804.
- , 1983, “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, 91, 249–265.
- Heaton, J., 1993, “The Interaction Between Time-Nonseparable Preferences and Time Aggregation,” *Econometrica*, 61, 353–385.
- Hodrick, R., 1992, “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 5, 357–386.
- Jagannathan, R., and Z. Wang, 1996, “The Conditional CAPM and the Cross-section of Expected Returns,” *Journal of Finance*, 51, 3–54.

- Lettau, M., and S. Ludvigson, 2001, “Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying,” *Journal of Political Economy*, forthcoming.
- Merton, R., 1973, “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 41, 867–887.
- Newey, W., and K. West, 1987, “A Simple, Positive Semi-definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.

Table 1: Summary Statistics

Portfolio	Cash Flows		Returns	
	Mean	Std.	Mean	Std.
s1	0.0369	0.1639	0.1359	0.2967
s2	0.0344	0.1232	0.1260	0.2512
s3	0.0280	0.0945	0.1220	0.2251
s4	0.0213	0.0996	0.1216	0.2245
s5	0.0229	0.1057	0.1211	0.2090
s6	0.0160	0.0832	0.1103	0.1953
s7	0.0181	0.0767	0.1117	0.1886
s8	0.0164	0.0987	0.1077	0.1743
s9	0.0136	0.0721	0.1033	0.1644
s10	0.0060	0.0633	0.0945	0.1686
b1	0.0091	0.1262	0.0907	0.2050
b2	0.0188	0.1288	0.0965	0.1622
b3	0.0135	0.1338	0.0962	0.1670
b4	0.0163	0.1198	0.0918	0.1672
b5	0.0326	0.0925	0.1151	0.1718
b6	0.0278	0.0816	0.1151	0.1655
b7	0.0261	0.0837	0.1173	0.1846
b8	0.0398	0.0936	0.1396	0.2069
b9	0.0382	0.1271	0.1400	0.2053
b10	0.0439	0.2202	0.1514	0.2359
i1	0.0275	0.0842	0.1062	0.1769
i2	0.0168	0.1725	0.1236	0.2406
i3	0.0073	0.1433	0.0907	0.1962
i4	0.0195	0.0741	0.1073	0.1947
i5	-0.0095	0.1086	0.1048	0.2254
i6	-0.0032	0.0619	0.0956	0.2021
i7	0.0034	0.1213	0.1067	0.2202
i8	0.0188	0.0900	0.1097	0.1812
i9	0.0003	0.0425	0.0874	0.1599
i10	0.0218	0.0785	0.1105	0.1922

Table 1 presents summary statistics for the data used in the paper. The table presents real mean returns and cash flow growth rates for a set of 30 portfolios. Portfolios are sorted into deciles on the basis of market capitalization (s1-s10), book-to-market (b1-b10), and ten industry groups (i1-i10). Data are sampled at the annual frequency over the period 1949 through 2001 and are converted to real using the PCE deflator.

