

Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium

MARTIN LETTAU and JESSICA A. WACHTER*

ABSTRACT

We propose a dynamic risk-based model that captures the value premium. Firms are modeled as long-lived assets distinguished by the timing of cash flows. The stochastic discount factor is specified so that shocks to aggregate dividends are priced, but shocks to the discount rate are not. The model implies that growth firms covary more with the discount rate than do value firms, which covary more with cash flows. When calibrated to explain aggregate stock market behavior, the model accounts for the observed value premium, the high Sharpe ratios on value firms, and the poor performance of the CAPM.

THIS PAPER PROPOSES A DYNAMIC RISK-BASED MODEL that captures both 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 a high ratio of price to fundamentals (growth stocks) have low expected returns relative to assets with a low ratio of price to fundamentals (value stocks). This finding by itself is not necessarily surprising, as it is possible that the premium on value stocks represents compensation for bearing systematic risk. However, Fama and French (1992) and others show that the capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) cannot account for the value premium: While the CAPM predicts that expected returns should rise with the beta on the market portfolio, value stocks have higher expected returns yet do not have higher betas 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 endogenously have high price ratios, while firms with cash flows weighted more to the present have low price ratios. Analogous to long-term bonds, growth firms are high-duration

*Lettau is at the Stern School of Business at New York University. Wachter is at the Wharton School at the University of Pennsylvania. The authors thank Andrew Abel, Jonathan Berk, John Campbell, David Chapman, John Cochrane, Lars Hansen, Leonid Kogan, Sydney Ludvigson, Anthony Lynch, Stijn Van Nieuwerburgh, an anonymous referee, and seminar participants at the 2004 National Bureau of Economic Research Summer Institute, the 2005 Society of Economic Dynamics meetings, the 2005 Western Finance Association Meetings, Duke University, New York University, Pennsylvania State University, University of British Columbia, University of Chicago, University of Pennsylvania, and Yale University for helpful comments.

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. Two properties of the stochastic discount factor account for the model's ability to fit the data. First, the price of risk varies, implying that at some times investors require a greater return per unit of risk than at others. Second, variation in the price of risk is not perfectly linked to variation in aggregate fundamentals. We show that the correlation between aggregate dividend growth and the price of risk crucially determines the ability of the model to fit the cross section.

We require that our model match not only the cross section of assets based on price ratios, but also aggregate dividend and stock market behavior. First, we assume that log dividend growth is normally distributed with a time-varying mean and calibrate the dividend process to fit conditional and unconditional moments of the aggregate dividend process in the data. Firms are distinguished by their cash flows, which we specify as stationary shares of the aggregate dividend. This modeling strategy, also employed by Menzly, Santos, and Veronesi (2004), ensures that the economy is stationary, and that firms add up to the market. Second, we choose stochastic discount factor parameters to fit the time series of aggregate stock market returns. These choices imply that expected excess returns on equity are time varying in the model, that there is excess volatility, and that excess returns are predictable. We find that the model can match unconditional moments of the aggregate stock market and produce dividend and return predictability close to that found in the data.

To test whether our model can capture the value premium, we sort firms into portfolios in simulated data. We find that risk premia, risk-adjusted returns, and Sharpe ratios increase in the value decile. The value premium (the expected return on a strategy that is long the extreme value portfolio and short the extreme growth portfolio) is 5.1% in the model compared with 4.9% in the data when portfolios are formed by sorting on book-to-market. Moreover, the CAPM alpha on the value-minus-growth strategy is 6.0% in the model, compared with 5.6% in the data. These results do not arise because value stocks are more risky according to traditional measures: Rather, standard deviations and market betas increase slightly in the value decile and then decrease, implying that the extreme value portfolio has a lower standard deviation and beta than the extreme growth portfolio. Our model therefore matches both the magnitude of the value premium and the outperformance of value portfolios relative to the CAPM that obtain in the data.

In its focus on explaining the value premium through cash flow fundamentals, our model is part of a growing literature that emphasizes the cash flow dynamics of the firm and how these relate to discount rates. In particular, in a model in which firms have assets in place as well as real growth options, Berk, Green, and Naik (1999) show that acquiring an asset with low systematic risk leads to a decrease in the firm's book-to-market ratio and lower future returns. More recently, Gomes, Kogan, and Zhang (2003) explicitly link risk premia to characteristics of firm cash flows in general equilibrium and Zhang

(2005) shows how asymmetric adjustment costs and a time-varying price of risk interact to produce value stocks that suffer increased risk during downturns. These models endogenously derive patterns in the cross section of returns from cash flows, but they do not account for the classic finding of Fama and French (1992) that value stocks outperform, and growth stocks underperform, relative to the CAPM.

Our model for the stochastic discount factor builds on the work of Brennan, Wang, and Xia (2004) and Brennan and Xia (2006) and is closely related to essentially affine term structure models (Dai and Singleton (2003), Duffee (2002)). As Brennan et al. show, their model for the stochastic discount factor implies that claims to single dividend payments are exponential-affine in the state variables, which allows for economically interpretable closed-form expressions for prices and risk premia. Motivated by these expressions, Brennan et al. empirically evaluate whether expected returns on a cross-section of assets can be explained by betas with respect to discount rates. Here we make use of similar analytical methods to address a different goal, namely, endogenously generating a value premium based on the firm's underlying cash flows.

Our paper also builds on work 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 may have high betas because of the duration of their cash flows, even if the risk of these cash flows is mainly idiosyncratic. Berk, Green, and Naik (2004) value a firm with large research and development expenses and show how discount rate and cash flow risk interact to produce risk premia that change over the course of a project. Their model endogenously generates a long duration for growth stocks. Leibowitz and Kogelman (1993) show that accounting for the sensitivity of the value of long-run cash flows to discount rates can reconcile various measures of equity duration. Dechow, Sloan, and Soliman (2004) 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 betas. 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 (2004) decompose the market return into news about cash flows and news about discount rates. They show that growth stocks have higher betas with respect to discount rate news than do value stocks, 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, further, that growth stocks seem particularly subject to such discount rate risk. Our model shows how these contributions can be parsimoniously tied together with those discussed in the paragraphs above.

Finally, this paper relates to the large and growing body of empirical research that explores the correlations of returns on value and growth stocks with sources of systematic risk. This literature explores conditional versions of traditional models (Jagannathan and Wang (1996), Lettau and Ludvigson

(2001a), Petkova and Zhang (2005), Santos and Veronesi (2006)) and identifies new sources of risk that covaries more with value stocks than with growth stocks (Lustig and Van Nieuwerburgh (2005), Piazzesi, Schneider, and Tuzel (2005), Yogo (2006)). Another strand of literature relates observed returns of value and growth stocks to aggregate market cash flows or macroeconomic factors (Campbell, Polk, and Vuolteenaho (2003), Liew and Vassalou (2000), Parker and Julliard (2005), Vassalou (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 dividends on value and growth portfolios directly (Bansal, Dittmar, and Lundblad (2005), Cohen, Polk, and Vuolteenaho (2003), and Hansen, Heaton, and Li (2004)) and finds evidence that the cash flows of value stocks covary more with aggregate cash flows. The results in these papers raise the question of why the observed covariation leads to the value premium. By explicitly linking firms' cash flow properties and risk premia, this paper takes a step toward answering this question.

The paper is organized as follows. Section I updates evidence that portfolios formed by sorting on 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, Sharpe ratio, and alpha tend to increase in the value decile. The differences in expected returns and alphas between value and growth portfolios are statistically and economically large.

