Bad Beta, Good Beta

Citation

Campbell, John Y., and Tuomo Vuolteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* 94(5): 1249-1275.

Published Version

<http://dx.doi.org/10.1257/0002828043052240>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3122489>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

HIER

Harvard Institute of Economic Research

Discussion Paper Number 2016

Bad Beta, Good Beta

by

John Y. Campbell
and
Tuomo Vuolteenaho

August 2003

Harvard University
Cambridge, Massachusetts

This paper can be downloaded without charge from:

<http://post.economics.harvard.edu/hier/2003papers/2003list.html>

The Social Science Research Network Electronic Paper Collection:

<http://ssrn.com/abstract=343780>

Bad Beta, Good Beta

John Y. Campbell and Tuomo Vuolteenaho¹

First draft: August 2002

This draft: August 2003

¹Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, USA, and NBER. Email john_campbell@harvard.edu and t_vuolteenaho@harvard.edu. We would like to thank Michael Brennan, Joseph Chen, Randy Cohen, Robert Hodrick, Matti Keloharju, Owen Lamont, Greg Mankiw, Lubos Pastor, Antti Petajisto, Christopher Polk, Jay Shanken, Andrei Shleifer, Jeremy Stein, Sam Thompson, Luis Viceira, and seminar participants at Chicago GSB, Harvard Business School, the Kellogg School, and the NBER Asset Pricing meeting for helpful comments. We are grateful to Ken French for providing us with some of the data used in this study. All errors and omissions remain our responsibility. This material is based upon work supported by the National Science Foundation under Grant No. 0214061 to Campbell.

Abstract

This paper explains the size and value “anomalies” in stock returns using an economically motivated two-beta model. We break the CAPM beta of a stock with the market portfolio into two components, one reflecting news about the market’s future cash flows and one reflecting news about the market’s discount rates. Intertemporal asset pricing theory suggests that the former should have a higher price of risk; thus beta, like cholesterol, comes in “bad” and “good” varieties. Empirically, we find that value stocks and small stocks have considerably higher cash-flow betas than growth stocks and large stocks, and this can explain their higher average returns. The poor performance of the CAPM since 1963 is explained by the fact that growth stocks and high-past-beta stocks have predominantly good betas with low risk prices.

JEL classification: G12, G14, N22

1 Introduction

How should a rational investor measure the risks of stock market investments? What determines the risk premium that will induce a rational investor to hold an individual stock at its market weight, rather than overweighting or underweighting it? According to the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965), a stock's risk is summarized by its beta with the market portfolio of all invested wealth. Controlling for beta, no other characteristics of a stock should influence the return required by a rational investor.

It is well known that the CAPM fails to describe average realized stock returns since the early 1960's, if a value-weighted equity index is used as a proxy for the market portfolio. In particular, small stocks and value stocks have delivered higher average returns than their betas can justify. Adding insult to injury, stocks with high past betas have had average returns no higher than stocks of the same size with low past betas. These findings tempt investors to tilt their stock portfolios systematically towards small stocks, value stocks, and stocks with low past betas.²

We argue that returns on the market portfolio have two components, and that recognizing the difference between these two components can eliminate the incentive to overweight value, small, and low-beta stocks. The value of the market portfolio may fall because investors receive bad news about future cash flows; but it may also fall because investors increase the discount rate or cost of capital that they apply to these cash flows. In the first case, wealth decreases and investment opportunities are unchanged, while in the second case, wealth decreases but future investment opportunities improve.

These two components should have different significance for a risk-averse, long-term investor who holds the market portfolio. Such an investor may demand a higher premium to hold assets that covary with the market's cash-flow news than to hold assets that covary with news about the market's discount rates, for poor returns driven by increases in discount rates are partially compensated by improved prospects for future returns. To properly measure risk for this investor, the single

²Seminal early references include Banz (1981) and Reinganum (1981) for the size effect, and Graham and Dodd (1934), Basu (1977, 1983), Ball (1978), and Rosenberg, Reid, and Lanstein (1985) for the value effect. Fama and French (1992) give an influential treatment of both effects within an integrated framework and show that sorting stocks on past market betas generates little variation in average returns.

beta of the Sharpe-Lintner CAPM should be broken into two different betas: a cash-flow beta and a discount-rate beta. We expect a rational investor who is holding the market portfolio to demand a greater reward for bearing the former type of risk than the latter. In fact, an intertemporal capital asset pricing model (ICAPM) of the sort proposed by Merton (1973) suggests that the price of risk for the discount-rate beta should equal the variance of the market return, while the price of risk for the cash-flow beta should be γ times greater, where γ is the investor's coefficient of relative risk aversion. If the investor is conservative in the sense that $\gamma > 1$, the cash-flow beta has a higher price of risk.

An intuitive way to summarize our story is to say that beta, like cholesterol, has a “bad” variety and a “good” variety. The required return on a stock is determined not by its overall beta with the market, but by its bad cash-flow beta and its good discount-rate beta. Of course, the good beta is good not in absolute terms, but in relation to the other type of beta.

We test these ideas by fitting a two-beta ICAPM to historical monthly returns on stock portfolios sorted by size, book-to-market ratios, and market betas. We consider not only a sample period since 1963 that has been the subject of much recent research, but also an earlier sample period 1929-1963 using the data of Davis, Fama, and French (2000). In the modern period, 1963:7-2001:12, we find that the two-beta model greatly improves the poor performance of the standard CAPM. The main reason for this is that growth stocks, with low average returns, have high betas with the market portfolio; but their high betas are predominantly good betas, with low risk prices. Value stocks, with high average returns, have higher bad betas than growth stocks do. In the early period, 1929:1-1963:6, we find that value stocks have higher CAPM betas and proportionately higher bad betas than growth stocks, so the single-beta CAPM adequately explains the data.

The ICAPM also explains the size effect. Over both subperiods, small stocks outperform large stocks by approximately 3% per annum. In the early period, this performance differential is justified by the moderately higher cash-flow and discount-rate betas of small stocks relative to large stocks. In the modern period, small and large stocks have approximately equal cash-flow betas. However, small stocks have much higher discount-rate betas than large stocks in the post-1963 sample. Even though the premium on discount-rate beta is low, the magnitude of the beta spread is sufficient to explain most of the size premium.

Our two-beta model also casts light on why portfolios sorted on past CAPM betas

show a spread in average returns in the early sample period but not in the modern period. In the early sample period, a sort on CAPM beta induces a strong post-ranking spread in cash-flow betas, and this spread carries an economically significant premium, as the theory predicts. In the modern period, however, sorting on past CAPM betas produces a spread only in good discount-rate betas but no spread in bad cash-flow betas. Since the good beta carries only a low premium, the almost flat relation between average returns and the CAPM beta estimated from these portfolios in the modern period is no puzzle to the two-beta model.

All these findings are based on the first-order condition of a long-term investor who is assumed to hold a value-weighted stock market index. We show that there exists a coefficient of risk aversion that makes the investor content to hold equities at their value weights, rather than systematically tilting her portfolio towards value stocks, small stocks, or stocks with low past betas. For an investor with this degree of risk aversion, the high average returns on such stocks are appropriate compensation for their risks in relation to the value-weighted index. An investor with a lower risk aversion coefficient would find value, small, and low-past-beta stocks attractive and would wish to overweight them, while an investor with a higher risk aversion coefficient would wish to underweight these stocks.

Our model explains why stocks with high cash-flow betas may offer high average returns, given that long-term investors are fully invested in equities at all times, or, in a slight generalization of the model, maintain a constant allocation to equities. Our model does not explain why long-term investors would wish to keep their equity allocations constant. If the equity premium is time-varying, it is optimal for a long-term investor with a fixed coefficient of relative risk aversion to invest more in equities at times when the equity premium is high (Campbell and Viceira 1999, Kim and Omberg 1996). We could generalize the model to allow a time-varying equity weight in the investor's portfolio, but this would not be consistent with general equilibrium if all investors have the same preferences. Thus our model cannot be interpreted as a representative agent general equilibrium model of the economy. Our achievement is merely to show that the prices of risk for value, small, and low-past-beta stocks are sufficient to deter investment in these stocks by conservative long-term investors who eschew market timing.

In developing and testing the two-beta ICAPM, we draw on a great deal of related literature. The idea that the market's return can be attributed to cash-flow and discount-rate news is not novel. Campbell and Shiller (1988a) develop a loglin-

ear approximate framework in which to study the effects of changing cash-flow and discount-rate forecasts on stock prices. Campbell (1991) uses this framework and a vector autoregressive (VAR) model to decompose market returns into cash-flow news and discount-rate news. Empirically, he finds that discount-rate news is far from negligible; in postwar US data, for example, his VAR system explains most stock return volatility as the result of discount-rate news. Campbell and Mei (1993) use a similar approach to decompose the market betas of industry and size portfolios into cash-flow betas and discount-rate betas, but they do not estimate separate risk prices for these betas.

The insight that long-term investors care about shocks to investment opportunities is due to Merton (1973). Campbell (1993) solves a discrete-time empirical version of Merton's ICAPM, assuming that asset returns are homoskedastic and that a representative investor has the recursive preferences proposed by Epstein and Zin (1989, 1991). The solution is exact in the limit of continuous time if the representative investor has elasticity of intertemporal substitution equal to one, and is otherwise a loglinear approximation. Campbell writes the solution in the form of a K -factor model, where the first factor is the market return and the other factors are shocks to variables that predict the market return. Campbell (1996) also tests this model on industry portfolios, but finds that in his specification the innovation to discount rates is highly correlated with the innovation to the market itself; thus his multi-beta model is hard to distinguish empirically from the CAPM. Li (1997), Hodrick, Ng, and Sengmueller (1999), Lynch (1999), Brennan, Wang, and Xia (2001, 2003), Ng (2002), Guo (2002), and Chen (2003) also explore the empirical implications of Merton's model.

The two papers that are closest to ours in their focus are Brennan, Wang, and Xia (2003) and Chen (2003). Brennan et al. model the riskless interest rate and the Sharpe ratio on the market portfolio as continuous-time AR(1) processes. They estimate the parameters of their model using bond market data, and explore the model's implications for the value and size effects in US equities since 1953. They have some success in explaining these effects, but they do not relate the risk prices for interest rate and Sharpe ratio shocks to the underlying preferences of investors. Chen (2003) extends the framework of Campbell (1993) to allow for heteroskedastic asset returns, and estimates a VAR-GARCH model to describe the dynamics of stock returns. Given the variables he includes in his model, he finds little evidence that growth stocks are valuable hedges against shocks to investment opportunities.

Recently, however, several authors have found that high returns to growth stocks, particularly small growth stocks, seem to predict low returns on the aggregate stock market. Eleswarapu and Reinganum (2003) use lagged 3-year returns on an equal-weighted index of growth stocks, while Brennan, Wang, and Xia (2001) use the difference between the log book-to-market ratios of small growth stocks and small value stocks to predict the aggregate market. In this paper we use a measure similar to that of Brennan et al. (2001) and find that indeed growth stock returns have high covariances with declines in market discount rates.

It is natural to ask why high returns on small growth stocks should predict low returns on the stock market as a whole. This is a particularly important question since time-series regressions of aggregate stock returns on arbitrary predictor variables can easily produce meaningless data-mined results. One possibility is that small growth stocks generate cash flows in the more distant future and therefore their prices are more sensitive to changes in discount rates, just as coupon bonds with a high duration are more sensitive to interest-rate movements than are bonds with a low duration (Cornell 1999). Another possibility is that small growth companies are particularly dependent on external financing and thus are sensitive to equity market and broader financial conditions (Ng, Engle, and Rothschild 1992, Perez-Quiros and Timmermann 2000). A third possibility is that episodes of irrational investor optimism (Shiller 2000) have a particularly powerful effect on small growth stocks.

Our finding that value stocks have higher cash-flow betas than growth stocks is consistent with the empirical results of Cohen, Polk, and Vuolteenaho (2002). Cohen et al. measure cash-flow betas by regressing the multi-year return on equity (ROE) of value and growth stocks on the market's multi-year ROE. They find that value stocks have higher ROE betas than growth stocks. There is also evidence that value stock returns are correlated with shocks to GDP-growth forecasts (Liew and Vassalou 2000, Vassalou 2003). These empirical findings are consistent with Brainard, Shapiro, and Shoven's (1991) suggestion that "fundamental betas" estimated from cash flows could improve the empirical performance of the CAPM. This sensitivity of value stocks' cash-flow fundamentals to economy-wide cash-flow fundamentals plays a key role in our two-beta model's ability to explain the value premium.

There are numerous competing explanations for the size and value effects. At the most basic level the Arbitrage Pricing Theory (APT) of Ross (1976) allows any pervasive source of common variation to be a priced risk factor. Fama and French

(1993) show that small stocks and value stocks tend to move together as groups, and introduce an influential three-factor model, including a market factor, size factor, and value factor, to describe the size and value effects in average returns. As Fama and French recognize, ultimately this falls short of a satisfactory explanation because the APT is silent about what determines factor risk prices; in a pure APT model the size premium and the value premium could just as easily be zero or negative.

Jagannathan and Wang (1996) point out that the CAPM might hold conditionally, but fail unconditionally. If some stocks have high market betas at times when the market risk premium is high, then these stocks should have higher average returns than are explained by their unconditional market betas. Lettau and Ludvigson (2001) and Zhang and Petkova (2002) argue that value stocks satisfy these conditions, although Lewellen and Nagel (2003) argue that time-varying betas cause only a very modest increase in average returns.

