Downside Risk*

Andrew Ang†
Columbia University, USC and NBER

Joseph Chen‡
USC

Yuhang Xing§
Rice University

This Version: 26 July 2004

*This paper is a substantial revision of an earlier paper titled "Downside Correlation and Expected Stock Returns." The authors thank Brad Barber, Geert Bekaert, Alon Brav, John Cochrane, Randy Cohen, Kent Daniel, Bob Dittmar, Rob Engle, Wayne Ferson, David Hirschleifer, N. Jegadeesh, Jonathan Lewellen, Qing Li, Terence Lim, Toby Moskowitz, Ľuboš Pástor, Akhtar Siddique, Rob Stambaugh, and Zhenyu Wang. We especially thank Cam Harvey (the editor) and Bob Hodrick for detailed comments. We thank seminar participants at Columbia University, Koç University, NYU, USC, the European Finance Association, the Five Star Conference, an NBER Asset Pricing Meeting, the Texas Finance Festival, and the Western Finance Association for helpful discussions. We thank two referees whose comments greatly improved the paper. The authors acknowledge funding from a Q-Group research grant.

†Marshall School of Business at USC, Hoffman Hall 701, Los Angeles, CA 90089-1427; ph: (213) 740-5615; fax: (213) 740-6650; email: aa610@columbia.edu; WWW: http://www.columbia.edu/∼aa610.

‡Marshall School of Business at USC, Hoffman Hall 701, Los Angeles, CA 90089-1427; ph: (213) 740-6509; fax: (213) 740-6650; email: joe.chen@marshall.usc.edu; WWW: http://www-rcf.usc.edu/∼josephsc.

§Jones School of Management, Rice University, Rm 230, MS 531, 6100 Main Street, Houston, TX 77004; ph: (713) 348-4167; email: yxing@rice.edu; WWW: http://www.ruf.rice.edu/∼yxing

Abstract

Agents who place greater weight on the risk of downside losses than they attach to upside gains demand greater compensation for holding stocks with high downside risk. We show that the cross-section of stock returns reflects a premium for downside risk. Stocks that covary strongly with the market when the market declines have high average returns. We estimate that the downside risk premium is approximately 6% per annum and demonstrate that the compensation for bearing downside risk is not simply compensation for market beta. Moreover, the reward for downside risk is not explained by coskewness or liquidity risk, and is also not explained by momentum, book-to-market, size, or other cross-sectional characteristics.

1 Introduction

If an asset tends to move downward in a declining market more than it moves upward in a rising market, it is an unattractive asset to hold because it tends to have very low payoffs precisely when the wealth of investors is low. Investors who are more sensitive to downside losses, relative to upside gains, require a premium for holding assets that covary strongly with the market when the market declines. Hence, assets with high sensitivities to downside market movements have high average returns in equilibrium. In this article, we show that the cross-section of stock returns reflects a premium for bearing downside risk. In contrast, we fail to find a significant discount for stocks that have high covariation conditional on upside movements of the market.

The reason that stocks with large amounts of downside risk have high average returns is intuitive. As early as Roy (1952), economists have recognized that investors care differently about downside losses than they care about upside gains. Markowitz (1959) advocates using semi-variance as a measure of risk, rather than variance, because semi-variance weights downside losses differently from upside gains. More recently, the behavioral framework of Kahneman and Tversky's (1979) loss aversion preferences, and the axiomatic approach taken by Gul's (1991) disappointment aversion preferences, allow agents to place greater weights on losses relative to gains in their utility functions. Hence in equilibrium, agents who are averse to downside losses demand greater compensation, in the form of higher expected returns, for holding stocks with high downside risk.

According to the Capital Asset Pricing Model (CAPM), a stock's expected excess return is proportional to its market beta, which is constant across down-markets and up-markets. As Bawa and Lindenberg (1977) suggest, a natural extension of the CAPM that takes into account the asymmetric treatment of risk is to specify separate downside and upside betas. We compute downside (upside) betas over periods where the excess market return is below (above) its mean. However, despite the intuitive appeal of downside risk, which closely corresponds to how individual investors actually perceive risk, there has been little empirical research into how downside risk is priced in the cross-section of stocks returns.

The paucity of research on a downside risk premium may be due to weak improvements over the CAPM that early researchers found by allowing betas to differ across the downside and the upside. For example, in testing a model with downside betas, Jahankhani (1976) fails to find any improvement over the traditional CAPM. Harlow and Rao (1989) find more support for downside and upside betas, but they only evaluate downside risk relative to the CAPM in a maximum likelihood framework and test whether the return on the zero-beta asset is the same across all assets. They do not directly demonstrate that assets that covary more with the market, conditional on market downturns, have higher average returns.[1] Hence, it is not surprising that recently developed multi-factor models, like the Fama and French (1993) three-factor model, account for the failure of the CAPM by emphasizing the addition of different factors, rather than incorporating asymmetry in the factor loadings across down-markets and up-markets.

Our strategy for finding a premium for bearing downside risk in the cross-section is as follows. First, we directly show at the individual stock level that stocks with higher downside betas contemporaneously have higher average returns. Second, we claim that downside beta may be a risk attribute because stocks that have high covariation with the market when the market declines exhibit high average returns over the same period. This contemporaneous relationship between factor loadings and risk premium is the foundation of a cross-sectional risk-return relationship, and has been exploited from the earliest tests of the CAPM (see, among others, Black, Jensen and Scholes, 1979; Gibbons, 1982). Fama and French (1992) also seek, but fail to find, a relationship between post-formation market betas and realized average stock returns.

Third, we differentiate the reward for holding high downside risk stocks from other known cross-sectional effects. In particular, Rubinstein (1973), Friend and Westerfield (1980), Kraus and Litzenberger (1976 and 1983), and Harvey and Siddique (2000) show that agents dislike stocks with negative coskewness, so that stocks with low coskewness tend to have high average returns. Downside risk is different from coskewness risk because downside beta explicitly conditions for market downside movements in a non-linear fashion, while the coskewness statistic does not explicitly emphasize asymmetries across down and up markets, even in settings where coskewness may vary over time (as in Harvey and Siddique, 1999). Since the coskewness measure captures, by definition, some aspects of downside covariation, we are especially careful to control for coskewness risk in assessing the premium for downside beta. We also control for the standard list of known cross-sectional effects, including size and book-to-market factor loadings and characteristics (Fama and French, 1993; Daniel and Titman, 1997), liquidity risk factor loadings (Pástor and Stambaugh, 2003), and past return characteristics (Jegadeesh and

[1] Pettengill, Sundaram and Mathur (1995) and Isakov (1999) estimate the CAPM by splitting the full sample into two subsamples that consist of observations where the realized excess market return is positive or negative. Naturally, they estimate a positive (negative) market premium for the subsample with positive (negative) excess market returns. In contrast, our approach examines premiums for asymmetries in the factor loadings, rather than estimating factor models on different subsamples. Price, Price and Nantell (1982) demonstrate that skewness in U.S. equity returns causes downside betas to be different from unconditional betas, but do not relate downside betas to average stock returns.

Titman, 1993). Controlling for these and other cross-sectional effects, we estimate that the cross-sectional premium is approximately 6% per annum.

Finally, we check if past downside betas predict future expected returns. We find that, for the majority of the cross-section, high past downside beta predicts high future returns over the next month, similar to the contemporaneous relationship between realized downside beta and realized average returns. However, this relation breaks down among stocks with very high volatility. This is due to two reasons. First, the past downside betas of stocks with very high volatility contain substantial measurement error making past downside beta a poor instrument for future downside risk exposure Second, there is a confounding effect of anomalously low average returns exhibited by stocks with very high volatility (see Ang, Hodrick, Xing and Zhang, 2003). Fortunately, the proportion of the market where past downside beta fails to provide good forecasts of future downside market exposure is small (less than 4% in terms of market capitalization). Confirming Harvey and Siddique (2000), we find that past coskewness predicts future returns, but the predictive power of past coskewness is not because past coskewness captures future exposure to downside risk. Hence, past downside beta and past coskewness are different risk loadings.

The rest of this paper is organized as follows. In Section 2, we present a simple model to show how a downside risk premium may arise in the cross-section. The equilibrium setting uses a representative agent with the kinked disappointment aversion utility function of Gul (1991) that places larger weight on downside outcomes. Section 3 demonstrates that stocks with high downside betas have high average returns over the same period that they strongly covary with declining market returns. In Section 4, we show that past downside beta also cross-sectionally predicts returns for the vast majority of stocks. However, the predictive downside beta relation breaks down for stocks with very high volatility, which have very low average returns. Section 5 concludes.

2 A Simple Model of Downside Risk

In this section, we show how downside risk may be priced cross-sectionally in an equilibrium setting. Specifically, we work with a rational disappointment aversion (DA) utility function that embeds downside risk following Gul (1991). Our goal is to provide a simple motivating example of how a representative agent with a larger aversion to losses, relative to his attraction to gains, gives rise to cross-sectional prices that embed compensation for downside risk.[2]

[2] While standard power, or CRRA, utility does produce aversion to downside risk, the order of magnitude of a downside risk premium, relative to upside risk, is economically negligible because CRRA preferences are locally

Our simple approach does not rule out other possible ways in which downside risk may be priced in the cross-section. For example, Shumway (1997) develops an equilibrium behavioral model based on loss averse investors. Barberis and Huang (2001) use a loss averse utility function, combined with mental accounting, to construct a cross-sectional equilibrium. However, they do not relate expected stock returns to direct measures of downside risk. Aversion to downside risk also arises in models with constraints that bind only in one direction, for example, binding short-sales constraints (Chen, Hong and Stein, 2002) or wealth constraints (Kyle and Xiong, 2001).

Rather than considering models with one-sided constraints or agents with behavioral biases, we treat asymmetries in risk in a rational representative agent framework that abstracts from the additional interactions from one-sided constraints. The advantage of treating asymmetric risk in a rational framework is that the disappointment utility function is globally concave, whereas optimal finite portfolio allocations for loss aversion utility may not exist (see Ang, Bekaert and Liu, 2004). Our example with disappointment utility differs from previous studies, because existing work with Gul (1991)'s first order risk aversion utility concentrates on the equilibrium pricing of downside risk for only the aggregate market, usually in a consumption setting (see, for example, Bekaert, Hodrick and Marshall, 1997; Epstein and Zin, 1990 and 2001; Routledge and Zin, 2003).

Gul (1991)'s disappointment utility is implicitly defined by the following equation:

$$U(\mu_W) = \frac{1}{K} \left(\int_{-\infty}^{\mu_W} U(W) dF(W) + A \int_{\mu_W}^{\infty} U(W) dF(W) \right), \tag{1}$$

where $U(W)$ is the felicity function over end-of-period wealth W, which we choose to be power utility, that is $U(W) = W^{(1-\gamma)}/(1-\gamma)$. The parameter $0 < A \leq 1$ is the coefficient of disappointment aversion, $F(\cdot)$ is the cumulative distribution function for wealth, μ_W is the certainty equivalent (the certain level of wealth that generates the same utility as the portfolio allocation determining W) and K is a scalar given by:

$$K = Pr(W \leq \mu_W) + A Pr(W > \mu_W). \tag{2}$$

Outcomes above (below) the certainty equivalent μ_W are termed "elating" ("disappointing") outcomes. If $0 < A < 1$, then the utility function (1) down-weights elating outcomes relative to disappointing outcomes. Put another way, the disappointment averse investor cares more about downside versus upside risk. If $A = 1$, disappointment utility reduces to the special case of standard CRRA utility, which is closely approximated by mean-variance utility.

mean-variance.

To illustrate the effect of downside risk on the cross-section of stock returns, we work with two assets x and y. Asset x has three possible payoffs u_x, m_x and d_x, and asset y has two possible payoffs u_y and d_y. These payoffs are in excess of the risk-free payoff. Our set-up has the minimum number of assets and states required to examine cross-sectional pricing (the expected returns of x and y relative to each other and to the market portfolio, which consists of x and y), and to incorporate higher moments (through the three states of x). The full set of payoffs and states are given by:

Payoff	Probability
(u_x, u_y)	p_1
(m_x, u_y)	p_2
(d_x, u_y)	p_3
(u_x, d_y)	p_4
(m_x, d_y)	p_5
(d_x, d_y)	p_6

The optimal portfolio weight for a DA investor is given by the solution to:

$$\max_{w_x, w_y} U(\mu_W) \tag{3}$$

where the certainty equivalent is defined in equation (1), w_x (w_y) is the portfolio weight of asset x (y), and end of period wealth W is given by:

$$W = R_f + w_x x + w_y y, \tag{4}$$

where R_f is the gross risk-free rate. An equilibrium is characterized by a set of asset payoffs, corresponding probabilities, and a set of portfolio weights so that equation (3) is maximized and the representative agent holds the market portfolio ($w_x + w_y = 1$) with $0 < w_x < 1$ and $0 < w_y < 1$.

The equilibrium solution even for this simple case is computationally non-trivial because the solution to the asset allocation problem (3) entails solving for both the certainty equivalent μ_W and for the portfolio weights w_x and w_y simultaneously. In contrast, a standard portfolio allocation problem for CRRA utility only requires solving the FOC's for the optimal w_x and w_y. We extend a solution algorithm for the optimization (3) developed by Ang, Bekaert and Liu (2004) to multiple assets. Appendix A describes our solution method and details the values used in the calibration. Computing the solution is challenging because for certain parameter values equilibrium cannot exist. This is because non-participation may be optimal for low A under DA utility (see Ang, Bekaert and Liu, 2004). This is unlike the asset allocation problem under standard CRRA utility, where agents always optimally hold risky assets that have strictly positive risk premia.