Section II presents our model for aggregate dividends and the stochastic discount factor. As a first step toward solving for prices of the aggregate market and firms, we solve for prices of claims to the aggregate dividend n periods in the future (zero-coupon equity). Because zero-coupon equity has a well-defined maturity, it provides a convenient window through which to view the role of duration in our model. The aggregate market is the sum of all the zero-coupon equity claims. We then introduce a cross section of long-lived assets, defined by their shares in the aggregate dividend. These assets are themselves portfolios of 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 III discusses the time-series and cross-sectional implications of our model. We calibrate the model to the time series of 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 find that the parameters necessary to fit the time series imply risk premia, Sharpe ratios, and alphas for zero-coupon equity that are increasing in maturity. In contrast, CAPM betas and volatilities are nonmonotonic, and thus do not explain the increase in risk premia. This suggests that our 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. When we sort the resulting assets into portfolios, our model can explain the observed value premium.

Section IV discusses 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 (2004), growth stocks have greater betas with respect to discount rates than do 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. Thus, even though long-horizon equity is riskier according to standard deviation and market beta, it is not seen as risky by investors because it loads on risks that investors do not mind bearing.

I. Evidence on the Value Premium

Much of the previous literature shows that portfolios of stocks with high ratios of prices to fundamentals have low future returns compared to stocks with low ratios of prices to fundamentals.¹ In this section, we update 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 increase in the measure of value. In our model, firms are distinguished by their cash flows, thus earnings, dividends, and cash flows are equivalent. For this reason, it is of interest to investigate whether the value effect is apparent in portfolios formed according to different measures of value.

Table I reports summary statistics for portfolios of firms sorted into deciles on each of the three characteristics described above and on book-to-market. Data, available from the website of Ken French, are monthly, from 1952 to 2002. We compute excess returns by subtracting monthly returns on the 1-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 1,200 in the case of means and $\sqrt{12} \times 100$ in the case of standard deviations). Each panel reports results for the earnings-to-price ratio, the cash-flow-to-price ratio, the dividend yield, and the book-to-market ratio.

¹ See Graham and Dodd (1934), Basu (1977, 1983), Ball (1978), Rosenberg, Reid, and Lanstein (1985), Jaffe, Keim, and Westerfield (1989), and Fama and French (1992). Cochrane (1999) surveys recent literature on the value effect.

Table I
Summary Statistics for Growth and Value Portfolios

Portfolios are formed by sorting firms into deciles on the dividend yield (D/P), the earnings yield (E/P), the ratio of cash flow to prices (C/P), and the book-to-market ratio (B/M). Moments are in annualized percentages (multiplied by 1,200 in the case of means and $\sqrt{12} \times 100$ in the case of standard deviations). The data are monthly and span the 1952 to 2002 period.

Portfolio	G	Growth to Value								V	V-G
	1	2	3	4	5	6	7	8	9	10	10-1
Panel A: 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
Panel B: 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
Panel C: 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
Panel D: 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

Panel A of Table I shows that for all measures except the dividend yield, the mean excess return is higher for the upper deciles (value) than for the lower deciles (growth). Panel B shows that 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 by sorting on the dividend yield. Panel C shows that the standard deviation of the excess return tends to decrease in the decile number, and thus move in the opposite direction of the mean return. Finally, Panel D shows that the Sharpe ratio increases in the decile number. For example, when portfolios are formed by sorting on the earnings-to-price ratio, the bottom decile (growth) has a Sharpe ratio of 0.24. The Sharpe ratio increases as the earnings-to-price ratio increases and the top decile (value) has a Sharpe ratio of 0.72. Thus value stocks not only deliver high returns, they deliver high returns per unit of standard deviation.

Table II
Correlation of Returns on Extreme Value and Growth Portfolios

Portfolios are formed by sorting firms into deciles on the dividend yield (D/P), the earnings yield (E/P), the ratio of cash flow to prices (C/P), and the book-to-market ratio (B/M). The data are monthly and span the 1952 to 2002 period.

	E/P	C/P	D/P	B/M
Panel A: Top Decile (Value)				
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
Panel B: Bottom Decile (Growth)				
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

The results in Table I suggest that portfolios formed by sorting on earnings-to-price, cash-flow-to-price, the dividend yield, and book-to-market may be closely related. This is confirmed in Table II, which shows the correlation of the bottom and top deciles. 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 I, which shows that portfolios formed by sorting on the dividend yield behave somewhat differently from portfolios formed by sorting on the other variables.

Following the same format as Table I, Table III shows alphas, standard errors on alphas, betas, standard errors on betas, and R^2 statistics when portfolios are formed by sorting on each measure of value. Alpha is the intercept from an ordinary least squares (OLS) regression of portfolio excess returns on excess returns of the value-weighted CRSP index, multiplied by 1,200. 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 A of this table confirms the classic result that value stocks have high alphas relative to the CAPM. Moreover, the story is consistent across all sorting variables, including the dividend yield: Alphas are negative for growth stocks, positive for value stocks, and increasing in the decile number. As Panel C shows, betas tend to decline in the decile number, except for the extreme value portfolio. Thus, value stocks have positive alphas relative to the CAPM, and relatively low betas.

To summarize, this section shows that, in the data, value stocks have higher expected excess returns and higher Sharpe ratios than do growth stocks. Value

Table III
Performance of Growth and Value Portfolios Relative to the CAPM

Intercepts and slope coefficients are calculated from OLS time-series regressions of excess portfolio returns on the excess return on the value-weighted CRSP index. Portfolios are formed by sorting firms into deciles on the dividend yield (D/P), the earnings yield (E/P), the ratio of cash flow to prices (C/P), and the book-to-market ratio (B/M). Intercepts are in annualized percentages (multiplied by 1,200). The data are monthly and span the 1952 to 2002 period.

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

Portfolio	G	Growth to Value								V	V-G
	1	2	3	4	5	6	7	8	9	10	10-1
Panel A: 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
Panel B: 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
Panel C: 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
Panel D: 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
Panel E: 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

stocks have large positive alphas while growth stocks have negative alphas. Moreover, value stocks do not have higher standard deviations or higher betas than do growth stocks. Thus, any explanation of the value premium must 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 hold not only when value is defined by the book-to-market ratio, but also when value is defined by the earnings-to-price or cash-flow-to-price ratios.

II. The Model

This section presents our model. The first subsection discusses our assumptions on aggregate cash flows and the stochastic discount factor. The second subsection solves for prices on equity that pays the aggregate dividend in a fixed number of years; we refer to these claims as “zero-coupon equity,” and they form the building blocks of our more complex assets. The third subsection describes the market portfolio.

A. Dividend Growth and the Stochastic Discount Factor

The model has three shocks, namely, a shock to dividend growth, a shock to expected dividend growth, and a shock to the preference variable. We let ϵ_{t+1} denote a 3×1 vector of independent standard normal shocks that are independent of variables observed at time t . Let D_t denote the aggregate dividend in the economy at time t , and $d_t = \ln D_t$. The aggregate dividend is assumed to evolve according to

$$\Delta d_{t+1} = g + z_t + \sigma_d \epsilon_{t+1}, \quad (1)$$

where z_t follows the AR(1) process

$$z_{t+1} = \phi_z z_t + \sigma_z \epsilon_{t+1}, \quad (2)$$

with $0 \leq \phi_z < 1$. The conditional mean of dividend growth is $g + z_t$. Row vectors σ_d and σ_z multiply the shocks on dividend growth and z_{t+1} . The conditional standard deviation of Δd_{t+1} equals $\|\sigma_d\| = \sqrt{\sigma_d \sigma_d'}$. 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 (2004) and by Campbell (1999).

