

Cointegration and Consumption Risks in Asset Returns

Ravi Bansal

Duke University and NBER

Robert Dittmar

University of Michigan

Dana Kiku

University of Pennsylvania

We argue that the cointegrating relation between dividends and consumption, a measure of long-run consumption risks, is a key determinant of risk premia at all investment horizons. As the investment horizon increases, transitory risks disappear, and the asset's beta is dominated by long-run consumption risks. We show that the return betas, derived from the cointegration-based VAR (EC-VAR) model, successfully account for the cross-sectional variation in equity returns at both short and long horizons; however, this is not the case when the cointegrating restriction is ignored. Our evidence highlights the importance of cointegration-based long-run consumption risks for financial markets. (*JEL G1, G12*)

How does the riskiness of equity returns change with investment horizon? We show that at long horizons, risks are dominated by long-run consumption risks in dividends, while at short horizons, additional price risks may also matter. The cointegrating relation between dividends and consumption (i.e., the long-run consumption beta of dividends) provides a measure of long-run risks in dividends. This cointegrating relation has important conceptual and empirical implications for measuring risks in asset returns at all horizons. Empirically, we document that consumption betas determined by the cointegration-based vector-autoregression (EC-VAR) can account for the cross section of mean equity returns at both short and long horizons. This, however, is not the case when the cointegration restriction between consumption and dividends is ignored. Hence, our evidence highlights the economic importance of long-run

An earlier draft of this paper was circulated with the title "Long-Run Risks and Equity Returns." We would like to thank Lars Hansen, George Tauchen, Jessica Wachter, Vivian Wang, Amir Yaron, and the seminar participants at Arizona State University, Duke University, Simon Fraser University, Stanford University, University of British Columbia, University of Michigan, University of North Carolina, University of Pennsylvania, Vanderbilt University, Federal Reserve Board of Governors Conference on Risk Premiums: Time Variation and Macroeconomic Links, and 2007 American Finance Association Meetings for helpful comments. The usual disclaimer applies. Address correspondence to Ravi Bansal, Fuqua School of Business, Duke University, Durham, NC 27708; telephone: 919-660-7758; e-mail: ravi.bansal@duke.edu.

risks in explaining expected equity returns across different investment horizons.

Our focus on the cointegrating relation is motivated by Bansal and Yaron (2004); Bansal, Dittmar, and Lundblad (2001, 2005); Hansen, Heaton, and Li (2006); and Bansal, Gallant, and Tauchen (2007), who show that long-run risks can be important in explaining risk premia. Cointegration, as made clear by the Granger and Engle (1987) representation theorem, has sharp implications for predictability. In the context of dividends and consumption, the representation theorem implies that the deviation of the level of dividends from consumption (the error correction variable) is important for predicting dividend growth rates and returns at all horizons. As the error correction variable alters the information set used to predict future returns, it also significantly influences the return innovation and, hence, the conditional consumption betas.

The empirical literature frequently imposes the implicit assumption of no cointegration by modeling returns with a standard vector autoregression (VAR). We show that this traditional approach leads to a significant deterioration in the ability of consumption-based models to explain risk premia at both short and long horizons relative to the cointegration-based specification. In this “standard” VAR, consumption and dividend growth rates are stationary; however, the levels of consumption and dividends are not cointegrated. Deviations of dividends from consumption contain a unit root and, thus, the two series can drift far apart—there is no error-correction mechanism that ties dividends and consumption together in the long run. Cointegration, and the implied EC-VAR framework, ties these two series together in the long run; for this reason, the error correction term can be quite important for predicting future returns and growth rates. The inclusion of the error correction mechanism can significantly alter the transition dynamics of returns and asset betas relative to the standard growth-rate-based VAR specification.

We first document that the error correction term in the dividend-consumption cointegrating relation contains important information for predicting future dividend growth and returns. Imposing cointegration, we are able to predict on average 11.5% of the variation in one-year returns, compared to 7.5% when we do not impose cointegration. This difference is even starker at longer horizons: at the 10-year horizon, the EC-VAR specification results in an average 44.0% adjusted R^2 , compared to 9.9% for the standard growth-rate VAR specification. That is, at longer horizons, we are able to predict far more variation in returns using the cointegration specification than when we do not impose cointegration on the return dynamics. This predictability evidence suggests that cointegration has important implications for measuring return innovations and, consequently, conditional consumption betas.

Using the EC-VAR, we find that the cointegrating relation indeed influences the conditional consumption betas and resulting prices of risk significantly at all investment horizons. The estimated market price of consumption risks is always positive and significant. For example, at the short horizon, the market

price of risk is 1.19 (SE = 0.41); at the long horizon, it continues to be highly significant at 0.72 (SE = 0.25). Our conditional consumption betas account for about 75% of the cross-sectional variation in risk premia at the one-year horizon, and are able to explain over 85% of the cross section of mean returns at long horizons. In contrast, VAR models used traditionally in the literature ignore the cointegrating relation in measuring conditional betas and, hence, are unable to account for the differences in risk premia across assets. The unconditional consumption CAPM also fails to explain the variation in average returns across assets. We conduct a series of Monte Carlo experiments and show that our empirical evidence is statistically significant and robust.

A rich array of specifications for dividends have been used in earlier asset-pricing models. In Campbell and Cochrane (1999) and Bansal and Yaron (2004) models, consumption and dividend growth rates have the same unconditional mean but are not cointegrated. Cochrane, Longstaff, and Santa-Clara (2006) developed a theoretical model where the average growth rates of aggregate dividends differ across sectors and analyze the implications for risk premia. Menzly, Santos, and Veronesi (2004) consider a specification where dividends and consumption across all sectors have unit cointegration, while Bansal, Dittmar, and Lundblad (2001); and Hansen, Heaton, and Li (2006) entertain specifications that permit heterogeneity in the cointegrating relation between dividends and consumption. The economic implications of these alternative specifications of dividends for the risk-return relation are empirically evaluated in the current paper. We find the heterogeneity in the cointegrating relation to be empirically important for understanding risk premia across assets at all investment horizons. Specifications, which do not allow for this heterogeneity, are not supported in the data. The empirical evidence in this paper, therefore, suggests that developing dynamic models that permit heterogeneity in the cointegrating relation would be very valuable.

The unique dimension of our paper is to show that after accounting for the long-run relation between dividends and consumption, conditional consumption betas contain important information about risk premia at all investment horizons. At long horizons, transitory risks vanish and only long-run risks drive risk compensations. Our approach, therefore, allows us to analyze the size of the compensation for long-run versus transitory risks in accounting for risk premia. Quantitatively, we find that long-run consumption risks are the dominant source of risk premia at all investment horizons. Conceptually, we provide a framework for linking the cointegration parameter to the conditional consumption betas by horizon and show that they can explain the cross-sectional variation in expected returns at both short and long horizons. While we focus on cointegration and its implications for measuring and understanding expected returns at various horizons, Hansen, Heaton, and Li (2006) aim to characterize the evolution of risk prices and the corresponding present-value relation from the perspective of the long-run risk model.

Our study of cointegration-based long-run consumption risks and expected returns at various horizons differs from earlier work on risk and return. Bansal, Dittmar, and Lundblad (2001, 2005) focus on cash-flow betas and inquire if these betas can account for short-horizon expected returns. Parker and Julliard (2005) measure covariance risk between current returns and future consumption, while Jagannathan and Wang (1996) and Lettau and Ludvigson (2001) rely on time-varying betas to empirically explain the cross-sectional pattern in one-period expected returns. Our approach goes beyond a single-period risk-return trade-off, allowing us to characterize the composition of dividend and price risks in asset returns and evaluate their relative importance across different investment horizons.

The rest of the paper is organized as follows. In Section 1, we discuss a brief, simple theoretical framework and present our econometric specification for return dynamics. Section 2 describes the data. Section 3 contains the results of our empirical analysis. Monte Carlo evidence that supports our data findings is discussed in Section 4. Finally, Section 5 provides concluding remarks.

1. Cointegration and Risk Premia

In this section, we discuss a simple theoretical framework that we employ and implications of cointegration for risk measures and risk premia. Throughout our discussion, we utilize a Taylor series approximation for log returns (see Campbell and Shiller, 1988). Let $z_t = p_t - d_t$ represent the log of the price-dividend ratio. The return approximation is given by

$$r_{t+1} = \kappa_0 + d_{t+1} - d_t + \kappa_1 z_{t+1} - z_t, \tag{1}$$

where the log-linearization constants κ_0 and κ_1 are

$$\kappa_1 = \frac{\exp(\bar{z})}{1 + \exp(\bar{z})}, \tag{2}$$

$$\kappa_0 = \log(1 + \exp(\bar{z})) - \kappa_1 \bar{z},$$

and \bar{z} is the mean of the log price-dividend ratio. Using $\Delta z_{t+s} = z_{t+s} - z_t$, Equation (1) also implies that:

$$r_{t+1} = \kappa_0 + \Delta d_{t+1} + \Delta z_{t+1} + (\kappa_1 - 1)z_{t+1}. \tag{3}$$

We further assume that the log price-dividend ratio z_t and dividend growth rates Δd_{t+1} are covariance stationary processes. To model long-horizon returns, it is useful to define the summed level of the price-dividend ratio as $z_{t+1 \rightarrow t+s} \equiv \sum_{j=1}^s z_{t+j}$, the accumulated change in it as $\Delta z_{t+1 \rightarrow t+s} \equiv z_{t+1 \rightarrow t+s} - z_{t \rightarrow t+s-1}$, and accumulated s -period growth rate as $\Delta d_{t+1 \rightarrow t+s} \equiv \sum_{j=1}^s \Delta d_{t+j}$. Using the return expression (3), it follows that s -horizon

compounded return, $r_{t+1 \rightarrow t+s} \equiv \sum_{j=1}^s r_{t+j}$, is completely characterized by these three components.

