

The Factor Structure of Time-Varying Discount Rates

Victoria Atanasov, Ilan Cooper, Richard Priestley, and Junhua Zhong*

January 2018

Abstract

Discount rate variation is driven by a short run business cycle component and a longer run trend component. Unlike long horizon investors, short term investors are not compensated with the improved long term prospects for poor returns driven by long term discount rate shocks. Thus, heterogeneity of investors' horizon implies that long term discount rate news command a higher price of risk. We test a three factor model and find that both discount rate shocks, as well as cash flow shocks, are priced. Consistent with the intuition above, long term discount rates earn a substantially higher price of risk.

*Atanasov is from the Finance Department, University of Mannheim, email: atanasov@uni-mannheim.de. Cooper and Priestley are from the Department of Finance, BI Norwegian Business School, emails: ilan.cooper@bi.no and richard.priestley@bi.no. Zhong is from the Research Center of Applied Finance, Dongbei University of Finance and Economics, email: zhongjunhua@dufe.edu.cn. Parts of this research were conducted while Atanasov was at the Department of Finance, VU University Amsterdam and Zhong at the Department of Finance, BI Norwegian Business School. We would like to thank Jules van Binsbergen (SFS Cavalcade discussant), Nicole Branger, Jaewon Choi, Amit Goyal, Stefano Giglio, Yufeng Han (MFA discussant), Dana Kiku, Hugues Langlois (EFA discussant), Paulo Maio, Ernst Maug, Michael Semenishev (Swiss Society for Financial Market Research discussant), Andrea Tamoni, Erik Theissen, Michael Weber, and conference participants at the 2016 SFS Finance Cavalcade, the 2016 European Finance Association annual meeting, the 2016 Midwest Finance Association annual meeting, the 2016 Swiss Society for Financial Market Research annual conference, and seminar participants at the Brown Bag Seminar at the University of Mannheim and the Macro and Econometrics Seminar at the University of Heidelberg for helpful comments and suggestions. We are grateful to Kenneth R. French, Amit Goyal, Martin Lettau, and Robert J. Shiller for graciously making their data available.

JEL Classification: G12, G14

Keywords: discount rates, risk factors, term structure of equity risk premium

1 Introduction

Discount rates are time-varying (Cochrane (2011)). This statement is supported by a host of variables that have been shown to predict returns.¹ The extant literature has focused on which variables better predict returns, why certain variables predict returns, and methodologies to assess in-sample and out-of sample predictability at different horizons. These questions have addressed the issue of how much discount rates vary over time.

In this paper, we focus on the information content of different predictor variables of the short end and long end of the term structure of discount rates, and subsequently ask a new question: how does this translate into a factor structure for expected returns? This question is raised in Cochrane's (2011) American Finance Association Presidential address where he asks "what is the factor structure of time varying expected returns?" and he emphasizes that "As we pursue the multivariate forecasting question using the large number of additional forecasting variables, we should look at pricing implications, and not just run short-run R2 contests". This paper is an attempt to address this question.

If there exists discount rate variation it will lead to state-variable hedging. Campbell and Vuolteenaho (2004) estimate a two factor intertemporal capital asset pricing model (ICAPM) where news about discount rates and cash flows is important in describing the

¹Stock return predictability is now established as a fact. Campbell and Shiller (1988), Fama and French (1988), and Lamont (1998) and Rangvid (2005) document stock return predictability using either dividends, earnings or GDP, scaled by prices. Campbell (1987), Fama and French (1989), Hodrick (1992), and Keim and Stambaugh (1986) show that stock returns are predictable using interest rate variables. Lettau and Ludvigson (2001) show that aggregate consumption, asset holdings, and labor income share a common trend, but may deviate substantially from one another in the short term, yielding stock return predictability. Cooper and Priestley (2009) show that the output gap can predict stock returns at business cycle frequencies. Santos and Veronesi (2006), Piazzesi, Schneider and Tuzel (2006), and Lustig and Van Nieuwerburgh (2006) develop models where either labor income or housing wealth variables can predict stock returns.

cross section of stock returns. However, the empirical evidence suggests that discount rate variation is different in the short-run and in the long-run. Figure 1 provides a simple means to understand this. Motivated by Cochrane (2011), we plot the actual returns and fitted values from a regression of returns on lagged values of Lettau and Ludvigson's (2001) consumption-wealth ratio, cay , and the log dividend price ratio, dp , using one-quarter, ten-year, fifteen-year, and twenty-year returns. Up to the five-year horizon, cay dominates the return forecast and compared to forecasts based on dp , adds volatility to predicted returns. By contrast, at longer horizons cay loses its forecasting power whereas dp gains in importance and the volatility of stock return becomes lower. This suggests a term structure of discount rates where there are more volatile business cycle variations in discount rates at short horizons and smoother variations in the long term trend of discount rates.

What are the pricing implications of this pattern in discount rate variation? The ICAPM assumes long term investors and implies that any component of the term structure of discount rates will have the same price of risk.² However, if investors have heterogeneous investment horizons then the marginal investor is not a Merton's ICAPM long term investor. If investors care about discount rate changes at different horizons, then revisions in forecasts of both future short term and long term returns would be priced differently. Both short term discount rate shocks and long term discount rate shocks lead to a fall in stock prices. This fall in prices is compensated by a rise in expected market returns. While short term discount rate shocks lead to an expected quick recovery of prices, long term discount rate shocks might be perceived as riskier by short term investors, because prices are expected to rise only in the long term, beyond their investment horizon. Thus, the presence of short term investors implies that exposure to

²Campbell and Vuolteenaho (2004) note that Merton's (1973) intertemporal CAPM (ICAPM) suggests that the price of risk of both short term discount rates and long term discount rates should equal the variance of the market return, while the price of risk of cash flows should be γ times larger, where γ is the investor's relative risk aversion.

long term discount rate shocks should be compensated by a larger risk premium than exposure to short term discount rate shocks.

Recent evidence uncovered by Betermier, Calvet, and Soldini (2017) supports our conjecture that short term investors demand a larger compensation for exposure to long term discount rate fluctuations. Specifically, these authors present evidence that over the life-cycle households climb the "value ladder". That is, they gradually shift from growth stocks (with long duration) to value stocks (with short duration) as their investment horizon shortens. Being assets with long duration, growth stocks are more sensitive to discount rate shocks. Thus, short term investors appear to be more averse to exposure to discount rate shocks because these investors do not benefit from the improvement in investment opportunities driven by (especially long term) discount rate shocks.

The different pricing implications of discount rate variations at different horizons gives rise to a three factors in the cross section of expected excess returns: a short run discount rate news factor, a long run discount rate news factor, and a cash flow news factor. This paper's contribution is twofold. First, as its main contribution, the paper investigates empirically if discount rate variation at different horizons is a source of different prices of risk. Second, the paper examines the implications of the cross sectional findings for the term structure of 'zero coupon' equity, namely dividend strips. That is, we investigate and confirm that our evidence of different prices of risk for short term discount rates news and long term discount rate news can be theoretically reconciled with recent evidence of a downward sloping term structure of 'zero coupon' equity ((Giglio, Maggiori, and Stroebel (2015), Binsbergen and Kojen (2012, 2017), Weber (2017)).

The paper's main findings can be summarized as follows. There are two separate sources of discount rate risk that emanate from short term return predictability and long term return predictability. Both discount rate news components derived from a VAR for stock returns that picks out these two effects separately are important in describing

the cross section of stock returns. We find that news about long term discount rates commands a price of risk of 5.1% per annum which is economically and statistically greater than the price of risk associated with short term discount rate news. The latter is rewarded with a premium of approximately 3% per annum. The estimated risk premium attached to cash flow news is 5.7% per annum. Our findings imply that the "good beta" of Campbel and Vuolteenaho (2004) is "better", that is, it entails a lower risk premium, when it is a beta with respect to short term discount rate shocks.

Using data from derivatives markets to recover the prices of dividend strips on the aggregate stock market, van Binsbergen, Brandt, and Kojen (2012) find that short-term dividends have a higher risk premium than long-term dividends. At first glance, our findings might be viewed as being consistent only with an upward sloping term structure of dividend strips and not with a downward sloping one. However, we derive and simulate an augmented version of the risk-based model of Lettau and Wachter (2007) and find that this is not necessarily the case. The augmented model includes shocks to dividend growth, expected dividend growth, short term discount rates, and long term discount rates. The model replicates our results that the price of long-term discount rate risk is higher than its short-term counterpart. The simulations also produce a downward sloping term structure of 'zero coupon' equity risk premia and their Sharpe ratios, consistent with the empirical evidence in van Binsbergen, Brandt, and Kojen (2012).

The value of the market portfolio is the sum of all dividend strip values. The discount rate shocks we identify in our empirical framework are, in essence, shocks to the expected holding period returns on the market portfolio. It is important to note that the shocks we identify are distinct from shocks to discount rates on dividend strips. The expected market return is a weighted average of expected dividend strip returns. Unexpected changes in the expected market return can therefore stem from shocks to the value of individual strips, which themselves are caused by either changes in the expected dividends, or by

changes in the discount rates on strips. For example, a positive shock to the two-year holding period return on the market could be a consequence of a positive shock to the one-year strip discount rate but also to the two-year strip expected return, the three-year strip and so on. Thus, our findings that the price of risk of shocks to long term market discount rates is higher than that of short term market discount rates is consistent with any shape of the term structure of dividend strips, and in particular with a downward sloping term structure, as our simulations also show.

Interestingly, our main result, namely that shocks to long term discount rates command a higher risk premium than shocks to short term discount rates does not imply an upward sloping term structure of expected holding period returns for the market portfolio either. This is evident in our model simulation which produce a rather flat term structure of the market's holding period returns.

Our empirical findings are reinforced when we sort stocks into portfolios based on their exposures to different types of news. For example, the average return differential between the 5th (high exposure) and the 1st (low exposure) quintiles of stocks sorted on exposure to longer term discount rate news is about 4.4% and the corresponding number for short term discount rate news is close to 3% in annual terms. Double sorting portfolios by their exposure to short and long term discount rate risk confirms that exposure to long term discount rate risk is associated with higher expected returns than exposure to short term discount rate risk.

We also find that the cross sectional estimates of the prices of risk differ across size and book to market samples. The price of risk of both discount rate news components is high for small and high book to market firms, and low for large and low book to market firms. Within each category of stocks sorted on their characteristics, we find a pervasive long term discount rate price of risk which is economically and statistically higher than its short term counterpart. Interestingly, we find a particularly strong relation between

firm size and short term discount rate risk. While the price of risk associated with short term discount rate news is greater for the lowest market equity firm quintile than it is in the cross section of all stocks, it is typically statistically indistinguishable from zero for the other market equity firm quintiles. Our results hence suggest that the price of risk associated with short term discount rate news is mostly due to small firms in the sample. This result is consistent with the findings of Perez-Quiros and Timmermann (2000) who show that small firms are more sensitive to credit market conditions in recessions, that is, at business cycle frequencies, when discount rates are high. In contrast, exposure to long run discount rate shocks is priced across all size quintiles.

The economic magnitude of different news components can be judged according to their expected return contributions. These are obtained by multiplying the risk premium for each factor with the appropriate beta. We find that each of the three factors has a positive and statistically significant return contribution. Similar to the patterns we observe in the risk premia, the long term discount rate news has a higher expected return contribution than the short term discount rate news. This pattern becomes more evident when we consider portfolios of stocks sorted on their characteristics. The impact of the long term discount rate factor is higher than that of the short term discount rate factor in each size and book to market equity quintile apart from the 1st (lowest) market equity and 5th (highest) book-to-market equity quintiles of stocks. This indicates that small and value firms have smaller sensitivity to long term discount rate news as opposed to short term discount rate news. In this vein, Lettau and Wachter (2007) show that firms with cash flows weighted more to the present have a low ratio of price to fundamentals and high expected returns relative to firms with cash flows weighted more to the future.

Since discount rates are not observable, we draw on recent popular studies and employ a VAR approach to model the empirical proxies for long term and short term news.³

³Well known examples include Campbell (1991), Campbell and Ammer (1993), Campbell and Mei

Despite its numerous advantages and intuitive appeal, this approach may be sensitive to the choice of sample periods or predictive variables (Chen and Zhao (2009)). We address this concern in a battery of robustness checks and consider nine alternative combinations of various state variables, employ alternative news estimation methods, and examine different sub samples. We find that our results are consistent and remain generally upheld.

The literature has thus far, to the best of our knowledge, ignored the asset pricing implications of the empirical fact that stock returns are predictable at short and long horizons with different variables, suggesting a term structure in discount rates. We exploit this in order to show that there are two discount rate factors and one cash flow factor driving the cross section of stock returns. A thorough understanding of discount rate behavior is important to connect the findings that stock returns are predictable at short and long horizons with risk factors that describe the cross section of returns. Our findings provide new insights into discount rate variation and its impact on asset prices.

In a related paper to ours Weber (2016) creates a firm-level cash flow duration measure using the balance sheet, and finds that stocks with a high cash flow duration earn substantially lower average returns than stocks with short cash flow duration, especially following periods of high investor sentiment. Weber's methodology is very different from ours as we look at shocks to the expected market returns and test the prices of risk of these shocks. As we explain above, our findings do not imply any particular shape for the term structure of dividend strips.

The remainder of the paper is organized as follows. To motivate our analysis, Section 2 shows that there are different variables which predict returns at short horizons and at long horizons. Section 3 outlines the methodology to extract short-term and long-term

(1993), Campbell (1996), Campbell and Vuolteenaho (2004), Campbell, Giglio, and Polk (2013), Campbell, Giglio, Polk, and Turley (2017), and Petkova (2006) among others.

discount-rate news components. Section 4 describes the data. Section 5 discusses our empirical findings, Section 6 contains a brief overview of the robustness checks. Section 7 briefly describes the dynamic risk-based model, and provides the simulative results. Section 8 concludes.

2 Short-Term and Long-Term Return Predictors

In this section, we examine the predictive power of two theoretically motivated economy-wide instruments, namely the log dividend price ratio, dp , and the consumption-wealth ratio, cay . We first outline the theoretical underpinning for the forecasting power of these two variables. Then we examine empirically their forecasting power for US equity stock returns over the period from 1952 to 2013.

