Why is long-horizon equity less risky? A duration-based explanation of the value premium *

Martin Lettau
NYU, CEPR, NBER

Jessica A. Wachter
NYU, University of Pennsylvania, NBER

Preliminary

Comments Welcome

First draft: June 29, 2004

*Lettau: Department of Finance, Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012-1126; Email: mlettau@stern.nyu.edu; Tel: (212) 998-0378; http://www.stern.nyu.edu/~mlettau. Wachter: The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia, PA 19104; Tel: (215) 898-7634; Email: jwachter@wharton.upenn.edu; http://finance.wharton.upenn.edu/~jwachter/. We thank participants in the NYU reading group for helpful comments. Lettau acknowledges financial support from the National Science Foundation.

Why is long-horizon equity less risky?
A duration-based explanation of the value premium

Abstract

This paper proposes a dynamic risk-based model that captures the high expected returns on value stocks relative to growth stocks, and the failure of the capital asset pricing model to explain these expected returns. To model the difference between value and growth stocks, we introduce a cross-section of long-lived firms distinguished by the timing of their cash flows. Firms with cash flows weighted more to the future have high price ratios, while firms with cash flows weighted more to the present have low price ratios. We model how investors perceive the risks of these cash flows by specifying a stochastic discount factor for the economy. The stochastic discount factor implies that shocks to aggregate dividends are priced, but that shocks to the time-varying price of risk are not. As long-horizon equity, growth stocks covary more with this time-varying price of risk than value stocks, which covary more with shocks to cash flows. When the model is calibrated to explain aggregate stock market behavior, we find that it can also account for much of the observed value premium, the high Sharpe ratios on value stocks relative to growth stocks, and the outperformance of value (and underperformance of growth) relative to the CAPM.

1 Introduction

This paper proposes a dynamic risk-based model that captures the high expected returns on value stocks relative to growth stocks, and the failure of the capital asset pricing model to explain these expected returns. The value premium, first noted by Graham and Dodd (1934) is the finding that assets with high ratios of price to fundamentals (growth stocks) have low expected returns relative to assets with low ratios of prices to fundamentals (value stocks). Fama and French (1992) show that the capital asset pricing model of Sharpe (1964) and Lintner (1965) cannot account for the value premium. The CAPM predicts that expected returns should rise with the β with the market portfolio; however, value stocks have higher expected returns yet do not appear to have higher βs than growth stocks.

To model the difference between value and growth stocks, we introduce a cross-section of long-lived firms distinguished by the timing of their cash flows. Firms with cash flows weighted more to the future have high price ratios, while firms with cash flows weighted more to the present have low price ratios. Drawing an analogy to bonds, we can say that growth firms are high-duration assets while value firms are low-duration assets.

We model how investors perceive the risks of these cash flows by specifying a stochastic discount factor for the economy, or equivalently, an intertemporal marginal rate of substitution for the representative agent. Our model for the stochastic discount factor shares some of the features of the external habit formation of Campbell and Cochrane (1999). As in the model of Campbell and Cochrane (1999), the riskfree rate is constant. Moreover, we allow the price of risk to vary, implying that at certain times, investors require a greater return per unit of risk to hold equities than at others. A key difference is the relation between the price of risk and the aggregate dividend. In the model of Campbell and Cochrane (1999) they are tightly linked: a shock to the aggregate dividend moves agents closer to their habit level and raises the return they require for bearing risk.[1] In our model the return investors require for bearing risk moves independently of the aggregate dividend. We show that the correlation between the aggregate dividend and the price of risk determines, in large part, the ability of the model to fit the cross-section.

[1] In the benchmark case of Campbell and Cochrane (1999), aggregate dividends and aggregate consumption are the same. Campbell and Cochrane also examine a case where dividends are correlated with consumption. Because of this correlation, there is still a link between dividends and the price of risk.

We require our model to explain not only the cross-section of assets based on price ratios, but also aggregate stock market behavior. To insure our model has reasonable aggregate properties, we specify a process for the aggregate dividend. Firms are distinguished by the pattern of shares of this dividend that they pay each period. We specify share processes that are stationary, and explore the robustness of our results to different models of the share process. This modeling strategy, also employed by Menzly, Santos, and Veronesi (2004) and Santos and Veronesi (2004), ensures that the economy is stationary, and that all future dividends are marketed. We assume that log dividend growth is normally distributed with a time-varying mean. We calibrate the dividend process to match conditional and unconditional moments of the dividend process in the data. Stochastic discount factor parameters are chosen to fit the time series of aggregate stock market returns. Expected excess returns on equity are time-varying in the model, implying excess volatility and return predictability. Despite the simplicity of our model, we find that it can match unconditional moments of the aggregate stock market. In addition, the model produces predictability of dividends and returns close to that found in the data.

To test whether our model can explain the value premium, we sort firms into portfolios in simulated data. We find that risk premia, risk-adjusted returns, and Sharpe ratios increase as portfolios move from growth to value. We find a value premium (the return on a strategy that is long the extreme value portfolio and short the extreme growth portfolio) between 3.4% and 5.2% (depending on the share process) compared with a value premium of 4.9% in the data when portfolios are formed on the basis of book-to-market. We find an α on the value-minus-growth strategy of between 4.7% and 6.2%, compared with 5.6% in the data. These results do not arise because value stocks are more risky according to traditional measures; standard deviations and market βs increase slightly and then decrease, implying that the extreme value portfolio has a lower standard deviation and β than the extreme growth portfolio. Our model therefore accounts for a large fraction of the value premium noted in the data.

Our paper builds on previous literature that uses the concept of duration to better understand the cross-section of stock returns. Using the decomposition of returns into cash flow and discount rate components proposed by Campbell and Mei (1993), Cornell (1999) shows that growth companies, such as Amgen, whose cash flows are mainly idiosyncratic, may have high βs because of the duration of these cash flows, and the induced sensitivity of

prices to market-wide changes in discount rates. Dechow, Sloan, and Soliman (2002) measure cash flow duration of value and growth portfolios; they find that empirically, growth stocks have higher duration than value stocks and that this contributes to their higher βs. Brennan and Xia (2003) show in a theoretical model that βs increase in the maturity of the cash flows. Santos and Veronesi (2004) develop a model that links time variation in betas to time-variation in expected returns through the channel of duration, and show that this link is present in industry portfolios. Campbell and Vuolteenaho (2003) decompose the market return into news about cash flows and news about discount rates. They show that growth stocks have higher βs with respect to discount rate news than growth stock βs, consistent with the view that growth stocks are high duration assets. These papers all show that discount-rate risk is an important component of total volatility, and that growth stocks seem particularly subject to this discount-rate risk.

This paper also relates to the large and growing body of literature that explores the correlations of returns on value and growth stocks with sources of systematic risk. This literature looks at either conditional versions of traditional models (Jagannathan and Wang (1996), Lettau and Ludvigson (2001), Zhang and Petkova (2002)), or identifies a new source of risk that covaries more with value stocks than with growth stocks (Lustig and Nieuwerburgh (2002), Piazzesi, Schneider, and Tuzel (2002), Yogo (2003)). Another strand of literature relates observed returns on value and growth stocks to aggregate market returns or macroeconomic factors (Brennan, Wang, and Xia (2003), Liew and Vassalou (2000), Vassalou (2003), Campbell, Polk, and Vuolteenaho (2003)). The results in these papers raise the question of what it is, fundamentally, about the cash flows of value and growth stocks that produces the observed patterns in returns. Other work examines the dividends on value and growth portfolios directly: Cohen, Polk, and Vuolteenaho (2002) and Bansal, Dittmar, and Lundblad (2003) offer evidence that the cash flows of value stocks covary more with aggregate cash flows. The results in these papers raise the question of why this observed covariation leads to the value premium. Gomes, Kogan, and Zhang (2003) propose a general equilibrium model of firms that produces a cross-section of book-to-market ratios. However, the innovation of our model is to simultaneously capture the behavior of the aggregate stock market, and to produce a spread between risk-adjusted returns on value and growth portfolios on the order of magnitude seen in the data.

The paper is organized as follows. Section 2 organizes and updates the evidence that

portfolios formed on the basis of prices scaled by fundamentals produce spreads in expected returns. We show that when value is defined by book-to-market, earnings-to-price, or cash-flow-to-price, the expected return tends to increase as portfolios move from growth to value, the standard deviation and β with the market tend to decrease, and the Sharpe ratio and the α tends to increase. The differences in expected returns and αs between value and growth portfolios are statistically and economically large.

Section 3 presents our model for aggregate dividends and the stochastic discount factor. As a first step to solving for prices on the aggregate market and firms, we solve for prices on zero-coupon equity. These are claims to the aggregate dividend m-periods in the future. Because zero-coupon equity has a well-defined maturity, it provides a convenient window through which to view the role of duration in the model. Moreover, as the model has similarities to essentially affine term structure models (Dai and Singleton (2003), Duffee (2002)), the prices and risk premia on zero-coupon equity have interpretable, closed-form expressions. The aggregate market is the sum of all of the zero-coupon equity claims. We then introduce a long-lived cross-section of assets, defined by their share prices of the aggregate dividends. These assets are themselves portfolios of the zero-coupon equity, and together their cash flows and market values sum up to the cash flows and market values of the aggregate market.

Section 4 studies the implication of our model for the time series and the cross-section. We calibrate the model using the time series of the aggregate returns, dividends, and the price-dividend ratio. After choosing parameters to match aggregate time-series facts, we examine the implications for zero-coupon equity. We see that the parameters necessary to fit the time series imply risk premia, Sharpe ratios, and αs for zero-coupon equity that are increasing in the maturity. βs and volatility are non-monotonic, and thus do not explain the increase in risk premia. This shows that the model has the potential to explain the value premium. We then choose parameters of the share process to approximate the distribution of dividend, earnings, and cash flow growth found in the data, and produce realistic distributions of price ratios. We examine several functional forms for the shares. When share processes are calibrated in this way, and the resulting assets are sorted into portfolios, our model can explain a large fraction of the observed value premium. We also show that our model generates "rational moment effects" similar to those described in Johnson (2002). That is, our model can partially replicate the empirical finding of Jegadeesh and Titman (1993)

that stocks with high returns over the previous year ("winners") have substantially higher returns than stocks with low returns over the previous year ("losers"). Moreover, our model implies that winners exhibit positive risk-adjusted returns relative to the CAPM, while losers exhibit negative risk-adjusted returns, thus partially replicating the findings of Jegadeesh and Titman (2001).

Section 5 describes the intuition for our results. We show that the covariation of asset returns with the shocks depends on the duration of the asset. Consistent with the results of Campbell and Vuolteenaho (2003), growth stocks have greater betas with respect to discount rates than value stocks. This is the duration effect: because cash flows on growth stocks are further in the future, their prices are more sensitive to changes in discount rates. Growth stocks also have greater betas with respect to changes in expected dividend growth. Value stocks, on the other hand, have greater betas with respect to shocks to near-term dividends. The price investors put on bearing the risk in each of these shocks determines the rates of return on value and growth stocks. While shocks to near-term dividends are viewed as risky by investors, shocks to expected future dividends are hedges under our calibration. Moreover, though discount rates vary over time, shocks to discount rates are independent of shocks to dividends and are therefore not priced directly. Even though long-horizon equity is riskier according to standard deviation and market β, it is not seen as risky by investors because it loads on risks investors do not mind bearing.

2 Evidence on the value premium

Much previous literature has shown that portfolios of stocks with high prices relative to fundamentals have low future returns relative to stocks with low prices relative to fundamentals.[2] In this section, we update and organize this evidence by running statistical tests on portfolios formed on ratios of market to book value, price to earnings, price to dividends, and price to cash flow. We show that in all cases, the sorting produces differences in expected returns that cannot be attributed to market beta. Moreover, the alpha relative to the CAPM tends to increasing in the measure of value, as shown by Fama and French (1992) for the book-to-market ratio. In our model, firms are distinguished on the basis of their cash flows,

[2]See Graham and Dodd (1934), Basu (1977, 1983) , Ball (1978), Rosenberg, Reid, and Lanstein (1985), and Fama and French (1992).

thus earnings, dividends, and cash flows are equivalent. For this reason, it is especially of interest to investigate whether the value effect is apparent in portfolios formed according to different measures of value.

Table 1 shows summary statistics for portfolios of firms sorted into deciles on the basis of the three characteristics above, as well as on the basis of book-to-market. Data, available from the website of Ken French, are monthly, from 1952 to 2002. Excess returns are computed by subtracting monthly returns on the one-month Treasury Bill from the portfolio return. The first panel reports the mean excess return, the second the standard error on the mean, the third the standard deviation of the return, and the fourth the Sharpe ratio. Means and standard deviations are in annual percentage terms (multiplied by 1200 in the case of means and $\sqrt{12} \times 100$ in the case of standard deviations). Each panel first reports results for portfolios sorted on the two measures correspond most closely to the dividend-price ratios in the model. These are the earnings-to-price and the cash-flow to price measures. Each panel also reports results for portfolios sorted on the basis of the dividend yield and the book-to-market ratio.

Panel 1 shows that for all measures except the dividend-yield, the mean excess return increases as one moves from the bottom scaled-price decile (growth stocks) to the top scaled-price decile (value stocks). The increase is usually, but not always, monotonic. As shown in Panel 2, the average return on the portfolio that is long the extreme value portfolio and short the extreme growth portfolio is highly statistically significant, again except when portfolios are formed on the basis of the dividend yield.

Panel 3 shows that the standard deviation of the excess return tends to decrease as one moves from the bottom decile to the stop. This holds for all four scaled-price measures. Finally, Panel 4 shows that the Sharpe ratio increases as one moves from the bottom decile to the top across all four scaled-price measures. For example, when portfolios are formed on the basis of the earnings-to-price ratio, the bottom decile (growth) has a Sharpe ratio of 0.24. The Sharpe ratio increases steadily as the earnings-to-price ratio increases; the top decile has a Sharpe ratio of 0.72. Value stocks not only deliver high returns; they deliver high returns per unit of standard deviation.

The results in Table 1 suggest that portfolios formed on the basis of earnings-to-price, cash-flow-to-price, the dividend yield, and book-to-market, may be closely related. This is confirmed in Table 2, which shows the correlation of the top growth decile and the top value

decile. For the bottom decile (growth), the correlations are 0.93 or above; for the top decile (value), the correlations are 0.74 or above. In both cases, deciles formed by sorting on the dividend-yield are less highly correlated with the deciles formed by sorting on the other three variables than the deciles formed by sorting on the other three variables are with each other. This is consistent with the results in Table 1, which shows that results based on sorting on the dividend-yield were somewhat different than the other variables.

Following the same format as Table 1, Table 3 shows alphas, standard errors on alphas, betas, standard errors on betas, and R^2 statistics when portfolios are formed on the basis of each measure of value. Alpha is the intercept from an OLS regression of excess returns on the portfolio on excess returns on the return on the value-weighted NYSE-AMEX-NASDAQ index and multiplied by 1200. Beta is the slope from this regression. The alpha for the portfolio that is long the extreme value portfolio and short the extreme growth portfolio is statistically significant for all four sorting variables. Panel 1 of this table confirms the classic result that value stocks have high alphas relative to the CAPM. The story is consistent across all sorting variables, including the dividend-yield: alphas are negative for growth stocks, rise as we go from growth to value, and are negative for value stocks. As Panel 3 shows, betas tend to decline as we move from growth to value; except for the extreme growth portfolio, betas are below one for portfolios in the low deciles and greater than one for portfolios in the high deciles, across all sorting variables. Thus value stocks have positive alphas relative to the CAPM, and relatively low betas.

This section shows that, in the data, value stocks have higher expected excess returns and higher Sharpe ratios than growth stocks. Value stocks have large positive alphas relative to the CAPM, while growth stocks have negative alphas. Moreover, value stocks do not have higher standard deviations or higher betas than growth stocks. Thus any story that explains the value premium needs to take into account the fact that value stocks do not appear to be riskier than growth stocks according to traditional measures of risk. These empirical results not only hold when value is defined by the book-to-market ratio, they hold when value is defined according to the earnings-to-price or cash-flow-to-price ratios.

3 The model

This section presents our model. The first subsection discusses the assumptions on aggregate cash flows and on the pricing kernel. The second subsection solves for prices on equity that pays the aggregate dividend in a fixed number of years from now; we refer to these claims as "zero-coupon equity" and they form the fundamental building blocks of our more complex assets. Interpretable, closed-form expressions are available for prices and conditional risk premia for zero-coupon equity. The third subsection describes how zero-coupon equity aggregates up to the market. The fourth subsection discusses the model for long-lived assets defined by their shares in the aggregate dividend. These assets, like the aggregate market, are portfolios of the zero-coupon equity and their prices can be determined accordingly. Thus the intuition for risk premia and price variation for zero-coupon equity can be transferred to these long-lived assets.