1.1 Risk premia

We assume, as in Lucas (1978), that a representative agent prices assets across various investment horizons. The Euler equation associated with this agent will therefore determine the risk-return relation at all horizons. To keep the analysis simple, we further assume that the representative agent's preferences are characterized by the standard power utility of aggregate consumption. Hence, the log of the one-step-ahead stochastic discount factor is $m_t = \ln(\mathcal{Q}) - \gamma \Delta c_t$, where Δc_t is the growth rate of log aggregate consumption, γ is the parameter of risk aversion, and \mathcal{Q} determines time preferences. The standard asset-pricing condition for any horizon s is

$$E_t[\exp(m_{t+1 \rightarrow t+s} + r_{t+1 \rightarrow t+s})] = 1, \quad (4)$$

where $m_{t+1 \rightarrow t+s} = \sum_{j=1}^s m_{t+j}$. We assume that the log of the stochastic discount factor and log returns are jointly conditionally normal and homoskedastic, and Δc_t is a covariance stationary process. Our preference and distributional assumptions are similar to those used by Hansen and Singleton (1983). In this specification, as in Breeden (1979), risk premia are determined by consumption betas.¹

Let $\eta_{t+1} \equiv \Delta c_{t+1} - E_t[\Delta c_{t+1}]$ be the one-step-ahead innovation in consumption growth and $\eta_{t+1 \rightarrow t+s}$ be the innovation in the cumulative growth in consumption, $\Delta c_{t+1 \rightarrow t+s}$. The conditional variance of the cumulative consumption growth corresponds to the variance of this s -period innovation, which we denote as $\sigma_{c,s}^2$. Similarly, $\eta_{r,t+1 \rightarrow t+s}$ is the innovation in s -horizon return, $r_{t+1 \rightarrow t+s}$, $\sigma_{r,s}^2$ denotes its conditional variance, and $\text{Cov}[\eta_{t+1 \rightarrow t+s}, \eta_{r,t+1 \rightarrow t+s}]$ corresponds to the conditional covariance between consumption and return innovations for horizon s . Given our distributional assumption of homoskedasticity, all conditional second moments are constants. The asset pricing condition in Equation (4) along with distributional assumptions implies that the conditional expected return at any horizon s is given by

$$\begin{aligned} \frac{1}{s} E_t[r_{t+1 \rightarrow t+s} + 0.5\sigma_{r,s}^2 - sr_{t,s}^f] &= \frac{1}{s} \gamma \text{Cov}[\eta_{t+1 \rightarrow t+s}, \eta_{r,t+1 \rightarrow t+s}] \\ &= \frac{1}{s} \gamma \sigma_{c,s}^2 \frac{\text{Cov}[\eta_{t+1 \rightarrow t+s}, \eta_{r,t+1 \rightarrow t+s}]}{\sigma_{c,s}^2} \\ &= \lambda_s \beta_s, \end{aligned} \quad (5)$$

¹ It is straightforward to extend the presented framework to a more general framework, such as in Bansal and Yaron (2004). This, however, would entail estimation of a larger set of parameters in the cross section of assets. To avoid this, we confine our attention to the case of power utility, which is sufficient for addressing questions about the cross-sectional dispersion in mean returns that are the focus of this paper.

where $r_{t,s}^f$ is the yield on an s -period discount bond known at date t . $\lambda_s \equiv \frac{1}{s} \gamma \sigma_{c,s}^2$ is the price of risk that is determined by the agent's risk aversion and the conditional variance of consumption growth for the corresponding investment horizon, and β_s is the conditional consumption-risk measure for horizon s . We scale expression (5) by s to ensure that moments exist and to provide the interpretation of the risk premium as per unit of time.

The expected return is obtained by taking unconditional expectations of Equation (5). Let μ_r be the one-period-ahead geometric-mean return:

$$E \left\{ \frac{1}{s} E_t [r_{t+1 \rightarrow t+s} + 0.5 \sigma_{r,s}^2] \right\} = \mu_r + \frac{0.5 \sigma_{r,s}^2}{s} = E[r_{t,s}^f] + \beta_s \lambda_s. \quad (6)$$

As shown in expression (6), expected returns per unit of time are determined by the market price of horizon risk, λ_s , and the horizon-dependent beta, β_s .

If returns are Gaussian *i.i.d.* processes, then conditional expected returns per unit time at all horizons will be identical, as $\sigma_{r,s}^2 = s \sigma_{r,1}^2$. On the other hand, if there is predictable variation in returns, then $\sigma_{r,s}^2$ need not equal $s \sigma_{r,1}^2$, leading to variation in expected returns per unit time across different investment horizons. In Section 1.2, we provide an empirical method for measuring $\sigma_{r,s}^2$, which allows us to construct expected returns at various horizons.

1.2 Betas by horizon

The return expression (3) implies that the innovation in s -horizon return is determined by the innovations to $\Delta d_{t+1 \rightarrow t+s}$, $\Delta z_{t+1 \rightarrow t+s}$, and $z_{t+1 \rightarrow t+s}$, which we denote as $\eta_{d,t+1 \rightarrow t+s}$, $\eta_{\Delta z,t+1 \rightarrow t+s}$, and $\eta_{z,t+1 \rightarrow t+s}$, respectively. The covariance of these components with consumption growth will determine the overall risk compensation. Specifically, note that the asset beta β_s can be expressed as the sum of three covariances:

$$\beta_s = \frac{\text{Cov}(\eta_{t+1 \rightarrow t+s}, \eta_{d,t+1 \rightarrow t+s})}{\sigma_{c,s}^2} + \frac{\text{Cov}(\eta_{t+1 \rightarrow t+s}, \eta_{\Delta z,t+1 \rightarrow t+s})}{\sigma_{c,s}^2} + (\kappa_1 - 1) \frac{\text{Cov}(\eta_{t+1 \rightarrow t+s}, \eta_{z,t+1 \rightarrow t+s})}{\sigma_{c,s}^2}, \quad (7)$$

which we write compactly as

$$\beta_s = [\beta_{d,s} + \beta_{\Delta z,s} + (\kappa_1 - 1) \beta_{z,s}]. \quad (8)$$

Thus, the asset beta is determined by risks in both dividend growth rates and price-dividend ratios. We refer to the first component $\beta_{d,s}$ as the cash-flow beta, $\beta_{\Delta z,s}$ as the price change beta, and $\beta_{z,t}$ as the level beta associated with the risk in the cumulative price-dividend ratio. Equation (8) suggests that the joint dynamics of consumption growth, dividend growth, and price-dividend ratios contribute to the determination of assets' risk. We model these dynamics and discuss their implications in Sections 1.3 and 1.4.

The expression for risk exposures [Equation (8)] also suggests that risks faced by a long-run investor are quite different than those faced by a single-period investor. While at short horizons, transitory movements in price-dividend ratios may be important, their contribution to asset betas and risk compensations is virtually washed out in the long run. As the change in the price-dividend ratio represents a change in a stationary variable, its covariance with consumption growth [see the second term in Equation (7)] is swamped gradually by the variance of the cumulative growth in consumption; in the limit (i.e., when $s = \infty$), $\beta_{\Delta z, s} = 0$. Further, as the approximation parameter κ_1 is close to one,² the final term in expression (8), $(\kappa_1 - 1)\beta_{z, s}$, is quite close to zero. Consequently, for a long-run investor, an asset's risk is dominated by the long-run covariance risk in dividends, which is the cointegration parameter between dividends and aggregate consumption measures. In particular, as the investment horizon increases,

$$\lim_{s \rightarrow \infty} \beta_s \equiv \beta_{lr} = \delta + (\kappa_1 - 1)\beta_{z, lr}, \quad (9)$$

where δ is the parameter of cointegration between dividends and consumption [see Equation (19) in Section 1.3]. In our empirical work, we show that the long-run risk measure δ has important information about expected returns at short horizons as well.

1.3 Cointegration specification

A number of recent papers (e.g., Bansal, Dittmar, and Lundblad, 2001; and Hansen, Heaton, and Li, 2006) suggest that consumption and dividends are stochastically cointegrated (see Campbell and Perron, 1991). In this section, we examine the implications of cointegration for the calculation of the risk measures outlined in Section 1.2.

We focus on the per-share dividend series traditionally used in the literature. Dividends per share are constructed as follows. The total return per dollar invested is

$$R_{t+1} = H_{t+1} + Y_{t+1}, \quad (10)$$

where H_{t+1} is the price appreciation and Y_{t+1} is the dividend yield. The level of dividends per share and the price per share can then be computed as

$$\begin{aligned} D_{t+1} &= Y_{t+1}V_t, \\ V_{t+1} &= H_{t+1}V_t, \end{aligned} \quad (11)$$

given V_0 is the price per share of portfolio at the initiation date. Note that the present value of dividends per share is price per share. These dividend series correspond to the series commonly used in empirical work, as in Campbell and

² In the data, the value of the log-linearization parameter κ_1 varies between 0.96 and 0.99.

Shiller (1988); Bansal, Dittmar, and Lundblad (2005); and Hansen, Heaton, and Li (2006), among others.

If two nonstationary variables are cointegrated, a linear combination of the variables is stationary. We express the relation between two such variables—dividends and consumption—as

$$d_t = \tau_0 + \tau t + \delta c_t + \epsilon_{d,t}, \quad (12)$$

where the level of dividends and consumption in logs is d_t and c_t , respectively, and $E[\epsilon_{d,t}] = 0$. From Equation (12), it also follows that $\tau = \mu_d - \delta\mu_c$, where μ_c and μ_d are the average growth of consumption and dividends, respectively. Substituting for τ , the above equation can equivalently be stated as

$$d_t - \mu_d t = \tau_0 + \delta(c_t - \mu_c t) + \epsilon_{d,t}. \quad (13)$$

That is, the cointegration parameter δ can be estimated via the projection of deterministically detrended dividends on detrended consumption (as in Bansal, Dittmar, and Lundblad, 2001). The Granger and Engle (1987) representation theorem states that the error correction variable $\epsilon_{d,t}$ will forecast dividend growth rates in future, and consequently, perhaps returns. As we show below, this predictability has important implications for the calculations of risk measures in standard models.

Several features of our cointegration specification are worthy of further discussion. In particular, the specification in Equation (12) includes a time trend and the cointegration parameter δ is unrestricted as in Bansal, Dittmar, and Lundblad (2001)—we refer to this as our “preferred specification.” In addition to our preferred specification, Hansen, Heaton, and Li (2006) also entertain a specification with $\tau = 0$ and unrestricted δ , while Menzly, Santos, and Veronesi (2004) consider a specification with $\tau = 0$ and $\delta = 1$. In our empirical analysis, we evaluate the implications of these alternative specifications as well.

It is important to note that it is not time trend but the cross-sectional heterogeneity in δ that is important for capturing differences in risks across assets. Note that without the time trend in Equation (12), the cointegration parameter simply equals the ratio of average dividend growth to average growth in aggregate consumption (see Hamilton, 1994). In this case, cross-sectional differences in δ s will tautologically reflect differences in average dividend growth (and average capital gains) and, therefore, average *ex post* returns. Inclusion of the time trend purges the effect of mean growth rates in dividends on the cointegration parameter, and ensures that long-run risk measures do not mechanically replicate cross-sectional differences in *ex post* average returns.

1.4 Deriving mean returns and betas by horizon

This section provides the details for estimating assets’ consumption betas for different investment horizons. We first estimate the cointegrating relation (12) between dividends and consumption via ordinary least squares (OLS) that is,

by regressing the detrended portfolio's cash flows onto the stochastic trend in consumption. Using the resulting cointegrating residual $\epsilon_{d,t}$, we model its dynamics jointly with the portfolio's price-dividend ratio z_t and consumption growth Δc_t , via the first-order error correction VAR (EC-VAR) structure

$$X_t = AX_{t-1} + Gu_t, \quad (14)$$

where $X'_t = (\Delta c_t \ \epsilon_{d,t} \ z_t \ \Delta d_t \ \Delta z_t)$, A is a 5×5 matrix of coefficients, G is a 5×3 matrix, and u is a 3×1 matrix of shocks, $u'_t = (\eta_t \ \eta_{\epsilon,t} \ \eta_{z,t})$.³ All the variables are demeaned throughout our discussion. While details of the EC-VAR specification are provided in the Appendix, in this section we highlight some of its salient features. The first three variables form the basis of the EC-VAR process. The last two variables provide no additional information; they are included in the EC-VAR to facilitate the description and discussion of consumption betas by horizon. The dynamics of Δd_t and Δz_t are derived from the dynamics of the first three variables by exploiting $\Delta d_t = \Delta \epsilon_{d,t} + \delta \Delta c_t$ and $\Delta z_t = z_t - z_{t-1}$.

