

Long Run Risks and Equity Returns

Ravi Bansal, Robert Dittmar, and Dana Kiku *

June 2005

Comments Welcome

*Bansal (email: ravi.bansal@duke.edu) is affiliated with the Fuqua School of Business, Duke University, Dittmar (rdittmar@umich.edu) is at the Stephen M. Ross School of Business, University of Michigan, and Kiku (email: dak6@duke.edu) is at the Department of Economics, Duke University. We would like to thank seminar participants at the University of British Columbia and Vanderbilt University for helpful comments. The usual disclaimer applies.

Long Run Risks and Equity Returns

Abstract

We argue that investor concerns about the exposure of asset returns to permanent movements in consumption levels are a key determinant of risk and return relation in assets markets. We show that as the investment horizon increases, (i) the return's systematic risk exposure (consumption beta) almost converges to the long-run relation between dividends and consumption, (ii) return volatility is increasingly dominated by dividend shocks. We find that most of the differences in risk premia, at short and long horizons, is due to the heterogeneity in the exposure to permanent risks in consumption. The long-run cross-sectional relation between risk and return provides a measure of the compensation for permanent risks in consumption. We find that the market compensation for these risks is large relative to that for transitory movements in consumption.

1 Introduction

The concerns of a short-run and long-run investor can be quite different, and these differences in concerns can provide new insights about the sources of risks in asset markets. Long-run investors care about permanent (or long-run) risks, while short-run investors care about both transitory and long-run risks. We show that indeed the nature of risks in the long run differ from those in the short run, and so do the implications for the risk and return relation. The exposure of dividend levels to permanent risks in consumption, a measure of long-run risks, plays an important role in determining the risk and return relation, both at long and short horizons. Our approach of focusing on risks at various horizons also provides a novel way to measure the market compensation for permanent risks in consumption. We find that, quantitatively, long-run risks in consumption are very important (relative to transitory risks) for interpreting the risk compensation in equity markets at all horizons.

We confine attention to risks in aggregate consumption as in Lucas (1978), Breeden (1979), Hansen and Singleton (1982) and others. Our intuition is based on considering the concerns of a long-run investor. We show that, as the investor's investment horizon increases, the return's consumption beta essentially converges to the long-run dividend beta (i.e., the cointegration parameter between dividends and consumption). The long-run dividend beta measures the exposure of dividends to permanent movements in consumption. Consequently, the cross-sectional risk and return relation at the very long horizon provides a way to measure the compensation for permanent movements in consumption. At short horizons, both transitory movements in consumption as well as long-run risks impact the risk-return relation.

Time diversification alters the mean return to equity portfolios by horizon.¹ Our results show that the typical profile of mean returns is declining with horizon. At short horizons, price shocks explain most of the return variance, however, as the horizon increases, the percentage of return variation due to price shocks decreases and the return variance is dominated by dividend shocks. This effect is most pronounced for value and

¹When all returns and growth are *i.i.d.*, then long-run and short-run expected returns coincide and there are no horizon effects. However, in the data both returns and dividends, as documented in a wide range of papers, are predictable making the *i.i.d.* case empirically implausible.

small-market capitalization stocks. The increasing importance of dividend shocks in explaining return volatility shows that for a long-term investor, movements in dividends are the key source of risk, while at short horizons, price shocks play a more important role. Further, as transitory shocks wash out over longer horizons, and permanent shocks do not, the permanent exposure of dividends to consumption (i.e., the cointegration parameter) is the key concern in determining systematic risks compensation for a long-run investor. Our analysis shows that the long-run relation between consumption and dividends (the cointegration parameter) alters not only the consumption betas in the long run, but also significantly alters the consumption betas at short horizons. Hence, an appropriate specification of the long-run relation is important even for correctly measuring short-horizon consumption betas.

We measure the long-run relation between consumption and dividends via stochastic cointegration. Dividend growth rates, price-dividend ratios, and consumption growth are modelled using an error-correction specification of a vector autoregression (EC-VAR) model. Consumption growth, as in Bansal and Yaron (2004) and Hansen, Heaton, and Li (2004), is assumed to contain both transitory and permanent predictable components; it is not an *i.i.d* process. From the EC-VAR model we compute time-horizon profiles of mean asset returns and the exposure of the asset to consumption risks (consumption beta). We show that both term-structures are not flat — the mean return and the consumption beta change considerably with the investment horizon. As the horizon increases, the latter essentially (in practice almost exactly) converges to the cointegration parameter between dividends and consumption. Bansal, Dittmar, and Lundblad (2001), in their working paper version, also use stochastic cointegration to measure long run dividend betas, they do not, however, explore the risk return relation by horizon. Menzly, Santos, and Veronesi (2004) present a model where the cointegration parameter is restricted to one — we explore the implications of this restriction for expected returns as well.

At all horizons, our consumption betas are able to explain over 60% of the cross-sectional differences in mean returns. This result by itself shows that the consumption risk measures have economic content and can account for considerable cross-sectional heterogeneity in mean returns for a challenging menu of assets which includes size and book-to-market sorted portfolios. The long-run betas explain over 85% of the cross-section of mean returns. The nature of the long run relation between dividends and

consumption alters not only the risk measures at long horizons, it also impacts the risk-return relation at short horizons. Heterogeneity in the long-run dividend betas impacts the beta at short horizons. This highlights the importance of exploring asset pricing implications not only at the short horizon (as is typically done) but also at long horizons. Inclusion of the cointegrating residual between dividends and consumption significantly alters the predicted mean of returns and dividend growth rates at all horizons and increases the predictive time-series R^2 's. For this reason, including the cointegrating residual significantly alters the return innovation and hence the measured consumption betas at all horizons.

For each horizon, the cross-sectional slope coefficients from projecting the mean returns on the consumption betas measures the risk compensation for consumption risk. The regression of the long-run mean returns on the long-run betas provides a market measure of risk-compensation solely for permanent risks (long-run risks). At all but the very long horizon, this risk compensation reflects compensation for both transitory and permanent shocks in consumption. The difference between the compensation for consumption risks at shorter horizons versus those at the very long run provides a measure of the market compensation for transitory risks in consumption. We find that the compensation for transitory risks is small relative to that for permanent risks, and declines with horizon. These measures of the market compensation for business cycle and permanent risks relate to the voluminous literature on measuring the cost of business cycles initiated by Lucas (1987). Our approach does not measure the total cost for business cycles; it does show, however, that the market offers much higher risk compensation for long-run risks relative to transitory (business cycle) risks in consumption. This is consistent with the arguments presented in Alvarez and Jermann (2004a) that the cost of business cycle risks is small relative to the overall consumption uncertainty. Alvarez and Jermann (2004b) use asset pricing bounds to document that the variance of the consumption based pricing kernel is dominated by permanent shocks, suggesting that compensation for long run-risks should dominate in asset markets.

The rest of the paper is organized as follows. Section 2 discusses risk and return relations across different investment horizons. Section 3 provides an econometric framework for estimating the betas. Empirical results are reported and discussed in Section 4. Finally, Section 5 concludes.

2 Long and Short Run Risks

In this section, we develop a simple, general model of asset returns. Throughout our discussion, we utilize a Taylor series approximation to log returns (see Campbell and Shiller (1988)). Let $z_t = p_t - d_t$ represent the log price-dividend ratio. The return approximation is given by

$$r_{t+1} = \kappa_0 + d_{t+1} - d_t + \kappa_1 z_{t+1} - z_t. \quad (1)$$

Using $\Delta z_{t+1} = z_{t+1} - z_t$, the above equation also implies,

$$r_{t+1} = \kappa_0 + \Delta d_{t+1} + \Delta z_{t+1} + (\kappa_1 - 1)z_{t+1}. \quad (2)$$

Define the summed dividend yield as $z_{t+1 \rightarrow t+s} \equiv \sum_{j=1}^s z_{t+j}$, and the accumulated change in this yield as $\Delta z_{t+1 \rightarrow t+s} \equiv z_{t+1 \rightarrow t+s} - z_{t \rightarrow t+s-1}$. Finally, let the yield on an s -period discount bond at date t be given by $r_{t,s}^f$. It is further assumed that the log price dividend ratio, z_t , and dividend growth rates, Δd_{t+1} , are covariance stationary processes.

2.1 Risk Premia

For simplicity, we assume power utility preferences for the representative agent. The logarithm of the representative agent's marginal utility for period $t+s$ is $s \ln(\beta) - \gamma c_{t+s}$, where γ is the risk aversion parameter, and β is the subjective discount factor. The one-step-ahead stochastic discount factor, therefore, is $m_t = \ln(\beta) - \gamma \Delta c_t$. This yields the pricing restriction,

$$E_t[\exp(m_{t+1 \rightarrow t+s} + r_{t+1 \rightarrow t+s})] = 1, \quad (3)$$

with $m_{t+1 \rightarrow t+s} = \sum_{j=1}^s m_{t+j}$ and $r_{t+1 \rightarrow t+s} = \sum_{j=1}^s r_{t+j}$. We assume that the log of intertemporal marginal rate of substitution (IMRS) and the log returns are conditionally jointly normal and homoskedastic. We further assume that Δc_t is a covariance stationary process.

Exploiting the pricing condition (3), the per-unit time expected excess return for this

agent who holds the equity for s periods in a log-normal world is given by

$$\frac{1}{s}E_t[r_{t+1 \rightarrow t+s} + 0.5\sigma_{r,s}^2 - sr_{t,s}^f] = \gamma \frac{\sigma_{\Delta c,s}^2}{s} \frac{\text{cov}_t[c_{t+s} - E_t(c_{t+s}), r_{t+s} - E_t(r_{t+s})]}{\sigma_{\Delta c,s}^2} \quad (4)$$

The conditional variance of the return $r_{t+1 \rightarrow t+s}$ for horizon s is $\sigma_{r,s}^2$, similarly $\sigma_{\Delta c,s}^2$ is the conditional variance of the accumulated consumption growth process for horizon s , $\sum_{j=1}^s \Delta c_{t+j}$. Scaling by s ensures that the moments exist and provides the interpretation that all equity premia are per unit of time. The risk-return relation, by horizon, can be compactly stated as

$$\frac{1}{s}E_t[r_{t+1 \rightarrow t+s} + 0.5\sigma_{r,s}^2 - sr_{t,s}^f] = \beta_s \gamma \frac{\sigma_{\Delta c,s}^2}{s} \quad (5)$$

where β_s is the horizon-dependent asset risk.

The unconditional mean is simply derived by taking expectations of expression (5). Let μ_r be the one-period-ahead arithmetic mean return,

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

We refer to the market price of risk for horizon s as $\lambda_s \equiv \gamma \frac{\sigma_{\Delta c,s}^2}{s}$. As shown in expression (6), unconditional mean returns per unit of time depend on the horizon both through the market price of horizon risk, λ_s , and through the horizon-dependent beta, β_s . Expected returns for all horizons are identical only in the case when returns are *i.i.d.*. Expression (6) provides a term structure for the risk-return tradeoff. At all horizons, this term structure is determined by the horizon-dependent covariance of the asset's return with consumption, β_s .