We directly specify the stochastic discount factor for this economy. In particular we assume that the price of risk is driven by a single state variable x_t that follows the AR(1) process

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

with $-1 \leq \phi_x < 1$. As above, σ_x is a 1×3 vector. This specification for the price of risk is used in a continuous-time setting by Brenetal et al. (2004). However, for simplicity, we assume that the real risk-free rate, denoted $r^f = \ln R^f$, 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 (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 the dividend shock is priced. This specification also allows us to compare our model to the external habit formation models of Campbell and Cochrane (1999) and Menzly et al. (2004), in which the only shock to the stochastic discount factor comes from aggregate consumption. The assumption that only dividend risk is priced implies that shocks to z_t and x_t will only be priced insofar as they correlate with Δd_{t+1} .

The specification of x_t and r^f and the fact that only dividend risk is priced together imply that the stochastic discount factor equals

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

where

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

The conditional 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 risk-free 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 as to how to interpret the variable x_t . In the models of Campbell and Cochrane (1999) and Menzly et al. (2004), 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. The corresponding assumption here is $\sigma_x/\|\sigma_x\| = -\sigma_d/\|\sigma_d\|$. However, we depart from these papers by assuming that shocks to x_{t+1} are uncorrelated with shocks to Δd_{t+1} and 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.

B. Prices of Zero-Coupon Equity

The building blocks of the long-lived assets in our economy are “zero-coupon” equity.² 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 one-period return on zero-coupon equity that matures in n periods. That is,

² The method of separating the aggregate dividend into its zero-coupon components and using affine term structure techniques to value each component is also applied in Ang and Liu (2004), Bakshi and Chen (1996), Bekaert, Engstrom, and Grenadier (2004), Johnson (2002), Wachter (2006), and Wilson (2003).

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

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

$$E_t[M_{t+1}R_{n,t+1}] = 1, \quad (6)$$

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

$$P_{nt} = E_t[M_{t+1}P_{n-1,t+1}], \quad (7)$$

with boundary condition

$$P_{0t} = D_t, \quad (8)$$

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

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

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

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

Matching coefficients on the constant, z_t , and x_t implies that

$$A(n) = A(n-1) - r^f + g + B_x(n-1)(1 - \phi_x)\bar{x} + \frac{1}{2}V_{n-1}V'_{n-1}, \quad (11)$$

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

and

$$B_z(n) = \frac{1 - \phi_z^n}{1 - \phi_z}, \quad (13)$$

where

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

$B_x(0) = 0$, and $A(0) = 0$. This confirms the conjecture (9).³

³ The fact that price-dividend ratios are exponential affine in the state variables invites a comparison to the affine term structure literature, wherein bond prices are exponential affine in the state variables. In fact, this model is related to the essentially affine class of term structure models explored in continuous time by Dai and Singleton (2003) and Duffee (2002) and in discrete time by Ang and Piazzesi (2003). Our model is essentially affine rather than affine because the stochastic discount factor is quadratic, as a result of the homoskedastic price of risk.

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

The behavior of B_x is more complicated. In our benchmark case of $\sigma_x \sigma'_d = 0$, $B_x(n) < 0$ for all n . An increase in x_t leads to an increase in risk premia and a decrease in prices.⁴ We explore the intuition behind $B_x(n)$ further in Section III. 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 risk-free 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 that arises because we are taking the expectation of a log-normal variable.

In order to understand risk premia on more complex assets, it is helpful to understand risk premia on zero-coupon equity. Define $r_{n,t+1} = \ln R_{n,t+1}$. To gain an understanding of the model, we compute $\ln E_t[R_{n,t+1}/R^f] = E_t[r_{n,t+1} - r^f] + \frac{1}{2}\sigma_t(r_{n,t+1})\sigma_t(r_{n,t+1})'$.⁵ It follows from (9) that $r_{n,t+1}$ can be written as

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

where

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

Thus, returns are conditionally log-normally distributed, and we can rewrite the conditional Euler equation (6) as

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

Solving for the expectation and taking logs produces the relation

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

As (16) shows, 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 examine below, this term becomes important. Also determining risk premia is the loading on z_t , $B_z(n)$, and the price of z_t -risk, which is given by $\|\sigma_d\|^{-1}\sigma_z\sigma'_d x_t$. In what follows, similar reasoning can be used to understand risk premia of the aggregate market and of firms, both of which are portfolios of these underlying assets.

⁴ Alternatively, it might be the case 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.

⁵ When we match the simulated model to the data, we compute $E[R_{t+1} - R^f]$.

C. Aggregate Market

The aggregate market is the claim to all future dividends. Accordingly, its price–dividend ratio is the sum of the price to aggregate dividend ratios of zero-coupon equity. That is,

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

The Appendix 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

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

In sum, this section describes the model for the stochastic discount factor and the aggregate dividend. The following section calibrates the model and describes its implications for equity returns.

III. 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) and the price–dividend ratio for the aggregate market from (17).

We calibrate the model to the annual data set of Campbell (1999), which begins in 1890, updating Campbell’s data (which end in 1995) through the end of 2002. To ensure 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 (described below) are constructed analogously to annual price–dividend ratios in the Campbell data set: We divide the price by the current dividend plus the previous three quarters of dividends on the asset.

Section A describes the calibration of our model to the aggregate time series. Section B gives the model’s implications for the behavior of the aggregate market and dividend growth and discusses the fit to the data. Section C gives the implications for prices and returns on zero-coupon equity. While zero-coupon equity has no analogue in the data, it allows us to illustrate the properties of the model in a stark way. Section D discusses the calibration of the share process that determines the prices of long-lived assets (“firms”), and describes implications of the model for portfolios formed by sorting on scaled price ratios.

A. Calibration

Following Menzly et al. (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 risk-free 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. All variables are adjusted for inflation. The Data Appendix of Campbell (1999) contains more details on data construction.

We set r^f equal to 1.93%, the sample mean of the risk-free rate. Similarly, we set g equal to 2.28%, which is the average dividend growth in the sample. Calibrating the process z_t , which determines expected dividend growth, is less straightforward as, strictly speaking, this process is unobservable to the econometrician. However, Lettau and Ludvigson (2005) show that if consumption growth follows a random walk and if the consumption–dividend ratio is stationary, the consumption–dividend ratio captures the predictable component of dividend growth. The consumption–dividend ratio can therefore be identified with z_t up to an additive and multiplicative constant.⁶ In our annual sample, the consumption–dividend ratio has a persistence of 0.91 and a conditional correlation with dividend growth of -0.83 ; these are, respectively, our values for ϕ_z and the correlation between z_t and Δd_t . We set $\|\sigma_d\|$ to match the unconditional standard deviation of annual dividend growth in the data.⁷ 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 growth 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. Moreover, the autocorrelation of dividend growth is relatively low (-0.09). We show that $\|\sigma_z\| = 0.0016$ (0.0032 per annum) produces similar results in simulated data.

The 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 $\phi_x = 0.87^{\frac{1}{4}} = 0.966$, as 0.87 is the autocorrelation of the price–dividend ratio in annual data. We set $\|\sigma_x\|$ to 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. This produces Sharpe ratios for the cross section that are close to those in the data. Setting the maximum Sharpe ratio $\sqrt{e^{\bar{x}^2} - 1}$ equal to 0.70 implies $\bar{x} = 0.625$. As we discuss in the subsequent section, this produces an average Sharpe ratio for the market that is 0.41, which is somewhat higher than the data equivalent of 0.33. However, expected stock

⁶ An equivalent way of writing down our model would be to specify a consumption process 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.

⁷ The model is simulated at a quarterly frequency and aggregated up to an annual frequency. Because dividend growth is slightly mean reverting, and because the variance of z_t is small, this results in an unconditional annual standard deviation of dividend growth very close to that in the data.