Note that the EC-VAR specification reduces to the standard growth-rate-based VAR if the error correction term $\epsilon_{d,t}$ is excluded from the set of variables that predict future growth rates and returns. In this “standard” VAR, c_t and d_t are each integrated (i.e., the growth rates of consumption and dividends are stationary); however, c_t and d_t are not cointegrated. Thus, deviations of dividends from consumption contain a unit root and these two series can arbitrarily drift far apart—there is no error correction mechanism that ties dividends and consumption together in the long run. Cointegration, and the implied EC-VAR framework, ties the long-run dynamics of the two series together. For this reason, the error-correction term $\epsilon_{d,t}$ can be quite important for predicting future dividend growth rates and returns and its inclusion can significantly alter the transition dynamics of returns. In sum, there can be large differences between the standard VAR and the EC-VAR implications for risk measures and expected returns, which we subsequently highlight in our empirical work.

Using the recursive structure of the EC-VAR, we can compute the conditional variance of returns $\sigma_{r,s}^2$, the conditional variance of consumption growth $\sigma_{c,s}^2$, and the conditional beta β_s , for any given horizon s . Specifically, let $B_j = B_{j-1} + A^{j-1}$, $j = 1, 2, \dots$, and $B_0 = 0$. The horizon- s covariance matrix of the above variables satisfies the recursion

$$\Sigma_s^* = B_s \Sigma_e B_s' + \Sigma_{s-1}^*, \quad (15)$$

where $\Sigma_e = G \Sigma_u G'$, and $\Sigma_0^* = 0$. As s increases, Σ_s^* grows without a bound; hence we consider $\Sigma_s \equiv \frac{\Sigma_s^*}{s}$. In the long-run limit, this covariance matrix

³ We also have considered specifications that include more lags in the vector-error-correction model. As this does not change the results materially, we only report evidence based on the first-order EC-VAR.

becomes

$$\Sigma_{lr} = [I - A]^{-1} \Sigma_e [I - A]^{-1'}. \quad (16)$$

Details of the derivation of the covariance matrix are provided in the Appendix.

The s -period covariance matrix Σ_s allows us to calculate the s -period beta of an asset using the appropriate covariance and variance terms. For a given horizon s , the covariance risk in the asset is

$$\beta_s = \frac{\Sigma_s(1, 4)}{\Sigma_s(1, 1)} + \frac{\Sigma_s(1, 5)}{\Sigma_s(1, 1)} + (\kappa_1 - 1) \frac{\Sigma_s(1, 3)}{\Sigma_s(1, 1)}, \quad (17)$$

where $\Sigma_s(i, j)$ is the (i, j) -element of the covariance matrix Σ_s . Note that computations in Equation (17) correspond to various components in Equation (8). In particular, the first component corresponds to the cash-flow risk $\beta_{d,s}$, the second component $\beta_{\Delta z,s}$ corresponds to the price change risk, and $\beta_{z,s}$ corresponds to the level risk.

There are several implications of expression (17) that are worthy of note. First, the cointegration parameter will be an important determinant of the asset's risk at any investment horizon. To illustrate this point, consider the beta for a one-period investment. Given Equation (17),

$$\beta_1 = \delta + [\Sigma_1(1, 2) + \Sigma_1(1, 5) + (\kappa_1 - 1)\Sigma_1(1, 3)]\Sigma_1(1, 1)^{-1}. \quad (18)$$

Notice that the cointegration parameter is one of the four components of the consumption beta at the one-period horizon. Further, the presence of cointegration not only alters betas directly (via the first term), it also alters the beta through the other terms (via its impact on the transition matrices A and G). Thus, the parameter of cointegration that reflects long-run consumption risks in the asset's cash flows is an important determinant of the risk-return relation at all investment horizons.

At long horizons, cash-flow cointegration risks are especially important. As noted above, variation in transitory components in returns dies out in the long run. Thus, variation in the error correction term and the growth in the price-dividend ratio will vanish in the limit. This means that, in the long run lr , the beta will become

$$\beta_{lr} = \delta + (\kappa_1 - 1)\Sigma_1(1, 3)\Sigma_1(1, 1)^{-1}. \quad (19)$$

That is, the beta becomes the cointegration parameter plus a term related to the accumulated price-dividend-level risk.

Expression (17) characterizes conditional consumption betas by horizon and highlights important implications of the cointegrating relation for risk measures. For example, when δ is assumed to be 1, as in Menzly, Santos, and Veronesi (2004), the cross-sectional dispersion in betas will be significantly reduced at

both short and long horizons. Further, when cointegration is not imposed, $\epsilon_{d,t}$ is not a predictive variable for dividend growth rates and returns. Consequently, the term δ is absent from the expression for betas at any horizon. The restriction of cointegration and the forecasting ability of the resulting error-correction variable for future growth rates and returns are empirical issues, which we investigate in Section 3.

The profile of expected returns by horizon can be read from the EC-VAR specification we have described, by adding half the variance on return innovation for a given horizon $\sigma_{r,s}^2/s$ to the mean log return. We use betas and mean returns by horizon, implied by the EC-VAR dynamics, to analyze the cross-sectional relation in risk and return at various investment horizons. Specifically, we investigate the relation by considering cross-sectional regressions for different horizons s :

$$E \left[r_i + 0.5 \frac{\sigma_{r_i,s}^2}{s} \right] = \lambda_{0,s} + \lambda_{1,s} \beta_{i,s}. \quad (20)$$

Evidence on the sources of betas at various horizons and the explanatory power of these betas for the cross section of mean returns is presented in Section 3.

2. Data

The portfolios employed in our empirical tests sort firms on dimensions that lead to cross-sectional dispersion in measured risk premia. The particular characteristics that we consider are firms' market value and book-to-market ratio. Our rationale for examining portfolios sorted on these characteristics is that size and book-to-market-based sorts are the basis for factor models used in Fama and French (1993) to explain the risk premia on other assets. Consequently, understanding the risk premia on these assets is an economically important step toward understanding the risk compensation of a wider array of assets.

We construct the set of portfolios formed on the basis of market capitalization by ranking all firms covered by CRSP on the basis of their market capitalization at the end of June of each year using NYSE capitalization breakpoints. We form annual returns on these portfolios over the period 1929 through 2002. In Table 1, we present means and standard deviations of market-value-weighted returns for size decile portfolios. The data evidence a substantial size premium over the sample period; the mean real annual return on the lowest decile firms is 13.45%, contrasted with a return of 7.58% for the highest decile.

Book-to-market portfolios are formed by ranking firms on their book-to-market ratios as of the end of June of each year using NYSE book-to-market breakpoints. Book values are computed using Moody's data prior to 1955 and Compustat data in the post-1955 period. 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

Table 1
Summary statistics

Portfolio	Cash flow growth		Returns	
	Mean	SD	Mean	SD
S1	0.0657	0.3969	0.1345	0.3388
S2	0.0829	0.4181	0.1235	0.2911
S3	0.0269	0.3518	0.1278	0.3163
S4	0.0238	0.2637	0.1256	0.2957
S5	0.0215	0.2463	0.1199	0.2797
S6	0.0168	0.2019	0.1106	0.2642
S7	0.0195	0.1684	0.1076	0.2517
S8	0.0120	0.1549	0.0994	0.2308
S9	0.0092	0.1303	0.0904	0.2181
S10	0.0034	0.1053	0.0758	0.1906
B1	-0.0026	0.1634	0.0701	0.2157
B2	0.0179	0.1598	0.0834	0.1902
B3	0.0041	0.1562	0.0727	0.1904
B4	0.0009	0.2194	0.0771	0.2224
B5	0.0160	0.1552	0.0963	0.2225
B6	0.0146	0.2461	0.0989	0.2305
B7	0.0165	0.2667	0.1069	0.2484
B8	0.0548	0.2194	0.1265	0.2187
B9	0.1009	0.3486	0.1360	0.2287
B10	0.0467	0.6799	0.1337	0.3315

This table presents descriptive statistics for the returns and cash flow growth rates on the 20 characteristic-sorted portfolios used in estimation. The portfolios examined are portfolios formed on market capitalization (S1–S10) and book-to-market ratio (B1–B10). Capitalization portfolios are formed by sorting NYSE, AMEX, and NASDAQ firms by their market capitalization as of June of each year (using NYSE breakpoints), and holding the capitalization decile constant for 1 year. Book-to-market portfolios are formed by sorting NYSE, AMEX, and NASDAQ firms based on their market capitalization as of June of each year divided by their book value as of the most recent fiscal-year-end available. Returns are value-weighted. The cash flow-growth rates are constructed by taking the first difference of the logarithm of dividend series. The data are converted to real using the PCE deflator. The data are sampled at the annual frequency, and cover the period 1929 through 2002, for a total of 74 observations.

value of equity at calendar year $t - 1$. Average-value-weighted portfolio returns are also presented in Table 1. The data evidence a book-to-market spread of similar magnitude to the size spread; the highest book-to-market firms earn average real-annual returns of 13.37%, whereas the lowest book-to-market firms average 7.01%.

We utilize the dividends paid on these value-weighted portfolios to explore the relations between portfolio cash flows and consumption. Our construction of the dividend series is standard; details of the construction can be found in Campbell and Shiller (1988); and Bansal, Dittmar, and Lundblad (2005). We construct the level of cash dividends per share, D_t , for the size and book-to-market portfolios on a monthly basis as described above. From this series, we construct the annual levels of dividends by summing the cash flows within the year. 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-flow series. Summary statistics for the cash-dividend growth rates of the portfolios under consideration are presented in Table 1. Earlier work shows that

Table 2
Cointegration parameters and ACF

Portfolio	δ	ACF (1)	ACF (5)	ACF (10)	Unit root
S1	9.62 (6.07)	0.90	0.34	-0.02	-2.09
S2	9.71 (5.06)	0.64	0.41	-0.01	-5.86*
S3	6.69 (2.86)	0.75	0.41	0.17	-3.32
S4	6.21 (2.37)	0.81	0.30	0.11	-2.75
S5	4.41 (1.70)	0.70	0.17	0.05	-3.65*
S6	4.38 (1.88)	0.84	0.28	0.20	-2.41
S7	2.42 (1.12)	0.72	0.24	0.14	-3.41
S8	2.38 (0.77)	0.65	-0.03	-0.16	-3.47
S9	2.38 (0.96)	0.81	0.22	0.03	-3.45
S10	0.82 (0.89)	0.79	0.24	-0.04	-3.71*
B1	-0.27 (1.28)	0.73	0.16	-0.13	-3.98*
B2	-2.59 (1.41)	0.72	-0.05	0.17	-3.72*
B3	-0.11 (0.79)	0.61	0.07	-0.02	-3.87*
B4	0.82 (1.91)	0.64	0.17	0.13	-3.77*
B5	2.79 (1.59)	0.91	0.61	0.17	-1.83
B6	4.83 (1.75)	0.68	-0.01	0.05	-5.07*
B7	6.36 (2.20)	0.76	0.20	0.07	-3.12
B8	9.70 (3.07)	0.79	0.21	0.16	-3.80*
B9	12.54 (5.95)	0.67	0.17	0.03	-4.02*
B10	10.25 (4.55)	0.77	0.42	0.09	-1.88

*Indicates significance at the 10% level (MacKinnon critical value = -3.59).

