

VU Research Portal

Three Essays in Asset pricing and Portfolio Choice

Botshekan, M.

2012

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Botshekan, M. (2012). *Three Essays in Asset pricing and Portfolio Choice*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Chapter 2

Cash Flow and Discount Rate Risk in Up and Down Markets: What Is Actually Priced?¹

2.1 Introduction

The capital asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) has since long been the most well-used work-horse model to understand the origins of expected returns.² An important contribution to the ability of the CAPM to explain the cross-section of stock returns was made by Campbell and Vuolteenaho (2004). Using a return decomposition method originally proposed by Campbell and Shiller (1988) and Campbell (1991), they show that the beta of the basic CAPM can be disentangled into a discount rate risk and a cash flow risk related beta component. Campbell and Vuolteenaho argue that in an economy with many long-term investors cash flow risk should carry a larger premium than discount rate risk: for long-term investors the negative impact of surprise increases in discount rates on current realized returns is partially compensated by higher expected returns. Using their two-fold beta decomposition, Campbell and Vuolteenaho (2004) succeed in partially explaining cross-sectional phenomena such as the size and book-to-market stock return premia. For instance, growth stocks tend to have high betas for the market portfolio, but these betas are related to discount rate risk and therefore carry a lower premium. By contrast, the betas for value stocks mainly relate to cash flow risk and therefore carry a larger compensation, resulting in higher expected returns.

In the current paper, we propose a new four-beta decomposition of the CAPM to

¹This chapter is based on Botshekan, Kraeussl, and Lucas (2012) which is forthcoming in *Journal of Financial and Quantitative Analysis*.

²The CAPM has of course seen numerous extensions, such as additional pricing factors like size, value, and momentum (Fama and French (1993), Fama and French (1996), Jegadeesh and Titman (1993), and Carhart (1997)); liquidity (Amihud (2002), Pastor and Stambaugh (2003), and Acharya and Pedersen (2005)); preference-based factors such as the downside betas of Ang, Chen, and Xing (2006), and the co-skewness of Friend and Westerfield (1980) and Harvey and Siddique (2000); and factors relating to deviations from market equilibria, see Lettau and Ludvigson (2001).

enhance our understanding of the cross-section of expected returns. The motivation for this extension lies in the literature on asymmetric preferences. Following the seminal work of Kahneman and Tversky (1979), a large number of papers has shown that typical decision makers are loss averse: the negative experience of a loss looms about twice as large as the positive experience of a similarly sized gain. The notion that preferences to losses versus gains may be different has a long history in finance as well. Markowitz (1959) already suggested to replace the variance as a (symmetric) risk measure of returns by the asymmetric semi-variance. This idea has been extended to lower partial moments and to an equilibrium context, see for example Hogan and Warren (1974), Bawa and Lindenberg (1977), and Harlow and Rao (1989).

Empirically, the importance of downside risk is supported by Ang, Chen, and Xing (2006). They define up and down betas by conditioning a stock's co-variation with the market on up and down markets. Using standard asset pricing tests, they find that equity risk premia correlate with downside betas, but not as much with upside betas. Their findings suggest that investors care more about the downside risk properties of stocks than about their general covariance properties.

A very similar line of reasoning applies to the “good” and “bad” beta model of Campbell and Vuolteenaho (2004). If the market goes down, loss averse investors experience a disproportionately large increase in marginal utility due to their asymmetric, kinked utility function, see for example the model in Ang et al. (2006). This by itself causes stocks with a higher covariation with downside market movements to require larger expected returns in equilibrium. For long-term, loss-averting investors, downside market movements due to bad cash flow news are worse than downside market movements due to unexpected discount rate increases. The intuition follows along the same lines as in the original paper by Campbell and Vuolteenaho (2004). As a result, if a sufficiently large fraction of the investor population consists of long term loss averters, assets that are exposed to downside cash flow shocks carry the largest premium in equilibrium. To test this conjecture, we define a four-beta model, where we measure a stock return's co-variation with cash flow and discount rate news separately in up and down markets.

Using our new four-beta decomposition and U.S. stock returns over the period 1963–2008, we investigate how the four components of beta are priced in the cross-section of stocks. We use Fama and MacBeth (1973) regressions with time-varying betas to obtain risk premia estimates. We find that both downside cash flow risk and downside discount rate risk are significantly priced and typically carry the largest premia. The upside pricing factors are less in magnitude and less robust. In particular and in line with our expectations, the downside cash flow risk is most consistently priced over different sub-periods in our sample. The magnitude, statistical significance, and even sometimes the sign of the other components is much more sensitive to the period used.

Interestingly, we find a strong relation between company size and downside cash flow

risk. For small stocks, we obtain the largest estimated premia for the downside risk components. By contrast, moving to larger companies, the priced components of risk become more symmetric (both upside and downside) and are cash flow related. Such a pattern can only be established in our proposed four-beta return decomposition and suggests that investors may take a different attitude towards risk compensation for small versus large stocks. If we control however for book-to-market rather than for size, no such pattern can be found. Both growth and value companies in our sample carry significant premia for all four risk components, with the premia related to downside risk dominating the upside risk premia.

A crucial step in our whole analysis is the direct construction of discount rate news via a vector autoregression (VAR) model for returns. The constructed discount rate news factor is combined with the returns to back out the cash flow news factor. Chen and Zhao (2009) criticize this decomposition approach and argue that it can be highly sensitive to the variables used in the VAR model. In particular, it matters whether discount rate news is modeled (via expected returns) and cash flow news is backed out, or whether one goes the other way around. Campbell, Polk, and Vuolteenaho (2010), Chen (2010), and Engsted, Pedersen, and Tanggaard (2010) argue that the sensitivity to the decomposition sequence can be reduced considerably by including the dividend yield as one of the state variables in the VAR model. We follow this approach in our paper by including the dividend yield as a state variable in the VAR model. Still, to account for the criticism as voiced in Chen and Zhao (2009), we also test explicitly whether our results are robust to the decomposition method used. We do so by constructing direct measures of cash flow news. We confirm that the decomposition method to some extent affects the size of estimated premia. However, we still find that the downside cash flow and downside discount rate components carry the largest compensation, thus confirming our baseline results.

Estimated risk premia are only one component of required returns. The latter are obtained by multiplying each risk premium by its appropriate beta and summing over all different risk factors. To obtain insight into the economic impact of the different risk components on average returns, we therefore also investigate the significance of the time-varying risk premia estimates multiplied by their time varying betas. In contrast to the results for the premia alone, we find that the discount rate related components of expected returns are largest. This implies that though investors charge a higher price for downside cash flow risk exposure, the sensitivity of the average stock to this risk factor is smaller than the sensitivity to discount rate news. The impact of the downside risk components, however, remains consistently statistically significant and positive.

As a final test of our model, we investigate whether our betas also have out-of-sample predictive power. We carry out a recursive analysis of estimating a VAR model for returns, computing the risk factors and risk factor sensitivities, and forecasting returns

out-of-sample. For one-month out-of-sample forecasts, we do not find significant results, as one-month returns are very noisy signals of expected returns. Using five-year out-of-sample average returns, results are very clear: downside cash flow risk is the only beta component that has a statistically significant price out of sample. The price of 4.5% per annum is smaller than for the in-sample results (6% p.a.), but surprisingly close.