Table IV
Parameters of the Model

Model parameters are calibrated to aggregate data starting in 1890 and ending in 2002. The model is simulated at a quarterly frequency. The unconditional mean of dividend growth g , the risk-free rate r_f , the persistence variables ϕ_x and ϕ_z , and the conditional standard deviations $\|\sigma_d\|$, $\|\sigma_z\|$, and $\|\sigma_x\|$, 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 conditional mean of dividend growth is determined by the log consumption–dividend ratio in the data. The parameter $\|\sigma_z\|$ is set to match the autocorrelation and predictability of dividend growth in the data, $\|\sigma_x\|$ is set to match the volatility of the price–dividend ratio, and ϕ_x is set to match the persistence of the price–dividend ratio.

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\ $	0.0032
$\ \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	
σ_d	[0.0724, 0, 0]
σ_z	[−0.0013, 0.0009, 0]
σ_x	[0, 0, 0.12]

returns are measured with noise, and 0.41 is still below the Sharpe ratio of post-war data.

To determine the vectors σ_d , σ_z , σ_x , we assume without loss of generality that the 3×3 matrix $[\sigma'_d, \sigma'_z, \sigma'_x]'$ is lower triangular. Thus $\epsilon_{1,t+1} = \epsilon_{d,t+1}$, so that the first element of σ_d equals $\|\sigma_d\|$ and the second and third elements equal zero. The vector σ_z has nonzero first and second elements determined by $\|\sigma_z\|$ and $\sigma_d\sigma'_z$, and zero third element. We focus on the case in which 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 IV summarizes these parameter choices.

Given our parameter choices, it is possible to infer the process for x_t based on the observed price–dividend ratio and consumption–dividend ratio. The consumption–dividend ratio can be used to construct an empirical proxy for z_t .⁸ For each time-series observation on the price–dividend ratio and z_t , we find a corresponding x_t by numerically solving (17). Figure 1 plots the resulting series for x_t , along with several macroeconomic time series that recent theory suggests should be related to aggregate risk aversion. These macroeconomic

⁸ Specifically, the consumption–dividend ratio is demeaned, divided by its standard deviation, and multiplied by the standard deviation of z_t .




Figure 1. Implied time series for x and macroeconomic variables. Macroeconomic variables are my (the deviation from the cointegration relationship between human wealth and outstanding home mortgages as in Lustig and Van Nieuwerburgh (2005)), α (the share of nonhousing consumption in total consumption as in Piazzesi, Schneider, and Tuzel (2005)), and cay (the consumption-wealth ratio of Lettau and Ludvigson (2001)). All series are demeaned and standardized. The annual data span the 1947 to 2002 period.

Table V
Results from Contemporaneous OLS Regressions
of x on Macroeconomic Variables

The variable my is the deviation from the cointegration relationship between human wealth and outstanding home mortgages as in Lustig and Van Nieuwerburgh (2005), cay is the consumption-wealth ratio of Lettau and Ludvigson (2001), and α is the share of nonhousing consumption in total consumption as in Piazzesi, Schneider, and Tuzel (2005). The annual data span the period 1947 to 2002.

	β	t -Statistics	R^2
my	2.80	6.13	0.54
cay	21.32	3.44	0.28
α	29.30	6.19	0.30

time series are: my , the deviation from the cointegration relationship between human wealth and outstanding home mortgages constructed by Lustig and Van Nieuwerburgh (2005); α , the share of non-housing consumption in total consumption constructed by Piazzesi et al. (2005); and cay , the consumption-wealth ratio of Lettau and Ludvigson (2001b). All series are demeaned and standardized. Figure 1 shows that all three series are positively correlated with x_t . Long-run fluctuations in x_t appear to be related to long-run fluctuations in both my and α , while cay (which is constructed using data on prices as well as macroeconomic quantities) also picks up short-run fluctuations in x_t .

Table V shows results of contemporaneous regressions of the implied x_t on the variables described above. This table confirms that x_t is positively and significantly related to all three macroeconomic-based risk aversion

measures.

B. Implications for the Aggregate Market and Dividend Growth

Table VI presents statistics from simulated data, and the corresponding statistics computed from actual 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 are set so that the model fits these parameters. The model produces a mean price–dividend ratio equal to 20.1, compared to 25.6 in the data. Matching this statistic is a common difficulty for models of this type: Campbell and Cochrane (1999), for example, 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 vs. 19.4% in the data), though it produces an equity premium that is slightly higher than in the data (7.9% in the model vs. 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 autocorrelation of excess returns is slightly positive (0.03). In our model, the autocorrelation is slightly negative (−0.02). The autocorrelation of dividend growth is small and negative (−0.04), just as in the data (−0.09).

Table VII 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. In our sample, as elsewhere (e.g., Campbell and Shiller (1988), Cochrane (1992), Fama and French (1989), Keim and Stambaugh (1986)), high price–dividend ratios predict low returns. The coefficients rise with the horizon. The R^2 s start small, at 0.05 at an annual horizon, and rise to 0.31 at a horizon of 10 years. The t -statistics, computed using autocorrelation- and

Table VI
Simulated Moments for the Aggregate Market and Dividend Growth

The model is simulated for 50,000 quarters. Returns, dividends, and price ratios are aggregated to an annual frequency. The data are annual and span the period 1890 to 2002.

	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%

Table VII
Long Horizon Regressions—Excess Returns

Excess returns are regressed on the lagged price–dividend ratio in annual data from 1890 to 2002 and in data simulated from the model. Specifically, we run the regression

$$\sum_{i=1}^H r_{t+i}^m - r_{t+i}^f = \beta_0 + \beta_1(p_t - d_t) + \epsilon_t$$

in the data and in the model. 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 data coefficients using the standard t -test at the 5% level are highlighted in boldface.

	Horizon in Years					
	1	2	4	6	8	10
Panel A: Full 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]
Panel B: 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]
Panel C: 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]

heteroskedasticity-adjusted standard errors, are significant at the 5% level. The simulated data exhibit the same pattern. The R^2 s start at 0.06 and rise to 0.28. We conclude that the model generates a reasonable amount of return predictability.⁹

Table VIII 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 predicted by the price–dividend ratio, contrary to what might be expected from a dividend discount model. This result also holds in our data: The coefficients from a regression of dividend growth on the price–dividend ratio are always insignificant and are accompanied by small R^2 statistics. In contrast, the consumption–dividend ratio predicts dividend growth in actual

⁹ Lettau and Ludvigson (2005) find evidence that excess returns are predictable by expected dividend growth, as well as by the price–dividend ratio. This effect can be captured in our model by allowing shocks to x_t to be positively correlated with shocks to z_t . Because introducing this positive correlation has very little effect on our cross-sectional results, for simplicity we focus on the case of zero correlation.

Table VIII
Long Horizon Regressions—Dividend Growth

Aggregate dividend growth is regressed on lagged values of the price–dividend ratio and the consumption–dividend ratio in annual data from 1890 to 2002 and in data simulated from the model. 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 data coefficients using the standard t -test at the 5% level are highlighted in boldface.

	Horizon in Years					
	1	2	4	6	8	10
Panel A: Data						
	$\sum_{i=1}^H \Delta d_{t+i} = \beta_0 + \beta_1(p_t - d_t) + \epsilon_t$					
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]
	$\sum_{i=1}^H \Delta d_{t+i} = \beta_0 + \beta_1(c_t - d_t) + \epsilon_t$					
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]
Panel B: Model						
	$\sum_{i=1}^H \Delta d_{t+i} = \beta_0 + \beta_1(p_t - d_t) + \epsilon_t$					
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]
	$\sum_{i=1}^H \Delta d_{t+i} = \beta_0 + \beta_1 z_t + \epsilon_t$					
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]

data. 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 10 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 s that range from 2% to 9% at a horizon of 10 years. By contrast, dividends are predictable by z_t . Here, the R^2 s 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.