This table presents the estimates of cointegration (CI) parameters for the 20 portfolios sorted by market capitalization (S1-S10) and book-to-market ratio (B1-B10). CI parameter estimates are obtained by regressing the deterministically detrended log level of portfolio dividends on the detrended log level of aggregate consumption. Aggregate consumption is defined as seasonally adjusted real consumption of nondurables plus services. Consumption data are taken from the NIPA tables available from the Bureau of Economic Analysis. The three columns labeled "ACF" report sample autocorrelation function of the cointegrating residuals for lags 1, 5, and 10. The final column presents a unit root test for null that the error correction residuals contain a unit root. Test statistics that exceed the 10% critical threshold are indicated with an asterisk.

alternative measures of dividends, such as including repurchases, do not affect the results.⁴

3. Empirical Results

In this section, we investigate the implications of the preceding framework for the measurement of assets' risks. We first analyze the cointegration of assets' dividends with consumption and investigate the implications of the cointegrating relation for the predictability of assets' growth rates and returns. We then compute the profile of consumption betas and expected returns for different investment horizons, implied by our EC-VAR framework, and analyze the cross-sectional implications of the model.

3.1 Cointegration evidence

In Table 2, we present point estimates of the cointegration parameters between portfolios' cash flows and consumption, the sample autocorrelation functions

⁴ Bansal, Dittmar, and Lundblad (2005) show that alternative dividend measures, which include share repurchases, do not make a big difference to their cash-flow risk measures. We find that in this paper the same is true for the empirical evidence.

(ACF) of the cointegrating residuals, and unit root tests of the stationarity of the cointegrating residuals. As discussed earlier, we estimate cointegration parameters via OLS by regressing the deterministically detrended dividends on detrended consumption. We first note that, for the majority of portfolios analyzed, the sample autocorrelations of the resulting cointegrating residuals exhibit a relatively rapid decline. This supports our assumption that the long-run dynamics of portfolios' dividends and aggregate consumption are governed by the same permanent component that can be eliminated by the appropriate linear combination of the levels. In addition, the large cross-sectional variation in the estimated cointegration parameters, presented in the first column, suggests that assets' cash flows differ in their exposures to this low-frequency component. The unit-root tests suggest cointegration in 10 of the 20 portfolios, and a number of portfolios' test statistics are close to the MacKinnon critical value of -3.59 . Because of the low power of this test in samples of the size that we analyze (74 observations), we conclude that the results presented in Table 2 provide reasonable support for the cointegration specification.

We now examine the point estimates of the cointegration parameters more closely. Note that for the size portfolios, the parameters exhibit a near-monotonic decreasing pattern across the market capitalization deciles. The estimate for the small-size portfolio is 9.62 compared to 0.82 for the large-size portfolio, mirroring the pattern in observed risk premia. For the book-to-market portfolios, the cointegration parameters exhibit an increasing relation in the book-to-market decile. The point estimate for the highest book-to-market portfolio is 10.25, compared to -0.27 for the growth portfolio. Again, this result is broadly consistent with the pattern of observed risk premia. Indeed, we find that the cointegration parameter by itself explains about 81% of the cross-sectional variation in average one-period returns on size and book-to-market-sorted assets. The price of risk of 0.486 (SE = 0.053) is positive and statistically significant. This evidence suggests that long-run risks embodied in assets' cash flows are able to account for a significant portion of the differences in risk premia. The fact that small and high book-to-market stocks have large exposures to permanent risks in consumption implies that the performance of these firms is linked to the permanent risks in the economy while that of large and low book-to-market portfolio firms is not. As consumption is largely dominated by permanent shocks, risks in large market-capitalization and low book-to-market stocks are largely unrelated to the long-run evolution of the economy. This, as we document further in Section 3.2, is the exact reason why large and low book-to-market portfolios should bear a low *ex ante* risk premium and why they already highlight the importance of long-run risks.

3.2 Predictability evidence

As stated above, cointegration implies that dividend growth rates are predicted by the cointegrating residuals. That is, the current deviations of an asset's cash flows from their long-run relation with consumption should forecast the

Table 3
Predictability evidence: dividend growth

Portfolio	1 year		Horizon 5 years		10 years	
	EC-VAR	VAR	EC-VAR	VAR	EC-VAR	VAR
S1	0.00	0.04	0.14	0.20	0.21	0.22
S2	0.11	0.00	0.15	0.01	0.14	0.01
S3	0.44	0.52	0.10	0.25	0.05	0.29
S4	0.11	0.10	0.09	0.19	0.07	0.21
S5	0.11	0.07	0.13	0.20	0.10	0.24
S6	0.15	0.16	0.12	0.21	0.10	0.22
S7	0.12	0.08	0.21	0.22	0.17	0.25
S8	0.12	0.06	0.22	0.22	0.17	0.28
S9	0.17	0.15	0.24	0.29	0.17	0.30
S10	0.19	0.03	0.43	0.15	0.35	0.19
B1	0.27	0.06	0.55	0.21	0.46	0.29
B2	0.25	0.08	0.52	0.31	0.37	0.37
B3	0.24	0.12	0.54	0.40	0.48	0.48
B4	0.39	0.25	0.48	0.51	0.42	0.39
B5	0.14	0.19	0.24	0.31	0.21	0.24
B6	0.23	0.25	0.21	0.37	0.11	0.24
B7	0.34	0.37	0.15	0.34	0.07	0.22
B8	0.17	0.14	0.16	0.26	0.09	0.22
B9	0.13	0.24	0.24	0.40	0.15	0.26
B10	0.07	0.10	0.17	0.22	0.19	0.27

This table presents the adjusted R^2 for dividend projections implied by the EC-VAR specification and the growth-rate VAR model that does not assume the long-run relation between assets' cash flows and consumption. The entries are reported for the 20 portfolios sorted by market capitalization (S1–S10) and book-to-market ratio (B1–B10). Data are sampled at the annual frequency, expressed in real terms, and cover the period 1929–2002.

dynamics of dividend growth rates while dividends are moving back toward the equilibrium. For example, if dividends are unusually high today, dividend growth is expected to fall in order for cash flows to adjust to the stochastic trend in consumption. Given the approximation for the log-return in Equation (3), the predictability of dividend growth rates potentially translates into return predictability. The variation in the cointegrating residuals, therefore, may also be able to account for the variation in expected future returns.

We explore the ability of our EC-VAR specification to predict future dividend growth rates and returns at various horizons. To emphasize the importance of the cointegrating relation, we compare the adjusted R^2 values for dividend growth and return projections implied by the EC-VAR model outlined in Section 1.4 with the corresponding \bar{R}^2 values from the growth-rate VAR specification. The latter simply excludes the cointegrating residuals from a set of VAR variables. Results for dividend growth are presented in Table 3. We examine results at horizons of 1, 5, and 10 years. As shown in Table 3, at the one-year horizon, the differences between the EC-VAR specification and the simple growth-rate VAR are quite stark. The mean-(median)-adjusted R^2 for the EC-VAR specification is 0.19 (0.16), compared to 0.15 (0.11) for the growth-rate specification. At longer horizons, the inclusion of the error correction term becomes less important; at the 10-year horizon, the mean-(median)-adjusted R^2 for the EC-VAR and

Table 4
Predictability evidence: returns

Portfolio	1 year		Horizon 5 years		10 years	
	EC-VAR	VAR	EC-VAR	VAR	EC-VAR	VAR
S1	0.17	0.06	0.40	0.04	0.36	0.01
S2	0.13	0.11	0.53	0.35	0.59	0.44
S3	0.10	0.03	0.40	0.08	0.53	0.13
S4	0.15	0.10	0.44	0.18	0.50	0.20
S5	0.16	0.15	0.36	0.21	0.49	0.26
S6	0.11	0.07	0.37	0.17	0.47	0.18
S7	0.09	0.08	0.29	0.17	0.46	0.23
S8	0.08	0.05	0.32	0.14	0.53	0.21
S9	0.04	0.01	0.28	0.12	0.46	0.16
S10	0.03	0.02	0.23	0.15	0.43	0.24
B1	0.04	0.00	0.19	0.03	0.32	0.04
B2	0.07	0.00	0.23	0.04	0.36	0.06
B3	0.23	0.14	0.27	0.06	0.43	0.07
B4	0.06	-0.01	0.26	-0.02	0.44	-0.03
B5	0.04	0.02	0.24	0.07	0.37	0.06
B6	0.08	0.06	0.26	0.03	0.37	0.00
B7	0.15	0.06	0.44	0.11	0.42	0.05
B8	0.17	0.11	0.49	0.21	0.41	0.18
B9	0.20	0.06	0.42	0.03	0.35	-0.01
B10	0.20	0.09	0.48	0.19	0.51	0.17

This table presents the adjusted R^2 for return projections implied by the EC-VAR specification and the growth-rate VAR model that does not assume the long-run relation between assets' cash flows and consumption. The entries are reported for the 20 portfolios sorted by market capitalization (S1–S10) and book-to-market ratio (B1–B10). Data are sampled at the annual frequency, expressed in real terms, and cover the period 1929–2002.

VAR specification are 0.20 (0.17) and 0.26 (0.24), respectively. This is not surprising, as the error correction term captures transitory variation in dividend growth rates.

Further, asset return predictability is also altered by the cointegration between dividends and consumption. Return projections' \bar{R}^2 values for horizons 1, 5, and 10 years implied by the EC-VAR model, as well as the alternative growth-rate specification, are reported in Table 4. The EC-VAR specification, on average, is able to explain an average (median) of 11.5% (10.5%) of return variation at the one-year horizon, and 44.0% (43.5%) of the variation in 10-year returns. Excluding the cointegrating residual significantly lowers the predictability of asset returns and alters the conditional mean of returns. We illustrate this point in Figure 1 by plotting 1- and 10-year returns predicted by the EC-VAR specification, along with the forecasts implied by the alternative VAR model. Predicted conditional means are displayed for the top and bottom market capitalization and book-to-market portfolios. It can be seen that the two specifications produce quite different predictions of future expected returns, especially at longer horizons. That is, the cointegrating residual, included in the error-correction specification, contains distinct information about future returns beyond that in the growth-rate-based model. Return innovations, therefore,

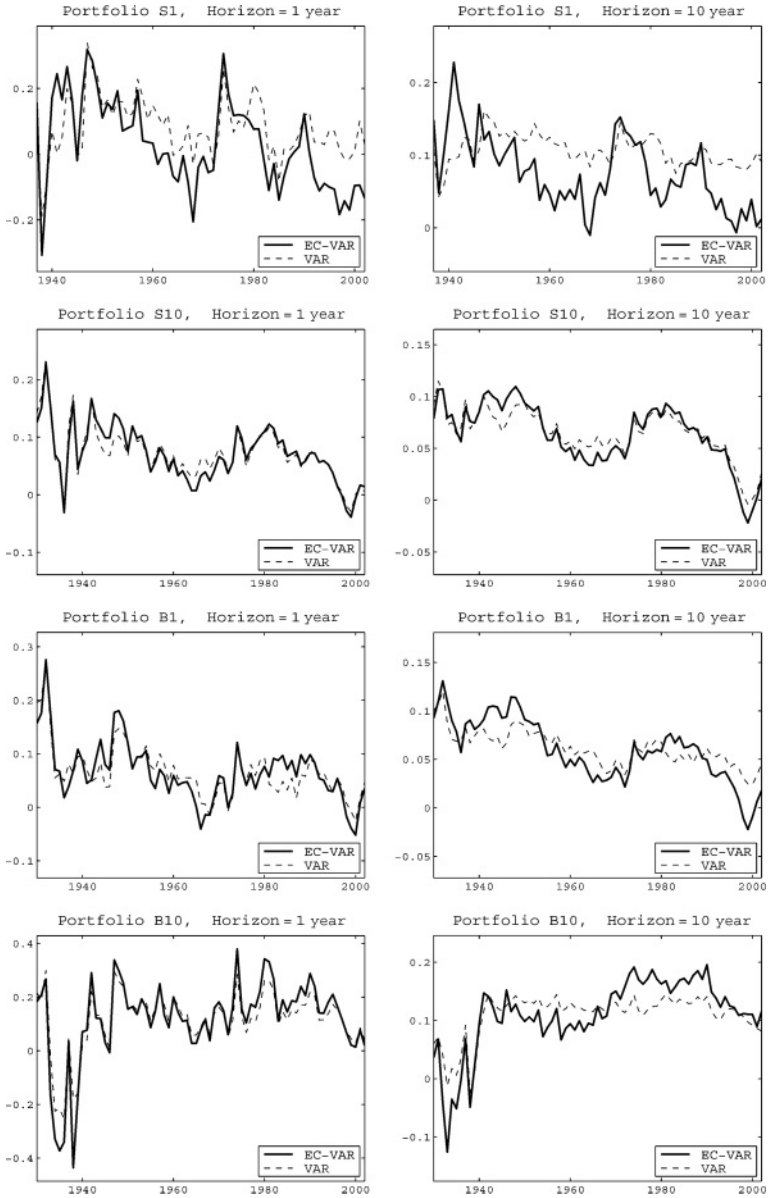


Figure 1
Expected returns

This figure plots 1- and 10-year returns predicted by the EC-VAR model and the alternative growth-rate VAR specification. The latter ignores the error correction information in predicting returns. Expected returns are plotted for small (S1) and large (S10) market capitalization firms, and low (B1) and high (B10) book-to-market portfolios.