3.1 Dividend growth and the pricing kernel

The model has three shocks: a shock to dividend growth, a shock to expected dividend growth, and a shock to the preference variable x_t. To model these shocks in a parsimonious fashion, we let ϵ_{t+1} denote a 3×1 vector of independent normal shocks that have unit standard deviation and are independent of any variables observed at time t. Let D_t denote the aggregate dividend in the economy at time t, and $d_t = \log D_t$. The aggregate dividend is assumed to follow the process

$$\Delta d_{t+1} = g + z_t + \sigma_d \epsilon_{t+1} \tag{1}$$

where z_t follows an AR(1) process:

$$z_{t+1} = \phi_z z_t + \sigma_z \epsilon_{t+1} \tag{2}$$

with $0 \leq \phi_z < 1$. The conditional mean of dividend growth is $g + \phi z_t$. Multiplying the shocks on dividend growth and z_{t+1} are 1×3 vectors σ_d and σ_z. The first component of σ_d is the loading on the first shock, the second component is the loading on the second shock, etc. Thus the conditional standard deviation of Δd_{t+1} is

$$\|\sigma_d\| = \sqrt{\sigma_d \sigma_d'}.$$

Similarly, the conditional standard deviation of z_t equals

$$\|\sigma_z\| = \sqrt{\sigma_z \sigma_z'},$$

while the conditional covariance is given by $\sigma_d \sigma_z'$. This model for dividend growth is also explored by Bansal and Yaron (2003), and by Campbell (1999).

We directly specify the pricing kernel for this economy. It is assumed that the price of risk is driven by a single state variable x_t that follows an AR(1) process:

$$x_{t+1} = (1 - \phi_x)\bar{x} + \phi_x x_t + \sigma_x \epsilon_{t+1} \tag{3}$$

with $0 \leq \phi_x < 1$. As above, σ_x is a 1×3 vector. This specification for the price of risk is used in a continuous-time setting by Brennan, Wang, and Xia (2003). For simplicity, we assume that the real riskfree rate, denoted $r^f = \log R^f$, is constant. Lastly, we need to make an assumption about which risks in the economy are priced. We could follow the affine term structure literature (see, e.g. Duffie and Kan (1996)) and allow all three shocks to be priced. For simplicity, and to reduce the number of degrees of freedom, we assume that only dividend risk is priced. This allows us to compare our models to the external habit formation models of Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004), where the one shock to the pricing kernel comes from aggregate consumption. The assumption that only dividend risk is priced implies that shocks to z_t and shocks to x_t will only be priced insofar as they correlate with Δd_{t+1}.

This specification of x_t, r^f, and the fact that only dividend risk is priced completely pins down the pricing kernel. We set

$$M_{t+1} = \exp\left\{-r^f - \frac{1}{2}x_t^2 - x_t \epsilon_{d,t+1}\right\} \tag{4}$$

where

$$\epsilon_{d,t+1} = \frac{\sigma_d}{\|\sigma_d\|}\epsilon_{t+1}.$$

The log-normality of M_{t+1} implies that

$$\begin{aligned} \ln E_t[M_{t+1}] &= -r^f - \frac{1}{2}x_t^2 + \frac{1}{2}x_t^2 \sigma_d \sigma_d' \|\sigma_d\|^{-2} \\ &= -r^f. \end{aligned}$$

Therefore, it follows from no-arbitrage that r^f is indeed the riskfree rate. The maximum Sharpe ratio will be achieved by the asset that is most negatively correlated with M_{t+1}. Following the same argument as in Campbell and Cochrane (1999), we note that the maximum

Sharpe ratio is given by

$$\frac{\sigma_t(M_{t+1})}{E_t[M_{t+1}]} = \sqrt{e^{x_t^2} - 1} \approx |x_t|.$$

The question naturally arises of how to interpret the variable x_t. In the models of Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004), time-variation in the price-of-risk is a decreasing function of the surplus consumption ratio. Conditionally, the price-of-risk is perfectly negatively correlated with consumption growth (and hence aggregate dividend growth). The corresponding assumption here would be to set $\sigma_x/\|\sigma_x\| = -\sigma_d/\|\sigma_d\|$. However, we depart from these papers by assuming that shocks to x_{t+1} are uncorrelated with shocks to Δd_{t+1} and shocks to z_{t+1}. In our model, shocks to x_{t+1} can be interpreted as shocks to preferences or changes in sentiment. These shocks are uncorrelated with changes in fundamentals. Below, we explain the implications for security returns of this departure from habit formation.

3.2 Prices of zero-coupon equity

The fundamental building-blocks of the long-lived assets in our economy are "zero-coupon" equity.[3] Let P_{nt} be the price of an asset that pays the aggregate dividend n periods from now. In this subsection, we solve for the price of zero-coupon equity in closed form. Let $R_{n,t+1}$ denote the return on zero-coupon equity maturing in n periods. That is:

$$R_{n,t+1} = \frac{P_{n-1,t+1}}{P_{nt}}. \tag{5}$$

The returns $R_{n,t+1}$ form a term structure of equities, analogous to the term structure of interest rates. No-arbitrage implies the following Euler equation:

$$E_t\left[M_{t+1}R_{n,t+1}\right] = 1 \tag{6}$$

which implies that P_{nt} and $P_{n-1,t}$ satisfy the recursive relation:

$$P_{nt} = E_t\left[M_{t+1}P_{n-1,t+1}\right] \tag{7}$$

with boundary condition

$$P_{0t} = D_t \tag{8}$$

[3]The notion of breaking the aggregate dividend into its zero-coupon claims, and using affine term structure techniques to calculate the value of these claims is also applied in Ang and Liu (2003), Bakshi and Chen (1996), Bekaert, Engstrom, and Grenadier (2004), Johnson (2002), Wachter (2003), and Wilson (2003).

since equity maturing today must be worth the aggregate dividend. We conjecture that a solution to (7) and (8) satisfies:

$$\frac{P_{nt}}{D_t} = F(x_t, z_t, n) = \exp\left\{A(n) + B_x(n)x_t + B_z(n)z_t\right\}. \tag{9}$$

By the boundary condition, it must be that $A(0) = B_x(0) = B_z(0) = 0$. Substituting (9) into (7) produces

$$E_t\left[M_{t+1}\frac{D_{t+1}}{D_t}F(x_{t+1}, z_{t+1}, n-1)\right] = F(x_t, z_t, n). \tag{10}$$

Matching coefficients on z_t, x_t and the constant implies that

$$B_z(n) = \frac{1-\phi_z^n}{1-\phi_z}, \tag{11}$$

while $B_x(n)$ and $A(n)$ satisfy

$$B_x(n) = B_x(n-1)\left(\phi_x - \sigma_x\frac{\sigma_d'}{\|\sigma_d\|}\right) - \left(\sigma_d + B_z(n-1)\sigma_z\right)\frac{\sigma_d'}{\|\sigma_d\|} \tag{12}$$

$$A(n) = A(n-1) - r^f + g + B_x(n-1)(1-\phi_x)\bar{x} + \frac{1}{2}V_{n-1}V_{n-1}', \tag{13}$$

where

$$V_{n-1} = \sigma_d + B_z(n-1)\sigma_z + B_x(n-1)\sigma_x,$$

and $B_x(0) = 0$, $A(0) = 0$. This confirms the conjecture (9).[4]

Note that $B_z > 0$ for all n. This makes sense; the higher is z_t, the higher is expected dividend growth, hence the higher is the price of equity that pays aggregate dividend growth in the future. Because expected dividend growth is persistent, and because D_{t+n} cumulates shocks between t and n, the greater is n, the greater the effect of changes in z_t on the price. B_z is increasing in n, and converges to $1/(1-\phi_z)$ as n approaches infinity.

The economics behind B_x are more complicated. For the case we focus on x_{t+1} is independent of Δd_{t+1}, so $\sigma_x\sigma_d' = 0$. Moreover, in our calibration $\sigma_z\sigma_d'$ is small compared to $\sigma_d\sigma_d'$.

[4]The fact that price-dividend ratios are exponential affine in the state variables invites a comparison to the affine term structure literature, where bond prices are exponential affine in the state variables. In fact, this model is related to the essentially affine class of continuous-time term structure models explored by Dai and Singleton (2003) and Duffee (2002). Our model is essentially affine rather than affine because the stochastic discount factor is quadratic, as a result of the homoscedastic price-of-risk variable. Ang and Piazzesi (2003) examine a discrete-time essentially affine term structure model.

It follows that $B_x(n) < 0$ for all n. This makes sense: an increase in x_t leads to an increase in risk premia and a decrease in prices.[5] We further explore the intuition behind $B_x(n)$ in Section 4. Finally, A_n is a constant term that determines the level of price-dividend ratios. The level depends on the average growth rate of dividends less the riskfree rate, as well as on the average level of the price of risk ($\bar{x}$). The remaining term, $\frac{1}{2}V_{n-1}V_{n-1}'$ is a Jensen's inequality adjustment, and arises because we are taking the expectation of a log-normal variable.

In order to understand risk premia on the more complicated assets, it is helpful to understand risk premia on zero-coupon equity. To gain an understanding of the model, we compute $\log E_t[R_{n,t+1}/R^f]$, following Campbell (1999).[6] Define $r_{n,t+1} = \log R_{n,t+1}$. It follows from (9) that $r_{n,t+1}$ can be written as

$$r_{n,t+1} = E_t[r_{n,t+1}] + \sigma_t(r_{n,t+1})\epsilon_{t+1} \tag{14}$$

where

$$\sigma_t(r_{n,t+1}) = V_{n-1} = \sigma_d + B_x(n-1)\sigma_x + B_z(n-1)\sigma_z \tag{15}$$

Therefore returns are conditionally log-normally distributed, and we can re-write the conditional Euler equation (6) as

$$E_t\left[\exp\left\{-r^f - \frac{1}{2}x_t^2 - x_t\epsilon_{d,t+1} + E_t[r_{n,t+1}] + \sigma_t(r_{n,t+1})\epsilon_{t+1}\right\}\right] = 1$$

Taking logs of both sides and solving for the expectation produces the relation

$$\begin{aligned}
\log E_t[R_{n,t+1}/R^f] &= E_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_t(r_{n,t+1})\sigma_t(r_{n,t+1})' \\
&= \sigma_t(r_{n,t+1})\frac{\sigma_d'}{\|\sigma_d\|}x_t \\
&= (\sigma_d + B_x(n-1)\sigma_x + B_z(n-1)\sigma_z)\frac{\sigma_d'}{\|\sigma_d\|}x_t
\end{aligned} \tag{16}$$

Risk premia on zero-coupon equity depend on the loadings on each of the sources of risk, multiplied by the "price" of each source of risk. In our base case, the term $\sigma_x\sigma_d'$ disappears, so the loading on shocks to x_t, $B_x(n)$, is not relevant for risk premia on zero-coupon equity. In other cases we will examine, this term becomes important. Also determining risk premia

[5] In an alternative setting, it might be that $(\sigma_d + B_z(n-1)\sigma_z)\sigma_d < 0$. In this case, an increase in x_t would decrease risk premia and increase prices.

[6] When we match the simulated model to the data, we will compute $E[R_{t+1} - R^f]$.

is the loading on z_t, $B_z(n)$, and the price of z_t-risk, given by $\|\sigma_d\|^{-1}\sigma_z\sigma_d' x_t$. In what follows, similar reasoning can be used to understand the price of risk of the aggregate market and of firms, all of which are portfolios of these underlying assets.

3.3 Aggregate market

The aggregate market is the claim to all future dividends. Accordingly, its price-dividend ratio is the sum of the ratios of price to aggregate dividends of the zero-coupon equity described in the section above:

$$\frac{P_t^m}{D_t} = \sum_{n=1}^{\infty} \frac{P_{nt}}{D_t} = \sum_{n=1}^{\infty} \exp\left\{A(n) + B_x(n)x_t + B_z(n)z_t\right\} \tag{17}$$

Appendix B gives necessary and sufficient conditions on the parameters such that (17) converges for all x_t and z_t. The return on the aggregate market equals:

$$\begin{aligned} R_{t+1}^m &= \frac{P_{t+1}^m + D_{t+1}}{P_t^m} \\ &= \frac{(P_{t+1}^m/D_{t+1}) + 1}{P_t^m/D_t}\frac{D_{t+1}}{D_t} \end{aligned} \tag{18}$$

3.4 Firms

Zero-coupon equity illustrates how duration matters to risk premia in a particularly stark way. However, there is no obvious analogue of zero-coupon equity in the data. Instead, in the data there are long-lived securities that pay a sequence of cash flows over time. We want to construct a cross-section of securities that sum up to the aggregate market portfolio. Moreover, we want to insure that no one security comes to dominate the market portfolio over time; that is, the cross-sectional distribution of dividends, returns, and ratios of prices to aggregate dividends should be stationary. In order to accomplish this, we follow Menzly, Santos, and Veronesi (2004) and specify the share each security has in the aggregate dividend process D_{t+1}. The continuous-time framework of Menzly, et al. allows them to specify the share process as stochastic, yet still keep shares between 0 and 1. This is more difficult in discrete time, and for this reason we adopt the simplifying assumption that the share process is deterministic.

Suppose there are N long-lived "firms" in the economy. Define an N-vector of shares, s_i, such that $s_i \geq 0$ and $\sum_{i=1}^N s_i = 1$. At time t, we define firm i as the asset that pays

dividend $s_i D_t$ today, a dividend of $s_{i+1} D_{t+1}$ next period, etc. We specify s_i as a function of i for $1 \leq i \leq N$, and set

$$s_i = s_{(i-1 \mod N)+1}$$

for all other j. By this definition, firm i becomes firm $i+1$ next period. For example, at time t, firm 1 pays dividend $s_1 D_t$ and has ex-dividend price:

$$P^F_{1,t} = s_2 P_{1,t} + s_3 P_{2,t} \overbrace{+ \cdots +}^{N-4 \text{ terms}} s_N P_{N-1,t} + s_1 P_{N,t} + s_2 P_{N+1,t} + \cdots$$

At time $t+1$, this firm is now firm 2, pays dividend $s_2 D_{t+1}$, and has ex-dividend price:

$$P^F_{2,t+1} = s_3 P_{1,t+1} + s_4 P_{2,t+1} \overbrace{+ \cdots +}^{N-5 \text{ terms}} s_N P_{N-2,t+1} + s_1 P_{N-1,t+1} + s_2 P_{N,t+1} + s_3 P_{N+1,t+1} + \cdots$$

Equation (6) implies that these prices are consistent with no-arbitrage:

$$\begin{aligned} E_t[M_{t+1}\left(s_2 D_{t+1} + P^F_{2,t+1}\right)] &= s_2 P_{1,t} + E_t[M_{t+1}\left(s_3 P_{1,t+1} + \cdots + s_N P_{N-2,t+1} + s_1 P_{N-1,t+1} + \cdots\right)] \\ &= s_2 P_{1,t} + s_3 P_{2,t} + \cdots + s_N P_{N-1,t} + s_1 P_{N,t} + s_2 P_{N+1,t} + \cdots \\ &= P^F_{1,t} \end{aligned}$$

More generally, firm $k < N$ pays dividend $s_k D_t$ and has price

$$P^F_{k,t} = s_{k+1} P_{1,t} \overbrace{+ \cdots +}^{N-(k+2) \text{ terms}} s_N P_{N-k,t} + s_1 P_{N-k+1,t} + s_2 P_{N-k+2,t}$$

while firm N pays dividend $s_N D_t$ and has price

$$P^F_{N,t} = s_1 P_{1,t} + s_2 P_{2,t} \overbrace{+ \cdots +}^{N-3 \text{ terms}} s_N P_{N,t} + s_1 P_{N+1,t} + s_2 P_{N+2,t}$$

and so forth. Note that firm N becomes firm 1 next period. The same argument as above shows that these prices are consistent with no-arbitrage. This structure insures that the economy is stationary, that in each period the sum of the dividends across all firms equals the aggregate dividend, and that all future dividends are marketed as of date t. Beyond these requirements, the key element of this structure is that it generates dispersion in when firms pay dividends. Other models of firms which generate such dispersion, such that the distribution of firms is stationary and sums to the market should yield results similar to the ones we describe below.

Given prices for firms, the ratio of price over the one-period dividend equals

$$\frac{P^F_{k,t}}{D^F_{k,t}} = \frac{P^F_{k,t}}{s_k D_t}. \tag{19}$$

Because $P^F_{k,t}/D_t$ is a function of the state variables x_t and z_t, the price-dividend ratio for the firm is also a function of the state variables. Returns on the firm are given by

$$\begin{aligned} R^F_{k,t+1} &= \frac{P^F_{k+1,t+1} + D^F_{k+1,t+1}}{P^F_{k,t}} \\ &= \frac{(P^F_{k+1,t+1}/D^F_{k+1,t+1}) + 1}{P^F_{k,t}/D^F_{k,t}} \frac{D_{t+1}}{D_t} \frac{s_{k+1}}{s_k}. \end{aligned} \tag{20}$$

Note that all firms in this economy are ex-ante identical in the share of the aggregate dividend they will pay, they are simply "out of phase" with each other. Because the firms are ex-ante identical, the market values are very close. A more complex model would be required to account for differences in firm size.

4 Implications for Equity Returns

To study implications for the aggregate market and the cross-section, we simulate 50,000 quarters from the model. Given simulated data on shocks ϵ_{t+1} and state variables x_{t+1} and z_{t+1}, we compute ratios of prices to aggregate dividends for zero-coupon equity from (9), the price-dividend ratio for the aggregate market from (17), and the price-dividend ratio for firms from (19). Returns can then be computed using (5) for zero-coupon equity, (18) for the market, and (20) for firms.

As discussed below, we calibrate the model to annual data set of Campbell (1999) that begins in 1890 and ends in 2002. So that our simulated values are comparable to the annual values in the data, we aggregate up to an annual frequency. Annual flow variables (returns, dividend growth) are constructed by compounding their quarterly counterparts. Price-dividend ratios for the market and for firms are constructed analogously to annual price-dividend ratios in the Campbell data set. We divide the price by the current dividend on the asset, plus the previous three quarters of dividends on the asset.