Adrian and Franzoni (2002) and Lewellen and Shanken (2002) consider the possibility that investors do not know the risk characteristics of stocks but must learn about them over time. Adrian and Franzoni, for example, suggest that investors tended to overestimate the market betas of value and small stocks as these betas trended downwards during the 20th Century. This led investors to demand higher average returns for such stocks than are justified by their average market risks.

Roll (1977) emphasizes that tests of the CAPM are misspecified if one cannot measure the market portfolio correctly. While Stambaugh (1982) and Shanken (1987) find that CAPM tests are insensitive to the inclusion of other financial assets, more recent research has stressed the importance of human wealth whose return can be proxied by revisions in expected future labor income (Campbell 1996, Jagannathan and Wang 1996, Lettau and Ludvigson 2001).

Finally, the value effect can also be interpreted in behavioral terms. Lakonishok, Shleifer, and Vishny (1994), for example, argue that investors irrationally extrapolate past earnings growth and thus overvalue companies that have performed well in the past. These companies have low book-to-market ratios and subsequently underperform once their earnings growth disappoints investors. Supporting evidence is provided by La Porta (1996), who shows that high long-term earnings forecasts of stock market analysts predict low stock returns while low forecasts predict high returns, and by La Porta et al. (1997), who show that the underperformance of stocks with low book-to-market ratios is concentrated on earnings announcement dates. Brav, Lehavy, and Michaely (2002) show that analysts' price targets imply high subjec-

tive expected returns on growth stocks, consistent with the hypothesis that the value effect is due to expectational errors.

In this paper we do not consider any of these alternative stories. We assume that unconditional betas are adequate proxies for conditional betas, we use a value-weighted index of common stocks as a proxy for the market portfolio, and we test an orthodox asset pricing model based on the first-order conditions of a rational investor who knows the parameters of the model. Our purpose is to clarify the extent to which deviations from the CAPM's cross-sectional predictions can be rationalized by Merton's (1973) intertemporal hedging considerations that are relevant for long-term investors. This exercise should be of interest even if one believes that investor irrationality has an important effect on stock prices, because even in this case one should want to know how a rational investor will perceive stock market risks. Our analysis has obvious relevance to long-term institutional investors such as pension funds, which maintain stable allocations to equities and wish to assess the risks of tilting their equity portfolios towards particular types of stocks.

The organization of the paper is as follows. In Section 2, we estimate two components of the return on the aggregate stock market, one caused by cash-flow shocks and the other by discount-rate shocks. In Section 3, we use these components to estimate cash-flow and discount-rate betas for portfolios sorted on firm characteristics and risk loadings. In Section 4, we lay out the intertemporal asset pricing theory that justifies different risk premia for bad cash-flow beta and good discount-rate beta. We also show that the returns to small and value stocks can largely be explained by allowing different risk premia for these two different betas. Section 5 concludes.

2 How cash-flow and discount-rate news move the market

A simple present-value formula points to two reasons why stock prices may change. Either expected cash flows change, discount rates change, or both. In this section, we empirically estimate these two components of unexpected return for a value-weighted stock market index. Consistent with findings of Campbell (1991), the fitted values suggest that over our sample period (1929:1-2001:12) discount-rate news causes much more variation in monthly stock returns than cash-flow news.

2.1 Return-decomposition framework

Campbell and Shiller (1988a) develop a loglinear approximate present-value relation that allows for time-varying discount rates. They do this by approximating the definition of log return on a dividend-paying asset, $r_{t+1} \equiv \log(P_{t+1} + D_{t+1}) - \log(P_t)$, around the mean log dividend-price ratio, $(\overline{d_t - p_t})$, using a first-order Taylor expansion. Above, P denotes price, D dividend, and lower-case letters log transforms. The resulting approximation is $r_{t+1} \approx k + \rho p_{t+1} + (1 - \rho)d_{t+1} - p_t$, where ρ and k are parameters of linearization defined by $\rho \equiv 1 / (1 + \exp(\overline{d_t - p_t}))$ and $k \equiv -\log(\rho) - (1 - \rho)\log(1/\rho - 1)$. When the dividend-price ratio is constant, then $\rho = P/(P + D)$, the ratio of the ex-dividend to the cum-dividend stock price. The approximation here replaces the log sum of price and dividend with a weighted average of log price and log dividend, where the weights are determined by the average relative magnitudes of these two variables.

Solving forward iteratively, imposing the “no-infinite-bubbles” terminal condition that $\lim_{j \rightarrow \infty} \rho^j(d_{t+j} - p_{t+j}) = 0$, taking expectations, and subtracting the current dividend, one gets

$$p_t - d_t = \frac{k}{1 - \rho} + E_t \sum_{j=0}^{\infty} \rho^j [\Delta d_{t+1+j} - r_{t+1+j}], \quad (1)$$

where Δd denotes log dividend growth. This equation says that the log price-dividend ratio is high when dividends are expected to grow rapidly, or when stock returns are expected to be low. The equation should be thought of as an accounting identity rather than a behavioral model; it has been obtained merely by approximating an identity, solving forward subject to a terminal condition, and taking expectations. Intuitively, if the stock price is high today, then from the definition of the return and the terminal condition that the dividend-price ratio is non-explosive, there must either be high dividends or low stock returns in the future. Investors must then expect some combination of high dividends and low stock returns if their expectations are to be consistent with the observed price.

While Campbell and Shiller (1988a) constrain the discount coefficient ρ to values determined by the average log dividend yield, ρ has other possible interpretations as well. Campbell (1993, 1996) links ρ to the average consumption-wealth ratio. In effect, the latter interpretation can be seen as a slightly modified version of the former. Consider a mutual fund that reinvests dividends and a mutual-fund investor

who finances her consumption by redeeming a fraction of her mutual-fund shares every year. Effectively, the investor's consumption is now a dividend paid by the fund and the investor's wealth (the value of her remaining mutual fund shares) is now the ex-dividend price of the fund. Thus, we can use (1) to describe a portfolio strategy as well as an underlying asset and let the average consumption-wealth ratio generated by the strategy determine the discount coefficient ρ , provided that the consumption-wealth ratio implied by the strategy does not behave explosively.

Campbell (1991) extends the loglinear present-value approach to obtain a decomposition of returns. Substituting (1) into the approximate return equation gives

$$\begin{aligned} r_{t+1} - E_t r_{t+1} &= (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \\ &= N_{CF,t+1} - N_{DR,t+1}, \end{aligned} \quad (2)$$

where N_{CF} denotes news about future cash flows (i.e., dividends or consumption), and N_{DR} denotes news about future discount rates (i.e., expected returns). This equation says that unexpected stock returns must be associated with changes in expectations of future cash flows or discount rates. An increase in expected future cash flows is associated with a capital gain today, while an increase in discount rates is associated with a capital loss today. The reason is that with a given dividend stream, higher future returns can only be generated by future price appreciation from a lower current price.

These return components can also be interpreted as permanent and transitory shocks to wealth. Returns generated by cash-flow news are never reversed subsequently, whereas returns generated by discount-rate news are offset by lower returns in the future. From this perspective it should not be surprising that conservative long-term investors are more averse to cash-flow risk than to discount-rate risk.

2.2 Implementation with a VAR model

We follow Campbell (1991) and estimate the cash-flow-news and discount-rate-news series using a vector autoregressive (VAR) model. This VAR methodology first estimates the terms $E_t r_{t+1}$ and $(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j}$ and then uses r_{t+1} and equation (2) to back out the cash-flow news. This practice has an important advantage – one

does not necessarily have to understand the short-run dynamics of dividends. Understanding the dynamics of expected returns is enough.

We assume that the data are generated by a first-order VAR model

$$z_{t+1} = a + \Gamma z_t + u_{t+1}, \quad (3)$$

where z_{t+1} is a m -by-1 state vector with r_{t+1} as its first element, a and Γ are m -by-1 vector and m -by- m matrix of constant parameters, and u_{t+1} an i.i.d. m -by-1 vector of shocks. Of course, this formulation also allows for higher-order VAR models via a simple redefinition of the state vector to include lagged values.

Provided that the process in equation (3) generates the data, $t + 1$ cash-flow and discount-rate news are linear functions of the $t + 1$ shock vector:

$$\begin{aligned} N_{CF,t+1} &= (e1' + e1'\lambda) u_{t+1} \\ N_{DR,t+1} &= e1'\lambda u_{t+1}. \end{aligned} \quad (4)$$

The VAR shocks are mapped to news by λ , defined as $\lambda \equiv \rho\Gamma(I - \rho\Gamma)^{-1}$. $e1'\lambda$ captures the long-run significance of each individual VAR shock to discount-rate expectations. The greater the absolute value of a variable's coefficient in the return prediction equation (the top row of Γ), the greater the weight the variable receives in the discount-rate-news formula. More persistent variables should also receive more weight, which is captured by the term $(I - \rho\Gamma)^{-1}$.

2.3 VAR data

To operationalize the VAR approach, we need to specify the variables to be included in the state vector. We opt for a parsimonious model with the following four state variables. First, the excess log return on the market (r_M^e) is the difference between the log return on the Center for Research in Securities Prices (CRSP) value-weighted stock index (r_M) and the log risk-free rate. The risk-free-rate data are constructed by CRSP from Treasury bills with approximately three month maturity.

Second, the term yield spread (TY) is provided by Global Financial Data and is computed as the yield difference between ten-year constant-maturity taxable bonds and short-term taxable notes, in percentage points.

Third, the price-earnings ratio (PE) is from Shiller (2000), constructed as the price of the S&P 500 index divided by a ten-year trailing moving average of aggregate earnings of companies in the S&P 500 index. Following Graham and Dodd (1934), Campbell and Shiller (1988b, 1998) advocate averaging earnings over several years to avoid temporary spikes in the price-earnings ratio caused by cyclical declines in earnings. We avoid any interpolation of earnings in order to ensure that all components of the time- t price-earnings ratio are contemporaneously observable by time t . The ratio is log transformed.

Fourth, the small-stock value spread (VS) is constructed from the data made available by Professor Kenneth French on his web site.³ The portfolios, which are constructed at the end of each June, are the intersections of two portfolios formed on size (market equity, ME) and three portfolios formed on the ratio of book equity to market equity (BE/ME). The size breakpoint for year t is the median NYSE market equity at the end of June of year t . BE/ME for June of year t is the book equity for the last fiscal year end in $t - 1$ divided by ME for December of $t - 1$. The BE/ME breakpoints are the 30th and 70th NYSE percentiles.

At the end of June of year t , we construct the small-stock value spread as the difference between the $\log(BE/ME)$ of the small high-book-to-market portfolio and the $\log(BE/ME)$ of the small low-book-to-market portfolio, where BE and ME are measured at the end of December of year $t - 1$. For months from July to May, the small-stock value spread is constructed by adding the cumulative log return (from the previous June) on the small low-book-to-market portfolio to, and subtracting the cumulative log return on the small high-book-to-market portfolio from, the end-of-June small-stock value spread.

Our small-stock value spread is similar to variables constructed by Asness, Friedman, Krail, and Liew (2000), Cohen, Polk, and Vuolteenaho (2003), and Brennan, Wang, and Xia (2001). Asness et al. use a number of different scaled-price variables to construct their measures, and also incorporate analysts' earnings forecasts into their model. Cohen et al. use the entire CRSP universe instead of small-stock portfolios to construct their value-spread variable. Brennan et al.'s small-stock value-spread variable is equal to ours at the end of June of each year, but the intra-year values differ because Brennan et al. interpolate the intra-year values of BE using year t and year $t + 1$ BE values. We do not follow their procedure because we wish to avoid using any future variables that might cause spurious forecastability of stock

³http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

returns.

These state-variable series span the period 1928:12–2001:12. Table 1 shows descriptive statistics and Figure 1 the time-series evolution of the state variables, excluding the market return. The variables in Figure 1 are demeaned and normalized by their sample standard deviation.

The black solid line in Figure 1 plots the evolution of PE , the log ratio of price to ten-year moving average of earnings. Our sample period begins only months before the stock market crash of 1929. This event is clearly visible from the graph in which the log price-earnings drops by an extraordinary five sample standard deviations from 1929 to 1932. Another striking episode is the 1983–1999 bull market, during which the price-earnings ratio increases by four sample standard deviations.

While the price-earnings ratio and its historical time-series behavior are well known, the history of the small-stock value spread is perhaps less so. Recall that our value-spread variable is the difference between value stocks’ log book-to-market ratio and growth stocks’ log book-to-market ratio. Thus a high value spread is associated with high prices for growth stocks relative to value stocks. Similar to figures shown by Cohen, Polk, and Vuolteenaho (2003) and Brennan, Wang, and Xia (2001), the post-war variation in VS appears positively correlated with the price-earnings ratio, high overall stock prices coinciding with especially high prices for growth stocks. The pre-war data appear quite different from the post-war data, however. For the first two decades of our sample, the value spread is negatively correlated with the market’s price-earnings ratio. The correlation between VS and PE is $-.48$ in the period 1928:12–1963:6, and $.57$ in the period 1963:7–2001:12. If most value stocks were highly levered and financially distressed during and after the Great Depression, it makes sense that their values were especially sensitive to changes in overall economic prospects, including the cost of capital. In the post-war period, however, most value stocks were probably stable businesses with relatively low financial leverage, no growth options, and thus probably little dependence on external equity-market financing. We will return to this changing sensitivity of value and growth stocks to various economy-wide shocks in Section 3.