There are several studies that tried to develop asset pricing models based on the return decomposition approach of Campbell and Vuolteenaho (2004) to explain cross-sectional differences in average returns, see for example, Chen and Zhao (2010), Da and Warachka (2009), Koubouros, Malliaropulos, and Panopoulou (2007), Koubouros, Malliaropulos, and Panopoulou (2010), and Maio (2009). To the best of our knowledge, however, no one has tried to disentangle the pricing properties of cash flow and discount rate news in up and down markets. The closest in this respect is the recent work by Campbell, Giglio, and Polk (2010). These authors estimate the different magnitudes of discount rate and cash flow news in two particularly bad market settings: the burst of the tech bubble and the stock market downturn of 2000-2001, and the financial crisis of 2007-2008. They conclude that the 2000-2001 crisis is mainly driven by bad cash flow news, whereas the more recent financial turmoil has a large bad discount rate news component to it. In contrast to Campbell et al. (2010), our paper does not study the composition over time of the news factors themselves, but rather focuses on the different pricing properties of discount rate and cash flow news in different market settings and over a longer period of time.

The remainder of this paper is organized as follows. Section 2.2 provides the background to our four-beta return decomposition model and introduces the methodology used for the empirical tests. Section 2.3 describes the data. Section 2.4 discusses the empirical results and robustness checks. Section 2.5 concludes.

2.2 Methodology

2.2.1 Downside and Upside Betas

Following the seminal work of Kahneman and Tversky (1979), there is sufficient empirical evidence supporting the view that typical investors are loss averse, i.e., their disutility of a large loss is higher than the positive utility of a similarly sized gain. Asymmetric preferences were already used in the early finance literature to provide alternatives to the standard CAPM, which is based on the symmetric concept of variance. Markowitz (1959), for example, introduced the notion of semi-variance as a measure of risk. The notion was exploited and extended in asset pricing theory by Hogan and Warren (1974), Bawa and Lindenberg (1977), and Harlow and Rao (1989).

Harlow and Rao (1989) use the expected market return to distinguish between up and

down markets. Their equilibrium framework gives rise to a downside beta, defined as

$$\beta_{i,D} = \frac{\text{E}[(R_{it} - \mu_i)(R_{mt} - \mu_m) | R_{mt} < \mu_m]}{\text{E}[(R_{mt} - \mu_m)^2 | R_{mt} < \mu_m]}, \quad (2.1)$$

where R_i and R_m are the return on stock i and on the market portfolio, with expectations μ_i and μ_m , respectively. Analogously, the upside beta can be defined as

$$\beta_{i,U} = \frac{\text{E}[(R_{it} - \mu_i)(R_{mt} - \mu_m) | R_{mt} \geq \mu_m]}{\text{E}[(R_{mt} - \mu_m)^2 | R_{mt} \geq \mu_m]}. \quad (2.2)$$

Ang et al. (2006) show that the cross-section of stock returns reflects a downside risk premium of approximately 6% p.a. They investigate whether the upside beta, downside beta, or both have a premium in the cross-section and find that risk premia mainly reflect a stock's downside and not its upside beta. They rationalize their findings by appealing to an economy with loss-averse agents. Such agents assign greater weight to the downside movements of the market than to upside movements. In this way, Ang et al. (2006) argue that downside risk is a separate risk attribute from other well-known risk premium determinants such as size, book-to-market, momentum, and liquidity.

2.2.2 Cash Flow and Discount Rate Betas

Campbell and Vuolteenaho (2004) take a different perspective and decompose the market return into two components related to cash flow risk and discount rate risk, respectively. Using these two components, the beta of a stock can be decomposed analogously. Part of beta is due to co-variation of the individual stock's return with the market's discount rate news factor. This is the so-called discount rate beta. The other part is due to co-variation with the market's cash flow factor and is called the cash flow beta. Campbell and Vuolteenaho label the discount rate beta as 'good', and the cash flow beta as 'bad'. Their terminology stems from the fact that discount rate news has two offsetting effects. If discount rates increase unexpectedly, current prices decrease and realized returns are negative. For long-term investors, however, these wealth decreases are partially offset by increases in expected returns, as the investment opportunity set has improved.

Campbell and Vuolteenaho argue that the presence of many long-term investors in the market causes a higher premium for assets that co-vary more with the market's cash flow news than with the discount rate factor. They also show that different loadings to cash flow news and discount rate news explain part of the size and value premia puzzles. The main reason is that while growth stocks (which have low average returns) have high betas for the market portfolio, these betas are predominantly 'good' betas with low risk premia. Value stocks, by contrast, have high average returns, but also higher 'bad' betas than growth stocks. Similarly, small stocks have considerably higher cash flow betas than large stocks, which is in line with the higher average realized returns for these stocks.

Our approach to decompose market returns in their discount rate and cash flow components is similar to Campbell and Vuolteenaho (2004) and uses the return decomposition technique of Campbell and Shiller (1988) and Campbell (1991). Campbell and Shiller (1988) use a log-linear approximation of the present value relation for stock prices that allows for time-varying discount rates. They obtain the following return decomposition

$$\begin{aligned} r_{m,t+1} - E_t r_{m,t+1} &\approx (E_{t+1} - E_t) \sum_{i=0}^{\infty} \rho^i \Delta d_{t+1+i} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{m,t+1+j} \\ &\equiv N_{CF,t+1} - N_{DR,t+1}, \end{aligned} \quad (2.3)$$

where r_{mt} is the log market return at time t , d_t is the log dividend paid by the stock at time t , Δ denotes the first difference operator, E_t denotes the rational expectations operator given the information set available at time t , and ρ is a linearization parameter defined as $\rho \equiv 1/(1 + \exp(\overline{dp}))$, where \overline{dp} is the average log dividend price ratio. We follow Campbell and Vuolteenaho (2004) and assume an annual value of $\rho = 0.95$. The factor $N_{CF,t+1}$ denotes news about future cash flows, i.e., the change in the discounted sum of current and future expected dividend growth rates. Similarly, $N_{DR,t+1}$ denotes news about future discount rates, i.e., the change in the discounted sum of future expected returns.

Following the decomposition of the market return into two separate news factors, we can define two separate betas. The cash flow beta is given by

$$\beta_{i,CF} = \frac{\text{cov}(R_{i,t}, N_{CF,t})}{\text{var}(u_{mt})}, \quad (2.4)$$

and the discount rate beta by

$$\beta_{i,DR} = \frac{\text{cov}(R_{i,t}, -N_{DR,t})}{\text{var}(u_{mt})}, \quad (2.5)$$

where $u_{mt} = r_{mt} - E_{t-1} r_{mt} = N_{CF,t} - N_{DR,t}$ is the unexpected market return at time t .

The key step in operationalizing (2.3) and calculating (2.4) and (2.5) is to postulate a model for expected returns $E_t[r_{t+j}]$ for $j = 0, 1, \dots$. We follow the standard approach as in Campbell and Vuolteenaho (2004) and assume the data is generated by a vector autoregression (VAR), so that the discount rate and cash flow news can be backed out directly from the VAR residuals.