also differ across the two specifications, and most importantly, so will the consumption betas measured from the two alternative models.

As a robustness check, we have also examined direct projections of multi-period compounded dividend growth rates and returns on the EC-VAR versus growth-rate VAR information sets. We do not entertain this method beyond 5–7 years as the number of independent observations in such direct projections decreases rapidly with the horizon. Predictability evidence from these projections is very similar to that discussed above and, for brevity, is not reported. In sum, we find cointegrating residuals to be a significant predictor of both future growth rates and future returns at short and intermediate horizons.⁵

The results of this section underscore the importance of cointegration in the measurement of risk and return. As emphasized in this section, temporary deviations of cash flows from the permanent component in consumption, that is, cointegration residuals, contain valuable information for predicting dividend growth rates and returns, and thus represent an important component in the calculation of expected returns and betas by horizon. We turn to this point in Section 3.3 and analyze how the risk-return relation changes with the investment horizon.

3.3 Betas and expected returns by horizon

Mean returns for the portfolios at the 1-, 5-, and 10-year horizons, implied by the EC-VAR, are presented in Table 5. As shown in the table, the general pattern observed in expected returns is preserved across the various horizons. Small firm portfolios tend to earn higher mean returns than do large firm portfolios, and low book-to-market portfolios earn lower expected returns than high book-to-market portfolios. Further, at all horizons, the mean returns exhibit considerable cross-sectional variation; however, the dispersion in mean returns at the 10-year horizon is slightly less than that at the one-year horizon.

We now explore the implications of cointegration for the determination of assets' consumption risks. Table 6 presents betas at various horizons for each of the portfolios. Similarly to mean returns, risk measures implied by the EC-VAR exhibit substantial cross-sectional variation. At the one-year horizon, the small-firm portfolio beta (4.12) exceeds the large-firm portfolio beta (1.54), and the high book-to-market beta (3.89) exceeds the low book-to-market beta (1.81). Further, there is a generally declining pattern in the size dimension and increasing pattern in the book-to-market dimension that is consistent with the pattern observed in mean returns. As the horizon increases, the cross-sectional pattern in these risk measures generally remains the same, although the precision of the estimates suffers. At the 10-year horizon, the small-firm

⁵ In particular, at the 1-year horizon, *t*-statistics on the cointegrating residual in dividend growth projections are significant for 11 out of 20 portfolios. As the horizon increases, the predictive power of the error correction variable improves considerably: at the 5-year horizon, robust *t*-statistics in dividend growth and return projections become statistically significant for virtually all the portfolios.

Table 5
Conditional mean returns

Portfolio	1 year		Horizon 5 years		10 years	
	EC-VAR	VAR	EC-VAR	VAR	EC-VAR	VAR
S1	11.74	11.72	10.65	11.69	10.74	11.67
S2	11.09	10.98	10.02	10.14	9.97	9.78
S3	11.71	11.57	10.39	10.98	10.02	10.58
S4	11.40	11.32	10.22	10.50	10.07	10.16
S5	10.78	10.75	9.97	10.14	9.58	9.72
S6	10.14	10.12	9.19	9.38	8.93	9.10
S7	9.86	9.85	9.13	9.24	8.68	8.84
S8	9.19	9.20	8.58	8.76	8.14	8.44
S9	8.47	8.47	7.83	7.97	7.44	7.66
S10	7.19	7.21	6.74	6.77	6.40	6.51
B1	6.61	6.64	6.32	6.37	5.99	6.19
B2	7.92	7.94	7.85	7.72	7.58	7.55
B3	6.60	6.70	6.68	6.63	6.36	6.45
B4	7.22	7.30	6.90	7.09	6.40	7.02
B5	9.09	9.10	8.38	8.57	8.02	8.29
B6	9.18	9.20	8.73	9.01	8.48	8.82
B7	9.96	9.84	9.05	9.22	9.20	9.06
B8	11.76	11.53	11.28	10.97	11.79	10.81
B9	12.89	12.45	12.80	12.36	13.57	12.40
B10	11.83	11.63	10.76	10.95	11.05	10.98

This table presents mean returns for investment horizons of 1, 5, and 10 years for each of the 20 portfolios sorted by market capitalization (S1–S10) and book-to-market ratio (B1–B10). Mean returns for a given horizon s are computed as $\overline{\ln(r_t)} + 0.5 \frac{\hat{\sigma}_{r_t, s}^2}{s}$. Data are sampled at the annual frequency and cover the period 1929–2002.

beta (6.54) continues to exceed the large-firm beta (0.34) and the high book-to-market beta (4.33) exceeds the low book-to-market beta (−0.83).

As argued in Section 1, the cointegration parameter is one of the components of the overall consumption beta at each horizon, along with risks arising from transitory fluctuations in dividends and prices. The relative importance of price risks and short-run dividend risks, however, decreases over time, and in the long run, systematic risks in returns should be dominated by permanent risks in assets' cash flows. We find that this theoretical proposition is strongly supported in the data. While at the one-year horizon the correlation between assets' betas and cointegration parameters is about 0.87, by the five-year horizon it already exceeds 0.90, reaching virtually 1 in the limit. Thus, the contribution of the cash-flow component of risks to the beta dominates in the long run. Moreover, even at short horizons, long-run consumption risks in dividends are an important determinant of assets' betas.

To illustrate the importance of the cointegrating relation between dividends and consumption for measuring expected returns and conditional consumption betas, we compare the evidence discussed above with that implied by the alternative growth-rate-based VAR specification. We find that the magnitude of mean returns for the growth-rate VAR specification, reported in Table 5, is somewhat higher relative to the EC-VAR specification but their cross-sectional

Table 6
Consumption betas by horizon

Portfolio	Unconditional	Horizon					
		1 year		5 years		10 years	
		EC-VAR	VAR	EC-VAR	VAR	EC-VAR	VAR
S1	0.71 (1.47)	4.12 (2.43)	1.77 (2.10)	4.51 (4.66)	-1.46 (3.94)	6.54 (4.26)	-1.55 (4.01)
S2	0.80 (1.38)	2.09 (1.12)	0.89 (1.09)	1.82 (2.65)	-1.24 (2.42)	4.21 (3.02)	0.13 (2.37)
S3	0.52 (1.36)	4.14 (1.29)	2.95 (1.98)	2.14 (2.38)	0.04 (2.44)	3.38 (2.29)	0.35 (2.86)
S4	0.77 (1.13)	3.52 (1.21)	2.66 (1.50)	2.09 (1.97)	0.79 (1.94)	3.56 (2.40)	1.59 (1.76)
S5	0.36 (1.21)	2.76 (0.99)	2.36 (1.22)	0.99 (1.72)	0.39 (1.52)	2.05 (2.12)	1.14 (1.37)
S6	0.64 (1.10)	3.07 (0.81)	2.58 (1.07)	1.42 (1.51)	0.71 (1.87)	2.50 (1.83)	1.41 (1.56)
S7	0.33 (1.10)	2.37 (0.86)	2.20 (0.82)	0.46 (1.37)	0.14 (1.28)	1.00 (1.67)	0.61 (1.18)
S8	-0.31 (1.14)	1.62 (0.68)	1.43 (0.76)	-0.12 (1.35)	-0.44 (1.62)	0.35 (1.68)	-0.13 (1.64)
S9	0.13 (1.09)	1.58 (0.77)	1.38 (0.80)	0.60 (1.41)	0.27 (1.62)	1.02 (1.59)	0.62 (1.49)
S10	0.69 (0.83)	1.54 (0.56)	1.64 (0.46)	0.31 (1.09)	0.51 (1.09)	0.34 (1.07)	0.67 (1.04)
B1	0.82 (0.97)	1.81 (0.54)	2.16 (0.33)	-0.58 (1.40)	0.31 (1.48)	-0.83 (1.32)	0.14 (1.63)
B2	-0.18 (0.84)	0.16 (0.56)	0.65 (0.39)	-1.69 (0.91)	-0.76 (0.98)	-2.05 (0.86)	-0.86 (1.06)
B3	-0.33 (0.84)	-0.09 (0.37)	0.32 (0.37)	-1.79 (0.77)	-1.34 (0.85)	-1.70 (0.84)	-1.33 (0.98)
B4	0.29 (1.10)	1.48 (1.29)	1.94 (1.50)	-0.67 (1.99)	-0.07 (2.10)	-0.59 (2.06)	-0.28 (2.22)
B5	0.27 (1.11)	1.94 (0.86)	1.67 (0.96)	1.18 (1.58)	0.81 (1.79)	1.60 (1.69)	1.10 (1.57)
B6	2.24 (1.01)	3.18 (1.49)	2.64 (1.97)	2.75 (2.42)	1.91 (1.96)	3.27 (2.37)	2.17 (1.77)
B7	0.21 (1.24)	2.74 (0.98)	1.46 (1.68)	2.67 (1.53)	1.01 (1.59)	4.22 (1.63)	1.70 (1.69)
B8	0.84 (1.23)	4.34 (1.83)	2.18 (2.14)	4.39 (2.72)	0.37 (2.22)	6.36 (2.69)	0.98 (2.34)
B9	-0.39 (1.54)	5.47 (2.31)	2.16 (2.91)	6.11 (3.75)	-1.55 (2.72)	8.32 (3.43)	-2.03 (2.79)
B10	0.14 (1.63)	3.89 (1.06)	2.87 (1.03)	2.14 (2.60)	0.54 (3.15)	4.33 (3.38)	1.49 (3.54)

This table presents consumption betas for investment horizons of 1, 5, and 10 years for each of the 20 portfolios sorted by market capitalization (S1–S10) and book-to-market ratio (B1–B10). In columns labeled “EC-VAR,” betas are measured using the error correction specification for consumption and asset returns. Columns labeled “VAR” present betas measured using a growth-rate VAR omitting the error correction information. These consumption betas are estimated as in Equation (17), using the covariance matrices implied by the relevant time-series model. The column labeled “Unconditional” represents the standard consumption beta. Robust standard errors are reported in parentheses. The number of lags used in the Newey and West (1987) covariance estimator is 8.

pattern by horizon is very comparable to that implied by the EC-VAR. However, the betas in the VAR specification, reported in Table 6, significantly differ from the EC-VAR-based betas. For comparison, in Table 6 we also present the unconditional C-CAPM betas. Neither the unconditional betas nor those based on the VAR reflect the cross-sectional differences in mean returns on size and book-to-market-sorted portfolios. This evidence underscores the importance of conditioning information contained in the cointegrating residual in computing assets’ exposures to consumption risks.