In this simple model, the standard beta with respect to the market portfolio (β) is not a sufficient statistic to describe the risk-return relationship of an individual stock. In our calibration, an asset's expected returns increase with β, but β does not fully reflect all risk. This is because the representative agent cares in particular about downside risk, through $A < 1$. Hence, measures of downside risk have explanatory power for describing the cross-section of expected returns. One measure of downside risk introduced by Bawa and Lindenberg (1977) is the downside beta β^-:

$$\beta^- = \frac{\text{cov}(r_i, r_m | r_m < \mu_m)}{\text{var}(r_m | r_m < \mu_m)}, \tag{5}$$

where r_i (r_m) is security i's (the market's) excess return, and μ_m is the average market excess return. We also compute a relative downside beta $(\beta^- - \beta)$, relative to the standard CAPM β, where $\beta = \text{cov}(r_i, r_m)/\text{var}(r_m)$.

Figure 1 shows various risk-return relationships holding in our DA cross-sectional equilibrium. The mean excess return increases with β^-. We define the CAPM α as the excess return of an asset not accounted for by the asset's CAPM beta, $\alpha = \text{E}(r_i) - \beta\text{E}(r_m)$. The CAPM α is also increasing with β^- or $(\beta^- - \beta)$. Hence, higher downside risk is remunerated by higher expected returns.

The bottom right-hand panel of Figure 1 plots the CAPM α versus coskewness, defined as:

$$\text{coskew} = \frac{\text{E}[(r_i - \mu_i)(r_m - \mu_m)^2]}{\sqrt{\text{var}(r_i)}\text{var}(r_m)}, \tag{6}$$

where μ_i is the average excess return of asset i. Harvey and Siddique (2000) predict that lower coskewness should be associated with higher expected returns. The coskewness measure can be motivated by a third-order Taylor expansion of a general Euler equation:

$$\text{E}_t\left[\frac{U'(W_{t+1})}{U'(W_t)} r_{i,t+1}\right] = 0, \tag{7}$$

where W is the total wealth of the representative agent, and $U'(\cdot)$ can be approximated by:

$$U' = 1 + WU''r_m + \frac{1}{2}W^2U'''r_m^2 + \ldots. \tag{8}$$

The Taylor expansion in equation (8) is necessarily only an approximation. In particular, since the DA utility function is kinked, polynomial expansions of U, such as the expansions used by Bansal, Hsieh and Viswanathan (1993), may not be good global approximations if the kink is large (or A is very small).[3] Nevertheless, measures like coskewness based on the Taylor approximation for the utility function should also have some explanatory power for returns.

[3] Taylor expansions have been used to account for potential skewness and kurtosis preferences in asset allocation problems by Guidolin and Timmerman (2002), Jondeau and Rockinger (2003), and Harvey, Liechty, Liechty and Müller (2003).

Downside beta and coskewness may potentially capture different effects. Note that for DA utility, both downside beta and coskewness are approximations because the utility function does not have an explicit form (equation (1) implicitly defines DA utility). Since DA utility is kinked at an endogenous certainty equivalent, skewness, and other centered moments may not effectively capture aversion to risk across upside and downside movements in all situations. This is because they are based on unconditional approximations to a non-smooth function. In contrast, the downside beta in equation (5) conditions directly on a downside event, that the market return is less than its unconditional mean. In Figure 1, our calibration shows that lower coskewness is compensated by higher expected returns. However, the Appendix describes a case where CAPM α's may increase as coskewness increases, the opposite predicted by the Taylor expansion in equation (8).

Our simple example shows how compensation to downside risk may arise in equilibrium. Of course, our setting, having only two assets, is simplistic. Nevertheless, the model provides motivation to ask if downside risk demands compensation in the cross-section of US stocks, and if such compensation is different in nature from compensation for risk based on measures of higher moments (such as the Harvey-Siddique, 2000, coskewness measure). As our model shows, the compensation for downside risk is in addition to the reward already incurred for standard, unconditional risk exposures, such as standard unconditional exposure to the market factor. In our empirical work, we investigate a premium for downside risk also controlling for other known cross-sectional effects such as the size and book-to-market effects explored by Fama and French (1992 and 1993), the liquidity effect of Pástor and Stambaugh (2003), and the momentum effect of Jegadeesh and Titman (1993).

3 Downside Risk and Realized Returns

In this section, we document that stocks that covary strongly with the market, conditional on down moves of the market, have high average returns over the same period. We document this phenomenon by looking at patterns of realized returns for portfolios sorted on downside risk in Section 3.1. In Section 3.2, we examine the reward to downside risk controlling for other cross-sectional effects by using Fama-MacBeth (1973) regressions. Section 3.3 conducts robustness tests. We disentangle the different effects of coskewness risk and downside beta exposure in Section 3.4.

In our empirical work, we concentrate on presenting the results of equal-weighted portfolios. While a relationship between factor sensitivities and returns should hold for both an average stock (equal-weighting) or an average dollar (value-weighting), we focus on computing equal-weighted portfolios because past work on examining non-linearities in the cross-section has found risk due to asymmetries to be bigger among smaller stocks. For example, the coskewness effect of Harvey and Siddique (2000) is strongest for equal-weighted portfolios.[4] We work with equally-weighted portfolios to emphasize the differences between downside risk and coskewness. Nevertheless, we also examine the robustness of our findings using value-weighted portfolios. We also concentrate only on NYSE stocks to minimize the illiquidity effects of small firms, but also consider all stocks on NYSE, AMEX and NASDAQ in robustness tests.

3.1 Unconditional, Downside, and Upside Betas

Research Design

If there is a cross-sectional relation between risk and return, then we should observe patterns between average realized returns and the factor loadings associated with the exposures to different sources of risk. For example, a standard unconditional CAPM implies that stocks that covary strongly with the market should have contemporaneously high average returns over the same period. In particular, the CAPM predicts an increasing relationship between realized average returns and realized factor loadings, or contemporaneous expected returns and market betas. More generally, an unconditional multi-factor model implies that we should observe patterns between average returns and sensitivities to different sources of risk over the same time period used to compute the average returns and the factor sensitivities.

Our research design follows Black, Jensen and Scholes (1972), Fama and French (1992), Jagannathan and Wang (1996), and others, and focuses on the contemporaneous relation between realized factor loadings and realized average returns. Both Black, Jensen and Scholes (1972) and Fama and French (1992) form portfolios based on pre-formation factor loadings, but perform their asset pricing tests using the post-ranking factor loadings that are computed using the returns of their constructed portfolios over the full sample. For example, Fama and French form 25 portfolios ranked on the basis of pre-formation size and market betas, and compute ex-post factor loadings for these 25 portfolios over the full sample. At each month, they assign the post-formation beta of a stock in a Fama-MacBeth (1973) cross-sectional regression to be the ex-post market factor loading of one of the 25 portfolios to which the stock belongs during that month. Our work differs from Fama and French (1992) in one important way. Rather

[4] In their paper, Harvey and Siddique (2000) state that they use value-weighted portfolios. From correspondence with Cam Harvey, the coskewness effects arise most strongly in equal-weighted, rather than value-weighted, portfolios.

than forming portfolios based on pre-formation regression criteria and then examining post-formation factor loadings, we directly sort stocks on the realized factor loadings at the end of the period and then compute realized average returns for these portfolios over the same period that actual covariation with the factors occurred. Whereas pre-formation factor loadings reflect both actual variation in factor loadings as well as measurement error effects, post-formation factor dispersion occurs almost exclusively from actual covariation of stock returns with risk factors. Hence, our approach has greater power.

A number of studies, including Fama and MacBeth (1973), Shanken (1990), Ferson and Harvey (1991), Pástor and Stambaugh (2003), among others, compute predictive betas formed using conditional information available at time t, and then examine returns over the next period. As noted by Daniel and Titman (1997), in settings where the covariance matrix is stable over time, pre-formation factor loadings are good instruments for the future expected (post-formation) factor loadings. However, work by Fama and French (1997), Ang and Chen (2003), and Lewellen and Nagel (2003) suggest that market risk exposures may be time-varying. If pre-formation betas are weak predictors of future betas, then using pre-formation betas as instruments will have low power to detect ex-post covariation between factor loadings and realized returns. We examine the relation between pre-formation estimates of factor loadings with post-formation realized factor loadings in Section 4.

Empirical Results

In Table 1, we investigate patterns between realized average returns and realized betas. While the vast majority of cross-sectional asset pricing studies use a horizon of one month, we work in intervals of twelve months, from t to $t + 12$. This is because to compute a downside beta, we need a sufficiently large number of observations to condition on periods of down markets. One month of daily data provides too short a window for obtaining reliable estimates of downside variation. Other authors, including Kothari, Shanken and Sloan (1995) also work with annual horizons.

At the end of the year, $t + 12$, we compute a stock's beta β, downside beta β^-, and upside beta β^+. The downside beta β^- is described in equation (5), while the upside β^+ takes the same form as equation (5), except we condition on movements of the market excess return above its average value:

$$\beta^+ = \frac{\text{cov}(r_i, r_m | r_m > \mu_m)}{\text{var}(r_m | r_m > \mu_m)}. \tag{9}$$

At the beginning of the year, at time t, we sort stocks into five quintiles based on their β, β^- or β^+ over the next twelve months. In the column labelled "Realized Return," Table 1 reports the average realized excess return from time t to $t+12$ in each equally-weighted quintile portfolio. The table also reports the average cross-sectional realized β, β^- or β^+ of each quintile portfolio. These average returns and betas are computed over the same 12-month period. Hence, Table 1 shows relationships between contemporaneous factor loadings and returns. Although we use a 12-month horizon, we evaluate 12-month returns at a monthly frequency. This use of overlapping information is more efficient, but induces moving average effects. To adjust for this, we report t-statistics of differences in average excess returns between quintile portfolio 5 (high betas) and quintile portfolio 1 (low betas) using 12 Newey-West (1987) lags.[5] The sample period is from July 1963 to December 2001, with our last twelve-month return period starting in January 2001. As part of our robustness checks (below), we also examine non-overlapping sample periods.

Panel A of Table 1 shows a monotonically increasing pattern between realized average returns and realized β. Quintile 1 (5) has an average excess return of 3.5% (13.9%) per annum, and the spread in average excess returns between quintile portfolios 1 and 5 is 10.4% per annum, with a corresponding difference in contemporaneous market betas of 1.36. Our results are consistent with the earliest studies testing the CAPM, like Black, Jensen and Scholes (1972), who find a reward for holding higher beta stocks. However, this evidence per se does not mean that the CAPM holds, because the CAPM predicts that no other variable other than beta should explain a firm's expected return. Nevertheless, it demonstrates that bearing high market risk is rewarded with high average returns. Panel A also reports the positive and negative components (β^- and β^+) of beta. By construction, higher β^- or higher β^+ must also mean higher unconditional β, so high average returns are accompanied by high β^-, β^+ and regular β. Note that for these portfolios sorted by realized β, the spread in realized β^- and β^+ is also similar to the spread in realized β. In the remainder of the panels in Table 1, we decompose the reward for market risk into downside and upside components.

Panel B shows that stocks with high contemporaneous β^- have high average returns. Stocks in the quintile with the lowest (highest) β^- earn 2.7% (14.5%) per annum in excess of the risk-free rate. The average difference between quintile portfolio 1 and 5 is 11.8%, which is statistically significant at the 1% level. These results are consistent with agents disliking downside risk and avoiding stocks that covary strongly when the market dips. Hence, stocks with high β^- should carry a premium in order to entice agents to hold them. A second explanation is that agents have no particular emphasis on downside risk versus upside risk, and high β^- stocks earn high returns simply because, by construction, high β^- stocks have

[5] The theoretical number of lags required to absorb all the moving average error effects is 11, but we include an additional lag for robustness.

high values of unconditional β. The average β^- spread between quintile portfolios 1 and 5 is very large (0.19 to 1.92), but sorting on β^- also produces variation in β and β^+. However, the variation in β or β^+ is not as disperse as the variation in β^-. Another possible explanation is that sorting on high contemporaneous covariance with the market mechanically produces high contemporaneous returns. However, this concern is not applicable to our downside risk measure since we are picking out precisely those observations for which stocks have already had very low returns. In Panels C and D, we demonstrate that it is the reward for downside risk alone that is behind the pattern of high β^- stocks earning high returns.

Panel C shows a smaller spread (relative to the spreads for β and β^- in Panels A and B) for average realized excess returns for stocks sorted on realized β^+. We find that low (high) β^+ stocks earn, on average, 5.7% (9.8%) per annum in excess of the risk-free rate. This pattern of high returns to high β^+ loadings is inconsistent with agents having strong preferences for upside risk. Since β^+ only measures exposure to a rising market, stocks that rise more when the market return increases should be more attractive and, on average, earn low returns. We do not observe this pattern. Instead, the increasing pattern of returns in Panel C may be due to the patterns of β or β^-, which increase from quintile portfolios 1 to 5. From the CAPM, high β implies high returns, and if agents dislike downside risk, high β^- also implies high returns.

Finally, in Panel D, we sort stocks by realized relative downside beta, defined as:

$$\text{Relative } \beta^- = (\beta^- - \beta). \tag{10}$$

Relative downside beta controls for the effect of unconditional market exposure. Panel D shows that stocks with high realized β^- have high average returns. The difference in average excess returns between portfolios 1 and 5 is 6.6% per annum and is highly significant with a robust t-statistic of 7.7. We can rule out that this pattern of returns is attributable to unconditional beta because the β loadings are flat over portfolios 1 to 5. Hence, the high realized returns from high relative β^- is produced by the exposure to downside risk, measured by high β^- loadings.