The VAR model is given by

$$z_{t+1} = a + \Gamma z_t + u_{t+1}, \quad (2.6)$$

$$r_{m,t+1} = e_1' z_{t+1}, \quad (2.7)$$

where z_{t+1} is a $k \times 1$ state vector with $r_{m,t+1}$ as its first element, a is a $k \times 1$ vector of constants, Γ is an $k \times k$ matrix of coefficients, e_1 is the first column from the $k \times k$ unit matrix I_k , and u_{t+1} is a vector of serially independent random shocks. The first element

of u_{t+1} thus equals the unexpected market return at time $t + 1$, $e'_1 u_{t+1} = u_{m,t+1}$. By recursively substituting (2.6) in (2.3), we obtain the cash flow and discount rate factors as

$$N_{DR,t+1} = e'_1 \Lambda u_{t+1}, \quad \text{and} \quad N_{CF,t+1} = e'_1 (\mathbf{I}_k + \Lambda) u_{t+1}, \quad (2.8)$$

with $\Lambda = \rho \Gamma (\mathbf{I}_k - \rho \Gamma)^{-1}$.

The VAR approach is the dominant method in the return decomposition literature. Chen and Zhao (2009) argue that the results based on the VAR methodology are sensitive to the decision to forecast expected returns explicitly while treating cash flow components as residuals, as in (2.8). Campbell et al. (2010), however, argue that when the VAR contains the dividend-price ratio as a state variable, there is little difference between (i) an approach that backs out the cash flow news component from a directly modeled discount rate news component, and (ii) an approach that backs out the discount rate news component from a modeled cash flow component. The argument was already made more generally by Ang and Liu (2007) and also particularly pointed out in the context of return decompositions by Chen (2010): return, dividend growth, and dividend yield are related by a (linearized) accounting identity, such that one can use each combination of two variables to back out the third. Chen (2010) therefore recommends that the dividend yield should always be included as a state variable in the VAR model. The findings are confirmed by Engsted et al. (2010), who show that the VAR model has to include the dividend-price ratio in order for the decomposition to be independent of which news component is treated as a residual.

Based on the above arguments, we also include the dividend yield in our VAR model. However, to still check the robustness of our results to the decomposition method used, we also provide results based on alternative methods of return decomposition that use direct cash flow modeling (see Section 2.4).

2.2.3 The Four-Beta Model

The decomposition of Campbell and Vuolteenaho (2004) does not make a distinction between upside and downside risk. The arguments based on asymmetric preferences by investors are, however, equally applicable in a context where we disentangle cash flow and discount rate risk. In particular, given the pricing results in Ang et al. (2006) as well as in Campbell and Vuolteenaho (2004), it is unclear whether downside risk is priced higher than upside risk, or whether cash flow risk is priced higher than discount risk, or any combination of these. In particular, we would like to test for the price of downside risk, cash flow risk, and discount rate risk, while controlling for the other types of risk. In order to do this, we propose a new four-fold beta model. The aim of this model is to isolate the relative importance of the cash flow and discount rate news components in up and down markets. This allows us to better pinpoint the origins of risk premia in the cross-section of stock returns. The new model distinguishes four different betas: a downside cash flow

(*DCF*) beta, a downside discount rate (*DDR*) beta, an upside cash flow (*UCF*) beta, and an upside discount rate (*UDR*) beta. Following the earlier definitions, the betas are defined as:

$$\beta_{i,DCF} = \text{E}[(R_{it} - \mu_i)N_{CF,t} | u_{mt} < 0] / \text{E}[u_{mt}^2 | u_{mt} < 0], \quad (2.9)$$

$$\beta_{i,DDR} = -\text{E}[(R_{it} - \mu_i)N_{DR,t} | u_{mt} < 0] / \text{E}[u_{mt}^2 | u_{mt} < 0], \quad (2.10)$$

$$\beta_{i,UCF} = \text{E}[(R_{it} - \mu_i)N_{CF,t} | u_{mt} \geq 0] / \text{E}[u_{mt}^2 | u_{mt} \geq 0], \quad (2.11)$$

$$\beta_{i,UDR} = -\text{E}[(R_{it} - \mu_i)N_{DR,t} | u_{mt} \geq 0] / \text{E}[u_{mt}^2 | u_{mt} \geq 0]. \quad (2.12)$$

By differentiating between the covariance of returns with the discount rate factor and cash flow factor in up and down markets, respectively, we can control for risk factors in both dimensions simultaneously. Note that the definitions in (2.9) through (2.12) are completely analogous to (2.4) and (2.5). The main difference is that we have conditioned the expectations on the unexpected market return u_{mt} being positive or negative. As the unexpected market return has zero mean by construction, zero is also the natural cut-off point to distinguish up from down markets. Also note that by construction, the discount rate and cash flow factors have zero means as they are directly based on the innovations u_t in (2.8).

The four different betas in (2.9) through (2.12) can now be used in standard asset pricing tests. In particular, we test the relative importance of estimated premia for different components in our new four-beta model,

$$\text{E}_t[R_{i,t+1}^e] = \alpha_1 + \lambda_{DCF} \cdot \beta_{i,DCF} + \lambda_{DDR} \cdot \beta_{i,DDR} + \lambda_{UCF} \cdot \beta_{i,UCF} + \lambda_{UDR} \cdot \beta_{i,UDR}, \quad (2.13)$$

where $R_{i,t+1}^e$ denotes the excess return (over the risk-free rate) for asset i , α_i is the intercept for asset i , and λ_j is the price of risk for $\beta_{i,j}$ for $j = DCF, DDR, UCF, UDR$. We benchmark our results to the simpler two-way decompositions of beta of Ang et al. (2006) and Campbell and Vuolteenaho (2004).

In our empirical analysis in Section 2.4 we follow Black, Jensen, and Scholes (1972), Gibbons (1982), and Ang et al. (2006) by testing the contemporaneous relationship between betas and the realized average returns (as a proxy for expected returns). We perform Fama-MacBeth regressions with time-varying betas estimated over 60-months rolling windows from July 1963 to December 2008. In this way, we can compute a time-series of estimated risk premia corresponding to the time-varying betas. The test then considers whether the time-series mean of risk premia is positive and significantly different from zero. We use overlapping windows to estimate the betas, and heteroskedasticity and autocorrelation consistent (HAC) standard errors for our pricing tests, see Andrews (1991).

Table 2.1: VAR Parameter Estimates for the Return Decomposition Model

This table shows the OLS estimates of the vector autoregressive (VAR) model (2.6). The dependent variables are the log excess market return ($R_{m,t}^e$), the short-term interest rate (SR_t), and the dividend yield (DY_t). Standard errors are given in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

	Intercept	$R_{m,t}^e$	SR_t	DY_t	$R^2\%$
$R_{m,t+1}^e$	-0.004 (0.005)	0.089** (0.038)	-0.199*** (0.067)	0.539*** (0.162)	3.04
SR_{t+1}	0.001 (0.001)	0.007* (0.004)	0.992*** (0.007)	-0.012 (0.017)	97.5
DY_{t+1}	0.000* (0.000)	-0.015*** (0.001)	0.004*** (0.002)	0.987*** (0.004)	99.2

2.3 Data

A return decomposition based on a VAR model should contain state variables with sufficient predictive ability. As argued by Campbell et al. (2010), Engsted et al. (2010), and Chen (2010), it is particularly important to include dividend yields in the analysis to reduce the sensitivity of the results to the precise VAR model used. We therefore specify the following three variables in our VAR model: (i) the log excess market return defined as the log of the CRSP value weighted market index minus the log of the three-month Treasury bill rate; (ii) the three-month Treasury bill rate itself; and (iii) the dividend yield on the S&P 500 composite price index calculated from data provided on Robert Shiller's website. In their original paper, Campbell and Vuolteenaho (2004) also stress the importance of the small-stock value spread as an important element of their VAR model. Over the sample period used in our paper (1963–2008), however, the variable turns out to be statistically insignificant and we exclude it from the further analyses.² The ability of the dividend yield to predict excess expected returns has been largely accepted and documented in the finance literature, see for example Campbell (1991), Cochrane (1992, 2008), and Lettau and Ludvigson (2001). Ang and Bekaert (2007) point out that this is best visible at short horizons by specifying the short-term interest rate as an additional regressor. They are more skeptical about the predictive power of dividend yields in the long-term. We therefore also include the short-term interest rate in our analysis.