Section 4.1 describes the calibration of our model to the aggregate time series. Section 4.2 shows the implications for the behavior of the aggregate market and dividend growth and discusses the fit to the data. Section 4.3 discusses implications for prices and returns

on zero-coupon equity. While zero-coupon equity have no analogue in the data, they are a useful construct in that they allow us to illustrate the properties of the model in a stark way. Section 4.4 discusses the calibration of the share processes which determine the prices of long-lived assets ("firms"), and describes implications of the model for portfolios formed on the basis of scaled-price ratios.

4.1 Calibration

Following Menzly, Santos, and Veronesi (2004), we calibrate the model to provide a reasonable fit to aggregate data. We then ask whether the model can match moments of the cross-section. In order to accurately capture the characteristics of our persistent processes, we use the century-long annual data set of Campbell (1999), which we update through 2002. The riskfree rate is the return on 6-month commercial paper purchased in January and rolled over in July. Stock returns, prices, and dividends are for the S&P 500 index. More details on data construction are contained in the Data Appendix for Campbell (1999). All variables are adjusted for inflation.

We set r^f equal to the mean of the riskfree rate in the sample, which is 1.93% in our data set. The average dividend growth in the sample is 2.28%, therefore this is our value for g. Less straightforward is the determination of the process z_t, the expected mean of dividend growth. This process, strictly speaking, is unobservable to the econometrician. However, Lettau and Ludvigson (2002) show that if consumption growth follows a random walk and if the consumption-dividend ratio is stationary, the consumption-dividend ratio captures all the predictability in dividend growth. Therefore the consumption-dividend ratio can be identified with z_t up to an additive and multiplicative constant.

For the purposes of calibration, we adopt the set-up of Lettau and Ludvigson (2002) and calibrate the autocorrelation of z_t and the correlation between shocks to expected dividend growth and shocks to z_t using the consumption-dividend ratio.[7] In our annual sample, the consumption-dividend ratio has a persistence of 0.91 and a conditional autocorrelation of -0.83. This still leaves the conditional standard deviations $\|\sigma_d\|$ and $\|\sigma_z\|$. We set the

[7] An equivalent way of writing down our model would be to assume a process, called consumption, that follows a random walk, and model the consumption-dividend ratio as an AR(1) process. Note however that consumption plays no special role in our model.

conditional standard deviation $\|\sigma_d\|$ to match unconditional standard deviation of annual dividend growth in the data. Our empirical results imply a standard deviation of z_t that is small relative to the standard deviation of dividend growth. Despite the fact that dividend is predictable at long horizons by the consumption-dividend ratio, the consumption-dividend ratio has very little predictive power for dividend growth at short horizons (with an R^2 of 3%). Moreover, the autocorrelation of dividend growth is relatively low (-.09%). We show that $\|\sigma_z\| = .0016$ (.0032 per annum) produces similar results in simulated data.

Remaining parameters are $\bar{x}$, ϕ_x, and $\|\sigma_x\|$. Because the variance of expected dividend growth is small, the autocorrelation of the price-dividend ratio is primarily determined by the autocorrelation of x. We therefore set $\phi_x = 0.87^{\frac{1}{4}} = 0.966$, to match the autocorrelation of 0.87 of the price-dividend ratio is annual data. We choose $\|\sigma_x\|$ to equal 0.12, or 0.24 per annum, to match the volatility of the log price dividend ratio. We choose $\bar{x}$ so that the maximal Sharpe ratio, when x_t is at its long-run mean, is 0.70, which is in the middle of the range of the maximum Sharpe ratios formed on the basis of scaled-price ratios. From the calculations in Section 3.1, $\bar{x}$ solves $\sqrt{e^{\bar{x}^2} - 1} = 0.70$. This translates into $\bar{x} = 0.625$. As discussed in the subsequent section, this produces an average Sharpe ratio for the market that is 0.41, somewhat higher than the data equivalent of 0.33. However, expected stock returns are measured with noise, and 0.41 is still below the Sharpe ratio in postwar data.

To link the conditional standard deviation of Δd_{t+1}, z_{t+1}, and, x_{t+1}, and the conditional correlation of Δd_{t+1} and z_{t+1} with the vectors σ_d, σ_z, σ_x, we assume, without loss of generality, that the 3×3 matrix

$$\begin{bmatrix} \sigma_d \\ \sigma_z \\ \sigma_x \end{bmatrix}$$

is lower triangular. Thus

$$\epsilon_{1,t+1} = \epsilon_{d,t+1}$$

so that σ_d has a nonzero first element equal to $\|\sigma_d\|$ and zero second and third elements. σ_z has a nonzero first and second element and zero third element. The first two elements are identified by $\|\sigma_z\|$ and the covariance $\sigma_d \sigma_z'$. We focus on the case where x_{t+1} is independent of Δd_{t+1} and z_{t+1}, so the first and second elements of σ_x equal zero, and the third equals $\|\sigma_x\|$. Table 4 summarizes these parameter choices.

4.2 Implications for the Aggregate Market and Dividend Growth

Table 5 presents statistics from simulated data, and the corresponding statistics computed from the data. The volatility of the price-dividend ratio is fit exactly, and the autocorrelation of the price-dividend ratio is very close (0.87 in the data versus 0.88 in the model). This is not a surprise because $\|\sigma_x\|$ and ϕ_x were set so that the model fits these parameters. The model produces a mean of the price-dividend ratio which is a bit too low, 20.1 in the model versus 25.6 in the data. Matching this statistic is a common difficulty for models of this type: for example, Campbell and Cochrane (1999) also find an average price-dividend ratio of 18.2. As they explain, this statistic is poorly measured due to the persistence of the price-dividend ratio. The model fits the volatility of equity returns (19.2% in the model versus 19.4% in the data), though it produces an equity premium that is a bit too high (7.9% in the model versus 6.3% in the data). As with the mean of the price-dividend ratio, the average equity premium is measured with noise. In the long-annual data set, the annual auto-correlation of returns is slightly positive (.03). In our model, the auto-correlation is slightly negative (-.02). The autocorrelation of dividend growth is small and negative (-.03), just as in the data (-.09).

Table 6 reports the results of long-horizon regressions of continuously compounded excess returns on the log price-dividend ratio in the model and in the data. As shown by Campbell and Shiller (1988), and Fama and French (1989), and replicated in our sample, high price-dividend ratios predict low excess returns. The coefficients rise with the horizon. The R^2 start small, at 0.05 at an annual horizon, and rise to 0.31 at a horizon of ten years. The t-statistics, using auto-correlation and heteroscedasticity-adjusted standard errors, are significant at the 5% level. The simulated data exhibits the same pattern. The coefficients rise with the horizon. The R^2 start at 0.06 and rise to 0.28. We conclude that the model generates a reasonable amount of return predictability.

Table 6 reports the results of long-horizon regressions of dividend growth on the price-dividend ratio. As Campbell and Shiller (1988) show dividend growth is not predictable by returns, contrary to what might be expected from a dividend-discount model. This result also holds true in our data set: the coefficients from a regression of dividend growth on the price-dividend ratio are always insignificant and have small R^2 statistics. In contrast, the consumption-dividend ratio predicts dividend growth. The coefficients are significant, and

the adjusted R^2 statistics start at 3% for an annual horizon and rise to 25% for a horizon of ten years.

Our model replicates both of these findings. Despite the fact that the mean of dividends is time-varying, dividends are only slightly predictable by the price-dividend ratio. A regression of simulated dividend growth on the simulated price-dividend ratio produces R^2 that range from 2% to 9% at a horizon of 10 years (in the data, the adjusted R^2 range from 0 to 5%). By contrast, dividends are predictable by z_t. Here, the R^2 statistics range from 4% to 24%, close to the values in the data. We conclude our model captures the pattern of dividend predictability found in the data.

4.3 Prices and Returns on Zero-Coupon Equity

Figure 1 plots the solution for $A(n)$, $B_z(n)$ and $B_x(n)$ as a function of n for the parameter values given above. $A(n)$ is steadily decreasing. This is a necessary feature for convergence of the solution for all x_t and z_t, and it makes economic sense: the further the payoff is in the future, the lower the value of the security when the state variables are at their long-run means. What generates the decrease is the average price-of-risk $\bar{x}$ and the riskfree rate r^f. Counteracting this decrease is average dividend growth g and the Jensen's inequality term. The net effect is that $A(n)$ is decreasing in n.

In contrast, $B_z(n)$ is positive, increasing in n and asymptotes to a value of $1/(1-\phi_z)$. The intuition for this variable was explained in Section 3.1. As discussed in Section 3.1, $B_x(n)$ is negative. This implies that an increase in the price of risk x_t leads to a decrease in valuations. Note that $B_x(n)$ is non-monotonic in n. It starts at 0, decreases to below -1, then increases, and eventually converges to a value near -0.6. It is not surprising that $B_x(n)$ initially decreases in maturity. This is the duration effect: the longer the maturity, the more sensitive is the asset to changes in the discount rate. More curious is the fact that B_x rises after a maturity of 50 quarters. This is because the duration effect is countered by the increase in $B_z(n)$. Because expected dividend growth and dividend growth are negatively correlated, shocks to expected dividend growth act as a hedge. Moreover, as the plot of B_z shows, expected dividend growth becomes more important the longer the maturity of the equity. Therefore equity that pays in the far future is less sensitive to changes in x_t than equity that pays in the medium term, though both are more sensitive than short-horizon

equity.

Figure 2 plots the ratios of price to aggregate dividends for zero-coupon equity as a function of maturity (Equation 9). The top panel sets z_t to be two long-run standard deviations $(2\|\sigma_z\|/(1-\phi_z^2)^{1/2})$ below its long-run mean and plots the price-dividend ratio for x_t at its long run mean and two long-run standard deviations $(2\|\sigma_x\|/(1-\phi_x^2)^{1/2})$ above and below the long-run mean. Not surprisingly, prices are increasing in expected dividend growth z_t for all values of x_t and for all values of the maturity. Moreover, for all values of the maturity and all values of z_t, prices decrease in x_t. The higher are conditional expected returns, the lower are prices.

For most values of z_t and x_t, prices decline with maturity. Generally, the further in the future the asset pays the aggregate dividend, the less it is worth today. Exceptions occur when x_t is two standard deviations below its long-run mean. In this cases, the premium for holding risky securities is negative in the short term, so short-horizon payoffs are discounted by more than long-horizon payoffs. Because x_t reverts back to its long-run mean, this effect is transitory and only holds at the short end of the equity "yield curve". The greater is z_t, the longer the effect persists. This makes sense: when expected dividend growth is high, equity that pays the aggregate dividend further in the future will go up in price more than equity that pays the aggregate dividend in the present. When z_t is two standard deviations above its long-run mean and x_t is two standard deviations below its long-run mean, the price of zero-coupon equity increases with maturity out to about 7 years, and then decreases again.

Figure 3 shows statistics for annual returns on zero-coupon equity. Annual returns are calculated by compounding quarterly returns are computed using (5). The top panel shows that the risk premium $ER_{i,t+1} - R^f$ decreases monotonically with maturity. The effect is economically large: for equity that pays a dividend in the next two years, the risk premium is 18% while the risk premium declines to 4% for equity that pays a dividend 40 years from now.

The second panel of Figure 3 plots the volatility of annual returns. The volatility initially increases with maturity, and then begins to decrease monotonically at a maturity of ten years. For long-horizon equity, increased risk premia are not accompanied by increased standard deviation. The third panel of Figure 3 shows that the unconditional Sharpe ratio declines monotonically in maturity from a value of 0.8 to a value of 0.2. Even for short-horizon equity,

the volatility increases less than the mean. These results suggest that the model has the potential to explain the patterns described in Table 1, if indeed value firms are stocks that pay dividends in the near future. Firms are portfolios of zero-coupon equity. Those firms that have more weight in lower-maturity equity will have higher expected returns, higher Sharpe ratios, and possibly lower variance, than firms that have more weight in equity of greater maturity.

Figure 4 shows the results of regressing simulated returns on zero-coupon equity on simulated returns on the market portfolio. The top panel plots the regression α, the middle panel the β, and the last panel the R^2 from the regression. As in Figure 3, returns are annual. The first panel shows that the α relative to the CAPM is decreasing in maturity over most of the range, and increases very slightly for long-duration equity. For the shortest-duration equity the α is as high as 11%. The α falls below zero for equity maturing in 5 or more years. The α never falls below -5%. Thus the model produces relatively large positive αs and relatively small negative αs, just as in the data, as shown in Table 3.

The second panel of Figure 4 shows the β. The β first increases, and then, beginning with a maturity of about ten years, decreases slowly as a function of maturity. The βs for zero-coupon equity lie in a relatively narrow range; the lowest β (for very long horizon equity) is about 0.7. The highest β (for equity of about ten years) is 1.5. The β for the shortest-horizon equity is about 0.9. This plot shows that at least for short-horizon equity, high αs are not necessarily accompanied by high βs. This also gives the model the potential to match the data described in Table 3.

While the simplicity of zero-coupon equity makes it a convenient way to illustrate the properties of the model, it does not have a direct interpretation in terms of value and growth. The price-dividend ratio is not well-defined because zero-coupon equity only pays dividends during a single quarter. For this reason, we turn to a model of firms: long-lived assets that have nonzero cash flows in every period.

4.4 Implications for the Cross-Section of Returns

This section shows the implications of the model for portfolios formed on the basis of price ratios. Following Menzly, Santos, and Veronesi (2004), we exogenously specify a share process for cash flows on long-lived assets. In each year of simulated data, we sort these assets into

deciles based on the ratio of price to dividends (or equivalently, earnings or cash flows) and form portfolios of the assets within each decile. This follows the procedure used in empirical studies of the cross-section (e.g. Fama and French (1992)). We then perform statistical analysis on the portfolio returns.

Specifying the share process

The model for firms is decidedly simple and somewhat mechanical. Our objective is to create dispersion in the timing of cash flows across firms in a way that captures the cross-sectional distribution of cash flow growth in the data. Our modeling strategy accomplishes this in a particularly simple way.

As described in Section 3.4, specifying a model for firms is equivalent to specifying a sequence $s_1, \ldots, s_N$ of dividend shares, such that $\sum_{i=1}^{N} s_i = 1$. As explained in Section 3.4, firm i pays dividend $s_i D_{i,t}$ at time t, $s_{i+1} D_{i,t+1}$ at time $t+1$, etc. At time $t+1$, firm i becomes firm $i+1$. We choose two simple functional forms for s_i to illustrate the implications of our model for a cross-section of long-lived assets.

The first functional form assumes that the rate of growth in the shares falls linearly in i:

$$\frac{s_{i+1}}{s_i} = g_0 \left(1 - \frac{i-1}{N/2-1}\right) \tag{21}$$

for $i = 1, \ldots, N-1$. The value of s_1 is set so the weights sum to 1. For the first 100 quarters, the firm's share in the aggregate dividend process grows, each year at a slower rate. After 100 quarters, the share in the aggregate dividend process declines. Eventually, the firm's share in the aggregate dividend process becomes negligible. At this point the share begins to grow again, as described in Section 3.4. A possible interpretation is that the capital invested in the dying firm is used to fund a new, growing firm.

As a robustness check, we examine a second functional form. The second functional form assumes that the rate of growth in the shares is a positive constant g_1 for the first half of the firm's life-cycle, and $-g_1$ for the second half of the firm's life-cycle. In other words,

$$\frac{s_{i+1}}{s_i} = g_1 \tag{22}$$

for $i = 1, \ldots, N/2-1$, 0 for $i = N/2$, and

$$\frac{s_{i+1}}{s_i} = -g_1 \tag{23}$$

for $i = N/2 + 1, \ldots, N - 1$. Once again, the value of s_1 is set so that the weights sum to 1. We calibrate each specification for firm shares using data on the cross-section of cash flows and price-dividend ratios. Here and throughout, we set the number of firms to be 200, implying a 200-quarter, or equivalently, a 50-year life cycle for a firm.

We first describe the calibration for the model with linearly declining growth. The parameter that controls the shrinking growth rate, g_0, is set so that the cross-section of dividend growth rates matches the cross-section of firms in the sample. Because data on earnings and cash flows are not available prior to 1952, we construct the cross-section for data from 1952 to 2002.[8] The parameter g_0 is set to be 0.75/4, for a per-annum value of 0.75. The specification (21) then implies that firms start their life cycle with an annual growth rate of 75%, growth declines to 0% after 25 yeas, and reaches -75% after another 25 years to complete one cycle. The top panel of Figure 5 plots the implied expected growth rate of dividends for firms in the model, as well as the cross-section of average growth rates in earnings, dividends, and cash flows. Because the firms in our model have no debt, the dividends in our model may be better analogues to earnings and cash flows in the data, rather than dividends themselves. As the top panel of Figure 5 shows, the linear model with g_0 set at 0.75/4 does a reasonable job of fitting the cross-section of growth rates in the data.

The second panel of Figure 5 shows the distribution of firm price-dividend ratios in the model, and price ratios in the data. The linearly declining growth model for shares produces more high price-dividend ratio firms than there are in the data. These firms have high price-dividend ratios because they have extremely low current dividends. It is possible to construct a model that fits the distribution more closely by imposing a lower bound on the dividend share, at the cost of a somewhat more complicated functional form for s_i. This modification is pursued in Appendix A.