C. Prices and Returns on Zero-Coupon Equity

Figure 2 plots the solutions for $A(n)$, $B_z(n)$, and $B_s(n)$ as a function of n for the parameter values given above. $A(n)$ is decreasing in n , as is necessary for convergence of the market price–dividend ratio. This is also sensible economically: The further the payoff is in the future, the lower the value of the security when




Figure 2. Model solution. Given the parameter values in Table IV, the top left panel shows the solution for A defined by (11), the top right panel shows the solution for B_z defined by (13) and scaled by $1 - \phi_z$, and the bottom panel shows the solution for B_x defined by (12).

the state variables are at their long-run means. What generates the decrease is the positive average price of risk $\bar{x}$ and the risk-free rate r^f , counteracted by average dividend growth g and the Jensen's inequality term.

Given that we describe the behavior of $B_z(n)$ in Section A, here we focus on $B_x(n)$. For all values of n , $B_x(n)$ is negative, indicating that an increase in the price of risk x_t leads to a decrease in valuations. Also, $B_x(n)$ is nonmonotonic, starting at zero, decreasing to below -1 , then increasing, eventually converging to a value near -0.5 . It is not surprising that $B_x(n)$ initially decreases in maturity. This is the duration effect: The longer is the maturity, the more sensitive the price is to changes in the discount rate. More curious is the fact that $B_x(n)$ rises after a maturity of 10 years. 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. Hence, 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 3 plots the ratios of price to aggregate dividends for zero-coupon equity as a function of maturity n . The top panel sets z_t to be two unconditional standard deviations ($2\|\sigma_z\|/(1 - \phi_z^2)^{1/2}$) below its unconditional mean, the middle




Figure 3. Ratios of prices to aggregate dividends for zero-coupon equity. The top panel shows ratios of prices to aggregate dividends as a function of maturity for $z = -2\|\sigma_z\|/\sqrt{1 - \phi_z^2}$, the middle panel for $z = 0$, and the lower panel for $z = 2\|\sigma_z\|/\sqrt{1 - \phi_z^2}$. Each panel shows price ratios for $x = \bar{x} - 2\|\sigma_x\|/\sqrt{1 - \phi_x^2}$, $x = \bar{x}$, and $x = \bar{x} + 2\|\sigma_x\|/\sqrt{1 - \phi_x^2}$.

panel to the unconditional mean of zero, and the bottom panel to two standard deviations above the mean. Each panel plots the price–dividend ratio for x_t at its unconditional mean and two unconditional standard deviations ($2\|\sigma_x\|/(1 - \phi_x^2)^{1/2}$) around the mean. Prices are increasing in z_t for all values of x_t and n , and decreasing in x_t for all values of z_t and n . That is, higher expected dividend growth and lower risk premia imply higher 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 the mean. In this case, 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 $\bar{x}$, 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 because z_t raises the price of long-run equity relative to short-run equity.

Figure 4 presents statistics for annual returns on zero-coupon equity. The top panel shows that the risk premium $ER_{i,t+1} - R^f$ declines with maturity.




Figure 4. Summary statistics for zero-coupon equity. The top panel shows risk premia $E[R_{nt} - R^f]$ on zero-coupon equity over the risk-free rate. The middle panel shows the standard deviation of returns on zero-coupon equity. The bottom panel shows the Sharpe ratio (the risk premium divided by the standard deviation). Returns are simulated at a quarterly frequency and aggregated to an annual frequency.

The effect is economically large: The risk premium is 18% for equity that pays a dividend 2 years from now and 4% for equity that pays a dividend 40 years from now.

The second panel of Figure 4 shows that the return volatility initially increases with maturity, and then decreases at maturities greater than 10 years. The third panel of Figure 4 shows that the unconditional Sharpe ratio decreases monotonically in maturity. These results suggest that the model has the potential to explain the patterns described in Table I. Firms that have more weight in low-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 5 shows the results of regressing simulated zero-coupon equity returns on simulated market returns. The top panel shows the regression alpha, the middle panel the beta, and the last panel the R^2 . As in Figure 4, returns are annual. The first panel shows that the alpha relative to the CAPM is decreasing in maturity over most of the range, increasing only slightly for long-duration equity. For the shortest-duration equity the alpha is as high as 11%. The alpha




Figure 5. CAPM regressions for zero-coupon equity. The top panel shows the intercept from time-series regressions of excess zero-coupon equity returns on the excess market return, the middle panel shows the slope coefficient, and the bottom panel shows the R^2 . Statistics are shown as a function of maturity. Returns are simulated at a quarterly frequency and aggregated to an annual frequency.

falls below zero for equity maturing in 5 or more years, but remains above -5% . Thus, the model produces relatively large positive alphas and relatively small negative alphas, just as in the data.

The second panel of Figure 5 shows the regression beta. The beta first increases, and then, beginning with a maturity of about 10 years, decreases slowly as a function of maturity. The betas for zero-coupon equity lie in a relatively narrow range; the lowest beta (for very long horizon equity) is about 0.7, and the highest beta (for equity of about 10 years) is 1.5. The beta for the shortest-horizon equity is about 0.9. This plot shows that at least for short-horizon equity, high alphas are not necessarily accompanied by high betas. These results suggest that the model has the potential to explain the patterns described in Table III.

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, that is, long-lived assets that have nonzero cash flows in every period.

D. Implications for the Cross Section of Returns

This section shows the implications of the model for portfolios formed by sorting on price ratios. Following Menzly et al. (2004), we exogenously specify a share process for cash flows on long-lived assets. For 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.

D.1. Specifying the Share Process

In order to assess the quantitative implications of the model, we specify long-lived assets with well-defined ratios of prices to dividends that together sum up to the market portfolio. Moreover, we require that the cross-sectional distribution of dividends, returns, and price ratios be stationary. In order to accomplish this, we follow Lynch (2003) and Menzly et al. (2004) in specifying the share each security has in the aggregate dividend process D_{t+1} . The continuous-time framework of Menzly et al. allows the authors to specify the share process as stochastic, yet still keep shares between zero and one. This is more difficult in discrete time; for this reason we adopt the simplifying assumption that the share process is deterministic.

Consider N sequences of dividend shares s_{it} , for $i = 1, \dots, N$. For convenience, we refer to each of these N sequences as a firm, though they are best thought of as portfolios of firms in the same stage of the life cycle. As our ultimate goal is to aggregate these firms into portfolios based on price-dividend ratios, this simplification does not affect our results. Firm i pays s_{it} of the aggregate dividend at time t , $s_{i,t+1}$ of the aggregate dividend at time $t + 1$, etc. Shares are such that $s_{it} \geq 0$ and $\sum_{i=1}^N s_{it} = 1$ for all t (so that the firms add up to the market). Because firm i pays a dividend sequence $s_{i,t+1} D_{t+1}, s_{i,t+2} D_{t+2}, \dots$, no-arbitrage implies that the ex-dividend price of firm i equals

$$P_{it}^F = \sum_{n=1}^{\infty} s_{i,t+n} P_{nt},$$

where P_{nt} is the price of zero-coupon equity maturing at time $t + n$.