Table 2.1 shows the VAR parameter estimates. Both the short-term interest rate and the dividend yield are highly persistent and have a statistically significant impact on stock returns. As expected, higher interest rates have a negative impact on returns, while the relation between dividend yields and returns is positive.

Using the VAR model from Table 2.1, we construct the cash flow ($N_{CF,t}$) and discount

²An online appendix to this paper is available that replicates most of the results of this paper using an extended six-variable rather than a three-variable VAR system. The six variables include the same three as in the current paper, as well as the term spread, the credit spread, and the small-stock value spread.

Table 2.2: Variance-Covariance Matrix of Cash Flow and Discount Rate News

This table shows the variance-covariance matrix of the unexpected market return (u_{mt}) and its two components, cash flow (CF) news and discount rate (DR) news, using the three-variable VAR model from Table 2.1. The VAR model includes the excess market return R_{mt} (above the risk-free rate), the short (3-month) rate SR_t , and the S&P500 dividend yield DY_t .

	u_{mt}	$N_{CF,t}$	$N_{DR,t}$
u_{mt}	0.0018	0.0006	0.0012
$N_{CF,t}$	0.0006	0.0007	-0.0001
$N_{DR,t}$	0.0012	-0.0001	0.0013
Mean	0.0013	0.0082	-0.0069

rate ($N_{DR,t}$) news factors from the VAR residuals using equation (2.8). The variance-covariance matrix of the news factors is presented in Table 2.2. The variance of DR news is almost twice the size of the CF news variance. Campbell (1991) finds similar results with the discount rate news being the dominant component of market return variance.

The test assets we use in our pricing regressions are individual stocks and not portfolios. The use of portfolios in the cross-sectional Fama-MacBeth regressions is fairly standard to mitigate the errors-in-variables problem caused by the use of estimated rather than true betas. However, this advantage comes at the cost of a significant loss of efficiency due to the reduced cross-sectional spread of estimated betas. This is particularly relevant in our current context as our model tries to identify four separate beta-related pricing components. Using portfolios as test assets then results in too much multicollinearity in the cross-sectional estimation step of the Fama-MacBeth procedure.³ As a result, the risk premia estimates would become unstable.

Ang, Liu, and Schwarz (2008) show analytically and empirically that the conclusions drawn from individual versus portfolio test assets can differ substantially due to the trade-off between bias and efficiency. They also indicate that the use of individual stocks as test assets generally permits better asset pricing tests and estimates of risk premia. We therefore follow their conclusion that there is no particular reason to create portfolios when just two-pass cross-sectional regression coefficients are estimated. Instead, it is preferable to run the asset pricing tests in such cases based on individual stocks. All tests presented in the next section are therefore based on all individual common stocks traded on the NYSE, AMEX, and NASDAQ exchanges over the period July 1963 to December 2008 (share codes 10 or 11 in the CRSP database). In our robustness checks we vary the sample period as well as the sampling frequency to see whether our baseline results remain valid. For the analyses based on monthly data, we use data from the CRSP-Compustat merged database in WDRS. For the analyses based on quarterly data, we take all available data from the CRSP database.

³Correlations between estimated up, down, cash flow, and discount rate betas for portfolio test assets are typically in excess of 95%.

2.4 Empirical Results

2.4.1 Baseline Results

Table 2.3 presents our baseline results. The first 60-months window spans from July 1963 to June 1968 and the last one from January 2004 to December 2008. We thus have 486 overlapping 60-months windows in total. The number of stocks in each cross-section varies from 383 in earlier periods to 3,703 in later periods. In order to ensure that extreme outliers do not drive our findings, we winsorize returns in each window at the 1% and 99% level.

Column I in Table 2.3 shows that the standard beta has a significant and positive premium. When we decompose the beta in an up and a down beta as in column II, we see that both betas carry a significant premium at the 1% and 5% level, respectively. The average premium for the downside beta is almost six times that for the upside beta. To get a clearer impression on the contribution of downside betas, we include column II-B. The regressors in the cross-sectional steps of the Fama-MacBeth procedure are taken as β_{it} and $(\beta_{it,D} - \beta_{it,U})$ rather than $\beta_{it,U}$ and $\beta_{it,D}$. We see a similar effect as before: the traditional beta is priced significantly, but on top of that the additional contribution of downside betas is priced as well.

Model III presents the results for the cash flow and discount rate beta model. Both cash flow and discount rate betas are priced significantly. In contrast to Campbell and Vuolteenaho (2004), there appears to be no significant difference between the two premia.

Model IV presents the results for our new four-beta model. The downside cash flow (*DCF*) and downside discount rate (*DDR*) betas carry the largest premia and are significant at the 1% level. The upside cash flow (*UCF*) and upside discount rate (*UDR*) betas are also significant at the 5% and 10% level, respectively, but the size of the *DCF* and *DDR* premia are about three times as high as the *UCF* and *UDR* premia. This implies that both cash flow (*CF*) and discount rate (*DR*) betas are priced more in down than in up markets. In line with our intuition, the downside CF beta carries the largest premium. From Ang et al. (2006) we would expect investors to charge higher premia for downside risk. From Campbell and Vuolteenaho (2004), on the other hand, we would expect a larger premium for CF betas. Our results show that both of these effects have explanatory power in the cross-section, and exposure to downside CF news carries the largest premium. It is also clear that the four-fold beta decomposition provides additional information here: in the standard two-fold decomposition of CF versus DR (model III), we find no significant difference in premia.

If we again alter the specification to test for the additional effect of downside cash flow and discount rate risk (model IV-B), the results are confirmed. The difference between (symmetric) cash flow and discount rate risk is much clearer than in model III after we allow for an additional downside risk component. Also, both downside risk components are

Table 2.3: Baseline Risk Premia Estimates

This table shows the time-series averages and their HAC standard errors (in parentheses) of the Fama-MacBeth premia estimates λ_t^i , where t denotes the 60-months rolling window and j denotes the risk factor, being downside (D), upside (U), additional downside ($D - \beta$) cash flow (CF), discount rate (DR), downside cash flow (DCF), downside discount rate (DDR), upside cash flow (UCF), upside discount rate (UDR), additional downside cash flow risk ($DCF - CF$), and additional downside discount rate risk ($DDR - DR$), respectively. The sample consists of monthly returns for all listed companies on the NYSE, AMEX, and NASDAQ exchanges from July 1963 to December 2008 (546 months), using the CRSP-Compustat merged database in WRDS. There are 486 sixty-months overlapping estimation windows in the sample. Stocks with one or more missing data points in a specific estimation window are deleted from the cross-sectional regression for that cross-sectional window. The number of stocks in each cross-sectional regression varies from 383 to 3,703. Returns in each window have been winsorized at the 1% level and 99% level. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