Figure 6 shows corresponding results for the model with constant growth in shares. When calibrated with an annual growth rate of 20% ($g_1 = .20/4 = .05$), the model with constant growth produces a reasonable fit to the cross-section of dividend, earnings, and cash-flow

[8] Adrian and Franzoni (2002) Ang and Chen (2003), and Campbell and Vuolteenaho (2003) show that value stocks have higher betas in the pre-war period, so the CAPM performs better. By matching the cross-section to the postwar data, we choose a harder target. We also assume that agents observe the parameters in the economy. Lewellen and Shanken (2002) show that introducing learning into a traditional model can help in understanding value premia.

growth in the data, despite its simplicity. A drawback of the constant growth model is the frequency of very high or very low dividend growth rates. This aspect of the distribution is better matched with the linear-growth model. However, the constant-growth model produces a superior fit tho the cross-section of price ratios, as shown in the second panel of Figure 6.

We conclude that the linearly declining growth and the constant growth model produce reasonable cross-sections of dividend growth and price-dividend ratios. As we will show, our results for portfolio statistics are very similar, regardless of which model we use. This suggests that our results will be robust to a range of cross-sectional distributions for firm dividend growth.

Portfolio returns

At the start of each year in the simulation, we sort firms according to their price-dividend ratio into deciles. We then form equally-weighted portfolios of the firms in each decile. As firms move through their life-cycle, they slowly shift (on average) from the growth category to the value category, and then revert back eventually to the growth category. Of course, this process is not deterministic, just as it is not deterministic in the data. Shocks have different impacts on the price-dividend ratio, the variable that determines the appropriate decile for the firm.

Having sorted the firms into deciles at the beginning of each "year", we compute statistical tests on returns over the year. Table 8 shows results for the linearly declining share growth specification; Table 9 shows results for constant share growth specification. Because the results are quantitatively similar, we discuss only Table 8.

The first panel of Table 8 shows the expected excess return, the standard deviation, and the Sharpe ratio for each portfolio. These simulation results should be compared to the numbers in Table 1, which shows corresponding results for the data. The expected return on extreme growth stocks is 4.44% per annum, while for extreme value stocks, it is 9.65% per annum.[9] A similar spread is generated in the data: the lowest book-to-market stocks have a

[9]Here and throughout this section, we compare the statistics on annual returns in the model to statistics on monthly returns in the data. The monthly data statistics are annualized as described in Section 2. We choose this approach because it corresponds most closely to the approach taken in the empirical literature on the value premium. Data results for annual returns are very similar to those in Tables 1-3 (except for standard errors).

premium of 5.67 percent, while the highest have a premium of 10.55%. The model generates reasonable volatilities; between 19% and 17%. The volatilities for book-to-market sorted portfolios vary between 18% and 15% in the data. The model predicts that the volatility decreases as one moves from the growth to the value portfolio, so growth stocks are actually less risky than value stocks, despite the lower return. This effect is also found in the data. The model predicts that the Sharpe ratio for the extreme growth stocks is 0.23, and for the extreme value stocks is 0.57, and that the Sharpe ratio increases monotonically in the decile number. In comparison, the low book-to-market portfolio has a Sharpe ratio of 0.32 in the data, the Sharpe ratio increases steadily as a function of decile, and the high book-to-market portfolio is 0.57. To summarize, in the model, value stocks have high expected returns, low volatility, and high Sharpe ratios, just as in the data. The magnitude of the difference between value and growth is comparable to that in the data.

The second panel of Table 8 shows αs and βs relative to the CAPM. Annual excess portfolio returns are regressed on excess returns on the aggregate market (the claim to all future dividends). α, β and the R^2 are reported for each decile. The second panel shows that the model can replicate the classic result of Fama and French (1992): value portfolios have positive αs relative to the CAPM, while growth portfolios have negative αs. Moreover, value portfolios tend to have lower βs than growth portfolios. Our model predicts αs that rise from -2.86 for the extreme growth portfolio, to 3.30 for the extreme value portfolio. In the data, the lowest book-to-market portfolio has an α of -1.66, while the highest book-to-market portfolio has an α of 3.97. Thus the model generates αs of the correct magnitude, as well as a sizable spread between value and growth. Moreover, αs in the model are asymmetric: growth αs are smaller in absolute value than value α, as in the data.

The third panel of Table 8 runs regressions of portfolio returns on the market and on a high-minus-low factor (HML), generated by the return on a short position in the extreme growth portfolio and a long position in the extreme value portfolio. The purpose is to see whether the model analogue to the high-minus-low factor of Fama and French (1992) describes the cross-section of returns in the model, as it does in the data. We find that the αs nearly disappear when we add HML to the regression; they are two orders of magnitude smaller than the αs relative to the CAPM. Not surprisingly, the loading on the HML factor increases in the decile; growth portfolios have a negative loading on the HML factor, while value portfolios have a positive loading.

The final panel of Table 8 compares the unconditional correlations of the portfolios with the stochastic discount factor, the shock to dividends, the shock to expected dividend growth orthogonal to dividend growth, and the shock to the variable controlling the Sharpe ratio. The correlation with the stochastic discount factor is lower, the greater the decile. This is consistent with the result that the Sharpe ratio increases in the decile; it follows from the unconditional Euler equation that the lower the correlation with the stochastic discount factor, the higher the Sharpe ratio.

The final panel also examines differential unconditional correlations with the shocks. Value stocks have the highest correlations with shocks to dividend growth (correlations range from .11 for the extreme growth portfolio to .26 for the extreme value portfolio). In contrast, growth stocks have a higher correlation with the independent component of the shock to expected dividend growth (correlations range from .21 for the extreme growth portfolio to .17 for the extreme value portfolio). A shock to expected dividend growth raises the valuation of all portfolios (as in the present value models of Campbell and Shiller (1988) and Vuolteenaho (2002)), and returns on growth stocks, which pay dividends in the distant future. Finally, all portfolios are negatively correlated with shocks to the Sharpe ratio variable x_t. A positive shock to x_t raises expected returns, and thus lowers prices and realized returns. Growth stocks are more negatively correlated with x_t than value stocks (correlations range from -.41 for the extreme growth portfolio to -.29 for the extreme value portfolio). This is because of duration; growth stocks pay dividends further in the future, and thus are impacted more by a shock to expected returns.

Finally, the model can explain part of the momentum puzzle described in Jegadeesh and Titman (1993) and Jegadeesh and Titman (2001). Table 10 shows that when returns are sorted based on past one-year returns, the "winner" portfolio exhibit expected returns of 11.2% per annum, while the "loser" portfolios exhibits expected returns of 7.7% per annum. The spread between the winner and loser portfolio is not as large as in the data; for the one-year momentum strategy Jegadeesh and Titman (1993) report a spread of 8.2%, rather than the 3.5% we find. However, it is economically significant. These results are related to findings of Johnson (2002). Johnson shows that time-varying expected dividend growth can lead to momentum effects in a model with constant discount rates and a single long-lived asset. Here we show that the momentum effects remain when discount rates are time-varying, and when there are multiple long-lived assets. Because we have a market portfolio, we can evaluate

whether these momentum effects appear even after adjusting for risk using the CAPM. In fact, Table 10 shows that the α relative to the CAPM on the winner-minus loser portfolio is 3.6%. This is smaller than the 12% reported by Jegadeesh and Titman (2001) (for the six-month momentum strategy), but nonetheless is economically significant.

This section has shown that the model replicates the principle aspects of the value premium in the data: value stocks have high expected returns, low volatility, high Sharpe ratios, high alphas, and low betas, relative to growth stocks. Value and growth stocks also differ in their correlations with the underlying shocks to the economy. As the next section shows, these differences in correlations are the root of the model's ability to explain the value premium.

5 Model Intuition

What explains the model's ability to capture the value premium? As hinted in the previous section, the value premium arises from the differential correlations of value and growth portfolios on underlying shocks.

Figure 8 presents another look at this issue. Rather than correlations, Figure 8 contains βs from unconditional regressions of portfolio returns on the three shocks, and the R^2 from the unconditional regressions. Three results are apparent. First, β_d is positive and increasing as the portfolios go from growth to value, second β_z is positive and decreasing as portfolios go from growth to value, and β_x is negative and decreasing as firms go from growth to value. The R^2s follow the same direction as the absolute value of the βs. Shocks to dividends account for an increasing amount of the total variance, shocks to z represent a decreasing amount of the variance, and shocks to x represent an increasing amount of the variance as portfolios go from growth to value. It should not be surprising that the R^2 fail to sum to 1. The three shocks account for 100% of the *conditional* variance of returns. Figure 8, however, plots the results from unconditional regressions. In this section, we first discuss the intuition for the differential loadings on the shocks. We then connect these loadings to risk premia.

The results in Figure 8 arise from the different loadings of zero-coupon equity on the underlying shocks. As (15) shows, the loading of returns on the shocks to x and z is determined by B_x and B_z (the loading on the shocks to d is always the same conditionally). For the parameter values we estimate, B_z is positive and increasing in the maturity n while B_x

is negative, and first decreases, and then, at high maturities, increases in n (Figure 1). In the case of B_z the increase occurs because long-horizon dividend growth matters more for assets that pay dividends in the future; for B_x, the pattern is accounted for by the duration effect, and because the risk price on expected dividend growth is negative. The intuition for these patterns is explained further in Section 3 and 4.3.

The loading on z_t for zeros is echoed for portfolios, because portfolios are formed from firms that are weighted averages of the underlying zeros. For high price-dividend firms (growth firms), the share of aggregate dividends is increasing. These firms place greater weight on higher-maturity equity than on lower-maturity equity. Thus their returns load more on B_z than the returns of value firms, which place more weight on lower maturity equity. The loading in x_t is a bit different for portfolios than for zero-coupon equity. There is no evidence of the non-monotonicity seen in Figure 1. This is because, all else equal, equity that pays further in the future is worth less, medium horizon equity may have a greater weight than long-horizon equity, even for growth firms. The non-monotonicity in B_x occurs sufficiently far out in the equity term structure that it does not appear in the loadings on portfolios.

The loadings of portfolios on various shocks present an intriguing link with the empirical results of Campbell and Vuolteenaho (2003). Using the vector auto-regression methodology of Campbell (1991), Campbell and Vuolteenaho decompose unexpected market returns into changes in expectations of future discount rates and changes in expectations of future dividend growth rates. Changes in expected discount rates are computed using the vector auto-regression; changes in expected growth rates are the residual: the variation in market returns not explained by variation in discount rates. When the market return is decomposed in this way, they show that βs with respect to changes in expected discount rates decrease in magnitude as portfolios go from growth to value, while βs with respect to cash flows increase in magnitude as portfolios go from growth to value.

While not precisely analogous, shocks to x_t are similar in spirit to news about discount rates in the VAR framework of Campbell and Vuolteenaho (2003). It is therefore encouraging that our model produces βs with respect to shocks to x_t that increase in magnitude as portfolios go from growth to value. The analogue to the dividend news term is less clear in our model. Campbell and Vuolteenaho compute this as a residual, though in our model Figure 8 shows that the residual variance is not accounted for by shocks to current or expected

future dividends. More of the residual is accounted for by Δd_{t+1} than by shocks to z_{t+1}; thus it is also encouraging that value stocks load more on Δd_{t+1} than growth stocks.

Figure 8 shows that value and growth portfolios have different loadings on the underlying shocks in the economy. How this translates into risk premia depends on the prices of risk of shocks of these shocks. Equation (16) provides an illustration of how conditional risk premia on zero-coupon equity vary based on loadings on different shocks. As discussed in Section 4.1 we estimate that shocks to expected dividend growth z_t are negatively correlated with shocks to realized dividend growth, based on dividend and consumption data. This empirical result implies that expected dividend growth has a negative risk price; because it is negatively correlated with shocks to realized dividend growth it serves as a hedge and reduces risk premia.

We assume that shocks to x_t carry a zero risk price; shocks to x_t are assumed to be uncorrelated with shocks to realized dividends or expected dividend growth. This assumption represents a departure from the models of Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004), where shocks to the price of risk are, conditionally, perfectly negatively correlated with shocks to aggregate dividends. What role does this assumption play in our analysis?

To answer this question, consider the equation for the conditional risk premium for equity that matures next period, versus the conditional risk premium for equity that matures two periods from now. By (16), equity that matures next period has a risk premium of

$$\log E_t[R_{1,t+1}/R^f] = \|\sigma_d\| x_t.$$

Equity that matures two periods from now has a risk premium of

$$\log E_t[R_{2,t+1}/R^f] = \left(1 - \rho_{dx}\|\sigma_x\| + \rho_{dz}\frac{\|\sigma_z\|}{\|\sigma_d\|}\right)\|\sigma_d\| x_t$$

where

$$\rho_{dx} = \frac{\sigma_d \sigma_x'}{\|\sigma_d\|\|\sigma_x\|}$$

the conditional correlation between Δd_{t+1} and x_{t+1}, and

$$\rho_{dz} = \frac{\sigma_d \sigma_z'}{\|\sigma_d\|\|\sigma_z\|}$$

the conditional correlation between Δd_{t+1} and z_{t+1}. The risk premium on equity that matures next period is equal to the quantity of risk – the standard deviation of dividends – multiplied

by the price of risk x_t. For equity maturing two periods from now, there is also the risk due to changes in x_t and changes in z_t. The latter effect will be relatively small because σ_z is a small fraction of σ_d. Whether long-horizon equity has a lower risk premium than short-horizon equity (as is necessary to explain the value premium) is due in large part to the sign of the correlation of dividend growth with x_t. In particular, $\rho_{dx} < 0$ leads to relatively higher premia for long-horizon equity, while $\rho_{dx} > 0$ leads to relatively lower premia for long-horizon equity.

We make this statement precise by solving the model under three different possibilities for ρ_{dx}. Figure 7 plots risk premia on zero-coupon equity when $\rho_{dx} = 0$ (our base case), $\rho_{dx} < 0$, and $\rho_{dx} > 0$. For $\rho_{dx} = 0$, Panel A shows that risk premia decrease in maturity, as long as $x_t > 0$ (as it is most of the time). The reason for this decrease is the negative correlation between Δd_{t+1} and z_{t+1}. In contrast, for $\rho_{dx} < 0$, Panel B shows that risk premia generally increase in the maturity. Long-horizon equity (i.e. growth stocks) have greater risk premia than short-horizon equity. This occurs even though ρ_{dz} is negative, as in Panel A. Even a modest correlation of -.5 between dividends and the price-of-risk overrides the effect of ρ_{dz}. The case of $\rho_{dx} < 0$ is of special interest because it corresponds to the correlation between the price of risk and aggregate dividends in external habit models. In the models of Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004), shocks to the aggregate dividend (which is identified with consumption) increase surplus consumption, and therefore lower the amount of return investors demand for taking on risk. Indeed, in a term structure context, Wachter (2003) shows that the model of Campbell and Cochrane (1999) implies that long-horizon assets exhibit greater risk premia than short horizon assets for exactly this reason. Long-horizon assets load more negatively on the shock to discount rates; if discount rates are negatively correlated with consumption (or dividends) then long-horizon assets will command greater risk premia.

An alternative to both the external habit models, and to our approach is to set the correlation between d_{t+1} and x_{t+1} to be positive. This case is illustrated in Panel C. Under this assumption, risk premia fall more dramatically in the maturity than under the assumption that d_{t+1} and x_{t+1} are uncorrelated and the premium for short-horizon equity is greater.

These results at first suggest that a model that seeks to explain the value premium should set $\rho_{dx} > 0$, rather than $\rho_{dx} = 0$ as we assume. However, we argue that this is not a good choice. One reason is the economic interpretation. For $\rho_{dx} = 0$, shocks to the price of

risk are preference shocks uncorrelated with cash flows in the economy. It is not hard to imagine that prices might be influenced by such shocks. In contrast, $\rho_{dx} > 0$ makes less sense economically. This says, in effect, that investors demand a greater risk premia on equities when they are better off, not when they are worse off. This assumption seems unappealing.

Also important are the aggregate time-series implications of the sign of ρ_{dx}. We are able to calibrate our model to match the time series of aggregate stock returns, as well the cross-section of value and growth portfolios. We accomplished this because our model produced reasonable risk premia in the aggregate. For $\rho_{dx} > 0$, this may not be the case. Figure 7 shows that the greater is ρ_{dx}, the lower are risk premia in the economy, for all but the shortest-maturity equity. This makes sense; as an asset that pays cash flows in the future, equity must load negatively on x_t. If investors view x_t-risk as a hedge ($\rho_{dx} > 0$), this makes equity less risky. On the other hand, if x_t moves in the same direction as dividends ($\rho_{dx} < 0$), equity becomes more risky. Explaining the level of the equity premium is therefore easiest when $\rho_{dx} < 0$ and hardest when $\rho_{dx} > 0$. The assumption that $\rho_{dx} < 0$ is part of what enables Campbell and Cochrane (1999) and Menzly, Santos, and Veronesi (2004) to explain the high variance and the high premium commanded by stock returns, with comparatively little variance in fundamentals.

Faced with this tension between the time series and the cross section, we choose the economically reasonable assumption of $\rho_{dx} = 0$. This section has shown that ρ_{dx}, in combination with the duration effect and the correlation between current and future dividend growth, makes long-horizon equity less risky than short-horizon equity. It creates a large premium on value stocks, while at the same time limiting their risk. We hope that future work will reveal micro-economic foundations that determine this important parameter.