The term yield spread (TY) is a variable that is known to track the business cycle, as discussed by Fama and French (1989). The term yield spread is very volatile during the Great Depression and again in the 1970’s. It also tracks the value spread closely, with a correlation of $.42$ over the full sample as shown in Table 1. Because long-bond yields are relatively stable, TY is mostly driven by the volatile short end of the term

structure, making the variable negatively correlated with the overall level of interest rates. Since growth stocks are assets with a high duration, as emphasized by Cornell (1999), it is not surprising that high prices for growth stocks coincide with low interest rates and thus a high term yield spread.

2.4 VAR parameter estimates

Table 2 reports parameter estimates for the VAR model. Each row of the table corresponds to a different equation of the model. The first five columns report coefficients on the five explanatory variables: a constant, and lags of the excess market return, term yield spread, price-earnings ratio, and small-stock value spread. OLS standard errors are reported in square brackets below the coefficients. For comparison, we also report in parentheses standard errors from a bootstrap exercise. Finally, we report the R^2 and F statistics for each regression. The bottom of the table reports the correlation matrix of the equation residuals, with standard deviations of each residual on the diagonal.

The first row of Table 2 shows that all four of our VAR state variables have some ability to predict excess returns on the aggregate stock market. Market returns display a modest degree of momentum; the coefficient on the lagged excess market return is .094 with a standard error of .034. The term yield spread positively predicts the market return, consistent with the findings of Keim and Stambaugh (1986), Campbell (1987), and Fama and French (1989). The smoothed price-earnings ratio negatively predicts the return, consistent with Campbell and Shiller (1988b, 1998) and related work using the aggregate dividend-price ratio (Rozeff 1984, Campbell and Shiller 1988a, and Fama and French 1988, 1989). The small-stock value spread negatively predicts the return, consistent with Eleswarapu and Reinganum (2003) and Brennan, Wang, and Xia (2001). Overall, the R^2 of the return forecasting equation is about 2.6%, which is a reasonable number for a monthly model.

The remaining rows of Table 2 summarize the dynamics of the explanatory variables. The term spread is approximately an AR(1) process with an autoregressive coefficient of .88, but the lagged small-stock value spread also has some ability to predict the term spread. This should not be surprising given the contemporaneous correlation of these two variables illustrated in Figure 1. The price-earnings ratio is highly persistent, with a root very close to unity, but it is also predicted by the lagged market return. This predictability may reflect short-term momentum in stock

returns, but it may also reflect the fact that the recent history of returns is correlated with earnings news that is not yet reflected in our lagged earnings measure. Finally, the small-stock value spread is also a highly persistent AR(1) process.

The persistence of the VAR explanatory variables raises some difficult statistical issues. It is well known that estimates of persistent AR(1) coefficients are biased downwards in finite samples, and that this causes bias in the estimates of predictive regressions for returns if return innovations are highly correlated with innovations in predictor variables (Stambaugh 1999). There is an active debate about the effect of this on the strength of the evidence for return predictability (Ang and Bekaert 2001, Campbell and Yogo 2002, Lewellen 2003, Torous, Valkanov, and Yan 2003).

For our sample and VAR specification, the four predictive variables in the return prediction equation are jointly significant at a better than 5% level. Our unreported experiments show that the joint significance of the return-prediction equation at 5% level survives bootstrapping excess returns as return shocks and simulating from a system estimated under the null with various bias adjustments. However, the statistical significance of the one-period return-prediction equation does not guarantee that our news terms are not materially affected by the above-mentioned small-sample bias.

As a simple way to assess the impact of this bias, we have generated 2500 artificial data series using the estimated VAR coefficients and have reestimated the VAR system 2500 times. The difference between the average coefficient estimates in the artificial data and the original VAR estimates is a simple measure of finite-sample bias. We find that there is some bias in the VAR coefficients, but it does not have a large effect on our estimates of cash-flow and discount-rate news. The reason is that the bias causes some overstatement of short-term return predictability (the $e1'\rho\Gamma$ component of $e1'\lambda$) but an understatement of the persistence of the VAR, and thus an understatement of the long-term impact of predictability [the $(I - \rho\Gamma)^{-1}$ component of $e1'\lambda$]. These two effects work against each other. The one variable that is moderately affected by bias is the value spread, whose role in predicting returns is biased downwards. Since this bias works against us in explaining the average returns on value and growth stocks, we do not attempt to correct it. Instead we use the estimated VAR as a reasonable representation of the data and ask what it implies for cross-sectional asset pricing puzzles, and for risks relevant to a long-horizon investor.

Table 3 summarizes the behavior of the implied cash-flow news and discount-rate news components of the market return. The top panel shows that discount-rate

news has a standard deviation of about 5% per month, much larger than the 2.5% standard deviation of cash-flow news. This is consistent with the finding of Campbell (1991) that discount-rate news is the dominant component of the market return. The table also shows that the two components of return are almost uncorrelated with one another. This finding differs from Campbell (1991) and particularly Campbell (1996); it results from our use of a richer forecasting model that includes the value spread as well as the aggregate price-earnings ratio.

Table 3 also reports the correlations of each state variable innovation with the estimated news terms, and the coefficients $(e1' + e1'\lambda)$ and $e1'\lambda$ that map innovations to cash-flow and discount-rate news. Innovations to returns and the price-earnings ratio are highly negatively correlated with discount-rate news, reflecting the mean reversion in stock prices that is implied by our VAR system. Market return innovations are weakly positively correlated with cash-flow news, indicating that some part of a market rise is typically justified by underlying improvements in expected future cash flows. Innovations to the price-earnings ratio, however, are weakly negatively correlated with cash-flow news, suggesting that price increases relative to earnings are not usually justified by improvements in future earnings growth.

Figure 2 illustrates the VAR model’s view of stock market history in relation to NBER recessions. Each dotted line in the figure corresponds to the trough of a recession as defined by the NBER. The top panel reports a trailing exponentially-weighted moving average of the market’s cash-flow news, while the bottom panel reports the same moving average of the market’s discount-rate news. It is clear from the figure that in some recessions our model attributes stock market declines to declining cash flows (e.g. 1991), in others to increasing discount rates (e.g. 2001), and in others to both types of news (e.g. the Great Depression and the 1970’s). We might call the first type of recession a “profitability recession”, the second type a “valuation recession”, and the third type a “mixed recession”. A valuation recession is characterized by a declining price-earnings ratio, a steepening yield curve, and larger declines in growth stocks than in value stocks. Profitability and valuation recessions, as opposed to mixed recessions, will be particularly influential observations when we estimate cash-flow and discount-rate betas, because these are episodes in which cash-flow and discount-rate news do not move closely together.

We set $\rho = .95^{1/12}$ in Table 3 and use the same value throughout the paper. Recall that ρ can be related to either the average dividend yield or the average consumption wealth ratio, as discussed on page 8. An annualized ρ of .95 corresponds to an average

dividend-price or consumption-wealth ratio of -2.94 (in logs) or 5.2% (in levels), where wealth is measured after subtracting consumption. We picked the value .95 because approximately 5% consumption of the total wealth per year seems reasonable for a long-term investor, such as a university endowment.

As a robustness check, we have estimated the VAR over subsamples before and after 1963. The coefficients that map state variable innovations to cash-flow and discount-rate news are fairly stable, with no changes in sign. Also, the value spread has greater predictive power in the first subsample than in the second. This is reassuring, since it indicates that the coefficient on this variable is not just fitting the last few years of the sample during which exceptionally high prices for growth stocks preceded a market decline. Given the stability of the VAR point estimates in the two subsamples and the unfortunate statistical fact that the coefficients of our monthly return-prediction regressions are estimated imprecisely (a problem that is magnified in shorter subsamples), we proceed to use the full-sample VAR-coefficient estimates in the remainder of the paper.

3 Measuring cash-flow and discount-rate betas

We have shown that market returns contain two components, both of which display substantial volatility and which are not highly correlated with one another. This raises the possibility that different types of stocks may have different betas with the two components of the market. In this section we measure cash-flow betas and discount-rate betas separately. We define the cash-flow beta as

$$\beta_{i,CF} \equiv \frac{\text{Cov}(r_{i,t}, N_{CF,t})}{\text{Var}(r_{M,t}^e - E_{t-1}r_{M,t}^e)} \quad (5)$$

and the discount-rate beta as

$$\beta_{i,DR} \equiv \frac{\text{Cov}(r_{i,t}, -N_{DR,t})}{\text{Var}(r_{M,t}^e - E_{t-1}r_{M,t}^e)}. \quad (6)$$

Note that the discount-rate beta is defined as the covariance of an asset's return with *good* news about the stock market in form of *lower-than-expected* discount rates, and that each beta divides by the total variance of unexpected market returns, not

the variance of cash-flow news or discount-rate news separately. This implies that the cash-flow beta and the discount-rate beta add up to the total market beta,

$$\beta_{i,M} = \beta_{i,CF} + \beta_{i,DR}. \quad (7)$$

Our estimates show that there is interesting variation across assets and across time in the two components of the market beta.

3.1 Test-asset data

We construct two sets of portfolios to use as test assets. The first is a set of 25 *ME* and *BE/ME* portfolios, available from Professor Kenneth French’s web site. The portfolios, which are constructed at the end of each June, are the intersections of five portfolios formed on size (*ME*) and five portfolios formed on book-to-market equity (*BE/ME*). *BE/ME* for June of year t is the book equity for the last fiscal year end in the calendar year $t - 1$ divided by *ME* for December of $t - 1$. The size and *BE/ME* breakpoints are NYSE quintiles. On a few occasions, no firms are allocated to some of the portfolios. In those cases, we use the return on the portfolio with the same size and the closest *BE/ME*.

The 25 *ME* and *BE/ME* portfolios were originally constructed by Davis, Fama, and French (2000) using three databases. The first of these, the CRSP monthly stock file, contains monthly prices, shares outstanding, dividends, and returns for NYSE, AMEX, and NASDAQ stocks. The second database, the COMPUSTAT annual research file, contains the relevant accounting information for most publicly traded U.S. stocks. The COMPUSTAT accounting information is supplemented by the third database, Moody’s book equity information hand collected by Davis et al.

Daniel and Titman (1997) point out that it can be dangerous to test asset pricing models using only portfolios sorted by characteristics known to be related to average returns, such as size and value. Characteristics-sorted portfolios are likely to show some spread in betas identified as risk by almost any asset pricing model, at least in sample. When the model is estimated, a high premium per unit of beta will fit the large variation in average returns. Thus, at least when premia are not constrained by theory, an asset pricing model may spuriously explain the average returns to characteristics-sorted portfolios.

To alleviate this concern, we follow the advice of Daniel and Titman and construct

a second set of 20 portfolios sorted on past risk loadings with VAR state variables (excluding the price-smoothed earnings ratio PE , since high-frequency changes in PE are so highly collinear with market returns). These portfolios are constructed as follows. First, we run a loading-estimation regression for each stock in the CRSP database:

$$\sum_{j=1}^3 r_{i,t+j} = b_0 + b_{r_M} \sum_{j=1}^3 r_{M,t+j} + b_{VS}(VS_{t+3} - VS_t) + b_{TY}(TY_{t+3} - TY_t) + \varepsilon_{i,t+3}, \quad (8)$$

where $r_{i,t}$ is the log stock return on stock i for month t . The regression (8) is reestimated from a rolling 36-month window of overlapping observations for each stock at the end of each month. Since these regressions are estimated from stock-level instead of portfolio-level data, we use a quarterly data frequency to minimize the impact of infrequent trading.

Our objective is to create a set of portfolios that have as large a spread as possible in their betas with the market and with innovations in the VAR state variables. To accomplish this, each month we perform a two-dimensional sequential sort on market beta and another state-variable beta, producing a set of ten portfolios for each state variable. First, we form two groups by sorting stocks on $\hat{b}_{VS}$. Then, we further sort stocks in both groups to five portfolios on $\hat{b}_{r_M}$ and record returns on these ten value-weight portfolios. To ensure that the average returns on these portfolio strategies are not influenced by various market-microstructure issues plaguing the smallest stocks, we exclude the smallest (lowest ME) five percent of stocks of each cross-section and lag the estimated risk loadings by a month in our sorts. We construct another set of ten portfolios in a similar fashion by sorting on $\hat{b}_{TY}$ and $\hat{b}_{r_M}$. We refer to these 20 return series as risk-sorted portfolios. Both the 25 size- and book-to-market-sorted returns and the 20 risk-sorted returns are measured over the period 1929:1–2001:12.

3.2 Empirical estimates of cash-flow and discount-rate betas

We estimate the cash-flow and discount-rate betas using the fitted values of the market's cash-flow and discount-rate news. Specifically, we use the following beta estimators:

$$\hat{\beta}_{i,CF} = \frac{\widehat{\text{Cov}}(r_{i,t}, \hat{N}_{CF,t})}{\widehat{\text{Var}}(\hat{N}_{CF,t} - \hat{N}_{DR,t})} + \frac{\widehat{\text{Cov}}(r_{i,t}, \hat{N}_{CF,t-1})}{\widehat{\text{Var}}(\hat{N}_{CF,t} - \hat{N}_{DR,t})} \quad (9)$$

$$\hat{\beta}_{i,DR} = \frac{\widehat{\text{Cov}}(r_{i,t}, -\hat{N}_{DR,t})}{\widehat{\text{Var}}(\hat{N}_{CF,t} - \hat{N}_{DR,t})} + \frac{\widehat{\text{Cov}}(r_{i,t}, -\hat{N}_{DR,t-1})}{\widehat{\text{Var}}(\hat{N}_{CF,t} - \hat{N}_{DR,t})} \quad (10)$$

Above, $\widehat{\text{Cov}}$ and $\widehat{\text{Var}}$ denote sample covariance and variance. $\hat{N}_{CF,t}$ and $\hat{N}_{DR,t}$ are the estimated cash-flow and expected-return news from the VAR model of Tables 2 and 3.