	I	II	II-B	III	IV	IV-B	V	VI	VI-B	VII	VIII	VIII-B
α	0.299*** (0.064)	0.273*** (0.064)	0.273*** (0.064)	0.314*** (0.063)	0.293*** (0.063)	0.293*** (0.063)	0.706*** (0.210)	0.674*** (0.210)	0.674*** (0.210)	0.787*** (0.208)	0.768*** (0.203)	0.768*** (0.203)
λ	0.474*** (0.057)		0.490*** (0.057)				0.516*** (0.047)		0.533*** (0.047)			
λ_D		0.420*** (0.051)						0.386*** (0.039)				
λ_U		0.071** (0.034)						0.148*** (0.036)				
$\lambda_{D-\beta}$			0.269*** (0.075)						0.150*** (0.071)			
λ_{CF}				0.507*** (0.073)		0.655*** (0.078)				0.603*** (0.076)		0.733*** (0.077)
λ_{DR}				0.526*** (0.087)		0.466*** (0.088)				0.591*** (0.075)		0.546*** (0.074)
λ_{DCF}					0.525*** (0.062)						0.480*** (0.054)	
λ_{DDR}					0.378*** (0.066)						0.385*** (0.049)	
λ_{UCF}					0.130** (0.065)						0.253*** (0.058)	
λ_{UDR}					0.088* (0.048)						0.161*** (0.052)	
λ_{DCF-CF}						0.323*** (0.108)						0.153* (0.089)
λ_{DDR-DR}						0.195*** (0.083)						0.118 (0.081)
Size							-0.060*** (0.015)	-0.059*** (0.015)	-0.059*** (0.015)	-0.066*** (0.014)	-0.065*** (0.014)	-0.065*** (0.014)
B/M							0.323*** (0.026)	0.327*** (0.026)	0.327*** (0.026)	0.315*** (0.026)	0.319*** (0.025)	0.319*** (0.025)
R^2	0.072	0.082	0.082	0.084	0.102	0.102	0.144	0.151	0.151	0.153	0.167	0.167

priced significantly, with the price of additional downside risk for the cash flow component dominating in size.

To investigate whether our baseline results are robust to size and book-to-market effects, we re-specify our models I to IV by adding the Fama and French (1992) size and book-to-market factors to the cross-sectional regressions.⁴ To account for influential observations, we also winsorize the size and book-to-market controls at the 1% and 99% level. Columns V to VIII-B of Table 2.3 show that most of the premia estimates are robust to controlling for size and book-to-market effects. There appears to be a mild shift downwards in the *DCF* premium, and an upward shift in the *UCF* and *UDR* premia (model VIII). All shifts fall well within the two standard error bands. The main difference occurs if we re-specify our model to measure the additional effect of the downside risk components (models VI-B and VIII-B). In model VI-B, the size of the additional downside risk premium is somewhat smaller, and its significance drops from 1% to 5%. If we further refine the model to distinguish between the additional downside cash flow and downside discount rate components in model VIII-B, only the additional downside cash flow component stays significant at the 10% level. The inter-relation between firm size and downside risk compensation is further investigated later on in this section. Consistent with Fama and French (1992), we find a robust and significantly negative premium for size, and a significantly positive premium for book-to-market.

We have also investigated whether there are differences between the values of downside betas (or downside incremental betas) across industry groups. To save space, the full results are not reported here, but are available upon request. The differences between downside betas across industries are statistically significant for many industries. The economic significance, however, is limited. The differences decrease further if we consider the incremental downside betas (models II-B, IV-B, VI-B, and VIII-B in Table 2.3). In that case, most of the industries have incremental downside betas that are not statistically significantly different from each other.

To investigate the time-series properties of the premia estimates, we re-estimate our four-beta model over the different decades in our sample. Each of the four sub-periods describes a different episode of the stock market. During the 1970s, the U.S. economy was hit by several recessions, including the two major oil price crises. During the 1980s, the U.S. economy suffered by the savings and loans crisis. In the 1990s, U.S. equity experienced a strong bull market. This rally led to the burst of the tech bubble in early 2000 followed by the financial crises at the end of our sample period 2007-2008. Panel A of Table 2.4 presents the results.

Comparing the premia for the *DCF* and the *UCF* beta, we find that the *DCF* beta

⁴Size is the log market capitalization at the start of each 60-months window. With respect to the book-to-market factor, we follow Fama and French (1992): for January till June of year t , we take the book value as of end-December of year $t - 2$, and for July till December of year t , we take the book value as of end-December of year $t - 1$. The book value is then divided by the current market value of equity.

is robustly priced in all four subsamples. The *UCF* beta, however, is only significantly priced during the stock market rally of the 1990s. The *DDR* and *UDR* premia show opposite and trending results over time. Sensitivity to *DDR* news is priced high in the cross-section at the start of our sample and during the 1980s. Over the 1990s and 2000s, however, the price declined and even becomes insignificant during the last decade. By contrast, the sensitivity to *UDR* news carries a negative price in the early years of the sample, but gradually increases over time to a positive and significant premium in the 2000s. During this last decade, the *UDR* premium is even the largest of the four premia.

Overall, our subsample analysis indicates that downside betas are priced more robustly than upside betas, which is consistent with our previous results over the whole sample period of 1963 to 2008. Considering the *UDR* beta, we obtain mixed evidence of positive and negative premia in different periods. The only beta component that is robustly priced throughout all subsamples remains the *DCF* beta, followed by the *DDR* beta.

To control further for possible size and book-to-market effects, we test our factor models using five subsamples constructed by sorting the data with respect to size and book-to-market, respectively. First, we sort our sample based on market capitalization (respectively book-to-market) at the beginning of each 60-months window of the Fama-MacBeth estimation procedure and divide the cross-section into five quintiles. Then, we compute our estimate of the premium by running the cross-sectional regressions for each of the five quintiles separately. The process is repeated for all estimation windows.

Panel B of Table 2.4 shows a clear effect of size on the estimated premia for the four-beta model. The *DDR* beta premium is lower for the largest two quintiles, and the decline is statistically significant. For the *DCF* premium, the decrease for large cap companies is much less strong, though also statistically significant. For the *UDR* and *UCF* premia, we see a much more constant pattern across size quintiles. In particular, there is no statistically significant difference between the premia estimates for large versus small companies, though the *UCF* premium shows a mild increase for increasing company size.

Comparing the relative magnitudes of the different premia, we see that for small companies the downside components are the dominant pricing ingredients. For large companies, however, it is predominantly the cash flow component that is relevant. In particular, the impact of the cash flow component appears more symmetric, with the magnitude of the premia for *DCF* and *UCF* being roughly the same. This suggests that the notion of downside risk is much more relevant for small companies, irrespective of whether this is due to downside cash flow or due to downside discount rate risk. If well-established companies are considered, a much more symmetric notion of stock market risk appears to apply, mainly relating to *CF* rather than to *DR* news.