6 Conclusion

This paper has proposed a parsimonious model of the stochastic discount factor that accounts for both the aggregate time series behavior of the stock market, and for the relative risk and return of value and growth stocks. At the root of the model is a dividend process calibrated to match the dividend process in the data, and a stochastic discount factor with a single factor proxying for investors' time-varying preference for risk. This factor, called x_t, is analogous to the inverse of surplus consumption in the external habit models of Campbell

and Cochrane (1999) and Menzly, Santos, and Veronesi (2004). When x_t is high, investors demand a greater return for risk. Time-varying preferences for risk allows the model to capture excess volatility and return predictability found in the data. Our specification for x_t allows for interpretable closed-forms solutions for asset prices and risk premia.

A key difference between our model and external habit models is that x_t does not arise from fluctuations in the aggregate dividend. This may seem like a small detail but it is key to the model's ability to explain how value stocks can have both higher returns and less risk than growth stocks. In our model, growth and value stocks differ based on their cash flows. Growth stocks have more of their cash flows in the future. They are high-duration assets, and thus their returns covary more with the price of risk x_t. We have shown that for growth stocks to have relatively low returns, it must be the case that investors do not fear shocks to x_t. This only occurs if the conditional correlation of the price of risk with dividend growth is zero or positive. We assume the correlation is zero. In contrast, external habit models assume a correlation of negative one. Shocks to the price of risk are feared as much, if not more, than shocks to cash flows.

Our proposed resolution to the value puzzle is risk-based. Value stocks, as short-horizon equity, vary more with fluctuations in cash flows – the fluctuations that investors fear the most. Growth stocks, as long-horizon equity, vary more with fluctuations in discount rates, which are independent of cash flows and which investors do not fear. We have shown that the value puzzle is only puzzling from the point of view of models where risk preferences, and thus discount rates, are constant. When risk preferences vary, and vary independently from consumption and dividends, we have shown that the value premium is not a puzzle after all.

7 Appendix

A Including a Lower Bound in the Share Specification

Including a lower bound in the share specifications (21) and (22)–(23) allows for a superior fit to the cross-section of price ratios. For a model with shares s_i, we can define a new share process $\hat{s}_i$ such that

$$\hat{s}_i = \underline{\mathrm{s}} + s_i(1 - N\underline{\mathrm{s}})$$

where N is the number of firms, and $0 \leq \underline{\mathrm{s}} < 1/N$. Then the sequence $\hat{s}_1, \ldots, \hat{s}_N$ is positive, sums to 1, and thus forms a well-defined share process. Moreover, $\underline{\mathrm{s}}$ will be a lower bound on the share of the aggregate dividend.

Figure 9 shows the results for linearly declining growth, when the share process has a lower bound of $\underline{\mathrm{s}} = 0.1\%$. This change implies fewer firms with very high growth rates, and fewer firms with higher price dividend ratios. The distribution of dividend growth and the distribution of price ratios provide a closer match to the data, as shown in Figure 9. Introducing a lower bound to the constant growth case has a similar effect, as shown in Figure 10. Tables 11 and 12 show that the behavior of the cross-section of returns is similar to the cases when there are no lower bound. Qualitatively, the effects are similar to the case where there is no lower bound. The α for the extreme value portfolio is 2.6% for the linear growth case, and 2.7% for the constant growth case, while the α for the extreme growth stocks is -2.2% for the linear growth case and -2.0% for the constant growth case.

B Convergence of the Price-Dividend Ratio

Because x_t and z_t can take on both positive and negative values, a necessary (but not sufficient) condition for (17) to converge for all values of x_t and z_t is that $B_x(n)$ and $B_z(n)$ approach finite values as $n \to \infty$. B_z converges if and only if:

$$|\phi_z| < 1 \tag{24}$$

Let

$$\lambda = \sigma_d / \|\sigma_d\|$$

Assuming (24) holds, B_x converges if and only if

$$|\phi_x - \sigma_x \lambda| < 1 \tag{25}$$

Given (24)

$$\lim_{n\to\infty} B_z(n) = \frac{1}{1-\phi_z} \equiv \bar{B}_z$$

Define $\bar{B}_x$ to be the solution to

$$\bar{B}_x = \bar{B}_x(\phi_x - \sigma_x\lambda) - (\sigma_d + \frac{\sigma_z}{1-\phi_z})\lambda$$

Then

$$\bar{B}_x = -\frac{(\sigma_d + \sigma_z/(1-\phi_z))\,\lambda}{1-(\phi_x - \sigma_x\lambda)}$$

Given (24) and (25), it follows that

$$\lim_{n\to\infty} B_x(n) = \bar{B}_x$$

and

$$\lim_{n\to\infty} V_n = \sigma_d + \frac{\sigma_z}{1-\phi_z} + \bar{B}_x\sigma_x \equiv \bar{V}$$

Finally, let

$$\bar{A} = -r + g + \bar{B}_x(1-\phi_x)\bar{x} + \frac{1}{2}\bar{V}\bar{V}$$

It follows from the recursion for A_n that for N sufficiently large

$$A(n) \approx \bar{A}n + \text{constant}$$

for $n \geq N$, and therefore

$$\sum_{n=N}^{\infty} \exp\left\{A(n) + B_z(n)z_t + B_x(n)x_t\right\} \approx \exp\left\{\text{constant} + \bar{B}_z z_t + \bar{B}_x x_t\right\} \sum_{n=N}^{\infty} \exp\left\{\bar{A}n\right\}.$$

It follows that necessary and sufficient conditions for convergence are (24), (25), and

$$-r + g + \bar{B}_x(1-\phi_x)\bar{x} + \frac{1}{2}\bar{V}\bar{V} < 0 \tag{26}$$

References

Adrian, Tobias, and Francesco Franzoni, 2002, Learning about beta: An explanation of the value premium, unpublished paper, MIT.

Ang, Andrew, and Joseph Chen, 2003, CAPM over the long run: 1926-2001, Unpublished paper, Columbia University and University of Southern California.

Ang, Andrew, and Jun Liu, 2003, How to discount cash flows with time-varying expected returns, forthcoming, *Journal of Finance*.

Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745–787.

Bakshi, and Chen, 1996, Inflation, asset prices, and the term structure of interest rates in a monetary economy, *Review of Financial Studies* 9, 241–275.

Ball, R., 1978, Anomolies in relationships between securities' yields and yield-surrogates, *Journal of Financial Economics* 6, 103–126.

Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2003, Interpreting risk premia across size, value, and industry portfolios, Unpublished paper, Duke University, Indiana University, and University of Michigan.

Bansal, Ravi, and Amir Yaron, 2003, Risks for the Long-Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance,* forthcoming.

Basu, Sanjoy, 1977, The Investment Performance of Common Stocks in Relation to Their Price to Earnings Ratios: A Test of the Efficient Market Hypothesis, *Journal of Finance* 32, 663–82.

Basu, Sanjoy, 1983, The relationship between earnings yield, market value, and return for NYSE common stocks: Further Evidence, *Journal of Financial Economics* 12, 129–156.

Bekaert, Geert, Eric Engstrom, and Steve Grenadier, 2004, Stock and bond returns with moody investors, Working paper, Columbia University, Stanford University, and University of Michigan.

Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2003, Estimation and test of a simple model of intertemporal asset pricing, forthcoming, *Journal of Finance*.

Brennan, Michael J., and Yihong Xia, 2003, Risk and valuation under an intertemporal capital asset pricing model, forthcoming, *Journal of Business*.

Campbell, John Y., 1991, A Variance Decomposition for Stock Returns, *Economic Journal* 101, 157–179.

Campbell, John Y., 1999, Asset prices, consumption, and the business cycle, in J.B. Taylor, and M. Woodford, eds.: *Handbook of Macroeconomics, Volume I* (North-Holland, Amsterdam).

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, 205–251.

Campbell, John Y., and Jianping Mei, 1993, Where do betas come from? Asset pricing dynamics and the sources of systematic risk, *Review of Financial Studies* 6, 567–592.

Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2003, Growth or glamour?, Unpublished paper, Harvard University and Northwestern University.

Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 58, 495–514.

Campbell, John Y., and Tuomo Vuolteenaho, 2003, Bad beta, good beta, Working Paper, Harvard University.

Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2002, Does risk or mispricing explain the cross-section of stock prices?, Unpublished paper, Harvard University and Northwestern University.

Cornell, Bradford, 1999, Risk, duration, and capital budgeting: New evidence on some old questions, *Journal of Business* 72, 183–200.

Dai, Qiang, and Kenneth Singleton, 2003, Term structure dynamics in theory and reality, *Review of Financial Studies* 16, 631–678.

Dechow, Patricia M., Richard G. Sloan, and Mark T. Soliman, 2002, Implied equity duration: A new measure of equity risk, Unpublished paper, University of Michigan.

Duffee, Gregory R., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 369–443.

Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379–406.

Fama, Eugene F., and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 29, 23–49.

Fama, Eugene F., and Kenneth R. French, 1992, The Cross-Section of Expected Returns, *Journal of Finance* 47, 427–465.

Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium Cross Section of Returns, *Journal of Political Economy* 111, 693–732.

Graham, Benjamin, and David L. Dodd, 1934, *Security Analysis*. (McGraw Hill New York) 1 edn.

Jagannathan, Ravi, and Zhenyu Wang, 1996, The Conditional CAPM and the Cross-Section of Expected Returns, *Journal of Finance* 51, 3–54.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jegadeesh, Narasimhan, and Sheridan Titman, 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *Journal of Finance* 56, 699–720.

Johnson, Timothy C., 2002, Rational momentum effects, *Journal of Finance* 57, 585–608.

Lettau, Martin, and Sydney C. 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 C. Ludvigson, 2002, Expected Returns and Expected Dividend Growth, NBER Working Paper No. 9605.

Lewellen, Jonathan, and Jay Shanken, 2002, Learning, asset-pricing tests, and market efficiency, *Journal of Finance* 57, 1113–1145.

Liew, Jimmy, and Maria Vassalou, 2000, Can Book-to-Market, Size and Momentum Be Risk Factors that Predict Economic Growth?, *Journal of Financial Economics* 57, 221–245.

Lintner, J., 1965, Security Prices, Risk and Maximal Gains from Diversification, *Journal of Finance* 20, 587–615.

Lustig, Hanno, and Stijn Van Nieuwerburgh, 2002, Housing Collateral, Consumption Insurance and Risk Premia, Unpublished manuscript, University of Chicago.

Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding Predictability, *Journal of Poitical Economy* 112, 1–47.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2002, Housing, Consumption, and Asset Pricing, Unpublished paper, UCLA.

Rosenberg, Barr, Kenneth Reid, and Ronald Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management* 11, 9–17.

Santos, Tano, and Pietro Veronesi, 2004, Conditional betas, Unpublished paper, Columbia University and University of Chicago.

Sharpe, W., 1964, Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance* 19, 425–444.

Vassalou, Maria, 2003, News related to future GDP growth as a risk factor in equity returns, *Journal of Financial Economics* 68, 47–73.

Vuolteenaho, Tuomo, 2002, What drives firm-level stock returns?, *Journal of Finance* 57, 233–264.

Wachter, Jessica A., 2003, A consumption-based model of the term structure of interest rates, Working paper, University of Pennsylvania.

Wilson, Mungo, 2003, A model of demand for dividend strips, Working Paper, Harvard University.

Yogo, Motohiro, 2003, A Consumption-Based Explanation of the Cross Section of Expected Stock Returns, Unpublished paper, Harvard University.

Zhang, Lu, and Ralitsa Petkova, 2002, Is value riskier than growth?, Unpublished paper, University of Rochester.

Table 1: Growth vs. Value Portfolios

	G				Growth to Value					V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
					Mean Excess Return (% per year)						
E/P	4.71	5.02	6.97	7.04	7.00	9.18	9.94	11.18	11.68	12.95	8.25
C/P	5.05	6.07	6.49	6.73	8.48	7.72	8.85	9.18	11.47	11.81	6.77
D/P	7.35	6.41	7.28	7.41	6.49	7.60	7.73	9.49	8.84	7.45	0.10
B/M	5.67	6.55	6.98	6.51	8.00	8.33	8.27	10.08	9.98	10.55	4.88
					Standard Error of Mean						
E/P	0.78	0.64	0.62	0.59	0.62	0.61	0.60	0.61	0.65	0.73	0.62
C/P	0.76	0.64	0.61	0.63	0.62	0.60	0.60	0.60	0.61	0.69	0.59
D/P	0.78	0.69	0.66	0.64	0.62	0.60	0.59	0.58	0.56	0.56	0.69
B/M	0.71	0.64	0.64	0.62	0.59	0.59	0.59	0.61	0.63	0.74	0.61
					Standard Deviation of Excess Return (% per year)						
E/P	19.35	15.93	15.49	14.78	15.43	15.04	14.87	15.29	16.11	18.11	15.40
C/P	18.99	15.95	15.24	15.75	15.43	14.95	14.96	14.98	15.14	17.24	14.57
D/P	19.36	17.11	16.31	15.85	15.43	15.00	14.58	14.37	13.93	13.83	17.08
B/M	17.77	15.89	15.82	15.42	14.65	14.73	14.74	15.11	15.71	18.46	15.15
					Sharpe Ratio						
E/P	0.24	0.32	0.45	0.48	0.45	0.61	0.67	0.73	0.73	0.72	0.54
C/P	0.27	0.38	0.43	0.43	0.55	0.52	0.59	0.61	0.76	0.69	0.46
D/P	0.38	0.37	0.45	0.47	0.42	0.51	0.53	0.66	0.63	0.54	0.01
B/M	0.32	0.41	0.44	0.42	0.55	0.57	0.56	0.67	0.64	0.57	0.32

Notes: Summary statistics for 10 growth to value portfolios. "D/P" are deciles computed from sorting firms based on their dividend yields, "E/P" are sorts based on the earnings yields, "C/P" are sorts based on the ratio of cash flows to price and "B/M" are sorts on the book-to-market ratio. All data are from Ken French's website. The data are monthly and span the 1952 to 2002 period.

Table 2: Correlation of Value and Growth Portfolios

	D/P	E/P	C/P	B/M
Top Growth Decile				
E/P	1.00	0.98	0.93	0.96
C/P	0.98	1.00	0.93	0.97
D/P	0.93	0.93	1.00	0.94
B/M	0.96	0.97	0.94	1.00
Top Value Decile				
E/P	1.00	0.94	0.76	0.85
C/P	0.94	1.00	0.74	0.85
D/P	0.76	0.74	1.00	0.75
B/M	0.85	0.85	0.75	1.00

Notes: Correlations among of growth and value portfolios formed from sorting on different variables. "D/P" are deciles computed from sorting firms based on their dividend yields, "E/P" are sorts based on the earnings yields, "C/P" are sorts based on the ratio of cash flows to price and "B/M" are sorts on the book-to-market ratio. All data are from Ken French's website. The data are monthly and span the 1952 to 2002 period.

Table 3: Growth vs. Value Portfolios

CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$

	G				Growth to Value					V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
					α_i (% per year)						
E/P	-3.09	-1.62	0.69	0.95	0.74	3.25	4.08	5.33	5.60	6.22	9.31
C/P	-2.70	-0.54	0.19	0.24	2.33	1.79	3.01	3.46	5.75	5.34	8.04
D/P	-0.58	-0.73	0.62	0.98	0.44	1.77	2.03	4.11	3.96	3.44	4.01
B/M	-1.66	-0.17	0.33	0.22	2.12	2.37	2.59	4.30	4.05	3.97	5.63
					Standard Error of α_i						
E/P	1.12	0.74	0.86	0.75	0.86	0.95	0.95	1.07	1.18	1.38	2.14
C/P	1.03	0.78	0.76	0.80	0.94	0.93	0.98	1.06	1.11	1.28	2.01
D/P	1.03	0.80	0.88	0.88	1.00	1.00	0.96	1.07	1.19	1.47	2.05
B/M	0.90	0.65	0.69	0.84	0.86	0.83	1.01	1.07	1.15	1.53	2.12
					β_i						
E/P	1.18	1.01	0.95	0.92	0.95	0.90	0.89	0.89	0.92	1.02	-0.16
C/P	1.17	1.00	0.95	0.98	0.93	0.90	0.89	0.87	0.87	0.98	-0.19
D/P	1.20	1.08	1.01	0.97	0.92	0.88	0.86	0.82	0.74	0.61	-0.59
B/M	1.11	1.02	1.01	0.95	0.89	0.90	0.86	0.87	0.90	1.00	-0.11
					Standard Error of β_i						
E/P	0.02	0.01	0.02	0.01	0.02	0.02	0.02	0.02	0.02	0.03	0.04
C/P	0.02	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.04
D/P	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.04
B/M	0.02	0.01	0.01	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.04
					R^2						
E/P	0.83	0.89	0.84	0.87	0.84	0.80	0.80	0.75	0.73	0.71	0.02
C/P	0.85	0.88	0.87	0.87	0.81	0.80	0.78	0.75	0.73	0.72	0.04
D/P	0.86	0.89	0.85	0.84	0.79	0.77	0.78	0.72	0.63	0.43	0.27
B/M	0.87	0.91	0.90	0.85	0.83	0.84	0.76	0.75	0.73	0.65	0.01

Notes: Estimates from CAPM regressions of 10 growth to value deciles. "D/P" are deciles computed from sorting firms based on their dividend yields, "E/P" are sorts based on the earnings yields, "C/P" are sorts based on the ratio of cash flows to price and "B/M" are sorts on the book-to-market ratio. All data are from Ken French's website. The data are monthly and span the 1952 to 2002 period.