Below, we more formally analyze the relation between mean returns and risk measures in the cross section across various investment horizons.

3.4 Cross-sectional risk and return

In this section, we investigate the cross-sectional risk-return relation and explore the ability of various specifications to account for cross-sectional differences in mean returns across various investment horizons. We report and discuss evidence based on our preferred EC-VAR specification, as well as alternative specifications: the standard growth-rate-based VAR and the unconditional C-CAPM. Table 7 presents results of estimating the cross-sectional regression (20)

Table 7
Cross-sectional regressions by horizon

	Unconditional	1 year		Horizon 5 years		10 years	
		EC-VAR	VAR	EC-VAR	VAR	EC-VAR	VAR
$\lambda_{1,s}$	0.51	1.19	1.28	0.73	-0.31	0.65	-0.07
SE	(2.24)	(0.41)	(1.57)	(0.32)	(0.40)	(0.24)	(0.44)
\bar{R}^2	-0.04	0.75	0.22	0.73	-0.03	0.84	-0.05


This table presents results for cross-sectional regressions for a set of 10 portfolios sorted by market capitalization and 10 portfolios sorted by book-to-market ratio. The first row labeled $\lambda_{1,s}$ reports the estimated prices of risk. Consumption risk for different investment horizons is measured by the corresponding consumption beta. In columns labeled "EC-VAR," betas are measured using the error correction specification for consumption and asset returns. Columns labeled "VAR" present betas measured using a growth-rate VAR omitting the error correction information. Consumption betas are estimated as in Equation (17), using the covariance matrices implied by the relevant time series model. The column labeled "Unconditional" represents the standard consumption beta. All risk prices are expressed in annual percentage terms. Robust standard errors, reported in parentheses, are computed by estimating time-series and cross-sectional parameters in one step via GMM. The number of lags used in the Newey and West (1987) covariance estimator is 8.

for different investment horizons. The market prices of risk are estimated jointly with the time-series parameters via one-step GMM. The reported standard errors of the cross-sectional parameters, therefore, are robust to the estimation error in betas.⁶


As shown in Table 7, at the one-year horizon, betas implied by our EC-VAR specification explain 75% of the cross-sectional variation in mean returns with a positive price of risk of 1.19 (SE = 0.41). This explanatory power is maintained at the 5- and 10-year horizons, with adjusted R^2 of 0.73 and 0.84, respectively, and prices of risk of 0.73 (SE = 0.32) and 0.65 (SE = 0.24), respectively. At the very long horizon, i.e., $s = \infty$, the estimate of the market price of risk remains strongly significant at 0.72 (SE = 0.25), and the cross section of long-run consumption betas accounts for a sizeable portion of the variation in long-run risk premia ($\bar{R}^2 = 87\%$). Fits for the one-year horizon and the very long horizon are plotted in Figure 2. Thus, at all horizons, the EC-VAR specification explains most of the cross-sectional variation in mean returns across assets. This evidence manifests the empirical importance of the cointegrating relation between dividends and consumption in determining assets' risk premia not only at long but also at short horizons.

Table 7 also provides the cross-sectional risk-return trade-off comparison of the EC-VAR relative to the VAR-based specification. As we might expect, given the estimates in the preceding section, the VAR betas have almost no power in explaining the cross-sectional variation in mean returns. In particular, at the one-year horizon, it explains just about 22% of the cross-sectional variation in mean returns and implies positive, but insignificant, price of risk of 1.278 (SE = 1.57). As the horizon extends, however, the ability of the VAR specification deteriorates substantially. At both 5- and 10-year horizons, it produces

⁶ In addition, we have constructed bootstrap distributions of time-series and cross-sectional parameters of interest. As bootstrap-based standard errors are very similar to the above-discussed standard errors, for brevity, they are not reported.



(a) Short-run horizon ($s=1$)



(b) Long-run horizon

Figure 2
Cross-sectional fit

This figure displays the fit from the cross-sectional regressions for the investment horizon of 1 year, as well as in the long-run limit. The figures plot fitted expected returns, implied by the model, against mean returns.

negative prices of risks and completely fails to account for the differences in risk premia across assets, as indicated by adjusted R^2 of -0.02 and -0.05 , respectively.⁷

In addition to the conditional beta regressions, Table 7 presents the cross-sectional evidence for the unconditional C-CAPM for the one-year horizon. In contrast to our preferred EC-VAR specification, the unconditional C-CAPM is not able to account for the variation of average returns across the portfolios—the market price of risk is insignificant and the \bar{R}^2 is negative. As the horizon increases, the cross-sectional fit of these unconditional betas remains quite low. For example, at the five-year horizon, the explanatory power of the C-CAPM is only 23%. We do not look beyond the five-year horizon, as the number of independent observations in such multiperiod regressions shrinks rapidly with horizon.

To ensure our results are robust, we consider standard misspecification tests as in Jagannathan and Wang (1998). Specifically, in addition to our EC-VAR betas, we also include commonly used portfolio attributes: size and book-to-market. We find that the EC-VAR betas remain highly significant, while t -statistics on portfolio-specific characteristics are largely insignificant. In particular, at the first-year horizon, the robust t -statistic on the EC-VAR beta is 2.5; for size and book-to-market characteristics, the corresponding statistics are -1.9 and 0.4 , respectively. As the horizon increases, the explanatory power of the EC-VAR betas increases, while the significance of both characteristics diminishes. For example, at the 10-year horizon, the t -statistic for our betas is 2.6, and for size and book-to-market attributes are -0.6 and 0.2 , respectively. In sum, this finding shows that our EC-VAR-based betas are very important for capturing the dispersion in risks across assets.

Our empirical evidence highlights the importance of the cointegration-based specification in understanding the risk-return trade-off at all investment horizons. The EC-VAR relative to alternative specifications, such as the growth-rate-based VAR, incorporates the error correction term $\epsilon_{d,t}$ as a predictive variable. This error correction state variable alters the information set used to predict future returns relative to other specifications of return dynamics, and hence significantly alters the return innovation and conditional betas.

It is worth noting that the implicit assumption (the null hypothesis) behind the EC-VAR specification is that the error correction term is stationary—all the eigenvalues of the A matrix are inside the unit circle. At the point estimates, this critical restriction is satisfied for all 20 portfolios. If this restriction were not satisfied, then our betas could not be constructed for the entire cross section

⁷ We find that this evidence is fairly robust if we instead employ a double-sorted set of portfolios. In particular, using nine portfolios sorted on size and book-to-market characteristics delivers virtually the same magnitudes of market prices of risks as for single-sorted collections of assets, which we have discussed in detail. The same is true for the cross-sectional \bar{R}^2 values—while the EC-VAR is able to account for more than 50% of the cross-sectional variation in risk premia across various investment horizons, the adjusted R^2 values are virtually zero in the growth-rate VAR setting.

of assets and the cross-sectional price of risks could not be computed at all horizons. Hence, for evaluating the cross-sectional implications, the null hypothesis that EC-VAR specification leads to stationary dynamics is an important input. In contrast, in the standard VAR, the error correction mechanism is absent as dividends and consumption are not cointegrated and deviations between them contain a permanent component. This, as discussed earlier, is the key economic difference between these two specifications.

3.5 Long-run risk compensation

Since consumption growth is a covariance-stationary process, its level satisfies a Beveridge and Nelson (1981) decomposition. That is, the consumption process can be presented as a sum of a deterministic trend, a random walk component, and a transitory (stationary) component:

$$c_t = \mu_c t + T_t + S_t, \quad (21)$$

where T_t is the stochastic trend of log consumption and S_t is a transitory (or short-run) component. The covariance of the return with consumption (i.e., the beta) can, therefore, be broken into two parts: the covariance with trend shocks ($\beta_{T,s}$) and the covariance with transitory shocks in consumption ($\beta_{S,s}$):

$$\beta_{T,s} = \frac{\text{Cov}_t(r_{t+1 \rightarrow t+s}, T_{t+1 \rightarrow t+s})}{\sigma_{T,s}^2}, \quad \text{and} \quad \beta_{S,s} \equiv \frac{\text{Cov}_t(r_{t+1 \rightarrow t+s}, S_{t+1 \rightarrow t+s})}{\sigma_{S,s}^2}. \quad (22)$$

The overall beta is just a weighted combination of the two risk measures:

$$\beta_s \equiv \beta_{T,s} \frac{\sigma_{T,s}^2}{\sigma_{c,s}^2} + \beta_{S,s} \frac{\sigma_{S,s}^2}{\sigma_{c,s}^2}. \quad (23)$$

Consequently, at each horizon, the market price of risk reflects the premium for both very long risks and short-run fluctuations in consumption:

$$\frac{1}{s} E_t [r_{t+1 \rightarrow t+s} + 0.5\sigma_{r,s}^2 - sr_{t,s}^f] \equiv \lambda_s \beta_s = \lambda_{T,s} \beta_{T,s} + \lambda_{S,s} \beta_{S,s}. \quad (24)$$

As the horizon increases, transitory consumption shocks die out, and the transitory risk-compensation shrinks to zero. Thus, the total risk-compensation in the long-run limit (when $s = \infty$) provides a measure of the compensation solely for long-run consumption risks, i.e., $\lambda_\infty = \lambda_{T,\infty}$. We isolate and estimate this long-run compensation by considering the long-run risk-return relation in the cross section of assets. Subtracting it from the overall risk-compensation for a given horizon s allows us to construct the time-horizon profile of risk premia for transitory consumption risks.

We find that the compensation for long-run risks in consumption is about 60% of the overall compensation at the one-year horizon. Specifically, the estimate of the long-run risk compensation is 0.72 per annum compared to




Figure 3
Consumption risks compensation by horizon

This figure plots the market price of risk along with the profile of implied compensations for transitory (or short-run) consumption risk for investment horizons up to 15 years. The difference between the total and short-run compensations represents the premium for long-run consumption risks. The prices of risks are expressed in annual percentage terms.

1.19 at the one-period horizon. Figure 3 plots the compensation for short-run consumption risks, along with the total market price of risks for investment horizons up to 15 years. The compensation for transitory fluctuations is small relative to the premium for long-run risks and exhibits a rapid decline as the time horizon grows, starting at about 50 basis points at the first horizon and falling to zero by the 5th year. That is, long-run fluctuations in consumption are the dominant source of the premium for consumption risks in asset markets.