2.1 Theoretical Motivation

The theoretical rationale for the predictive ability of dividend yields dates back to the late 1980s. Using a first-order Taylor expansion, Campbell and Shiller (1988) approximate the log one-period return, $r_t = \ln(P_{t+1} + D_{t+1}) - \ln(P_t)$, where P_t is the price, and D_t is the dividend, around the mean log dividend-price ratio, $(\overline{d_t - p_t})$, to show that log returns can be approximated as $r_t \approx k + \rho p_{t+1} + (1 - \rho)d_{t+1} - p_t$, where ρ and k are parameters in the linearization defined by $\rho \equiv 1 / (1 + \exp(\overline{d_t - p_t}))$ and $k \equiv -\log \rho - (1 - \rho) \log(1/\rho - 1)$ and lowercase letters are used for logs. Solving forward this identity iteratively and imposing a transversality condition that $\lim_{j \rightarrow \infty} \rho^j (d_{t+j} - p_{t+j}) = 0$, the authors manifest that the log price-dividend ratio is determined by the expected discounted value of future dividend growth and returns:

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

where E_t denotes a rational expectation formed at the end of period t and Δ is a one-period backward difference. This approximation becomes more accurate at longer horizons.

Equation (1) says that the log price-dividend ratio is high when dividends are expected to grow rapidly, or when future stock returns are expected to be low. Time variability in dividend yields can hence be attributed to the variation in expected cash flow growth rates or expected future risk premia. Because the dividend yield is a weak forecaster of dividend growth, changes in dividend yield are commonly related to revisions in expectations about future returns (Campbell (1991) and Cochrane (1992)). Campbell and Shiller (1988), Fama and French (1988) and many others document the forecasting potential of dividend yields for excess returns, in particular at longer horizons.⁴

Campbell and Mankiw (1989) apply a similar approximation to the intertemporal budget constraint of a consumer, $W_{t+1} = (1 + R_{t+1})(W_t - C_t)$, where C_t is consumption, W_t is wealth, and R_t is a time-varying risky return. Assuming that the consumption-wealth ratio is stationary, Campbell and Mankiw (1989) rewrite the budget constraint as $\Delta w_{t+1} \approx \kappa + r_{t+1} + (1 - 1/\delta)(c_t - w_t)$, where δ is the steady-state ratio of new investment to total wealth, $(W - C)/W$, and κ is a constant. By solving this difference equation forward, the following expression for the log consumption-wealth ratio obtains:

$$c_t - w_t \approx E_t \sum_{j=1}^{\infty} \delta^j (r_{t+j} - \Delta c_{t+j}) + \frac{\delta \kappa}{1 - \delta}. \quad (2)$$

This equation highlights the forward-looking nature of the consumption-wealth ratio by

⁴Keim and Stambaugh (1986), Fama and French (1988) and Goyal and Welch (2003) among others find that aggregate dividend yields strongly predict equity and bond returns. Ang and Bekaert (2007) warn, however, that the statistical evidence for return predictability can depend critically on the choice of standard errors and show that the predictive ability of the dividend yield is considerably enhanced in bivariate regressions with the short rate.

relating it to future expected returns on aggregate wealth (or on the market portfolio) and consumption growth rates.

Despite its theoretical purity, the formulation in equation (2) is not straightforward to apply in the empirical sense because aggregate wealth and human wealth, in particular, cannot be directly observed. Lettau and Ludvigson (2001) overcome this obstacle by linking the stock of human wealth to labor income. They show that consumption, asset wealth and labour income share a common trend, but may deviate from each other in the short run. The residual from a cointegrating relation between these variables labelled as *cay*, captures the predictive component for future returns. The economic intuition is simple. When future returns are expected to be high, investors who wish to maintain smooth consumption intertemporally will increase their consumption out of current wealth and labour income, which will shift the level of consumption above its common trend with the wealth components. Lettau and Ludvigson (2001) show that *cay* is a strong predictor of both real stock returns and excess returns over a Treasury bill rate. Specifically, *cay* outperforms the dividend yield, the dividend payout ratio, and several other popular forecasting variables at short and intermediate horizons.

2.2 Predictive Regressions

Previous work indicates that *cay* captures short-term discount rate movements and has explanatory power for short-horizon returns, but very little for long-term returns and long-term dividend growth (Lettau and Ludvigson (2005)). Cochrane (2011) conjectures that *cay* helps predict one-year returns without much changing long-term expected returns, if it has an offsetting effect on returns with horizons longer than one-year. Cochrane (2011) also shows that shocks to *cay* bring about a shift in expected returns from the distant future to the near term and labels it the shift to the term structure of risk premia.

In contrast to *cay*, the dividend-price ratio dp , matches expected long-horizon returns

very well (Cochrane, 2011). It captures the long-term movement in equity returns, driven possibly by innovations in technology or demographic changes (see, for example, Lochester (2009) and Favero, Gozluklu, and Tamoni (2011)).

As formalized in equation (1), the dividend-price ratio is interpreted as reflecting the outlook for dividends and/or the rate at which future dividends are discounted to today's price. Changes in dp forecast significant persistent changes in expected stock returns. Campbell, Lo, and MacKinlay (1997) emphasize that dp is a better proxy for expectations of long-horizon returns than for expectations of short-horizon returns.

In the following, we investigate the relative predictive power of cay and dp for quarterly returns and returns at longer horizons. We consider univariate regressions of the form

$$r_{m,t+1}^{e,(h)} = a_0^{(h)} + a_1^{(h)} \times cay_t + \varepsilon_{t+1}^{(h)} \quad (3)$$

and

$$r_{m,t+1}^{e,(h)} = b_0^{(h)} + b_1^{(h)} \times dp_t + \varepsilon_{t+1}^{(h)}. \quad (4)$$

In equations (3) and (4), $r_{m,t+1}^{e,(h)}$ denotes the log excess return on the value-weighted CRSP index over the risk-free rate at time $t + 1$ over a horizon of h quarters, $a_1^{(h)}$ and $b_1^{(h)}$ are h horizon slope coefficients, and $a_0^{(h)}$ and $b_0^{(h)}$ are the respective intercepts.

Table 1 summarizes the results of long-horizon predictive regressions. For each time horizon, the table reports the regression coefficients along with their Hansen-Hodrick (1980) corrected t -statistics in parentheses. Our estimates in Column I of Table 1 illustrate that cay is a strong predictor of future market returns at short and medium term horizons. First, the forecasting power of cay is pronounced already for quarterly returns as indicated by the first row in Column I: Variation in the consumption-wealth ratio explains roughly 1.30% of changes in future market returns with a t -statistic of 2.01. As the horizon increases, the predictive ability of cay improves both in economic and

statistical terms. The coefficient of *cay* increases from 0.62 for quarterly returns to 6.79 for five-year returns and reaches its maximum at horizons of between five and six years. This increase in predictive power is associated with a rise in the adjusted $\overline{R^2}$ statistic. Second, although *cay* has a dramatic effect on short- to medium-term return forecasts, its predictive power diminishes for longer horizon returns. The coefficients in Column I become smaller for horizons beyond five years, and there are similar patterns in the *t*-values and the measures of fit. At horizons of about 20 years, the coefficient in the predictive regression in equation (3) becomes economically and statistically close to zero. Furthermore, the point estimates go the "wrong" way for longer period returns of more than 20 years. Then high consumption-wealth ratio signals low returns, in contrast to the theoretical prediction.

Column II of Table 1 presents the results of predictive regressions with the dividend-price ratio. At the quarterly horizon, *dp* captures less variation in expected returns than *cay*, and its statistical significance is low with a *t*-ratio of 1.59. The forecasting potential of the dividend-price ratio becomes more pronounced at longer horizons. The slope coefficients and adjusted $\overline{R^2}$ s rise with the horizons and the predictive power of *dp* persists also for holding period returns of more than 20 years. This result is in stark contrast to the evidence reported for *cay* whose forecasting power diminishes after five years. For example, *cay* explains up to 22% of the variation in future returns at short- and medium-term horizons whereas *dp* can capture up to 50% of changes in expected stock returns at horizons of about 18-20 years. Comparing these univariate forecasting regressions, *cay* tends to outperform *dp* at short horizons, they have similar forecasting power at medium-term horizons, and *dp* outperforms *cay* in the long horizon regressions.⁵

The different roles of *cay* and *dp* in forecasting stock returns are captured in Fig-

⁵We find similar evidence in forecasts based on bivariate regressions with *cay* and *dp* (unreported). The adjusted $\overline{R^2}$ s increase from slightly more than 2% for quarterly returns to close to 55% at horizons of about 15 years, and stagnate in the range of 40-50% for longer-term returns.

ure 1 which plots the actual and fitted excess returns for 1-, 40-, 60-, and 80-quarter horizons, respectively. The forecasting power of *cay* and *dp* is different in the following respects: *cay* helps to forecast business cycle frequency "wiggles", but has no effect on the "trend", similar to evidence reported in Cochrane (2011). This kind of predictive power of *cay* persists up to the horizon of five to six years. However, when the horizon is prolonged, the difference between *cay* and *dp* in terms of their predictive power becomes pronounced even stronger. At horizons of beyond 15 years, *cay* has virtually a zero effect on return forecasts. Over 15-year horizons and longer, adding *cay* does not improve the predictability relative to that based on *dp* alone. Based on unreported results, the fitted returns generated by univariate *dp* forecasts and bivariate forecasts with *cay* and *dp*, merely overlap.

The empirical evidence in this section provides us with a yardstick to correctly capture short-term and long-term variation in future expected market returns. In the following, we examine the impact of the time-varying structure of market discount rates on risk premia in the stock market.

3 Cash-Flow, Short-Term and Long-Term Discount-Rate Risks

3.1 Permanent and Transitory Return Components

A simple present value formula states that changes in asset prices should be associated with changes in expectations of future cash flows, time-varying discount rates, or both (Campbell and Shiller (1988)). Elaborating on this insight, Campbell (1991) shows that the unexpected market return can be decomposed into news about future dividend growth (cash flows) and news about future returns (discount rates):

$$\begin{aligned}
r_{m,t+1} - E_t r_{m,t+1} &\approx (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{m,t+1+j} \\
&\equiv \eta_{cf,t+1} - \eta_{dr,t+1},
\end{aligned} \tag{5}$$

where the term $\eta_{cf,t+1}$ denotes news about future cash flows, i.e. revisions in expectations about future dividend growth, and the term $\eta_{dr,t+1}$ denotes news about future discount rates, i.e. revisions in expectations about future returns. As argued by Campbell and Vuolteenaho (2004), these two components can be interpreted approximately as permanent and transitory shocks to wealth. Returns generated by cash flow news are associated with a capital gain/loss and are not reversed. By contrast, returns generated by discount rate news are offset in the future because a rise in discount rates leads to a capital loss and vice versa. Therefore, a conservative long-term investor is generally more adverse to cash-flow risk exposure than to risk stemming from unexpected changes in discount rates.

Campbell (1991) suggests using a vector autoregressive (VAR) model to measure cash flow and discount rate news in the data. Empirically, this approach implies that we first estimate the terms $E_t r_{m,t+1}$ and $(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{m,t+1+j}$ and then use the actual return realizations to back out the cash flow news in equation (5).

In our analysis, we follow Campbell and Vuolteenaho (2004) and assume that the data are generated by a first-order vector autoregressive (VAR) model

$$\mathbf{z}_{t+1} = \mathbf{a} + \mathbf{\Gamma} \mathbf{z}_t + \mathbf{u}_{t+1} \tag{6}$$

where \mathbf{z}_{t+1} is a m -by-1 state vector with r_{t+1} as its first element, \mathbf{a} and $\mathbf{\Gamma}$ are m -by-1 vector and m -by- m companion matrix of constant parameters, and \mathbf{u}_{t+1} is an i.i.d. m -by-1

vector of shocks.⁶

Given the process in (6), the discount-rate news can be written as

$$\eta_{dr,t+1} = \mathbf{e}\mathbf{1}'\rho\mathbf{\Gamma}\mathbf{\Phi}\mathbf{u}_{t+1}, \quad (7)$$

where $\mathbf{\Phi} \equiv (\mathbf{I} - \rho\mathbf{\Gamma})^{-1}$, and $\mathbf{e}\mathbf{1}$ denotes an m -by-1 vector whose first element is unity and the remaining elements are all zero. $\mathbf{e}\mathbf{1}'\rho\mathbf{\Gamma}\mathbf{\Phi}$ captures the long-run significance of each individual VAR shock to expected returns. The greater the absolute value of a variable's coefficient in the return predictive regression, i.e. the first row of $\mathbf{\Gamma}$, the greater the weight of this variable in the discount-rate news formula. Furthermore, the term $\mathbf{\Phi}$ guarantees that more persistent variables receive more weight.

The cash-flow news can be backed out as a residual:

$$\eta_{cf,t+1} = (\mathbf{e}\mathbf{1}' + \mathbf{e}\mathbf{1}'\rho\mathbf{\Gamma}\mathbf{\Phi}) \mathbf{u}_{t+1}. \quad (8)$$

The VAR approach has three main advantages. First, as noted by Campbell (1996) it links the vast time-series literature on predictability to the cross-sectional literature. Second, one does not necessarily need to understand short-term dynamics of dividends which are notoriously difficult to predict. Third, this approach yields results that are almost identical to those derived from forecasting cash flows directly based on the same information set and provided that the dividend yield is included in the set of state variables and the VAR generates reliable return forecasts (Engsted, Pedersen, and Tanggaard (2012) and Campbell, Giglio, and Polk (2013)).

⁶As discussed by Campbell and Shiller (1988), the assumption that the VAR is first-order is not restrictive, since this formulation allows for higher-order models by stacking lagged values into the state vector.

3.2 The Intertemporal CAPM

Assuming an infinitely lived investor with recursive preferences proposed by Epstein and Zin (1989, 1991), Campbell (1993) derives an approximate discrete-time version of Merton's (1973) ICAPM. The model's central pricing relation is based on a first-order Euler equation that relies on time preference parameter δ and distinguishes between the coefficient of relative risk aversion γ and the elasticity of intertemporal substitution (EIS) φ . When asset returns and consumption growth are jointly conditionally homoscedastic and lognormally distributed, consumption can be substituted out from the budget constraint, and risk premia can be approximated as a function of the coefficient of relative risk aversion and a discount coefficient ρ . This approximation is sufficiently accurate if $\varphi = 1$, $\rho = \delta$ and the optimal consumption-wealth ratio is conveniently constant. Assuming further that a well-diversified market portfolio is a good proxy for the optimal portfolio of a long-horizon investor, the risk premium on any asset i satisfies

$$E_t(r_{i,t+1}) - r_{t+1}^f + \frac{\sigma_{i,t}^2}{2} = \gamma Cov_t(r_{i,t+1}, r_{m,t+1} - E_t(r_{m,t+1})) + (1 - \gamma)Cov_t(r_{i,t+1}, -\eta_{dr,t+1}), \quad (9)$$

where r^f is the risk-free rate, σ_i^2 denotes the asset's variance, i.e. the Jensen's inequality effect, and η_{dr} is the same as above. In the case of $\gamma = 1$, equation (9) reduces to the static CAPM framework.