These beta estimators deviate from the usual regression-coefficient estimator in two respects. First, we include one lag of the market's news terms in the numerator. Adding a lag is motivated by the possibility that, especially during the early years of our sample period, not all stocks in our test-asset portfolios were traded frequently and synchronously. If some portfolio returns are contaminated by stale prices, market return and news terms may spuriously appear to lead the portfolio returns, as noted by Scholes and Williams (1977) and Dimson (1979). In addition, Lo and MacKinlay (1990) show that the transaction prices of individual stocks tend to react in part to movements in the overall market with a lag, and the smaller the company, the greater is the lagged price reaction. McQueen, Pinegar, and Thorley (1996) and Peterson and Sanger (1995) show that these effects exist even in relatively low-frequency data (i.e., those sampled monthly). These problems are alleviated by the inclusion of the lag term.

Second, as in (5) and (6), we normalize the covariances in (9) and (10) by $\widehat{\text{Var}}(\hat{N}_{CF,t} - \hat{N}_{DR,t})$ or, equivalently by the sample variance of the (unexpected) market return, $\widehat{\text{Var}}(r_{M,t}^e - E_{t-1}r_{M,t}^e)$. Under the maintained assumptions, $\hat{\beta}_{i,M} = \hat{\beta}_{i,CF} + \hat{\beta}_{i,DR}$ is equal to the portfolio i 's Scholes-Williams (1977) beta on unexpected market return. It is also equal to the so-called "sum beta" employed by Ibbotson Associates, which is the sum of multiple regression coefficients of a portfolio's return on contemporaneous and lagged unexpected market returns.⁴

⁴Scholes and Williams (1977) include an additional lead term, which captures the possibility that the market return itself is contaminated by stale prices. Under the maintained assumption that our news terms are unforecastable, the population value of this term is zero.

The Scholes-Williams beta formula also includes a normalization. The sum of the three regression coefficients is divided by one plus twice the market's autocorrelation. Since the first-order autocorrelation of our news series is zero under the maintained assumptions, this normalization factor is identically one.

"Sum beta" uses multiple regression coefficients instead of simple regression coefficients. Under the maintained assumption that the news terms are unforecastable, the explanatory variables in the

When we apply this estimation technique to our test-asset returns and our estimated market's cash-flow and discount-rate news series, we find dramatic differences in the beta estimates between the first half of our 1929:1–2001:12 sample and the second half. Accordingly, we report betas separately for two subsamples, 1929:1–1963:6 and 1963:7–2001:12. We choose to split the sample at 1963:7, because that is when COMPUSTAT data become reliable and because most of the evidence on the book-to-market anomaly is obtained from the post-1963:7 period. Unlike the thoroughly mined second subsample, the first subsample is relatively untouched and presents an opportunity for an out-of-sample test.

The top half of Table 4 shows the estimated betas for the 25 size and book-to-market portfolios over the period 1929:1–1963:6. The portfolios are organized in a square matrix with growth stocks at the left, value stocks at the right, small stocks at the top, and large stocks at the bottom. At the right edge of the matrix we report the differences between the extreme growth and extreme value portfolios in each size group; along the bottom of the matrix we report the differences between the extreme small and extreme large portfolios in each BE/ME category. The top matrix displays cash-flow betas, while the bottom matrix displays discount-rate betas. In square brackets after each beta estimate we report a standard error, calculated conditional on the realizations of the news series from the aggregate VAR model.

In the pre-1963 sample period, value stocks have both higher cash-flow and higher discount-rate betas than growth stocks. An equal-weighted average of the extreme value stocks across size quintiles has a cash-flow beta .16 higher than an equal-weighted average of the extreme growth stocks. The difference in estimated discount-rate betas is .22 in the same direction. Similar to value stocks, small stocks have higher cash-flow betas and discount-rate betas than large stocks in this sample (by .18 and .36 respectively, for an equal-weighted average of the smallest stocks across value quintiles relative to an equal-weighted average of the largest stocks). In summary, value and small stocks were unambiguously riskier than growth and large stocks over the 1929:1–1963:6 period.

A partial exception to this statement involves the smallest growth portfolio, which is particularly risky and has both cash-flow and discount-rate betas that exceed those of the smallest value portfolio. This small growth portfolio is well known to present a particular challenge to asset pricing models, for example the three-factor model of

multiple regression are uncorrelated, and thus the multiple regression coefficients are equal to simple regression coefficients.

Fama and French (1993) which does not fit this portfolio well. Recent evidence on small growth stocks by Lamont and Thaler (2001), Mitchell, Pulvino, and Stafford (2002), D’Avolio (2002) and others suggests that the pricing of some small growth stocks is materially affected by short-sale constraints and other limits to arbitrage. This may help to explain the unusual behavior of the small growth portfolio.

The bottom half of Table 4 shows the cash-flow and discount-rate betas for the risk-sorted portfolios. Both cash-flow betas and discount-rate betas are high for stocks that have had high market betas in the past. Thus, in the early sample period, sorting stocks by their past market betas induces a spread in both cash-flow betas and discount-rate betas. Sorting stocks by their value-spread or term-spread sensitivity induces only a relatively modest spread in either beta.

The patterns are completely different in the post-1963 period shown in Table 5. In this subsample, value stocks still have slightly higher cash-flow betas than growth stocks, but much lower discount-rate betas. The difference in cash-flow betas between the average across extreme value portfolios and the average across extreme growth portfolios is a modest .09. What is remarkable is that the pattern of discount-rate betas reverses in the modern period, so that growth stocks have significantly higher discount-rate betas than value stocks. The difference is economically large (.45) and statistically significant. Recall that cash-flow and discount-rate betas sum up to the CAPM beta; thus growth stocks have higher market betas in the modern period, but their betas are disproportionately of the “good” discount-rate variety rather than the “bad” cash-flow variety.

The changes in the risk characteristics of value and growth stocks that we identify by comparing the periods before and after 1963 are consistent with recent research by Franzoni (2002). Franzoni points out that the market betas of value stocks and small stocks have declined over time relative to the market betas of growth stocks and large stocks. We extend his research by exploring time changes in the two components of market beta, the cash-flow beta and the discount-rate beta.

What economic forces have caused these changes in betas? We suspect that the changing characteristics of value and growth stocks and small and large stocks are related to these patterns in sensitivities. Our first subsample is dominated by the Great Depression and its aftermath. Perhaps in the 1930’s value stocks were fallen angels with a large debt load accumulated during the Great Depression. The higher leverage of value stocks relative to that of growth stocks could explain both the higher cash-flow and expected-return betas of value stocks from 1929–1963. In general, low

leverage and strong overall position of a company may lead to a low cash-flow beta, and high leverage and weak position to a high cash-flow beta.

We also hypothesize that future investment opportunities, long duration of cash flows, and dependence on external equity finance lead to a high discount-rate beta. For example, if a distressed firm needed new equity financing simply to survive after the Great Depression, and if the availability and cost of such financing is related to the overall cost of capital, then such a firm's value is likely to have been very sensitive to discount-rate news. Similarly, new small firms with a negative current cash flow but valuable investment opportunities are likely to be very sensitive to discount-rate news. In the modern subsample, the growth portfolio probably contains a higher proportion of young companies following the initial-public-offering (IPO) wave of the 1960's, the inclusion of NASDAQ firms in our sample during the late 1970's, and the flood of technology IPOs in the 1990's.

The increase in growth stocks' discount-rate betas may also be partially explained by changes in stock market listing requirements. During the early period, only firms with significant internal cash flow made it to the Big Board and thus our sample. This is because, in the past, the New York Stock Exchange had very strict profitability requirements for a firm to be listed on the exchange. The low- BE/ME stocks in the first half of the sample are thus likely to be consistently profitable and independent of external financing. In contrast, our post-1963 sample also contains NASDAQ stocks and less-profitable new lists on the NYSE. These firms are listed precisely to improve their access to equity financing, and many of them will not even survive – let alone achieve their growth expectations – without a continuing availability of inexpensive equity financing.

Finally, it is possible that our discount-rate news is simply news about investor sentiment. If growth investing has become more popular among irrational investors during our sample period, growth stocks may have become more sensitive to shifts in the sentiment of these investors.

Our risk-sorted portfolios also have different betas in the second subsample. Sorting on market risk while controlling for other state variables induces a spread in only the discount-rate beta in the second subsample.

4 Pricing cash-flow and discount-rate betas

We have shown that in the period since 1963, there is a striking difference in the beta composition of value and growth stocks. The market betas of growth stocks are disproportionately composed of discount-rate betas rather than cash-flow betas. The opposite is true for value stocks.

Motivated by this finding, we next examine the validity of a long-horizon investor's first-order condition, assuming that the investor holds a 100% allocation to the market portfolio of stocks at all times. We ask whether the investor would be better off adding a margin-financed position in some of our test assets (such as value or small stocks), as a short-horizon investor's first-order condition would suggest.

Our main finding is that the long-horizon investor's first-order condition is not violated by our test assets and that the difference in beta composition can largely explain the high returns on value and low returns on growth stocks relative to the predictions of the static CAPM. The extreme small-growth portfolio remains an outlier even in our model, but the returns on this portfolio are not sufficiently anomalous to cause a statistical rejection of the model.

4.1 An intertemporal asset pricing model

Campbell (1993) derives an approximate discrete-time version of Merton's (1973) intertemporal CAPM. The model's central pricing statement is based on the first-order condition for an investor who holds a portfolio p of tradable assets that contains all of her wealth. Campbell assumes that this portfolio is observable in order to derive testable asset-pricing implications from the first-order condition.

Campbell considers an infinitely lived investor who has the recursive preferences proposed by Epstein and Zin (1989, 1991):

$$U(C_t, E_t(U_{t+1})) = \left[(1 - \delta) C_t^{\frac{1-\gamma}{\theta}} + \delta (E_t(U_{t+1}^{1-\gamma}))^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \quad (11)$$

where C_t is consumption at time t , $\gamma > 0$ is the relative risk aversion coefficient, $\psi > 0$ is the elasticity of intertemporal substitution, $0 < \delta < 1$ is the time discount factor, and $\theta \equiv (1 - \gamma)/(1 - \psi^{-1})$. These preferences are a generalization of power utility, formalized with an objective function (U) that retains the desirable scale-independence

of the power utility function. Deviating from the power-utility model, however, the Epstein-Zin preferences relax the restriction that the elasticity of intertemporal substitution must equal the reciprocal of the coefficient of relative risk aversion. In the Epstein-Zin model, the elasticity of intertemporal substitution, ψ , and the coefficient of relative risk aversion, γ , are both free parameters.

Campbell assumes that all asset returns are conditionally lognormal, and that the investor's portfolio returns and its two components are homoskedastic. The assumption of lognormality can be relaxed if one is willing to use Taylor approximations to the true Euler equations, and the model can be extended to allow changing variances as discussed by Chen (2003). Empirically, changes in volatility seem to be much less persistent than changes in expected returns, and thus they generate relatively modest intertemporal hedging effects on portfolio demands (Chacko and Viceira 1999). For this reason we continue to assume constant variances in the empirical work of this paper.

Campbell derives an approximate solution in which risk premia depend only on the coefficient of relative risk aversion γ and the discount coefficient ρ , and not directly on the elasticity of intertemporal substitution ψ . The approximation is accurate if the elasticity of intertemporal substitution is close to one, and it holds exactly in the limit of continuous time if the elasticity equals one. In the $\psi = 1$ case, $\rho = \delta$ and the optimal consumption-wealth ratio is conveniently constant and equal to $1 - \rho$. Thus our choice of $\rho = .95^{1/12}$ implies that at the end of each month, the investor chooses to consume .43% of her wealth if $\psi = 1$.⁵

Under these assumptions, the optimality of portfolio strategy p requires that the risk premium on any asset i satisfies

$$\begin{aligned} E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} &= \gamma \text{Cov}_t(r_{i,t+1}, r_{p,t+1} - E_t r_{p,t+1}) \\ &\quad + (1 - \gamma) \text{Cov}_t(r_{i,t+1}, -N_{p,DR,t+1}), \end{aligned} \quad (12)$$

where p is the optimal portfolio that the agent chooses to hold and $N_{p,DR,t+1} \equiv (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{p,t+1+j}$ is discount-rate or expected-return news on this portfolio.

The left hand side of (12) is the expected excess log return on asset i over the riskless interest rate, plus one-half the variance of the excess return to adjust for

⁵Schroder and Skiadas (1999) examine this case in a continuous-time framework which eliminates the need for approximations if $\psi = 1$.

Jensen's Inequality. This is the appropriate measure of the risk premium in a log-normal model. The right hand side of (12) is a weighted average of two covariances: the covariance of return i with the return on portfolio p , which gets a weight of γ , and the covariance of return i with negative of news about future expected returns on portfolio p , which gets a weight of $(1 - \gamma)$. These two covariances represent the myopic and intertemporal hedging components of asset demand, respectively. When $\gamma = 1$, it is well known that portfolio choice is myopic and the first-order condition collapses to the familiar one used to derive the pricing implications of the CAPM.