4. Robustness of Evidence

4.1 Alternative cointegration specifications

In this section, we discuss the implications of various restrictions on the cointegration specification presented above. In particular, we focus on three alternative specifications relative to our preferred cointegration specification where both δ and τ are unrestricted. In the first, we estimate Equations (12) and (14) under the restriction $\delta = 1$ for all assets and τ is unrestricted. In the second specification, we impose the restriction that $\tau = 0$ (no time trend) but δ is

Table 8
Cross-sectional regressions by horizon with cointegration restrictions

	1 year		Horizon 5 years		10 years	
	$\delta = 1$	$\tau = 0$	$\delta = 1$	$\tau = 0$	$\delta = 1$	$\tau = 0$
$\lambda_{1,s}$	1.58	1.55	0.79	1.20	1.12	1.12
SE	(0.39)	(0.28)	(0.54)	(0.31)	(0.41)	(0.19)
\bar{R}^2	0.45	0.62	0.06	0.43	0.26	0.64

This table presents results for cross-sectional regressions for a set of 10 portfolios sorted by market capitalization and 10 portfolios sorted by book-to-market ratio. The first row labeled $\lambda_{1,s}$ reports the estimated prices of risk. Consumption risk for different investment horizons is measured by the corresponding consumption beta. Betas are measured under two restrictions on the cointegrating relation between dividends and consumption:

$$d_t = \tau_0 + \tau t + \delta c_t + \epsilon_{d,t}$$

In the left columns, we restrict the cointegration parameter, $\delta = 1$, for all assets while still allowing for different time trends. In the right columns, we suppress the time trend in the cointegration specification, $\tau = 0$, but do not restrict the cointegration parameter δ . The beta for a given horizon is calculated from the error correction VAR model as in Equation (17). All risk prices are expressed in annual percentage terms. Robust standard errors, reported in parentheses, are computed by estimating time-series and cross-sectional parameters in one step via GMM. The number of lags used in the Newey and West (1987) covariance estimator is 8.

unrestricted. In the third case, we impose a joint restriction of $\tau = 0$ and $\delta = 1$ across all assets.⁸

Results for cross-sectional regressions incorporating these restrictions are presented in Table 8. The first set of columns presents results for the first specification, where $\delta = 1$ for all assets (while still allowing for differences in time trends). The results at the one-year horizon are somewhat weaker relative to our preferred case; the price of risk is 1.58 (SE = 0.39) and the betas explain 45% of the cross-sectional variation in average returns, as indicated by the \bar{R}^2 . However, as the horizon increases, the explanatory power of the specification deteriorates rapidly. At the two- and five-year horizons, the prices of risk are no longer statistically significant and the explanatory power of the regression is near zero. These results indicate that allowing for heterogeneity in the long-run risk in dividends is important for capturing variation in risk premia not only in the long run, but also in the short term.

In the second set of columns, we report results for the second specification, where δ is unrestricted and $\tau = 0$ (no time trend) for all assets. These results represent an improvement over the case in which the cointegration parameter is restricted to be one, and are mostly comparable to our preferred specification. At the one-year horizon, the estimated price of risk is positive and statistically significant and \bar{R}^2 is of 62%. The price of risk and the explanatory power

⁸ In the time-series dimension, these restrictions are rejected sharply for the majority of assets. In particular, according to Park (1992) test statistics, single hypotheses of $\tau = 0$ and $\delta = 1$ are rejected for 13 and 14 portfolios, respectively. The joint null of no time-trend and a unit cointegration parameter is rejected for all 20 assets (for all but one portfolio, at the 1% level). Empirical evidence provided in this section reinforces this analysis. We show that imposing restrictions on the long-run dynamics of assets' dividends and consumption significantly limits the ability of the cointegration-based betas to account for the cross-sectional variation in risk premia at all investment horizons.

of the regression remain high as the horizon lengthens. Finally, we also have considered a specification with $\tau = 0$ and $\delta = 1$ for all assets as in Menzly, Santos, and Veronesi (2004). This specification, however, is sharply rejected in the data and, for brevity, we do not report the detailed evidence.

A comparison of our preferred case (unrestricted time trend and cointegration parameter), with the second case where the time trend is eliminated, highlights the fact it is not the time trend *per se* but cross-sectional heterogeneity that is important for capturing differences in risks across assets at both long and short horizons. This is further reinforced by the evidence that when $\delta = 1$, the specification cannot account for the differences in risks at various horizons. As mentioned previously, in the second case ($\tau = 0$ and δ unrestricted), the cointegration parameter equals the ratio of the mean growth rate in the portfolio dividends to the mean growth of aggregate consumption, and hence may tautologically reflect average *ex post* returns. Including the time trend purges the effect of mean growth rates in dividends on the cointegration parameter, and ensures that long-run risk measures (that is, δ_s) do not mechanically replicate cross-sectional differences in *ex post* average returns. Despite this, if one chooses to restrict $\tau = 0$ for all assets and estimate δ_s under this restriction, this specification produces cross-sectional results that are comparable to our preferred case. This, once again, underscores the fact that it is appropriate heterogeneity in δ that is critical and not merely the inclusion or exclusion of time trends in the cointegration specification. Below, we further highlight the importance of the cross-sectional dispersion in long-run dividend risks δ_s , using Monte Carlo simulations.

4.2 Monte Carlo analysis

The results presented in Section 4.1 show that the price of risk and \bar{R}^2 are highly significant for our preferred EC-VAR specification. In this section, by means of Monte Carlo simulations, we show that our empirical evidence is robust to alternative specifications and is not likely to be an outcome of a lucky draw. We consider four different Monte Carlos, which highlight different aspects of our empirical evidence.

As noted earlier, the cross-sectional heterogeneity in the cointegration parameter is important. Our first Monte Carlo experiment (MC-1) is designed to evaluate if this heterogeneity could arise in a setup where, in fact, there is none due to either small sample errors or overfitting in small samples. In this experiment, the population value of the cointegration parameter is set at 1 for all assets. We ask the question if such an economy is capable of replicating the cross-sectional evidence that we find across various investment horizons. We simulate all the data, of the sample length, from the EC-VAR specification that imposes the restriction $\delta = 1$. Using the simulated data, we estimate the cointegration parameter and the unrestricted EC-VAR, as we have done in the data. We then use the constructed betas to run cross-sectional regressions and to compute the prices of consumption risk and the \bar{R}^2 by horizon. Table 9 reports

Table 9
Monte Carlo 1: alternative $\delta = 1$

	2.5%	5%	50%	95%	97.5%
1-year horizon					
$\lambda_{1,1}$	0.10	0.16	0.43	0.68	0.74
t -stat	0.37	0.67	1.83	2.68	3.74
\bar{R}^2	-0.04	-0.01	0.24	0.54	0.58
10-years horizon					
$\lambda_{1,1}$	-0.07	0.00	0.25	0.54	0.59
t -stat	-0.61	-0.21	1.05	2.54	3.59
\bar{R}^2	-0.05	-0.05	0.07	0.38	0.45
Cross-section with δ					
$\lambda_{1,1}$	0.01	0.05	0.22	0.47	0.53
t -stat	0.05	0.33	1.52	2.47	3.13
\bar{R}^2	-0.05	-0.05	0.12	0.37	0.41

This table reports the first Monte Carlo experiment (MC-1). It presents distributions of cross-sectional statistics based on asset betas computed using an unrestricted EC-VAR; the data are simulated from a specification where the cointegration parameter δ is restricted to be one for all assets. We simulate samples of 74 annual time-series observations of consumption growth, dividend growth, and price-dividend from the EC-VAR, using estimates of A and G , and the covariance matrix Σ . Using these simulated data, we reestimate the EC-VAR without imposing the restriction. See Section 4.2 for further details.

the cross-sectional evidence from this simulation exercise. The table shows that the conditional consumption betas in this case fail to explain the cross-sectional differences in mean returns across almost all investment horizons—the cross-sectional \bar{R}^2 values are very small and the cross-sectional slope coefficients are mostly insignificant. This Monte Carlo experiment, therefore, suggests that our empirical evidence is unlikely to come from a world where long-run cross-sectional dividend heterogeneity is absent. In addition, it underscores the importance of cross-sectional differences in the cointegration parameter for explaining differences in mean returns across all horizons.

In the second Monte Carlo (MC-2), we highlight the importance of the EC-VAR specification relative to the standard VAR specification. As is well known (see Granger and Engle, 1987), imposing a growth-rate VAR structure when dividends and consumption are cointegrated entails loss of information that would otherwise emanate from the error correction mechanism. This can substantially affect the transition dynamics of dividend growth rates and returns and, consequently, the asset’s consumption betas that interest us. The null model that we simulate from is the EC-VAR; we then estimate the standard growth-rate-based VAR (using the same specification as discussed earlier) to measure asset betas, price of risk, and the cross-sectional \bar{R}^2 . The evidence reported in Table 10 shows that the VAR loses considerable information and, on average, cannot account for the cross-sectional differences in mean returns. In all, this suggests that if the data indeed have an EC-VAR structure, then, as documented in our empirical section, the standard-VAR specification (commonly used in empirical work) will fail to measure asset risks.

The third Monte Carlo (MC-3) sets the standard growth-rate VAR as the null model, and asks what one would find if one estimated our preferred EC-VAR

Table 10
Monte Carlo 2

	2.5%	5%	50%	95%	97.5%
	1-year horizon				
$\lambda_{1,1}$	-0.24	-0.13	0.39	0.96	1.11
t -stat	-0.90	-0.63	1.05	2.87	3.33
\bar{R}^2	-0.10	-0.10	0.06	0.35	0.41
	10-years horizon				
$\lambda_{1,10}$	-0.05	0.03	0.51	1.03	1.15
t -stat	-0.43	-0.09	1.42	3.43	3.95
\bar{R}^2	-0.09	-0.07	0.14	0.43	0.49

This table reports distributions of cross-sectional statistics from the second Monte Carlo experiment (MC-2). We simulate from our preferred unrestricted EC-VAR but asset betas are computed from a standard VAR specification. The cross-sectional prices of risks and other statistics are computed using the betas constructed from the growth-rate-based VAR. See Section 4.2 for further details.

model.⁹ It is important to note that under the null of the VAR specification, consumption and dividend growth rates are stationary. However, the level of dividend is not cointegrated with the consumption level as they do not share a common stochastic trend. A stationary error correction variable does not exist as the deviation between dividends and consumption contains a unit root (see Granger and Newbold, 1974; Phillips, 1986; and Hamilton, 1994). As the number of time-series observations increases, this unit root will become easier to detect (technically, one of the roots of A matrix will be outside the unit circle for all the assets). This implies that our EC-VAR-based consumption betas will not exist in large enough samples. Our focus here, though, is on the modest sample length we observe in the data. We simulate sample lengths of data for 20 portfolios and then attempt to estimate our preferred EC-VAR specification.

We find that in each sample we simulate, the A matrix for at least one of the portfolios implies an explosive path for the estimated error-correction variable. On average, about three assets have explosive dynamics and, thus, the EC-VAR conditional betas do not exist for the whole cross section. Hence, the cross-sectional regression of mean returns on the asset betas cannot be conducted. That is, in this set-up, the distribution of various cross-sectional statistics of interest does not exist even in our modest-length samples.¹⁰ In addition, we have examined the Monte Carlo distribution of the first-order autocorrelation [ACF(1)] of the estimated cointegration residual from MC-3. We find that the ACF(1) coefficient, even in our modest samples, is very high relative to the data, suggesting why some portfolios in a given draw have explosive dynamics. More specifically, for all the portfolios but one, the estimates of the first-order autocorrelation coefficient in the data, reported in Table 2, are well in the left

⁹ Notice that a special case of the null is the *i.i.d.* growth rates specification. This nested case brings us back to the unconditional C-CAPM specification that, as we have shown, is strongly rejected in the data.