Using simple expected returns, $E_t(R_{i,t+1} - R_{t+1}^f)$, instead of log returns, $E_t(r_{i,t+1}) - r_{t+1}^f + \frac{\sigma_{i,t}^2}{2}$, and exploiting the approximation in equation (5), Campbell and Vuolteenaho (2004) show that asset returns can be related to two fundamental sources of risk captured by permanent and transitory news components. In unconditional terms, the expected risk premium on asset i obeys

$$E[R_i^e] = \gamma\sigma_m^2\beta_{i,cf} + \sigma_m^2\beta_{i,dr}, \quad (10)$$

where $E[R_i^e]$ denotes the average excess return on stock i in excess of the risk-free rate and σ_m^2 is the variance of the market portfolio. Here the cash-flow beta is defined as:

$$\beta_{i,cf} \equiv \frac{Cov(\eta_{cf,t+1}, R_{i,t+1}^e)}{Var(\eta_{m,t+1})} \quad (11)$$

and the discount-rate beta as:

$$\beta_{i,dr} \equiv \frac{Cov(-\eta_{dr,t+1}, R_{i,t+1}^e)}{Var(\eta_{m,t+1})}. \quad (12)$$

Both betas add up to the traditional market beta measured as:

$$\beta_{i,m} \equiv \frac{Cov(\eta_{m,t+1}, R_{i,t+1}^e)}{Var(\eta_{m,t+1})}, \quad (13)$$

where $\eta_{m,t+1} = r_{m,t+1} - E_t(r_{m,t+1}) = \eta_{cf,t+1} - \eta_{dr,t+1}$ denotes the unexpected market return at time $t + 1$. Note that sensitivities in (11) and (12) are defined as rescaled slope coefficients following Campbell and Mei (1993). Furthermore, the discount-rate beta is defined as the sensitivity of an asset's return to the good news about the market, i.e. lower-than-expected discount rates, such that we expect a positive premium as a compensation for high beta.

3.3 Two Discount Rate News Components

Motivated by the evidence presented in Section 2, we further decompose the total discount rate news into news about future long-term discount rates and news about future short-term discount rates:

$$\eta_{dr,t+1} = \eta_{ldr,t+1} + \eta_{sdr,t+1}. \quad (14)$$

We base the distinction between these two discount rate news components on a threshold horizon h . We measure the long-term news component, defined as a sum of future unexpected returns from the horizon h up to infinity, as:⁷

$$\eta_{ldr,t+1} = \mathbf{e}\mathbf{1}'\rho^h\mathbf{\Gamma}^h\mathbf{\Phi}\mathbf{u}_{t+1}. \quad (15)$$

and retrieve the short-term discount-rate news, defined as a sum of future unexpected returns from now up to the horizon $h - 1$, residually:

$$\eta_{sdr,t+1} = \mathbf{e}\mathbf{1}'(\rho\mathbf{\Gamma} - \rho^h\mathbf{\Gamma}^h)\mathbf{\Phi}\mathbf{u}_{t+1}. \quad (16)$$

Note that this formulation reduces to a two-beta ICAPM for h equal to unity. Our baseline empirical findings are reported for h equal to one year or four quarters.⁸

We define the asset's long-term discount-rate beta by its sensitivity to the long-term discount rate news:

$$\beta_{i,ldr} \equiv \frac{Cov(-\eta_{ldr,t+1}, R_{i,t+1}^e)}{Var(\eta_{m,t+1})}, \quad (17)$$

and the asset's short-term discount rate beta by its sensitivity to the short-term discount rate news:

$$\beta_{i,sdr} \equiv \frac{Cov(-\eta_{sdr,t+1}, R_{i,t+1}^e)}{Var(\eta_{m,t+1})}. \quad (18)$$

Clearly, these two beta components add up to the discount rate beta, $\beta_{i,dr} = \beta_{i,ldr} + \beta_{i,sdr}$ as in Campbell and Vuolteenaho (2004). Furthermore, the overall market beta can then be written as $\beta_{i,m} = \beta_{i,cf} + \beta_{i,ldr} + \beta_{i,sdr}$.

The distinction between these two discount rate beta components raises the possibility that discount rate risks can have different cross sectional implications at short term and

⁷Chen and Zhao (2009) implement an analogous approach to study fixed-term maturity bonds.

⁸In the Appendix, we show that none of our conclusions are affected by this choice. Our results remain upheld for alternative values of h .

at long term horizons. High expected returns associated with high discount rate betas can thus be attributed to high long term discount rate betas, high short term discount rate betas, or both. Our analysis in Section 5 addresses this issue empirically.

4 Data

In order to implement the beta decomposition, we need to construct empirical proxies for news about cash flows and discount rates. Since these news components are not observable, the traditional approach is to predict them from observable state variables. In the next section, we describe the state variables used in the estimation of the VAR.

4.1 State Variables Choice

To operationalize the VAR approach, we need to specify the variables in the state vector. We opt for a parsimonious system with five state variables: the log excess market return measured by the excess log return on the CRSP value-weight index; the log dividend-price ratio on the S&P 500 index⁹ (dp); the log consumption-wealth ratio (cay) of Lettau and Ludvigson (2001); the short-term interest rate, measured by the one-month Treasury bill rate in percent; and the default yield spread, measured as the yield difference between BAA and AAA Moody's corporate bonds in percentage points. The data on the excess market return and the short term rate are available on the website of Ken French. The series for the aggregate consumption-wealth ratio is provided by Martin Lettau, while the data for default spread are obtained from the Federal Reserve Bank of St. Louis. We estimate the VAR over the longest available sample period spanning from 1952Q1 to 2013Q3, restricted by the availability of cay series.

In general, the implementation of the return decomposition framework requires the

⁹Online data is available on <http://www.econ.yale.edu/~shiller/data.htm>.

excess equity market return as a necessary component of the state vector. The other variables are optional. Our choice of dp and cay is motivated by theoretical reasons and their return forecasting ability (see also, Section 2). We include these variables in the system because they are important determinants of future stock returns and capture the time-varying dynamics of returns over different time horizons. In particular, we include dp in the system following Campbell, Polk, and Vuolteenaho (2010) and Engsted, Pedersen, and Tanggaard (2012) who emphasize that for the equity return decomposition in equation (5) to be valid, the dividend yield is indispensable. In addition, we consider the consumption-wealth ratio since this variable is a better forecaster of future returns at short and intermediate horizons than is the dividend yield and several other traditionally used variables (e.g. Lettau and Ludvigson (2001)). Next, there are two reasons to include the short-term interest rate in the VAR: First, previous studies document a strong predictive ability of the short rate for future excess returns (see, for instance, Fama and Schwert (1977) and Shiller and Beltratti (1992)). Secondly, in line with Ang and Bekaert (2007) we find that the forecasting potential of the dividend yield and the overall significance of the return predictive regression are considerably enhanced upon the inclusion of the short rate in a set of regressors. Finally, the default spread is an important state variable in the ICAPM because it contains information about future corporate profits and it helps describe time-variation in the investment opportunity set faced by investors (e.g. Fama and French (1989) and Campbell, Giglio, and Polk (2013)).

4.2 VAR Parameter Estimates

Table 2 reports the benchmark characteristics of the first-order VAR model including the log excess market return, the log dividend-price ratio, the log consumption-wealth ratio, the short-term rate, and the default spread. The VAR is estimated using OLS

and employing $\rho = 0.95^{1/4}$ for quarterly data.¹⁰ Each row of Table 2 corresponds to a different dependent variable listed in the header of the row. The first five columns give coefficients on the explanatory variables listed in the column header; the last column gives the adjusted \bar{R}^2 statistics. In parentheses are two t -statistics for each coefficient estimate. OLS t -statistics are reported in the upper row; Newey-West (1987) t -statistics are reported in the bottom row.

The first row shows the stock market return forecasting equation when lags of returns, price-dividend ratio, consumption-wealth ratio, short-term rate, and default spread are applied as regressors. The \bar{R}^2 statistic for the return equation is 7.11% over the full sample. According to unadjusted OLS t -statistics, all forecasting variables except for the lagged market return contain significant information about expected stock returns. In line with previous findings, the dividend yield and consumption-wealth ratio positively predict the market returns with t -statistics of 2.69 and 3.06, respectively. Higher past short-term rates are associated with lower returns similar to Ang and Bekaert (2007). The coefficient on the default spread is positive and statistically significant. Fama and French (1989) document similar evidence and argue that default spread tracks business cycle conditions closely. Taking into account the serial correlation and heteroskedasticity has no strong impact on standard errors for the variables in the system apart from the default spread.

The remaining rows provide evidence of strong interaction between the state variables. The autoregressive coefficients of the dividend yield, consumption-wealth ratio, and short term rate are all very close to unity, but they can be all explained to some extent by the other variables in the system. For example, past returns, past consumption-wealth ratio realizations, and past short term rate are strong predictors of future dividend yields. The short term rate and the default spread forecast next period changes in the *cay* residual.

¹⁰The results do not alter qualitatively for other plausible linearization parameter values.

Lagged default spread is a significant determinant of current short term rate, while the past short term rate is important in forecasting the default spread. The forecasting power of the VAR system is relatively high with $\overline{R^2}$ s varying from 87.38% in the *cay* forecasting regression to 96.77% in the predictive regression for *dp*. High persistence in the data might be a challenge to correct statistical inference and coefficient interpretation. However, advocates of stock return predictability argue that expected returns contain a slow-moving time-varying component whose persistence implies that the predicting variables should be persistent as well.¹¹

Using the VAR to calculate cash flow news produces a series which is almost unrelated to discount rate news. For example, the long-term discount rate news has a correlation of -0.08 with the cash-flow news, while the short-term discount rate news covaries positively with unexpected cash flow changes with a correlation of 0.09. Short-term and long-term discount rate news are negatively related with correlation of -0.53.

4.3 Individual Stocks as Test Assets

In asset pricing tests, stocks are often grouped into portfolios to mitigate the errors-in-variables bias caused by the estimation of betas. However, Ahn, Conrad, and Dittmar (2009), and Lewellen, Nagel, and Shanken (2010) note that the particular method of portfolio grouping can have a dramatic impact on the results. To avoid these problems, several recent studies employ a universe of individual stocks as test assets in cross-sectional re-

¹¹The appendix to Campbell and Vuolteenaho (2004) shows that persistence in the data is a priori unlikely to affect the performance of the ICAPM. In our case, the Kendall (1954) bias reduces the variability of discount rates and bias correction would most likely strengthen the significance of discount rate news terms. Furthermore, there are negligible correlations between return innovations and innovations in *cay*, *i*, and *def* of the order of -0.15, which imply a restrained Stambaugh (1999) bias for these variables. The upward bias in the return forecasting regression on the dividend yield works against the Kendall (1954) bias with unclear total outcome.

gressions. For instance, Ang, Liu, and Schwarz (2010) advocate the use of individual stocks in cross sectional tests of asset pricing models on statistical grounds. They show analytically and empirically that creating portfolios diminishes dispersion in betas and generates larger standard errors in cross-sectional risk premium estimates (see also the discussions in Avramov and Chordia, 2006, and Boons, 2016).

Litzenberger and Ramaswamy (1979), and Ang, Liu, and Schwarz (2010) emphasize that creating portfolios destroys information contained in the cross-section of betas and results in efficiency losses in the estimation of risk premia. This turns out to be a relevant empirical case for our three-fold beta decomposition because the use of portfolios produces a high degree of multicollinearity in cross-sectional regressions similar to the evidence reported in Botshekan, Kraeussl, and Lucas (2012). The cross sectional tests presented below are therefore based on individual common stocks with share codes 10 or 11 traded on the NYSE, AMEX, and NASDAQ exchanges over the period 1963Q3-2013Q3 from the Center of Research for Security Prices (CRSP) database in Wharton Research Data Services (WRDS) for which Compustat data on book equity is available.

We measure market equity (ME) as the total market capitalization at the firm level, i.e. stock price times shares outstanding. Balance sheet data is obtained from the annual Compustat database. We define book equity (BE) as stockholders' common equity plus balance-sheet deferred taxes and investment tax credit (if available) minus the book value of preferred stock. Based on availability, we use the redemption value, liquidation value or par value (in that order) for the book value of preferred stock. Book-to-market equity (BE/ME) is then book common equity for the fiscal year ending in calendar year t-1 divided by market equity at the end of December of t-1.

5 Empirical Findings

5.1 Baseline Risk Premium Estimates

To obtain the empirical estimates of risk premia associated with long-term and short-term discount rate news, we run cross-sectional regressions of individual stock excess returns on betas defined in Section 3. To form a basis for comparison, we first consider a standard single-beta CAPM:

$$E[R_i^e] = \lambda_0 + \lambda_m \beta_{i,m}, \quad (19)$$

and a two-beta ICAPM variant introduced by Campbell and Vuolteenaho (2004):

$$E[R_i^e] = \lambda_0 + \lambda_{cf} \beta_{i,cf} + \lambda_{dr} \beta_{i,dr}. \quad (20)$$

We evaluate the ability of short-term and long-term discount rate risks to capture the cross-section of stock returns by estimating a simple three-beta empirical ICAPM representation which distinguishes between short-term and long-term discount rate risks:

$$E[R_i^e] = \lambda_0 + \lambda_{cf} \beta_{i,cf} + \lambda_{ldr} \beta_{i,ldr} + \lambda_{sdr} \beta_{i,sdr}. \quad (21)$$

In representations (19)-(21), λ_0 is the intercept and λ_j is the price of risk factor j .

Table 3 presents our baseline cross-sectional estimates of prices of risk in percent per annum. For each model, we compute time-varying risk loadings recursively in 40-quarter overlapping rolling time-series estimation windows following definitions (11)-(13), (17) and (18). We then run recursive cross-sectional regressions of average returns in each estimation window on the betas computed over the same rolling window. In this way, we can compute a time series of estimated risk premia corresponding to the time-varying betas. We use heteroskedasticity and autocorrelation consistent (HAC) t -statistics of Newey and West (1987) to test if the time-series mean of the price of risk is significantly

different from zero.

We follow Botshekan, Kraeussl, and Lucas (2012) and Boguth and Kuehn (2013) and work with overlapping time-series windows of the length of ten years. A ten-year horizon gives a reasonable total number of estimation windows, while maintaining a sufficiently large number of observations within each window to compute reliable estimates of betas. There are 162 overlapping estimation windows in total. The first window covers the period 1963Q3-1973Q2, while the last window corresponds to the 2003Q4-2013Q3 period. We exclude stocks with one or more missing data points in a specific estimation window from the cross-sectional regression for that window. Furthermore, to ensure that our results are not driven by extreme outliers, we follow Botshekan, Kraeussl, and Lucas (2012) and winsorize returns in each estimation window at the 1% and 99% levels.

Column I of Table 3 shows that the standard market beta carries a positive and statistically significant price of risk of about 4.5% annually. The estimate of the associated constant term is 6.24% per annum. For comparison, Lewellen, Nagel, and Shanken (2010) report estimates of the zero-beta rate between 6.8% and 14.32% in annual terms for a number of asset pricing models which are often applied in empirical work.