We specify a simple model for shares. Let $\underline{s}$ be the lowest share of a firm in the economy, and assume without loss of generality that firm 1 starts at $\underline{s}$, namely $s_{11} = \underline{s}$. We assume that the share grows at a constant rate g_s until reaching $s_{1,N/2+1} = (1 + g_s)^{N/2} \underline{s}$ and then shrinks at the rate g_s until reaching $s_{1,N+1} = \underline{s}$ again. At this point the cycle repeats. All firms are ex ante identical, but are “out of phase” with one another. Firm 1 starts out at $\underline{s}$, firm 2 at $s_{21} = (1 + g_s) \underline{s}$, Firm $N/2$ at $s_{N/2,1} = (1 + g_s)^{N/2-1} \underline{s}$, and Firm N at $s_{N1} = (1 + g_s) \underline{s}$. The variable

$\bar{s}$ is such that the shares sum to one for all t .¹⁰ We set the number of firms to 200, implying a 200-quarter, or equivalently, 50-year life cycle for a firm. While this model for firms is simple and somewhat mechanical, it accomplishes our objective of creating dispersion in the timing of cash flows across firms in a straightforward way.¹¹

The parameter that determines the growth in the share process, g_s , is set to 5%, implying an annual growth rate of 20%. We choose this value so that the cross-sectional distribution of dividend growth rates in the model matches that 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 2002.¹² The top panel of Figure 6 plots the implied cross-section of average growth rates of dividends for firms in the model, as well as the cross-section of average growth rates in earnings, dividends, and cash flows in the sample. 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. The bottom panel of Figure 6 shows the distribution of firm price–dividend ratios in the model, and price ratios in the data. While the overall fit is reasonable, the model 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 dividends. It is possible to construct models that fit the dividend growth and price ratio distributions more closely by assuming growth is linearly decreasing or imposing a greater lower bound on the dividend share. As Lettau and Wachter (2005) show, the asset pricing implications of these alternative models are very similar to the present constant growth model.

D.2. Portfolio Returns

At the start of each year in the simulation, we sort firms into deciles by their price–dividend ratio. We then form equal-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. This process is not deterministic because

¹⁰ That is, $\sum_{i=1}^N s_{it} = \bar{s} + (1 + g_s)^{N/2} \bar{s} + 2 \sum_{i=1}^{N/2-1} (1 + g_s)^i \bar{s} = 1$.

¹¹ This model, like those of Gomes, Kogan, and Zhang (2003) and Menzly, Santos, and Veronesi (2004), assumes that firms are infinitely lived. Alternatively, one could specify that a firm pays dividends for one N -period cycle and then exits at the same time that a new firm enters. This entry and exit specification still implies that at any time t , the aggregate dividend is D_t ; however, the market will be a claim to only a fraction of future dividends. Allowing entry and exit in this way implies cross-sectional results that are stronger than those in the present model: In particular, alphas are smaller for growth firms and larger for value firms than when firms are infinitely lived.

¹² Adrian and Franzoni (2002), Ang and Chen (2005), and Campbell and Vuolteenaho (2004) show that value stocks have higher betas in the pre-war period, so the CAPM performs better. By matching the cross section to the post-war 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.




Figure 6. Cross-sectional distributions in the model and in the data. The top panel illustrates the distribution of annual growth rates of dividends, earnings, and cash flows across all firms for the 1952 to 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 the data simulated from the model. The bottom panel illustrates the corresponding distribution of various price multiples in actual and simulated data.

shocks have differential impacts on price–dividend ratios of firms at different stages of the life cycle.

Having sorted the firms into deciles at the beginning of each “year,” we compute statistical tests on returns over the year. The first panel of Table IX 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 I, which show corresponding results for the data. The expected excess

Table IX
Performance of Growth and Value Portfolios in the Model

In each simulation year, firms are sorted into deciles on the price–dividend ratio. Returns are calculated over the subsequent year. Intercepts and slope coefficients are from OLS time-series regressions of excess portfolio returns on the excess market return, and on the excess market return together with the return on a portfolio short the extreme growth decile and long the extreme value decile (HML).

Portfolio	G	Growth to Value								V	V–G
	1	2	3	4	5	6	7	8	9	10	10–1
Panel A: Summary Statistics											
$ER^i - R^f$	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
Panel B: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \epsilon_{it}$											
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
Panel C: $R_t^i - R_t^f = \alpha_i + \beta_i(R_t^m - R_t^f) + \gamma_i HML_t + \epsilon_{it}$											
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

return on the extreme growth portfolio is 5.0% per annum, while for the extreme value portfolio it is 10.1% per annum.¹³ A similar spread occurs in the data: The lowest book-to-market stocks have a premium of 5.7%, while the highest have a premium of 10.6%. The model generates volatilities between 19% and 17%; the volatilities for book-to-market-sorted portfolios vary between 18% and 15% in the data. Moreover, the model predicts that value portfolios have lower volatilities than growth portfolios despite their higher returns, as is the case in the data. The model predicts that the Sharpe ratio rises from 0.26 for the extreme growth portfolio to 0.58 for the extreme value portfolio. In the data, the lowest book-to-market portfolio has a Sharpe ratio of 0.32 while the highest book-to-market portfolio has a Sharpe ratio of 0.57. To summarize, the model implies that value stocks have high expected returns, low volatility, and high Sharpe ratios, just as in the data, and further, the magnitude of the difference between value and growth is comparable to that in the data.

¹³ 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 I. 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 I–III (except for standard errors).

The second panel of Table IX shows alphas and betas relative to the CAPM. Annual excess portfolio returns are regressed on excess returns on the aggregate market. Alpha, beta, and the R^2 are reported for each decile. As this panel shows, the model can replicate the classic result of Fama and French (1992): Value portfolios have positive alphas relative to the CAPM, while growth portfolios have negative alphas. Moreover, value portfolios tend to have lower betas than growth portfolios. Our model predicts alphas that rise from -2.6 for the extreme growth portfolio to 3.4 for the extreme value portfolio. In the data, the lowest book-to-market portfolio has an alpha of -1.7 , while the highest book-to-market portfolio has an alpha of 4.0 . Thus, the model generates alphas of the correct magnitude, as well as a sizable spread between value and growth. Moreover, alphas in the model are asymmetric: Growth alphas are smaller in absolute value than are value alphas, as in the data.

The third panel of Table IX shows results of regressing portfolio returns on the market return and on a high-minus-low factor (HML) equal to the return on a portfolio short the extreme growth decile and long the extreme value decile. The purpose of this test is to see whether the model analogue to the high-minus-low Fama-French factor describes the cross section of returns in the model, as it does in the data. When we add HML to the regression, the alphas are indeed two orders of magnitude smaller than the alphas relative to the CAPM.

D.3. Relation to Conditional Factor Models

The previous discussion shows that the model replicates the high expected returns, low volatility, high Sharpe ratios, and high alphas of value stocks relative to growth stocks. The model also generates testable predictions. Because only the innovation to dividends is priced, expected returns on stocks should be determined by their conditional correlation with the aggregate dividend process. According to the model, a conditional CAPM does not hold because innovations to market returns are not perfectly conditionally correlated with innovations to dividends. Moreover, a conditional dividend CAPM should provide a better fit to the cross section than a conditional CAPM.