¹⁰ In particular, the fraction of draws without at least one portfolio having explosive dynamics is zero; the percentage of draws with at least one explosive portfolio is 17%, with at least two explosive portfolios is 35%, and with at least three explosive portfolios is 25%. Additionally, probabilities of observing at least four, five, and six explosive portfolios are 16%, 5%, and 2%; for seven portfolios and above, it is essentially zero.

tail (below the 2.5% percentile) of the Monte Carlo distribution. This evidence, and the fact that in the data not a single asset has an explosive path at our point estimates, suggests that the population model of a standard growth-rate-based VAR is an unlikely description of the observed data. Thus, our findings based on the EC-VAR specification are not a lucky draw.

Finally, to ensure that we do not have a lot of degrees of freedom with the considered collection of assets, we have constructed a Monte Carlo-based distribution for the cross-sectional \bar{R}^2 and t -statistics (MC-4). We simulate consumption growth rates of the sample size we observe and replace the observed consumption data with these simulated draws; the rest of the data (returns, pd ratios, etc.) are identical to what we observe. The EC-VAR-based consumption betas estimated using this simulated consumption should provide no information about the cross-sectional distribution of mean returns as the simulated consumption is just a random factor (this is our counterpart to the exercise considered in Lewellen, Nagel, and Shanken, 2006). We find that our point estimate for the \bar{R}^2 and the t -stat on the slope coefficient lie above the 95% probability cutoff of the corresponding Monte Carlo distributions. This Monte Carlo evidence, therefore, shows that the empirical findings, using our asset menu, are both statistically and economically significant and cannot simply be regarded as good luck.

To summarize, the Monte Carlo results corroborate the empirical evidence presented in the paper and show that these findings are robust against many alternatives.

5. Conclusion

We show that the long-run relation between consumption and dividends, that is, the cointegration, is important for understanding the dynamics of asset returns and their risk compensations across investment horizons. Cointegration measures the long-run covariance consumption risk in dividends and implies that deviations between dividends and consumption (i.e., the error correction) are temporary. An important implication of cointegrating relation is that returns can be characterized by an EC-VAR, in which returns can be predicted by the error correction variable. Hence, the error correction mechanism alters the transition dynamics of returns and, hence, the conditional consumption betas by horizon in interesting ways relative to a commonly used standard VAR setup. We demonstrate that the cointegration vastly improves the consumption-based model's ability to explain risk premia on size and book-to-market-sorted assets at both short and long horizons.

At the one-year horizon, the cross-sectional explanatory power rises from an adjusted- R^2 of less than zero for the unconditional C-CAPM and 22% for a growth-rate-based conditional C-CAPM, to 75% for a C-CAPM based on a cointegration specification. At long horizons, these differences in favor of the EC-VAR framework are even more dramatic. At a conceptual level, we

show that the conditional consumption betas at long horizons are determined largely by the cointegration parameter between dividends and consumption. Alternative dividend growth models, which do not impose cointegration or impose unit cointegration across assets, are not supported in the data. This points to the importance of cross-sectional differences in long-run consumption risks for explaining the risk-return trade-off.

Our approach allows us to assess the importance of long-run consumption risks relative to the overall consumption uncertainty. At all but the very long horizon, the cross-sectional slope coefficient from projecting mean returns on consumption betas reflects the risk compensation for both short- and long-run fluctuations in aggregate consumption. In the limit, on the other hand, it provides a market measure of risk compensation solely for long-run consumption risks. We find it to be a dominant source of the total risk premium, which, once again, points strongly toward the importance of long-run consumption risks in asset markets. This, along with the cross-sectional evidence, suggests that a more primitive production-based multisector model, which yields differences in long-run cointegration risks across assets, would be a very valuable contribution to the literature.

Appendix

In this appendix, we provide the details of the EC-VAR structure employed in the paper and the calculation of the horizon-dependent covariance matrix. Given estimates of the parameters and residuals in the cointegrating relation (12) between dividends and consumption, we model the dynamics of the resulting cointegrating residual $\epsilon_{d,t}$ jointly with the portfolio's price-dividend ratio z_t and consumption growth Δc_t by the following EC-VAR structure:

$$\begin{pmatrix} \Delta c_t \\ \epsilon_{d,t} \\ z_t \\ \Delta d_t \\ \Delta z_t \end{pmatrix} = \begin{pmatrix} \rho_c & 0 & 0 & 0 & 0 \\ a_{\epsilon c} & \rho_\epsilon & a_{\epsilon z} & a_{\epsilon d} & 0 \\ a_{z c} & a_{z \epsilon} & \rho_z & a_{z d} & 0 \\ a_{\epsilon c} + \delta \rho_c & (\rho_\epsilon - 1) & a_{\epsilon z} & a_{\epsilon d} & 0 \\ a_{z c} & a_{z \epsilon} & (\rho_z - 1) & a_{z d} & 0 \end{pmatrix} \begin{pmatrix} \Delta c_{t-1} \\ \epsilon_{d,t-1} \\ z_{t-1} \\ \Delta d_{t-1} \\ \Delta z_{t-1} \end{pmatrix} + \begin{pmatrix} \eta_t \\ \eta_{\epsilon,t} \\ \eta_{z,t} \\ \eta_{\epsilon,t} + \delta \eta_t \\ \eta_{z,t} \end{pmatrix}. \quad (\text{A1})$$

The first three variables form the basis of the EC-VAR process. The last two variables provide no additional information; they are included in the EC-VAR to simplify the computation of long-run covariances. The dynamics of Δd_t and Δz_t are derived from the dynamics of the first three variables by exploiting $\Delta d_t = \Delta \epsilon_{d,t} + \delta \Delta c_t$ and $\Delta z_t = z_t - z_{t-1}$.

Denoting $X'_t = (\Delta c_t \ \epsilon_{d,t} \ z_t \ \Delta d_t \ \Delta z_t)$, we can rewrite the EC-VAR compactly as

$$X_t = AX_{t-1} + Gu_t, \tag{A2}$$

where the matrix A is defined above, G is a 5×3 matrix, and u is a 3×1 matrix of shocks, $u'_t = (\eta_t \ \eta_{\epsilon,t} \ \eta_{z,t})$, that is,

$$X_t - AX_{t-1} = Gu_t \equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \delta & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t \\ \eta_{\epsilon,t} \\ \eta_{z,t} \end{pmatrix}. \tag{A3}$$

Given this structure and a horizon $s \geq 1$, the innovation in the sum of s consecutive X can be extracted as follows:

$$\sum_{j=1}^s X_{t+j} - E_t \left[\sum_{j=1}^s X_{t+j} \right] \equiv \zeta_{t,t+s}, \tag{A4}$$

where $\zeta_{t,t+s}$ is

$$\zeta_{t,t+s} = \sum_{j=1}^s B_j e_{t+1+s-j} \tag{A5}$$

with $e_t \equiv Gu_t$, $B_j = B_{j-1} + A^{j-1}$, and $B_0 = 0$; for $j = 1, \dots, s$.

Exploiting the fact that the errors are identically distributed and uncorrelated, the covariance matrix of $\zeta_{t,t+s}$ for any given horizon s is

$$\Sigma_s^* = B_s \Sigma_e B'_s + \Sigma_{s-1}^*, \tag{A6}$$

where $\Sigma_e = G \Sigma_u G'$ and $\Sigma_0^* = B_0 \Sigma_e B'_0 = 0$. As s increases, Σ_s^* grows without bound; hence we consider $\Sigma_s \equiv \frac{\Sigma_s^*}{s}$, that is, the covariance matrix of $\zeta_{t,t+s}$ scaled by the horizon. Given Σ_e and B_s , the evolution of Σ_s is given by

$$\Sigma_s = \frac{1}{s} B_s \Sigma_e B'_s + \left(1 - \frac{1}{s} \right) \Sigma_{s-1}. \tag{A7}$$

Equation (A7) provides a direct recursive algorithm for the construction of the covariance matrix of interest. For large s , the long-run matrix is determined by the limit of B_s , which is $[I - A]^{-1}$, i.e.,

$$\Sigma_{lr} = [I - A]^{-1} \Sigma_e [I - A]^{-1'}. \tag{A8}$$

References

Bansal, R., R. F. Dittmar, and C. Lundblad. 2001. Consumption, Dividends, and the Cross Section of Equity Returns. Working Paper, Duke University.

Bansal, R., R. F. Dittmar, and C. Lundblad. 2005. Consumption, Dividends, and the Cross Section of Equity Returns. *Journal of Finance* 60:1639–72.

- Bansal, R., R. A. Gallant, and G. Tauchen. 2007. Rational Pessimism, Rational Exuberance, and Asset Pricing Models. *Review of Economic Studies* 74:1005–33.
- Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481–509.
- Beveridge, S., and C. R. Nelson. 1981. A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle. *Journal of Monetary Economics* 7:151–74.
- Breeden, D. T. 1979. An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities. *Journal of Financial Economics* 7:265–96.
- Campbell, J., and R. Shiller. 1988. Stock Prices, Earnings, and Expected Dividends. *Journal of Finance* 43:661–76.
- Campbell, J. Y., and J. H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205–55.
- Campbell, J. Y., and P. Perron. 1991. Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots. *NBER Macroeconomics Annual*: 141–201.
- Cochrane, J., F. A. Longstaff, and P. Santa-Clara. 2006. Two Trees. Working Paper, University of Chicago and NBER.
- Fama, E., and K. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33:3–56.
- Granger, C. W., and R. Engle. 1987. Cointegration and Error-Correction: Representation, Estimation, and Testing. *Econometrica* 55:251–76.
- Granger, C. W., and P. Newbold. 1974. Spurious Regressions in Econometrics. *Journal of Econometrics* 2:111–20.
- Hamilton, J. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Hansen, L. P., J. C. Heaton, and N. Li. 2006. Consumption Strikes Back? Working Paper, University of Chicago.
- Hansen, L. P., and K. Singleton. 1983. Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns. *Journal of Political Economy* 91:249–65.
- Jagannathan, R., and Z. Wang. 1996. The Conditional CAPM and the Cross Section of Expected Returns. *Journal of Finance* 51:3–54.
- Jagannathan, R., and Z. Wang. 1998. An Asymptotic Theory for Testing Beta-Pricing Models Using Cross-Sectional Regression. *Journal of Finance* 53:1285–309.
- Lettau, M., and S. Ludvigson. 2001. Resurrecting the (C)CAPM: A Cross-Sectional Test when Risk Premia Are Time-Varying. *Journal of Political Economy* 109:1238–87.
- Lewellen, J., S. Nagel, and J. Shanken. 2006. A Skeptical Appraisal of Asset-Pricing Tests. NBER Working Paper 12360.
- Lucas, R. 1978. Asset Prices in an Exchange Economy. *Econometrica* 46:1429–46.
- Menzly, L., T. Santos, and P. Veronesi. 2004. Understanding Predictability. *Journal of Political Economy* 112(1):1–47.
- Newey, W., and K. West. 1987. A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55:703–8.
- Park, J. Y. 1992. Canonical Cointegrating Regressions. *Econometrica* 60:119–43.
- Parker, J. A., and C. Julliard. 2005. Consumption Risk and the Cross-Section of Expected Returns. *Journal of Political Economy* 113(1):185–222.
- Phillips, P. C. 1986. Understanding Spurious Regressions in Econometrics. *Journal of Econometrics* 33:311–40.