In summary, Table 1 demonstrates that downside risk is rewarded in the pattern of cross-sectional returns. Stocks with high β^- loadings earn high average returns over the same period that is not mechanically driven by high unconditional CAPM betas. In contrast, stocks that covary strongly with the market conditional on positive moves of the market do not command significant discounts.

3.2 Fama-MacBeth Regressions

While Table 1 establishes a relationship between β^- and average returns, it does not control for the effects of other known patterns in the cross-section of stock returns. A long literature from

Banz (1981) onwards has shown that various firm characteristics also have explanatory power in the cross-section. The size effect (Banz, 1981), the book-to-market effect (Basu, 1982), the momentum effect (Jegadeesh and Titman, 1991), exposure to coskewness risk (Harvey and Siddique, 2000), exposure to cokurtosis risk (Dittmar, 2002), and exposure to aggregate liquidity risk (Pástor and Stambaugh, 2003), all imply different patterns for the cross-section of expected returns. We now demonstrate that downside risk is different from all of these effects by performing a series of cross-sectional Fama and MacBeth (1973) regressions at the firm level, over the sample period from July 1963 to December 2001.

We run Fama-MacBeth regressions of excess returns on firm characteristics and realized betas with respect to various sources of risk. We use a 12-month horizon for excess returns to correspond to the contemporaneous period over which our risk measures are calculated. Since the regressions are run using a 12-month horizon, but at the overlapping monthly frequency, we compute the standard errors of the coefficients by using 12 Newey-West (1987) lags. Table 2 reports the results listed by various sets of independent variables in regressions I-VI. We also report means and standard deviations to help us gauge economic significance. We regress realized firm returns over a 12-month horizon (t to $t+12$) on realized market beta, downside beta and upside beta, (β, β^-, and β^+) computed over the same period. Hence, these regressions demonstrate a relationship between contemporaneous returns and betas. We control for log-size, book-to-market ratio, and past 12-month excess returns of the firm at the beginning of the period t. We also include realized standard deviation of the firm excess returns, coskewness (equation 6), and cokurtosis as control variables. All of these are also computed over the period from t to $t+12$. We define cokurtosis in a similar manner to coskewness in equation (6) (see Dittmar, 2002):

$$\text{cokurt} = \frac{\mathrm{E}[(r_i - \mu_i)(r_m - \mu_m)^3]}{\sqrt{\text{var}(r_i)}\text{var}(r_m)^{3/2}}, \tag{11}$$

where r_i is the firm excess return, r_m is the market excess return, μ_i is the average excess stock return and μ_m is the average market excess returns. Scott and Horvath (1980) and Dittmar (2002) argue that stocks with high cokurtosis should have low returns because investors dislike even-numbered moments. Finally, we also include the Pástor and Stambaugh (2003) historical liquidity beta at time t to proxy for liquidity exposure.

In order to avoid putting too much weight on extreme observations, we Winsorize all independent variables at the 1% and 99% levels.[6] Winsorization has been performed in cross-

[6] For example, if an observation for the firm's book-to-market ratio is extremely large and above the 99th percentile of all the firms' book-to-market ratios that month, we replace that firm's book-to-market ratio with the book-to-market ratio corresponding to the 99th percentile. The same is done for firms whose book-to-market ratios

sectional regressions by Knez and Ready (1997), among others, and ensures that extreme outliers do not drive the results. It is particularly valuable for dealing with the book-to-market ratio, because extremely large book-to-market values are sometimes observed due to low prices, particularly before a firm delists.

We begin with Regression I in Table 2 to show the familiar, standard set of cross-sectional return patterns. While market beta carries a positive coefficient, the one-factor CAPM is overwhelmingly rejected because beta is not a sufficient statistic to explain the cross-section of stock returns. Small stocks and stocks with high book-to-market ratios also cause a firm to have high average returns, confirming Fama and French (1992). The coefficient on past returns is positive (0.017) but only weakly significant. The very large and highly significant negative coefficient (-8.43) on the firm's realized volatility of excess returns confirms the anomalous finding of Ang et al. (2003), who find that stocks with high return volatilities have low average returns. Consistent with Harvey and Siddique (2000), stocks with high coskewness have low returns. Cokurtosis is only weakly priced, with a small and insignificant coefficient of 0.015.

In Regressions II-VI, we separately examine the downside and upside components of beta and show that downside risk is priced.[7] We turn first to Regression II, which reveals that downside risk and upside risk are treated asymmetrically. The coefficient on downside risk is positive (0.069) and highly significant, confirming the portfolio sorts in Table 1. The coefficient on β^+ is negative (-0.029), but lower in magnitude than the coefficient on β^-. We did not obtain a negative reward for β^+ in Table 1, Panel C, because in Table 1, we did not control for the effect of downside and upside beta separately.

Regression III shows that the reward for both downside and upside risk is robust to controlling for size, book-to-market and momentum effects. In particular, in regression III, the Fama-MacBeth coefficients for β^- and β^+ remain almost unchanged from their Regression II estimates at 0.064 and -0.025, respectively. However, once we account for coskewness risk in Regression IV, the coefficient on β^+ becomes very small (0.003) and becomes statistically insignificant, with a t-statistic of 0.22. Controlling for coskewness also brings down the coefficient on β^- from 0.064 in Regression III (without the coskewness risk) to 0.028. Nevertheless, the premium for downside risk remains positive and statistically significant. In Regression V, where we add controls for realized firm volatility and realized firm cokurtosis, the coefficient on downside risk remains consistently positive at 0.062, and also remains highly

lie below the 1%-tile of all firms' book-to-market ratios that month.

[7] By construction, β is a weighted average of β^- and β^+. If we put β and β^- on the RHS of Regressions II-VI, the coefficients on β^- are the same to three decimal places as those reported in Table 2. Similarly, if we specify β and $\beta+$ to be regressors, the coefficients on β^+ are almost unchanged.

statistically significant.

Finally, Regression VI investigates the reward for downside and upside risk controlling for the full list of firm characteristics and realized factor loadings. We lose five years of data in constructing the Pástor-Stambaugh historical liquidity betas, so this regression is run from January 1967 to December 2001. The coefficient on β^- is 0.056, with a robust t-statistic of 5.25. In contrast, the coefficient on β^+ is statistically insignificant, whereas the premium for coskewness is significantly negative, at -0.188. Importantly, while β^- and coskewness risk measure downside risk, the coefficients on both β^- and coskewness are statistically significant. This shows that the risk exposure of downside beta is different from coskewness risk.

To help interpret the economic magnitudes of the premia reported in the Fama-MacBeth regressions, the last column of Table 2 reports the time-series averages of the cross-sectional mean and standard deviation of each of the factor loadings or characteristics. The average market beta is less than one (0.88) because we are focusing on NYSE firms, which tend to be relatively large firms with low market betas. The average downside beta is 0.97, with a cross-sectional standard deviation of 0.64. This implies that for a downside risk premium of 6% per annum, a two standard deviation move across stocks in terms of β^- corresponds to a change in expected returns of 7.6% per annum. While the premium on coskewness appears much larger in magnitude, at approximately -19% per annum, coskewness is not a beta and must be carefully interpreted. A two standard deviation movement across coskewness changes expected returns by $2 \times 0.19 \times 0.18 = 6.7\%$ per annum, which is the same order as magnitude as the effect of downside risk.

The consistent message from these regressions is that the reward for downside risk β^- is always positive at approximately 6% per annum and statistically significant. High downside beta leads to high average returns, and this result is robust to controlling for other firm characteristics and risk characteristics. In particular, downside beta risk remains significantly positive in the presence of coskewness risk controls. On the other hand, the reward for upside risk (β^+) is fragile. A priori, we expect the coefficient on β^+ to be negative, but it often flips sign and is insignificant when we control for other variables in the cross-section.

3.3 Robustness

In Table 3, we subject our results to a battery of robustness checks. We report the robustness checks for realized β^- in Panel A and for realized relative β^- in Panel B. In each panel, we report average 12-month excess returns of quintile portfolios sorted by realized β^-, or realized relative β^-, over the same 12-month period. The table reports the differences in average excess returns between quintile portfolios 1 and 5 with robust t-statistics.[8]

The first column of Table 3 shows that value-weighting the portfolios preserves the large spreads in average returns for sorts by β^- and relative β^-. In Table 1, the 5-1 spread in average returns from equal-weighting the β^- portfolios is 11.8% per annum, which reduces to 7.1% per annum when the portfolios are value-weighted in Table 3. Similarly, the 5-1 spread in relative β^- portfolios in Panel B reduces from 6.6% to 4.0%. In both cases, the spreads remain highly significant at the 1% level. Thus, while the effect of value-weighting reduces the effect of downside risk, the premium associated with β^- remains both economically large and highly statistically significant.[9]

In the second column, labelled "All Stocks," we use all stocks listed on NYSE, AMEX and NASDAQ, rather than restricting ourselves to stocks listed on the NYSE. We form equal-weighted quintile portfolios at the beginning of the period, ranking on realized betas, using breakpoints calculated over NYSE stocks. Naturally, using all stocks increases the average excess returns, as many of the newly included stocks are small. The 5-1 spreads in average returns also increase substantially. For the quintile β^- (relative β^-) portfolios, the 5-1 difference increases to 15.2% (8.6%). By limiting our universe to NYSE stocks, we deliberately understate our results to avoid confounding influences of illiquidity and non-synchronous trading.[10]

One concern about the 12-month horizon of Tables 1 and 2 is that they use overlapping observations. While this is statistically efficient, we examine the effect of using non-overlapping 12-month periods in the last column of Table 3. Our 12-month periods are from the beginning of January to the end of December each calendar year. With non-overlapping samples, it is not necessary to control for the moving average errors with robust t-statistics, but we have fewer observations. Nevertheless, the results show that the point estimates of the 5-1 spreads are still statistically significant at the 1% level. Not surprisingly the point estimates remain roughly unchanged from Table 1. Hence, our premium for downsize beta is robust to value-weighting, using the whole stock universe and using non-overlapping observations.

[8] We also examined robustness of the horizon period, but do not report the results to save space. When we examine realized betas and realized returns over a 60-month horizon using monthly frequency returns, we find the same qualitative patterns as using a 12-month horizon. In particular, the difference in average returns between quintile portfolios 1 and 5 sorted by realized β^- is still statistically significant at the 5% level.

[9] If we remove the smallest quintile of stocks from our analysis, we also obtain results similar to Table 1.

[10] To further control for the influence of illiquidity, we have also repeated our exercise using control for non-synchronous trading in a manner analogous to using a Scholes-Williams (1977) correction to compute the downside betas. Although this method is ad hoc, using this correction does not change our results.

3.4 Downside Beta Risk and Coskewness Risk

The Fama-MacBeth (1973) regressions in Section 3.2 demonstrate that both downside beta and coskewness have predictive power for the cross-section. Since both β^- and coskewness capture the effect of asymmetric higher moments, we now measure the magnitude of the reward for exposure to downside beta, while explicitly controlling for the effect of coskewness. Table 4 presents the results of this exercise.

To control for the effect of coskewness, we first form five portfolios sorted on coskewness. Then, within each coskewness quintile, we sort stocks into five portfolios based on β^-. These portfolios are equally-weighted and both coskewness and β^- are computed over the same 12-month horizon for which we examine realized excess returns. After forming the 5×5 coskewness and β^- portfolios, we average the realized excess returns of each β^- quintile over the five coskewness portfolios. This characteristic control procedure creates a set of quintile β^- portfolios with near-identical levels of coskewness risk. Thus, these quintile β^- portfolios control for differences in coskewness.

Panel A of Table 4 reports average excess returns of the 25 coskewness $\times \beta^-$ portfolios. The column labelled "Average" reports the average 12-month excess returns of the β^- quintiles controlling for coskewness risk. The row labelled "High-Low" reports the differences in average returns between the first and fifth quintile β^- portfolios within each coskewness quintile. The last row reports the 5-1 quintile difference for the β^- quintiles that control for the effect of coskewness exposure. The average excess return of 7.6% per annum in the bottom right entry of Panel A is the 5-1 difference in average returns between stocks with low and high β^-, controlling for coskewness risk. This difference has a robust t-statistic of 4.16. Hence, coskewness risk cannot account for the reward for bearing downside beta risk.

In Panel A, the patterns within each coskewness quintile moving from low β^- to high β^- stocks (reading down each column) are very interesting. As coskewness increases, the differences in average excess returns due to different β^- loadings decrease. The effect is quite pronounced. In the first coskewness quintile, the difference in average returns between the low and high β^- quintiles is 14.6% per annum. The average return difference in the low and high β^- portfolios decreases to 2.1% per annum for the quintile of stocks with the highest coskewness.

The reason for this pattern is as follows. As defined in equation (6), coskewness is effectively the covariance of a stock's return with the square of the market return, or with the volatility of the market. A stock with negative coskewness tends to have low returns when market volatility is high. These are also usually, but not always, periods of low market returns. Volatility of the market treats upside and downside risk symmetrically, so both extreme upside and extreme downside movements of the market have the same volatility. Hence, the prices of stocks with large negative coskewness tend to decrease when the market falls, but the prices of these stocks may also decrease when the market rises. In contrast, downside beta concentrates only on the former effect by explicitly considering only the downside case. When coskewness is low, there is a wide spread in β^- because there is large scope for market volatility to represent both large negative and large positive changes. This explains the large spread in average returns across the β^- quintiles for stocks with low coskewness.

The small 2.1% per annum 5-1 spread for the β^- quintiles for the highest coskewness stocks is due to the highest coskewness stocks exhibiting little asymmetry in their return distributions across up and down markets. The distribution of coskewness across stocks is skewed towards the negative side and is negative on average. Across the low to high coskewness quintiles in Panel A, the average coskewness ranges from -0.41 to 0.09. Hence, the quintile of the highest coskewness stocks have little coskewness. This means that high coskewness stocks essentially do not change their behavior across periods where market returns are stable or volatile. Furthermore, the range of β^- in the highest coskewness quintile is also smaller. The small range of β^- for the highest coskewness stocks explains the low 2.1% spread for the β^- quintiles in the second last column of Panel A.

