

Downside Risk

Andrew Ang

Columbia University and NBER

Joseph Chen

University of Southern California

Yuhang Xing

Rice University

Economists have long recognized that investors care differently about downside losses versus upside gains. Agents who place greater weight on downside risk demand additional compensation for holding stocks with high sensitivities to downside market movements. We show that the cross section of stock returns reflects a downside risk premium of approximately 6% per annum. Stocks that covary strongly with the market during market declines have high average returns. The reward for bearing downside risk is not simply compensation for regular market beta, nor is it explained by coskewness or liquidity risk, or by size, value, and momentum characteristics. (*JEL* C12, C15, C32, G12)

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 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, in an economy with agents placing greater emphasis on downside risk than upside gains, assets with high sensitivities to downside market movements have high

Portions of this manuscript previously circulated in an earlier article titled "Downside Correlation and Expected Stock Returns." The authors thank Enrique Arzuc Brad Barber, Geert Bekaert, Alon Brav, John Cochrane, Randy Cohen, Qiang Dai, Kent Daniel, Bob Dittmar, Rob Engle, Wayne Ferson, Eric Ghysels, John Heaton, David Hirshleifer, N. Jegadeesh, Gautam Kaul, Jonathan Lewellen, Qing Li, Terence Lim, Toby Moskowitz, Ľuboš Pástor, Adam Reed, Akhtar Siddique, Rob Stambaugh, and Zhenyu Wang. We especially thank Cam Harvey (the editor) and Bob Hodrick for detailed comments. We also thank two anonymous referees whose comments greatly improved the article. We thank seminar participants at Columbia University, Koç University, London Business School, London School of Economics, Morgan Stanley, NYU, PanAgora Asset Management, UNC, USC, the American Finance Association, the Canadian Investment Review Risk Management Conference, the Conference on Financial Market Risk Premiums at the Federal Reserve Board, the European Finance Association, the Five Star Conference, an Inquire Europe meeting, a LA Society of Financial Analysts meeting, an NBER Asset Pricing meeting, a Q-Group meeting, the Texas Finance Festival, the Valuation in Financial Markets Conference at UC Davis, and the Western Finance Association for helpful discussions. The authors acknowledge funding from a Q-Group research grant. Address correspondence to Andrew Ang, Columbia Business School, New York NY 10027 or e-mail: aa610@columbia.edu.

average returns. In this article, we show that the cross-section of stock returns reflects a premium for bearing downside risk.

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 semivariance as a measure of risk, rather than variance, because semivariance measures downside losses rather than 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 periods of high and low market returns. As Bawa and Lindenberg (1977) suggest, a natural extension of the CAPM that takes into account the asymmetric treatment of risk is to specify asymmetric downside and upside betas. We compute downside (upside) betas over periods when the excess market return is below (above) its mean. We show that stocks with high downside betas have, on average, high unconditional average returns. We also find some mixed evidence that there exists a discount for stocks that have high covariation conditional on upside movements of the market.

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 stock returns. Early researchers found little evidence of a downside risk premium because they did not focus on measuring the downside risk premium using all individual stocks in the cross section. For example, Jahankhani (1976) fails to find any improvement over the traditional CAPM by using downside betas, but his investigation uses portfolios formed from regular CAPM betas. Similarly, Harlow and Rao (1989) 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. All of these early authors do not directly estimate a downside risk premium by demonstrating that assets which covary more when the market declines have higher average returns.¹

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.

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 is 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 premia 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 (1972) and Gibbons (1982)]. More recently, Fama and French (1992) also seek, but fail to find, a relationship between postformation market betas from an unconditional CAPM and realized average stock returns over the same period. Our study differs from these earlier tests by examining a series of short one-year samples using daily data, rather than a single long sample using monthly data. This strategy provides greater statistical power in an environment where betas may be time-varying [see comments by Ang and Chen (2005) and Lewellen and Nagel (2005)].

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, 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 nonlinear fashion, whereas 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 coskewness captures 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) and Daniel and Titman (1997)], liquidity risk factor loadings [Pástor and Stambaugh (2003)], and past return characteristics [Jegadeesh and 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. We attribute this to two effects. First, the future downside covariation of very volatile stocks is difficult to predict using past betas—the

average one-year autocorrelation of one-year betas for very volatile stocks is only 17.3% compared to 43.5% for a typical stock. This is not surprising, because high volatility increases measurement error. Second, stocks with very high volatility exhibit anomalously low returns [see Ang, Hodrick, Xing, and Zhang (2006)]. Fortunately, the proportion of the market where past downside beta fails to predict future returns 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 article is organized as follows. In Section 1, we present a simple model to show how a downside risk premium may arise in a cross-sectional equilibrium. The framework uses a representative agent with the kinked disappointment aversion utility function of Gul (1991), which places larger weight on downside outcomes. Section 2 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 3, we examine the predictive ability of past downside risk loadings. Section 4 concludes.

1. 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

However, 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 aversion 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 (2001) and Hong and Stein (2003)] or wealth constraints [Kyle and Xiong (2001)].

² While standard power, or CRRA, utility also produces aversion to downside risk, the order of magnitude of a downside risk premium, relative to upside potential, is economically negligible because CRRA preferences are locally mean-variance.

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 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 and provides solvable portfolio-allocation problems, whereas optimal finite portfolio allocations for loss aversion utility may not exist [see Ang, Bekaert, and Liu (2005)]. Our example with disappointment utility differs from previous studies, because existing work with Gul's (1991) 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, 2001) and Routledge and Zin (2003)]. While a full equilibrium analysis of downside risk would entail using consumption data, in our simple example and in our empirical work, we measure aggregate wealth by the market portfolio, similar to a CAPM setting.