2.2 Market Prices of Transitory and Long Run Risks

As discussed in the previous section, the sole source of risk in our model is consumption growth, assumed to be a covariance stationary process. Under this assumption, it then follows that the level of consumption satisfies a Beveridge and Nelson (1981) decomposition; that is, log consumption, c_t , equals the sum of a deterministic trend, a random

walk component (or stochastic trend), S_t , and a transitory (stationary) component, B_t ,

$$c_t = \mu_c t + \Psi(1)S_t + B_t. \quad (7)$$

The parameter $\Psi(1)$ determines the long-run variance of consumption (see Hamilton (1994)). Thus, two components comprise consumption; a long-run, permanent (trend) component and a short-run, transitory (stationary) component. As a result, the covariance of returns (beta) can be broken into two parts: the covariance with permanent shocks in consumption and the covariance with transitory shocks,

$$\begin{aligned} \beta_s &\equiv \frac{Cov_t(r_{t+1 \rightarrow t+s}, c_{t+1 \rightarrow t+s})}{\sigma_{\Delta c, s}^2} \\ &= \frac{Cov_t(r_{t+1 \rightarrow t+s}, \Psi(1)S_{t+1 \rightarrow t+s}) + Cov_t(r_{t+1 \rightarrow t+s}, B_{t+1 \rightarrow t+s})}{\sigma_{\Delta c, s}^2}. \end{aligned}$$

We express the asset beta with respect to trend shocks as $\beta_{\tau, s}$ and the asset beta with respect to transitory (or business cycle) shocks as $\beta_{bc, s}$.²

Given the above expression, the risk premium can be decomposed as follows,

$$\frac{1}{s} E_t[r_{t+1 \rightarrow t+s} + 0.5\sigma_{r, s}^2 - sr_{t, s}^f] = \beta_{\tau, s} \gamma \frac{\sigma_{\tau, s}^2}{s} + \beta_{bc, s} \gamma \frac{\sigma_{bc, s}^2}{s}. \quad (8)$$

where $\sigma_{\tau, s}^2$ and $\sigma_{bc, s}^2$ are the variances of cumulative shocks in the trend and transitory component of consumption respectively. An asset with a unit exposure to the trend risks, and a zero exposure to the business cycle risks would receive $\lambda_{\tau, s} \equiv \frac{\gamma \sigma_{\tau, s}^2}{s}$ as risk compensation. Similarly, $\lambda_{bc, s} \equiv \frac{\gamma \sigma_{bc, s}^2}{s}$ corresponds to the business cycle risk compensation for an asset with a zero trend and a unit business cycle risk. An asset with a unit exposure (unit beta) to both trend and cycle would have a unit consumption beta and receives the overall risk compensation of $\lambda_{\tau, s} + \lambda_{bc, s}$.

Notice that as horizon increases, transitory consumption shocks die out, and the transitory risk compensation shrinks to zero. The total risk compensation in the long-run

²That is,

$$\beta_{\tau, s} \equiv \frac{Cov_t(r_{t+1 \rightarrow t+s}, \Psi(1)S_{t+1 \rightarrow t+s})}{\sigma_{\tau, s}^2} \quad \text{and} \quad \beta_{bc, s} \equiv \frac{Cov_t(r_{t+1 \rightarrow t+s}, B_{t+1 \rightarrow t+s})}{\sigma_{bc, s}^2}.$$

limit, therefore, depends solely on the compensation for the trend risks in consumption,

$$\lim_{s \rightarrow \infty} \frac{1}{s} E_t[r_{t+1 \rightarrow t+s} + 0.5\sigma_{r,s}^2 - sr_{t,s}^f] = E[r_{t,lr} + 0.5\sigma_{r,lr}^2 - r_{t,lr}^f] = \lambda_{lr}\beta_\tau, \quad (9)$$

where $r_{t,lr}^f$ is the yield on the long-run bond, $\sigma_{\Delta c,lr}^2$ is the long-run variance of consumption, and $\sigma_{r,lr}^2$ is the long-run variance of the return. λ_{lr} measures the compensation for permanent risks, and given the Beveridge-Nelson decomposition, is equal to $\gamma\Psi(1)^2\sigma_\eta^2$. As follows from equation (9), the compensation for permanent risks in consumption can be isolated and estimated by considering the long-run risk-return relation in the cross-section of assets. This estimate, then, allows us to construct the time-profile of the compensation for the transitory (i.e., business cycle) risks by simply subtracting the limiting compensation (i.e., compensation for permanent risks) from the total risk compensation for a given horizon s :

$$\lambda_{bc,s} = \lambda_s - \lambda_{lr}. \quad (10)$$

where λ_s is the s -period market price of risk in the preceding section. The sources of the risks compensated by these terms are discussed in the following section.

2.3 Sources of Long and Short-Run Risks in Returns

The return approximation, (2), shows that returns can be decomposed into three components: growth in dividends, growth in price-dividend ratios, and the level of the price-dividend ratio. The covariance of these components with the long- and short-run components of consumption will determine the asset's long- and short-run risk compensation. Specifically, note that the asset beta, β_s , can be expressed as the sum of three covariances:

$$\begin{aligned} \beta_s &= \left[\frac{\text{cov}_t[c_{t+s} - E_t(c_{t+s}), d_{t+s} - E_t(d_{t+s})]}{\sigma_{\Delta c,s}^2} \right] \\ &+ \left[\frac{\text{cov}_t[c_{t+s} - E_t(c_{t+s}), \Delta z_{t+1 \rightarrow t+s} - E_t(\Delta z_{t+1 \rightarrow t+s})]}{\sigma_{\Delta c,s}^2} \right] \\ &+ \left[(\kappa_1 - 1) \frac{\text{cov}_t[c_{t+s} - E_t(c_{t+s}), z_{t+1 \rightarrow t+s} - E_t(z_{t+1 \rightarrow t+s})]}{\sigma_{\Delta c,s}^2} \right], \end{aligned} \quad (11)$$

which we express compactly as

$$\beta_s = [\beta_{d,s} + \beta_{\Delta z,s} + (\kappa_1 - 1)\beta_{z,s}]. \quad (12)$$

Thus, the asset beta can be considered the sum of three betas, related to cash flow growth and price-dividend ratios.

As stated, it is the covariance of dividend growth rates, changes in price-dividend ratios, and/or the level of price-dividend ratios with the long-run component of consumption that drives long-run risk compensation. Thus, consider the beta in the long-run limit. As s goes to infinity, the second term in equation (12) will converge to zero, as Δz_t is the difference of a stationary process (z_t is $I(0)$), and the last term will converge to a constant. As a practical matter, since the coefficient κ_1 is very close to one, the contribution of this last term will be quite small in the limit. Consequently, the only channel in returns that can feasibly drive long-run risk compensation is dividend growth, Δd_{t+s} , and its associated risk measure, $\beta_{d,s}$. We can operationalize this long-run relation in consumption and dividends by considering an economic setting in which consumption and dividends are cointegrated. That is, dividends may temporarily wander from a long-run relation with consumption, but will converge to this relation as the horizon grows. The cointegration parameter, δ , corresponds to the long-run dividend beta.

Consider the expression for the long run risk premium, that is the premium as s goes to infinity is,

$$E[r_{t,lr} + 0.5\sigma_{r,lr}^2 - r_{t,lr}^f] = [\delta + 0 + (\kappa_1 - 1)\beta_{z,lr}]\gamma\sigma_{\Delta c,lr}^2. \quad (13)$$

The long-run consumption beta of the asset, therefore, is $\beta_{lr} \equiv [\delta + 0 + (\kappa_1 - 1)\beta_{z,lr}]$. As mentioned before, since κ_1 is very close to one, the effect of the last term is quite small. That is, the beta of the asset in the long run approximately equals the cointegration parameter between dividends and consumption. Differences in risk premia in the long run, therefore, are likely to be dominated by differences in the cointegration relation between dividends and consumption.

This indicates that most of the risk in the long run is related to the risk in the permanent or near-permanent component in consumption and its shared component in dividends. More generally, as horizon increases, the effect of the price change on the asset

beta and the systematic risks should diminish, and a larger proportion of the risks would be associated with dividend risks. Comparing the risk premium in expressions (13) and (12) suggests that the short-horizon betas are affected by both business cycle risks and the permanent shocks to dividends and consumption; the long-horizon betas, however, essentially isolate the exposure of returns to the permanent component in consumption.

2.4 Alternative Preferences

In this section, we briefly discuss the impact of altering the assumed preference structure on our results. Our economic model is based on a simple CRRA specification of preferences. Several alternative preference structures have emerged in recent years to assist our understanding of the relation between agents' consumption decisions and asset prices.

We first consider Epstein and Zin (1989) preferences, for which the logarithm of the intertemporal marginal rate of substitution is given by,

$$m_{t+1} = \theta \ln \beta - \gamma \Delta c_{t+1} - (1 - \theta) \ln \left(\frac{1 + W_{c,t+1}}{W_{c,t}} \right). \quad (14)$$

The parameter $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$, where γ , as before, is the coefficient of relative risk aversion, and ψ is the intertemporal elasticity of substitution (IES); $W_{c,t}$ is the aggregate wealth to aggregate consumption ratio which is endogenous in the model and determined by the process for consumption growth. With these generalized preferences the predictable variation in consumption may significantly increase the compensation for long run risks — that is, agents with these preferences may require much larger compensation for permanent risks in consumption, relative to the power specification.³

Specifically, assume (as in Bansal and Yaron (2004)) that aggregate consumption growth follows an ARMA(1,1) process,

$$\Delta c_{t+1} = (1 - \rho)\mu_c + \rho\Delta c_t + \eta_{t+1} - \omega\eta_t. \quad (15)$$

³Bansal and Yaron (2004) account for the equity premium and other asset market puzzles by exploiting this channel. This provides motivation for long-run risks being a potentially important driver in asset markets. In general, all state variables that govern the consumption growth dynamics will determine the risk premia on the asset (see, for example, Bansal and Yaron (2004)). Hansen, Heaton, and Li (2004) also consider the generalized preferences to emphasize the importance of long-run risks.

Alternately stated,

$$\Delta c_{t+1} = \mu_c + x_t + \eta_{t+1}, \quad (16)$$

$$x_{t+1} = \rho x_t + (\rho - \omega)\eta_{t+1}, \quad (17)$$

where $\mu_c + x_t$ is the expected consumption growth rate. Bansal and Yaron (2004) show that the given dynamics of the state variable implies that the log of the IMRS is

$$-m_{t+1} = \bar{m} + \frac{1}{\psi}x_t + \lambda\eta_{t+1}, \quad (18)$$

where

$$\lambda = \gamma + (1 - \theta) \frac{1 - \frac{1}{\psi}}{1 - \kappa_1 \rho} (\rho - \omega) > 0, \quad (19)$$

and \bar{m} represents the unconditional mean of m .