We can rewrite equation (12) to relate the risk premium to covariance with cash-flow news and discount-rate news. Since $r_{p,t+1} - E_t r_{p,t+1} = N_{p,CF,t+1} - N_{p,DR,t+1}$, we have

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \text{Cov}_t(r_{i,t+1}, N_{p,CF,t+1}) + \text{Cov}_t(r_{i,t+1}, -N_{p,DR,t+1}). \quad (13)$$

Multiplying and dividing by the conditional variance of portfolio p 's return, $\sigma_{p,t}^2$, we obtain

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \sigma_{p,t}^2 \beta_{i,CF_p,t} + \sigma_{p,t}^2 \beta_{i,DR_p,t}. \quad (14)$$

This equation delivers our prediction that “bad beta” with cash-flow news should have a risk price γ times greater than the risk price of “good beta” with discount-rate news, which should equal the variance of the return on portfolio p .

In our empirical work, we begin by assuming that portfolio p is fully invested in a value-weighted equity index. This assumption implies that the risk price of discount-rate news should equal the variance of the value-weighted index, about 5% in the early subsample and 2.5% in the modern subsample. The only free parameter in equation (14) is then the coefficient of relative risk aversion, γ .

An alternative assumption would be that portfolio p places a weight w on the value-weighted index and $(1 - w)$ on Treasury bills. If the real Treasury-bill return is constant, this would imply that the variance of portfolio p is w^2 times the variance of the index return, while the cash-flow and discount-rate betas of test asset i with portfolio p are $(1/w)$ times the cash-flow and discount-rate betas with the index return. Under this alternative the risk prices for both cash-flow and discount-rate betas are w times smaller, but the risk price for the cash-flow beta is still γ times the risk price for the discount-rate beta. The risk prices of the two betas can be used to identify the two free parameters w and γ .

4.2 Empirical estimates of risk premia

Would an all-stock investor be better off holding stocks at market weights or overweighting value and small stocks? We examine the validity of an unconditional version of the first-order condition (14) relative to the market portfolio of stocks. We modify (14) in three ways. First, we use simple expected returns, $E_t[R_{i,t+1} - R_{rf,t+1}]$, on the left-hand side, instead of log returns, $E_t[r_{i,t+1}] - r_{rf,t+1} + \sigma_{i,t}^2/2$. In the log-normal model, both expectations are the same, and by using simple returns we make our results easier to compare with previous empirical studies. Second, we condition down equation (13) to derive an unconditional version of (14) to avoid estimation of all required conditional moments. Finally, we change the subscript p to M and use all-stock investment in the market portfolio of stocks as the reference portfolio, reflecting the fact that we test the optimality of the market portfolio of stocks for the long-horizon investor. These modifications yield:

$$E[R_i - R_f] = \gamma \sigma_M^2 \beta_{i,CFM} + \sigma_M^2 \beta_{i,DRM} \quad (15)$$

We assume that the log real risk-free rate is approximately constant. We make this assumption mainly because monthly inflation data are unreliable, especially over our long 1928:12-2001:12 sample period. This assumption is unlikely to have a major impact on our tests, since we focus on stock portfolios. The main practical implication of the constant-real-rate assumption is that cash-flow and discount-rate news computed from excess CRSP value-weight index returns are identically equivalent to news terms computed from real CRSP value-weight index returns.

We use 45 test assets, 25 size- and book-to-market sorted portfolios and 20 risk-sorted portfolios, on the left hand side of the unconditional first-order condition (15). We evaluate the performance of the traditional CAPM that restricts cash-flow and discount-rate betas to have the same price of risk, our two-beta intertemporal asset pricing model that restricts the price of discount-rate risk to equal the variance of the market return, and an unrestricted two-beta model that allows free risk prices for cash-flow and discount-rate betas. As discussed above, the unrestricted model can be interpreted as a slight generalization of our model that allows the rational investor's portfolio to include Treasury bills as well as equities.

Each model is estimated in two different forms: one with a restricted zero-beta rate equal to the Treasury-bill rate, and one with an unrestricted zero-beta rate following Black (1972). The first specification includes Treasury bills in the set of

alternative assets available to the investor, while the second assumes that the investor is considering only reallocations of the portfolio among alternative types of equities. Thus in the first specification we ask the model to explain the unconditional equity premium as well as the premia to value stocks, small stocks, and risk-sorted stocks; in the second specification we remove the equity premium from the set of phenomena to be explained.

Table 6 reports results for the early sample period 1929:1–1963:6. The table has six columns, two specifications for each of our three asset pricing models. The first nine rows of Table 6 are divided into three sets of four rows. The first set of four rows corresponds to the zero-beta rate (in excess of the Treasury-bill rate), the second set to the premium on cash-flow beta, and the third set to the premium on discount-rate beta. With each set, the first row reports the point estimate in fractions per month, and the second row annualizes this, multiplying by 1200 to ease the interpretation of the estimate. The third and fourth rows present two alternative standard errors of the monthly estimate.

These parameters are estimated from a cross-sectional regression

$$\overline{R}_i^e = g_0 + g_1 \hat{\beta}_{i,CF} + g_2 \hat{\beta}_{i,DR} + e_i, \quad (16)$$

where bar denotes time-series mean and $\overline{R}_i^e \equiv \overline{R}_i - \overline{R}_{rf}$ denotes the sample average simple excess return on asset i . The implied risk-aversion coefficient can be recovered as g_1/g_2 .

Standard errors are produced with a bootstrap from 2500 simulated realizations. Our bootstrap experiment samples test-asset returns and VAR errors, and uses the OLS VAR estimates in Table 2 to generate the state-variable data. We partition the VAR errors and test-asset returns into two groups, one for 1929:1–1963:6 and another for 1963:7–2001:12, which enables us to use the same simulated realizations in subperiod analyses. The first set of standard errors (labelled A) conditions on estimated news terms and generates betas and return premia separately for each simulated realization, while the second set (labelled B) also estimates the VAR and the news terms separately for each simulated realization. Standard errors B thus incorporate the considerable additional sampling uncertainty due to the fact that the news terms as well as betas are generated regressors.

Below the premia estimates, we report the $\hat{R}^2$ statistic for a cross-sectional regression of average returns on our test assets onto the fitted values from the model. The

regression $\hat{R}^2$ is computed as

$$\hat{R}^2 = 1 - \frac{\hat{e}'\hat{e}}{(\bar{R}_i^e - \sum_i \bar{R}_i^e)' (\bar{R}_i^e - \sum_i \bar{R}_i^e)} , \quad (17)$$

which allows for negative $\hat{R}^2$ for poorly fitting models estimated under the constraint that the zero-beta rate equals the risk-free rate.

Although the regression $\hat{R}^2$ is intuitive and transparent, it gives equal weight to each asset included in the set of test assets even though some assets may be more volatile than others. To address this concern we also report a composite pricing error and its 5% critical value. The composite pricing error is computed as $\hat{e}'\hat{\Omega}^{-1}\hat{e}$, where $\hat{e}$ is the vector of estimated residuals from regression (16) and $\hat{\Omega}$ is a diagonal matrix with estimated return volatilities on the main diagonal. The weighting matrix, $\hat{\Omega}^{-1}$, in the composite pricing error formula places less weight on noisy observations yet it is independent of the specific pricing model. We avoid using a freely estimated variance-covariance matrix of test asset returns for $\hat{\Omega}$ because with 45 test assets, we are concerned that the inverse of this matrix would be poorly behaved. Hodrick and Zhang (2001) discuss related alternative methods for assessing the performance of asset pricing models.

Two alternative 5% critical values for the composite pricing error are produced with a bootstrap method similar to the one we have described above, except that the test-asset returns are adjusted to be consistent with the pricing model before the random samples are generated. Critical values A condition on estimated news terms, while critical values B take account of the fact that news terms must be estimated.

Table 6 shows that in the 1929:1–1963:6 period, the traditional CAPM explains the cross-section of stock returns reasonably well, and is comparable to the restricted two-beta model and the two-beta model with unrestricted risk prices. The cross-sectional R^2 statistics are about 40% for models with zero-beta rates equal to the Treasury-bill rate, and around 45% for models with unrestricted zero-beta rates. None of the models in the table come close to being rejected at the 5% level.

Figure 3 provides a visual summary of these results. The figure plots the predicted average excess return on the horizontal axis and the actual sample average excess return on the vertical axis. For a model with a 100% estimated R^2 , all the points would fall on the 45-degree line displayed in each graph. The triangles in the figures denote the 24 Fama-French portfolios and asterisks the 20 risk-sorted portfolios. All

the models generate nearly identical scatter plots.

The good performance of the CAPM in the 1929–1963 period is due to the fact that in this period, the bad cash-flow beta is roughly a constant fraction of the CAPM beta across assets. Thus our tests cannot discriminate between the static and intertemporal CAPM models in this period.

Results are very different in the 1963:7–2001:12 period. Table 7 shows that in this period, the CAPM fails disastrously to explain the returns on the test assets. When the zero-beta rate is left a free parameter, the cross-sectional regression picks a negative premium for the CAPM beta and implies a near-zero estimated R^2 . When the zero-beta rate is constrained to the risk-free rate, the CAPM $\hat{R}^2$ falls to -60%, i.e., the model has larger pricing error than the null hypothesis that all portfolios have equal expected returns. The static CAPM is easily rejected at the 5% level by both sets of critical values.

The two-beta model with a restricted risk price for discount-rate news explains almost 50% of the cross-sectional variation in average returns across our test assets. The model performs almost as well with a restricted zero-beta rate, equal to the Treasury bill rate, as it does with an unrestricted Treasury bill rate. This indicates that both the unconditional equity premium and the premia on alternative equity portfolios can be rationalized by the same coefficient of risk aversion. The estimated risk price for cash-flow beta is high at 58% per year with a restricted zero-beta rate and 69% per year with an unrestricted zero-beta rate. There are large standard errors on these estimates, but they are statistically distinguishable from the low risk price on discount-rate news. The model is not rejected at the 5% level by either set of critical values.

The critical values for the restricted intertemporal model with a restricted zero-beta rate are particularly large, an order of magnitude larger than those for the other models in the table. This is due to the fact that this model pins down both the zero-beta rate and the risk price for discount-rate news, and thus it pins down the total return generated by a unit of discount-rate beta. Since estimated discount-rate betas are noisy, estimates of this model can behave extremely badly even if the model is true.

The two-beta model with an unrestricted risk price assigns an even lower risk price to discount-rate beta than the variance of the market return. This would be consistent with a modified model in which a conservative rational investor holds a

portfolio that contains Treasury bills as well as equities. The implied share of equities in the portfolio is 60% in the model with a restricted zero-beta rate, and slightly below 40% in the model with an unrestricted zero-beta rate. This model generates cross-sectional R^2 statistics slightly above 50%. A visual summary of these results is provided by Figure 4.

Another way to evaluate the performance of our model is to compare it to less theoretically structured models. We do this in two ways. First, we compare our restricted ICAPM model to a model whose factors are the four innovations from our VAR system, with unrestricted risk prices. In the modern sample the four unrestricted risk prices line up almost perfectly with those implied by our restricted model. Second, we compare the two-beta model to the influential three-factor model of Fama and French (1993). The Fama-French model includes three risk factors, one each for the market, small stocks, and value stocks. We estimate the betas for each test asset from simple returns using Ibbotson’s sum-beta methodology with one lag and then regress the average excess test-asset returns on the estimated betas. In the early subsample, the cross-sectional R^2 statistic for the Fama-French three-factor model is 10 percentage points higher than that for our two-beta model with an unconstrained zero-beta rate, and 1 percentage point higher with a zero-beta rate constrained to the risk-free rate. In the modern subsample, the Fama-French model outperforms the two-beta model by 30 and 26 percentage points, respectively. This difference in explanatory power is not statistically significant, as the restrictions of our model are not rejected by our composite pricing error test. Given that the Fama-French model has three freely estimated betas and thus two additional degrees of freedom (since the premium on discount-rate beta is restricted to equal the variance of the market’s return in our model), we consider the relative performance of the two-beta ICAPM to be a success.

Although the two-beta model is generally quite successful in explaining the cross-section of average returns, the model cannot price the extreme small-growth portfolio. In the first subsample, the extreme small-growth portfolio has an annualized average return that is 8.8 percentage points lower than the model’s prediction. In the second subsample, this return on this portfolio is 3.2 percentage points lower than the model’s prediction. These pricing-error calculations use the model specification with the zero-beta rate constrained to the risk-free rate. In both subsamples, these pricing errors are economically large and not meaningfully smaller than the pricing errors of the Sharpe-Lintner CAPM for this portfolio (9.9 percentage points in the first and 7.25 percentage points in the second subsample).

One concern about these results might be that the estimated preference parameters appear rather different in the first and second subsamples. The point estimate of risk aversion, in the model with a restricted zero-beta rate and risk price for discount-rate news, is 3.6 in the first subsample and almost 24 in the second subsample. Even if betas and the variance of the market return have changed over time, one would hope that the underlying preferences of investors have remained stable. To address this concern, we have estimated a version of our model that allows for changing betas and variances across the two subsamples, but imposes a constant coefficient of relative risk aversion. This model is not rejected at the 5% level, and the implied risk aversion coefficient is approximately six. Also, if we allow for different risk-aversion coefficients for the subsamples, we cannot reject the hypothesis that the two parameters are the same.