In column II of Table 3, we break the total market beta into cash-flow and discount-rate components. Both sources of risk are associated with statistically significant and positive prices of risk of 5.92% and 3.58% per annum respectively. Therefore, in line with the prediction of Merton's (1973) ICAPM and the empirical analysis in Campbell and Vuolteenaho (2004), our estimates indicate that sensitivity to the long-lived permanent shock through cash flow news is rewarded with a higher price of risk than sensitivity to the short-lived shock to discount rates. The row "Diff." reports results of a t -test for differences in estimated cash-flow and discount-rate risk premia. It shows that the discount-rate price of risk is 2.34 percentage points lower than the cash-flow price of risk. This difference is significant with a t -statistic of -2.64.

Next, Column III of Table 3 presents the results of the three-beta ICAPM which distinguishes between short-term and long-term discount rate risks. Our estimates suggest that all three sources of risk have explanatory power in the cross section of stock returns. In particular, the cash-flow news is associated with the largest price of risk of 5.7% per annum in our specification. The price of risk for long-term discount rate news of 5.07% per annum exceeds the price of risk for short-term discount rate news of 3.03% per annum. The difference in long-term and short-term discount rate prices of risk is statistically different from zero and around 2.04 percentage points.

The estimates in Column IV of Table 3 suggest that the three-factor model of Fama and French (1993) is not a particularly well suited tool to describe returns on individual stocks. The market risk premium is estimated with a large standard error and the estimate of risk premium for HML is negative. Ang, Liu, and Schwarz (2010) and Jegadeesh and Noh (2013) similarly report a negative price of risk on HML.

To investigate the sensitivity of our ICAPM specifications to size and book-to-market effects, we augment the models in Columns I-III with characteristics. The Size and Value controls are measured by the log market capitalization and log book-to-market ratio in the first quarter of each rolling window. For consistency, we winsorize the Size and Value controls at the 1% and 99% levels in each window. The estimates in Columns I-A, II-A and III-A of Table 3 suggest that the overall cross-sectional patterns are unaffected by individual stock characteristics. There is a mild downward shift in the total market, cash-flow, and discount rate risk premia, and an upward shift in the wedge between long-term and short-term discount rate risk premia.

The results show that there is a distinct role for both short term and long term discount rate news and that long term discount rate news carries a higher price of risk than short term discount rate news. Both of these news terms command a lower price of risk than cash flow news. In an ICAPM framework, this makes sense for a conservative

long term investor since cash flow news has a permanent effect on stock returns. Higher than expected short term discount rates are reversed quickly and while initially bad news for long term investors, they are compensated by higher future returns quickly. In contrast, bad long term discount rate news today lowers stock prices today and it takes a long time before the higher future returns are realized. This is especially detrimental for short term investors.

5.2 Portfolio Sorts

An alternative approach to examine the relation between estimated risk loadings and expected returns is to group the estimates cross-sectionally and form portfolios according to their risk exposures. If there is a cross-sectional relation between short- and long-term discount rate risks and returns, then we should observe patterns between average realized returns and these discount rate risks loadings. In particular, the ICAPM implies that stocks with higher realized loadings on discount rate risks have higher average returns over the same period. This approach has an important advantage relative to Fama-MacBeth regressions. While the errors-in-variables problem generates biased standard errors in the second-stage regressions, it leads here to conservative statistical inference. If the betas are estimated with noise, the portfolio formation procedure will be somewhat less accurate because some stocks are assigned to a wrong portfolio, which diminishes the cross-sectional dispersion in returns across portfolios. However, since the inference is based solely on portfolio returns, the measurement error will eventually lower the statistical significance (see also, Boguth and Kuehn (2013)).

5.2.1 Single Sorted Portfolios

We first sort all stocks in our sample into quintiles based on their estimated betas computed following definitions (11), (17) and (18) in an independent fashion.¹² For consistency with the Fama-MacBeth regressions, we work with overlapping time-series windows of the length of ten years.¹³

Over every ten-year period, we compute the cash-flow, short-term and long-term discount rate risk exposures of every stock on a quarterly basis. At the beginning of each rolling window and for each risk characteristic separately, we rank individual stocks into quintiles based on their realized β_{cf} , β_{ldr} and β_{sdr} over the next ten years. The first portfolio contains stocks with the lowest betas, whereas the fifth portfolio contains stocks with the highest betas.

Table 4 reports various summary statistics of these portfolios. Columns "EW" and "VW" show the average equal- and value-weighted realized returns over ten-year windows in each portfolio. These average returns are computed over the same ten-year period as the respective betas. Hence, Table 4 shows the relation between contemporaneous factor loadings and returns, since a contemporaneous relation between risk measures and risk premia is the foundation of a traditional cross-sectional risk-return trade-off (e.g. Fama and MacBeth (1973), Ang, Chen, and Xing (2006) and references therein).

Panel A of Table 4 documents a monotonically increasing pattern between realized average returns and realized cash-flow betas. This result is not affected by the weighting scheme of portfolio returns. Portfolio 1 (5) has an average equal-weighted excess return of 8.85% (14.12%) per annum. The spread in average returns between these portfolios is 5.27% per annum with a t -statistic of 6.95.¹⁴ It is interesting to note that this difference

¹²Our experiments with decile portfolios lead to similar evidence.

¹³We have experimented with using other intervals but this had no impact on our results.

¹⁴To adjust for the induced moving average effects, we report HAC t -statistics with an optimal lag length.

corresponds closely to the estimate of the cash flow price of 5.70% per annum in our cross sectional regressions reported in Table 3. We observe a similar monotonic pattern in value-weighted returns on cash flow beta sorted portfolios which yield a spread in returns on extreme portfolios of 6.70% per annum with a t -statistic of 6.98.

Cross-sectional differences in returns might be not very surprising if the betas with respect to a risk factor under consideration covary with other variables which are known to explain returns. Table 4, Panel A shows that there is little evidence to support this hypothesis in our sample. For each portfolio, we report the average market share in percent (Mkt Share), average log market capitalization (Size), and average book-to-market (B/M) characteristics. Portfolio 1 has on average higher market capitalization and book-to-market ratio relative to portfolio 5, but there appears no systematic relation between firm characteristics known to predict future returns and average returns on the cash flow sensitivities sorted portfolios.

Table 4, Panel B shows that stocks with high contemporaneous β_{ldr} have high average returns. An equally-weighted portfolio of stocks in the quintile with the lowest (highest) long-term discount rate betas earns 9.39% (13.81%) per annum in excess of the risk-free rate. Differences in returns between quintile portfolios 5 and 1 are statistically significant and economically of the order of 4.42% per annum which is very similar to the estimated price of long run discount rate risk of about 5.1% per annum. This pattern is more pronounced for value-weighted portfolios with strictly monotonically increasing returns. Stocks with high β_{ldr} must carry a premium in order to compensate investors for their low returns in times of high long-term discount rates.

The portfolios sorted by long term discount rate news show a clearer relationship to firm characteristics. For example, portfolio 1 with the lowest beta with respect to long term discount rate news has a B/M ratio of 0.99, whereas portfolio 5 with the highest long term discount rate beta has a B/M ratio of 0.84. This result reminds of Campbell

and Vuolteenaho (2004) who document that growth stocks have higher total discount rate betas than value stocks over the post-1963 period.

In Panel C of Table 4, we sort stocks by their realized β_{sdr} . Our estimates show that stocks with high β_{sdr} have high average returns. The difference in average returns between portfolios 5 and 1 is 3.05% per annum for equal-weighted returns and 2.95% for value-weighted returns. These differences in returns are highly significant with robust t -statistics of 5.08 and 3.89, respectively. Again, we notice that these return differentials are tightly related to the estimated price of risk for assets' exposure to short term discount rate risk documented in Table 4. In sum, the evidence in Panel C of Table 4 demonstrates that bearing the short-term discount rate risk component is rewarded with high average returns. There is a modest spread in the B/M ratios of the firms in the portfolios where the low short-term discount rate beta portfolio has a B/M of 0.87 and the high short-term discount rate beta having a B/M of 0.93. We also find that stocks with the highest short term discount rate betas (portfolio High) tend to have lower market equity than stocks with the lowest short term discount rate betas (portfolio Low). However, we find no monotonic relation between firm size and exposure to short term discount rate news here.

5.2.2 Double Sorted Portfolios

In Table 5, we report average portfolio returns of independent double sorts. First, in each quarter, we sort stocks into terciles based on cash-flow beta, short-term discount rate beta, and long-term discount rate beta over the next ten years independently and compute average returns for each portfolio over the same ten-year period. Thus, we construct three bins of stocks sorted by each risk characteristic such that the first bin contains stocks with the lowest respective beta and the third bin contains stocks with the highest respective beta. We then construct 3x3 portfolios which are the intersections of

each two categories of the beta sorts.¹⁵

Panel A of Table 5 presents returns on portfolios sorted by cash-flow and long-term discount rate betas, Panel B presents returns on portfolios sorted by cash-flow and short-term discount rate betas, and Panel C summarizes the performance of portfolios sorted by long-term discount rate and short-term discount rate betas. The general patterns we observe are similar for equal-weighted and value-weighted portfolio schemes (omitted for brevity).

The evidence presented in Panels A and B of Table 5 suggests that stocks with high cash flow risk are systematically rewarded with higher excess returns than stocks with high long term or short term discount rate risk. For instance, the spreads in returns on High versus Low cash flow beta sorted portfolios vary from 4.36% to 2.68% per annum for Low versus High long term discount rate beta stocks (Panel A) and from 4.33% to 4.34% per annum for Low versus High short term discount rate beta stocks (Panel B). The respective spreads on High versus Low long term discount rate beta sorted portfolios vary from 3.02% to 1.33% per annum (Panel A), and on High versus Low short term discount rate beta sorted portfolios from 1.37% to 1.38% per annum (Panel B). These results give further support to the Merton's (1973) ICAPM idea that agents should receive a greater compensation for exposure to permanent versus transitory shocks.

Moreover, our double-sorting exercise confirms that there is a higher compensation associated with long run discount rate risk as opposed to short run discount rate risk. This result emerges from the fact that within the same cash flow risk category, spreads in returns on long term discount rate betas tend to exceed spreads in returns on short term discount rate betas (Panels A and B). This intuition is reaffirmed in double sorts based on the two discount rate beta components: Panel C of Table 5 shows that after controlling for short term discount rate risk, the spreads in returns on High versus Low

¹⁵Our results remain generally upheld in 5x5 double sorts (not reported).

long term discount rate beta sorted portfolios vary from 4.37% to 3.99% per year. By contrast, after controlling for long term discount rate risk, the spread in returns on High versus Low short term discount rate beta sorted portfolios are lower and vary from 3.39% to 3.01% per year.

Two conclusions can be drawn from these results. First, stocks whose returns are more positively related to cash-flow, short-term and long-term discount rate risks earn higher average returns, consistent with our evidence that each fundamental risk source is rewarded with a positive risk premium. This makes sense economically, as stocks with high betas are a risky investment for risk-averse agents who would like to hedge against unexpected drops in the market return. In particular, our portfolio sorts tend to show a monotonically increasing pattern for each risk characteristic along each dimension. Differences in returns on portfolios with high versus low risk exposures are statistically significant at the 1% level in each case. Second, cash flow risk exposure is rewarded with the highest excess returns, while the compensation for long term discount rate risk is greater than the compensation for the short term discount rate risk.

The results presented in Section 5.2 are consistent with the evidence from Fama-MacBeth (1973) regressions documented in Section 5.1 in two dimensions. First, each of our three risk factors related to cash flow, long term and short term discount rate news, is rewarded with positive excess returns. Second, we find that fluctuations in the long term discount rates are perceived as a more severe source of risk than fluctuations in the short term discount rates and the former are hence rewarded with a higher premium than the latter.

5.3 Size and Book to Market

To investigate size and book-to-market effects, we test our three-beta specification employing different cross-sectional sub-samples sorted on firms' characteristics. In Panel

A of Table 6, we form five quintiles of stocks sorted by size. In the beginning of each ten-year rolling window, we sort all companies in the sample based on their market capitalization. We follow Fama and French (1993) and use NYSE breakpoints for market equity (ME) to allocate NYSE, Amex, and NASDAQ stocks to five groups. We then run cross-sectional Fama-MacBeth (1973) regressions for each quintile of stocks separately. This process is repeated recursively on a rolling window basis. In Panel B of Table 6, we proceed analogously but form five quintiles of stocks sorted by book-to-market equity (BE/ME).

Our estimates in Panel A of Table 6 show a clear effect of size on the estimated premia in the three-beta model. There is a strongly pronounced drop in the cash-flow price of risk from low market capitalization stocks to high market capitalization stocks. The cash-flow premium declines almost monotonically from small to large stocks both in economic and statistical terms with estimates of about 6.23% per annum for the lowest market capitalization stocks and less than 0.1% for the highest market capitalization stocks. This premium is not statistically significant in the fourth and fifth quintiles.

There are also large changes in the price of the discount rate news factors across firms of different size. Within each category of companies sorted on size, we find a pervasive long term discount rate risk premium which declines almost monotonically from small to large companies. Sensitivity to the long term discount rate news component is rewarded with a positive premium of 5.97% per annum among small stocks and 2.70% per annum among large stocks. There appears also a strong relation between firm size and short term discount rate risk. While the short term discount rate risk premium is about 3.98% per annum for the lowest market equity firm quintile, and thus exceeds its full sample estimate of about 3% per annum, it is statistically indistinguishable from zero for the firms in the second, third, and fifth market equity quintiles. This evidence is consistent with the findings of Perez-Quiros and Timmermann (2000) who show that small firms

are stronger affected by credit market conditions in recessions, i.e. at business cycle frequencies when discount rates are high. We find that differences in prices of long term and short term discount rate risks range from 1.88 to 4.25 percentage points across the five stock groups we examine and are statistically significant within each portfolio.

We also find that there is a substantial variation in the cross sectional estimates of the prices of risk across book to market samples. Panel B of Table 6 shows that the cash flow and long term discount rate news terms exhibit an increasing pattern from low book-to-market to high book-to-market firms. For example, the cash flow risk premium is 4.58% per annum for the stocks in the first bin and 8.12% per annum for the stocks in the fifth bin. The respective return premia for the long-term discount rate risk are 4.79% and 6.08% per annum. The relation between book-to-market and short-term discount rate risk is generally less clearly pronounced, but stocks in the lowest book-to-market bin are rewarded with a premium of 2.06% per annum whereas stocks in the highest book-to-market bin are rewarded with a higher premium of 3.14% per annum. Furthermore, the estimates in Panel B of Table 6 reinforce our main finding that there is a higher premium attached to long term versus short term discount rate news.