To derive the implications for the risk-return relation, it is convenient to focus on the process for marginal utility at time t , which is defined as $mu_t = \sum_{j=0}^{\infty} m_{t-j}$ and is comprised of a deterministic, a permanent and a transitory (business cycle) components:

$$-mu_t = -\bar{m} * t + \left[\lambda + \frac{1}{\psi} \frac{\rho - \omega}{1 - \rho} \right] \sum_{j=0}^{\infty} \eta_{t-j} - \left[\frac{1}{\psi} \frac{\rho - \omega}{1 - \rho} \right] \sum_{j=0}^{\infty} \rho^j \eta_{t-j}. \quad (20)$$

The above marginal utility process can be connected back to the trend and stationary components in consumption as,

$$-mu_t = -\bar{m} * t + \left[\lambda + \frac{1}{\psi} \frac{\rho - \omega}{1 - \rho} \right] S_t + \frac{1}{\psi} B_t. \quad (21)$$

The market prices of risks are driven by the preference parameters. When $\theta = 1$ we get the power utility specification, one unit shock to S or B impact the marginal utility of the agent exactly the same way. When θ is not equal to one, a small value of the IES parameter, ψ , increases investors concerns about business cycle fluctuations leading to a higher compensation for transitory risks. Increasing the risk aversion raises λ and the importance of permanent risks. Another specialized case, with IES equal to one is considered by Hansen, Heaton, and Li (2004).

Another alternative is the habit formation model of Campbell and Cochrane (1999).

In the short run shocks to the habits may alter the risk return relation, however in long run risk return implication of this model are identical to the power utility case as the long run variance of marginal utility is driven by the permanent component in consumption.

In summary, it is straightforward, and to extend the framework beyond power utility. In order to do so, one must estimate an larger set of parameters in the cross-section of assets. In order to avoid this complication, and still capture first order effects in the cross-section of assets, we have confined attention to the simple power utility case. As in Lucas (1987), by using power utility we are able to directly measure the role of permanent and transitory risks in observed aggregate consumption.

3 Estimation and Econometric Framework

3.1 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 I, 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 value of equity at calendar year $t - 1$. Average value-weighted portfolio returns are also presented in Table I. 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), Bansal, Dittmar, and Lundblad (2005), and Hansen, Heaton, and Li (2004). We construct the level of cash *dividends per share*, D_t , for the size and book-to-market portfolios on a monthly basis.⁴ 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 I. Earlier work shows that alternative measures of dividends, such as including repurchases, do not affect the results.⁵

3.2 Cointegration and Dividends

A critical element in our specification is the measurement of the long-run covariance between dividends per share and consumption. We rely on stochastic cointegration (see Campbell and Perron (1991)) as a way to measure the dividend’s long run consumption beta. That is, we utilize the following specification:

$$d_t = \tau_0 + \tau_1 t + \delta c_t + \epsilon_{d,t}. \tag{22}$$

⁴The dividends per share are constructed as follows. The total return per dollar invested is $R_{t+1} = H_{t+1} + Y_{t+1}$, where H_{t+1} is the price appreciation and Y_{t+1} is the dividend yield. Then, the level of the dividends per share can be computed as $D_{t+1} = Y_{t+1} V_t$, where $V_{t+1} = H_{t+1} V_t$ and $V_0 = 1$.

⁵Bansal, Dittmar, and Lundblad (2004) show that alternative dividend measures, which include share repurchases do not make a big difference to their cash-flow risk measures. We find the same is true for the empirical evidence in this paper.

The level of dividends and consumption in logs is d_t and c_t respectively. The long-run dividend beta measure is δ and $E[\epsilon_{d,t}] = 0$. From equation (3.2) it also follows that $\tau_1 = \mu_d - \delta\mu_c$, where μ_c and μ_d are the average growth of consumption and dividends respectively. Substituting for τ_1 , the above equation can equivalently be stated as

$$d_t - \mu_d t = \tau_0 + \delta(c_t - \mu_c t) + \epsilon_{d,t} \quad (23)$$

That is, the long run parameter δ can be estimated via the projection of deterministically detrended dividends on detrended consumption. Alternatively, δ can be estimated by regressing d_t on a time trend and $(c_t - \mu_c t)$; μ_d in this case is the coefficient on the time trend. Bansal, Dittmar, and Lundblad (2001) also consider this specification for measuring the long run dividend beta.

Our cointegration specification includes a time trend and does not restrict the cointegration coefficient, δ_i across assets. The validity of this specification in examining the cointegration of dividends and consumption is the subject of some debate. In particular, Hansen, Heaton, and Li (2004) consider specifications with and without a time trend and with and without restricting the cointegration parameter to unity, and Menzly, Santos, and Veronesi (2004) consider solely a specification with no time trend and unit cointegration parameter. In our empirical work, we also entertain these restrictions, and discuss their plausibility below.

Including the time trend in the estimating the cointegration parameter (long run dividend beta) and allowing for heterogeneity in the cointegration parameter are important for the following reasons. First, without the time trend in equation (), the cointegration parameter simply equals the ratio of average dividend growth to average aggregate consumption growth (see Hamilton (1994)). Cross-sectional differences in this long-run dividend beta (the cointegration parameter) will tautologically reflect differences in average dividend growth (and average capital gains), and therefore, average ex-post returns. Including the time trend purges the effects of mean growth dividend rates on the cointegration parameter, and ensures that the long-run dividend beta does not tautologically reflect cross-sectional differences in ex-post average returns. Hence, for the cross-sectional risk-return relation, it is important to include a time trend in measuring the long-run dividend beta.

Second, it may be economically appealing to restrict the cointegration parameter of

aggregate dividends on the market or a sector of the economy to one, as this implies that the average growth rate in these dividends will match that of aggregate consumption. However, there is no economic rationale for this restriction on the cointegration parameter for *dividends per share*, which we and other empirical studies employ. In order to reinforce this important point, we plot the log ratio of aggregate dividends to consumption in Figure 1(a) and the log ratio of aggregate dividends per share to consumption in Figure 1(b). Aggregate dividends at date t for the market are computed by multiplying the dividend yield at date t with the lagged market capitalization (K),

$$D_t^{agg} = Y_t K_{t-1}. \quad (24)$$

As shown in Figure 1(a), the ratio of aggregate dividends to consumption is quite stationary. Further, the average growth rate of aggregate dividends is 3.2%, which is comparable to that of aggregate consumption. Thus, restricting this series to be cointegrated with consumption with a parameter of one and omitting time trends appears to be quite reasonable.

In contrast, the ratio of market per share dividends, typically used in asset pricing, to consumption displays a dramatic decline over time, as shown in Figure 1(b). The average growth rate in the dividend per share series is 0.9%, which is considerably lower than that in consumption. This difference arises as the growth in the aggregate dividend series reflects the growth in market capital (capitalization), while the growth in dividends per-share reflects the growth in price per share. The average growth of market capital (capitalization) is 4.3%, while that in price per share is only about 2.0%. As discussed above, the average growth in price per share could simply be due to differences in the riskiness of different assets.⁶

We examine the implications of restricting the time trend and the cointegration parameter in more detail in our cross-sectional investigations below. At present, we note simply that omitting the time trend will result in a nearly tautological relation between dividend betas and consumption growth, and the unit cointegration parameter

⁶If we assume a uniform relation between dividends and consumption in the long run, and restrict the cointegration parameter to be one for all the assets, the sample ACF's (not reported) detect a much higher degree of persistence in the $d_{i,t} - c_t$ series for most of the portfolios. That is, the logarithms of the dividends to consumption ratios are much farther from being stationary than are the OLS residuals. This evidence is in favor of heterogeneity in the cointegration relations between dividends and consumption in the cross-section of assets.

will restrict the long-run dividend beta for all assets to be one. In sum, we conclude that omitting a time trend and restricting the cointegration parameters to unity across assets is not reasonable when using dividends per share.

3.3 Deriving The Term Structure of Betas

This section provides the details for estimating assets' consumption betas for different investment horizons. We first estimate the cointegrating relation (3.2) between dividends and consumption via OLS, that is by regressing the de-trended portfolio's cash flows onto the stochastic trend in consumption. Using the resulting cointegrating residual, $\hat{\epsilon}_{d,t}$, we model its dynamics jointly with the portfolio's price-dividend ratio, z_t , and consumption growth, Δc_t , in the first-order error-correction VAR (EC-VAR) structure:

$$X_t = AX_{t-1} + Gu_t, \quad (25)$$

where $X_t' = (\Delta c_t \ \hat{\epsilon}_{d,t} \ z_t \ \Delta d_t \ \Delta z_t)$, A is a 5×5 matrix of coefficients, G is a 5×3 matrix, and u is a three by one matrix of shocks, $u_t' = (\eta_t \ \eta_{\epsilon,t} \ \eta_{z,t})$. Details of the VAR specification are provided in the appendix; all growth rates are de-meanned throughout our discussion.

Using the recursive structure of the VAR, we can compute the long-run variance of consumption growth and the various pieces of long-run risks in equation (12). Specifically, let $B_j = B_{j-1} + A^{j-1}$ and $B_0 = 0$. The horizon- s covariance matrix of the above variables satisfies the recursion

$$\Sigma_s^* = B_s \Sigma_e B_s' + \Sigma_{s-1}^*, \quad (26)$$

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

$$\Sigma_{lr} = [I - A]^{-1} \Sigma_e [I - A]^{-1'}. \quad (27)$$

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 in the matrix. For a given

horizon s , the covariance risk in the asset is

$$\beta_s = [\Sigma_s(1, 4) + \Sigma_s(1, 5) + (\kappa_1 - 1)\Sigma_s(1, 3)]\Sigma_s(1, 1)^{-1}, \quad (28)$$

where $\Sigma_s(i, j)$ is the (i, j) -element of the covariance matrix Σ_s . At short horizons, the latter two terms (i.e., price and valuation-ratio risks) are significant determinants of an asset's risk; at long horizons, the first term (dividend risk) dominates. Alternatively stated, price risk has the dominant influence on the beta at short horizons, the dividend risk at long horizons. Notice that, in the limit, the first term equals δ , that is $\Sigma_{lr}(1, 4)\Sigma_{lr}(1, 1)^{-1} = \delta$, and the long-run beta is

$$\begin{aligned} \beta_{lr} &= [\Sigma_{lr}(1, 4) + 0 + (\kappa_1 - 1)\Sigma_{lr}(1, 3)]\Sigma_{lr}(1, 1)^{-1} \\ &= \delta + (\kappa_1 - 1)\Sigma_{lr}(1, 3)\Sigma_{lr}(1, 1)^{-1}. \end{aligned} \quad (29)$$

Consequently, at long horizons, the beta will be determined almost entirely by the cointegration parameter, as the second term in the beta expression converges to zero, and the third term, with κ_1 close to one, has a minimal effect.

There are several implications of the expression for horizon-dependent betas, (28), that are worthy of note. First, dispersion in long-run risks will significantly impact the consumption beta even at short horizons, despite the greater role for price-dividend risks. To illustrate this point, consider the beta for one-period investment. Given (28),

$$\beta_1 = \delta + [\Sigma_1(1, 2) + \Sigma_1(1, 5) + (\kappa_1 - 1)\Sigma_1(1, 3)]\Sigma_1(1, 1)^{-1}. \quad (30)$$

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). As the long-run beta is approximately equal to δ , equations (28) and (29) show that $\beta_s - \delta$ represents the effects of transitory consumption risks on the betas.