To evaluate these predictions, we compare pricing errors for an unconditional CAPM, an unconditional dividend CAPM, a conditional CAPM, and a conditional dividend CAPM in simulated and actual data. For the simulated data, the assets are the 10 portfolios formed on dividend-price ratios described above; for the actual data the assets are the 10 value-weighted book-to-market-sorted portfolios. Theoretically, the conditioning variables should be x_t and z_t . However, because innovations in the price-dividend ratio are driven by innovations to these variables, the price-dividend ratio works well as a conditioning variable in data simulated from the model. To estimate each factor model, we solve $\min_{\delta} [g(\delta)'g(\delta)]$, where $g(\delta) = E[\delta'f_tR_t - 1]$ and R is the vector of returns. For the CAPM, $f_t = [1, R_t^m]'$; for the dividend CAPM, $f_t = [1, \Delta d_t]'$; for the conditional CAPM, $f_t = [1, R_t^m, (d_{t-1} - p_{t-1})R_t^m, d_{t-1} - p_{t-1}]'$; and for the conditional dividend CAPM, $f_t = [1, \Delta d_t, (d_{t-1} - p_{t-1})\Delta d_t, d_{t-1} - p_{t-1}]'$, where R^m is the market return, Δd_t is log dividend growth and $d_t - p_t$ is the log dividend-price ratio on

Table X
Minimized Pricing Errors in the Data and in the Model

A factor model is estimated by minimizing $g(\delta)'g(\delta)$, where $g(\delta) = E[\delta'f_t R_t - 1]$ and f_t is the vector of factors at time t . In the data, the return vector R consists of the 10 value-weighted book-to-market sorted portfolios. In the model, R consists of the 10 portfolios formed by sorting firms into deciles on the price-dividend ratio. For the CAPM, $f_t = [1, R_t^m]'$; for the dividend CAPM, $f_t = [1, \Delta d_t]'$; for the conditional CAPM, $f_t = [1, R_t^m, (d_{t-1} - p_{t-1})R_t^m, d_{t-1} - p_{t-1}]'$; and for the conditional dividend CAPM, $f_t = [1, \Delta d_t, (d_{t-1} - p_{t-1})\Delta d_t, d_{t-1} - p_{t-1}]'$, where R^m is the market return, Δd_t is log dividend growth, and $d_t - p_t$ is the log dividend-price ratio. In the column "Data-Repurchases," dividends are adjusted for share repurchases. The table reports the annualized square root of the squared average pricing errors. The monthly data span the period 1952–1 to 2003–12.

	Data	Avg. Pricing Error Data-Repurchases	Model
CAPM	1.634%	1.634%	0.571%
Dividend CAPM	1.399%	1.076%	0.266%
Cond. CAPM	0.930%	0.687%	0.033%
Cond. Dividend CAPM	0.609%	0.492%	0.014%

the market. In the data, the value-weighted CRSP portfolio is used to proxy for the market.

Table X reports the annualized square root of the squared average pricing errors for each factor model. The first column reports the results from the data, the second column reports results for which the dividend growth process and the price–dividend ratio are adjusted for repurchases as in Boudoukh et al. (2007), and the last column reports data simulated from the model. In both the data and the model, the unconditional CAPM fares the worst, with the unconditional dividend CAPM performing better. Both conditional factor models perform better than either unconditional model in the data, a finding that the model replicates. Moreover, in both the model and the data, the conditional dividend CAPM implies the lowest pricing errors of all the factor models.

IV. Model Intuition

What explains the model's ability to capture the value premium? As we suggest in Section II, the value premium arises from the differential correlations of returns on value and growth portfolios with underlying shocks.

Figure 7 plots betas from unconditional regressions of portfolio returns on the three shocks, and the R^2 from the unconditional regressions. The coefficient on the dividend shock, β_d , is positive and greater for value portfolios than for growth portfolios. The coefficient on the shock to expected dividends, β_z , is also positive but smaller for value portfolios than for growth portfolios. While 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)), it especially affects the valuations of growth stocks, which pay dividends in the distant future. Finally, all portfolios are negatively correlated with shocks to




Figure 7. Regressions of portfolio returns on fundamental shocks. The top panels show results of regressing portfolio returns on the shock to dividends $2\sigma_d\epsilon_t$, the middle panels show results for regressions 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 show results for regressions on the shocks to the price of risk $2\sigma_x\epsilon_t$ (note that the shocks to the Sharpe ratio are uncorrelated with shocks to dividends and expected dividends). Data are simulated from the model at a quarterly frequency and aggregated to an annual frequency.

the Sharpe ratio variable x_t , as indicated by a negative β_x . A positive shock to x_t raises expected returns, and thus lowers prices and realized returns. Because of the duration effect, β_x is greater in magnitude for growth portfolios. The R^2 coefficients follow the same pattern as the magnitude of the β s.¹⁴

The patterns in Figure 7 can be traced back to the properties of zero-coupon equity. Growth firms place more weight on high-duration zero-coupon claims than do value firms, and thus they inherit the sensitivity of these high-duration claims to shocks to x_t and z_t . Interestingly, β_x does not inherit the nonmonotonicity of B_x in Figure 2. This is because, all else equal, equity that pays further in the future is worth less (Figure 3). Medium-horizon equity may therefore have a greater weight than long-horizon equity, even for growth firms.

The loadings of portfolios on various shocks present an intriguing link with the empirical results of Campbell and Vuolteenaho (2004). Using the

¹⁴ The R^2 s fail to sum to one because Figure 7 plots the results from unconditional regressions rather than conditional regressions.

vector autoregression (VAR) 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 VAR; changes in expected growth rates comprise the residual variation in market returns. Relative to value firms, growth firms have high betas with respect to news about discount rates, but low betas with respect to news about dividends.

While not precisely analogous, shocks to x_t are similar in spirit to news about discount rates in the Campbell and Vuolteenaho (2004) framework. It is therefore encouraging that our model produces betas with respect to shocks to x_t that are greater in magnitude for growth firms than for value firms. The analogue to dividend growth news is less clear in our model. Campbell and Vuolteenaho compute this as a residual, but Figure 7 shows that the residual variance is not accounted for by shocks to current or expected future dividends. Rather, we find that more of the residual is accounted for by Δd_{t+1} than by z_{t+1} . Thus, it is also encouraging that value portfolios load more on shocks to Δd_{t+1} than growth portfolios.

Figure 7 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 price of risk 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 we discuss in Section III, we estimate that shocks to expected dividend growth z_t are negatively correlated with shocks to realized dividend growth. 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 are uncorrelated with shocks to realized dividends, and thus carry a zero risk price. This assumption represents a departure from the models of Campbell and Cochrane (1999) and Menzly et al. (2004), in which shocks to the price of risk are perfectly negatively correlated with shocks to aggregate dividends. What role does this assumption play in our analysis?

To answer this question, consider the conditional risk premium for equity that matures next period:

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

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

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

where $\rho_{dx} = \frac{\sigma_d \sigma'_x}{\|\sigma_d\|\|\sigma_x\|}$ represents the conditional correlation between Δd_{t+1} and x_{t+1} , and $\rho_{dz} = \frac{\sigma_d \sigma'_z}{\|\sigma_d\|\|\sigma_z\|}$ represents 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 is small because σ_z is a small fraction of σ_d . Whether long-horizon equity has a lower risk premium than short-horizon equity depends in large part on the sign of the correlation of dividend growth with x_t . In particular, $\rho_{dx} < 0$ leads to relatively high premia for long-horizon equity, while $\rho_{dx} > 0$ leads to relatively low premia for long-horizon equity.