In Panel C of Table 6, we report estimates of the prices of risk for all, and all but the lowest market equity stocks, i.e. we exclude stocks in the lowest market equity quintile from the cross section. This exercise reveals two interesting properties of the data. First, in line with the evidence documented in Panel A of Table 6, we find that the risk premia are higher for each of the three risk factors in the cross section of all stocks compared to a smaller cross section of stocks which does not include small firms. Second, the estimated price of short term discount rate risk declines by more than a half and becomes insignificant once we take the small firms out of the sample. We find no such evidence when we exclude other stocks from the sample. We hence conclude that the short term discount rate risk premium is mostly due to small firms in the sample.

This exercise reveals a considerable variation in the cross sectional estimates of the prices of risk across size and book to market samples. The price of risk of both discount rate news components is high for small and high book to market firms, and low for large and low book to market firms. The longer term discount rate risk is priced most consistently across samples, while the short-term discount rate risk premium is largely due to low market equity firms. Finally, the long term discount rate risk premium is economically and statistically higher than the short term discount rate risk premium within any category of stocks we examine.

5.4 Factor Contributions

Our analysis has so far focused on the estimates of risk prices λ_j in the three-beta ICAPM where j denotes the cash-flow (*cf*), long-term discount rate (*ldr*), and short-term discount rate (*sdr*) risk components. The overall contributions of individual risk factors to the total expected excess return on stock i are, however, comprised of products of these premia and the associated risk exposures β_j^i 's of the return on stock i . For instance, high average returns can be associated with high prices of risk, high risk exposures or both. To gain insight into the relative importance of factor contributions to overall risk premium in stock markets we follow Botshekan, Kraeussl, and Lucas (2012) and perform the following analysis. For each ten-year overlapping rolling time-series estimation window t of the Fama-MacBeth (1973) procedure, we compute the product of the factor premium estimate and the cross-sectional average beta over that window $\lambda_{jt} \times \bar{\beta}_{jt}$, where the betas are computed following definitions (11), (17) and (18) over the window t . Table 7 shows the time-series averages of expected return contributions attributed to cash flow, long-term discount rate and short-term discount rate risk factors in percent per annum and their HAC t -statistics in parentheses for various sub-samples.

Panel A of Table 7 summarizes the results for the five quintiles of companies sorted

on their market capitalization in the beginning of each rolling window. Confirming our evidence from Table 6, we find that the expected return component of each risk factor is higher for small than for large stocks. For example, the contributions of the cash flow, long term discount rate and short term discount rate news factors are 2.17%, 1.46%, and 1.82% respectively for stocks in the lowest market equity portfolio, whereas the corresponding figures for stocks in the highest market equity portfolio are -0.12%, 0.48%, and 0.21%. Amongst the three risk factors, the long term discount rate news contributes most pervasively to the overall excess return within each quintile of size sorted stocks. The cash flow contributions are insignificant for stocks in the fourth and fifth size quintiles. Furthermore, the contributions of the short term discount rate risk factor are close to zero for stocks in the second, third, and fifth size quintiles. In addition, except for stocks in the lowest market equity category, we find that the relative contribution of the longer term discount rate risk factor is greater than that of the shorter term discount rate risk factor. Because long term discount rate news is generally associated with a higher price than short term discount rate news, this indicates that small firms have smaller sensitivity to long term discount rate news as opposed to short term discount rate news in line with Perez-Quiros and Timmermann (2000).

As regards book-to-market sorted portfolios, we find that each factor contributes significantly to the overall expected returns across all BE/ME bins. This evidence reinforces our findings reported above. Moreover, the contributions of the three factors are smaller for firms with lowest BE/ME ratios relative to firms with highest BE/ME ratios. For example, the expected return contributions of the cash flow, long term discount rate, and short term discount rate news components are 1.99%, 1.34%, and 0.68% respectively for stocks in the lowest BE/ME quintile, while the corresponding numbers for stocks in the highest BE/ME quintile are 2.41%, 1.43%, and 1.60%.

Contrasting with the evidence for size portfolios where the relative importance of cash

flow, long term discount rate and short term discount rate risk factors varies a lot across bins, for book-to-market sorted portfolios, we find that our three risk factors have a fairly constant impact on returns in relative terms. Specifically, the cash flow factor has the greatest impact on the overall expected return, followed by the long term discount rate factor. The short term discount rate factor has the lowest contribution to the expected returns across BE/ME portfolios apart from stocks in the highest BE/ME category. Since longer term news is rewarded with a higher price than shorter term news, this suggests that value firms are similar to small firms in that they have smaller sensitivity to the long term as opposed to the short term discount rate risk factor. This result echoes the findings of Lettau and Wachter (2007) who show that firms with cash flows weighted more to the present have a low ratio of price to fundamentals and high expected returns relative to assets with a high ratio of price to fundamentals.

Turning to Panel C of Table 7, our estimates show that in the full sample, all three risk factors contribute significantly to expected returns. Cash flow risks have the largest impact on expected returns. This result is in line with the key intuition of the Merton's (1973) ICAPM and our evidence from baseline Fama-MacBeth (1973) regressions presented in Table 3. Moreover, the contribution of the long term discount rate risk factor exceeds that of the short term discount rate risk factor. Differences in relative contributions of the two discount rate risk factors become stronger when we exclude small stocks from the sample. In this case, the contribution of the short term discount rate news goes down to about 0.35% per annum and becomes statistically not distinguishable from zero.

In sum, our evidence presented in Table 7 highlights three interesting observations. First, while each of our three factors has a positive and statistically significant return contribution in the full sample, there is substantial variation in the relative importance of risk components across size and book-to-market. Second, judged by the t -statistics and the economic magnitude of the estimates, the long term discount rate risk factor emerges

as a more pervasive contributor to the expected returns than the short term discount rate factor. Finally, our results indicate that small and value firms have smaller sensitivity to long term discount rate news as opposed to short term discount rate news.

6 Additional Robustness Checks

We have performed several additional tests to examine the robustness of our results. A detailed description and analysis are reported in the Appendix. This section provides a brief overview of the findings. The first subsection discusses the sensitivity of our results with respect to the choice of state variables in the system. The second subsection asks whether our conclusions are sensitive to the news computation method. The third subsection studies the evolution of risk prices in response to changes in the linearization parameter values and the values of the threshold time horizon parameter. Finally, the fourth subsection discusses the stability of our estimates over time and summarizes several further robustness checks.

6.1 Sensitivity to the Choice of State Variables

The analysis in this paper has hitherto relied on our basic VAR which includes the log excess market return, the log dividend-price ratio, the log consumption-wealth ratio, the short-term interest rate, and the default yield spread. However, Chen and Zhao (2009) warn that the loglinear return decomposition approach might be sensitive to the particular choice of state variables in the system. To address this concern, we consider nine alternative combinations of various return predictors to proxy for the discount rate and cash flow news. It turns out that our main conclusions with regard to the relation of the equity stock returns and the two market discount rate news components are almost unaffected by the variation in the state variables. Stock return sensitivities to the longer

term discount rate news are rewarded with a significant premium which is economically and statistically greater than that of the short term discount rate news.

Table 8 presents the details of this pricing exercise. Panel A of Table 8 gives an overview of the different combinations of the state variables that we consider. The plus signs indicate the selected variables in our baseline specification and in nine alternative models that we evaluate. Each model contains the excess market return as a necessary component (Campbell and Shiller (1988) and Campbell (1991)). In addition, Campbell, Polk, and Vuolteenaho (2010) and Engsted, Pedersen, and Tanggaard (2012) emphasize the importance of the asset price as a state variable in the VAR for the validity of the empirical return approximation. Hence, since we deal with equities, we include a proxy of the market dividend-price ratio or the market price-earnings ratio (interchangeably) as a second indispensable component in the system. The other variables are in general optional. Here we follow the related literature and consider alternative return predictors discussed for instance in Ang and Bekaert (2007) and Welch and Goyal (2008).

Our results in Panel B of Table 8 suggest that the economic estimates of risk premia indeed depend to some extent on the set of the underlying state variables. In particular, it is interesting to note that the choice of the VAR can shift the relative importance of the cash flow risks versus the discount rate risks as reported by Chen and Zhao (2009). Yet, our main conclusions remain almost unaffected by the variation in the vector of predictor variables. First, while sensitivities to the longer term discount rate news are rewarded with a positive and statistically significant risk premium in eight out of nine specifications that we examine, our estimates appear less conclusive about the short term discount rate risk premium: Different specifications suggest that the short term discount rate news can be positively, negatively or insignificantly priced. Secondly, the premium for the longer run discount rate news exceeds its short term counterpart in economic and statistical terms with the only exception of Specification VIII.

Overall, despite the differences in the magnitudes of the estimated prices of risk across the alternative VAR specifications, a problem common to this methodology, our main finding of a higher price of longer term discount rate risk is generally upheld. Please consult a separate Appendix on further discussion of these results and details on how we construct the data.

6.2 Directly Modeled Cash-Flow News

With a focus on innovations driven by discount rate news, this paper conventionally treats the cash flow news as a residual. In this vein, Engsted, Pedersen, and Tanggaard (2012) and Campbell, Polk, and Vuolteenaho (2010) argue that in a properly specified VAR, the risk premia are not necessarily affected by the decision to forecast returns or cash flows. However, Chen and Zhao (2009) and Chen, Da, and Zhao (2013) emphasize the importance of the news computation method for the relative significance of permanent and temporary shocks.

To address this criticism in our setup, we use a separate VAR system to model the cash flow news directly and independently of the discount rate news. This method acknowledges the fact that there is a noise component in stock returns which cannot be explained by permanent shocks to the dividend stream or transitory shocks associated with changes in discount rates.

In general, we find that the relative significance of the cash flows can change when we estimate them independently of the future returns. We also find that the noise betas are important in explaining cross sectional patterns in the data (see also, Garrett and Priestley (2012)). However, our results indicate that the news computation method has no material impact on the relative importance of the long-term and the short-term discount rate risks. Overall, our analysis further reinforces that the long run discount rate news is priced most consistently and pervasively across specifications and commands a higher

premium than short run discount rate news. We defer a more detailed discussion of this issue to the Appendix.

6.3 Sensitivity to Changes in ρ and h

In our benchmark specification, we follow Campbell and Vuolteenaho (2004) and assume $\rho = 0.95$ per year. We looked at the ρ parameter values between 0.93 and 0.97 but found that these changes have only a small impact on our estimates. Overall, the estimates of the longer term discount rate risk premium are more sensitive to changes in ρ because they are discounted more heavily. Higher values of ρ are typically associated with greater cash flow risk prices, lower long term discount rate prices, lower short term discount rate prices, and smaller differences between λ_{ldr} and λ_{sdr} . While higher values of ρ imply a flatter term structure of equity risk premium, our results are generally robust to reasonable variation in the parameter ρ .

A further important parameter in our decomposition is h , a threshold value to distinguish between the long term and the short term discount rates. The results reported in this paper are based on the value of h equal to four quarters. Our choice of the one-year horizon is motivated by practical reasons. We have experimented with lower and higher values of h but found that these changes had no significant impact on our conclusions. In general, the choice of h has no effect on the cash flow risk premium estimate and only a negligible impact on the short term discount rate risk premium estimate. Higher values of h imply larger estimates of the price of risk for the longer term discount rate news but lower its statistical significance. Accordingly, higher values of h are associated with economically larger but statistically less precisely measured differences between the two discount rate risk premia.

6.4 Further Sensitivity Tests

We have also investigated the sensitivity of our findings in the following dimensions: (i) the stability of prices of risk over different sub-samples, (ii) alternative sample periods to derive the news proxies, (iii) different rolling window lengths in the Fama and MacBeth (1973) estimation procedure, and (iv) alternative thresholds for winsorizing the data. We verify that none of these changes alter our main conclusions. The Appendix contains further details on these robustness checks.

7 Simulation

Our findings that short-term discount rate risk has a lower price of risk than the long-term discount rate risk does not contradict with a downward-sloping term structure of zero coupon equities' risk premium. In this Section we simulate a risk-based model to show this. Our empirical estimation captures shocks to the expected market holding period return rather than to discount rates on dividend strips. Before proceeding with the simulations, we would like to outline the difference between the two types of shocks.

The value of the market portfolio be the sum of the values on all dividend strips, namely the claims to aggregate dividends at each of the future dates, and is given by

$$P_t^m = \sum_{n=1}^{\infty} P_{nt}, \quad (22)$$

where P_{nt} denote the price of an asset that pays the aggregate dividend n periods from now and nothing in-between (a dividend strip). The market return is a weighted average of the returns on all dividend strip returns:

$$E_t(R_{m,t+1}) = \sum_{n=1}^{\infty} w_{tn} E_t(R_{n,t+1}), \quad (23)$$

where $w_{tn} = \frac{P_{nt}}{P_t^m}$. Note that shocks to the expected market returns can stem from shocks to the expected returns of any combination of the expected returns on the dividend strips (which constitute also shocks to the weights), as well as shocks to expected dividends which constitute shocks to the weights. Thus, a shock to the expected holding period on the market for a given horizon can be explained by an infinite number of combinations of shocks to discount rates and cash flows on the dividend strips.

To show this, we modify the dynamic risk-based model in Lettau and Wachter (2007), in particular their specification of the time-varying price of risk, and keep other ingredients intact. That is, our model has four shocks in total: a shock to dividend growth, a shock to expected dividend growth, and two components of shocks to the price of risk (One captures transient short term discount rate risk, and the other characterizes the long term discount rate risk, each of which takes a form of a mean-reverting process). The rest of this section summarizes the primary settings of this asset pricing model, and briefly discusses the results from simulation.¹⁶

We denote a 4×1 vector of *independent* standard normal shocks as ϵ_{t+1} . These shocks are independent of what econometricians can observe at time t . Let d_t denote the log aggregate dividend in the economy at time t . The aggregate dividend process evolves according to

$$\Delta d_{t+1} = g + z_t + \sigma_d \epsilon_{t+1} \quad (24)$$

in which z_t follows the AR(1) process

$$z_{t+1} = \phi_z z_t + \sigma_z \epsilon_{t+1} \quad (25)$$

with $0 \leq \phi_z < 1$. These specifications are the same as in Lettau and Wachter (2007).