Second, expression (28) highlights important effects in the beta at all horizons that will not arise if the cointegration parameter is restricted or cointegration is not imposed. Equation (28) characterizes the entire term structure of the betas; when δ is assumed

to be one, as in Menzly, Santos, and Veronesi (2004), the implications for risk and return are altered at all horizons. Dispersion in both short-run and long-run betas will be impacted by this restriction. Further, when cointegration is not imposed, then $\hat{\epsilon}_{d,t}$ is not a predictive variable for returns and dividend growth rates, and the permanent term δ is absent from the term structure of betas. This predictive power is an empirical consideration which we investigate in the next section. The expression shows that the cointegration of dividends and consumption, and heterogeneity in the long-run dividend beta are potentially important determinants of the risk-return relation at all horizons.

Although cointegration yields several potentially important insights into the risk-return relation, it is important to note the distinction between long-run risks and cointegration. The beta expression, (28) holds at horizon s for any specification of the VAR. Even if consumption and dividends are not cointegrated, a standard VAR representation will admit a horizon-dependent beta, β_s . Consequently, the long-run covariance between an asset's cash flows and consumption can still be the dominant influence on the long-run beta in the absence of cointegration.

We use the term structure of betas implied by expression (28) in investigating the cross-sectional relation in risk and return at various horizons. Specifically, we investigate the relation by considering cross-sectional regressions for different horizons s ,

$$E[r_i + 0.5\sigma_{r,s}^2] = \lambda_{0,s} + \lambda_{1,s}\beta_{i,s}. \quad (31)$$

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 4.

4 Empirical Results

In this section, we investigate the implications of the preceding framework for the measurement of assets' risks. We first perform a variance decomposition for returns, and discuss the evidence on the cointegration of dividends and consumption for the cross-section of portfolios used in our analysis. We then report consumption betas and expected returns across different investment horizons, implied by our cointegrated VAR framework, and analyze the cross-sectional implications of the model. Finally, using the

cross-sectional estimates we quantify the implied compensation for business cycle risks and compare its magnitude with that for permanent consumption risks.

4.1 Sources of Risk in the Short and Long Run

As we have discussed in the preceding sections, the consumption risk in equities can differ in the short and long run. In this section, we examine the potential sources of these risks in returns. As shown in the return decomposition, equation (2), variation in returns arises from variation in dividend growth rates and price dividend ratios. We analyze which of these sources dominate return variation in the long and the short run.

We perform a variance decomposition for returns using a VAR model which models dividend growth rates and log price to dividend ratios. Dividend growth is projected on its own lag, whereas the price-dividend ratio depends on one lag of the dividend growth, as well as its own lag. That is,

$$\begin{pmatrix} \Delta d_t \\ z_t \end{pmatrix} = \begin{bmatrix} a_{11} & 0 \\ a_{12} & a_{22} \end{bmatrix} \begin{pmatrix} \Delta d_{t-1} \\ z_{t-1} \end{pmatrix} + \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}. \quad (32)$$

The innovation in the price-dividend ratio, $u_{2,t}$, could be correlated with that in the dividend growth rates, $u_{1,t}$. To provide a clean interpretation of the role of price shocks versus dividend shocks we regress the price innovation onto the dividend innovation; the residual from this regression, therefore, is the price innovation orthogonal to the cash-flow shock. That is, we are assuming that dividend innovations lead to price movements, but price innovations do not lead to (contemporaneous) responses in dividends.

Table II shows the percentage of the return variance that is attributed to dividend shocks. As shown in the table, the contribution of dividends and price-dividend ratios to return variation varies considerably by horizon. In the short run, price risks dominate. On average, about 10% of the return variation over the over the one-year horizon is attributable to dividend variation, indicating that 90% of return variation is attributed to price risks. This pattern reverses as the horizon increases. By the five-year horizon, on average 32% of return variation is due to dividend shocks, and by the ten-year horizon, 64% of the variation is due to dividend shocks. In the long-run limit, virtually all of the variation in returns is attributed to dividend shocks.

These patterns suggest that short-run investors will concern themselves with price shocks and long-run investors will concern themselves with dividend shocks. However, even at the short horizon, dividend shocks play an important role. In particular, at the 1-year horizon, the variation in relatively high book-to-market portfolio returns (B6, B7, and B9) is composed of over 20% dividend shocks. By the five-year horizon, a clearer pattern emerges; small stocks and high book-to-market stocks' return variation is dominated by dividend shocks, whereas large and low book-to-market stocks' return variation remains largely attributable to price shocks. Thus, these patterns indicate that dividend risks are potentially important even at short horizons, and that there is, as in mean returns, cross-sectional variation in the contribution of dividend shocks to return variation.

The results in this section are suggestive of our claim that dividend risks play the dominant role in capturing long-run risk in asset returns. We build on this evidence by examining the long-run consumption risks in asset returns. In the following section, we discuss our particular approach to measuring these risks, namely cointegration in dividends and consumption.

4.2 Cointegration Evidence

In Table III, we present point estimates of the cointegration parameters between portfolios' cash flows and consumption, and report the sample autocorrelation functions (ACF) of the cointegrating residuals. As discussed earlier, we estimate cointegration parameters via OLS by regressing the deterministically de-trended dividends on de-trended consumption. We first note that, for the majority of portfolios analyzed, the sample autocorrelations of the resulting cointegration 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, 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.

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 pat-

tern across the market capitalization deciles. The parameter 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. The fact that high book-to-market stocks have large exposures to permanent risks in consumption suggests that the performance of these firms is linked to the permanent risks in the economy, while that of low book-to-market portfolio firms is not. As consumption is largely dominated by permanent shocks, the risks in low book-to-market stocks are largely unrelated to the long-run evolution of the economy. This, as we document further below, is exactly why the low book-to-market portfolio should bear a low ex-ante risk premium and already highlights the importance of long run risks.

As we discuss in Section 3.2, cointegration of dividends and consumption has important implications for several aspects of asset pricing. One of these implications is for the dynamics of returns and their components, price-dividend ratios and dividends. In particular, under the cointegration specification, the cointegration residual is an important predictor of future dividend growth. As such, this residual is a potentially important piece of our characterization of horizon-dependent mean returns. We examine the predictive power of this residual and its implications in the next section.

4.3 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 dynamics of dividend growth rates while dividends are moving back towards 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 trend in consumption. Given the approximation for the log return in equation (2), the predictability of the dividend growth rates directly translates into return predictability. The variation in the cointegrating residuals, therefore, should 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 importance of the cointegration relation, we compare the adjusted- R^2 's for dividend growth and return projections implied by the EC-VAR model outlined above with the corresponding \bar{R}^2 's from the growth-rates VAR specification. The latter simply excludes the cointegrating residuals from a set of VAR variables. We find that that the cointegration evidence does sharply alter the predictability of dividend growth rates. The average predictability of cash-flow growth rates in the cross-section of assets changes from 16% at the first horizon to about 30% at ten-year horizon. The \bar{R}^2 for dividend growth projections deteriorates when the cointegrating residuals are excluded from the VAR. ⁷

Further, asset return predictability is also altered by the cointegration between dividends and consumption. Return projections' \bar{R}^2 's for horizons one, five and ten years implied by the EC-VAR model as well as the alternative, growth-rates, specification are reported in Table IV. The EC-VAR specification, on average, is able to explain over 15% of returns variation at one year horizon, and around 50% of the variation in ten-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 2 by plotting one 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 rates model. Return innovations, therefore, also differ across the two specifications, and most importantly so do the consumption betas measured from the two alternate models.

The results of this section further underscore the importance of the cointegration specification in the measurement of risk and return. The specification has important implications for returns in that the cointegration parameter governs the long-run relation between risk and return, and is an important determinant of the short-run relation as well. Further, as emphasized in this section, temporary deviations of cash flows from the permanent component of consumption contain important information for predicting

⁷For brevity we have excluded the detailed tables with this evidence. This evidence is available from the authors upon request.

dividend growth rates and returns, and thus represent an important component in the calculation of the term structure of expected returns and betas.⁸ We turn to this term structure in the next section, and analyze how the risk-return relation changes with the investment horizon.

4.4 The Term Structure of Betas and Expected Returns

In this section, we explore the implications of the VAR structure for expected returns and the consumption risk inherent in the portfolio returns. In Table V, we present betas and expected returns at various horizons for each of the portfolios. At the short horizon ($s = 1$), the measured betas exhibit a pattern that generally matches that of risk premia across the characteristic deciles. The small firm beta (4.13) exceeds the large firm beta (1.54) and the high book-to-market beta (3.89) exceeds the low book-to-market beta (1.81), matching similar patterns in expected returns. The dispersion in these betas is reasonably large; the cross-sectional standard deviation of the one-period betas is 1.42, compared to a cross-sectional standard deviation of 1.92% in the average returns.

As the horizon increases, neither dispersion in the betas nor the expected returns decreases. As shown in the second and third column of Table V, the cross-sectional standard deviation of returns at the five-year horizon is 1.72% and at the ten-year horizon is 1.99%. As we have discussed previously, in order to match this pattern, there must be important sources of variation in either the dividend yield or cash flow component of returns at longer horizons. If not, the cross-sectional dispersion in estimated betas will shrink. Further, as most of the variability in returns at long horizons derives from cash flows, variability in cash flow risk must be high at long horizons in order to match return dispersion. The evidence in Table V indicates that cross-sectional dispersion in betas does indeed increase. At the five-year horizon, the cross-sectional standard deviation is 2.04, and at the ten-year horizon it is 2.81. We illustrate this pattern in Figure 3, plotting the time-horizon profiles of cross-sectional dispersion in consumption betas and mean returns. As is shown, the cross-sectional dispersion in both consumption and mean returns increases with horizon.

These results suggest the existence of considerable heterogeneity in dividend betas

⁸Our evidence on predictability of cash-flows is consistent with that in Lettau and Ludvigson (2005), Ang and Bekaert (2003), and Bansal, Khatchatrian, and Yaron (2004).

at long horizons. Since the long-run dividend beta is dominated by the parameter of cointegration with consumption, the evidence supports estimating dividend cointegration parameters without restrictions. If we suppress this heterogeneity by restricting the cointegration parameter to be unity for all assets, dispersion in both consumption betas and mean returns shrinks dramatically. This result is shown as well in Figure 3; dispersion in mean returns when we restrict $\delta_i = 1$ (for all i) is virtually flat across all investment horizons. Thus, differences in permanent consumption risks in dividends account for cross-sectional dispersion in consumption betas and risk premia at both short and long investment horizons.

The importance of dividend betas is further emphasized in Table VI. In this Table, we use the decomposition, (12), to extract the component of the overall consumption beta attributable to risks in dividends. In the Table, we show betas at the 1, 5, and 10-year horizons, as well as the cointegration parameter, reproduced from Table III for convenience. Comparing the one-period betas in Table V and Table III shows that even the one-period beta is dominated by the dividend portion of the overall consumption beta. By the 10-year horizon, the dividend beta has nearly converged to the cointegration parameter. Much of the cross-sectional relation between returns and risk appears to be driven by this component of the beta. These results point to the importance of the long-run risk in dividends and consumption even at shorter horizons. We explore this issue in more detail below.