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

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

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:

$$K = Pr(W \leq \mu_W) + A Pr(W > \mu_W). \quad (2)$$

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.

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 setup 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 is 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), end of period wealth W is given by:

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

and 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 nontrivial, because the solution to the asset allocation problem [Equation (3)] entails simultaneously solving for both the certainty equivalent μ_W and for the portfolio weights w_x and w_y . In contrast, a standard portfolio allocation problem for CRRA utility only requires solving the first-order conditions for the optimal w_x and w_y . We extend a solution algorithm for the optimization of Equation (3) developed by Ang, Bekaert, and Liu (2005) 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 because nonparticipation may be optimal for low A under DA utility. 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 regular beta with respect to the market portfolio (denoted by 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 (denoted by 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 (denoted by $\beta^- - \beta$), relative to the regular CAPM beta, 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. In the left figure of the top row, mean excess returns increase with β^- . To make sure that β^- is not merely reflecting the regular CAPM beta, we define the CAPM alpha as the excess return of an asset not accounted for by the regular CAPM beta, $\alpha = E(r_i) - \beta E(r_m)$. The figure in the right column in the top row of Figure 1 shows that the CAPM alpha is increasing with β^- . In the second row, we also find that CAPM alpha is also increasing with relative downside beta, $(\beta^- - \beta)$. Hence, higher downside risk is remunerated by higher expected returns not captured by the CAPM.

The right-hand figure in the middle row of Figure 1 plots the CAPM alpha versus coskewness, where coskewness is defined as:

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

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:

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

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

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

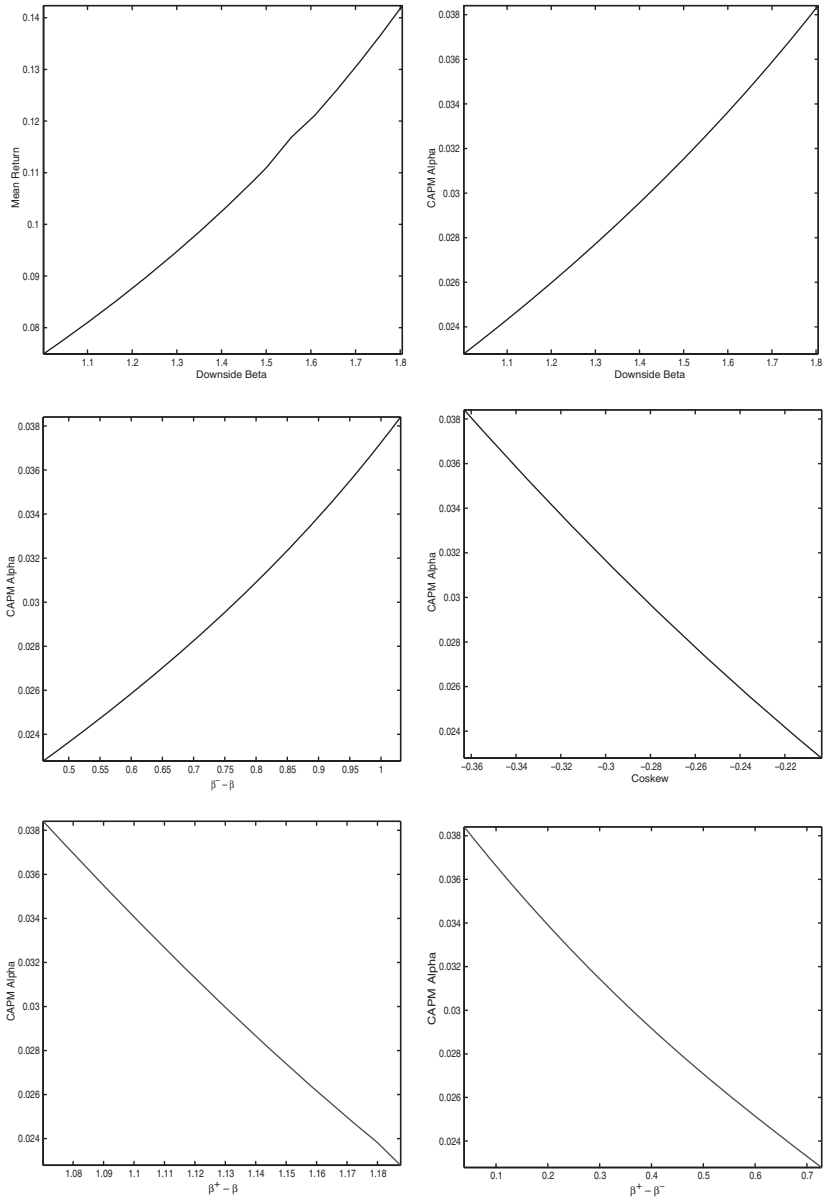


Figure 1
Risk-Return Relations in a Disappointment Aversion Cross-Sectional Equilibrium

This figure shows risk-return relations for an asset in the DA cross-sectional equilibrium. In the top row we have (i) a plot of the asset's mean excess return versus downside beta β^- ; in the second row, (ii) a plot of the asset's CAPM alpha versus downside beta, $(\beta^+ - \beta)$, and (iv) a plot of the asset's CAPM alpha versus coskewness; and in the bottom row, (v) a plot of the asset's CAPM alpha versus relative upside beta, $(\beta^+ - \beta^-)$, and (vi) a plot of the asset's CAPM alpha versus $(\beta^+ - \beta^-)$.