Table 2.4: Subsample Analysis

This table shows the premia estimates and their standard errors as in Table 2.3, but for different subsamples. Panel A shows the results for different decades. In Panel B, we sort all companies for each rolling window based on their market capitalization at the beginning of the period and construct 5 quintiles. In Panel C, we sort all companies based on their book-to-market value at the beginning of each rolling window. Premia are computed for each quintile. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

Panel A: Sample Periods					
	1970s	1980s	1990s	2000s	1963-2008
α	0.085 (0.076)	0.191** (0.092)	0.240** (0.115)	0.805*** (0.124)	0.293*** (0.063)
λ_{DCF}	0.246** (0.096)	1.007*** (0.086)	0.642*** (0.091)	0.209* (0.118)	0.525*** (0.062)
λ_{DDR}	0.473*** (0.149)	0.579*** (0.101)	0.194*** (0.068)	0.093 (0.111)	0.378*** (0.066)
λ_{UCF}	0.228 (0.199)	-0.032 (0.044)	0.432*** (0.069)	-0.009 (0.057)	0.130** (0.065)
λ_{UDR}	-0.315*** (0.055)	-0.015 (0.045)	0.159*** (0.031)	0.526*** (0.112)	0.088* (0.048)
Panel B: Size					
	Small	2	3	4	Large
α	0.327*** (0.094)	0.186** (0.078)	0.310*** (0.071)	0.451*** (0.050)	0.472*** (0.050)
λ_{DCF}	0.586*** (0.047)	0.573*** (0.061)	0.553*** (0.084)	0.355*** (0.110)	0.228** (0.106)
λ_{DDR}	0.620*** (0.061)	0.469*** (0.070)	0.325*** (0.057)	0.058 (0.069)	0.066 (0.105)
λ_{UCF}	0.186*** (0.050)	0.273*** (0.072)	0.251*** (0.067)	0.361*** (0.067)	0.285*** (0.078)
λ_{UDR}	0.149*** (0.028)	0.142** (0.061)	0.123 (0.075)	0.153** (0.074)	0.084 (0.069)
Panel C: Book-to-Market (B/M)					
	Low	2	3	4	High
α	-0.136 (0.085)	0.040 (0.066)	0.186*** (0.060)	0.389*** (0.062)	0.491*** (0.074)
λ_{DCF}	0.556*** (0.088)	0.462*** (0.078)	0.534*** (0.073)	0.553*** (0.064)	0.775*** (0.068)
λ_{DDR}	0.464*** (0.071)	0.547*** (0.068)	0.377*** (0.069)	0.308*** (0.073)	0.422*** (0.069)
λ_{UCF}	0.188*** (0.072)	0.153** (0.072)	0.210*** (0.066)	0.134** (0.061)	0.167*** (0.063)
λ_{UDR}	0.112* (0.062)	0.194*** (0.064)	0.225*** (0.053)	0.179*** (0.045)	0.069** (0.033)

Panel C of Table 2.4 displays the results for the book-to-market quintiles. In contrast to the results in Panel B, we do not observe a clear pattern. Only the *DCF* premium appears to be somewhat larger for the highest book-to-market quintile, and the difference with the other quintiles is significant at the 5% level. We do see the higher premia again

for the downside factors compared to the upside ones. The downside premia are two-fold up to five-fold their upside counterparts. The downside cash flow premium is higher than the downside discount rate premium. The difference is significant for the higher book-to-market quintiles. For the upside premia UDR and UCF , there is no such clear difference. Again, we conclude that downside cash flow risk is consistently priced and carries the largest premium, followed by downside discount rate risk. The upside risk factors are less consistently priced and smaller in magnitude.

Overall, both asymmetric preferences for downside versus upside risk as well as for long-term versus short-term risk play a major role in explaining the cross-section of stock returns. Our new four-beta model helps to isolate the effects of these different components on market risk premia. The baseline results show that both DCF and DDR betas are priced more robustly in the cross-section, while both UCF and UDR betas are not priced consistently. The only component that is priced robustly over all samples is the DCF beta. Downside betas have larger premia than their upside counterparts in most subsamples. However, downside risk particularly appears to be a concern for small stocks, while expected returns for larger stocks appear to be driven more by a symmetric notion of cash flow risk.

2.4.2 Robustness Analysis

So far, we have computed discount rate news (as the change in the discounted sum of future expected returns) directly, treating CF news as the residual outcome, i.e., as the unexpected market return minus the computed DR news factor. Chen and Zhao (2009) argue that such a definition of cash flow news influences the size of the premia estimates. Campbell et al. (2010), Chen (2010), and Engsted et al. (2010), however, show that the sensitivity of premia estimates and factor sensitivities to the decomposition method used is reduced considerably by including the dividend yield in the underlying VAR model. Still, to check the sensitivity of our results, we follow Chen and Zhao (2009) and investigate the robustness of our four-beta model to alternative decomposition methods. In particular, we build an additional VAR model to construct CF news directly, rather than as a residual. For more details, we refer to Chen and Zhao (2009).

The VAR model for dividend growth takes the lagged dividend growth rate and the lagged market excess returns as explanatory variables. To reduce seasonality issues while retaining a reasonable number of observations for the time-series regressions, we use quarterly rather than monthly (or annual) data from 1963:Q3 to 2008:Q4. The CF news component at time $t + 1$ is computed as

$$N_{CF,t+1}^{\text{dir}} = e_1' \Lambda_2 \nu_{t+1}, \quad (2.14)$$

where $\Lambda_2 = (\mathbf{I} - \rho \Gamma_2)^{-1} \Gamma_2$; Γ_2 is the coefficient matrix of the VAR model for dividend growth, ν_{t+1} denotes the vector of VAR residuals, and the first element in this second

specified VAR model is the dividend growth. We can compute the correlation between our direct estimate of cash flow news $N_{CF,t}^{\text{dir}}$ from (2.14) with our previous indirect estimate $N_{CF,t}$. As in Chen and Zhao (2009), the correlation between the two estimates is far from perfect. In our case the correlation is only 0.291. Part of this low correlation may be due to the simple VAR model used to construct the direct estimate of cash flow news, as the dividend growth rate is notoriously difficult to model. Despite this low correlation, the results presented below indicate that the consistent significance of downside cash flow news as a priced risk factor stays robust. The current analysis therefore provides a strong robustness check for our claims on the relevance of the downside cash flow in stock returns.

As a further robustness check, we also compute the results with an alternative computation for the discount rate news component. As mentioned earlier, we originally computed $N_{DR,t}$ directly, and computed $N_{CF,t}$ as the residual. With our new $N_{CF,t}^{\text{dir}}$ cash flow risk factor, we can also take the opposite perspective and define $N_{DR,t}$ as the residual. We do so by defining

$$-N_{DR,t}^{\text{res}} = u_{mt} - N_{CF,t}^{\text{dir}}, \quad (2.15)$$

with u_{mt} as the unexpected return from the VAR model for returns, see Section 2.2.2. The correlation between the indirect DR news factor $N_{DR,t}^{\text{res}}$ and the original direct DR news factor $N_{DR,t}$ is again not perfect with a value of 0.823. Interestingly, however, the construction of the discount rate news factor appears less sensitive to the decomposition method used.

Panel A in Table 2.5 presents the results for our three different decomposition methods. We use a 40-quarter rolling window to estimate different betas and average returns, resulting in 143 overlapping windows. Because we only use price data in this exercise, the number of stocks varies from 1,158 to 2,678 per cross-section, as we do not lose observations by matching CRSP price data with Compustat book value data.

We observe that the DCF , DDR , and UDR betas always have a positive and significant premium irrespective of the decomposition method used. The estimates of the DCF and DDR premia are larger than their UCF and UDR counterparts, implying the downside risk dimension is more important, irrespective of the decomposition method used. We also note that the UCF factor is not consistently priced across decomposition methods. This reinforces our conclusion regarding the price impact of downside risk.