In addition to implications for risk measures and mean returns, the assumption of cross-sectional homogeneity in cointegration relations has non-trivial implications for assets' Sharpe ratios. In Figure 4, we plot the Sharpe ratios of the market and extreme size and book-to-market portfolios by horizon. This Sharpe ratio is computed as a ratio of an asset's expected return, scaled by horizon, to the per-unit-time volatility of return.⁹ The Sharpe ratios implied by the unrestricted EC-VAR are plotted in Figure 4(a). A number of interesting implications emerge from the figure. First, small and value firms offer a better Sharpe ratio deal to short-term investors than do growth and large stocks. However, as the investment horizon increases, the volatility of large and growth stocks declines, whereas the volatility of small and value stocks increases. These patterns occur

⁹Note that we use mean returns instead of expected excess returns as is required by the definition of the Sharpe ratio. However, as the average slope of the real yield curve in practice is only slightly positive, the ratios that we present should not be critically affected by the term structure of interest rates.

because the long-run performance of small and value stocks are tightly related to low-frequency variation in aggregate consumption, while large and growth stocks' returns are not. Thus, at long horizons, large and growth stocks offer greater mean returns per unit of risk than do their small and value counterparts. Finally, it is also interesting to note that the market portfolio delivers the highest Sharpe ratio at long horizons, reaching a maximum of approximately 0.80 at the 35-year horizon and declining to 0.50 in the long run.

In contrast, restricting the long-run dynamics of assets' cash flows significantly changes the risk-return relation across horizons. To illustrate this point, in Figure 4(b) we plot the Sharpe ratios calculated under the assumption that all assets have a unit cointegration parameter with consumption. Since all assets' dividends are uniformly tied to long-run consumption variation, cross-sectional variation in risk fades as the horizon increases. Thus, assets that provide a better Sharpe ratio at a one-year horizon also exhibit higher Sharpe ratios at long horizons; value and small stocks consistently offer higher mean returns per unit of risk than growth and large firms. Additionally, for some portfolios, this restriction implies a very high magnitude of the Sharpe ratio. In particular, small-capitalization stocks earn a Sharpe ratio of 1.0 at the 50-year horizon and 1.5 in the long run. Although there is no established theoretical limit, several authors have suggested the existence of "good deal" bounds on Sharpe ratios (e.g. MacKinlay (1995) and Cochrane and Saa-Requejo (2000)). The high magnitudes of Sharpe ratios suggested by these results cast reasonable doubt on the reliability of these estimates and, therefore, the validity of imposing the unit cointegration restriction.

In the next section, we more formally analyze the relation between mean returns and risk measures in the cross-section. Additionally, we examine the implications of restrictions on the cointegration parameter and time trend for capturing cross-sectional relations between mean returns and risk measures. The results provide additional insight into the challenge of measuring an asset's long run risk.

4.5 Cross-Sectional Risk and Return

A wide variety of consumption risk measures have been proposed in the asset pricing literature that are able to some extent to account for the observed heterogeneity in risk premia. Some papers rely on constant beta models (for example, Parker and Jul-

liard (2005) consider what they call the ultimate consumption risk beta), others try to exploit the cross-sectional implications of conditional versions of the CCAPM (as in Jagannathan and Wang (1996), Lettau and Ludvigson (2001)). The aforementioned papers focus merely on short-horizon, one-period returns, as is typical in empirical studies. However, as we emphasize, the risk and return relation varies considerably across investment horizons. Thus, it is important to explore the cross-sectional implications of an asset pricing model not only in the short but also in the long run. We therefore investigate the cross-sectional risk-return relations at various investment horizons, and explore the ability of our consumption betas to explain the cross-sectional dispersion in risk premia at both short, as well as long ends of the investment horizon.

We estimate the cross-sectional regression in (31) for the portfolio menu at horizons of one, five, and ten years, as well as the long-run limit. These results are presented in Tables VII and VIII. 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.¹⁰ We present three specifications in the tables: the case in which a time trend is included in the cointegration specification and the cointegration parameters are not restricted; the case in which the time trend is suppressed and the cointegration parameters are not restricted; and the case in which we include a time trend but restrict the cointegration parameters to unity for all assets.

We benchmark our results by presenting regressions of the average return on the cointegration parameter in Panel A. In the limit, for κ_1 close to one, the consumption beta converges to this parameter. As shown in the table, this approach fits average returns quite well. The price of risk of 0.486 (S.E. = 0.053) is positive and statistically significant, and the cointegration parameter explains 81% of cross-sectional variation in average returns. This approximation to the long-run risk-return relation suggests that the long-run risks embodied in assets' cash flows explain a significant portion of the cross-sectional variation in risk premia.

In Panel B, we present results in which a time trend is included in the cointegration specification and the cointegration parameter is unrestricted. The results indicate that

¹⁰To resolve the estimation issues arising from a relatively small number of time-series observations, we use a block-diagonal structure of the GMM weighting matrix. The number of lags used in Newey-West var-cov estimator is set to 8.

the relation between average return and beta under this specification is both statistically and economically significant across various horizons. The estimates of the price of risk vary from 1.187 (S.E. 0.334) for one-year horizon to 0.650 (S.E. 0.236) for the investment horizon of ten years, converging to 0.718 (S.E. 0.225) in the long-run limit. The explanatory power of these betas for average returns ranges from an adjusted- R^2 of 0.62 in the two-period case to 0.87 in the long run. Note that, as the explanatory power (\bar{R}^2) of the cointegration coefficient for average returns is 0.81, that the majority of the betas' explanatory power, particularly in the long run, is due to long-run risks present in the cointegration of dividends and consumption. The fit of the one-period and long-run specification are plotted in Figure 5.

In Panel B, we also present results for cross-sectional regressions by horizon using *only* the dividend portion of the consumption beta, implied by expression (12). As shown in the table, most of the explanatory power of the consumption beta can be traced to the dividend component of the beta. At shorter horizons, $s = 1, 2$, it is apparent that the remaining components of the betas play some role in describing cross-sectional variation in mean returns; compare the \bar{R}^2 of 0.75 and 0.62 at the 1- and 2-year horizon, respectively, for the overall beta to 0.66 and 0.76 for the dividend beta. However, by the five-year horizon, the regression \bar{R}^2 suggests that virtually all of the cross-sectional variation explained by the beta is related to the dividend portion of the beta. Insofar as this beta is dominated by measures of long-run risk, this result highlights again the importance of long-run risk for explaining not only long-horizon, but also short-horizon mean return variation.

We present results for the restricted specifications in Table VIII. The first set of columns present results restricting the cointegration parameter to unity across all assets. The results at the one-period horizon are somewhat weaker, but similar to those for the unrestricted case; the price of risk is 1.583 (S.E. = 0.393) 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 accounting 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 the short-term. This specification does poorly at both long and short horizons in accounting for risk premia across assets.

In the second set of columns, we omit the time trend, but allow the cointegration parameter to vary across assets. These results represent an improvement over the case in which the cointegration parameter is restricted. Across all horizons, the estimated price of risk is positive and statistically significant, and the explanatory power of the regression varies from an \bar{R}^2 of 0.35 for the two-year horizon to 64% for the ten-year horizon. However, as mentioned previously, these results may occur because, in this case, the cointegration parameter tautologically reflects the mean growth rate in the portfolio dividends and hence average ex-post retruns. Thus, to avoid this potential complication, we stress the results obtained from including a time trend in the specification.

Our results are robust to different sample periods; over the post-war sample (1954-2002), we observe that consumption betas account for more than 60% of the cross-sectional variation in risk premia at short and long horizons. Further, the estimated prices of risk exhibit a similar decline across horizons. In summary, our results indicate that our specification matches the pattern in expected returns quite well at all horizons. Further, the results demonstrate that the long-run risk in returns, embodied in the cointegration parameter, is critical for understanding the relation between assets' risks and their average returns, regardless of the horizon.

As shown in the results, the price of consumption risk is declining with the investment horizon. We discuss in Section 2 the implications of the price of risk across horizons for understanding the cost of business cycles. In the following section, we briefly examine the implications of the results for the asset-market costs of these cycles.

4.6 Business Cycle Risks Compensation

Our last empirical consideration is the implied compensation for business cycle fluctuations in consumption. We construct the time-horizon profile of business cycle risk compensation as described in equation (10) using the cross-sectional estimates of the market price of risk at various investment horizons. Figure 6 plots the compensation for short-run consumption risks along with the total market price of risk. The compensation for business cycle fluctuations 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. The risk compensation in the long run, as shown in Table VII, is about 0.718 per annum and at very short horizon is 1.187 percent. The small and sharply declining business

cycle risk compensation again emphasizes the investor's concerns about low-frequency consumption risks in assets' payoffs. The bulk of the risks in consumption are due to concerns about permanent risks and not due to transitory risks in consumption. Note that evidence about the small size of the business cycle risks continues to be true even when alternative methods for measuring long run dividends betas are used, the compensation for business cycle risks (see Table VIII) in the no time trend specification is also similar to our leading specification.

These measures of the market compensation for business cycle and permanent risks relate to the voluminous literature on measuring the cost of business cycles initiated by Lucas (1987). Our unique evidence shows that the market offers much higher risk compensation for long-run risks relative to transitory (business cycle) risks in consumption. This is consistent with the arguments presented in Lucas (1987) and particularly Alvarez and Jermann (2004a) that cost of business cycle risks seem small relative to those associated with permanent risks.

In summary, the results of this section point strongly to the case for long-run cash-flow risks. As discussed above, the cross-sectional dispersion in average returns does not decrease with horizon. However, the variation in price-dividend ratios is small, and the variation in changes to these ratios must converge to zero in the long run. Consequently, the only avenue via which we can explain the variation in average returns at long horizons is through long-run cash-flow risk. In this section, we document that this cash-flow risk is indeed the dominant force in describing long-run returns; moreover, it has a strong explanatory power for variation in average returns at short horizons as well.

5 Monte Carlo Analysis

As in Bansal, Dittmar, and Lundblad (2001), Hansen, Heaton, and Li (2004) also measure the long run dividend beta's, they argue that that these risk measures are imprecisely estimated. Using a series of Monte Carlo experiments we evaluate the impact of this imprecise measurement on the cross-sectional evidence—which is the key focus of the paper. The main message of these Monte Carlo's is that indeed all beta's are imprecisely measured in the time series, however the cross-sectional evidence is very robust to this measurement error. For brevity we do not include the detailed tables (they are available

on request).

In the first Monte Carlo (MC-1) we carry out all the estimation steps in arriving at our empirical results. This allows us to analyze the empirical distributions of parameter estimates, t-statistics and adjusted- R^2 's under the null hypothesis that the EC-VAR model is correct.¹¹ One-step ahead mean returns are assumed, as per the model, to be determined by the asset beta's and the market prices of risk,

$$\mu_i = \lambda_{0,1} + \lambda_{1,1}\beta_{i,1}. \quad (33)$$

Our results, as in Hansen, Heaton, and Li (2004) show that the standard errors on the long run dividend betas are large. However, we also find that the mean estimate of this long-run dividend betas, as well as all the consumption betas, is very close to its population value; there is no bias in the estimate. This is important, as it shows that despite the time series uncertainty these betas measure the population values reasonably well. Despite large confidence intervals of the time-series estimates, the cross-sectional regressions produce significantly positive market prices of risk at all investment horizons. While the betas are the average returns are imprecisely estimated is the time series—the cross-sectional correlation between them continues to is very high. Hence, the cross-sectional evidence seems to be robust.