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.

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 nonsmooth 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 model shows that more negative coskewness is compensated by higher expected returns. However, the Appendix describes a case where CAPM alphas may increase as coskewness increases, which is the opposite of the relation predicted by the Taylor expansion.

With DA utility, a representative agent is willing to hold stocks with high upside potential at a discount, all else being equal. A stock with high upside potential relative to downside risk tends to pay off more when an investor's wealth is already high. Such stocks are not as desirable as stocks that pay off when the market decreases. Consider two stocks with the same downside beta, but with different payoffs in up markets. The stock that covaries more with the market when the market rises has a larger payoff when the market return is high. This stock does not need as high an expected return in order for the representative agent to hold it. Thus, there is a discount for stocks with high upside potential. To measure upside risk, we compute an upside beta (denoted by β^+) that takes the same form as Equation (5), except we condition on movements of the market excess return above its average value:

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

Regular, downside, and upside betas are, by construction, not independent of each other. To differentiate the effect of upside risk from downside risk, we introduce two additional measures. Similar to relative

³ Taylor expansions have been used to account for potential skewness and kurtosis preferences in asset allocation problems by Guidolin and Timmerman (2002), Harvey, Liechty, Liechty, and Müller (2003), and Jondeau and Rockinger (2006).

downside beta, we compute a relative upside beta (denoted by $\beta^+ - \beta$). We also directly examine the difference between upside beta and downside beta by computing the difference between the two, $(\beta^+ - \beta^-)$. In the last row of Figure 1, our model shows that, controlling for regular beta or downside beta, higher upside potential is indeed remunerated by lower expected returns in our model.

Our simple example illustrates one possible mechanism by which compensation to downside risk may arise in equilibrium and how downside versus upside risk may be priced differently. Of course, this example, having only two assets, is simplistic. Nevertheless, the model provides motivation to ask if downside risk demands compensation in the cross section of U.S. 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 in standard, unconditional risk exposures, such as the regular unconditional exposure to the market factor reflected in the CAPM beta. 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, 1993), the liquidity effect of Pástor and Stambaugh (2003), and the momentum effect of Jegadeesh and Titman (1993).

2. Downside Risk and Realized Returns

In this section, we document that stocks which strongly covary with the market, conditional on down moves of the market, have contemporaneously high average returns. We document this phenomenon by first looking at patterns of realized returns for portfolios sorted on downside risk in Section 2.1. Throughout, we take care in controlling for the regular beta and emphasize the asymmetry in betas by focusing on relative downside beta in addition to downside beta. In Section 2.2, we examine the reward to downside risk, controlling for other cross-sectional effects by using Fama-MacBeth (1973) regressions. We disentangle the different effects of coskewness risk and downside beta exposure in Section 2.3. Section 2.4 conducts various robustness tests. In Section 2.5, we show some additional usefulness of accounting for downside risk by examining if the commonly used Fama-French (1993) portfolios sorted by size and book-to-market characteristics exhibit exposure to downside risk.

2.1 Regular, Downside, and Upside Betas

2.1.1 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 exposure to risk. For

example, the CAPM implies that stocks that covary strongly with the market 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, a multifactor 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 MacBeth (1973), Fama and French (1992), Jagannathan and Wang (1996), and others and focuses on the contemporaneous relation between realized factor loadings and realized average returns. More recently, in testing factor models, Lettau and Ludvigson (2001), Bansal, Dittmar, and Lundblad (2005), and Lewellen and Nagel (2005), among others all employ risk measures that are measured contemporaneously with returns. While both Black, Jensen, and Scholes (1972) and Fama and French (1992) form portfolios based on preformation factor loadings, they continue to perform their asset pricing tests using postranking factor loadings, computed using the full sample. In particular, Fama and French (1992) first form 25 portfolios ranked on the basis of preformation size and market betas. Then, they compute ex post factor loadings for these 25 portfolios over the full sample. At each month, they assign the postformation beta of a stock in a Fama-MacBeth (1973) cross-sectional regression to be the ex post market factor loading of the appropriate size and book-to-market sorted portfolio to which that stock belongs during that month. Hence, testing a factor relation entails demonstrating a contemporaneous relationship between realized covariance between a stock return and a factor with the realized average return of that stock.

Our work differs from Fama and MacBeth (1973) and Fama and French (1992) in one important way. Rather than forming portfolios based on preformation regression criteria and then examining postformation factor loadings, we directly sort stocks on the realized factor loadings within a period and then compute realized average returns over the same period for these portfolios. Whereas preformation factor loadings reflect both actual variation in factor loadings and measurement error effects, postformation factor dispersion occurs almost exclusively from the actual covariation of stock returns with risk factors. Moreover, we estimate factor loadings using higher frequency data over shorter samples, rather than lower frequency data over longer samples. Hence, our approach has greater power.

A number of studies, including Fama and MacBeth (1973), Shanken (1992), Ferson and Harvey (1991), and Pástor and Stambaugh (2003),

compute predictive betas formed using conditional information available at time t and then examine returns over the next period. These studies implicitly assume that risk exposures are constant and not time-varying. Indeed, as noted by Daniel and Titman (1997), in settings where the covariance matrix is stable over time, preformation factor loadings are good instruments for the future expected (postformation) factor loadings. If preformation betas are weak predictors of future betas, then using preformation betas as instruments will also have low power to detect ex post covariation between factor loadings and realized returns. We examine the relation between preformation estimates of factor loadings with postformation realized factor loadings in Section 3.