Panel B of Table 4 repeats the same exercise as Panel A, except we examine the reward for coskewness controlling for different levels of β^-. Panel B first sorts stocks on coskewness before sorting on β^-, and then averages across the β^- quintiles. This exercise examines the coskewness premium controlling for downside exposure. Controlling for β^-, the 5-1 difference in average returns for coskewness portfolios is -6.2%, which is highly statistically significant with a t-statistic of 8.2. Moreover, there are large and highly statistically significant spreads for coskewness in every β^- quintile. Coskewness is able to maintain a high range within each β^- portfolio, unlike the diminishing range for β^- within each coskewness quintile in Panel A.

In summary, both downside beta risk and coskewness risk are different. The high returns to high β^- stocks are robust to controlling for coskewness risk, and vice versa. Downside beta risk is strongest for stocks with low coskewness. Coskewness does not differentiate between large market movements on the upside or downside. For stocks with low coskewness, downside beta is better able to capture the downside risk premium associated only with market declines than an unconditional coskewness measure.

4 Predicting Downside Risk

The previous section demonstrates a strong positive relationship between stocks that exhibit high downside risk and average returns for holding such stocks over the same period. While this is the essence of the relationship implied by a risk story, a remaining question is whether downside risk is predictable. That is, can we use information about how stocks have covaried with market down movements in the past to predict downside risk, and hence returns in the future?

In this section, we investigate the predictive relationship between past downside beta and future returns. There are two key questions we are interested in. First, we examine whether downside risk loadings computed using information up to time t have forecasting power for future returns. Second, we examine if stocks sorted on past downside risk provide ongoing exposure to downside risk in the future and, thus, compensation for exposure to risk. Section 4.1 examines returns of portfolios of stocks sorted by past β^- and past coskewness. Section 4.2 delves deeper into the analysis to find stocks that have persistent exposures to downside risk and maintain a risk-return tradeoff even when using past, rather than contemporaneous, risk measures. In Section 4.3, we conclude by conducting robustness tests.

4.1 Past Downside Risk and Future Returns

Research Design

In this section, our research design follows Fama and French (1993), Jegadeesh and Titman (1993), Harvey and Siddique (2000), and Pástor and Stambaugh (2003), among others, who sort stocks into portfolios based on pre-formation factor loadings or characterstics, and then examine monthly holding period returns. Fama and French (1993) show that sorting stocks into 25 portfolios based on size and book-to-market variables produces impressive spreads in post-formation factor loadings on size and book-to-market risk factors. Whereas the previous section finds a strong ex-post relationship between post-formation downside betas and realized returns, a corresponding strong predictive relationship between past downside betas and future realized returns will hold only if past downside betas are a good instrument to predict future (post-formation) exposure to downside risk.

We form portfolios sorted on past downside risk loadings, computed using daily returns over the previous 12 months. Since we form the portfolios using only information observable at the beginning of each month, we can examine a monthly holding period return, unlike the analysis in Section 3 where we had to use an annual horizon to obtain precise estimates for the realized covariation of stock returns with market returns during market downturns. Our conditioning here relies only on past information and not on future realized factor loadings. Hence, these portfolios sorted by past downside beta are investable portfolios. We report raw average holding period returns and average holding period returns adjusted for various characteristic effects like size, book-to-market, past returns, coskewness, and liquidity following Daniel, Grinblatt, Titman and Wermers (1997).

Empirical Results

Table 5 investigates the relationship between past downside risk measures and future returns.[11] At the beginning of each month t, we sort stocks into five quintiles based on their past β^- (Panel A) or coskewness (Panel B). In the column labelled "Average Next Month Return," we report the average realized excess return over the next month from t to $t+1$. The sample period is from July 1962 to January 2001, with our first twelve-month risk measurement period ending in June 1963 for the portfolio held in July 1963. Table 5 reports the differences in the excess returns of the quintile portfolios 1 and 5 in the row labelled "High-Low." We also report the cross-sectional realized β, β^- and β^+ of each quintile portfolios. These realized betas are computed over the following 12 months. The last column of Table 5 shows the cross-sectional realized coskewness of each quintile portfolios, also computed over the next 12 months.

Panel A reports a monotonically increasing pattern in the realized betas of stocks sorted on past downside betas. That is, stocks with low (high) β^- in the past continue to have low (high) β^- going forward. The difference in the realized β^- for portfolio 1 and 5 sorted on past β^- is 0.77. Hence, past β^-, on average, seems to be a good predictor of future β^-. The 12-month autocorrelation of β^- among stocks listed on NYSE is 0.43, which implies a monthly autocorrelation of 0.93, so individual firm β^- loadings are also persistent for the average stock. Panel A also shows that the past variation of β^- produces little variation in realized coskewness.

While the ex-post β^- loadings maintain their monotonic pattern, Panel A disappointingly shows a very weak relationship between past β^- and future returns. We observe a strictly increasing pattern moving from the first quintile β^- portfolio to the fourth quintile β^- portfolio from 0.59% to 0.84% per month. However, the highest β^- quintile portfolio has much lower average excess return of 0.70% per month . The weak pattern for past β^- and future returns is similar to the observation made by Fama and French (1992) and Jagannathan and Wang (1996), who find that sorting stocks into 100 portfolios based on past β and size generates almost no

[11] We find (but do not report the results to save space) that past β^+ has no predictive ability for future returns, even excluding the most volatile stocks as in Section 4.2. This is consistent with the results that realized β^+ has no ability to explain realized returns in the previous section.

patterns in future average returns along the past β dimension.

In contrast, Panel B shows the relation between past coskewness, future returns, and future risk attributes. The portfolios of stocks sorted by past coskewness do not exhibit large ex-post variation in coskewness. The average pre-ranking spread in coskewness between portfolios 1 and 5 is 0.50, while the average post-ranking spread in realized coskewness is only 0.05. Nevertheless, we observe a strong, strictly decreasing pattern in the returns of stocks sorted by past coskewness. The portfolio of the most negative coskewness stocks (quintile 1) has an average excess return of 0.84%, whereas the portfolio of the least negative coskewness stocks (quintile 5) has an average excess return of 0.57%. This pattern is consistent with Harvey and Siddique (2000). The strong predictive pattern of past coskewness and future returns does not have a relationship with downside beta. The future β^- loadings from past coskewness are almost flat. Hence, the predictive pattern for cross-sectional returns from past coskewness is not picking up β^- exposure.

It is puzzling why there is such a strong contemporaneous pattern between β^- and average returns but such a weak relationship between past β^- and future returns, despite the fact that the pre-ranking β^- sorts produce ex-post variation in β^-. We now investigate why the past β^- portfolios seem to fail to produce predictive return variation.

4.2 Can we Predict Downside Risk?

Why Doesn't Past Downside Risk Predict Future Returns?

The reason that past downside beta provides exposure to downside risk over the following period, but that past downside beta does not seem to provide investors with compensation for bearing this risk, is due to a strong pattern of decreasing future average returns with increasing stock volatility characteristics. Ang et al. (2003) show that stocks with high volatility (both total or idiosyncratic volatility) have extremely low future returns. High volatility stocks tend to be high beta stocks because, holding correlation between the market and the stock return constant, a high individual stock volatility implies a high β (since $\beta = \rho_{im}\sigma_i/\sigma_m$). The same is also true for β^-, so high volatility stocks tend to be have high β^- loadings. The Ang et al. (2003) volatility effect confounds the predictive ability of past β^- to forecast future expected returns, making it hard to predict downside risk for stocks that are very volatile.

We illustrate the confounding volatility effect with the downside risk expected return patterns in Figure 2. In Panel A, we show the relationship between contemporaneous β^-, individual stock volatility (σ), and excess returns. We plot realized β^- on the x-axis and realized

σ on the y-axis. Realized average excess returns are shown as contour lines, with the light (dark) lines indicating low (high) returns. Each circle represents approximately 0.5% of the number of stocks listed on the NYSE. First, as we move down the y-axis, from high to low σ, the lines become darker. This is the Ang et al. (2003) volatility effect, where stocks with high volatility have low returns. Second, moving from low β^- to high β^- along the x-axis, Panel A shows that the contour lines become darker. Hence, high β^- implies contemporaneously high average returns. In fact, if we average out the σ effect, this is just a pictorial representation of the portfolio sorts on realized β^- in Table 1.

In Panel B of Figure 2, we draw a similar graph to Panel A, except we use past β^-, past individual stock volatility σ, and the contour lines represent next month's future average excess returns. Again, moving down the y-axis from high σ to low σ, the lines become progressively darker, showing the volatility effect. There is a large light area in the upper right hand corner of stocks with high past σ, high past β^-, and low returns. If we average out the volatility effect (by averaging out σ along vertical lines), the high past β^- levels, from approximately 1.8 to 2.5, must include this region of low returns. It is this volatility effect that confounds the β^- expected return pattern. Hence, stocks with high past β^- have do not have high future returns because many of the high β^- stocks have high volatilities that are negatively correlated in the cross-section with low average returns.

Panel B also contains a long diagonal line emanating from the origin labelled "$\rho^- = 1.0$," which represents the limiting case for a past downside correlation of unity.[12] We can represent lines of different downside correlation by changing the angle of the diagonal line, holding fixed the origin, from perfectly horizontal (the y-axis), representing zero downside correlation, to the diagonal $\rho^- = 1.0$ line. For example, the line labelled "$\rho^- = 0.5$" represents the case of a downside correlation of a half. As we sweep the lines through the origin from the y-axis to the $\rho^- = 1$ line, we pass through regions that become increasingly dark. Notice the lines very close to the $\rho^- = 1$ line pass through the the darkest regions, with the highest average future returns. Hence, stocks with high downside correlation have high returns, which is the effect that Ang, Chen and Xing (2002) find. Note that the downside correlation effect holds both contemporaneously (Panel A) and predictively (Panel B).

To clearly see the confounding interaction between downside beta and volatility, we rewrite equation (5) as:

$$\beta^- = \rho^- \times \frac{\sigma_i^-}{\sigma_m^-}, \tag{12}$$

[12] In this analysis, we assume that $\sigma = \sigma^-$, where $\sigma^- = \text{var}(r_i | r_m < \mu_m)$. However, the results do not change significantly when we use σ^-.

where $\rho^- = \text{corr}(r_i, r_m | r_m < \mu_m)$ is downside correlation, and $\sigma_i^- = \sqrt{\text{var}(r_i | r_m < \mu_m)}$ and $\sigma_m^- = \sqrt{\text{var}(r_m | r_m < \mu_m)}$ represent stock volatility and market volatility, both conditional on down markets, respectively. High downside beta can be produced by high downside correlation, ρ^-, or by high downside volatility, σ_i^-. But, holding constant ρ^-, stocks with high volatility, or σ_i^-, tend to have low returns, which is exactly opposite to the high β^-, high average return effect that we wish to observe. In contrast, when σ_i^- is held constant, increasing downside correlation can only increase β^-. Hence, we tend to see high average future returns for stocks with high past downside correlation. Ang, Chen and Xing (2002) report that the difference in average future returns over the next month between decile portfolios 1 and 10 with the lowest and highest past ρ^- is approximately 5% per annum.

While downside correlation is immune from the volatility effect, it is an unsatisfactory measure of risk. Any economic model giving rise to downside risk, such as the model in Section 2, has implications for both magnitude and direction through a beta factor loading. Correlation lacks any sense of magnitude. Second, while downside correlation may produce strong predictive patterns in cross-sectional returns, it is not a coherent risk measure in the sense of Artzner, Delbaen, Eber and Heath (1999). In particular, the downside correlation of a portfolio of two assets may be greater or less than the the sum of the weighted downside correlations of the two individual assets.[13] In contrast, a downside beta of a portfolio is the weighted sum of the downside beta of its constituent assets, making it appropriate for a performance evaluation (see Pedersen and Satchell, 2000).

Fortunately, Panel B of Figure 2 shows us how to find a large subset of stocks where a predictive β^- relationship holds that are not dramatically affected by the Ang et al. (2003) volatility puzzle. If we draw a horizontal line around $\sigma = 0.6$ in Panel B and only consider stocks with past volatility below this cut-off, then moving from left to right along the x-axis produces increasingly darker contour lines. Hence, if we exclude stocks with the highest levels of past volatility, we can find a large subsample of stocks where past β^- predicts returns.

Not examining the predictive ability of downside risk on stocks with very high volatility has two important implications. First, we completely abstract from the Ang et al. (2003) volatility effect that we cannot resolve in this paper. Second, high stock volatility leads to much larger measurement error in the pre-formation downside betas. It is no accident that low future returns occur for stocks with high past β^- and high past volatility (the light area in the upper right hand corner of Panel B, Figure 2). Large measurement error makes the pre-formation β^- loadings

[13] Coskewness is also an incoherent risk measure according to Artzner et al. (1999), because coskewness coefficients defined by equation (6) are not additive. However, the coskewness statistic can be made additive by changing the normalization in the denominator to $\text{var}(r_m)^{3/2}$, rather than using $\sqrt{\text{var}(r_i)}\text{var}(r_m)$.

a more unreliable predictor of actual, post-formation, downside betas. Hence, not focusing on stocks with very high volatility allows us to construct more accurate ex-ante predictors of future ex-post downside risk loadings.