Another way to come at this issue is to estimate the preference parameters from a conditional model. We show the results of this exercise in Figure 5. We show the smoothed conditional premium on $\text{Cov}_t(r_{i,t+1}, N_{M,CF,t+1})$ and $\text{Cov}_t(r_{i,t+1}, -N_{M,DR,t+1})$, with the ICAPM predicting premium of γ on the former and unit premium on the latter. The graph is produced in three steps as follows. First, we run three sets of 45 time-series regressions on a constant, time trend, and the lagged VAR state variables, i.e., three regressions per test asset. The dependent variables in these regressions are simple excess return on the test assets $(R_{i,t}^e)$, $(N_{CF,t} + N_{CF,t-1})R_{i,t}^e$, and $(N_{DR,t} + N_{DR,t-1})R_{i,t}^e$. Second, each month we regress the forecast values of excess return on the forecast values of the two covariances, excluding the constant and thus restricting the zero-beta rate to equal the risk-free rate. Third, we plot the five-year moving averages of these cross-sectional regression coefficients in Figure 5.

The lower line in Figure 5 is the estimated risk price for the discount-rate beta, divided by the variance of market returns. If our ICAPM holds exactly, this should be a horizontal line with a height of one. The upper line is the estimated risk price for the cash-flow beta, again divided by the variance of market returns. According to our ICAPM, this should be a horizontal line with a height of γ . The traditional CAPM implies that both lines should have the same height. Figure 5 shows that the scaled price of discount-rate risk has a long-term average very close to one, with substantial variations around this average, while the scaled price of cash-flow risk has a long-term average around six, again with substantial shorter-term variations. During the period 1935–1955 the two lines are close to one another, illustrating the good performance of the CAPM in this period. For most of the period since 1960 the two lines have diverged substantially, but there is no sign of a trend or other

low-frequency instability in the risk prices.

In an effort to ascertain the robustness of our empirical results, we have experimented with various alternative specifications to those considered here. Our results are robust to many reasonable changes in our empirical methodology, but a few features of this methodology are essential to the good performance of the model. First, when we estimate betas in our monthly model we need to include at least one lagged news term in the regression in the manner of Scholes and Williams (1977). If we change the data frequency to quarterly, we no longer need to use any lagged news terms when we estimate betas. We attribute this in part to the fact that our aggregate VAR contains a price-earnings ratio whose earnings term is updated quarterly, so this source of news about aggregate cash flows is measured quarterly rather than monthly. Second, the loglinearization parameter ρ affects the relative volatility of the cash-flow and discount-rate news terms. If we vary ρ within the range 0.94 to 0.96 the results are very similar to those we report for $\rho = 0.95$, but outside this range the performance of our most restricted model, the ICAPM with a single free parameter, begins to deteriorate. Less restricted versions of the model, with a free zero-beta rate or free risk price for discount-rate beta, are relatively insensitive to the choice of ρ .

Finally, our results depend on the inclusion of the small-stock value spread in our aggregate VAR system. If we exclude this variable we no longer find a large difference between the cash-flow betas of value stocks and growth stocks. We have discussed various reasons why the small-stock value spread might predict market returns. Further motivation is provided by the ICAPM itself. We know that value stocks outperform growth stocks, particularly among smaller stocks, and that this cannot be explained by the traditional static CAPM. If the ICAPM is to explain this anomaly, then small growth stocks must have intertemporal hedging value that offsets their low returns; that is, their returns must be negatively correlated with innovations to investment opportunities. In order to evaluate this hypothesis it is natural to ask whether a long moving average of small growth stock returns predicts investment opportunities. This is exactly what we do when we include the small-stock value spread in our forecasting model for market returns. In short, the small-stock value spread is not an arbitrary forecasting variable but one that is suggested by the asset pricing theory we are trying to test.

4.3 Loose ends and future directions

A number of unresolved issues remain. First, we have used a model that assumes a constant variance for the market return and its two components. We can extend the model to allow for changing volatility of the market return, in the manner Chen (2003), but in this case we must measure news about volatility-adjusted discount rates rather than simply news about discount rates themselves. We believe that the properties of market discount-rate news will be fairly insensitive to any volatility adjustment, since movements in market volatility appear to be relatively short-lived. Related to this, we can allow for dynamically changing betas rather than assuming, as we have done here, that betas are constant over long periods of time. Ang and Chen (2003) and Franzoni (2002) discuss alternative methods for estimating the evolution of betas over time.

We have assumed that the rational long-term investor always holds a constant proportion of her assets in equities. But if expected returns on stocks vary over time while the risk-free interest rate and the volatility of the stock market are approximately constant, the long-term investor has an obvious incentive to strategically time the market. In future work we plan to extend the model to examine whether a long-term investor who strategically allocates wealth into stocks and bonds would be better off overweighting small and value stocks than holding the stock portion of her portfolio at market weights. With this extension it will be important to handle changing volatility correctly, since a strategic market-timing portfolio will be heteroskedastic even if the stock market portfolio is homoskedastic.

We have nothing to say about the profitability of momentum strategies. Although we have not examined this issue in detail, we are pessimistic about the two-beta model's ability to explain average returns on portfolios formed on past one-year stock returns, or on recent earnings surprises. Stocks with positive past news and high short-term expected returns are likely to have a higher fraction of their betas due to discount-rate betas, and thus are likely to have even lower return predictions in the ICAPM than the already-too-low predictions of the static CAPM.

Our model is silent on what is the ultimate source of variation in the market's discount rate. The mechanism that causes the market's overall valuation level to fluctuate would have to meet at least two criteria to be compatible with our simple intertemporal asset-pricing model. The shock to discount rates cannot be perfectly correlated with the shock to cash flows. Also, states of the world in which discount

rates increase while expected cash flows remain constant should not be states in which marginal utility is unusually high for other reasons. If marginal utility is very high in those states, the discount-rate risk factor will have a high premium instead of the low premium we detect in the data.

We have estimated the cash-flow and discount-rate betas of value and growth stocks from the behavior of their returns, without showing how these betas are linked to the underlying cash flows of value and growth companies. Similar to our decomposition of the market return, an individual firm's stock return can be split into cash-flow and discount-rate news. Through this decomposition, a stock's cash-flow and discount-rate betas can be further decomposed into two parts each, along the lines of Campbell and Mei (1993) and Vuolteenaho (2002), and this decomposition might yield interesting additional insights. For example, the hypothesis that growth stocks are equity-dependent companies predicts that at least some of the high covariance between growth stocks' returns and the market's discount-rate news is due to covariance between growth stocks' cash flows and the market's discount-rate news. A pure investor-sentiment hypothesis would probably predict that all of the higher discount-rate beta is due to covariance between growth stocks' expected returns and the market's discount-rate news. Preliminary results in Campbell, Polk, and Vuolteenaho (2003) suggest that the cash-flow properties of growth and value stocks are the main determinants of their betas with the cash-flow and discount-rate news on the aggregate stock market. Bansal, Dittmar, and Lundblad (2003) also model the cash flows of value and growth stocks in relation to their risks in a consumption-based asset pricing model.

Finally, our model has interesting implications for corporate finance, specifically for the methods used by corporations to calculate a cost of capital when evaluating investment projects. The two-beta model suggests that the most important determinant of the cost of capital is not the market beta of a project, but its cash-flow beta. This could be estimated using an econometric model, as we do here, but it is possible that simpler methods, such as estimating betas over long horizons or regressing returns on aggregate corporate profitability, would also provide useful estimates of cash-flow beta and thus of the cost of capital.

5 Conclusions

In his discussion of empirical evidence on market efficiency, Fama (1991) writes: “In the end, I think we can hope for a coherent story that (1) relates the cross-section properties of expected returns to the variation of expected returns through time, and (2) relates the behavior of expected returns to the real economy in a rather detailed way.” In this paper, we have presented a model that meets the first of Fama’s objectives and shows empirically that Merton’s (1973) intertemporal capital asset pricing model (ICAPM) helps to explain the cross-section of average stock returns.

We propose a simple and intuitive two-beta model that captures a stock’s risk in two risk loadings, cash-flow beta and discount-rate beta. The return on the market portfolio can be split into two components, one reflecting news about the market’s future cash flows and another reflecting news about the market’s discount rates. A stock’s cash-flow beta measures the stock’s return covariance with the former component and its discount-rate beta its return covariance with the latter component. Intertemporal asset pricing theory suggests that the “bad” cash-flow beta should have a higher price of risk than the “good” discount-rate beta. Specifically, the ratio of the two risk prices equals the risk aversion coefficient that makes an investor content to hold the aggregate market, and the “good” risk price should equal the variance of the return on the market.

Empirically, we find that value stocks and small stocks have considerably higher cash-flow betas than growth stocks and large stocks, and this can explain their higher average returns. The post-1963 negative CAPM alphas of growth stocks are explained by the fact that their betas are predominantly of the good variety. The model also explains why the sort on past CAPM betas induces a strong spread in average returns during the pre-1963 sample but little spread during the post-1963 sample. The post-1963 CAPM beta sort induces a post-ranking spread only in the good discount-rate beta, which carries a low premium. Finally, the model achieves these successes with the discount-rate premium constrained to the prediction of the intertemporal model.

Our model has important implications for rational investors. While we do not show that such investors should hold the market portfolio in preference to strategically timing the equity market, we do show that sufficiently risk-averse equity-only investors with a long investment horizon should view the high average returns on value stocks and small stocks as appropriate compensation for risk rather than a justification for systematic tilts towards these types of stocks.

Our two-beta model is, of course, not the first attempt to operationalize Merton's (1973) ICAPM. However, we hope that our model is an improvement over the specifications by Campbell (1996), Li (1997), Hodrick, Ng, and Sengmueller (1999), Lynch (1999), Ng (2002), Guo (2002), Brennan, Wang, and Xia (2003), Chen (2003) and others in two respects. First, our specification "works" in the sense that it has respectable explanatory power in explaining the cross-section of average asset returns with premia restricted to values predicted by the theory. Second, by restating the model in the simple two-beta form, with a close link to the static CAPM, we hope to facilitate the empirical implementation of the ICAPM in both academic research and practical applications.

References

- Adrian, Tobias and Francesco Franzoni, 2002, Learning about beta: An explanation of the value premium, unpublished paper, MIT.
- Ang, Andrew and Geert Bekaert, 2001, Stock return predictability: Is it there?, unpublished paper, Columbia University.
- Ang, Andrew and Joseph Chen, 2003, CAPM over the long run: 1926–2001, unpublished paper, Columbia University and USC.
- Asness, C., J. Friedman, R. Krail, and J. Lieu, 2000, Style timing: Value versus growth, *Journal of Portfolio Management*, Spring.
- 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, 2003, Interpreting risk premia across size, value, and industry portfolios, unpublished paper, Duke University and Indiana University.
- Banz, Rolf W., 1981, The relation between return and market value of common stocks, *Journal of Financial Economics* 9, 3–18.
- Basu, Sanjoy, 1977, Investment performance of common stocks in relation to their price-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.
- Black, Fischer, 1972, Capital market equilibrium with restricted borrowing, *Journal of Business* 45, 444–454.
- Brainard, William C., Matthew D. Shapiro, and John Shoven, 1991, Market value and fundamental value, in William C. Brainard, William D. Nordhaus, and Harold W. Watts eds. *Money, Macroeconomics, and Economic Policy: Essays in Honor of James Tobin*, MIT Press, Cambridge, MA.

- Brav, Alon, Reuven Lehavy, Roni Michaely, 2002, Expected return and asset pricing, unpublished paper, Duke University.
- Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2001, A simple model of intertemporal capital asset pricing and its implications for the Fama-French three-factor model, unpublished paper, UCLA.
- Brennan, Michael J., Ashley W. Wang, and Yihong Xia, 2003, Estimation and test of a simple model of intertemporal capital asset pricing, forthcoming *Journal of Finance*.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, John Y., 1993, Intertemporal asset pricing without consumption data, *American Economic Review* 83, 487–512.
- Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298–345.
- Campbell, John Y. and Jianping Mei, 1993, Where do betas come from? Asset price dynamics and the sources of systematic risk, *Review of Financial Studies* 6, 567–592.
- Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2003, Growth or glamour?, unpublished paper, Harvard University and Northwestern University.
- Campbell, John Y. and Robert J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y. and Robert J. Shiller, 1988b, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.
- Campbell, John Y. and Robert J. Shiller, 1998, Valuation ratios and the long-run stock market outlook, *Journal of Portfolio Management* 24(2), 11–26.

- Campbell, John Y. and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.
- Campbell, John Y. and Motohiro Yogo, 2002, Efficient tests of stock return predictability, unpublished paper, Harvard University.
- Chacko, George and Luis M. Viceira, 1999, Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, NBER Working Paper No. 7377.
- Chen, Joseph, 2003, Intertemporal CAPM and the cross-section of stock returns, unpublished paper, USC.
- Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2002, Does risk or mispricing explain the cross-section of stock prices?, unpublished paper, Harvard University.
- Cohen, Randolph B., Christopher Polk, and Tuomo Vuolteenaho, 2003, The value spread, *Journal of Finance* 58.
- Cornell, Bradford, 1999, Risk, duration, and capital budgeting: New evidence on some old questions, *Journal of Business* 72, 183–200.
- Daniel, Kent and Sheridan Titman, 1997, Evidence on the characteristics of cross sectional variation in stock returns, *Journal of Finance* 52, 1–33.
- Davis, J. L., Eugene F. Fama, and Kenneth R. French, 2000, Characteristics, covariances, and average returns: 1929 to 1997, *Journal of Finance* 55, 389–406.
- D’Avolio, Gene, 2002, The market for borrowing stock, *Journal of Financial Economics* 66, 271–306.
- Dimson, Elroy, 1979, Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics* 7, 197–226.
- Eleswarapu, Venkat R. and Marc R. Reinganum, 2003, The predictability of aggregate stock market returns: Evidence based on glamour stocks, forthcoming *Journal of Business*.