2.1.2 Empirical results. We investigate patterns between realized average returns and realized betas. While many cross-sectional asset pricing studies use a horizon of one month, we work in intervals of 12 months, from t to $t + 12$, following Kothari, Shanken, and Sloan (1995). Our choice of an annual horizon is motivated by two concerns. First, 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. We check the robustness of our results to using intervals of 24 months with weekly frequency data to compute downside betas. Second, Fama and French (1997), Ang and Chen (2005), and Lewellen and Nagel (2005) show that market risk exposures are time-varying. Very long time intervals may cause the estimates of conditional betas to be noisy. Fama and French (2005) also advocate estimating betas using an annual horizon.

Over every 12-month period, we compute the sample counterparts to various risk measures using daily data. We calculate a stock's regular beta, downside beta as described in Equation (5), and upside beta as described in Equation (9). We also compute a stock's relative downside beta, $\beta^- - \beta$, a stock's relative upside beta, $\beta^+ - \beta$, and the difference between upside beta and downside beta, $\beta^+ - \beta^-$. Since these risk measures are calculated using realized returns, we refer to them as realized β , realized β^- , realized β^+ , realized relative β^- , realized relative β^+ , and realized $\beta^+ - \beta^-$.

In our empirical work, we concentrate on presenting the results of equal-weighted portfolios and equal-weighted Fama-MacBeth (1973) regressions. 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 nonlinearities 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. In a series of robustness checks, we also examine if our findings hold using value-weighted portfolios or in value-weighted Fama-MacBeth regressions. We concentrate only on NYSE stocks to minimize the illiquidity effects of small firms, but also consider all stocks on the NYSE, AMEX, and NASDAQ in robustness tests.

At the beginning of the one-year period at time t , we sort stocks into five quintiles based on their realized β , realized β^- , realized β^+ , realized relative β^- , realized relative β^+ , or realized $\beta^+ - \beta^-$ over the next 12 months. In the column labeled "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 1-year 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 (1987a) lags.⁵ The sample period is from July 1963 to December 2001, with our last 12-month return period starting in January 2001. As part of our robustness checks (below), we also examine nonoverlapping 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 of beta (β^- and β^+). By construction, higher β^- or higher β^+ must also mean higher unconditional β , so high average returns are accompanied by

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

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.

Table 1
Returns of Stocks Sorted by Realized Factor Loadings

Panel A: Stocks sorted by realized β					Panel B: Stocks sorted by realized β^-				
Portfolio	Return	β	β^-	β^+	Portfolio	Return	β	β^-	β^+
1 Low β	3.52%	0.28	0.36	0.19	1 Low β^-	2.71%	0.40	0.19	0.42
2	6.07%	0.59	0.67	0.51	2	5.62%	0.63	0.61	0.62
3	7.58%	0.82	0.90	0.77	3	7.63%	0.83	0.89	0.79
4	9.48%	1.10	1.18	1.06	4	10.16%	1.06	1.23	0.99
5 High β	13.95%	1.64	1.72	1.63	5 High β^-	14.49%	1.49	1.92	1.34
High-low <i>t</i> -stat	10.43% [4.98]	1.36	1.36	1.44	High-low <i>t</i> -stat	11.78% [6.16]	1.09	1.72	0.92
Panel C: Stocks sorted by realized relative β^-					Panel D: Stocks sorted by realized β^+				
Portfolio	Return	β	β^-	β^+	Portfolio	Return	β	β^-	β^+
1 Low relative β^-	4.09%	0.98	0.56	1.12	1 Low β^+	5.73%	0.44	0.63	-0.04
2	7.69%	0.83	0.73	0.84	2	7.42%	0.62	0.73	0.45
3	8.53%	0.80	0.86	0.75	3	8.29%	0.82	0.90	0.76
4	9.56%	0.84	1.08	0.72	4	9.33%	1.05	1.10	1.12
5 High relative β^-	10.73%	0.98	1.60	0.71	5 High β^+	9.83%	1.49	1.46	1.85
High-low <i>t</i> -stat	6.64% [7.70]	0.00	1.04	-0.41	High-low <i>t</i> -stat	4.11% [2.62]	1.05	0.83	1.89
Panel E: Stocks sorted by realized relative β^+					Panel F: Stocks sorted by realized $\beta^+ - \beta^-$				
Portfolio	Return	β	β^-	β^+	Portfolio	Return	β	β^-	β^+
1 Low relative β^+	10.48%	0.91	1.19	0.25	1 Low $\beta^+ - \beta^-$	11.35%	0.94	1.45	0.37
2	9.51%	0.82	0.95	0.58	2	9.83%	0.83	1.03	0.62
3	8.53%	0.81	0.88	0.76	3	8.53%	0.80	0.87	0.75
4	7.72%	0.85	0.87	0.99	4	7.35%	0.84	0.78	0.95
5 High relative β^+	4.37%	1.03	0.95	1.56	5 High $\beta^+ - \beta^-$	3.55%	1.02	0.70	1.46
High-low <i>t</i> -stat	-6.11% [9.02]	0.12	-0.23	1.31	High-low <i>t</i> -stat	-7.81% [9.03]	0.08	-0.75	1.08

This table lists the equal-weighted average returns and risk characteristics of stocks sorted by realized betas. For each month, we calculate β , β^- , β^+ , relative β^- (given by $\beta^- - \beta$), relative β^+ (given by $\beta^+ - \beta$), and $\beta^+ - \beta^-$ with respect to the market of all stocks listed on the NYSE using daily continuously compounded returns over the next 12 months. 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 labeled “Return” reports the average return in excess of the one-month Treasury-bill rate over the next 12 months (which is the same period as the period used to compute β , β^- , and β^+). The row labeled “High-low” reports the difference between the returns of portfolio 5 and portfolio 1. The entry labeled “*t*-stat” in square brackets is the *t*-statistic computed using Newey-West (1987a) heteroskedastic-robust standard errors with 12 lags for the High-low difference. The columns labeled “ 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.