Table 2: Time Series Parameters

Port.	γ_i		$\rho_{1i} + \rho_{2i}$		b_i		$\beta_{d,i}$	
s1	0.229	(1.769)	0.335	(0.100)	5.027	(1.872)	7.809	(3.707)
s2	2.755	(1.355)	0.386	(0.147)	3.668	(1.231)	10.136	(4.083)
s3	1.073	(1.091)	0.192	(0.165)	3.780	(1.016)	6.298	(2.810)
s4	1.753	(1.022)	-0.086	(0.225)	2.038	(0.854)	4.578	(1.910)
s5	1.643	(1.058)	-0.155	(0.242)	0.852	(1.037)	3.261	(1.855)
s6	1.177	(0.771)	0.035	(0.203)	2.395	(0.695)	4.291	(1.865)
s7	1.334	(0.861)	0.154	(0.181)	1.307	(0.753)	3.579	(1.805)
s8	1.875	(0.911)	-0.440	(0.237)	1.156	(0.778)	3.682	(1.601)
s9	1.039	(0.638)	0.079	(0.172)	1.745	(0.598)	3.487	(1.385)
s10	-0.012	(0.630)	0.031	(0.222)	1.555	(0.420)	1.594	(1.222)
b1	0.478	(2.261)	0.159	(0.169)	2.978	(1.502)	4.282	(5.242)
b2	-2.341	(1.521)	-0.109	(0.217)	-0.052	(1.368)	-3.645	(2.988)
b3	-0.731	(1.490)	-0.371	(0.147)	2.410	(0.929)	0.664	(2.120)
b4	-0.087	(1.227)	-0.046	(0.236)	-0.345	(1.374)	-0.467	(2.283)
b5	0.430	(0.994)	0.160	(0.182)	0.112	(1.065)	0.790	(2.074)
b6	0.775	(1.203)	-0.192	(0.187)	0.522	(0.904)	1.626	(2.081)
b7	1.961	(0.966)	-0.085	(0.175)	0.456	(0.815)	3.411	(1.841)
b8	1.637	(0.923)	-0.010	(0.107)	3.952	(0.913)	6.443	(1.987)
b9	1.895	(1.742)	0.057	(0.134)	3.411	(1.510)	6.526	(3.406)
b10	4.536	(2.590)	0.148	(0.169)	7.026	(2.814)	15.131	(5.649)
i1	0.470	(0.824)	0.360	(0.245)	0.285	(0.549)	1.150	(1.789)
i2	0.527	(1.874)	-0.050	(0.155)	7.136	(1.409)	7.678	(3.813)
i3	2.552	(1.144)	-0.714	(0.398)	-0.011	(1.218)	3.900	(1.934)
i4	0.520	(0.634)	0.013	(0.218)	0.121	(1.038)	0.916	(1.572)
i5	1.846	(0.889)	-0.148	(0.273)	1.686	(1.048)	4.319	(2.029)
i6	-0.317	(0.619)	-0.140	(0.213)	0.456	(0.655)	-0.082	(1.243)
i7	1.199	(0.821)	-0.538	(0.242)	1.901	(1.018)	3.098	(1.666)
i8	0.423	(0.514)	0.350	(0.164)	1.315	(0.541)	2.630	(1.334)
i9	-0.244	(0.522)	0.400	(0.265)	0.957	(0.404)	1.141	(1.071)
i10	1.085	(0.652)	-0.037	(0.196)	0.551	(0.682)	2.192	(1.294)

Table 2 depicts the estimated time series parameters from the model

$$\begin{aligned}
g_{c,t+1} &= \rho_c g_{c,t} + \eta_{t+1} \\
g_{i,t+1} &= \gamma_i g_{c,t} + u_{i,t+1} \\
u_{i,t+1} &= b_i \eta_{t+1} + \rho_{1,i} u_{i,t} + \rho_{2,i} u_{i,t-1} + \zeta_{i,t+1} \\
\beta_{d,i} &= \frac{\kappa_{i,1} \gamma_i}{1 - \kappa_{i,1} \rho_c} + \frac{b_i}{1 - \kappa_{i,1} \rho_{1,i} - \kappa_{i,1}^2 \rho_{2,i}}
\end{aligned}$$

where $g_{c,t+1}$ represents the demeaned growth rate in aggregate consumption and $g_{i,t+1}$ represents the demeaned growth rate in the dividends paid on portfolio i . The parameter $\kappa_{i,1}$ represents a constant of approximation in the Campbell and Shiller (1988) expression for returns. Data used in the estimation are 30 portfolios sorted on size, book-to-market and industry and are sampled annually over the period 1949 through 2001. All quantities are demeaned and converted to real using the PCE deflator. Standard errors are presented in parentheses and are estimated using a HAC covariance matrix with one Newey-West lag. Standard errors for $\rho_{1i} + \rho_{2i}$ and $\beta_{d,i}$ are computed via the delta method.

Table 3: Cross-Sectional Regressions

Panel A: Annual Data: 1949-2001

	λ_0	λ_c	$\bar{R}^2$
Point Est.	9.882	0.357	0.574
t -stat	4.926	2.021	

Distribution of t -statistics and $\bar{R}^2$

	2.5%	5.0%	10.0%	20.0%	30.0%	40.0%	50.0%	60.0%	70.0%	80.0%	90.0%	95.0%	97.5%
t -stat	-2.294	-2.102	-1.870	-1.505	-1.110	-0.602	0.049	0.699	1.165	1.514	1.880	2.111	2.319
$\bar{R}^2$	-0.035	-0.035	-0.032	-0.020	-0.002	0.025	0.061	0.105	0.160	0.232	0.341	0.417	0.478