We also find that the λ_1 and the cross-sectional \bar{R}^2 are biased downward — the distribution of the risk price at one-year horizon is centered around 0.65, whereas the data have been simulated using the point estimate of 1.187; the median \bar{R}^2 also falls far below its theoretical value of 1. This suggests that our data estimates correspond to very conservative, low bounds of the true values of risk prices and \bar{R}^2 's. The bias in the risk price estimates disappears in the long run — that is, as the horizon increases, the empirical distributions of the cross-sectional estimates get centered around their population values. Our estimate of the market price of the permanent risks, $\hat{\lambda}_{lr}$, therefore, reveals the true compensation for the permanent risks in consumption. In contrast, the downward bias in the cross-sectional estimates at short horizons makes it much more difficult to uncover the true magnitude of the business cycle risk compensation.

¹¹Specifically, we simulate 5,000 samples of 74 annual time-series observations of consumption growth (and levels), dividend growth (and levels) and portfolios' attributes (i.e. pd-ratios, cointegrating residuals, and returns) from the cointegrated VAR system using point estimates of matrices A and G , and the covariance matrix Σ_u

Our second monte carlo (MC-2) experiment tried to confirm the importance of cash-flow risks in justifying the cross-sectional variation in risk premia. In this experiment there is no cross-sectional heterogeneity in the long-run dividend betas, all of these are assumed to be one in the null model. From an econometrician’s perspective we estimate in our monte carlos, a specification which does not impose that the cointegration parameter is one. We find that the return beta’s fail to explain the cross-sectional differences in mean returns in finite samples—the cross-sectional R^2 ’s are very small and the cross-sectional slope coefficients mostly insignificant. This evidence emphasizes the importance of long-run dividend risk heterogeneity in the population. If the population values of long run dividends betas are identical across assets, then consumption betas, at least using the method to we measure them in the data, will fail to account for the risk-return relation at various investment horizons. In the data, we account for the risk-return relation at both short and long horizons — this monte carlo exercise, therefore, suggests that our empirical evidence is unlikely to come from a world where long-run cross-sectional dividend heterogeneity is absent.

To summarize, the Monte Carlo results corroborate the empirical evidence presented in the paper. In particular, they highlight the importance of permanent consumption risks embodied in assets’ dividends to account for the risk premium heterogeneity observed in the data.

6 Conclusion

In this paper we show that the exposure of assets’ dividends to permanent risks in consumption is a key determinant of the risk compensation in assets markets. We consider the evolution of the risk-return relation and its determinants across different investment horizons. We show that, while at short horizons, the variation of an asset’s return is dominated by transitory price shocks, as the time horizon increases, it is dividend shocks that are the major source of the return volatility. The sources of return volatility, and, therefore, the nature of concerns of an investor change significantly by horizon.

We show that the long-run beta of the asset is largely equal to the long-run dividend beta — that is, the cointegration parameter between dividends and consumption. The exposure of assets’ cash flows to a low-frequency component in aggregate consumption

(i.e., the cointegration parameter between dividends and consumption), therefore, is a key concern of a long-run investor. This relation determines not only the long-run risk-return relation, but also the short-run relation between systematic risks and expected returns. Using the error-correction VAR framework, we construct the term-structures of consumption risk measures (betas) and mean returns, and use them to measure the market compensation for consumption risks across different investment horizons. We find that the market compensation for permanent consumption risks is much higher than that for the transitory (business cycle) fluctuations in consumption. That is, long-run risks seem to be the key concern in asset markets.

Appendix

In this appendix, we provide the details of the 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 (3.2) between dividends and consumption, we model the dynamics of the resulting cointegrating residual, $\hat{\epsilon}_{d,t}$, jointly with the portfolio's price-dividend ratio, z_t , and consumption growth, Δc_t , by the following VAR structure:

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

The first three variables form the basis of the VAR process. The last two variables provide no additional information; they are included into the 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 $\Delta d_t = \Delta \hat{\epsilon}_{d,t} + \hat{\delta}\Delta c_t$, and $\Delta z_t = z_t - z_{t-1}$.

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

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

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

$$X_t - AX_{t-1} = Gu_t \equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \hat{\delta} & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \eta_t \\ \eta_{\epsilon,t} \\ \eta_{z,t} \end{pmatrix}. \quad (\text{A-2})$$

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

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

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

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

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

$$\Sigma_s^* = B_s \Sigma_e B_s' + \Sigma_{s-1}^*, \quad (\text{A-3})$$

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

$$\Sigma_s = \frac{1}{s} B_s \Sigma_e B_s' + \left(1 - \frac{1}{s}\right) \Sigma_{s-1}, \quad (\text{A-4})$$

Equation (A-4) 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.,

$$\Sigma_{tr} = [I - A]^{-1} \Sigma_e [I - A]^{-1'}. \quad (\text{A-5})$$

References

- Alvarez, Fernando, and Urban J. Jermann, 2004a, Using asset prices to measure the cost of business cycles, *Journal of Political Economy* 112(6), 1223–1256.
- Alvarez, Fernando, and Urban J. Jermann, 2004b, Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth, Working paper, University of Chicago and the Wharton School of the University of Pennsylvania.
- Ang, Andrew, and Geert Bekaert, 2003, Stock return predictability: Is it there?, Working paper, Columbia University.
- Bansal, Ravi, Robert F. Dittmar, and Christian Lundblad, 2001, Consumption, dividends, and the cross-section of equity returns, Working paper, Duke University.
- Bansal, Ravi, Robert F. Dittmar, and Christian Lundblad, 2005, Consumption, dividends, and the cross-section of equity returns, *Journal of Finance*, *forthcoming*.
- Bansal, Ravi, Varoujan Khatchatrian, and Amir Yaron, 2004, Interpretable asset markets?, *European Economic Review*, *forthcoming*.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- 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, 151–174.
- Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.
- Campbell, John, and Robert Shiller, 1988, Stock Prices, Earnings, and Expected Dividends, *Journal of Finance* 43, 661–676.
- Campbell, John Y., and John H. Cochrane, 1999, By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior, *Journal of Political Economy* 107, 205–255.

- Campbell, John Y, and Pierre Perron, 1991, Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots, *NBER Macroeconomics Annual* pp. 141–201.
- Cochrane, John, and Jesus Saa-Requejo, 2000, Beyond Arbitrage: Good Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy* 108, 79–119.
- Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the intertemporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Hamilton, James, 1994, *Time Series Analysis*. (Princeton University Press, Princeton, NJ).
- Hansen, Lars, John Heaton, and Nan Li, 2004, Consumption strikes back?, Working paper, University of Chicago.
- Hansen, Lars Peter, and Kenneth Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 63, 767–804.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–54.
- Lettau, Martin, and Sydney Ludvigson, 2001, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.
- Lettau, Martin, and Sydney Ludvigson, 2005, Expected Returns and Expected Dividend Growth, *Journal of Financial Economics*, forthcoming.
- Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1446.
- Lucas, Robert E. Jr., 1987, *Models of Business Cycles*. (Basil Blackwell New York).
- MacKinlay, A. Craig, 1995, Multifactor Models do not Explain Deviations from the CAPM, *Journal of Financial Economics* 38, 3–28.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112(1), 1–47.

Parker, Jonathan A., and Christian Julliard, 2005, Consumption Risk and the Cross-Section of Expected Returns, *Journal of Political Economy* 113(1), 185–222.

Table I
Summary Statistics

Portfolio	Cash Flow Growth		Returns	
	Mean	Std. Dev.	Mean	Std. Dev.
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

Table I 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 one 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.

Table II
Variance Decomposition for Returns

Portfolio	H o r i z o n			
	1	5	10	LR
S1	0.06	0.56	0.84	1.00
S2	0.00	0.11	0.50	0.98
S3	0.00	0.30	0.68	0.99
S4	0.11	0.25	0.70	0.99
S5	0.06	0.14	0.58	0.98
S6	0.15	0.22	0.65	0.98
S7	0.04	0.09	0.46	0.96
S8	0.03	0.09	0.41	0.93
S9	0.02	0.11	0.49	0.95
S10	0.04	0.07	0.29	0.87
B1	0.03	0.14	0.37	0.94
B2	0.00	0.13	0.45	0.95
B3	0.06	0.08	0.38	0.93
B4	0.06	0.52	0.81	0.99
B5	0.11	0.30	0.68	0.96
B6	0.23	0.58	0.86	0.99
B7	0.24	0.65	0.92	0.99
B8	0.10	0.54	0.84	0.99
B9	0.26	0.87	0.96	1.00
B10	0.02	0.62	0.89	1.00

Table II reports the fraction of the return variance that is accounted for by dividend growth shocks. The percentages are reported for one-, five-, and ten-year horizons, as well as for the long-run limit, for the 20 portfolios sorted by market capitalization (S1-S10) and book-to-market ratio (B1-B10). Variance decomposition is performed by fitting a VAR(1) model to portfolio's cash-flow growth rates and price-to-dividend ratio. The coefficient matrix in the VAR model is assumed to be lower triangular. Variance decomposition is calculated by orthogonalizing dividend and pd-ratio shocks. The percentage of the return variance due to price shocks, therefore, is equal one minus the reported entries.

Table III
Cointegration Parameters and ACF

Portfolio	$\hat{\delta}$	ACF(1)	ACF(5)	ACF(10)
S1	9.62 (6.07)	0.90	0.34	-0.02
S2	9.71 (5.06)	0.64	0.41	-0.01
S3	6.69 (2.86)	0.75	0.41	0.17
S4	6.21 (2.37)	0.81	0.30	0.11
S5	4.41 (1.70)	0.70	0.17	0.05
S6	4.38 (1.88)	0.84	0.28	0.20
S7	2.42 (1.12)	0.72	0.24	0.14
S8	2.38 (0.77)	0.65	-0.03	-0.16
S9	2.38 (0.96)	0.81	0.22	0.03
S10	0.82 (0.89)	0.79	0.24	-0.04
B1	-0.27 (1.28)	0.73	0.16	-0.13
B2	-2.59 (1.41)	0.72	-0.05	0.17
B3	-0.11 (0.79)	0.61	0.07	-0.02
B4	0.82 (1.91)	0.64	0.17	0.13
B5	2.79 (1.59)	0.91	0.61	0.17
B6	4.83 (1.75)	0.68	-0.01	0.05
B7	6.36 (2.20)	0.76	0.20	0.07
B8	9.70 (3.07)	0.79	0.21	0.16
B9	12.54 (5.95)	0.67	0.17	0.03
B10	10.25 (4.55)	0.77	0.42	0.09

Table III 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 de-trended log level of portfolio's dividends on the de-trended 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 last three columns, labelled "ACF", report sample autocorrelation function of the cointegrating residuals for lags 1, 5 and 10.