It becomes also clear that the choice of the decomposition method influences the size of the premium estimates. Particularly the DCF premium, and to a lesser extent the DDR premium, is higher if a direct measure of cash flow news is used. The larger price for downside risk under the alternative decomposition methods is in line with our earlier results: the downside risk components, and the downside cash flow related parts in particular, carry the largest price.

Table 2.5: Robustness Analysis for Alternative Decomposition Methods

Panel A shows the Fama-MacBeth premia estimates λ_j and their HAC standard errors (in parentheses) for j equal to downside cash flow (DCF), downside discount rate (DDR), upside cash flow (UCF), and upside discount rate (UDR) risk, respectively. The estimates are based on three different decomposition methods for computing cash flow and discount rate news. The sample contains quarterly return data for all listed companies on the NYSE, AMEX, and NASDAQ exchanges over July 1963 to December 2008 (182 quarters). We use a 40-quarter rolling window to estimate betas and average returns. Stocks with one or more missing data points in a specific estimation window are deleted from the cross-sectional regression for that window. The number of stocks varies from 1,158 to 2,678 over the sample. Method I uses a direct measure for DR news and an indirect measure for CF news as in (2.8). Method II uses a direct measure for DR news and a direct measure for CF news as in (2.14). Method III uses an indirect measure for DR news and a direct measure for CF news as in (2.15). Panel B reports the time-series averages and their HAC standard errors of $\lambda_{jt} \cdot \bar{\beta}_{jt}$, where $\bar{\beta}_{jt}$ is the cross-sectional mean of beta for risk factor j over the 40-quarter rolling window t , and λ_{jt} is the premium estimate for risk factor j over the same window. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

	Panel A: Premium Estimates			Panel B: Expected Return Contributions ($\lambda \cdot \bar{\beta}$)		
	I	II	III	I	II	III
α	0.825*** (0.196)	1.309*** (0.226)	0.868*** (0.197)			
λ_{DCF}	1.931*** (0.190)	3.790*** (0.555)	2.487*** (0.514)	0.851*** (0.082)	0.258*** (0.051)	0.144*** (0.035)
λ_{DDR}	0.868*** (0.110)	1.391*** (0.140)	1.267*** (0.100)	0.596*** (0.078)	0.926*** (0.094)	1.325*** (0.113)
λ_{UCF}	0.769*** (0.143)	-0.133 (0.283)	0.098 (0.238)	0.315*** (0.069)	0.034* (0.018)	0.008 (0.015)
λ_{UDR}	0.601*** (0.081)	0.702*** (0.105)	0.582*** (0.096)	0.405*** (0.051)	0.465*** (0.065)	0.648*** (0.111)

2.4.3 Economic Significance

So far, we have focused on the premia estimates λ_j for $j = DCF, DDR, UCF, UDR$. The expected returns, however, are a composite of these premia and their associated β_{ij} s. For example, it might well be the case that the higher observed premia are partly offset by lower average levels of β for a particular segment of the stock market. In order to provide more insight into the economic magnitude of the product of betas and their premia, we perform the following analysis: for each window of the Fama-MacBeth procedure, we compute the product of the premium estimate and the cross-sectional average beta over that window. In this way, we obtain the contribution of the risk factor j to the overall expected return in the rolling window t . Subsequently, we compute the time-series averages of all these contributions and their HAC standard errors.

The most right column in Panel A of Table 2.6 shows that over the complete sample period 1963–2008 the expected return component $\lambda_j \times \beta_j$ is again largest for the downside components $j = DCF, DDR$. Moreover, the downside components are statistically significant, whereas the upside components UCF and UDR are not. In contrast to some of our earlier results for the premia λ_j (see Table 2.3), the product of betas and premia is

larger for the *DDR* factor than for the *DCF* component. The higher premium on average for the *DCF* that we observed in Table 2.3 is thus off-set by a lower average *DCF* beta compared to the *DDR* beta. [-0.1cm] Examining the results for the different decades in the sample, we observe that the expected return component for the *DDR* is stable over most of the sample, except during the 2000s. Over the last decade, the significance level drops to 10% and the component becomes slightly negative. The opposite holds for the *UDR* factor. Both results are in line with our earlier findings for risk premia only (see Table 2.4). The altered monetary regime over the 2000s combined with the bear markets during the burst of the tech bubble in the year 2000 and the 2007-2008 financial crisis significantly affect the premia estimates for discount rate news. [-0.1cm] Comparing the *DCF* and *DDR* expected return components, we find that the contribution of downside discount rates to expected returns is higher up to the year 2000. Over the most recent period, the downside cash flow is the dominant factor in expected returns. We also note that the *DCF* component is consistently and significantly positive, though its magnitude over the different periods differs. All other pricing components, by contrast, show statistically significant changes of sign over time. This supports our earlier view of *DCF* as the most consistently priced factor in the data.

In order to control for the possible size effect as documented in Table 2.4, we also calculate the expected return components using five size-sorted subsamples. Our empirical findings in Panel B of Table 2.6 indicate once more a substantial and significant size effect. For small companies, both *DCF* and *DDR* betas capture the expected returns, with the *DDR* component twice the size of its *DCF* counterpart. Again, we conclude that downside cash flow risk carries the largest premium, but firms' exposures to this risk factor on average are smaller, such that the downside discount rate risk makes up a larger fraction of the average returns.[-0.1cm]

For the large companies, average returns can neither be explained by downside or upside preferences, nor by the long-term versus short-term decomposition. Though the components are significant at (only) the 10% level, they are economically small.[-0.1cm]

Finally, computing the components across decomposition methods, we obtain the results in Panel B of Table 2.5. Though downside cash flow risk carries the largest premium, the dominant contribution of downside cash flow news sensitivity to average returns hinges on the choice of the decomposition method I, i.e., based on direct discount rate new modeling. If direct measures of cash flow news are used as in methods II and III, downside and upside discount rate news are more important contributors to average returns. In terms of statistical significance, only the *UCF* component appears not robust.

2.4.4 Betas and Future Returns

Our final test is the most challenging one and provides an out-of-sample robustness check. In testing factor models using Fama-Macbeth regressions, different proxies can be taken

Table 2.6: Expected Return Contribution for Monthly Data

The data and set-up used for this table are the same as for Table 2.3 and Table 2.4. This table shows the time-series averages and their HAC standard errors (in parentheses) of $\lambda_{jt} \cdot \bar{\beta}_{jt}$, where $\bar{\beta}_{jt}$ is the cross-sectional mean of beta for risk factor j over the 60-months rolling window t , and λ_{jt} is the premium estimate for risk factor j , with j equal to downside cash flow (*DCF*), downside discount rate (*DDR*), upside cash flow (*UCF*), and upside discount rate (*UDR*) risk, respectively. Panel A splits the sample period in different decades. Panel B uses size-sorted subsamples. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