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 unconditional 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 portfolios 1 and 5 is 11.8% per annum, 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, such as the DA representative agent described in Section 1. Stocks with high β^- must carry a premium in order to entice agents to hold them. An alternative explanation is that agents have no particular emphasis on downside risk versus upside potential. High β^- stocks may earn high returns simply because, by construction, high β^- stocks have high regular 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 dispersed 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 already have very low returns when the market declines. 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.

In panel C of Table 1, we sort stocks by realized relative downside beta ($\beta^- - \beta$). Relative downside beta focuses on the incremental impact of downside beta over the regular, unconditional market beta. Panel C shows that stocks with high realized relative β^- have high average returns. The difference in average excess returns between portfolios 5 and 1 is 6.6% per annum and is highly significant with a robust t -statistic of 7.70. We can rule out that this pattern of returns is attributable to regular beta because the β loadings are flat over portfolios 1 to 5. Hence, the high realized returns from high relative β^- are produced by the exposure to downside risk, measured by high β^- loadings.

Panel D shows a smaller spread for average realized excess returns for stocks sorted on realized β^+ , relative to the spreads for β and β^- in panels A and B. 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 a discount for stocks that have attractive upside exposure. 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 seems to be inconsistent with agents having strong preferences for upside potential; however, this measure does not control for the effects of regular β or for

the effects of β^- . Instead, the increasing pattern of returns in panel D may be due to the patterns of β or β^- , which increase from quintile portfolios 1 to 5. The spread in regular β is 1.05, while the spread in β^- is 0.83. From the CAPM, high β implies high returns, and if agents dislike downside risk, high β^- also implies high returns. Hence, we now turn to measures that control for these confounding effects when we look for an upside risk premium.

In panel E, we examine the effect of β^+ while controlling for regular β by sorting stocks according to realized relative upside beta, $(\beta^+ - \beta)$. Examining relative β^+ helps us to disentangle the effect of β^+ from the effects of regular β and β^- . Panel E shows that stocks with high realized relative β^+ have low returns. We find that high (low) relative β^+ stocks earn, on average, 4.4% (10.5%) per annum in excess of the risk-free rate. Furthermore, stocks sorted by relative upside beta produce a spread in β^+ while keeping the spread in regular β and β^- relatively flat. The differences in regular β and β^- across the highest quintile and the lowest quintile ($\beta^+ - \beta$) portfolios are relatively low at 0.12 and -0.23 , respectively. In contrast, we obtain a wide spread in β^+ of 1.31 between the highest quintile and the lowest quintile portfolios. This pattern of low returns to high relative β^+ stocks is consistent with agents accepting a discount for holding stocks with high upside potential, which would result from the DA agent equilibrium in Section 1.

Finally, we sort stocks by the realized difference between upside beta and downside beta ($\beta^+ - \beta^-$) in Table 1, panel F. We look at this measure to gauge the effect of upside risk relative to downside risk. In panel F, we observe a decreasing pattern in average realized excess returns with increasing $(\beta^+ - \beta^-)$. On average, we find that stocks with high (low) $(\beta^+ - \beta^-)$ earn 3.6% (11.4%) per annum in excess of the risk-free rate. While this direction is consistent with a premium for downside risk and a discount for upside potential, it is hard to separate the effects of downside risk independently from upside risk using these $(\beta^+ - \beta^-)$ portfolio sorts. The quintile portfolios sorted by $(\beta^+ - \beta^-)$ show little variation in regular β , but they show a decreasing pattern in β^- and an increasing pattern in β^+ . Hence, it is difficult to separate whether the patterns in realized returns arise because of exposure to downside losses or exposure to upside gains. Thus, our preferred measures to examine downside or upside risk are relative β^- and relative β^+ in panels C and E, which control for the effect of β .

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

discounts. However, these relations do not control for other known patterns in the cross section of stock returns, which we now investigate.

2.2 Fama-MacBeth Regressions

A long literature from Banz (1981) onward 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 (1983)], the momentum effect [Jegadeesh and Titman (1993)], the volatility effect [Ang et al. (2006)], exposure to coskewness risk [Harvey and Siddique (2000)], exposure to cokurtosis risk [Scott and Horvarth (1980) and 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 (1987a) 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 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 capture a contemporaneous relationship between average returns and factor loadings or characteristics. 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, realized coskewness, and realized 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):

$$\text{cokurt} = \frac{E[(r_i - \mu_i)(r_m - \mu_m)^3]}{\sqrt{\text{var}(r_i)\text{var}(r_m)^{3/2}}}, \quad (10)$$

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. Finally, we also include the Pástor and Stambaugh (2003) historical liquidity beta at time t to proxy for liquidity exposure.