Table 4: Parameters of the Model

$$
\begin{aligned}
\Delta d_{t+1} &= g + z_t + \sigma_d \epsilon_{t+1} \\
z_{t+1} &= \phi_z z_t + \sigma_z \epsilon_{t+1} \\
x_{t+1} &= (1-\phi_x)\bar{x} + \phi_x x_t + \sigma_x \epsilon_{t+1} \\
M_{t+1} &= \exp\left\{ -r^f - \frac{1}{2} x_t^2 - x_t \frac{\sigma_d}{\|\sigma_d\|} \epsilon_{t+1} \right\}
\end{aligned}
$$

Variable	Value
g	2.28%
r_f	1.93%
$\bar{x}$	0.625
ϕ_z	0.91
ϕ_x	0.87
$\|\sigma_d\|$	0.145
$\|\sigma_z\|$	0.0016
$\|\sigma_x\|$	0.24
Correlation of Δd and z shocks	-0.83
Correlation of Δd and x shocks	0
Correlation of z and x shocks	0

Implied volatility parameters

Variable	Values
σ_d	[0.0724, 0, 0]
σ_z	[-0.0013, 0.0009, 0]
σ_z	[0, 0, 0.12]

Notes: The model parameters are calibrated to the Campbell's (1999) century-long annual sample starting in 1889 and updated through 2002. The model is simulated at a quarterly frequency. The unconditional mean of dividend growth g, the riskfree rate r_f, the persistence variables ϕ_x and ϕ_z, and the conditional standard deviations $\|\sigma_d\|$, $\|\sigma_d\|$, $\|\sigma_d\|$, are in annual terms (i.e. $4g$, ϕ_x^4, $2\|\sigma_d\|$). Parameters g, r_f, and $\|\sigma_d\|$ are set to match their data counterparts. Parameters ϕ_z and the correlation between shocks to z and shocks to Δd are set to match their data counterparts, assuming that the the conditional mean of dividend growth is the log of the consumption-dividend ratio in the data.

Table 5: Aggregate Asset Pricing Implications

	Data	Model
$E(P/D)$	25.55	20.96
$\sigma(p-d)$	0.38	0.38
AC of $p-d$	0.87	0.88
$E[R^m - R^f]$	6.33%	7.87%
$\sigma(R^m - R^f)$	19.41%	19.19%
AC of $R^m - R^f$	0.03	-0.04
Sharpe ratio of market	0.33	0.41
AC of Δd	-0.09	-0.04
$\sigma(\Delta d_t)$	14.48%	14.43%

Notes: Aggregate moments for the parameter values in Table 4. The model is simulated for 50,000 quarters and statistics are aggregated to an annual frequency. The asset pricing data is from Campbell's (1999) century-long annual sample starting in 1889 and updated through 2002.

Table 6: Long Horizon Regressions – Excess Returns

	Horizon in Years					
	1	2	4	6	8	10
$\sum_{i=1}^{H} r^m_{t+i} - r^f_{t+i} = \beta_0 + \beta_1(p_t - d_t) + \epsilon_t$						
			Data			
β_1	**-0.12**	**-0.23**	**-0.37**	**-0.60**	**-0.86**	**-1.09**
t-stat	(-2.39)	(-2.44)	(-2.01)	(-2.24)	(-2.97)	(-3.54)
R^2	[0.05]	[0.08]	[0.10]	[0.16]	[0.25]	[0.31]
			Data up to 1994			
β_1	**-0.21**	**-0.39**	**-0.61**	**-0.89**	**-1.16**	**-1.34**
t-stat	(-3.45)	(-4.04)	(-3.17)	(-4.08)	(-5.81)	(-6.22)
R^2	[0.07]	[0.13]	[0.19]	[0.30]	[0.41]	[0.44]
			Model			
β_1	-0.11	-0.21	-0.36	-0.49	-0.58	-0.65
R^2	[0.06]	[0.11]	[0.18]	[0.23]	[0.26]	[0.28]

Notes: Long-horizon regressions of excess returns and dividend growth using Campbell's annual data set and the model simulated using parameter values in Table 2. For each data regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t-statistics (in parentheses) and adjusted R^2 statistics in square brackets. Significant coefficients using the standard t-test at the 5% level are highlighted in bold face.

Table 7: Long Horizon Regressions – Dividend Growth

	Horizon in Years					
	1	2	4	6	8	10
Data						
$\sum_{i=1}^{H} \Delta d_{t+i}^{m} = \beta_0 + \beta_1 (p_t - d_t) + \epsilon_t$						
β_1	0.02	-0.01	-0.04	-0.12	-0.23	-0.31
t-stat	(0.56)	(-0.23)	(-0.34)	(-0.85)	(-1.26)	(-1.61)
R^2	[-0.01]	[-0.01]	[-0.01]	[0.00]	[0.02]	[0.05]
$\sum_{i=1}^{H} \Delta d_{t+i}^{m} = \beta_0 + \beta_1 (c_t - d_t) + \epsilon_t$						
β_1	**0.10**	**0.18**	**0.34**	**0.56**	**0.65**	**0.68**
t-stat	(2.30)	(2.52)	(3.05)	(3.42)	(3.56)	(3.78)
R^2	[0.03]	[0.06]	[0.13]	[0.24]	[0.26]	[0.25]
Model						
$\sum_{i=1}^{H} \Delta d_{t+i}^{m} = \beta_0 + \beta_1 (p_t - d_t) + \epsilon_t$						
β_1	0.05	0.09	0.17	0.24	0.29	0.33
R^2	[0.02]	[0.03]	[0.06]	[0.08]	[0.09]	[0.09]
$\sum_{i=1}^{H} \Delta d_{t+i}^{m} = \beta_0 + \beta_1 z_t + \epsilon_t$						
β_1	3.73	7.09	13.19	18.13	22.23	25.81
R^2	[0.04]	[0.07]	[0.13]	[0.18]	[0.21]	[0.24]

Notes: Long-horizon regressions of excess returns and dividend growth using Campbell's annual data set and the model simulated using parameter values in Table 2. For each data regression, the table reports OLS estimates of the regressors, Newey-West (1987) corrected t-statistics (in parentheses) and adjusted R^2 statistics in square brackets. Significant coefficients using the standard t-test at the 5% level are highlighted in bold face.

Table 8: Growth vs. Value Portfolios: Linear Growth

	G				Growth to Value					V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
$ER^i - R^f$	4.44	4.55	4.85	5.29	6.00	6.91	7.91	8.79	9.37	9.65	5.20
$\sigma(R^i - R^f)$	19.22	19.32	19.48	19.52	19.41	19.05	18.47	17.82	17.24	16.91	8.33
Sharpe Ratio	0.23	0.24	0.25	0.27	0.31	0.36	0.43	0.49	0.54	0.57	0.62
CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$											
α_i	-2.86	-2.80	-2.59	-2.20	-1.49	-0.47	0.76	1.94	2.84	3.30	6.17
β_i	0.99	1.00	1.01	1.02	1.02	1.00	0.97	0.93	0.89	0.86	-0.13
R_i^2	0.96	0.96	0.97	0.98	0.99	1.00	1.00	0.98	0.95	0.94	0.09
CAPM & HML: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \gamma_i \text{HML}_t + \epsilon_{it}$											
α_i	0.04	0.03	0.03	0.02	0.02	0.01	0.01	0.02	0.03	0.04	0.00
β_i	0.93	0.94	0.95	0.97	0.99	0.99	0.99	0.97	0.95	0.93	0.00
γ_i	-0.47	-0.46	-0.43	-0.36	-0.25	-0.08	0.12	0.31	0.46	0.53	1.00
R_i^2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Unconditional Correlations											
$\rho(R^i, M)$	-0.08	-0.08	-0.09	-0.09	-0.11	-0.12	-0.14	-0.16	-0.18	-0.19	-0.19
$\rho(R^i, \epsilon^1)$	0.11	0.11	0.12	0.13	0.14	0.17	0.19	0.22	0.24	0.26	0.27
$\rho(R^i, \epsilon^2)$	0.21	0.21	0.21	0.20	0.19	0.19	0.18	0.18	0.17	0.17	-0.14
$\rho(R^i, \epsilon^3)$	-0.41	-0.41	-0.41	-0.41	-0.40	-0.39	-0.36	-0.34	-0.31	-0.29	0.35

Notes: This table presents results for 10 portfolios formed according to the price-dividend ratios of firms. Linearly declining growth of firm dividends from 75% p.a. to 0% p.a. No lower bound for dividends.

Table 9: Growth vs. Value Portfolios: Constant Growth

	G			Growth to Value						V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
$ER^i - R^f$	5.00	5.18	5.47	5.90	6.46	7.15	7.89	8.58	9.16	10.08	5.09
$\sigma(R^i - R^f)$	19.27	19.48	19.64	19.67	19.51	19.08	18.38	17.56	16.99	17.30	8.27
Sharpe Ratio	0.26	0.27	0.28	0.30	0.33	0.37	0.43	0.49	0.54	0.58	0.62
CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$											
α_i	-2.60	-2.52	-2.31	-1.93	-1.33	-0.50	0.52	1.59	2.48	3.38	5.98
β_i	1.00	1.01	1.02	1.03	1.02	1.00	0.97	0.92	0.88	0.88	-0.12
R_i^2	0.97	0.97	0.97	0.98	0.99	1.00	1.00	0.98	0.96	0.93	0.07
CAPM & HML: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \gamma_i \text{HML}_t + \epsilon_{it}$											
α_i	0.05	0.04	0.02	0.01	0.01	0.01	0.03	0.05	0.06	0.05	0.00
β_i	0.95	0.96	0.98	0.99	1.00	0.99	0.98	0.95	0.93	0.95	0.00
γ_i	-0.44	-0.43	-0.39	-0.32	-0.22	-0.09	0.08	0.26	0.40	0.56	1.00
R_i^2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Unconditional Correlations											
$\rho(R^i, M)$	-0.06	-0.06	-0.06	-0.07	-0.08	-0.09	-0.10	-0.12	-0.14	-0.15	-0.19
$\rho(R^i, \epsilon^1)$	0.09	0.09	0.10	0.11	0.12	0.14	0.17	0.19	0.22	0.24	0.29
$\rho(R^i, \epsilon^2)$	0.19	0.19	0.18	0.17	0.17	0.16	0.16	0.16	0.15	0.14	-0.16
$\rho(R^i, \epsilon^3)$	-0.41	-0.41	-0.41	-0.41	-0.40	-0.39	-0.37	-0.34	-0.32	-0.29	0.34

Notes: Constant 20% p.a. expected dividend growth in growth and decline phases. No lower bound for dividends.

Table 10: Momentum Portfolios

	L				Losers to Winners					W	W-L
Horizon H	1	2	3	4	5	6	7	8	9	10	10-1
1	7.68	7.58	7.54	7.75	8.34	9.21	9.98	10.53	10.97	11.22	3.54
2	7.74	7.64	7.61	7.83	8.39	9.15	9.87	10.46	10.91	11.18	3.44
5	7.92	7.80	7.76	7.95	8.42	9.06	9.73	10.33	10.79	11.03	3.12
10	8.09	7.96	7.89	8.04	8.44	9.01	9.63	10.21	10.66	10.90	2.81
$H = 1$					CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$						
α_i	-1.68	-1.75	-1.77	-1.63	-1.20	-0.47	0.37	1.05	1.59	1.91	3.59
β_i	0.96	0.95	0.95	0.96	0.98	1.00	0.99	0.97	0.96	0.95	-0.01
R_i^2	0.96	0.96	0.97	0.98	0.99	0.99	0.99	0.98	0.97	0.95	0.00

Notes: Momentum portfolios are formed by sorting firm returns based on their past 1-year return. The table presents mean returns over the subsequent H years.

Table 11: Growth vs. Value Portfolios: Linear Growth with Lower Bound

	G				Growth to Value					V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
$ER^i - R^f$	5.81	5.60	5.51	5.80	6.37	7.10	7.88	8.58	9.09	9.25	3.44
$\sigma(R^i - R^f)$	20.32	19.61	18.88	18.73	18.84	18.75	18.40	17.96	17.54	17.19	7.24
Sharpe Ratio	0.29	0.29	0.29	0.31	0.34	0.38	0.43	0.48	0.52	0.54	0.48
CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$											
α_i	-2.15	-2.07	-1.88	-1.56	-1.06	-0.32	0.60	1.53	2.26	2.59	4.73
β_i	1.07	1.03	0.99	0.99	1.00	1.00	0.98	0.95	0.92	0.89	-0.17
R_i^2	0.98	0.98	0.98	0.99	0.99	1.00	1.00	0.99	0.97	0.96	0.20
CAPM & HML: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \gamma_i \text{HML}_t + \epsilon_{it}$											
α_i	0.02	0.00	-0.04	-0.04	-0.02	0.00	0.00	0.01	0.02	0.02	0.00
β_i	0.99	0.95	0.93	0.93	0.96	0.98	1.00	1.00	1.00	0.99	0.00
γ_i	-0.46	-0.44	-0.39	-0.32	-0.22	-0.07	0.13	0.32	0.47	0.54	1.00
R_i^2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Unconditional Correlations											
$\rho(R^i, M)$	-0.07	-0.07	-0.07	-0.07	-0.08	-0.09	-0.10	-0.11	-0.12	-0.13	-0.12
$\rho(R^i, \epsilon^1)$	0.12	0.12	0.12	0.13	0.14	0.16	0.18	0.20	0.22	0.23	0.21
$\rho(R^i, \epsilon^2)$	0.17	0.18	0.19	0.19	0.18	0.17	0.16	0.16	0.15	0.15	-0.13
$\rho(R^i, \epsilon^3)$	-0.40	-0.40	-0.40	-0.39	-0.39	-0.38	-0.36	-0.34	-0.32	-0.31	0.39

Notes: Linearly declining growth of firm dividends from 75% p.a. to 0% p.a. Lower bound 0.1%.

Table 12: Growth vs. Value Portfolios: Constant Growth with Lower Bound

	G			Growth to Value						V	V-G
Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
$ER^i - R^f$	5.50	5.68	5.87	6.13	6.50	7.00	7.58	8.18	8.71	9.40	3.91
$\sigma(R^i - R^f)$	19.35	19.63	19.58	19.42	19.20	18.89	18.44	17.88	17.46	17.60	6.72
Sharpe Ratio	0.28	0.29	0.30	0.32	0.34	0.37	0.41	0.46	0.50	0.53	0.58
CAPM: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$											
α_i	-1.97	-1.90	-1.71	-1.40	-0.96	-0.36	0.40	1.24	1.99	2.69	4.67
β_i	1.01	1.03	1.03	1.02	1.01	1.00	0.97	0.94	0.91	0.91	-0.10
R_i^2	0.98	0.98	0.99	0.99	0.99	1.00	1.00	0.99	0.97	0.96	0.08
CAPM & HML: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \gamma_i \text{HML}_t + \epsilon_{it}$											
α_i	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.02	0.00
β_i	0.97	0.99	0.99	0.99	0.99	0.99	0.98	0.97	0.95	0.97	0.00
γ_i	-0.43	-0.41	-0.37	-0.31	-0.21	-0.08	0.08	0.26	0.42	0.57	1.00
R_i^2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Unconditional Correlations											
$\rho(R^i, M)$	-0.09	-0.09	-0.09	-0.10	-0.11	-0.12	-0.13	-0.14	-0.16	-0.17	-0.19
$\rho(R^i, \epsilon^1)$	0.15	0.15	0.15	0.16	0.17	0.18	0.20	0.22	0.24	0.26	0.27
$\rho(R^i, \epsilon^2)$	0.19	0.18	0.18	0.18	0.17	0.17	0.16	0.16	0.15	0.14	-0.16
$\rho(R^i, \epsilon^3)$	-0.39	-0.39	-0.38	-0.38	-0.37	-0.37	-0.35	-0.33	-0.31	-0.29	0.35

Notes: Constant 20% p.a. expected dividend growth in growth and decline phases. Lower bound for dividends 0.1% of total dividends.