Table IV
Predictability Evidence

	EC-VAR			VAR		
	<u>Horizon</u>			<u>Horizon</u>		
	1	5	10	1	5	10
S1	0.17	0.40	0.36	0.06	0.04	0.01
S2	0.13	0.53	0.59	0.11	0.35	0.44
S3	0.10	0.40	0.53	0.03	0.08	0.13
S4	0.15	0.44	0.50	0.10	0.18	0.20
S5	0.16	0.36	0.49	0.15	0.21	0.26
S6	0.11	0.37	0.47	0.07	0.17	0.18
S7	0.09	0.29	0.46	0.08	0.17	0.23
S8	0.08	0.32	0.53	0.05	0.14	0.21
S9	0.04	0.28	0.46	0.01	0.12	0.16
S10	0.03	0.23	0.43	0.02	0.15	0.24
B1	0.04	0.19	0.32	0.00	0.03	0.04
B2	0.07	0.23	0.36	0.00	0.04	0.06
B3	0.23	0.27	0.43	0.14	0.06	0.07
B4	0.06	0.26	0.44	-0.01	-0.02	-0.03
B5	0.04	0.24	0.37	0.02	0.07	0.06
B6	0.08	0.26	0.37	0.06	0.03	0.00
B7	0.15	0.44	0.42	0.06	0.11	0.05
B8	0.17	0.49	0.41	0.11	0.21	0.18
B9	0.20	0.42	0.35	0.06	0.03	-0.01
B10	0.20	0.48	0.51	0.09	0.19	0.17
Mrk	0.04	0.25	0.44	0.04	0.18	0.27

Table IV presents the adjusted- R^2 for return projections implied by the EC-VAR specification and the growth-rates 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), as well as for the market portfolio (Mrk).

Table V
Consumption Betas and Mean Returns by Horizon

Portfolio	H o r i z o n					
	— 1 —		— 5 —		— 10 —	
	β_1	μ_1	β_5	μ_5	β_{10}	μ_{10}
S1	4.13 (2.43)	11.69	4.51 (4.66)	10.61	6.54 (4.26)	10.70
S2	2.09 (1.12)	11.05	1.82 (2.65)	9.99	4.21 (3.02)	9.94
S3	4.14 (1.29)	11.66	2.14 (2.38)	10.36	3.38 (2.29)	10.00
S4	3.52 (1.21)	11.35	2.09 (1.97)	10.19	3.56 (2.40)	10.05
S5	2.76 (0.99)	10.74	0.99 (1.72)	9.94	2.05 (2.12)	9.57
S6	3.07 (0.81)	10.11	1.43 (1.51)	9.17	2.50 (1.83)	8.91
S7	2.37 (0.86)	9.83	0.46 (1.37)	9.10	1.00 (1.67)	8.67
S8	1.62 (0.68)	9.17	-0.12 (1.35)	8.56	0.35 (1.68)	8.13
S9	1.58 (0.77)	8.44	0.60 (1.41)	7.81	1.02 (1.59)	7.43
S10	1.54 (0.56)	7.17	0.31 (1.09)	6.73	0.34 (1.07)	6.39
B1	1.81 (0.54)	6.58	-0.58 (1.40)	6.30	-0.83 (1.32)	5.97
B2	0.16 (0.56)	7.89	-1.69 (0.91)	7.83	-2.05 (0.86)	7.57
B3	-0.09 (0.37)	6.58	-1.79 (0.77)	6.66	-1.70 (0.84)	6.34
B4	1.48 (1.29)	7.20	-0.67 (1.99)	6.87	-0.59 (2.06)	6.38
B5	1.94 (0.86)	9.06	1.18 (1.58)	8.36	1.60 (1.69)	8.00
B6	3.18 (1.49)	9.15	2.75 (2.42)	8.71	3.27 (2.37)	8.46
B7	2.74 (0.98)	9.93	2.67 (1.53)	9.03	4.22 (1.63)	9.18
B8	4.35 (1.83)	11.73	4.40 (2.72)	11.26	6.37 (2.69)	11.77
B9	5.49 (2.31)	12.86	6.12 (3.75)	12.78	8.32 (3.43)	13.54
B10	3.89 (1.06)	11.78	2.14 (2.60)	10.73	4.33 (3.38)	11.01
Cross- Sectional StDev	1.42	1.92	2.04	1.72	2.81	1.99

Table V presents consumption betas and 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). Consumption betas are estimated as in equation (28) using the covariance matrices implied by the estimated error-correction VAR model. Standard errors are reported in parentheses. Mean returns for a given horizon s are computed as $\overline{\ln(r_i)} + 0.5\hat{\sigma}_{r,s}^2$. The last row of the table reports standard deviations of each column.

Table VI
Dividend Betas and Mean Returns by Horizon

Portfolio	H o r i z o n			
	— 1 —	— 5 —	— 10 —	Coint.
	β_1	β_5	β_{10}	δ
S1	13.12 (1.58)	10.23 (3.60)	9.89 (3.68)	9.62
S2	7.76 (0.58)	9.76 (1.45)	10.43 (1.69)	9.71
S3	6.79 (0.76)	9.23 (0.84)	8.86 (0.76)	6.69
S4	8.59 (0.70)	8.45 (0.99)	7.63 (1.16)	6.21
S5	7.78 (0.60)	6.37 (0.65)	5.39 (0.61)	4.41
S6	6.93 (0.48)	6.46 (0.84)	5.45 (0.86)	4.38
S7	4.84 (0.51)	3.62 (0.59)	2.90 (0.47)	2.42
S8	4.03 (0.39)	3.15 (0.57)	2.66 (0.49)	2.38
S9	3.57 (0.40)	2.96 (0.53)	2.45 (0.38)	2.38
S10	2.39 (0.35)	0.78 (0.24)	0.34 (1.07)	0.82
B1	3.35 (0.40)	0.08 (0.62)	-0.29 (0.25)	-0.27
B2	0.18 (0.43)	-1.59 (0.38)	-2.11 (0.21)	-2.59
B3	0.42 (0.19)	0.06 (0.25)	-0.27 (0.17)	-0.11
B4	3.79 (0.93)	0.74 (0.42)	0.40 (0.34)	0.82
B5	4.08 (0.48)	3.16 (0.59)	2.62 (0.51)	2.79
B6	6.31 (0.92)	4.88 (0.47)	4.49 (0.36)	4.83
B7	6.07 (0.45)	7.59 (0.72)	6.58 (0.51)	6.36
B8	6.95 (1.01)	8.59 (1.45)	8.88 (1.32)	9.70
B9	9.62 (1.58)	9.48 (2.66)	10.43 (2.04)	12.54
B10	15.12 (0.67)	12.43 (1.12)	10.79 (1.34)	10.25

Table VI presents dividend 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). Dividend betas are estimated as in equation (28) using the covariance matrices implied by the estimated error-correction VAR model. Standard errors are reported in parentheses. Mean returns for a given horizon s are computed as $\overline{\ln}(r_i) + 0.5\hat{\sigma}_{r,s}^2$. The last row of the table reports standard deviations of each column.

Table VII
Cross-Sectional Regressions by Horizon

Panel A: $\bar{R}_i = \lambda_0 + \lambda_{1,\delta}\delta_i + \varepsilon_i$

$\lambda_{1,\delta}$	\bar{R}^2
0.486 (0.053)	0.81

Panel B: $\bar{R}_i = \lambda_{0,s} + \lambda_{1,s}\beta_s + \varepsilon_i$

Horizon	Return Betas		Dividend Betas	
	$\lambda_{1,s}$	\bar{R}^2	$\lambda_{1,s}$	\bar{R}^2
1	1.187 (0.334)	0.75	0.423 (0.143)	0.66
2	0.982 (0.295)	0.62	0.409 (0.139)	0.76
5	0.724 (0.264)	0.73	0.371 (0.146)	0.76
10	0.650 (0.236)	0.84	0.433 (0.169)	0.80
Long Run	0.718 (0.255)	0.87	0.779 (0.309)	0.85

Table VII 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. Panel A reports the estimate of the market price of risk and the cross-sectional R^2 , utilizing the parameter of cointegration between portfolio's cash-flows and aggregate consumption as a measure of consumption risk. In Panel B, consumption risk for different investment horizons is measured by the corresponding consumption beta. The beta for a given horizon is calculated from the error-correction VAR model as in equation (28). The betas used in the first set of results are calculated using all return components; the second set of results reflect betas calculated using only the dividend growth portion of returns. All risk prices are expressed in annual percentage terms. Robust standard errors, reported in parentheses in Panel B, are computed by estimating time-series and cross-sectional parameters in one step via GMM. The number of lags used in Newey-West covariance estimator is 8.

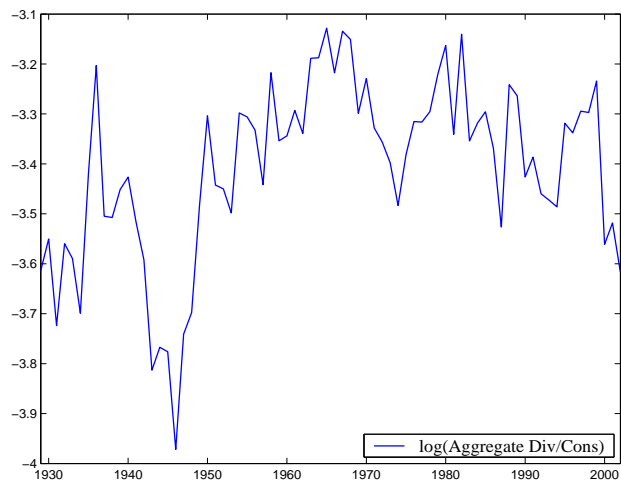
Table VIII
Cross-Sectional Regressions by Horizon with Cointegration Restrictions

	$\delta = 1$		No Time Trend	
Horizon	$\lambda_{1,s}$	\bar{R}^2	$\lambda_{1,s}$	\bar{R}^2
1	1.583 (0.393)	0.45	1.547 (0.275)	0.62
2	0.935 (0.586)	0.08	1.370 (0.411)	0.35
5	0.786 (0.540)	0.06	1.203 (0.310)	0.43
10	1.123 (0.406)	0.26	1.122 (0.192)	0.64
Long Run	0.804 (0.369)	0.16	1.095 (0.194)	0.62

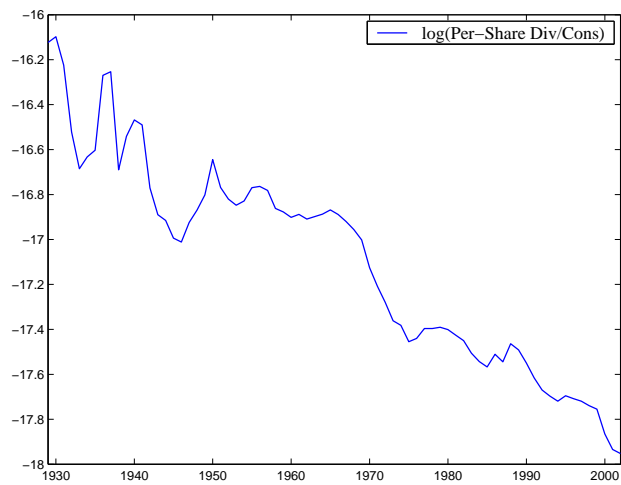
Table VIII presents results for cross-sectional regressions for a set of 10 portfolios sorted by market capitalization and 10 portfolios sorted by book-to-market ratio. Consumption risk for different investment horizons is measured by the corresponding consumption beta, measured using all return components. These components are measured under two restrictions on the cointegrating relation between dividends and consumption:

$$d_t = \tau_0 + \tau_1 t + \delta c_t + \epsilon_{d,t}.$$

In the left columns, we restrict the cointegration parameter, $\delta = 1$, for all assets. In the right columns, we suppress the time trend in the cointegration specification, $\tau_1 = 0$, 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 (28). All risk prices are expressed in annual percentage terms. Robust standard errors, reported in parentheses in Panel B, are computed by estimating time-series and cross-sectional parameters in one step via GMM. The number of lags used in Newey-West covariance estimator is 8.