Removing the Confounding Influence of the Volatility Puzzle

We examine the expected return patterns of past β^- portfolios when we exclude stocks with very high past volatility in Table 6. We first sort stocks based on their past total volatility σ, calculated using daily returns over the past year. We sort stocks into quintiles, octiles, deciles and demi-decile (5%-tile) groups according to past σ. Panel A of Table 6 excludes stocks that fall into the highest quintile, octile, decile or demi-decile of σ, and then re-sorts the remaining stocks into quintiles according to past β^-. We report the average excess returns of these β^- quintiles that exclude very volatile stocks. We also report the difference in returns between the lowest and the highest β^- portfolios (quintiles 1 and 5), as well as the difference in returns between quintile 1 and 4.

The first column in Panel A excludes stocks in the highest σ quintile and shows a clear monotonic, increasing pattern in the returns of stocks sorted by past β^-. The difference in average returns between quintile portfolios 1 and 5 is 0.34% per month, which is statistically significant with a t-statistic of 2.31. The difference between the fourth quintile portfolio and the first quintile portfolio is slightly smaller at 0.25% per month, but has approximately the same statistical significance. If we exclude fewer stocks and only exclude stocks in the highest volatility octile or decile, the return difference between highest β^- portfolio and the lowest β^- portfolio is about the same order of magnitude (roughly 0.31% per month), but the statistical significance is somewhat weaker, with p-values of 0.051 and 0.063, respectively. If we exclude only the stocks in the highest volatility demi-decile, the return difference between the lowest β^- portfolio and the high β^- portfolio becomes quite weak (0.22% per month). However, the return difference between quintiles 1 and 4 is unaffected for all the σ exclusions at 0.25% per month and is always statistically significant. This is consistent with Panel B of Figure 2, which shows that the volatility effect is strongest among stocks with the highest past β^-.

Panel B of Table 6 shows how highly volatile stocks interact with the predictive relation between past β^- and returns. For the overall sample and for each of the highest volatility groups, we report the average market capitalization, volatility, past β^-, and average returns adjusted for size and book-to-market using a characteristic control similar to Daniel et al. (1997). The quintile of stocks with the highest volatility constitutes, on average, only 3.9% of the market. Hence, by excluding the highest quintile of volatile stocks, we exclude stocks that represent only a small fraction of the market. The second row of Panel B merely repeats the volatility puzzle that high σ stocks have low average returns. We report size and book-to-market adjusted returns of stocks in the highest volatility groups, similar to Daniel et al. (1997).[14] Stocks in the highest volatility demi-decile underperform their benchmark portfolios by, on average, a very large 0.67% per month.

Stocks in the highest σ quintile have an average annualized volatility of 61%, which is significantly higher than average annualized volatility of 36% volatility for all stocks. The highest quintile σ stocks have an average past β^- of 1.44, but over the next 12 months, these stocks have an average β^- of only 1.25. In contrast, Table 1 shows that stocks in the highest realized β^- quintile have an average β^- of 1.92. Hence, high σ stocks tend to have high past β^-, but this does not imply that they continue to exhibit high β^- the following period. Moreover, the average 12-month autocorrelation of β^- for these volatile stocks is 26%, which is much lower than the 12-month autocorrelation for the whole sample, which is 44%. If we narrow our focus to stocks on the highest octile, decile or demi-decile volatility stocks, we find that these stocks tend to be even smaller in terms of market capitalization and exhibit even lower persistence of β^- across the formation period and future holding periods. Hence, stocks with very high volatility have past β^- loadings that are poor predictors of future (ex-post) β^-. This is not unexpected, since high stock volatility induces large measurement error in the past β^- loadings.

In summary, the volatility effect that stocks with very high volatility have very low returns confounds the predictive ability of past β^- to forecast future β^- exposure, and hence, future expected returns. Fortunately, these very volatile stocks constitute a small fraction of the total market capitalizations (up to only 4%), and if we focus on stocks not in the highest volatility quintile, we find a strong pattern between past β^- and future holding period returns.

4.3 Robustness Checks of Predictive Downside Risk

While we have established a predictive relation between β^- and future returns for a large proportion of the market, there remains a concern that past β^- may be proxying for size, book-to-market, momentum, coskewness, or liquidity effects and we are incorrectly attributing the predictive return patterns to downside beta. Our final Table 7 ensures that this is not the case. We focus our attention on size and book-to-market adjusted returns, but continue to focus on the first four volatility quintiles (96% of all stocks, on average).

[14] Since we use all stocks listed on NYSE, AMEX and NASDAQ in constructing size and book-to-market benchmark portfolios, the average adjusted returns of all stocks listed on NYSE does not sum up to 0.0%.

The first column of Table 7 shows the average size and book-to-market adjusted returns of portfolios sorted by past β^-. We observe a monotonically increasing relation between past β^- and next month returns. Controlling for size and book-to-market increases the average difference in returns between the lowest and highest β^- quintile to 0.44% per month from 0.34% per month in Table 6. This difference is highly statistically significant, with a t-statistic of 3.36.

In the next three columns, we control for additional past return characteristics: momentum, coskewness, and liquidity. To control for momentum, we use past 12-month returns. Our universe of stocks is the stocks in the first four volatility quintiles. We first sort stocks the stocks into quintiles based on past returns. Then, within each past return quintile, we further sort stocks into quintiles ranked on past β^-. We average the β^- quintiles across the momentum quintiles, and report book-to-market and characteristic-matched returns within each β^- quintile. A similar procedure is repeated to control for coskewness and liquidity, which is measured using the historical liquidity betas of Pástor and Stambaugh (2003). The spread in adjusted returns between β^- quintiles 1 and 5 remains significant at the 5% level, at over 0.30% per month, controlling for momentum, coskewness, and liquidity. Hence, our predictive pattern of returns for past β^- are not due to size, book-to-market, past return, coskewness, or liquidity effects.

5 Conclusion

The cross-section of stock returns reflects a premium for downside risk. Stocks that covary strongly with the market, conditional on market declines, have high average returns over the same period. This risk-return relationship is consistent with agents placing greater weight on downside risk than they place on upside gains. Agents with aversion to downside risk require a premium to hold assets that have high sensitivities to market downturns. Hence, stocks with high downside risk exposure, or high downside betas, have high average returns.

We find that the contemporaneous high average returns earned by stocks with high downside betas are not explained by a standard list of cross-sectional effects, including size and book-to-market, coskewness, liquidity risk, and past returns. The effect is also different from simple market beta. Controlling for these and other cross-sectional effects, we estimate that the cross-sectional premium for bearing downside beta risk is approximately 6% per annum. In particular, we find that the premium captured by downside beta is quite different from the coskewness effect of Harvey and Siddique (2000). Downside beta measures risk conditional only on market declines, whereas coskewness captures the covariation of a stock with both extreme upside and extreme downside movements of the market.

Past downside beta is a good predictor of future covariation with down market movements, except for stocks that are extremely volatile. For most stocks, past downside beta cross-sectionally predicts future returns. However, for stocks with very high volatility, the past downside beta predictive relationship breaks down. This is because the past downside beta of very volatile stocks contains substantial measurement error and poorly forecasts future downside exposure, and the returns of these stocks are confounded by the Ang et al. (2003) volatility effect, where stocks with very high volatility have extremely low returns. Fortunately, high volatility stocks constitute only a small fraction of the total market, so past downside beta forecasts future expected returns for the vast majority of stocks.

Appendix

A Solution of the Disappointment Aversion Asset Allocation Problem

Ang, Bekaert and Liu (2004) develop an algorithm for solving the portfolio allocation problem for DA utility that transforms the optimization problem in equation (3) into a series of standard CRRA problems under a transformed measure that involves the degree of disappointment aversion A. The simplicity of their algorithm relies crucially on the assumption of a discretized state space that is ordered by wealth. However, their set-up is only for a single risky asset. We extend their algorithm to a multiple asset case, by considering all possible combinations of the six states. This appendix outlines this numerical solution.

Epstein and Zin (1989 and 2001) show that the First Order Conditions (FOC) for equation (3) are given by:

$$\begin{aligned} \mathrm{E}\left[\frac{\partial U(W)}{\partial W} x \mathbf{1}_{\{W \le \mu_W\}}\right] + A \cdot \mathrm{E}\left[\frac{\partial U(W)}{\partial W} x \mathbf{1}_{\{W > \mu_W\}}\right] &= 0 \\ \mathrm{E}\left[\frac{\partial U(W)}{\partial W} y \mathbf{1}_{\{W \le \mu_W\}}\right] + A \cdot \mathrm{E}\left[\frac{\partial U(W)}{\partial W} y \mathbf{1}_{\{W > \mu_W\}}\right] &= 0, \end{aligned} \tag{A-1}$$

where $\mathbf{1}$ is an indicator function.

Over a discrete-state space over states (x_s, y_s) indexed by s, the definition of the certainty equivalent μ_W in equation (1) can be written as:

$$\mu_W^{1-\gamma} = \frac{1}{K}\left(\sum_{s:W_s \le \mu_W} p_s W_s^{1-\gamma} + \sum_{s:W_s > \mu_W} A\, p_s W_s^{1-\gamma}\right) \tag{A-2}$$

where wealth in state s is given by:

$$W_s = R_f + w_x x_s + w_y y_s,$$

and the FOC in equation (A-1) take the form:

$$\begin{aligned} \left[\sum_{s:W_s \le \mu_W} p_s W_s^{-\gamma} x_s\right] + A\left[\sum_{s:W_s > \mu_W} p_s W_s^{-\gamma} x_s\right] &= 0 \\ \left[\sum_{s:W_s \le \mu_W} p_s W_s^{-\gamma} y_s\right] + A\left[\sum_{s:W_s > \mu_W} p_s W_s^{-\gamma} y_s\right] &= 0 \end{aligned} \tag{A-3}$$

Ang, Bekaert and Liu (2004) note that equation (A-3) is a standard CRRA maximization problem with a changed probability measure, where the probabilities for wealth above the certainty equivalent are down-weighted. That is, defining the probabilities as:

$$\pi_i \equiv \frac{(p_1, ..., p_i, Ap_{i+1}, ..., Ap_N)'}{(p_1 + ... + p_i) + A(p_{i+1} + ... + p_N)}, \tag{A-4}$$

which allows equation (A-3) to be re-written as:

$$\sum_s \pi_s W_s^{-\gamma} x_s = 0 \quad \text{and} \quad \sum_s \pi_s W_s^{-\gamma} y_s = 0. \tag{A-5}$$

The algorithm starts with a state i, solves the standard CRRA problem with probability distribution $\{\pi_i\}$ for the optimal portfolio weights w_{xi}^* and w_{yi}^*, and then computes the certainty equivalent μ_{Wi}^*, given by:

$$\mu_{Wi}^* = \left(\sum_{s=1}^{N} (W_s^*)^{1-\gamma} \pi_{is}\right)^{\frac{1}{1-\gamma}}.$$

We must find the state i where:

$$\mu_{Wi}^* \in [R_f + w_{xi}^* x_i + w_{yi}^* y_i,\ R_f + w_{xi}^* x_{i+1} + w_{yi}^* y_{i+1}). \quad \text{(A-6)}$$

If this condition holds, then the optimal portfolio weights for x and y, w_x^* and w_y^*, have been found, so $w_x^* = w_{xi}^*$ and $w_y^* = w_{yi}^*$ and the optimal utility is given by $\mu_W^* = \mu_{Wi}^*$.

The condition (A-6) relies on ordering the states in increasing wealth. To modify this algorithm, we take all possible M orderings of the states. Then, we find state i of ordering j where the condition (A-6) is satisfied. This gives the solution to the DA asset allocation problem over the two assets x and y.

For our calibrations, we set $\gamma = 6$, $A = 0.8$ and set the gross-risk free rate to be $R_f = 1.05$. For a base-line case, we take $u_x = 0.25 + \mu$, $m_x = 0.16 + \mu$, $d_x = -0.25 + \mu$, $u_y = 0.40$ and $d_y = -0.15$. The 6 states have probabilities given by $p_1 = 0.15$, $p_2 = 0.20$, $p_3 = 0.15$, $p_4 = 0.25$ and $p_5 = 0.20$. In equilibrium, the value of $\mu = 0.0021$. This gives us equilibrium weights of $w_x^* = 0.5543$ and $w_y^* = 0.4457$. In this specification, the mean excess returns, standard deviations and betas of the two assets and the market are given by:

	mean	stdev	β
x	0.1168	0.1863	0.6944
y	0.1250	0.2750	1.3800
mkt	0.1200	0.1375	1.0000

To obtain the relations between the β, downside beta β^- and coskewness and alpha's, we alter x_d from -0.16 to -0.30. Figure 1 shows the risk-return relations for asset x.

With an alternative set of parameters, CAPM α's increase with increasing β^- but also increasing co-skewness. Here, $\gamma = 6$, $A = 0.7$, $R_f = 1.05$, $x_u = 0.50 + \mu$, $x_m = 0.20 + \mu$, $x_d = -0.30 + \mu$, $y_u = 0.35$ and $y_d = -0.7$. The probabilities are given by $p_1 = 0.10$, $p_2 = 0.20$, $p_3 = 0.20$, $p_4 = 0.20$ and $p_5 = 0.20$. If p_1 is altered between 0.08 and 0.10 and μ solved for each case to obtain equilibrium, then we have:

CAPM α	β	β^-	co-skew
0.0273	0.9567	2.9883	-0.1417
0.0317	0.8473	3.1752	-0.1087
0.0348	0.7326	3.3873	-0.0598

However, this case is unrealistic because the values of the β^-'s are extremely high.