Figure 1: Model Solution

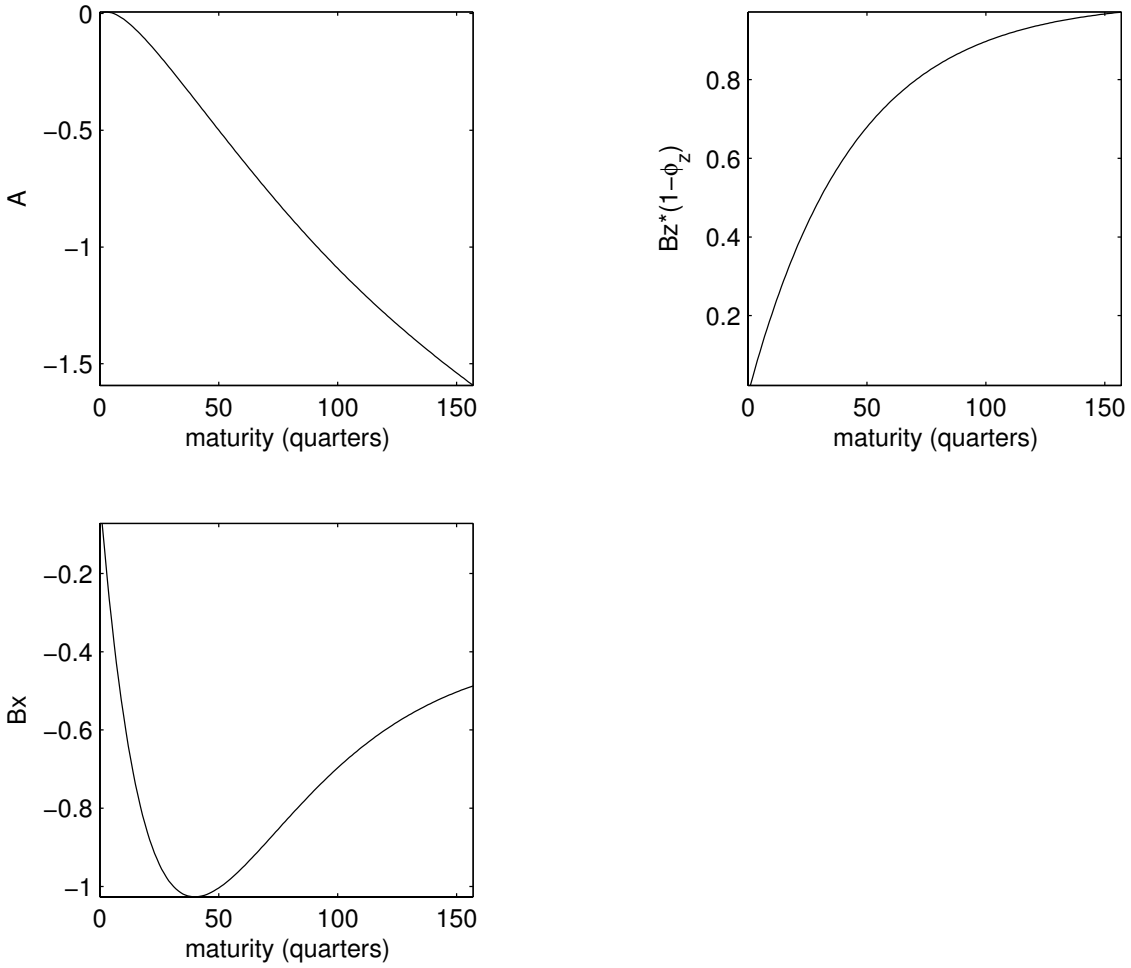

Notes: This figure shows the solutions of the system of algebraic equations (??)-(??) for the parameter values in Table 4.

Figure 2: Price/Dividend Ratio of Zero-Coupon Equity

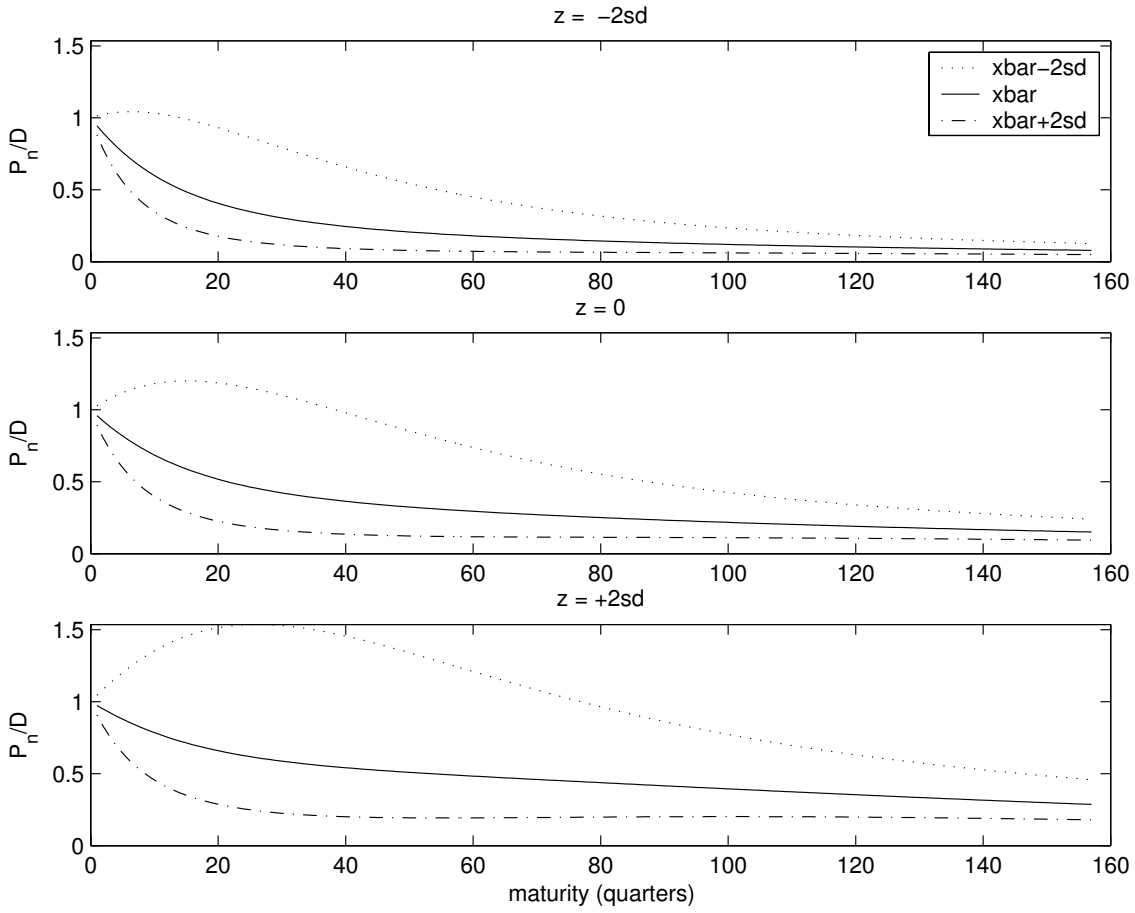

Notes: This figure plots the ratio of the price of n-period zero-coupon equity to the aggregate (quarterly) dividend as a function of maturity n for the parameter values in Table 4. The three panels show the zero-coupon curves for $z = -2\|\sigma_z\|/\sqrt{1-\phi_z^2}$, $z = 0$ and $z = 2\|\sigma_z\|/\sqrt{1-\phi_z^2}$. Each panel plots three graphs for $x = \bar{x} - 2\|\sigma_x\|/\sqrt{1-\phi_x^2}$, $x = \bar{x}$ and $x = \bar{x} + 2\|\sigma_x\|/\sqrt{1-\phi_x^2}$.

Figure 3: Properties of Zero-Coupon Equity

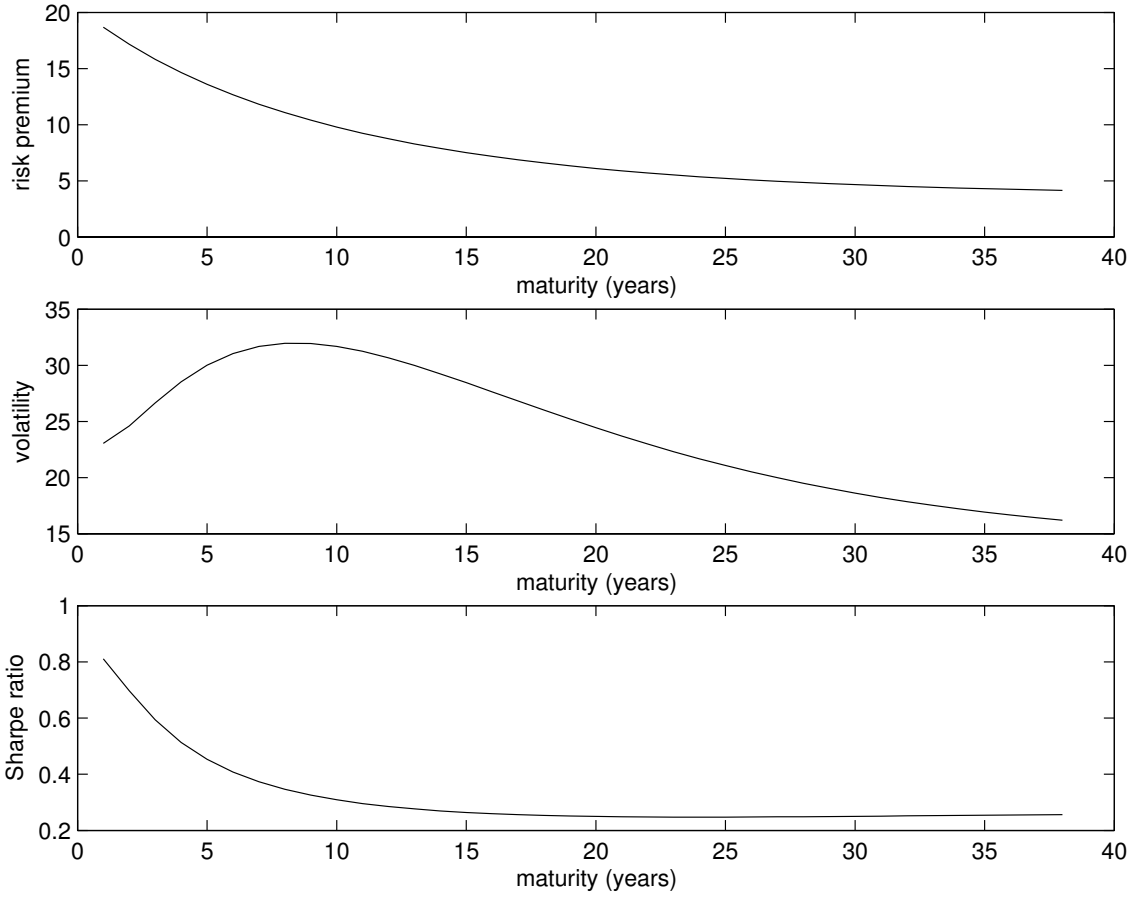

Notes: This figure shows return statistics for zero-coupon equity as a function of maturity. The top panel plots risk premia of annual level returns over the riskfree rate. The second panel plots the standard deviation of annual level returns. The third panel plot the Sharpe ratio (risk premia dividend by standard deviation).

Figure 4: CAPM Regressions for Zero-Coupon Equity

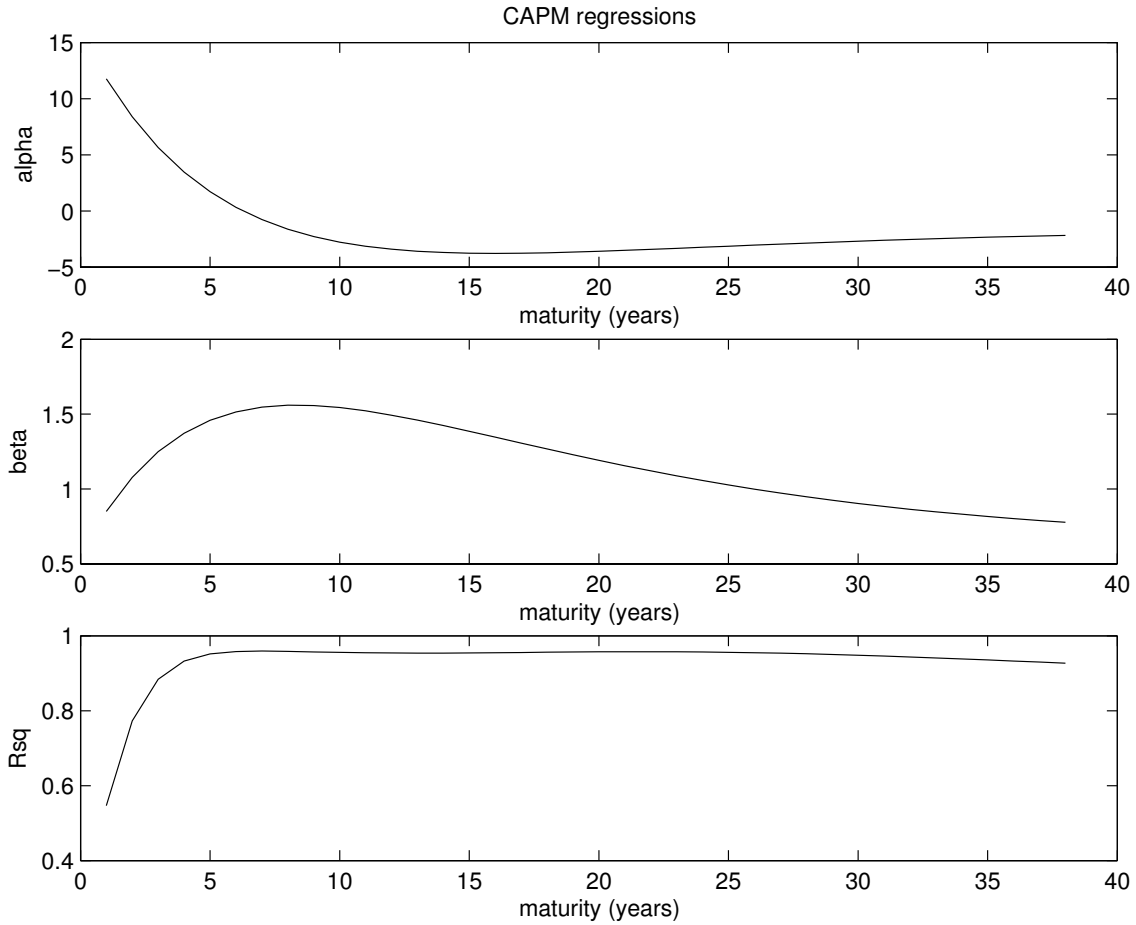

Notes: This figure shows results from unconditional CAPM regressions of annual returns of zero-coupon equity on the market return. The top panel plots $\alpha^{(n)}$, the middle panel plots $\beta^{(n)}$ and the bottom panel plots the $R^{2(n)}$ as a function of maturity.

Figure 5: Cross-Sectional Distribution – Linear Growth

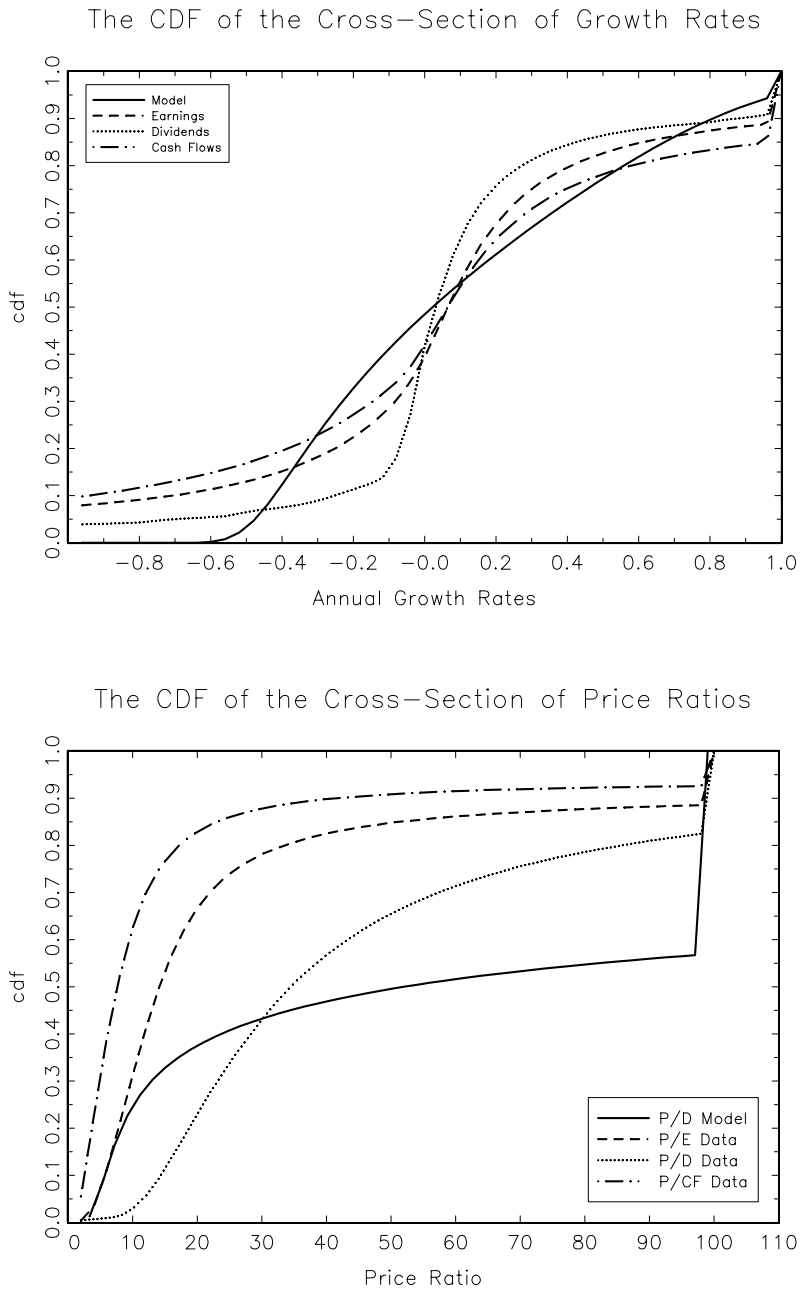

Notes: This table compares the cross-sectional distribution of growth rates of various cash flow measures and price ratios in the model to those in the data. The top panel plots the distribution of annual growth rates of dividends, earnings and cash flows (defined as in Fama and French) across all firms for the 1952 – 2002 period. Growth rates are censored at 100%. Firms that exit the sample are assigned a growth rate of -100%. The solid line is the distribution of annual dividend growth rates for all firms in a simulation of 50,000 quarters. The bottom panel plots the corresponding distribution of various price multiples in the data and in the simulated model.