Panel B: Quarterly Data: 1949:2-2001:4

	λ_0	λ_c	$\bar{R}^2$
Point Est.	2.426	0.078	0.498
t -stat	4.789	2.217	

Distribution of t -statistics and $\bar{R}^2$

	2.5%	5.0%	10.0%	20.0%	30.0%	40.0%	50.0%	60.0%	70.0%	80.0%	90.0%	95.0%	97.5%
t -stat	-2.433	-2.263	-2.028	-1.648	-1.223	-0.688	-0.024	0.626	1.167	1.595	2.000	2.217	2.383
$\bar{R}^2$	-0.036	-0.035	-0.033	-0.024	-0.009	0.012	0.040	0.076	0.126	0.184	0.282	0.361	0.426

Table 3 presents cross sectional regressions of average returns for 30 portfolios on the long-run beta, $\beta_{d,i}$ developed in the paper:

$$R_{i,t} = \lambda_0 + \lambda_c \beta_{d,i} + v_{i,t}$$

In Panel A, results are presented for annual data covering the period 1949-2001; results with quarterly data over the period 1949:2-2001:4 are presented in Panel B. Parameters are estimated via ordinary least squares (OLS); t -statistics are computed with HAC-adjusted standard errors. We also present the distribution of the t -statistics for the test $H_0 : \lambda_c = 0$ and the $\bar{R}^2$ generated by a Monte Carlo experiment of 10,000 replications. In the Monte Carlo, we simulate the demeaned consumption growth rate as

$$\hat{g}_{c,t+1} = \hat{\rho}_c \hat{g}_{c,t} + \hat{\eta}_{t+1}$$

where $\hat{\rho}_c$ is the autoregressive parameter for consumption estimated in the data and σ_η corresponds to the standard deviation of the residual in the data. The simulated consumption growth rates and observed dividend growth rates are used to estimate the time series parameters in model (9) and the cross-sectional parameters in the above specification. In the case of quarterly data, the simulated growth rates are also used to generate a trailing four quarter moving sum of growth rates used in the estimation of γ_i .

Table 4: Cross-Sectional Regressions

Panel A: Annual Betas					
Port.	β_i		Port.	β_i	
s1	1.301	(0.152)	b6	0.896	(0.063)
s2	1.221	(0.115)	b7	0.912	(0.107)
s3	1.138	(0.098)	b8	1.041	(0.133)
s4	1.141	(0.095)	b9	1.019	(0.121)
s5	1.110	(0.073)	b10	1.112	(0.130)
s6	1.057	(0.068)	i1	0.798	(0.119)
s7	1.060	(0.054)	i2	1.239	(0.074)
s8	0.972	(0.052)	i3	1.061	(0.090)
s9	0.942	(0.039)	i4	0.817	(0.104)
s10	0.974	(0.029)	i5	1.190	(0.081)
b1	1.098	(0.055)	i6	0.769	(0.150)
b2	0.911	(0.041)	i7	1.050	(0.118)
b3	0.951	(0.034)	i8	0.894	(0.077)
b4	0.877	(0.081)	i9	0.559	(0.127)
b5	0.895	(0.083)	i10	0.931	(0.092)

Panel B: Annual Data: 1949-2001

	λ_0	λ_c	λ_β	R^2
Point Est.	9.626	0.349	0.288	0.558
<i>t</i> -stat	2.667	2.492	0.068	

Panel C: Quarterly Data: 1949:2-2001:4

	λ_0	λ_c	λ_β	R^2
Point Est.	2.299	0.075	0.136	0.483
<i>t</i> -stat	3.020	2.417	0.159	

Table 4 presents cross sectional regressions of average returns for 30 portfolios on the long-run beta, $\beta_{d,i}$ developed in the paper:

$$R_{i,t} = \lambda_0 + \lambda_c \beta_{d,i} + \lambda_\beta \beta_{mkt,i} + v_{i,t}$$

Panel A depicts GMM estimates of betas using annual real returns data. In Panel B, results are presented for annual data covering the period 1949-2001; results with quarterly data over the period 1949:2-2001:4 are presented in Panel C. Parameters are estimated via ordinary least squares (OLS); *t*-statistics are computed with HAC-adjusted standard errors.