- Epstein, Lawrence and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–69.
- Epstein, Lawrence and Stanley Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical investigation, *Journal of Political Economy* 99, 263–286.
- Fama, Eugene F., 1991, Efficient capital markets: II, *Journal of Finance* 46, 1575–1617.
- Fama, Eugene F. and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–27.
- Fama, Eugene F. and Kenneth R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Fama, Eugene F. and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.
- Fama, Eugene F. and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Franzoni, Francesco, 2002, Where is beta going? The riskiness of value and small stocks, unpublished paper, MIT.
- Graham, Benjamin and David L. Dodd, 1934, *Security Analysis*, first edition, McGraw Hill, New York.
- Guo, Hui, 2002, Time-varying risk premia and the cross section of stock returns, unpublished paper, Federal Reserve Bank of St. Louis.
- Hodrick, Robert J., David Ng, and Paul Sengmueller, 1999, An international dynamic asset pricing model, *International Taxation and Public Finance* 6, 597–620.
- Hodrick, Robert J., and Xiaoyan Zhang, 2001, Evaluating the specification errors of asset pricing models, *Journal of Financial Economics* 62, 327–376.
- Jagannathan, Ravi and Z. Wang, 1996, The conditional CAPM and the cross-section of expected returns, *Journal of Finance* 51, 3–53.

- Keim, Donald and Robert Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357-390.
- Kim, Tong Suk, and Edward Omberg, 1986, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141-161.
- Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny, 1994, Contrarian investment, extrapolation, and risk, *Journal of Finance* 49, 1541-1578.
- Lamont, Owen, and Richard Thaler, 2003, Can the market add and subtract? Mispricing in tech stock carve-outs, *Journal of Political Economy* 111, 227-268.
- La Porta, Rafael, 1996, Expectations and the cross-section of returns, *Journal of Finance* 51, 1715-1742.
- La Porta, Rafael, Josef Lakonishok, Andrei Shleifer, and Robert W. Vishny, 1997, Good news for value stocks: Further evidence on market efficiency, *Journal of Finance* 52, 859-874.
- Lettau, Martin and Sydney Ludvigson, 2001, Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238-1287.
- Lewellen, Jonathan, 2003, Predicting returns with financial ratios, forthcoming *Journal of Financial Economics*.
- Lewellen, Jonathan and Stefan Nagel, 2003 The conditional CAPM does not explain asset-pricing anomalies, unpublished paper, MIT and LBS.
- Lewellen, Jonathan and Jay Shanken, 2002, Learning, asset-pricing tests, and market efficiency, *Journal of Finance* 57, 1113-1145.
- Li, Y., 1997, Intertemporal asset pricing without consumption data: Empirical tests, *Journal of Financial Research* 20, 53-69.
- 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, The valuation of risky assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13-37.

- Lo, Andrew and A. Craig MacKinlay, 1990, An econometric analysis of nonsynchronous trading, *Journal of Econometrics* 45, 181–212.
- Lynch, Anthony W., 1999, Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability, *Journal of Financial Economics* 62, 67–130.
- McQueen, Pinegar, and Thorley, 1996, Delayed reaction to good news and the cross-autocorrelation of portfolio returns, *Journal of Finance* 51, 889–919.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–87.
- Mitchell, Mark, Todd Pulvino, and Erik Stafford, 2001, Limited arbitrage in equity markets, *Journal of Finance* 57, 551–584.
- Ng, David, 2003, The international CAPM when expected returns are time-varying, forthcoming, *Journal of International Money and Finance*.
- Ng, Victor, Robert F. Engle, and Michael Rothschild, 1992, A multi-dynamic-factor model for stock returns, *Journal of Econometrics* 52, 245–266.
- Perez-Quiros, Gabriel and Allan Timmermann, 2000, Firm size and cyclical variations in stock returns, *Journal of Finance* 55, 1229–1262.
- Peterson, James D. and Gary C. Sanger, 1995, Cross-autocorrelations, systematic risk and the period of listing, unpublished paper, University of Notre Dame.
- Reinganum, Mark R., 1981, Misspecification of capital asset pricing: Empirical anomalies based on yields and market values, *Journal of Financial Economics* 9, 19–46.
- Roll, Richard, 1977, A critique of the asset pricing theory's tests: Part I, *Journal of Financial Economics* 4, 129–176.
- Rosenberg, Barr, Kenneth Reid, and Ronald Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management* 11, 9–17.
- Ross, Stephen A., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.

- Rozeff, M., 1984, Dividend yields are equity premiums, *Journal of Portfolio Management* 11, 68-75.
- Scholes, Myron and J. Williams, 1977, Estimating betas from nonsynchronous data, *Journal of Financial Economics* 5, 309-327.
- Schroder, Mark, and Costis Skiadas, 1999, Optimal consumption and portfolio selection with stochastic differential utility, *Journal of Economic Theory* 89, 68-126.
- Shanken, Jay, 1987, Multivariate proxies and asset pricing relations: Living with the Roll critique, *Journal of Financial Economics* 18, 91-110.
- Sharpe, William, 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425-442.
- Shiller, Robert J., 2000, *Irrational Exuberance*, Princeton University Press, Princeton, NJ.
- Stambaugh, Robert F., 1982, On the exclusion of assets from tests of the two parameter model, *Journal of Financial Economics* 10, 235-268.
- Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375-421.
- Torous, Walter, Rossen Valkanov, and Shu Yan, 2003, On predicting stock returns with nearly integrated explanatory variables, forthcoming *Journal of Business*.
- 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.
- Zhang, Lu and Ralitsa Petkova, 2002, Is value riskier than growth?, unpublished paper, University of Rochester.

Table 1: Descriptive statistics of the VAR state variables

The table shows the descriptive statistics of the VAR state variables estimated from the full sample period 1928:12-2001:12, 877 monthly data points. r_M^e is the excess log return on the CRSP value-weight index. TY is the term yield spread in percentage points, measured as the yield difference between ten-year constant-maturity taxable bonds and short-term taxable notes. PE is the log ratio of S&P 500's price to S&P 500's ten-year moving average of earnings. VS is the small-stock value-spread, the difference in the log book-to-market ratios of small value and small growth stocks. The small value and small growth portfolios are two of the six elementary portfolios constructed by Davis, Fama, and French (2000). "Stdev." denotes standard deviation and "Autocorr." the first-order autocorrelation of the series.

Variable	Mean	Median	Stdev.	Min	Max	Autocorr.
r_M^e	.004	.009	.056	-.344	.322	.108
TY	.629	.550	.643	-1.350	2.720	.906
PE	2.868	2.852	.374	1.501	3.891	.992
VS	2.653	1.522	.374	1.192	2.713	.992
Correlations	$r_{M,t+1}^e$	TY_{t+1}	PE_{t+1}	VS_{t+1}		
$r_{M,t+1}^e$	1	.071	-.006	-.030		
TY_{t+1}	.071	1	-.253	.423		
PE_{t+1}	-.006	-.253	1	-.320		
VS_{t+1}	-.030	.423	-.320	1		
$r_{M,t}^e$	.103	.065	.070	-.031		
TY_t	.070	.906	-.248	.420		
PE_t	-.090	-.263	.992	-.318		
VS_t	-.025	.425	-.322	.992		

Table 2: VAR parameter estimates

The table shows the OLS parameter estimates for a first-order VAR model including a constant, the log excess market return (r_M^e), term yield spread (TY), price-earnings ratio (PE), and small-stock value spread (VS). Each set of three rows corresponds to a different dependent variable. The first five columns report coefficients on the five explanatory variables, and the remaining columns show R^2 and F statistics. OLS standard errors are in square brackets and bootstrap standard errors in parentheses. Bootstrap standard errors are computed from 2500 simulated realizations. The table also reports the correlation matrix of the shocks with shock standard deviations on the diagonal, labeled “corr/std.” Sample period for the dependent variables is 1929:1-2001:12, 876 monthly data points.

	constant	$r_{M,t}^e$	TY_t	PE_t	VS_t	R^2 %	F
$r_{M,t+1}^e$	.062 [.020] (.026)	.094 [.033] (.034)	.006 [.003] (.003)	-.014 [.005] (.007)	-.013 [.006] (.008)	2.57	5.34
TY_{t+1}	.046 [.097] (.012)	.046 [.165] (.170)	.879 [.016] (.017)	-.036 [.026] (.031)	.082 [.028] (.036)	82.41	1.02×10^3
PE_{t+1}	.019 [.013] (.017)	.519 [.022] (.022)	.002 [.002] (.002)	.994 [.004] (.004)	-.003 [.004] (.005)	99.06	2.29×10^4
VS_{t+1}	.014 [.017] (.024)	-.005 [.029] (.028)	.002 [.003] (.003)	.000 [.005] (.006)	.991 [.005] (.008)	98.40	1.34×10^4
corr/std	$r_{M,t+1}^e$	TY_{t+1}	PE_{t+1}	VS_{t+1}			
$r_{M,t+1}^e$	.055 (.003)	.018 (.048)	.777 (.018)	-.052 (.052)			
TY_{t+1}	.018 (.048)	.268 (.013)	.018 (.039)	-.012 (.034)			
PE_{t+1}	.777 (.018)	.018 (.039)	.036 (.002)	-.086 (.045)			
VS_{t+1}	-.052 (.052)	-.012 (.034)	-.086 (.045)	.047 (.003)			

Table 3: Cash-flow and discount-rate news for the market portfolio

The table shows the properties of cash-flow news (N_{CF}) and discount-rate news (N_{DR}) implied by the VAR model of Table 2. The upper-left section of the table shows the covariance matrix of the news terms. The upper-right section shows the correlation matrix of the news terms with standard deviations on the diagonal. The lower-left section shows the correlation of shocks to individual state variables with the news terms. The lower right section shows the functions ($e1' + e1'\lambda, e1'\lambda$) that map the state-variable shocks to cash-flow and discount-rate news. We define $\lambda \equiv \rho\Gamma(I - \rho\Gamma)^{-1}$, where Γ is the estimated VAR transition matrix from Table 2 and ρ is set to .95 per annum. r_M^e is the excess log return on the CRSP value-weight index. TY is the term yield spread. PE is the log ratio of S&P 500's price to S&P 500's ten-year moving average of earnings. VS is the small-stock value-spread, the difference in log book-to-markets of value and growth stocks. Bootstrap standard errors (in parentheses) are computed from 2500 simulated realizations.

News covariance	N_{CF}	N_{DR}	News corr/std	N_{CF}	N_{DR}
N_{CF}	.00064 (.00022)	.00015 (.00037)	N_{CF}	.0252 (.004)	.114 (.232)
N_{DR}	.00015 (.00037)	.00267 (.00070)	N_{DR}	.114 (.232)	.0517 (.007)
Shock correlations	N_{CF}	N_{DR}	Functions	N_{CF}	N_{DR}
r_M^e shock	.352 (.224)	-.890 (.036)	r_M^e shock	.602 (.060)	-.398 (.060)
TY shock	.128 (.134)	.042 (.081)	TY shock	.011 (.013)	.011 (.013)
PE shock	-.204 (.238)	-.925 (.039)	PE shock	-.883 (.104)	-.883 (.104)
VS shock	-.493 (.243)	-.186 (.152)	VS shock	-.283 (.160)	-.283 (.160)

Table 4: Cash-flow and discount-rate betas in the early sample

The table shows the estimated cash-flow ($\hat{\beta}_{CF}$) and discount-rate betas ($\hat{\beta}_{DR}$) for the 25 ME - and BE/ME -sorted portfolios and 20 risk-sorted portfolios. “Growth” denotes the lowest BE/ME , “value” the highest BE/ME , “small” the lowest ME , and “large” the highest ME stocks. $\hat{b}_{VS}$, $\hat{b}_{TY}$, and $\hat{b}_{rM}$ are past return-loadings on value-spread shock, term-yield shock, and market-return shock. “Diff.” is the difference between the extreme cells. Standard errors [in brackets] are conditional on the estimated news series. Estimates are for the 1929:1-1963:6 period.