We make this statement precise by solving the model under three different values for ρ_{dx} . Figure 8 plots risk premia on zero-coupon equity when $\rho_{dx} = -0.5$, $\rho_{dx} = 0$ (our base case), and $\rho_{dx} = 0.5$. For $\rho_{dx} = 0$, Panel B 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 $\rho_{dx} = -0.5$, Panel A shows that risk premia generally increase with maturity. Long-horizon equity (i.e., growth stocks) have greater risk premia than do short-horizon equity. This occurs even though ρ_{dz} is negative; even a modest correlation of -0.5 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 et al. (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 (2006) shows that the model of Campbell and Cochrane (1999) implies that long-horizon assets exhibit greater risk premia than do 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 is to set the correlation between d_{t+1} and x_{t+1} to be positive, as illustrated in Panel C. Under this assumption, risk premia fall more dramatically in maturity than when 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, the sign of ρ_{dx} has time-series as well as cross-sectional implications. We are able to calibrate our model to match the time series of aggregate stock returns, as well as the cross section of value and growth portfolios, because our model produces reasonable risk premia in the aggregate. For $\rho_{dx} > 0$, this may not be the case. Figure 8 shows that the greater is ρ_{dx} , the lower are risk premia in the economy, for all but the shortest-maturity equity. 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 et al. (2004) to explain both the high variance and the high premium that stocks command, with comparatively little




Figure 8. Effect of ρ_{dx} on zero-coupon equity. In each panel, risk premia are shown for x_t equal to its unconditional mean $\bar{x}$ and two standard deviations on either side of the mean. In the top panel ρ_{dx} (the correlation of dividend shocks and shocks to x_t) equals -0.5 , in the middle panel ρ_{dx} equals 0, and in the bottom panel ρ_{dx} equals 0.5.

variance in fundamentals. Faced with this tension between the time series and the cross section, we choose to set the correlation between dividend growth and x_t to zero.

In summary, this section shows that setting ρ_{dx} to zero, 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 covariance with the market. We hope that future work will reveal microeconomic foundations that determine this important parameter.

V. Conclusion

This paper proposes a parsimonious model of the stochastic discount factor that accounts for both the aggregate time-series behavior of the stock market and the relative risk and return of value and growth stocks. At the root of the model is a dividend process calibrated to match the aggregate dividend process in the data, and a stochastic discount factor with a single factor, x_t , proxying for investors' time-varying preference for risk. Time-varying preferences for risk allow the model to capture the excess volatility and return predictability that obtain 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, which also feature time-varying preferences for risk, is that x_t does not arise from fluctuations in aggregate dividends. 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 lower betas than growth stocks. In our model, growth and value stocks differ based on the timing of 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 show 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 that the correlation is zero. In contrast, external habit models assume a perfect negative correlation so that shocks to the price of risk are feared as much as, if not more than, shocks to cash flows.

Our proposed resolution of 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. As we show, such a resolution accounts for the time-series behavior of the aggregate market, the relative returns of value and growth stocks, and the failure of the capital asset pricing model to explain these returns.

Appendix: 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 $n \rightarrow \infty$. We have that B_z converges if and only if

$$|\phi_z| < 1. \quad (\text{A1})$$

Let

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

Assuming (A1) holds, B_x converges if and only if

$$|\phi_x - \sigma_x \lambda| < 1. \quad (\text{A2})$$

Given (A1),

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

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

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

Then

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

Given (A1) and (A2), it follows that

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

and

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

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,

$$A(n) \approx \bar{A}n + \text{constant}, \quad n \geq N,$$

and therefore

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

It follows that necessary and sufficient conditions for convergence are (A1), (A2), and

$$-r + g + \bar{B}_x(1 - \phi_x)\bar{x} + \frac{1}{2}\bar{V}\bar{V} < 0. \quad (\text{A3})$$

REFERENCES

- Adrian, Tobias, and Francesco Franzoni, 2002, Learning about beta: An explanation of the value premium, Working paper, MIT.
- Ang, Andrew, and Joseph Chen, 2005, CAPM over the long run: 1926–2001, *Journal of Empirical Finance* 60, 1639–1672.
- Ang, Andrew, and Jun Liu, 2004, How to discount cash flows with time-varying expected returns, *Journal of Finance* 59, 2745–2783.
- 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, Gurdip S., and Zhiwu 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, Anomalies in relationships between securities' yields and yield-surrogates, *Journal of Financial Economics* 6, 103–126.
- Bansal, Ravi, Robert F. Dittmar, and Christian T. Lundblad, 2005, Consumption, dividends, and the cross-section of equity returns, *Journal of Finance* (forthcoming).
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- 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, 663–682.
- 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.
- Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.
- Berk, Jonathan B., Richard C. Green, and Vasant Naik, 2004, Valuation and return dynamics of new ventures, *Review of Financial Studies* 17, 1–35.
- Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 63 (forthcoming).
- Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2004, Estimation and test of a simple model of intertemporal asset pricing, *Journal of Finance* 59, 1743–1775.
- Brennan, Michael J., and Yihong Xia, 2006, Risk and valuation under an intertemporal capital asset pricing model, *Journal of Business* 79, 1–35.
- 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? Working 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, 2004, Bad beta, good beta, *American Economic Review* 94, 1249–1275.
- Cochrane, John H., 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243–280.
- Cochrane, John H., 1999, New facts in finance, *Economic Perspectives* Federal Reserve Bank of Chicago 23, 59–78.

- Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, The price is (almost) right, Working 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, 2004, Implied equity duration: A new measure of equity risk, *Review of Accounting Studies* 9, 197–228.
- 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).
- Hansen, Lars Peter, John C. Heaton, and Nan Li, 2004, Consumption strikes back, Working paper, University of Chicago.
- Jaffe, Jeffrey, Donald B. Keim, and Randolph Westerfield, 1989, Earnings yields, market values, and stock returns, *The Journal of Finance* 44, 134–148.
- Jagannathan, Ravi, and Zhenyu Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–54.
- Johnson, Timothy C., 2002, Rational momentum effects, *Journal of Finance* 57, 585–608.
- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Econometrics* 17, 357–390.
- Leibowitz, Martin L., and Stanley Kogelman, 1993, Resolving the equity duration paradox, *Financial Analysts Journal* January–February, 51–64.
- Lettau, Martin, and Sydney C. Ludvigson, 2001a, 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, 2001b, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–850.
- Lettau, Martin, and Sydney C. Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics* 76, 583–626.
- Lettau, Martin, and Jessica A. Wachter, 2005, Why is long-horizon equity less risky? A duration-based explanation of the value premium, NBER working paper # 11144.
- 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, John, 1965, Security prices, risk and maximal gains from diversification, *Journal of Finance* 20, 587–615.
- Lustig, Hanno, and Stijn Van Nieuwerburgh, 2005, Housing collateral, consumption insurance and risk premia: An empirical perspective, *Journal of Finance* 60, 1167–1219.
- Lynch, Anthony W., 2003, Portfolio choice with many risky assets, market clearing, and cash-flow predictability, Working paper, New York University.
- Menzly, Lior, Tano Santos, and Pietro Veronesi, 2004, Understanding predictability, *Journal of Political Economy* 112, 1–47.
- Newey, Whitney, and Kenneth West, 1987, A simple, positive-definite, heteroskedastic and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Parker, Jonathan A., and Christian Julliard, 2005, Consumption risk and the cross section of expected returns, *Journal of Political Economy* 113, 185–222.

- Petkova, Ralitsa, and Lu Zhang, 2005, Is value riskier than growth? *Journal of Financial Economics* 78, 187–202.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2005, Housing, consumption, and asset pricing, *Journal of Financial Economics* (forthcoming).
- 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, Working paper, Columbia University and University of Chicago.
- Santos, Tano, and Pietro Veronesi, 2006, Labor income and predictable stock returns, *Review of Financial Studies* 19, 1–44.
- 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., 2006, A consumption-based model of the term structure of interest rates, *Journal of Financial Economics* 79, 365–399.
- Wilson, Mungo, 2003, A model of demand for dividend strips, Working paper, Harvard University.
- Yogo, Motohiro, 2005, A consumption-based explanation of the cross section of expected stock returns, *Journal of Finance* 61, 539–580.
- Zhang, Lu, 2005, The value premium, *Journal of Finance* 60, 67–103.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.