(a) Aggregate Dividends to Consumption Ratio



(b) Per-Share Dividends to Consumption Ratio

Figure 1. Dividends to Consumption Ratio

Figure 1 plots the logarithm of the aggregate dividends to consumption ratio and the ratio of per-share dividends to consumption.

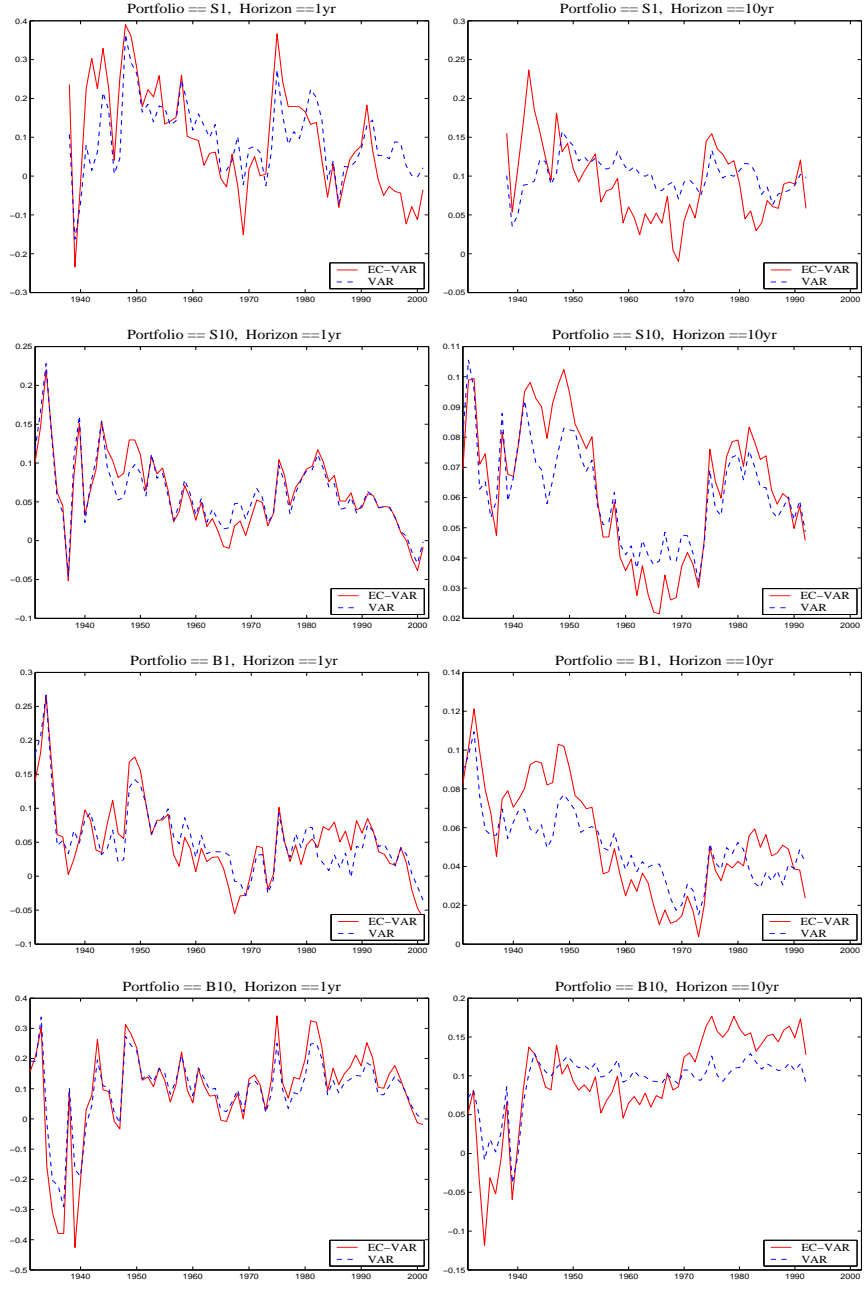


Figure 2. Expected Returns

Figure 2 plots one- and 10-year returns predicted by the EC-VAR model and the alternative, growth-rates VAR specification. The latter does not include the cointegrating residual. Expected returns are plotted for small- and large-size firms, as well as for growth and value portfolios.

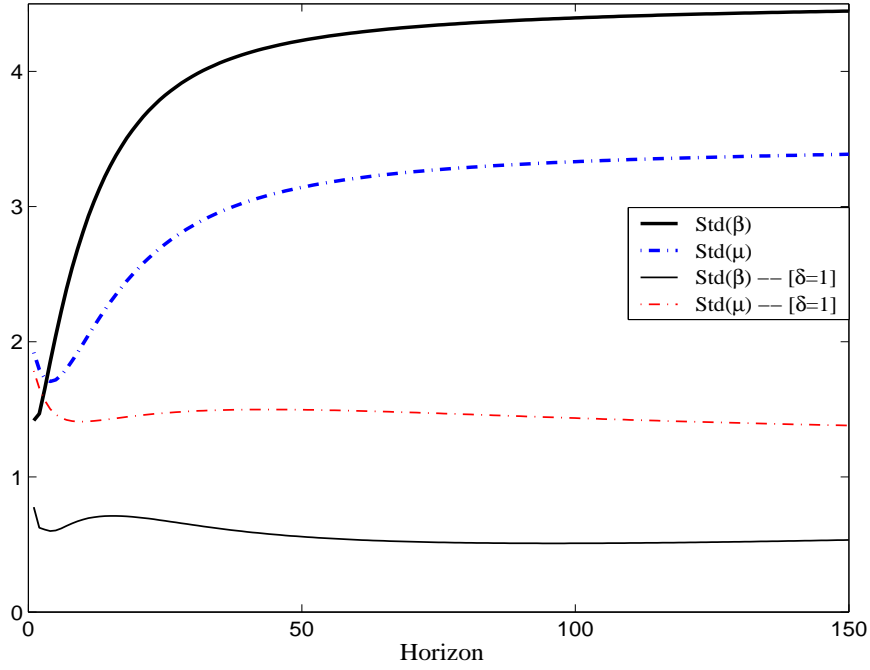




Figure 3. Cross-Sectional Dispersion of Betas and Mean Returns by Horizon

Figure 3 plots the cross-sectional standard deviations of consumption betas and mean returns across different investment horizons. Thick solid and dashed-dotted lines correspond to the statistics, implied by the unrestricted EC-VAR model. The two thin lines display the dispersion in risk measures and average returns, computed by restricting the cointegration parameter between portfolio's cash flows and consumption to be one.



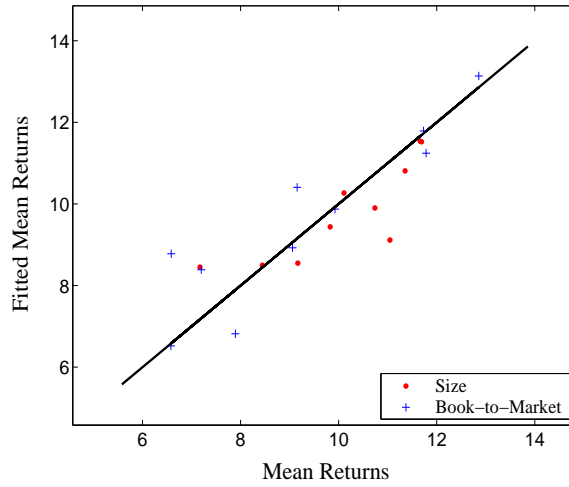
(a) EC-VAR



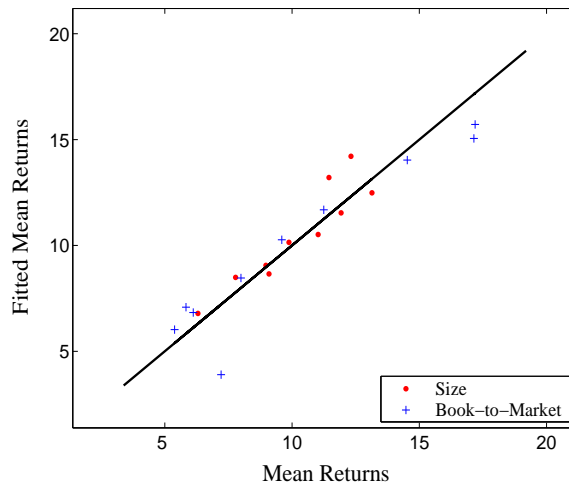
(b) EC-VAR ($\delta = 1$)

Figure 4. Sharpe Ratios by Horizon

Figure 4 plots the Sharpe ratios at various investment horizons for small and large market capitalization firms, growth and value stocks, as well as for the market portfolio. Panel (a) plots the SR's computed from the main EC-VAR specification, while Panel (b) corresponds to the specification with homogeneous long-run risks (that restricts the parameters of cointegration to be one for all the assets).



(a) Short-Run Horizon ($s=1$)



(b) Long-Run Horizon

Figure 5. Cross-Sectional Fit

Figure 5 displays the fit from the cross-sectional regressions for the investment horizon of one year, as well as in the long-run limit. The figures plot fitted expected returns, implied by the model, against mean returns.

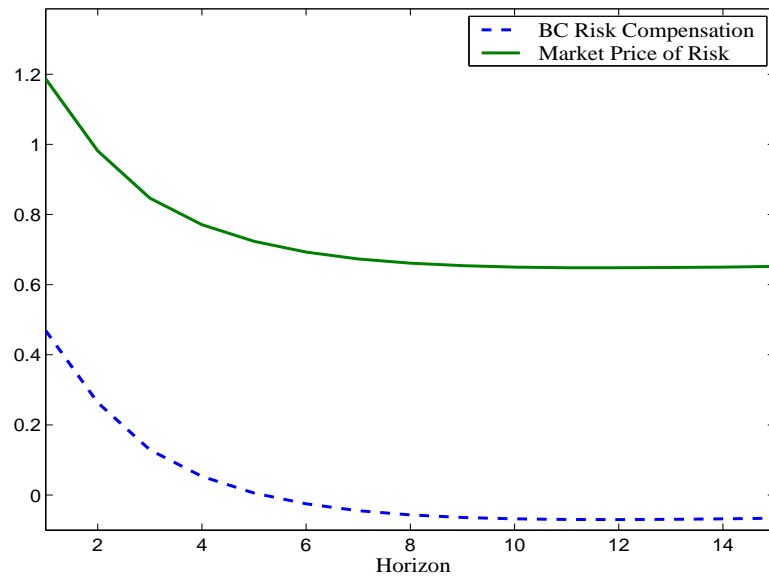


Figure 6. Risk Compensations by Horizon

Figure 6 plots the market price of risk along with the profile of implied business cycle risk compensations for investment horizons up to 15 years. The prices of risks are expressed in percentage terms.