$\hat{\beta}_{CF}$	Growth		2		3		4		Value		Diff.	
Small	.53	[.11]	.46	[.09]	.40	[.08]	.42	[.07]	.49	[.08]	-.04	[.07]
2	.30	[.06]	.34	[.06]	.36	[.06]	.38	[.06]	.45	[.08]	.16	[.04]
3	.30	[.06]	.28	[.05]	.31	[.06]	.35	[.06]	.47	[.08]	.18	[.04]
4	.20	[.05]	.26	[.05]	.31	[.05]	.35	[.07]	.50	[.09]	.30	[.05]
Large	.20	[.05]	.19	[.05]	.28	[.06]	.33	[.07]	.40	[.09]	.19	[.06]
Diff.	-.33	[.09]	-.26	[.06]	-.12	[.05]	-.09	[.04]	-.10	[.04]		
$\hat{\beta}_{DR}$	Growth		2		3		4		Value		Diff.	
Small	1.32	[.18]	1.46	[.19]	1.32	[.15]	1.27	[.14]	1.27	[.15]	-.06	[.15]
2	1.04	[.11]	1.15	[.11]	1.09	[.11]	1.25	[.11]	1.25	[.13]	.21	[.08]
3	1.13	[.10]	1.01	[.08]	1.08	[.09]	1.05	[.10]	1.27	[.12]	.14	[.06]
4	.87	[.07]	.97	[.08]	.97	[.09]	1.06	[.10]	1.36	[.13]	.49	[.10]
Large	.88	[.07]	.82	[.07]	.87	[.08]	1.06	[.09]	1.18	[.12]	.31	[.10]
Diff.	-.45	[.15]	-.64	[.15]	-.43	[.10]	-.21	[.08]	-.08	[.10]		
$\hat{\beta}_{CF}$	Lo $\hat{b}_{rM}$		2		3		4		Hi $\hat{b}_{rM}$		Diff.	
Lo $\hat{b}_{VS}$	.21	[.04]	.25	[.05]	.31	[.06]	.37	[.07]	.45	[.09]	.25	[.05]
Hi $\hat{b}_{VS}$	.15	[.03]	.19	[.04]	.25	[.06]	.28	[.06]	.37	[.08]	.22	[.05]
Lo $\hat{b}_{TY}$	.18	[.04]	.21	[.05]	.26	[.06]	.31	[.07]	.41	[.08]	.23	[.04]
Hi $\hat{b}_{TY}$	.16	[.04]	.21	[.04]	.27	[.05]	.32	[.06]	.40	[.08]	.24	[.05]
$\hat{\beta}_{DR}$	Lo $\hat{b}_{rM}$		2		3		4		Hi $\hat{b}_{rM}$		Diff.	
Lo $\hat{b}_{VS}$	.73	[.06]	.87	[.07]	1.04	[.09]	1.20	[.11]	1.46	[.13]	.73	[.09]
Hi $\hat{b}_{VS}$	.64	[.05]	.75	[.07]	.96	[.08]	1.09	[.09]	1.30	[.11]	.66	[.08]
Lo $\hat{b}_{TY}$	.73	[.06]	.85	[.07]	1.00	[.09]	1.17	[.10]	1.38	[.12]	.64	[.08]
Hi $\hat{b}_{TY}$	.65	[.06]	.76	[.06]	.88	[.08]	1.09	[.10]	1.34	[.12]	.69	[.09]

Table 5: Cash-flow and discount-rate betas in the modern sample

The table shows the estimated cash-flow ($\hat{\beta}_{CF}$) and discount-rate betas ($\hat{\beta}_{DR}$) for the 25 ME - and BE/ME -sorted portfolios and 20 risk-sorted portfolios. “Growth” denotes the lowest BE/ME , “value” the highest BE/ME , “small” the lowest ME , and “large” the highest ME stocks. $\hat{b}_{VS}$, $\hat{b}_{TY}$, and $\hat{b}_{rM}$ are past return-loadings on value-spread shock, term-yield shock, and market-return shock. “Diff.” is the difference between the extreme cells. Standard errors [in brackets] are conditional on the estimated news series. Estimates are for the 1963:7-2001:12 period.

$\hat{\beta}_{CF}$	Growth		2		3		4		Value		Diff.	
Small	.06	[.07]	.07	[.06]	.09	[.05]	.09	[.04]	.13	[.04]	.07	[.04]
2	.04	[.06]	.08	[.05]	.10	[.04]	.11	[.04]	.12	[.04]	.09	[.03]
3	.03	[.05]	.09	[.04]	.11	[.04]	.12	[.03]	.13	[.04]	.09	[.04]
4	.03	[.05]	.10	[.04]	.11	[.03]	.11	[.03]	.13	[.04]	.10	[.04]
Large	.03	[.04]	.08	[.03]	.09	[.03]	.11	[.03]	.11	[.03]	.09	[.03]
Diff.	-.03	[.05]	.02	[.05]	-.01	[.04]	.02	[.04]	-.01	[.04]		
$\hat{\beta}_{DR}$	Growth		2		3		4		Value		Diff.	
Small	1.66	[.13]	1.37	[.11]	1.18	[.10]	1.12	[.09]	1.12	[.10]	-.54	[.08]
2	1.54	[.11]	1.22	[.09]	1.07	[.08]	.96	[.08]	1.03	[.09]	-.52	[.08]
3	1.41	[.10]	1.11	[.08]	.95	[.08]	.82	[.07]	.94	[.09]	-.47	[.09]
4	1.27	[.09]	1.05	[.08]	.89	[.07]	.79	[.07]	.87	[.08]	-.41	[.09]
Large	1.00	[.07]	.87	[.07]	.74	[.06]	.63	[.07]	.68	[.07]	-.33	[.08]
Diff.	-.66	[.12]	-.50	[.11]	-.44	[.10]	-.49	[.09]	-.44	[.08]		
$\hat{\beta}_{CF}$	Lo $\hat{b}_{rM}$		2		3		4		Hi $\hat{b}_{rM}$		Diff.	
Lo $\hat{b}_{VS}$	.09	[.03]	.08	[.03]	.10	[.04]	.10	[.04]	.12	[.05]	.04	[.04]
Hi $\hat{b}_{VS}$	.06	[.03]	.06	[.03]	.07	[.04]	.05	[.05]	.06	[.06]	-.01	[.04]
Lo $\hat{b}_{TY}$	.06	[.03]	.04	[.03]	.08	[.04]	.08	[.04]	.06	[.06]	.00	[.04]
Hi $\hat{b}_{TY}$	.09	[.03]	.07	[.03]	.09	[.03]	.08	[.04]	.10	[.05]	.00	[.04]
$\hat{\beta}_{DR}$	Lo $\hat{b}_{rM}$		2		3		4		Hi $\hat{b}_{rM}$		Diff.	
Lo $\hat{b}_{VS}$	.57	[.06]	.77	[.06]	.88	[.07]	1.12	[.08]	1.40	[.09]	.82	[.08]
Hi $\hat{b}_{VS}$	.67	[.06]	.85	[.07]	1.06	[.07]	1.30	[.09]	1.58	[.11]	.91	[.11]
Lo $\hat{b}_{TY}$	.73	[.07]	.86	[.07]	1.05	[.07]	1.23	[.08]	1.60	[.12]	.87	[.10]
Hi $\hat{b}_{TY}$	.61	[.06]	.79	[.06]	.91	[.06]	1.11	[.07]	1.39	[.09]	.78	[.08]

Table 6: Asset pricing tests for the early sample

The table shows premia estimated from the 1929:1-1963:6 sample for an unrestricted factor model, the two-beta ICAPM, and the CAPM. The test assets are the 25 *ME*- and *BE/ME*- sorted portfolios and 20 risk-sorted portfolios. The second column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}). Estimates are from a cross-sectional regression of average simple excess test-asset returns (monthly in fractions) on an intercept and estimated cash-flow ($\hat{\beta}_{CF}$) and discount-rate betas ($\hat{\beta}_{DR}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporating full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5% critical value.

Parameter	Factor model		Two-beta ICAPM		CAPM	
R_{zb} less R_{rf} (g_0)	.0042	0	.0023	0	.0023	0
% per annum	4.98%	0%	2.76%	0%	2.74%	0%
Std. err. A	[.0032]	N/A	[.0024]	N/A	[.0028]	N/A
Std. err. B	(.0029)	N/A	(.0030)	N/A	(.0028)	N/A
$\hat{\beta}_{CF}$ premium (g_1)	.0173	.0069	.0083	.0148	.0051	.0067
% per annum	20.76%	8.22%	9.91%	17.80%	6.11%	8.00%
Std. err. A	[.0231]	[.0221]	[.0167]	[.0175]	[.0046]	[.0034]
Std. err. B	(.0266)	(.0248)	(.0221)	(.0442)	(.0046)	(.0034)
$\hat{\beta}_{DR}$ premium (g_2)	-.0003	.0066	.0041	.0041	.0051	.0067
% per annum	-.41%	7.93%	4.95%	4.95%	6.11%	8.00%
Std. err. A	[.0092]	[.0067]	[.0006]	[.0006]	[.0046]	[.0034]
Std. err. B	(.0088)	(.0071)	(.0006)	(.0006)	(.0046)	(.0034)
$\hat{R}^2$	48.08%	40.26%	45.85%	37.98%	44.52%	40.26%
Pricing error	.0117	.0126	.0119	.0133	.0127	.0126
5% critic. val. A	[.019]	[.024]	[.024]	[0.033]	[.021]	[.027]
5% critic. val. B	(.019)	(.024)	(.031)	(0.099)	(.021)	(.027)

Table 7: Asset pricing tests for the modern sample

The table shows premia estimated from the 1963:7-2001:12 sample for an unrestricted factor model, the two-beta ICAPM, and the CAPM. The test assets are the 25 *ME*- and *BE/ME*- sorted portfolios and 20 risk-sorted portfolios. The second column per model constrains the zero-beta rate (R_{zb}) to equal the risk-free rate (R_{rf}). Estimates are from a cross-sectional regression of average simple excess test-asset returns (monthly in fractions) on an intercept and estimated cash-flow ($\hat{\beta}_{CF}$) and discount-rate betas ($\hat{\beta}_{DR}$). Standard errors and critical values [A] are conditional on the estimated news series and (B) incorporating full estimation uncertainty of the news terms. The test rejects if the pricing error is higher than the listed 5% critical value.

Parameter	Factor model		Two-beta ICAPM		CAPM	
R_{zb} less R_{rf} (g_0)	.0009	0	-.0009	0	.0069	0
% per annum	1.05%	0%	-1.04%	0%	8.24%	0%
Std. err. A	[.0029]	N/A	[.0031]	N/A	[.0026]	N/A
Std. err. B	(.0033)	N/A	(.0031)	N/A	(.0026)	N/A
$\hat{\beta}_{CF}$ premium (g_1)	.0529	.0572	.0575	.0483	-.0007	.0051
% per annum	63.47%	68.59%	69.04%	57.92%	-.83%	6.10%
Std. err. A	[.0178]	[.0163]	[.0182]	[.0272]	[.0034]	[.0023]
Std. err. B	(.0325)	(.0444)	(.0262)	(.0423)	(.0034)	(.0023)
$\hat{\beta}_{DR}$ premium (g_2)	.0007	.0012	.0020	.0020	-.0007	.0051
% per annum	.88%	1.44%	2.43%	2.43%	-.83%	6.10%
Std. err. A	[.0033]	[.0031]	[.0002]	[.0002]	[.0034]	[.0023]
Std. err. B	(.0085)	(.0099)	(.0002)	(.0002)	(.0034)	(.0023)
$\hat{R}^2$	52.10%	51.59%	49.26%	47.41%	3.10%	-61.57%
Pricing error	.0271	.0269	.0272	.0275	.0592	.0875
5% critic. val. A	[.028]	[.042]	[.051]	[.314]	[.032]	[.046]
5% critic. val. B	(.030)	(.071)	(.051)	(.488)	(.032)	(.046)

Figure 1: Time-series evolution of the predictor variables.

This figure plots the time-series of three predictor variables: (1) The log ratio of price to a ten-year moving average of earnings, marked with a solid line; (2) the small-stock value spread, marked with line and squares; and (3) the term yield spread, marked with dashed line and triangles. All variables are demeaned and normalized by their sample standard deviations. The sample period is 1928:12-2001:12.

Figure 2: Cash-flow and discount-rate recessions.

This figure plots the cash-flow news and negative of discount-rate news, smoothed with a trailing exponentially-weighted moving average. The decay parameter is set to .08 per month, and the smoothed news series are generated as $MA_t(N) = .08N_t + (1 - .08)MA_{t-1}(N)$. The dotted vertical lines denote NBER business-cycle troughs. The sample period is 1929:1-2001:12.

Figure 3: Performance of the CAPM and ICAPM, 1929:1-1963:6.

The four diagrams correspond to (clockwise from the top left) the ICAPM with a free zero-beta rate, the ICAPM with the zero-beta rate constrained to the risk-free rate, the CAPM with a constrained zero-beta rate, and the CAPM with an unconstrained zero-beta rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns. The predicted values are from regressions presented in Table 6. Triangles denote the 25 ME and BE/ME portfolios and asterisks the 20 risk-sorted portfolios.

Figure 4: Performance of the CAPM and ICAPM, 1963:7-2001:12.

The four diagrams correspond to (clockwise from the top left) the ICAPM with a free zero-beta rate, the ICAPM with the zero-beta rate constrained to the risk-free rate, the CAPM with a constrained zero-beta rate, and the CAPM with an unconstrained zero-beta rate. The horizontal axes correspond to the predicted average excess returns and the vertical axes to the sample average realized excess returns. The predicted values are from regressions presented in Table 7. Triangles denote the 25 ME and BE/ME portfolios and asterisks the 20 risk-sorted portfolios.

Figure 5: Conditional risk premia for cash-flow and discount-rate betas.

We show the smoothed conditional premium on β_{CF} (top line) and β_{DR} (bottom line), both scaled by the market's conditional volatility. The horizontal lines are time-series averages. First, we run three sets of 45 time-series regressions on a constant, time trend, and the lagged VAR state variables, where the dependent variables are (1) excess return on the test assets ($R_{i,t}^e$), (2) $(N_{CF,t} + N_{CF,t-1})R_{i,t}^e$, and (3) $(N_{DR,t} + N_{DR,t-1})R_{i,t}^e$. Then, each month, we regress the fitted values of (1) on the fitted values of (2) and (3), and plot the five-year moving averages of these cross-sectional coefficients.