B Data and Portfolio Construction

We use data from the Center for Research in Security Prices (CRSP) to construct portfolios of stocks sorted by various characteristics of returns. We confine our attention to ordinary common stocks listed on NYSE, AMEX and NASDAQ, omitting ADRs, REITs, closed-end funds, foreign firms and other securities which do not have a CRSP share type code of 10 or 11. We use daily and monthly returns from CRSP for the period covering July 3rd, 1962 to December 31st, 2001, including NASDAQ data which is only available after 1972. We use the one-month Treasury bill rate from Ibbotson Associates the risk-free rate and take CRSP's value-weighted returns of all stocks as the market portfolio. All our returns are expressed as continuously compounded returns. We also use book value information found on COMPUSTAT.

For every twelve month period, we construct portfolios based on measures of risk between asset i's excess return, r_{it}, and the market's excess return, r_{mt}. We exclude stocks with more than five missing observations from our analysis. We first demean returns within each period, and denote $\tilde{r}_{it}$ as the demeaned excess return of asset i and $\tilde{r}_{mt}$ as the demeaned market excess return. We obtain estimates of the regular market β, denoted $\hat{\beta}$, and the individual stock volatility σ, denoted $\hat{\sigma}$, in the usual manner as:

$$\hat{\beta} = \frac{\sum \tilde{r}_{it}\tilde{r}_{mt}}{\sum \tilde{r}_{mt}^2}, \quad \text{and} \quad \hat{\sigma} = \sqrt{\frac{1}{T}\sum \tilde{r}_{it}^2}, \quad \text{(B-7)}$$

where T is the number of trading days in one period, which we take to be 12 months. We estimate downside beta and upside beta by conditioning on the observations above or below the sample mean of the market. That is, we calculate $\hat{\beta}^-$ and $\hat{\beta}^+$ as:

$$\begin{aligned} \hat{\beta}^- &= \frac{\sum_{\{r_{mt}<\hat{\mu}_m\}}(r_{it}-\hat{\mu}_i^-)(r_{mt}-\hat{\mu}_m^-)}{\sum_{\{r_{mt}<\hat{\mu}_m\}}(r_{mt}-\hat{\mu}_m^-)^2} \\ \hat{\beta}^+ &= \frac{\sum_{\{r_{mt}>\hat{\mu}_m\}}(r_{it}-\hat{\mu}_i^+)(r_{mt}-\hat{\mu}_m^+)}{\sum_{\{r_{mt}>\hat{\mu}_m\}}(r_{mt}-\hat{\mu}_m^+)^2}, \end{aligned} \quad \text{(B-8)}$$

where $\hat{\mu}_m = \frac{1}{T}\sum r_{mt}$; $\hat{\mu}_m^-$ and $\hat{\mu}_m^+$ are the sample means of the market conditional on downside and upside periods; and $\hat{\mu}_i^-$, and $\hat{\mu}_i^+$ are the sample means of the stock return conditional on downside and upside periods:

$$\hat{\mu}_m^- = \frac{1}{T^-}\sum_{\{r_{mt}<\hat{\mu}_m\}} r_{mt} \qquad \hat{\mu}_m^+ = \frac{1}{T^+}\sum_{\{r_{mt}>\hat{\mu}_m\}} r_{mt}$$
$$\hat{\mu}_i^- = \frac{1}{T^-}\sum_{\{r_{mt}<\hat{\mu}_m\}} r_{it} \qquad \hat{\mu}_i^+ = \frac{1}{T^-}\sum_{\{r_{mt}>\hat{\mu}_m\}} r_{it}, \tag{B-9}$$

where T^- is the number of observations where $r_{mt} < \hat{\mu}_m$ and T^+ is the number of observations where $r_{mt} > \hat{\mu}_m$.

We calculate higher-order moments of stock returns using continuously compounded daily returns over each twelve month period. Coskewness and cokurtosis are estimated as:

$$\widehat{\text{coskew}} = \frac{\frac{1}{T}\sum \tilde{r}_i \tilde{r}_m^2}{\sqrt{\frac{1}{T}\sum \tilde{r}_{it}^2}\left(\frac{1}{T}\sum \tilde{r}_{mt}^2\right)} \quad \text{and} \quad \widehat{\text{cokurt}} = \frac{\frac{1}{T}\sum \tilde{r}_i \tilde{r}_m^3}{\sqrt{\frac{1}{T}\sum \tilde{r}_{it}^2}\left(\frac{1}{T}\sum \tilde{r}_{mt}^2\right)^{3/2}}. \tag{B-10}$$

We also collect for each stock, market capitalizations, book-to-market ratio and past twelve month returns at the beginning of each twelve month period.

To calculate the liquidity betas for individual stocks, at the end of each month, we identify stocks listed on NYSE, AMEX and NASDAQ with at least five years of monthly returns. For each stock, we estimate a liquidity beta, β_i^L, by running the following regression using the most recent five years of monthly data, following Pástor and Stambaugh (2003):

$$r_{i,t} = \beta_i^0 + \beta_i^L L_t + \beta_i^M r_{m,t} + \beta_i^S SMB_t + \beta_i^H HML_t + \epsilon_{i,t}, \tag{B-11}$$

where L_t is the innovation in aggregate liquidity and SMB_t and HML_t are size and book-to-market factors of Fama and French (1993). The details of the construction of L_t is in Pástor and Stambaugh (2003).

Once portfolios are formed, we calculate the returns to holding these portfolios. Over every twelve month period, we collect the cumulative returns of each stock in excess of the one-month Treasury bill rate over the period. We also collect the excess stock return over the next one month, as well as stock returns in excess of size and book-to-market matched benchmark portfolios. These size and book-to-market adjusted returns are calculated in a manner similar to Daniel et al. (1997). Each month, stocks listed on NYSE, AMEX and NASDAQ are sorted into quintiles according to their beginning of period market capitalizations based on NYSE breakpoints. Then within each of these quintiles, stocks are further sorted into quintiles according to their book-to-market ratios based on NYSE breakpoints. For each 5×5 grouping, we calculate the return on an equal-weighted portfolio consisting of all stocks that fall into that grouping. For each stock, size and book-to-market adjusted returns are defined as the return in excess of the portfolio return of the 5×5 grouping to which the stock belongs. All of these returns are calculated with an adjustment for delisting by taking the delisting return at the time the stock is delisted. If a return is missing, we take the next available return.

References

[1] Ang, A., G. Bekaert and J. Liu, 2004, "Why Stocks May Disappoint," forthcoming *Journal of Financial Economics*.

[2] Ang, A., and J. Chen, 2003, "CAPM Over the Long-Run: 1926-2001," working paper, Columbia University.

[3] Ang, A., J. Chen and Y. Xing, 2002, "Downside Correlation and Expected Stock Returns," NBER working paper 8643.

[4] Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang, 2003, "The Cross-Section of Volatility and Expected Returns," working paper, Columbia University.

[5] Artzner, P., F. Delbaen, J. M. Eber, and D. Heath, 1999, "Coherent Measures of Risk," *Mathematical Finance*, 9, 3, 203-228.

[6] Bansal, R., D. A. Hsieh, and S. Viswanathan, 1993, "A New Approach to International Arbitrage Pricing," *Journal of Finance*, 48, 1719-1747.

[7] Banz, R., 1981. "The Relation between Return and Market Value of Common Stocks," *Journal of Financial Economics*, 9, 3-18.

[8] Barberis, N. and M. Huang, 2001, "Mental Accounting, Loss Aversion, and Individual Stock Returns," *Journal of Finance*, 56, 4, 1247-1292.

[9] Basu, S., 1983, "The Relationship between Earnings Yield, Market Value, and Return for NYSE Common Stocks: Further Evidence," *Journal of Financial Economics*, 12, 129-156.

[10] Bawa, V. S. and E. B. Lindenberg, 1977, "Capital Market Equilibrium in a Mean-Lower Partial Moment Framework," *Journal of Financial Economics*, 5, 189-200.

[11] Bekaert, G., R. J. Hodrick and D. A. Marshall, 1997, "The Implications of First-Order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics*, 40, 3-39.

[12] Black, F., M. Jensen, and M. Scholes, 1972, "The Capital Asset Pricing Model: Some Empirical Tests," in M. Jensen, ed., *Studies in the Theory of Capital Markets*, Praeger, New York.

[13] Chen, J., Hong, H., and J. Stein, 2002, "Breadth of Ownership and Stock Returns," *Journal of Financial Economics*, 66, 171-205.

[14] Daniel, K., M. Grinblatt, S. Titman and R. Wermers, 1997, "Measuring Mutual Fund Performance with Characteristic Based Benchmarks," *Journal of Finance*, 52, 3, 1035-1058.

[15] Daniel, K., and S. Titman, 1997, "Evidence on the Characteristics of Cross-Sectional Variation in Stock Returns," *Journal of Finance*, 52, 1-33.

[16] Dittmar, R., 2002, "Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross-Section of Equity Returns," *Journal of Finance*, 57, 1, 369-403.

[17] Epstein, L. G. and S. E. Zin, 1990, "First-Order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics*, 26, 3, 387-407.

[18] Epstein, L. G. and S. E. Zin, 2001, "The Independence Axiom and Asset Returns," *Journal of Empirical Finance*, 8, 537-572.

[19] Fama, E. F., and K. R. French, 1992, "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, 2, 427-465.

[20] Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3-56.

[21] Fama, E. F., and K. R. French, 1997, "Industry Costs of Equity," *Journal of Financial Economics*, 43, 153-193.

[22] Fama, E. F., and J. D. MacBeth, 1973, "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 71, 607-636.

[23] Ferson, W. E., and C. R. Harvey, 1991, "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 2, 385-415.

[24] Friend, I., and R. Westerfield, 1980, "Co-Skewness and Capital Asset Pricing," *Journal of Finance*, 35, 4, 897-913.

[25] Gibbons, M. R., 1982, "Multivariate Tests of Financial Models," *Journal of Financial Economics*, 10, 3-27.

[26] Guidolin, M., and A. Timmermann, 2002, "Optimal Portfolio Choice Under Regime Switching, Skew and Kurtosis Preferences," working paper, UCSD.

[27] Gul, F., 1991, "A Theory of Disappointment Aversion," *Econometrica*, 59, 3, 667-686.

[28] Harlow, W. and R. Rao, 1989, "Asset Pricing in a Generalized Mean-Lower Partial Moment Framework: Theory and Evidence," *Journal of Financial and Quantitative Analysis*, 24, 285-311.

[29] Harvey, C. R., J. C. Liechty, M. W. Liechty, and P. Müller, 2003, "Portfolio Selection with Higher Moments," working paper, Duke University.

[30] Harvey, C. R., and A. Siddique, 1999, "Autoregressive Conditional Skewness," *Journal of Financial and Quantitative Analysis*, 34, 4, 465-477.

[31] Harvey, C. R., and A. Siddique, 2000, "Conditional Skewness in Asset Pricing Tests," *Journal of Finance*, 55, 3, 1263-1295.

[32] Isakov, D., 1999, "Is Beta Still Alive? Conclusive Evidence from the Swiss Stock Market," *European Journal of Finance*, 5, 202-212.

[33] Jagannathan, R., and Z. Wang, 1996, "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance*, 51, 3-53.

[34] Jahankhani, A., 1976, "E-V and E-S Capital Asset Pricing Models: Some Empirical Tests," *Journal of Financial and Quantitative Analysis*, 11, 4, 513-528.

[35] Jegadeesh, N., and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48, 65-91.

[36] Jondeau, E., and M. Rockinger, 2003, "How Higher Moments affect the allocation of assets," *Finance Letters*, 1, 2.

[37] Kahneman, D., and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263-291.

[38] Knez, P. J., and M. J. Ready, 1997, "On the Robustness of Size and Book-to-Market in Cross-Sectional Regressions," *Journal of Finance*, 52, 4, 1355-1382.

[39] Kothari, S. P., J. Shanken, and R. G. Sloan, 1995, "Another Look at the Cross Section of Expected Stock Returns," *Journal of Finance*, 50, 1, 185-224.

[40] Kraus, A., and R. Litzenberger, 1976, "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance*, 31, 4, 1085-1100.

[41] Kraus, A., and R. Litzenberger, 1983, "On the Distributional Conditions for a Consumption-Oriented Three Moment CAPM," *Journal of Finance*, 38, 5, 1381-1391.

[42] Kyle, A. W., and W. Xiong, 2001, "Contagion as a Wealth Effect of Financial Intermediaries," *Journal of Finance*, 56, 1401-1440.

[43] Lewellen, J. and S. Nagel, 2003, "The Conditional CAPM does not Explain Asset-Pricing Anomalies," working paper, MIT.

[44] Markowitz, H., 1959, *Portfolio Selection*. New Haven, Yale University Press.

[45] Newey, W. K., and K. D. West, 1987, "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703-8.

[46] Pástor, L., and R. F. Stambaugh, 2003, "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, 3, 642-685.

[47] Pederson, C., and S. E. Satchell, 2000, "Small Sample Analysis of Performance Measures in the Asymmetric Response Model," *Journal of Financial and Quantitative Analysis*, 35, 3, 425-450.

[48] Pettengill, G. N., S. Sundarm, and I. Mathur, 1995, "The Conditional Relation between Beta and Returns," *Journal of Financial and Quantitative Analysis*, 30, 1, 101-116.

[49] Price, K., B. Price, and T. J. Nantell, 1982, "Variance and Lower Partial Moment Measures of Systematic Risk: Some Analytical and Empirical Results," *Journal of Finance*, 37, 3, 843-855.

[50] Routledge, B. R., and S. E. Zin, 2003, "Generalized Disappointment Aversion and Asset Prices," NBER working paper 10107.

[51] Roy, A. D., 1952, "Safety First and the Holding of Assets," *Econometrica*, 20, 431-449.

[52] Rubinstein, M., 1973, "The Fundamental Theory of Parameter-Preference Security Valuation," *Journal of Financial and Quantitative Analysis*, 8, 61-69.