The price of risk is exposed not only to short-term shocks, but also to long-term

¹⁶We refer interested readers to the online Appendix for further details on model set-up, derivation, calibration and simulation.

shocks. We assume that the *long-term component* of the price of risk is driven by a state variable \bar{x} which follows a highly persistent process

$$\bar{x}_{t+1} = \mu_{\bar{x}} + \phi_{\bar{x}}\bar{x}_t + \sigma_{\bar{x}}\epsilon_{t+1} \quad (26)$$

with $-1 \leq \phi_{\bar{x}} < 1$. The constant $\mu_{\bar{x}}$ determines the long-term mean which \bar{x} reverts to. The price of risk at time $t + 1$ is assumed to be governed by

$$x_{t+1} = (1 - \phi_x)\bar{x}_t + \phi_x x_t + \sigma_x \epsilon_{t+1} \quad (27)$$

with $-1 \leq \phi_x < 1$. It is a convex combination of \bar{x}_t and x_t plus shocks to the short-term discount rate coming in the current period. This process captures a rapidly mean-reverting shorter-run component of price of risks. The loading of x_{t+1} on \bar{x}_t is given by $(1 - \phi_x)$. This system has one more degree of freedom than the specification in Lettau and Wachter (2007).¹⁷

We assume that the stochastic discount factor (denoted as M_{t+1}) is an exponential-affine process. By doing so, we can accommodate the correlations among the risk factors driving asset returns, and get computationally tractable pricing relations that can be used in estimation (Singleton, 2006). It can be shown that M_{t+1} takes the form of

$$M_{t+1} = \exp \left\{ -r^f - \frac{1}{2}\phi_x^2 x_t^2 - \frac{1}{2}(1 - \phi_x)^2 \bar{x}_t^2 - \phi_x(1 - \phi_x)x_t \cdot \bar{x}_t - \phi_x x_t \epsilon_{d,t+1} - (1 - \phi_x)\bar{x}_t \epsilon_{d,t+1} \right\}, \quad (28)$$

where $r^f = \ln R^f$ denotes the log constant risk-free rate, and $\epsilon_{d,t+1} = \frac{\sigma_d}{\|\sigma_d\|}\epsilon_{t+1}$.

We focus on the prices of “zero-coupon” equities, the sum of which is identical to the price of the aggregate dividend in our economy. Recall that P_{nt} denote the price of an asset that pays the aggregate dividend n periods from now and nothing in-between. The

¹⁷Such a structure is also utilized in Drechsler and Yaron (2012), Bansal and Shaliastovich (2009), and Duffie, Pan, and Singleton (2000) which entertain a LRR model with 2-volatility processes.

very same asset's price at time $t+1$ is denoted by $P_{n-1,t+1}$. We conjecture that a solution to the pricing equation $P_{nt} = E_t[M_{t+1}P_{n-1,t+1}]$ should satisfy

$$\frac{P_{nt}}{D_t} = F(x_t, \bar{x}_t, z_t, n) = \exp\{A(n) + B_x(n)x_t + B_{\bar{x}}(n)\bar{x}_t + B_z(n)z_t\} \quad (29)$$

with a boundary condition, $P_{0t} = D_t$. The coefficients in equation (29), and henceforth the price series for the zero-coupon equities can be solved for in closed form.

The model is calibrated by using the quarterly data set described in Section 4. We consider two model setups: a *baseline* model in which shocks to cash flows are uncorrelated with shocks to either component of the price of risks, and a *full-fledged* model which takes into account pair-wise correlation between $\Delta d_{t+1}, z_{t+1}$ and x_{t+1}, \bar{x}_{t+1} . Panel A of Table 9 summarizes the parameter choices. Note that both components of the discount rate shocks are negatively correlated with the expected dividend growth rate shocks. Expected dividend growth and dividend growth shocks are negatively correlated ($\rho_{\Delta d, z} = -0.7800$), indicating that the hedging effect from shocks to expected dividend growth as reported in Lettau and Wachter (2007) also exists in our simulated models. The prices of risk for both components of discount rate shocks are positively correlated with the dividend growth rates, but the correlation is lower for the short-term component ($\rho_{\Delta d, x} = 0.42$). The opposite pattern applies to the correlation between the two components of the price of risks and the expected dividend growth, with $\rho_{z, x} = -0.49$ which is somewhat lower than $\rho_{z, \bar{x}}$.

We simulate 50,000 quarters from the system. Given simulated data on the shocks ϵ_{t+1} and state variables, we compute ratios of prices to aggregate dividends for zero-coupon equities according to (29). The quarterly return for a zero-coupon equity to mature in n quarters is computed as $R_{n,t+1} = P_{n-1,t+1}/P_{nt}$. Panel B of Table 9 reports the estimated price of risk for short- and long-term discount rate shocks separately. The price of risk for the short-term discount rate shocks based on the baseline model is 0.299, which

is lower than its long-term counterpart 0.453. By comparison, the full-fledged model generates a pair of slightly higher prices, but their ratios remain the same. The p -values at the bottom of Panel B computed by running Welch two-sample t-test indicate that the differences between the estimated price of risk for short- and long-term discount rate shocks are highly statistically significant. Compared with the average Sharpe ratio for the simulative market obtained by Lettau and Wachter (2007), which is 0.41 per annum, our estimates for the short-term components of the price of risks (0.299 p.a. based on the benchmark model and 0.304 p.a. based on the full-fledged model) are closer to the data equivalent of 0.33 as reported in Lettau and Wachter (2007). The gain in precision is largely attributed to disentanglement from the long-term component of the price of risk when setting up the dynamics of the price of risk for short-term discount rate shocks. Absence of this often leads to an upward-biased estimate for the price of risk.

This simulative result of different prices for the two components of discount rate risk coexists with the downward-sloping term structure of zero-coupon equities. The top panel of Figure 2 plots the risk premium for annualized holding period returns of individual zero-coupon equities. Both the baseline and full-fledged models produce a larger difference between the average returns of low and high duration assets relative to the Lettau and Wachter (2007) model. That is, decomposing the discount rate risk into the short- and long-term components renders the term structure of zero-coupon equities even more downward-sloping. The middle panel of Figure 2 illustrates that the return volatility initially increases with maturity and reaches its peak around 3 years. Then it monotonically decreases and quickly converges at maturities greater than 15 years. The hump-shaped pattern for return volatility reported by Lettau and Wachter (2007) is not obvious in our model. The bottom panel of Figure 2 plots the unconditional Sharpe ratios. Similar to Lettau and Wachter (2007) model, firms with more cash flow to be realized in the immediate future have higher Sharpe ratios than firms having more cash

flows in the far future.

The downward-sloping pattern for the term structure of risk premium for zero-coupon equities does not readily apply to the risk premia for holding the market portfolio over periods of various lengths. Figure 3 plots the annualized holding period returns for the CRSPVW market portfolio for various holding periods. In the data, the term structure of the market's holding period returns is somewhat downward sloping, albeit very non-monotonically. For instance, the annualized risk premium for holding the market portfolio for 20 years is on average higher than for merely 3 years. The other plots in Figure 3 are based on three simulative models, all of which generate monotonically decreasing term structure of the risk premium. Our simulative models (baseline and full-fledged) yield a pattern which matches the real data better than Lettau and Wachter (2007) model in a sense that the annualized risk premium drops rapidly in the first 3 years. Nevertheless, the mean risk premium in Lettau and Wachter (2007) model is closer to the real mean than our models in terms of magnitude.

8 Conclusion

The literature has thus far ignored the asset pricing implications of the empirical fact that stock returns are predictable at short and long horizons with different variables, suggesting a term structure in discount rates. We exploit this in order to show that there are two discount rate factors and one cash flow factor driving the cross section of stock returns. This enables us to connect the empirical findings that discount rate variation is driven by a short run business cycle component and a longer run trend component to risk factors in the cross section of stock returns. The decomposition of discount rates into short term and long term parts leads to state variable hedging of these two components and ICAPM logic implies a three factor model for expected returns.

In estimating a three factor ICAPM model, with both types of discount rate shocks and cash flow shocks, we find that while the price of risk of cash flow shocks is the highest, both types of discount rate shocks earn an economically meaningful and statistically significant risk premium. In particular, the price of risk of long term discount rate shocks is considerably higher than the price of risk of short term discount rate shocks. This finding is consistent with heterogeneity in investors' investment horizons. We also find that long term discount rate risk is priced consistently across samples while the short term discount rate risk premium is largely due to low market equity firms.

Our findings provide evidence on the types of risks investors require a premium for holding and make an important step towards understanding the factor structure of time varying expected returns.

References

- Ahn, D., Conrad, J., Dittmar, R., 2009. Basis assets. *Review of Financial Studies* 22, 5133-5174.
- Ang, A., Bekaert, G., 2007. Stock Return Predictability: Is It There?. *Review of Financial Studies* 20, 651-707.
- Ang, A., Chen, J., Xing, Y., 2006. Downside Risk. *Review of Financial Studies* 19, 1191-1239.
- Ang, A., Liu, J., Schwarz, K., 2010. Using Stocks or Portfolios in Tests of Factor Models?. Working Paper. Columbia University and University of Pennsylvania.
- Avramov, D., T. Chordia, 2006. Asset Pricing Models and Financial Market Anomalies, *Review of Financial Studies* 19, 1001-1040.
- Boons, M., 2016. State Variables, Macroeconomic Activity, and the Cross Section of Individual Stocks, *Journal of Financial Economics*, 119, 489-511.
- van Binsbergen, J., Brandt, M., Koijen, R., 2012. On the Timing and Pricing of Dividends. *American Economic Review* 102, 1596-1618.
- van Binsbergen, J. and R. Koijen, 2017, The term structure of returns: Facts and theory, *Journal of Financial Economics* 124, 1–21.
- Boguth, O., Kuehn, L.-A., 2013. Consumption Volatility Risk. *Journal of Finance* 68, 2589-2615.
- Botshekan, M., Kraeussl, R., Lucas, A., 2012. Cash Flow and Discount Rate Risk in Up and Down Markets: What Is Actually Priced?, *Journal of Financial and Quantitative Analysis* 47, 1279-1301.

- Brennan, M. J., Wang, A. W., and Xia, Y., 2004. Estimation and test of a simple model of intertemporal asset pricing. *Journal of Finance* 59, 1743-1775.
- Campbell, J.Y., 1987. Stock Returns and the Term Structure. *Journal of Financial Economics* 18, 373-399.
- Campbell, J.Y., 1991. A Variance Decomposition for Stock Returns. *Economic Journal* 101, 157-179.
- Campbell, J.Y., 1993. Intertemporal Asset Pricing Without Consumption Data. *American Economic Review* 83, 487-512.
- Campbell, J.Y., 1996. Understanding Risk and Return. *Journal of Political Economy* 104, 298-345.
- Campbell, J.Y., Ammer, J., 1993. What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns. *Journal of Finance* 48, 3-37.
- Campbell, J.Y., Chan, Y.L., Viceira, L.M., 2003. A Multivariate Model of Strategic Asset Allocation. *Journal of Financial Economics* 67, 41-80.
- Campbell, J.Y., Cochrane, J.H., 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107, 205-251.
- Campbell, J.Y., Giglio, S., Polk, C., 2013. Hard Times. *Review of Asset Pricing Studies* 3, 95-132.
- Campbell, J.Y., Giglio, S., Polk, C., Turley, R., 2017. An Intertemporal CAPM with Stochastic Volatility. *Journal of Financial Economics*, forthcoming.

- Campbell, J.Y., Lo, A.W., MacKinlay, A. C., 1997. *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press.
- Campbell, J.Y., Mankiw, G., 1989. Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence. In Olivier J. Blanchard and Stanley Fisher, eds.: *NBER Macroeconomics Annual*, Cambridge, MA, MIT Press.
- Campbell, J.Y., Mei, J., 1993. Where do Betas Come From? Asset Price Dynamics and the Sources of Systematic Risk. *Review of Financial Studies* 6, 567-592.
- Campbell, J.Y., Polk, C., Vuolteenaho, T., 2010. Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns. *Review of Financial Studies* 23, 305-344.
- Campbell, J.Y., Shiller, R.J., 1988. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1, 195-228.
- Campbell, J.Y., Vuolteenaho, T., 2004. Bad Beta, Good Beta. *American Economic Review* 94, 1249-1275.
- Chen, L., Zhao, X., 2009. Return Decomposition. *Review of Financial Studies* 22, 5213-5249.
- Chen, L., Da, Z., Zhao, X., 2013. What Drives Stock Price Movements? *Review of Financial Studies* 26, 841-876.
- Cochrane, J.H., 1992. Explaining the Variance of Price-Dividend Ratios. *Review of Financial Studies* 5, 243-280.
- Cochrane, J.H., 2008. The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies* 21, 1533-1575.
- Cochrane, J.H., 2011. Presidential Address: Discount Rates. *Journal of Finance* 66, 1047-1108.

- Drechsler, I. and Yaron, A., 2011. What's Vol Got To Do With It? Review of Financial Studies 24(1), 1-45.
- Duffie, D., J. Pan and K. Singleton. 2000. Transform Analysis and Asset Pricing for Affine Jump Diffusions. Econometrica 68: 1343-1376.
- Engsted, T., Pedersen, T.Q., Tanggaard, C., 2012. Pitfalls in VAR Based Return Decompositions: A Clarification. Journal of Banking and Finance 36, 1255-1265.
- Epstein, L.G., Zin, S.E., 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57, 937-969.
- Epstein, L.G., Zin, S.E., 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. Journal of Political Economy 99, 263-286.
- Fama, E.F., French, K.R., 1988. Permanent and Transitory Components of Stock Prices. Journal of Political Economy 96, 246-273.
- Fama, E.F., French, K.R., 1989. Business Conditions and Expected Returns on Stocks and Bonds. Journal of Financial Economics 25, 23-50.
- Fama, E.F., French, K.R., 1993. Common Risk Factors in the Returns on Stocks and Bonds. Journal of Financial Economics 33, 3-56.
- Fama, E.F., MacBeth, J.D., 1973. Risk, Return, and Equilibrium: Empirical Tests. Journal of Political Economy 81, 607-636.
- Fama, E.F., Schwert, G.W., 1977. Asset Returns and Inflation. Journal of Financial Economics 96, 115-146.

- Favero, C.A., Gozluklu, A.E., Andrea, T., 2011. Demographic Trends, the Dividend-Price Ratio, and the Predictability of Long-Run Stock Market Returns. *Journal of Financial and Quantitative Analysis* 46, 1493-1520.
- Garrett, I., Priestley, R., 2012. Dividend Growth, Cash Flow, and Discount Rate News. *Journal of Financial and Quantitative Analysis* 47, 1003-1028.
- Giglio, S., M. Maggiori, and J. Stroebe, 2015, Very long-run discount rates, *Quarterly Journal of Economics* 130, 1–53.
- Hansen, L., Hodrick, R., 1980. Forward Exchange Rates as Optimal Predictors of Future Spot Rates. *Journal of Political Economy* 88, 829-853.
- Hodrick, R., 1992. Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement. *Review of Financial Studies* 5, 357-386.
- Jegadeesh, N., Noh, J., 2013. Empirical Tests of Asset Pricing Models with Individual Stocks. Working Paper. Case Western University and Emory University.
- Keim, D.B., Stambaugh, R.F., 1986. Predicting Returns in the Stock and Bond Markets. *Journal of Financial Economics* 17, 357-390.
- Kendall, M., 1954. Note on Bias in the Estimation of Autocorrelation. *Biometrika* 41, 403-404.
- Lamont, O., 1998. Earnings and Expected Returns. *Journal of Finance* 53, 1563-1587.
- Lettau, M., Ludvigson, S.C., 2001. Consumption, Aggregate Wealth and Expected Stock Returns. *Journal of Finance* 66, 815-849.
- Lettau, M., Ludvigson, S.C., 2005. Expected Returns and Expected Dividend Growth. *Journal of Financial Economics* 76, 583-626.