Table 2
Fama-MacBeth Regressions

Model	I	II	III	IV	V	VI	Mean (Std Dev)
Intercept	0.300 [9.35]	0.044 [3.39]	0.054 [1.66]	0.046 [1.42]	0.246 [7.62]	0.257 [7.79]	
β	0.177 [8.19]						0.828 (0.550)
β^-		0.069 [7.17]	0.064 [7.44]	0.028 [2.68]	0.062 [6.00]	0.056 [5.25]	0.882 (0.739)
β^+		-0.029 [4.85]	-0.025 [4.15]	0.003 [0.22]	0.020 [2.33]	0.017 [1.91]	0.722 (0.842)
Log-size	-0.039 [8.82]		-0.007 [1.47]	-0.013 [3.03]	-0.034 [7.77]	-0.034 [7.39]	5.614 (1.523)
Bk-Mkt	0.017 [3.87]		0.024 [5.17]	0.023 [5.03]	0.017 [3.67]	0.018 [3.76]	0.768 (0.700)
Past Ret	0.017 [1.91]		0.063 [6.32]	0.053 [5.40]	0.020 [2.12]	0.015 [1.50]	0.085 (0.370)
Std Dev	-8.433 [10.7]				-5.781 [6.41]	-6.459 [7.04]	0.355 (0.174)
Coskewness	-0.229 [10.7]			-0.181 [4.31]	-0.196 [5.07]	-0.188 [4.59]	-0.179 (0.188)
Cokurtosis	0.015 [1.57]				0.045 [4.40]	0.047 [4.52]	2.240 (1.353)
β_L						-0.008 [0.93]	-0.166 (0.456)

This table shows the results of Fama-MacBeth (1973) 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 (1987a) 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 are all 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 the cross-sectional mean and standard deviation (in parentheses) of each independent variable in the last column.

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

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

We begin with regression I in Table 2 to show the familiar, standard set of cross-sectional return patterns. While the regular market beta carries a positive coefficient, the one-factor CAPM is rejected because β is not a sufficient statistic to explain the cross section of stock returns. The results of the regression confirm several CAPM anomalies found in the literature. For example, small stocks and stocks with high book-to-market ratios have high average returns [see Fama and French (1992)], while stocks with high past returns also continue to have high returns [see Jegadeesh and Titman, (1993)]. 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. (2006), 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. Finally, stocks with positive cokurtosis tend to have high returns, consistent with Dittmar (2002).

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 priced 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^- . These results are consistent with the positive premium on relative β^- and the discount on relative β^+ reported in panels C and E of Table 1.

Regression III shows that the reward for both downside and upside risk is robust to controlling for size, book-to-market, and momentum effects. Note that asymmetric beta risk does not remove the book-to-market or momentum effects. But, importantly, the Fama-MacBeth coefficients for β^- and β^+ remain almost unchanged from their regression II estimates at 0.064 and -0.025 , respectively. While neither β^- nor β^+ is a sufficient statistic to explain the cross section of stock returns, Table 2 demonstrates a robust reward for holding stocks with high (low) β^- (β^+) loadings controlling for the standard size, book-to-market, and past return effects.

In regressions IV–VI, accounting for additional measures of risk does not drive out the significance of β^- but drives out the significance of β^+ . 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

⁷ By construction, β is a weighted average of β^- and β^+ . If we place both β and β^- on the right-hand side of regressions II–VI and omit β^+ , the coefficients on β^- are the same to three decimal places as the β^- coefficients reported in Table 2. Similarly, if we specify both β and β^+ to be regressors, the coefficients on β^+ are almost unchanged.

risk control) to 0.028 in regression IV (with the coskewness risk control). 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 statistically significant, with a robust t -statistic of 6.00. The preference function of Section 1 that weights downside outcomes more than upside outcomes also produces a discount for upside potential. We observe these patterns in data only for regressions II and III, where the upside discount is about half the size in magnitude of the premium for downside risk. After the additional controls beyond size, book-to-market, and momentum are included in regressions IV and V, the discount for potential upside becomes fragile. Thus, the premium for downside risk dominates in the cross section.

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 . Since both β^- and coskewness risk measure downside risk, and the coefficients on both risk measures are statistically significant, we carefully disentangle the β^- and coskewness effects in Section 2.4.

To help interpret the economic magnitudes of the risk premia reported in the Fama-MacBeth regressions, the last column of Table 2 reports the time-series average of the cross-sectional mean and standard deviation of each of the factor loadings or characteristics. The average market beta is less than one (0.83) because we are focusing on NYSE firms, which tend to be skewed towards large firms with relatively low betas. The average downside beta is 0.88, with a cross-sectional standard deviation of 0.74. This implies that, for a downside risk premium of 6.9% per annum, a two-standard-deviation move across stocks in terms of β^- corresponds to a change in expected returns of $2 \times 0.069 \times 0.74 = 10.2\%$ 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 $2 \times 0.188 \times 0.19 = 7.1\%$ per annum, which is slightly less than, but of the same order of magnitude as, the effect of downside risk.

The consistent message from the regressions in Table 2 is that reward for downside risk β^- is always positive at approximately 6% per annum and statistically significant. High downside beta is compensated for by high average returns, and this result is robust to controlling for other

firm characteristics and risk characteristics, including upside beta. Moreover, 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 in data, it often flips sign and is insignificant when we control for other cross-sectional risk attributes. Thus, aversion to downside risk is priced more strongly, and more robustly, in the cross section than investors' attraction to upside potential.

2.3 Downside Beta Risk and Coskewness Risk

The Fama-MacBeth (1973) regressions in Table 2 demonstrate that both downside beta and coskewness have predictive power for the cross-section of stock returns. Since both β^- and coskewness capture the effect of asymmetric higher moments and downside risk, we now measure the magnitude of the reward for exposure to downside beta, while explicitly controlling for the effect of coskewness. Table 3 presents the results of this exercise.