[53] Scholes, M., and J. T. Williams, 1977, "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5, 3, 309-27.

[54] Scott, R. C., and P. A. Horvath, 1980, "On the Direction of Preference for Moments of Higher Order than the Variance," *Journal of Finance*, 35, 4, 915-919.

[55] Shumway, Tyler, 1997, "Explaining Returns with Loss Aversion," working paper, University of Michigan.

Table 1: Returns of Stocks Sorted by Realized β

Panel A: Stocks Sorted by Realized β

Portfolio	Realized Return	β	β^-	β^+
Low β	3.52%	0.28	0.36	0.19
2	6.07%	0.59	0.67	0.51
3	7.58%	0.82	0.90	0.77
4	9.48%	1.10	1.18	1.06
High β	13.95%	1.64	1.72	1.63
High-Low	10.43%	1.36	1.36	1.44
t-stat	[4.98]			

Panel B: Stocks Sorted by Realized β^-

Portfolio	Realized Return	β	β^-	β^+
Low β^-	2.71%	0.40	0.19	0.42
2	5.62%	0.63	0.61	0.62
3	7.63%	0.83	0.89	0.79
4	10.16%	1.06	1.23	0.99
High β^-	14.49%	1.49	1.92	1.34
High-Low	11.78%	1.09	1.72	0.92
t-stat	[6.16]			

Panel C: Stocks Sorted by Realized β^+

Portfolio	Realized Return	β	β^-	β^+
Low β^+	5.73%	0.44	0.63	-0.04
2	7.42%	0.62	0.73	0.45
3	8.29%	0.82	0.90	0.76
4	9.33%	1.05	1.10	1.12
High β^+	9.83%	1.49	1.46	1.85
High-Low	4.11%	1.05	0.83	1.89
t-stat	[2.62]			

Panel D: Stocks Sorted by Realized Relative β^-

Portfolio	Realized Return	β	β^-	β^+
Low Relative β^-	4.09%	0.98	0.56	1.12
2	7.69%	0.83	0.73	0.84
3	8.53%	0.80	0.86	0.75
4	9.56%	0.84	1.08	0.72
High Relative β^-	10.73%	0.98	1.60	0.71
High-Low	6.64%	0.00	1.04	0.41
t-stat	[7.70]			

The table lists the equal-weighted average returns and risk characteristics of stocks sorted by realized betas. For each month, we calculate β, β^-, β^+ and relative β^- (given by $(\beta^- - \beta)$) with respect to the market using daily continuously compounded returns over the next 12 months for all stocks listed on the NYSE. For each risk characteristic, we rank stocks into quintiles (1–5) and form equal-weighted portfolios at the beginning of each 12 month period. The number of stocks in each portfolio varies across time from 216 to 317 stocks. The column labelled "Realized Return" reports the average return in excess of the one-month T-bill rate over the next 12 months (which is the same period as the period used to compute β, β^-, and β^+). The row labelled "High-Low" reports the difference between the returns of portfolio 5 and portfolio 1. The entry labelled "t-stat" in square brackets is the t-statistic computed using Newey-West (1987) heteroskedastic-robust standard errors with 12 lags. The columns labelled "β", "β^-," and "β^+" report the time-series and cross-sectional average of equal-weighted individual stock betas over the 12-month holding period. The sample period is from July 1963 to December 2001, with the last 12-month period from January 2001 to December 2001, and observations are at a monthly frequency.

Table 2: Fama-MacBeth Regressions

Model	I	II	III	IV	V	VI	Mean (Std Dev)
Intercept	0.300	0.044	0.054	0.046	0.246	0.257	
	[9.35]	[3.39]	[1.66]	[1.42]	[7.62]	[7.79]	
β	0.177						0.884
	[8.19]						(0.499)
β^-		0.069	0.064	0.028	0.062	0.056	0.966
		[7.17]	[7.44]	[2.68]	[6.00]	[5.25]	(0.636)
β^+		-0.029	-0.025	0.003	0.020	0.017	0.829
		[4.85]	[4.15]	[0.22]	[2.33]	[1.91]	(0.703)
Log-Size	-0.039		-0.007	-0.013	-0.034	-0.034	5.633
	[8.82]		[1.47]	[3.03]	[7.77]	[7.39]	(0.756)
Bk-Mkt	0.017		0.024	0.023	0.017	0.018	0.863
	[3.87]		[5.17]	[5.03]	[3.67]	[3.76]	(1.603)
Past Ret	0.017		0.063	0.053	0.020	0.015	0.087
	[1.91]		[6.32]	[5.40]	[2.12]	[1.50]	(0.385)
Std Dev	-8.433				-5.781	-6.459	0.364
	[10.7]				[6.41]	[7.04]	(0.174)
Coskewness	-0.229			-0.181	-0.196	-0.188	-0.161
	[10.7]			[4.31]	[5.07]	[4.59]	(0.177)
Cokurtosis	0.015				0.045	0.047	2.406
	[1.57]				[4.40]	[4.52]	(1.222)
β_L						-0.008	-0.014
						[0.93]	(0.294)

This table shows the results of Fama-MacBeth regressions of 12-month excess returns on firm characteristics and realized risk characteristics. The sample period is from July 1963 to December 2001, with the last 12-month period from January 2001 to December 2001, and observations are at a monthly frequency (451 months) for all stocks listed on the NYSE. For Regression VI, the sample period is from January 1967 to December 2001 (397 months). The number of stocks in each regression varies across time from 1080 to 1582 stocks. The t-statistics in square brackets are computed using Newey-West (1987) heteroskedastic-robust standard errors with 12 lags. The firm characteristics are log of market capitalizations ("Log-Size"), book-to-market ratios ("Bk-Mkt"), and past 12-month excess returns ("Past Ret"), all computed at the beginning of each period. The realized risk characteristics are β, β^-, β^+, standard deviations ("Std Dev"), coskewness and cokurtosis calculated over the following 12-month period using daily continuously compounded returns. We also include the Pástor-Stambaugh (2003) liquidity beta, β_L, for January 1967 to January 2001. All independent variables are Winsorized at the 1% level and at the 99% within each month. We report time-series averages of cross-sectional mean and standard deviation of each independent variable in the last column.

Table 3: Robustness Checks of Realized β^- Portfolios

Panel A: Sorts by Realized β^-

Portfolio	Value-Weighted	All Stocks	Non-Overlapping
Low	2.62%	4.14%	3.69%
2	4.23%	7.23%	5.74%
3	6.04%	8.84%	8.33%
4	9.32%	11.35%	10.68%
High	9.76%	19.37%	16.15%
High-Low	7.14%	15.24%	12.46%
t-stat	[3.30]	[5.57]	[3.51]

Panel B: Sorts by Realized Relative β^-

Portfolio	Value-Weighted	All Stocks	Non-Overlapping
Low	3.19%	5.58%	5.67%
2	6.86%	8.66%	8.78%
3	7.07%	9.41%	8.42%
4	7.56%	10.86%	9.67%
High	7.18%	14.21%	12.05%
High-Low	3.99%	8.63%	6.38%
t-stat	[3.06]	[7.02]	[3.87]

We perform robustness checks of the results in Table 1. For each month, we calculate β^- and relative β^- using daily continuously compounded returns over the next 12 months. We report the results using realized β^- in Panel A and the results using realized relative β^- in Panel B. For each risk characteristic, we rank stocks into quintiles (1–5). In the first column of each panel, we form value-weighted portfolios using stocks listed on NYSE at the beginning of each 12 month period. In the second column, we use all stocks listed on NYSE, AMEX and NASDAQ and form equal-weighted portfolios at the beginning of each period, using quintile breakpoints based on NYSE stocks. In the last column, we compute the risk characteristics using stocks listed on NYSE and form equal-weighted portfolios at the beginning of each January using non-overlapping 12-month horizon observations. We report the average return in excess of the one-month T-bill rate over the next 12 months. The row labelled "High-Low" reports the difference between the returns of portfolio 5 and portfolio 1. For the columns labelled "Value-Weighted" and "All Stocks," the sample period is from July 1963 to December 2001, with the last 12-month period from January 2001 to December 2001, and observations are at a monthly frequency. For the column labelled "Non-Overlapping," the sample period is from January 1964 to December 2001, with the last 12-month period from January 2001 to December 2001. The number of stocks in each portfolio varies across time from 216 to 317 stocks, except for "All Stocks," where it varies from 289 to 2330 stocks. The entry labelled "t-stat" in square brackets is the t-statistic computed using Newey-West (1987) heteroskedastic-robust standard errors with 12 lags, except for the column labelled "Non-Overlapping," where 1 lag is used because there are no overlapping observations.

Table 4: Returns of Stocks Sorted by Realized βs, Controlling for Coskewness

Panel A: β^- Sorts Controlling for Coskewness

	Coskewness Quintiles					
Portfolio	Low	2	3	4	High	Average
Low β^-	7.21%	5.74%	4.03%	3.40%	0.22%	4.21%
2	10.55%	8.40%	6.94%	5.59%	2.61%	6.82%
3	13.63%	11.30%	8.30%	6.08%	3.76%	8.61%
4	15.63%	12.82%	9.35%	6.74%	2.56%	9.42%
High β^-	21.84%	15.85%	11.51%	6.81%	2.32%	11.67%
High-Low	14.64%	10.11%	7.48%	3.41%	2.10%	7.55%
t-stat	[5.62]	[5.22]	[3.91]	[1.87]	[1.32]	[4.16]

Panel B: Coskewness Sorts Controlling for β^-

	β^- Quintiles					
Portfolio	Low	2	3	4	High	Average
Low coskew	4.69%	7.15%	9.30%	12.59%	17.61%	10.27%
2	4.17%	6.19%	9.61%	12.33%	18.21%	10.10%
3	2.74%	6.51%	8.68%	11.31%	16.07%	9.06%
4	1.50%	5.24%	6.68%	9.16%	12.83%	7.08%
High coskew	0.41%	2.96%	3.86%	5.37%	7.65%	4.05%
High-Low	-4.28%	-4.18%	-5.45%	-7.22%	-9.96%	-6.22%
t-stat	[4.23]	[5.64]	[7.12]	[8.09]	[7.94]	[8.17]

This table examines the relation between β^- and coskewness. For each month, we compute β^- and coskewness with respect to the market of all stocks listed on NYSE using daily continuously compounded returns over the next 12 months. In Panel A, we first rank stocks into quintiles (1–5) at the beginning of each 12 month period based on coskewness over the next 12 months. Then, we rank stocks within each first-sort quintile into additional quintiles according to β^-, computed over the next 12 months. For each 5×5 grouping, we form an equal-weighted portfolio. In Panel B, we reverse the order so that we first sort on β^-, and then on coskewness. The sample period is from July 1963 to December 2001 and the number of stocks in each portfolio varies across time from 43 to 64 stocks. We report the average return in excess of the one-month T-bill rate over the next 12 months. For the column labelled "Average", we report the average return of stocks in each second sort quintile. This controls for coskewness (β^-) in Panel A (B). The row labelled "High-Low" reports the difference between the returns of portfolio 1 and portfolio 5. The entry labelled "t-stat" in square brackets is the t-statistic computed using Newey-West (1987) heteroskedastic-robust standard errors with 12 lags.

Table 5: Returns of Stocks Sorted by Past Asymmetry Measures

Panel A: Stocks Sorted by Past β^-

Portfolio	Average Next Month Return	Realized Statistics β	β^-	β^+	coskew
Low β^-	0.59%	0.54	0.61	0.48	-0.13
2	0.71%	0.70	0.77	0.65	-0.15
3	0.77%	0.85	0.93	0.80	-0.17
4	0.84%	1.02	1.11	0.96	-0.18
High β^-	0.70%	1.31	1.41	1.25	-0.18
High-Low	0.11%	0.77	0.80	0.77	-0.05
t-stat	[0.60]				

Panel B: Stocks Sorted by Past Coskewness

Portfolio	Average Next Month Return	Realized Statistics β	β^-	β^+	coskew
Low coskew	0.84%	0.91	1.01	0.85	-0.19
2	0.82%	0.90	0.99	0.84	-0.17
3	0.76%	0.89	0.97	0.83	-0.16
4	0.60%	0.87	0.95	0.82	-0.15
High coskew	0.57%	0.85	0.91	0.81	-0.13
High-Low	-0.28%	-0.06	-0.10	-0.04	0.05
t-stat	[2.76]				

The table reports the equal-weighted average returns and risk characteristics of stocks sorted by past β^- and coskewness. For each month, we compute β^- and coskewness with respect to the market of all stocks listed on NYSE using daily continuously compounded returns over the previous 12 months. For each risk characteristic, we rank stocks into quintiles (1–5) and form equal-weighted portfolios at the beginning of each month. The sample period is from July 1962 to January 2001. The number of stocks in each portfolio varies across time from 221 to 346 stocks. The column labelled "Return" reports the average return in excess of the one-month T-bill rate over the next month. The row labelled "High-Low" reports the difference between the returns of portfolio 1 and portfolio 5. The entry labelled "t-stat" is the simple OLS t-statistic in square brackets. The columns labelled "β", "β^-," and "β^+" report the time-series averages of equal-weighted cross-sectional averages of individual stock betas over the next 12 month period. The column labelled "coskew" reports the time-series averages of equal-weighted cross-sectional averages of individual stock coskewness over the next 12 months.