- Lettau, M., Wachter, J.A., 2007. Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium. *Journal of Finance* 62, 55-92.
- Lewellen, J., Nagel, S., Shanken, J., 2010. A Skeptical Appraisal of Asset Pricing Tests. *Journal of Financial Economics* 96, 175-194.
- Lintner, J., 1965. The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *Review of Economics and Statistics* 47, 13-37.
- Litzenberger, R.H., Ramaswamy, K., 1979. The effect of personal taxes and dividends on capital asset prices. *Journal of Financial Economics* 7, 163-196.
- Lochester, L., 2009. Expected Returns and the Business Cycle: Heterogeneous Goods and Time-Varying Risk Aversion. *Review of Financial Studies* 22, 5251-5294.
- Lustig, H., van Nieuwerburgh, S., 2006. The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street. *Review of Financial Studies* 21, 2097-2137.
- Merton, R.C., 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica* 41, 867-887.
- Newey, W.K., West, K.D., 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55, 703-708.
- Perez-Quiros, G., Timmermann, A., 2000. Cyclical Variations in Stock Returns. *Journal of Finance* 55, 1229-1262.
- Petkova, R., 2006. Do the Fama-French Factors Proxy for Innovations in Predictive Variables? *Journal of Finance* 61, 581-612.
- Piazzesi, M., Schneider, M., Tuzel, S., 2007. Housing, Consumption, and Asset Pricing. *Journal of Financial Economics* 83, 531-569.

- Rietz, T.A., 1988. The Equity Risk Premium: A Solution. *Journal of Monetary Economics* 22, 117-131.
- Santos, T., Veronesi, P., 2006. Labor Income and Predictable Stock Returns. *Review of Financial Studies* 19, 1-44.
- Sharpe, W.F., 1964. Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance* 19, 425-442.
- Shiller, R.J., Beltratti, A.E., 1992. Stock Prices and Bond Yields: Can Their Co-movements be Explained in Terms of Present Value Models?. *Journal of Monetary Economics* 30, 25-46.
- Singleton, K. J., 2006. *Empirical Dynamic Asset Pricing: Model Specification and Econometric Assessment*. Princeton University Press. ISBN-13: 978-0-691-12297-7.
- Stambaugh, R.F., 1999. Predictive Regressions. *Journal of Financial Economics* 54, 375-421.
- Wachter, J., 2006. A Consumption-based Model of the Term Structure of Interest Rates. *Journal of Financial Economics* 79, 365-399.
- Weber, M., 2016, Cash Flow Duration and the Term Structure of Equity Returns, forthcoming in *Journal of Financial Economics*.
- Welch, I., Goyal, A., 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21, 1455-1508.

Table 1: Summary of Forecasting Regressions

The table reports estimation results of long-horizon forecasting regressions of the form $r_{m,t+1}^{e,(h)} = a_0^{(h)} + a_1^{(h)} \times cay_t + \varepsilon_{t+1}^{(h)}$ in Column I and $r_{m,t+1}^{e,(h)} = b_0^{(h)} + b_1^{(h)} \times dp_t + \varepsilon_{t+1}^{(h)}$ in Column II. $r_{m,t+1}^{e,(h)}$ denotes the log excess return on the value-weighted CRSP return at time $t + 1$ over a horizon of h quarters, cay_t is the log consumption-wealth ratio of Lettau and Ludvigson (2001), and dp_t is the log dividend-price ratio on the S&P 500 index. The t -statistics in parentheses use the Hansen-Hodrick (1980) correction. $\overline{R^2}$ is the adjusted R^2 in %. The sample period is 1952Q1-2013Q3.

| h | I | | II | |
|-----|------------------|------------------|----------------|----------------------|
| | cay_t | $\overline{R^2}$ | dp_t | $\overline{R^2}(\%)$ |
| 1 | 0.62 (2.01) | 1.30 | 0.02 (1.59) | 0.86 |
| 4 | 2.37 (1.92) | 5.32 | 0.10 (1.74) | 4.84 |
| 12 | 5.08 (2.67) | 12.78 | 0.23 (2.51) | 15.84 |
| 16 | 5.96 (4.01) | 19.12 | 0.25 (3.32) | 19.62 |
| 20 | 6.79 (4.22) | 22.69 | 0.27 (4.70) | 21.06 |
| 40 | 6.08 (3.28) | 12.89 | 0.39 (4.75) | 28.71 |
| 60 | 2.09 (1.19) | 1.85 | 0.48 (4.32) | 34.83 |
| 80 | 0.32 (0.22) | -0.54 | 0.60 (5.82) | 49.02 |
| 100 | -2.08 (-1.26) | 3.24 | 0.39 (4.11) | 33.94 |
| 120 | -2.49 (-2.05) | 10.60 | 0.24 (3.58) | 28.51 |

Table 2: VAR Coefficient Estimates

The table shows OLS parameter estimates for a first order VAR model including the log excess market return (r_m), the log dividend-price ratio (dp), the log consumption-wealth ratio (cay), the short-term interest rate (i), and the default yield spread (def). All variables are mean-adjusted. Each row corresponds to a different dependent variable. The first five columns report coefficients on the explanatory variables listed in the column header; the last column shows the adjusted $\overline{R^2}$ statistics in %. In parentheses are two t -statistics for each coefficient estimate. The top statistic uses OLS standard errors; the bottom statistic uses the Newey-West (1987) correction. The sample period is 1952Q1-2013Q3.

| | $r_{m,t}$ | dp_t | cay_t | i_t | def_t | $\overline{R^2}(\%)$ |
|-------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|----------------------|
| $r_{m,t+1}$ | 0.05 (0.79) (0.67) | 0.04 (2.69) (2.57) | 0.95 (3.06) (3.19) | -0.03 (-3.56) (-3.07) | 0.03 (2.01) (1.29) | 7.11 |
| dp_{t+1} | -0.12 (-2.17) (-1.62) | 0.97 (77.16) (77.92) | -0.90 (-3.20) (-3.12) | 0.02 (3.09) (3.09) | -0.03 (-2.21) (-1.37) | 96.77 |
| cay_{t+1} | -0.00 (-0.02) (-0.02) | 0.00 (0.39) (0.46) | 0.92 (37.85) (41.42) | 0.00 (2.13) (2.43) | -0.00 (-2.19) (-1.97) | 87.38 |
| i_{t+1} | 0.20 (1.43) (1.41) | 0.02 (0.64) (0.65) | -1.27 (-1.79) (-1.85) | 0.99 (51.26) (46.87) | -0.11 (-3.28) (-3.18) | 93.51 |
| def_{t+1} | -0.73 (-5.21) (-5.13) | -0.01 (-0.44) (-0.56) | -0.84 (-1.19) (-1.39) | 0.06 (3.22) (2.64) | 0.82 (25.73) (17.16) | 78.70 |

Table 3: Baseline Cross-Sectional Regressions

The table reports Fama-MacBeth (1973) estimates of risk prices in percent per annum and their HAC t -statistics in parentheses for the single-beta CAPM (Column I), the two-beta ICAPM (Column II), the three-beta ICAPM (Column III), and the three-factor Fama-French model (Column IV). Specifications in Columns I-A, II-A and III-A represent augmented models with characteristics. The news series are computed from a VAR system in Table 2 following definitions (2)-(4), (10) and (11). We set the linearization parameter ρ to 0.95 per annum and the horizon h to one year. In quarterly recursive cross-sectional regressions, we regress average returns over 40-quarter overlapping rolling time-series estimation windows on time-varying risk loadings computed following definitions (6)-(8), (14) and (15) over the same rolling window. Stocks with one or more missing data points in a specific estimation window are deleted from the cross-sectional regression for that window. Returns in each window have been winsorized at the 1% and 99% levels. The log market capitalization (Size) and log book-to-market ratio (Value) are measured in the first quarter of each rolling window. The Size and Value controls have been winsorized at the 1% and 99% levels in each window. There are 162 overlapping estimation windows in total. The first window covers the period 1963Q3-1973Q2; the last window covers the period 2003Q4-2013Q3. $\overline{R^2}$ is average adjusted cross-sectional R^2 in percent. "Diff." reports results of a t -test for differences in estimated discount-rate and cash-flow risk premia in Columns II and II-A, and long-term and short-term discount-rate risk premia in Columns III and III-A.

| | I | II | III | IV | I-A | II-A | III-A |
|-----------------|-----------------|------------------|-----------------|------------------|------------------|------------------|------------------|
| λ_0 | 6.24 (11.23) | 6.24 (11.08) | 6.22 (10.53) | 10.57 (19.24) | 21.10 (12.90) | 21.16 (12.68) | 22.10 (13.36) |
| λ_m | 4.50 (11.66) | | | 0.75 (1.20) | 4.10 (11.62) | | |
| λ_{cf} | | 5.92 (8.79) | 5.70 (7.80) | | | 5.61 (9.07) | 4.68 (6.82) |
| λ_{dr} | | 3.58 (6.71) | | | | 3.54 (6.66) | |
| λ_{tdr} | | | 5.07 (11.64) | | | | 5.50 (12.88) |
| λ_{sdr} | | | 3.03 (4.15) | | | | 2.77 (3.41) |
| λ_{HML} | | | | -0.96 (-1.59) | | | |
| λ_{SMB} | | | | 0.21 (0.42) | | | |
| Size | | | | | -1.17 (-9.57) | -1.17 (-9.34) | -1.25 (-9.85) |
| Value | | | | | 1.12 (6.68) | 1.15 (6.63) | 1.16 (6.78) |
| Diff. | - | -2.34 (-2.64) | 2.04 (3.40) | - | - | -2.07 (-2.29) | 2.73 (3.77) |

Table 4: Characteristics of Single-Sorted Portfolios Formed on Risk Exposures

The table lists the average equal-weighted (EW) and value-weighted (VW) returns in percent per annum of portfolios sorted on estimated betas. The time-varying betas are computed recursively, on a quarterly basis, in 40-quarter overlapping rolling time-series estimation windows following definitions (12), (18) and (19). Stocks with one or more missing data points in a specific estimation window are deleted from the cross-sectional regression for that window. Returns in each window have been winsorized at the 1% and 99% levels. There are 162 overlapping estimation windows in total. The first window covers the period 1963Q3-1973Q2; the last window covers the period 2003Q4-2013Q3. For each risk characteristic, we rank individual stocks into quintiles and form portfolios of average returns in each estimation window. Portfolio Low contains stocks with the lowest betas, whereas portfolio High contains stocks with the highest betas. For each portfolio, we report the average market share in percent (Mkt Share), average log market capitalization (Size), average book-to-market (B/M) characteristics. The row labeled "High-Low" reports the difference between the returns of portfolio High and portfolio Low. The row labeled "t-stat." is the HAC t -statistic for the difference in returns on portfolio High and portfolio Low.

| Rank | EW | VW | Mkt Share | Size | B/M |
|---|---------|---------|-----------|------|------|
| Panel A: Portfolios Sorted by β_{cf} | | | | | |
| Low | 8.85 | 8.22 | 22.45 | 6.05 | 0.93 |
| 2 | 9.29 | 8.55 | 25.01 | 6.12 | 0.88 |
| 3 | 10.31 | 9.57 | 22.61 | 6.08 | 0.85 |
| 4 | 11.56 | 10.98 | 19.06 | 5.99 | 0.85 |
| High | 14.12 | 14.92 | 10.88 | 5.73 | 0.85 |
| High-Low | 5.27 | 6.70 | | | |
| t-stat. | (6.95) | (6.98) | | | |
| Panel B: Portfolios Sorted by β_{ldr} | | | | | |
| Low | 9.39 | 7.86 | 14.69 | 5.88 | 0.99 |
| 2 | 9.34 | 8.67 | 24.03 | 6.11 | 0.87 |
| 3 | 10.27 | 9.72 | 24.93 | 6.13 | 0.83 |
| 4 | 11.32 | 10.59 | 22.84 | 6.09 | 0.83 |
| High | 13.81 | 13.03 | 13.51 | 5.85 | 0.84 |
| High-Low | 4.42 | 5.17 | | | |
| t-stat. | (11.79) | (10.70) | | | |
| Panel C: Portfolios Sorted by β_{sdr} | | | | | |
| Low | 10.25 | 9.31 | 17.80 | 5.93 | 0.87 |
| 2 | 9.44 | 8.75 | 24.03 | 6.11 | 0.84 |
| 3 | 10.14 | 9.27 | 25.25 | 6.13 | 0.84 |
| 4 | 11.00 | 9.91 | 20.74 | 6.04 | 0.87 |
| High | 13.30 | 12.26 | 12.18 | 5.79 | 0.93 |
| High-Low | 3.05 | 2.95 | | | |
| t-stat. | (5.08) | (3.89) | | | |

Table 5: Returns on Double-Sorted Portfolios Formed on Risk Exposures

The table lists the average equal-weighted returns in percent per annum of independent double sorts on estimated cash-flow and long-term discount-rate betas in Panel A, cash-flow and short-term discount-rate betas in Panel B, and long-term discount-rate and short-term discount-rate betas in Panel C. The time-varying betas are computed recursively, on a quarterly basis, in 40-quarter overlapping rolling time-series estimation windows following definitions (12), (18) and (19). Stocks with one or more missing data points in a specific estimation window are deleted from the cross-sectional regression for that window. Returns in each window have been winsorized at the 1% and 99% levels. There are 162 overlapping estimation windows in total. The first window covers the period 1963Q3-1973Q2; the last window covers the period 2003Q4-2013Q3. For each risk characteristic, we first rank individual stocks into terciles and form portfolios of average returns in each estimation window. Portfolio Low contains stocks with the lowest betas, whereas portfolio High contains stocks with the highest betas. We then form 3x3 portfolios which are the intersections of respectively two categories of beta sorts. "High-Low" reports the difference between the returns of portfolio High and portfolio Low. "t-stat." is the HAC t -statistic for the difference in returns on portfolio High and portfolio Low.