To control for the effect of coskewness, we first form quintile portfolios sorted on coskewness. Then, within each coskewness quintile, we sort stocks into five equally weighted portfolios based on β^- . 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 3 reports average excess returns of the 25 coskewness $\times \beta^-$ portfolios. The column labeled "Average" reports the average 12-month excess returns of the β^- quintiles, controlling for coskewness risk. The row labeled "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 difference in average returns between the fifth and first β^- quintile portfolios that control 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

Table 3
Returns of Stocks Sorted by Realized Downside Beta and Coskewness

Panel A: β^- Sorts controlling for coskewness						
Coskewness Quintiles						
Portfolio	1 Low	2	3	4	5 High	Average
1 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%
5 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%
<i>t</i> -stat	[5.62]	[5.22]	[3.91]	[1.87]	[1.32]	[4.16]
Panel B: Coskewness sorts controlling for β^-						
β^- Quintiles						
Portfolio	1 Low	2	3	4	5 High	Average
1 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%
5 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%
<i>t</i> -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 the 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 Treasury bill rate over the next 12 months. For the column labeled “Average”, we report the average return of stocks in each second-sort quintile. This controls for coskewness (β^-) in panel A (B). The row labeled “High-low” reports the difference between the returns of portfolio 1 and portfolio 5. The entry labeled “*t*-stat” in square brackets is the *t*-statistic computed using Newey-West (1987a) heteroskedastic-robust standard errors with 12 lags for the High-low difference.

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. The distribution of coskewness across stocks is skewed towards the negative side and is negative on average. Across the low to the 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 3 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 8.17 . 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, 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 the 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.

2.4 Robustness Checks

We now show that our results do not depend on the way we have measured asymmetries in betas or the design of our empirical tests. In particular, we show that our results are robust to measuring asymmetries

with respect to different cutoff points across up markets and down markets. We also show that our results are robust to using longer-frequency data. Finally, we show that our results are not driven by using equal weighting, concentrating on NYSE stocks, or using overlapping portfolios by checking robustness with respect to value weighting, including all stocks listed on NYSE, AMEX, and NASDAQ, and using nonoverlapping portfolios.

We begin by using other cutoff points to determine up markets and down markets. Our measures of β^- and β^+ use returns relative to realized average market excess return. Naturally, realized average market returns vary across time and may have particularly low or high realizations. Alternatively, rather than using the average market excess return as the cutoff point between up markets and down markets, we can also use the risk-free rate or the zero rate of return as the cutoff point. We define downside and upside beta relative to the risk-free rate as:

$$\beta_{rf}^- = \frac{\text{cov}(r_i, r_m | r_m < r_{rf})}{\text{var}(r_m | r_m < r_{rf})} \quad \text{and} \quad \beta_{rf}^+ = \frac{\text{cov}(r_i, r_m | r_m > r_{rf})}{\text{var}(r_m | r_m > r_{rf})}. \quad (11)$$

We define downside and upside beta relative to the zero rate of return as:

$$\beta_0^- = \frac{\text{cov}(r_i, r_m | r_m < 0)}{\text{var}(r_m | r_m < 0)} \quad \text{and} \quad \beta_0^+ = \frac{\text{cov}(r_i, r_m | r_m > 0)}{\text{var}(r_m | r_m > 0)}. \quad (12)$$

We show the correlations among these risk measures in Table 4, which reports the time-series averages of the cross-sectional correlations of regular beta and the various downside and upside risk measures. Table 4 shows that β^- , β_{rf}^- , and β_0^- are all highly correlated with each other with correlations greater than 0.96. Similarly, we find that β^+ , β_{rf}^+ , and β_0^+ are also highly correlated with each other. Given these correlations, it is not surprising that reproducing Table 1 and Table 2 using either one of these alternative cutoff points yields almost identical results.⁸ Therefore, the finding of a downside risk premium is indeed being driven by emphasizing losses versus gains, rather than by using a particular cutoff point for the benchmark.

Table 4 also shows that the regular measure of beta is quite different from measures of downside beta and upside beta. The correlation of regular beta with downside beta or upside beta is 0.78 and 0.76, respectively. Therefore, downside beta and upside beta capture different aspects of risk, and are not simply reflective of the regular, unconditional market beta. Interestingly, the correlation between downside beta and upside beta is only around 0.46. Thus, a high downside risk exposure does not necessarily imply a high upside risk exposure. We examine

⁸ These results are available upon request.

Table 4
Correlations of Beta Measures

	β	β^-	β_{rf}^-	β_0^-	β^+	β_{rf}^+	β_0^+
β	1.000	0.779	0.770	0.762	0.760	0.769	0.776
β^-		1.000	0.971	0.967	0.464	0.444	0.439
β_{rf}^-			1.000	0.990	0.447	0.467	0.452
β_0^-				1.000	0.447	0.461	0.464
β^+					1.000	0.972	0.969
β_{rf}^+						1.000	0.991
β_0^+							1.000

This table reports the time-series averages of cross-sectional correlations of various risk characteristics for stocks listed on the NYSE. The risk characteristics are regular beta (β), downside beta (β^-), downside beta defined relative to the risk-free rate (β_{rf}^-), downside beta defined relative to a zero return (β_0^-), upside beta (β^+), upside beta relative to the risk-free rate (β_{rf}^+), and upside beta relative to a zero return (β_0^+). The regular downside and upside beta, β^- and β^+ , respectively, are calculated relative to the sample mean market return. All risk characteristics are computed using daily returns over the past 12 months. The sample period is from July 1963 to December 2001, with the last 12-month period from January 2001 to December 2001. There are a total of 451 observations at a monthly frequency.

further the cross-sectional determinants of future downside risk exposure in Section 3.