Table 6: Stocks Sorted by Past β^- Excluding Most Volatile Stocks

Panel A: Average Excess Returns

	Excluding Stocks of the Highest Quintile σ	Octile σ	Decile σ	5%-tile σ	No Exclusion
1 Low β^-	0.58%	0.58%	0.57%	0.58%	0.59%
2	0.69%	0.71%	0.72%	0.72%	0.71%
3	0.82%	0.80%	0.79%	0.77%	0.77%
4	0.82%	0.82%	0.83%	0.84%	0.84%
5 High β^-	0.92%	0.89%	0.88%	0.80%	0.70%
High-Low	0.34%	0.31%	0.31%	0.22%	0.11%
t-stat	[2.31]	[1.95]	[1.86]	[1.26]	[0.60]
Q4–Low	0.25%	0.25%	0.26%	0.26%	0.25%
t-stat	[2.28]	[2.17]	[2.31]	[2.21]	[2.06]

Panel B: Selected Characteristics of σ Portfolios

	All Stocks	Stocks of the Highest Quintile σ	Octile σ	Decile σ	5%-tile σ
Market Cap	100.0%	3.9%	1.9%	1.3%	0.4%
Size/Book-to-Market Adjusted Returns	-0.08%	-0.38%	-0.50%	-0.56%	-0.67%
Annualized Volatility	35.9%	61.0%	68.4%	72.5%	85.6%
Past β^-	0.99	1.44	1.47	1.48	1.44
12-month Autocorrelation of β^-	43.5%	25.8%	21.8%	20.5%	17.3%

In Panel A, each month, we calculate individual stock volatility σ and β^- with respect to the market of all stocks listed on the NYSE using daily continuously compounded returns over the previous 12 months. We first sort stocks according to σ into quintiles, octiles, deciles and demi-deciles (5%-tiles). Then for each σ sort, we exclude the stocks that fall into the highest quintile, octile, decile or demi-decile of σ. We rank the remaining stocks into equal-weighted quintiles (1–5) according to past β^-. We report the average excess return over the next 1-month. The row labelled " High-Low" ("Q4-Low") reports the difference between the average returns of portfolio 5 (portfolio 4) and portfolio 1. We report simple t-statistics in square brackets. The number of stocks in each portfolio varies across time and groupings from 177 to 346 stocks. Panel B reports selected average characteristics of stocks in each σ group. The first column reports the characteristics over the entire sample. The other columns report the characteristics within the highest σ groups. The row labelled "Market Cap" reports the time-series averages of cumulative market capitalization represented by the stocks in each group. The other rows report the returns adjusted for size and book-to-market using a characteristic control similar to Daniel et al. (1997), the annualized past volatility (σ), past β^-, autocorrelation of β^- between the past 12-months and the following 12-months. For each characteristic, we report the time-series averages of equal-weighted cross-sectional averages. The sample period is from July 1962 to January 2001.

Table 7: Characteristic Controls on Stocks Sorted by Past β^-

	Size/Bk-Mkt	Include Additional Controls for		
	Adjusted	Momentum	Coskewness	Liquidity
1 Low β^-	-0.25%	-0.21%	-0.21%	-0.17%
2	-0.09%	-0.07%	-0.07%	-0.02%
3	0.05%	0.04%	0.07%	0.05%
4	0.07%	0.10%	0.04%	0.10%
5 High β^-	0.20%	0.12%	0.15%	0.13%
High-Low	0.44%	0.32%	0.36%	0.30%
t-stat	[3.36]	[2.71]	[2.69]	[2.15]

The table reports robustness checks of the results in Table 6. For each month, we compute individual stock volatility σ and downside β^- with respect to the market of all stocks using daily continuously compounded returns over the previous 12 months. We first sort stocks according to σ into quintiles and exclude stocks that fall within the highest σ ranking. We rank the remaining stocks into quintiles (1–5) according to past β^- and form equal-weighted portfolios at the beginning of each month. The table reports characteristic-adjusted holding period returns over the next month of the β^- quintiles that exclude stocks in the highest σ quintile. In column labelled "Size/Bk-Mkt Adjusted", we report the average returns in excess of size and book-to-market matched benchmark portfolios. In the next three columns, we include additional controls for momentum (as measured by past 12 month returns), coskewness, and historical liquidity betas, computed following Pástor and Stambaugh (2003). For each additional control, we first perform a quintile sort based on the characteristic and then on past β^- excluding the highest σ quintile of stocks. Then, we average the β^- quintiles across the characteristic quintiles, and report book-to-market and characteristic-matched returns within each β^- quintile. The number of stocks in each portfolio varies across time from 177 to 277 stocks. The row labelled "High-Low" reports the difference between the returns of portfolio 5 and portfolio 1. The entry labelled "t-stat" in square brackets is the simple t-statistic. The sample period is from July 1962 to January 2001, except in the last column where the sample period is from January 1967 to January 2001.

Figure 1: Risk-Return Relations in a Disappointment Aversion Cross-Sectional Equilibrium

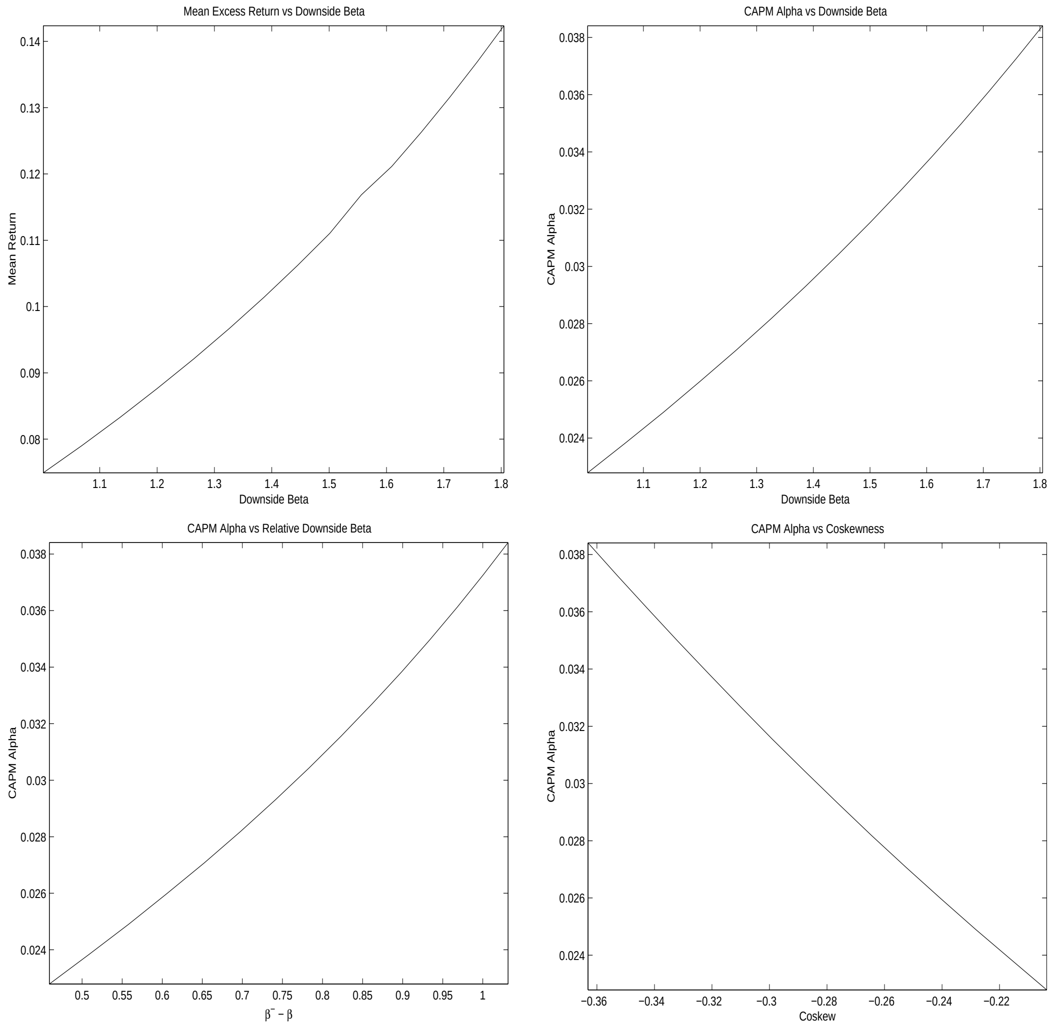

This figure shows risk-return relations for an asset in the DA cross-sectional equilibrium. Clockwise from the top left-hand corner we have: (i) a plot of the asset's mean excess return versus downside beta β^{-}, (ii) a plot of the asset's CAPM α versus downside beta β^{-}, (iii) a plot of the asset's CAPM alpha versus relative downside beta $(\beta^{-} - \beta)$, and (iv) a plot of the asset's CAPM alpha versus coskewness.

Figure 2: Distribution of Stocks According to β^- and σ versus Returns

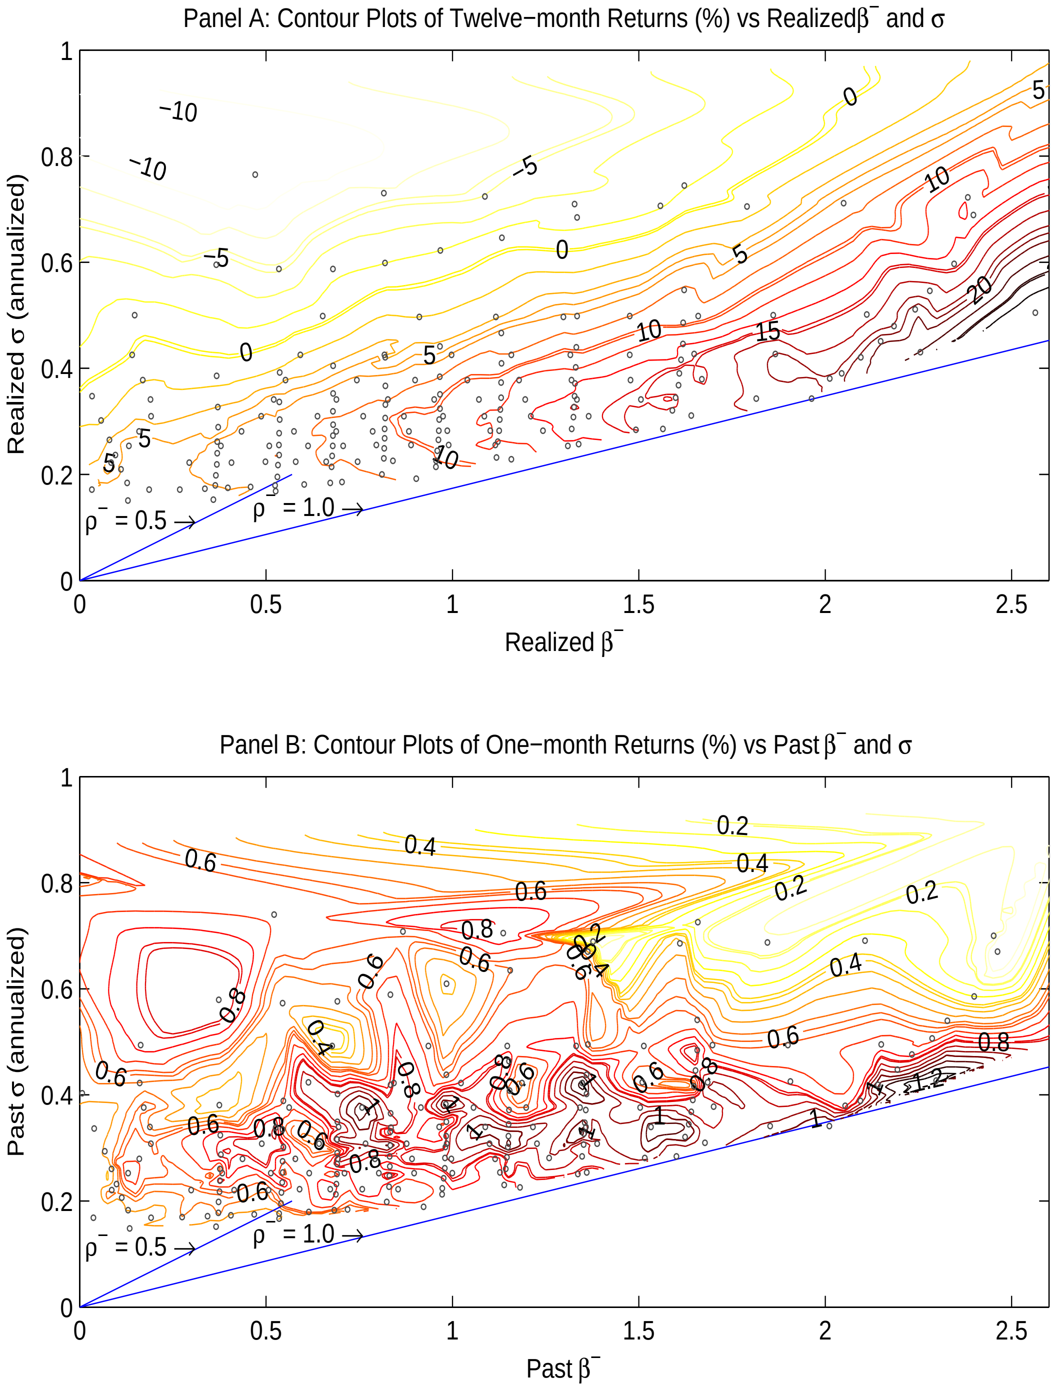

Panel A shows the average realized 12-month excess return of stocks in contour lines, according to realized β^- and realized individual stock volatility σ. The average realized returns, realized β^- and realized σ are computed using daily continuously compounded returns over the same 12 month period. Panel B shows the average future excess 1-month return, according to past β^- and past σ, computed over the previous 12-month period. Each marker, 'o', represents approximately 0.5% of the number of stocks listed on the NYSE. The contours represent low (high) average excess returns in light (dark) lines. The sample period is from July 1963 to December 2001. The y-axis numbers are annualized in both panels.