| Panel A: Portfolios Sorted by β_{cf} and β_{ldr} | | | | | |
|---|---------|--------|----------|----------|---------|
| Rank | Low CF | Med | High CF | High-Low | t-stat. |
| Low LDR | 8.25 | 9.47 | 12.61 | 4.36 | (6.55) |
| Med | 9.05 | 10.02 | 12.29 | 3.24 | (5.78) |
| High LDR | 11.26 | 11.65 | 13.94 | 2.68 | (4.67) |
| High-Low | 3.02 | 2.18 | 1.33 | | |
| t-stat. | (7.36) | (7.80) | (4.59) | | |
| Panel B: Portfolios Sorted by β_{cf} and β_{sdr} | | | | | |
| Rank | Low CF | Med | High CF | High-Low | t-stat. |
| Low SDR | 8.51 | 10.00 | 12.84 | 4.33 | (7.28) |
| Med | 8.70 | 9.89 | 12.39 | 3.69 | (8.71) |
| High SDR | 9.88 | 11.36 | 14.22 | 4.34 | (7.61) |
| High-Low | 1.37 | 1.37 | 1.38 | | |
| t-stat. | (2.74) | (3.39) | (2.70) | | |
| Panel C: Portfolios Sorted by β_{ldr} and β_{sdr} | | | | | |
| Rank | Low LDR | Med | High LDR | High-Low | t-stat. |
| Low SDR | 7.73 | 8.98 | 12.10 | 4.37 | (12.19) |
| Med | 8.54 | 9.88 | 12.56 | 4.02 | (11.57) |
| High SDR | 11.11 | 12.25 | 15.11 | 3.99 | (8.81) |
| High-Low | 3.39 | 3.27 | 3.01 | | |
| t-stat. | (6.72) | (6.91) | (4.81) | | |

Table 6: Size and Book-to-Market

The table reports Fama-MacBeth (1973) estimates of risk prices for the three-beta ICAPM in percent per annum and their HAC t -statistics in parentheses for various subsamples. In Panel A, we sort all companies in each rolling window based on their market capitalization and construct five quintiles. In Panel B, we sort all companies in each rolling window based on their book-to-market value and construct five quintiles. In Panel C, we consider all stocks and all but the lowest market equity stocks. "Diff." reports results of a t -test for differences in estimated long-term and short-term discount-rate risk premia.

| Panel A: Size Quintile Portfolios | | | | | |
|---|-----------------|-----------------|-----------------|-----------------|-----------------|
| | Small | 2 | Med | 4 | Large |
| λ_0 | 7.82 (12.05) | 8.10 (18.26) | 7.11 (13.48) | 5.53 (9.47) | 6.57 (15.73) |
| λ_{cf} | 6.23 (9.08) | 2.98 (2.99) | 3.25 (2.53) | 1.01 (1.22) | 0.09 (0.08) |
| λ_{ldr} | 5.97 (10.65) | 4.94 (8.27) | 4.21 (4.36) | 5.76 (7.17) | 2.70 (4.10) |
| λ_{sdr} | 3.98 (5.26) | 0.69 (0.66) | 0.94 (0.93) | 3.13 (3.62) | 0.83 (0.82) |
| Diff. | 2.00 (3.69) | 4.25 (3.65) | 3.27 (3.40) | 2.63 (2.61) | 1.88 (2.16) |
| Panel B: Book-to-Market Quintile Portfolios | | | | | |
| | Growth | 2 | Med | 4 | Value |
| λ_0 | 5.66 (9.09) | 4.96 (7.51) | 5.34 (7.12) | 6.39 (11.88) | 7.74 (13.16) |
| λ_{cf} | 4.58 (7.30) | 6.14 (6.66) | 5.74 (4.98) | 7.34 (9.01) | 8.12 (9.95) |
| λ_{ldr} | 4.79 (10.08) | 5.49 (10.27) | 5.99 (9.33) | 5.68 (7.21) | 6.08 (9.40) |
| λ_{sdr} | 2.06 (2.56) | 3.61 (4.32) | 4.36 (6.08) | 3.44 (3.93) | 3.14 (3.29) |
| Diff. | 2.73 (3.75) | 1.88 (2.72) | 1.63 (2.08) | 2.23 (3.44) | 2.93 (3.12) |

Panel C: All Stocks and All Apart From Small Stocks

| | All | Excl. Small |
|-----------------|-----------------|-----------------|
| λ_0 | 6.22 (10.53) | 6.30 (11.85) |
| λ_{cf} | 5.70 (7.80) | 3.64 (4.15) |
| λ_{ldr} | 5.07 (11.64) | 4.27 (7.39) |
| λ_{sdr} | 3.03 (4.15) | 1.25 (1.50) |
| Diff. | 2.04 (3.40) | 3.03 (3.46) |

Table 7: Expected Return Contributions

This table gives the time-series averages of expected return contributions attributed to each factor in the three-beta ICAPM in percent per annum and their HAC t -statistics in parentheses for various subsamples. We compute $\lambda_{jt} \times \bar{\beta}_{jt}$, where j denotes the cash-flow (cf), long-term discount rate (ldr) and short-term discount rate (sdr) risk factors, $\bar{\beta}_{jt}$ is the cross-sectional mean of beta for risk factor j computed following definitions (12), (18) and (19) over 40-quarter overlapping rolling time-series estimation window t and λ_{jt} is the respective premium estimate. In Panel A, we sort all companies in each rolling window based on their market capitalization and construct five quintiles. In Panel B, we sort all companies in each rolling window based on their book-to-market value and construct five quintiles. In Panel C, we consider all stocks and all but the lowest market equity stocks.

| Panel A: Size Quintile Portfolios | | | | | |
|--|--------|--------|--------|--------|---------|
| | Small | 2 | Med | 4 | Large |
| $\lambda_{cf} \times \bar{\beta}_{cf}$ | 2.17 | 1.07 | 1.12 | 0.30 | -0.12 |
| | (5.52) | (4.19) | (2.85) | (1.36) | (-0.50) |
| $\lambda_{ldr} \times \bar{\beta}_{ldr}$ | 1.46 | 1.53 | 1.08 | 1.80 | 0.48 |
| | (4.88) | (5.46) | (3.38) | (6.21) | (2.83) |
| $\lambda_{sdr} \times \bar{\beta}_{sdr}$ | 1.82 | 0.08 | 0.28 | 0.86 | 0.21 |
| | (3.96) | (0.21) | (0.90) | (3.09) | (0.74) |

| Panel C: Book-to-Market Quintile Portfolios | | | | | |
|---|--------|--------|--------|--------|--------|
| | Growth | 2 | Med | 4 | Value |
| $\lambda_{cf} \times \bar{\beta}_{cf}$ | 1.99 | 2.44 | 2.03 | 2.17 | 2.41 |
| | (7.10) | (7.66) | (6.70) | (7.30) | (5.82) |
| $\lambda_{ldr} \times \bar{\beta}_{ldr}$ | 1.34 | 1.58 | 1.56 | 1.32 | 1.43 |
| | (5.68) | (6.34) | (5.42) | (4.04) | (4.50) |
| $\lambda_{sdr} \times \bar{\beta}_{sdr}$ | 0.68 | 1.18 | 1.48 | 1.25 | 1.60 |
| | (2.37) | (4.48) | (5.70) | (3.52) | (2.80) |

| Panel C: All Stocks and All Apart From Small Stocks | | |
|---|--------|----------------|
| | All | Excl. Small |
| $\lambda_{cf} \times \bar{\beta}_{cf}$ | 2.25 | 1.33 |
| | (7.10) | (7.37) |
| $\lambda_{ldr} \times \bar{\beta}_{ldr}$ | 1.22 | 1.10 |
| | (6.08) | (5.43) |
| $\lambda_{sdr} \times \bar{\beta}_{sdr}$ | 1.14 | 0.35 |
| | (3.91) | (1.39) |

Table 8: Alternative VAR Specifications

We consider nine alternative VAR specifications to model discount rate news. Panel A describes our choice of the respective state vectors. The state variables include different combinations of the log excess market return (r_m), the log dividend-price ratio (dp), the log consumption-wealth ratio (cay), the short-term interest rate (i), the default yield spread (def), the log ten-year price-earnings ratio ($pe10$), the log one-year price-earnings ratio ($pe1$), the small stock value spread (vs), the term yield spread between long-term and short-term bonds (ty), stock variance ($svar$), inflation rate (cpi), the book-to-market spread (bm). The plus signs indicate the selected variables in our baseline specification and the alternative specifications we examine. Panel B reports Fama-MacBeth (1973) estimates of risk prices for the three-beta ICAPM in percent per annum and their HAC t -statistics in parentheses. "Diff." reports results of a t -test for differences in estimated long-term and short-term discount-rate risk premia.

| Panel A: Variables in VAR | | | | | | | | | | | | |
|---------------------------|-------|------|-------|-----|-------|--------|-------|------|------|--------|-------|------|
| Specification | r_m | dp | cay | i | def | $pe10$ | $pe1$ | vs | ty | $svar$ | cpi | bm |
| Baseline | + | + | + | + | + | | | | | | | |
| I | + | + | + | + | | | | | | | | |
| II | + | + | + | | + | | | | | | | |
| III | + | | + | + | | | + | | | | | |
| IV | + | | + | + | + | + | | | | | | |
| V | + | + | | | | | | + | | | | + |
| VI | + | | | | | + | | + | | | | + |
| VII | + | | | | + | | + | | | + | | |
| VIII | + | | | | | + | | | | | | + |
| IX | + | + | + | | | | | | + | | | + |

Panel B: Cross-Sectional Risk Premium Estimates

| | I | II | III | IV | V | VI | VII | VIII | IX |
|-----------------|-----------------|------------------|-----------------|-----------------|-----------------|-----------------|------------------|------------------|-----------------|
| λ_0 | 6.27 (10.96) | 6.35 (11.09) | 6.35 (11.11) | 6.23 (10.50) | 6.68 (11.69) | 6.65 (11.52) | 6.30 (11.19) | 6.40 (10.65) | 6.30 (11.05) |
| λ_{cf} | 5.28 (7.54) | 11.58 (9.72) | 3.85 (7.27) | 5.83 (7.82) | 7.50 (3.98) | 8.03 (4.16) | 6.15 (9.16) | 11.50 (7.46) | 7.27 (8.36) |
| λ_{ldr} | 5.62 (7.51) | 24.62 (6.35) | 6.35 (11.26) | 4.43 (8.70) | 7.95 (1.90) | 4.20 (2.04) | 212.28 (2.15) | -6.10 (-2.88) | 9.58 (8.88) |
| λ_{sdr} | 3.63 (7.00) | -2.48 (-1.58) | 4.52 (10.38) | 3.23 (4.76) | 2.06 (1.56) | 1.89 (1.48) | -4.59 (-1.72) | 0.51 (0.44) | 0.87 (0.83) |
| Diff. | 2.00 (2.29) | 27.10 (5.13) | 1.83 (2.90) | 1.20 (2.63) | 5.89 (1.90) | 2.31 (2.29) | 216.87 (2.15) | -6.61 (-4.42) | 8.71 (5.09) |

Table 9: Parameter choices and the Estimated Prices of Risk

Model parameters reported in Panel A are calibrated to the available data set spanning periods as long as possible. The unconditional mean of dividend growth g , the risk-free rate r_f , the persistence variables $\phi_{\bar{x}}$ and ϕ_x , and the standard variance of short- and long-term discount rate shocks ($\|\sigma_{\bar{x}}\|$ and $\|\sigma_x\|$) are estimated using data from 1872Q1 to 2015Q4. The persistence variables ϕ_z , the standard variance of Δd_{t+1} and z_{t+1} ($\|\sigma_d\|$ and $\|\sigma_z\|$), and the elements in the correlation matrix (the ρ 's, i.e., correlation coefficients between shocks of the state variables pair-wisely) are estimated using data from 1952Q2 to 2015Q4 due to the lack of data for *cay* prior to year 1952. All numbers reported are on quarterly frequency. Panel B reports the price of risks for short-term discount rates shocks ($E[\phi_x x_t]$) and the long-term discount rate shocks ($E[(1 - \phi_x)\bar{x}_t]$). The bottom line reports the p -values from the Welch two-sample tests.

Panel A: Parameter choices

| Variable | Baseline Model | Full-fledged Model |
|---------------------------|----------------|--------------------|
| g | 0.0125 | 0.0125 |
| r_f | 0.0088 | 0.0088 |
| $\mu_{\bar{x}}$ | 0.0264 | 0.0264 |
| ϕ_z | 0.9190 | 0.9190 |
| $\phi_{\bar{x}}$ | 0.9660 | 0.9660 |
| ϕ_x | 0.3980 | 0.3980 |
| $\ \sigma_d\ $ | 0.1250 | 0.1250 |
| $\ \sigma_z\ $ | 0.0095 | 0.0095 |
| $\ \sigma_{\bar{x}}\ $ | 0.0703 | 0.0703 |
| $\ \sigma_x\ $ | 0.0259 | 0.0259 |
| $\rho_{x,\bar{x}}$ | 0.6210 | 0.6210 |
| $\rho_{\Delta d,z}$ | - | -0.7800 |
| $\rho_{\Delta d,\bar{x}}$ | - | 0.6110 |
| $\rho_{\Delta d,x}$ | - | 0.4160 |
| $\rho_{z,\bar{x}}$ | - | -0.4470 |
| $\rho_{z,x}$ | - | -0.4860 |

Panel B: Price of the discount rate risk

| Prices | Baseline Model | Full-fledged Model |
|----------------------------|----------------|--------------------|
| $E[(1 - \phi_x)\bar{x}_t]$ | 0.453 | 0.460 |
| $E[\phi_x x_t]$ | 0.299 | 0.304 |
| Diff. | 0.154 | 0.156 |
| p -value | 0.000 | 0.000 |

Figure 1: Forecast and Actual Excess Return

The figure plots the fitted values (black, blue and red) from long-horizon forecasting regressions of log excess return on the value-weighted CRSP market index tabulated in Table 1. Actual returns (dotted) are plotted on the same date as their forecast. The horizons are in quarters ($h = 1, 40, 60, 80$).

Figure 2: Annualized Holding Period Returns For Zero-Coupon Equities

The top panel shows risk premia on zero-coupon equity over the constant risk-free rate. The middle panel shows the standard deviation of returns on zero-coupon equity. The bottom panel shows the corresponding Sharpe ratio. Returns on zero-coupon equities of varying maturities are simulated at quarterly frequency, and aggregated to an annual frequency.

Figure 3: Annualized Risk Premia For Holding Aggregate Market Portfolios

This figure plots the annualized risk premia by holding market portfolios over the period from 1 year to 50 years. All return series are computed as the exponential mean of holding period returns subtracting the corresponding risk-free rates. The annualized risk premia for holding CRSP value-weighted aggregate market portfolio (CRSPVW) are computed using actual annual return series spanning the period from year 1927 to 2015. The blue line plots the annualized risk premia based on the annual returns on market portfolio generated by the baseline model. And the dashed line plots the annualized risk premia based on the annual returns on market portfolio generated by the full-fledged model. The dotted line plots the simulative annualized risk premia for holding the aggregate market portfolio generated by the Lettau and Wachter (2007) model.