Figure 6: Cross-Sectional Distributions – Constant Growth

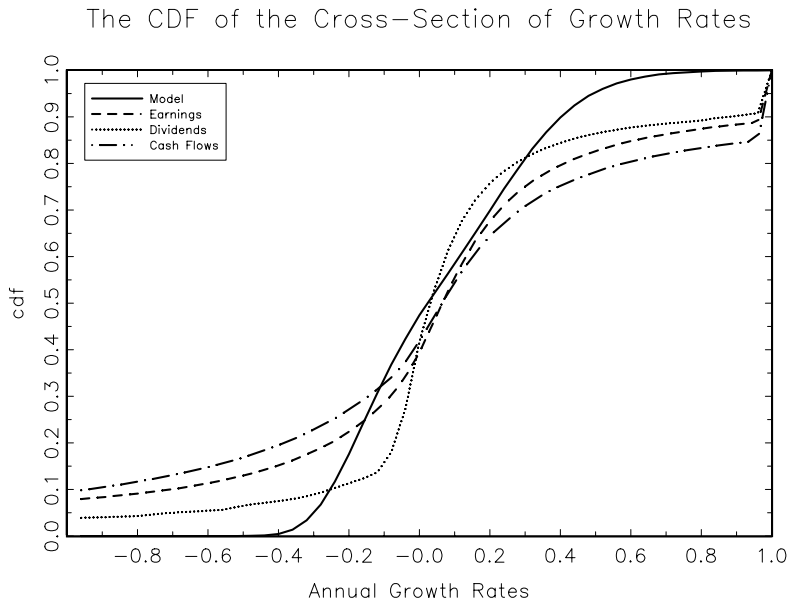

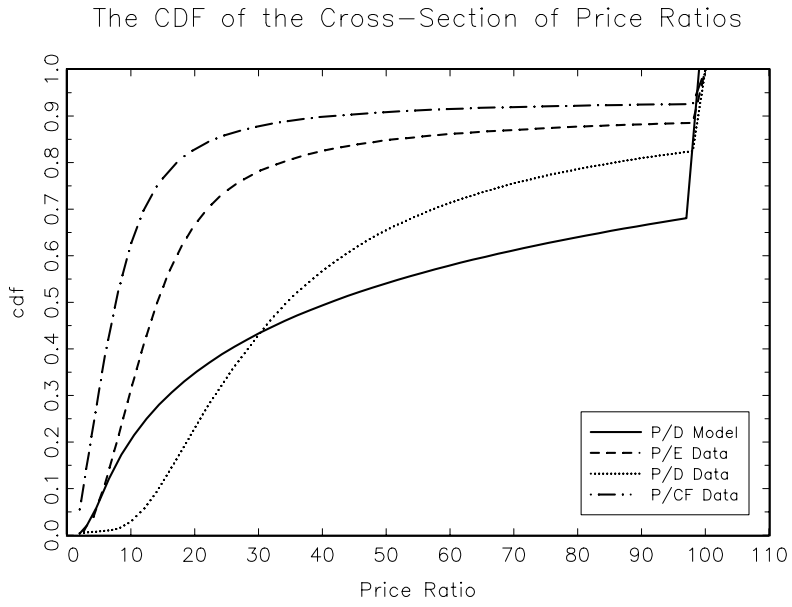

Notes: This table compares the cross-sectional distribution of growth rates of various cash flow measures and price ratios in the model to those in the data. The top panel plots the distribution of annual growth rates of dividends, earnings and cash flows (defined as in Fama and French) across all firms for the 1952 – 2002 period. Growth rates are censored at 100%. Firms that exit the sample are assigned a growth rate of -100%. The solid line is the distribution of annual dividend growth rates for all firms in a simulation of 50,000 quarters. The bottom panel plots the corresponding distribution of various price multiples in the data and in the simulated model.

Figure 7: Effect of ρ_{dx} on Zero-Coupon Equity

Panel A: $\rho_{dx} = 0$

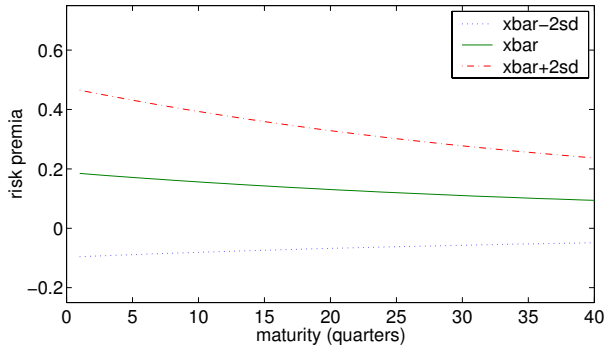

Panel B: $\rho_{dx} = -0.5$

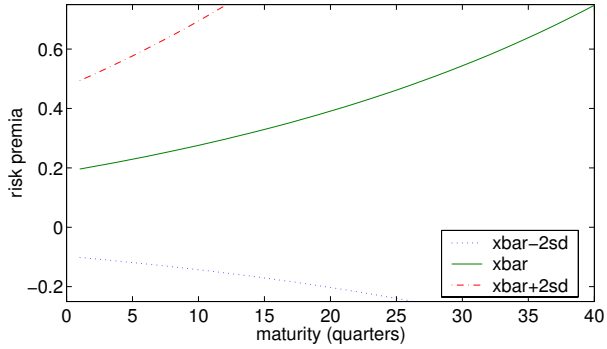

Panel C: $\rho_{dx} = 0.5$

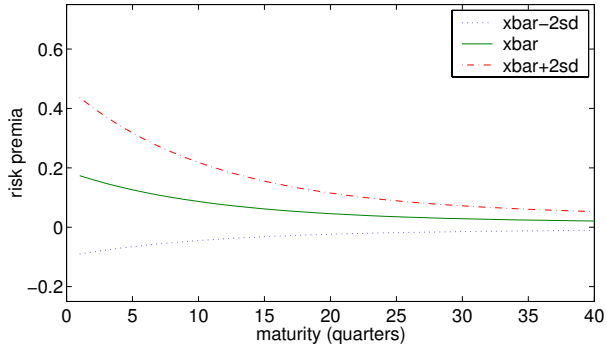

Notes: The three panels show the risk premium for zero-coupon equity as a function of the maturity for the parameter values in Table 4. Each panel has three graphs that plot the risk premium at the mean of x_t, $\bar{x}$, and $\bar{x} \pm 2\sigma_x$, respectively. The top panel assumes that the correlation of dividend shocks and shocks to x_t is zero, the middle panel assumes that the correlation of dividend shocks and shocks to x_t is -0.5 and the bottom panel assumes that the correlation of dividend shocks and shocks to x_t is 0.5. All other parameter values are identical.

Figure 8: Regressions on Fundamental Shocks

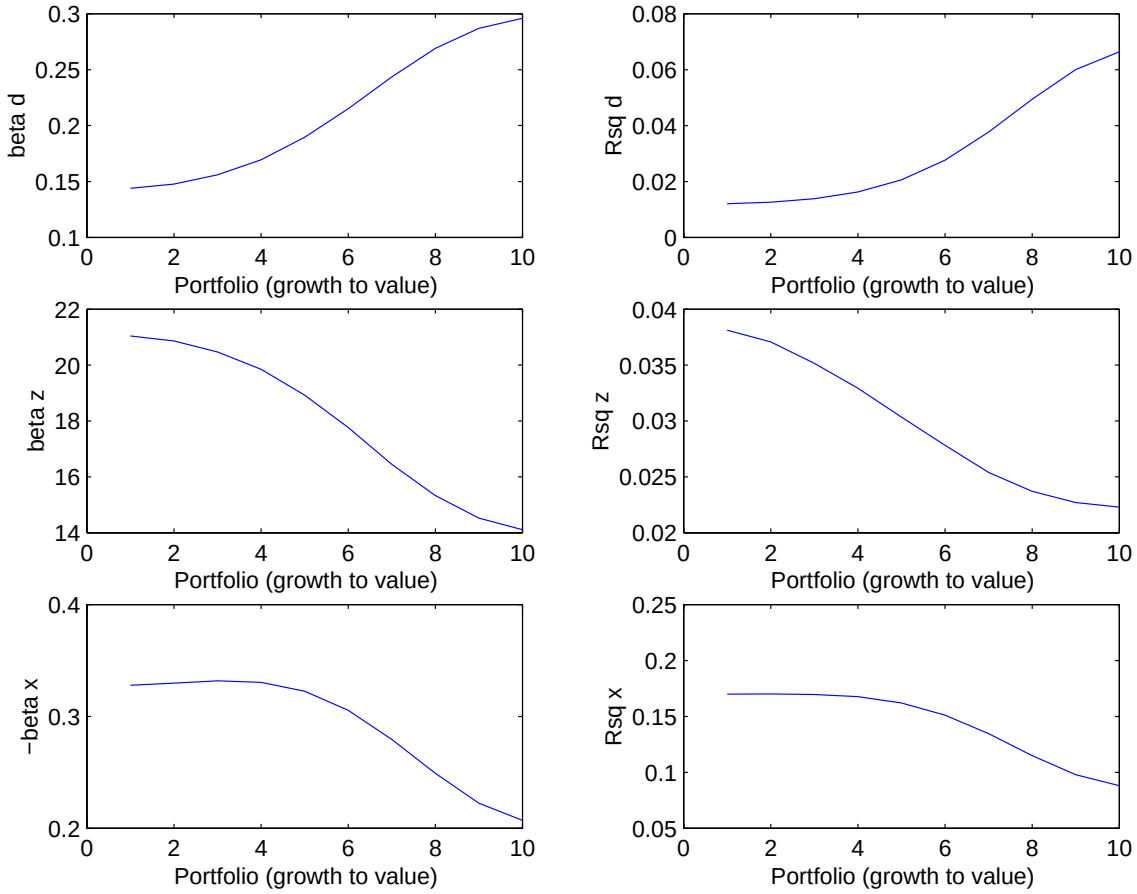

Notes: This figure plots the OLS coefficients and R^2 statistics from regressions of annual portfolio returns on the fundamental shocks. Firms are constructed assuming linearly decreasing dividend growth. The top panels regress portfolio returns on the shock to dividends $2\sigma_d\epsilon_t$, the middle panels regress portfolio returns on the shocks to the component of expected dividend growth that is uncorrelated with the shock to dividends $2\sigma_z(2)\epsilon_t(2)$, and the bottom panels regress firm returns on the shocks to the Sharpe ratio $2\sigma_x\epsilon_t$ (note that the shocks to the Sharpe ratio are uncorrelated with shocks to dividends and expected dividends).

Figure 9: Cross-Sectional Distributions – Linear Growth with with Lower Bound

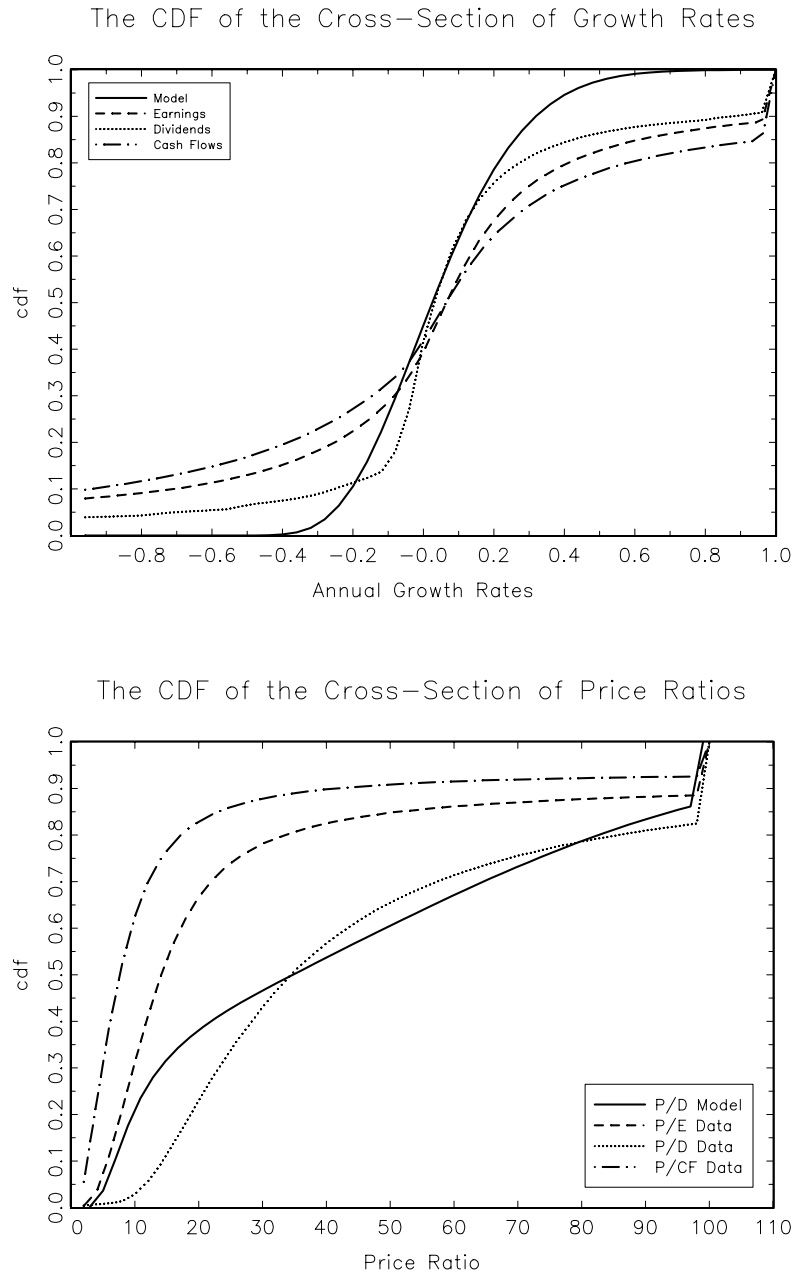

Notes: See notes of Table 7 for the linear growth case with lower bound of 0.1%.

Figure 10: Cross-Sectional Distributions – Constant Growth with Lower Bound

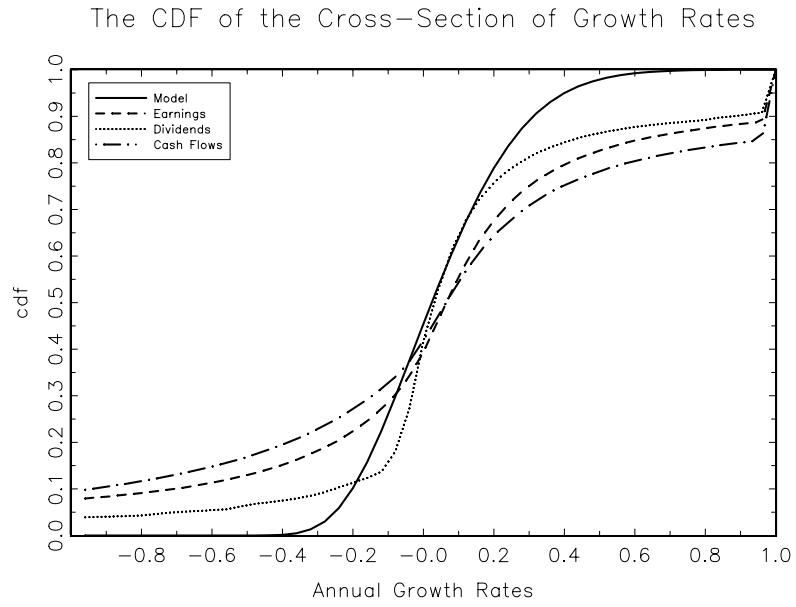

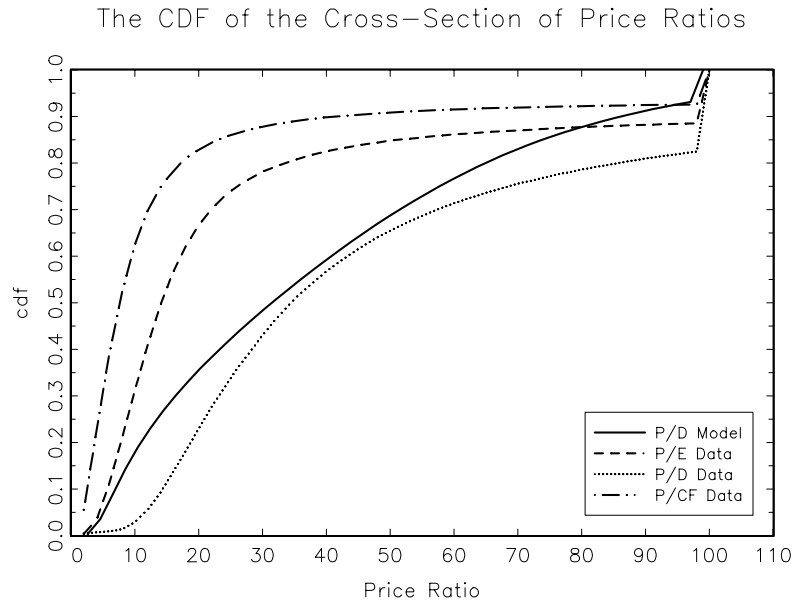

Notes: See notes of Table 7 for the constant growth case with lower bound of 0.1%.