Panel A: Sample Periods					
	1970s	1980s	1990s	2000s	1963-2008
λ_{DCF}	0.079* (0.043)	0.218*** (0.041)	0.087* (0.037)	0.523*** (0.081)	0.212*** (0.034)
λ_{DDR}	0.343*** (0.105)	0.561*** (0.060)	0.396*** (0.062)	-0.075* (0.048)	0.353*** (0.050)
λ_{UCF}	0.034* (0.064)	-0.023* (0.010)	0.064*** (0.018)	0.062*** (0.021)	0.028* (0.019)
λ_{UDR}	-0.070* (0.047)	-0.001* (0.026)	-0.045*** (0.015)	0.164*** (0.034)	0.023* (0.021)
Panel B: Size					
	Small	2	3	4	Large
λ_{DCF}	0.202*** (0.021)	0.196*** (0.027)	0.187*** (0.035)	0.096* (0.040)	0.055* (0.040)
λ_{DDR}	0.445*** (0.048)	0.367*** (0.051)	0.219*** (0.040)	0.010* (0.047)	-0.008* (0.056)
λ_{UCF}	0.014* (0.015)	0.052* (0.022)	0.055* (0.026)	0.102*** (0.024)	0.072* (0.028)
λ_{UDR}	0.051* (0.020)	-0.013* (0.040)	-0.022* (0.048)	0.012* (0.046)	-0.024* (0.040)

for the expected return. So far, we have followed the approach of Lettau and Ludvigson (2001) and Ang et al. (2006) and used in our cross-sectional regressions estimated betas and average returns that were measured over the same data window. We label this approach model I. Alternatively, however, we could use the betas over the estimation window to forecast the next month's out-of-sample return. We label this as model II. This is a substantial step, as one-month returns are obviously much noisier proxies of expected returns than five-year averages of monthly returns. One can thus expect the results to become less clear than in previous sections.

We also specify another model. The cross-sectional regressions in this model III use the betas from the estimation window to forecast the return average of the next 60-months out-of-sample returns. This proxy should be of similar quality as that of model I. The crucial difference, however, is that model I checks the in-sample predictive power of betas, whereas model III tests the out-of-sample predictive power. Table 2.7 present the results.

A possible criticism to this out-of-sample exercise is that a look-ahead-bias is evident due to the use of a VAR model. So far, our VAR model was estimated on the complete

Table 2.7: Current Betas and Future Expected Returns: A Rolling VAR

This table shows the time-series averages and their corresponding HAC standard errors (in parentheses) of Fama-MacBeth estimates of the premia for downside cash flow (*DCF*), downside discount rate (*DDR*), upside cash flow (*UCF*), and upside discount rate (*UDR*) risk. Model I uses the cross-sectional Fama-MacBeth regressions based on 60-months rolling window estimates of betas and average returns over the same rolling window. Model II uses the same betas, but uses the next month's return following the rolling window as its dependent variable. Model III uses the average of the next 60-months out-of-window returns as the dependent variable. Models IV to VI are similar to I to III, but also include size and book-to-market controls. The data is the same as for Table 2.3. The main difference is that we use for the return decomposition in model I one single VAR estimated over the whole period. But for the models II and III, we use an in-sample rolling VAR for each window of Fama-Macbeth regressions, which uses data available till the end of that window for the VAR estimation and corresponding return decomposition. ***, **, and * denote significance at the 1, 5, and 10 percent level, respectively.

	I	II	III	IV	V	VI
	60m in-sample	1m out-of-sample	60m out-of-sample	60m in-sample	1m out-of-sample	60m out-of-sample
α	0.293*** (0.063)	0.596*** (0.210)	0.700*** (0.059)	0.768*** (0.203)	0.199 (0.566)	1.104*** (0.195)
λ_{DCF}	0.525*** (0.062)	-0.180 (0.198)	0.343*** (0.061)	0.480*** (0.054)	-0.182 (0.183)	0.227*** (0.047)
λ_{DDR}	0.378*** (0.066)	0.080 (0.175)	-0.023 (0.041)	0.385*** (0.049)	0.182 (0.158)	0.054* (0.032)
λ_{UCF}	0.130** (0.065)	0.034 (0.165)	-0.044 (0.043)	0.253*** (0.058)	0.006 (0.149)	0.051 (0.038)
λ_{UDR}	0.088* (0.048)	0.014 (0.110)	-0.027 (0.032)	0.161*** (0.052)	0.009 (0.100)	0.006 (0.027)
Size				-0.065*** (0.014)	0.020 (0.036)	-0.048*** (0.015)
B/M				0.319*** (0.025)	0.175** (0.082)	0.212*** (0.019)

sample and the residuals were used to construct the cash flow and discount rate news factors. These factors, in turn, were used to compute the betas over the estimation window. Via the VAR model, the betas at a particular time therefore also contain information of future returns. To resolve this potential bias issue, we also recursively estimate the VAR model in our forecast test. As an example, for the estimation window January 1981 till December 1985, we use all the data up to December 1985 to estimate the VAR model and compute the news factors. These news factors are used to estimate the betas over the estimation window. The returns over the (future) period January 1986 (one-month return) or January 1986 to December 1990 (five-year average of monthly returns) are then regressed on the estimated betas. Finally, the estimation window is rolled forward by one month, and all computations are repeated.

The results for model I are the same as presented earlier. As expected for model II, the data is very noisy. As a result, hardly any of the estimates are statistically significant. The estimates of the downside risk components *DDR* and *DCF* have the correct sign. Their magnitude, however, has decreased substantially, whereas the standard errors have

gone up. The higher standard errors are a natural consequence of the fact that one-month returns are very noisy proxies of expected returns. The upside components have an estimated negative premium. For the upside discount rate factor, the estimate is even statistically significant.

If we take average returns over 60-months rather than over one-month out-of-sample as our dependent variable in the cross-sectional regressions (as in model III), results become clear again. The *DCF* beta is the only one that is priced significantly, and it carries a positive premium of around 3.5% p.a. This is smaller than the in-sample estimate of around 6% p.a., but still substantial. All point estimates of the other premia are very close to zero, and only *UCF* is significant at the 10% level. We also note that the standard errors for model III are again in line with those for model I.

These results are very supportive of the *DCF* factor as carrying the largest and most consistent premium, not only in-sample, but also out-of-sample. This finding is further confirmed if we add the size and book-to-market controls to the regressions as in models IV–VI. As expected, the size of the *DCF* premium is somewhat reduced, but it remains the only significant premium estimate among the four beta related premia.

Overall, in an out-of-sample context, the results once again support our major finding that investors want to be compensated for downside risk. The required compensation can mainly be attributed to the downside cash flow component.

2.5 Conclusions

Through a decomposition of the simple CAPM beta into four components (downside and upside cash flow betas and downside and upside discount rate betas) we show that we can increase our understanding of the types of risk that investors want to be compensated for. Using individual U.S. stock returns over the period 1963–2008, we find that downside cash flow (*DCF*) risk is most consistently priced in our sample. Downside discount rate (*DDR*) risk is next in line, followed by upside cash flow (*UCF*) risk and upside discount rate (*UDR*) risk, respectively. Our results survive a range of robustness tests with respect to sampling periods, use of controls, and use of decomposition methods. In particular, the *DCF* premium is the only robust premium of the four beta premia in an in-sample versus out-of-sample comparison.

Interestingly, we find that the downside risk compensation is closely linked to the small stock premium. For small-sized companies, the *DCF* and *DDR* beta appear the main pricing determinants. For larger companies, a much more symmetric notion of cash flow risk (both upside and downside) applies. It thus appears that downside risk is mainly relevant for small stocks. Apparently, differences in the investor base of large versus small stocks cause asymmetric preferences to be less relevant in the large stock segment.