We now turn to additional robustness checks to make sure our findings are being driven by variation in downside risk rather than by some statistical bias introduced by our testing method. In Table 5, we subject our results to a battery of additional robustness checks. Here, we check to see if our results are robust to excluding small stocks, using longer-frequency data to compute β^- and relative β^- , creating value-weighted portfolios, using all stocks, and using nonoverlapping annual observations.⁹ 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 (or 24-month) excess returns of quintile portfolios sorted by realized β^- , or realized relative β^- , over the same period. The table also reports the differences in average excess returns between quintile portfolios 1 and 5 with robust t -statistics.

One possible worry is that our use of daily returns introduces a bias due to nonsynchronous trading. Indeed, one of the reasons for our focus on just the NYSE is to minimize these effects. To further check the influence of very small stocks, the first column of Table 5 excludes from our sample stocks that fall within the lowest size quintile. When small stocks are

⁹ In addition, we conduct further robustness checks that are available upon request. In particular, to control for the influence of nonsynchronous trading, we also repeat our exercise using control for nonsynchronous 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. We also find that the point estimates of the premiums are almost unchanged when we exclude stocks that fall into the highest volatility quintile, but statistical significance increases dramatically.

Table 5
Robustness Checks of Realized Downside Beta Portfolios

Portfolio	Exclude small	Two-year weekly	Value-weighted	All stocks	Nonoverlapping
Panel A: Sorts by realized β^-					
1 Low	4.48%	7.08%	2.62%	4.14%	3.69%
2	6.25%	12.54%	4.23%	7.23%	5.74%
3	7.69%	16.86%	6.04%	8.84%	8.33%
4	10.33%	21.34%	9.32%	11.35%	10.68%
5 High	12.81%	29.35%	9.76%	19.37%	16.15%
High-low	8.34%	22.27%	7.14%	15.24%	12.46%
<i>t</i> -stat	[4.54]	[4.96]	[3.30]	[5.57]	[3.51]
Panel B: Sorts by realized relative β^-					
1 Low	5.07%	11.46%	3.19%	5.58%	5.67%
2	8.01%	15.78%	6.86%	8.66%	8.78%
3	8.66%	18.06%	7.07%	9.41%	8.42%
4	9.44%	20.48%	7.56%	10.86%	9.67%
5 High	10.37%	21.37%	7.18%	14.21%	12.05%
High-low	5.30%	9.91%	3.99%	8.63%	6.38%
<i>t</i> -stat	[6.46]	[4.29]	[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 weekly (Wednesday to Tuesday) continuously compounded returns over each 24-month period or daily continuously compounded returns over each 12-month period. We report the results using realized β^- in panel A and the results using realized relative β^- (given by $\beta^- - \beta$) in panel B. For each risk characteristic, we rank stocks into quintiles (1–5). In the first column of each panel, we form equal-weighted portfolios among NYSE stocks, but exclude stocks that fall in the lowest size quintile. In the second column, we form equal-weighted portfolios among NYSE stocks formed by ranking on β^- or $(\beta^- - \beta)$ computed using two years of weekly data. In the third column of each panel, we form value-weighted portfolios using stocks listed on the NYSE at the beginning of each 12-month period. In the fourth column, we use all stocks listed on the 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 the NYSE and form equal-weighted portfolios at the beginning of each January using nonoverlapping 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 labeled “High-low” reports the difference between the returns of portfolio 5 and portfolio 1. For the columns labeled “Exclude small,” “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, with observations at a monthly frequency. For the column labeled “Two-year weekly,” the sample period is from July 1963 to December 2001, with the last 24-month period spanning January 2000 to December 2001. For the column labeled “Nonoverlapping,” the sample period is from January 1964 to December 2001, with the last 12-month period lasting 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 labeled “*t*-stat” in square brackets is the *t*-statistic computed using Newey-West (1987a) heteroskedastic-robust standard errors with 24 lags for the second column, 1 lag for the last column, and 12 lags for all other columns for the High-low difference.

removed, the difference between quintiles 5 and 1 for the stocks sorted by realized β^- remains strongly statistically significant (with a robust *t*-statistic of 4.54) at 8.34% per annum, but is slightly reduced from the 5–1 difference of 11.8% per annum when small stocks are included in Table 1. Similarly, the 5–1 difference in average returns for relative β^- also remains highly significant.

The second column computes β^- and $(\beta^- - \beta)$ using two years of weekly data, rather than one year of daily data. We compute weekly returns from Wednesdays to Tuesdays, and use two years of data to ensure that we have a sufficient number of observations to compute the factor loadings. We report the contemporaneous 24-month realized return over the same 24-month period used to measure the downside risk loadings. The table shows that there is no change in our basic message: there exists a reward for exposures to downside risk and relative downside risk.¹⁰

In the third column of Table 5, we examine the impact of constructing value-weighted portfolios rather than equal-weighted portfolios. Using value weighting preserves the large spreads in average excess returns for sorts by β^- and relative β^- . In particular, the 5–1 spread of value-weighted quintile portfolios in realized returns from sorting on realized β^- is 7.1% per annum. Although this has reduced from 11.8% per annum using equal-weighted portfolios in Table 1, the difference remains statistically significant at the 1% level. Similarly, the 5–1 spread in relative β^- portfolios in panel B reduces from 6.6% per annum using equal weighting to 4.0% per annum with value weighting. This difference is also significant at the 1% level.

In the fourth column, labeled “All stocks,” we use all stocks listed on the 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 based on realized beta rankings. To keep our results comparable with our earlier results, we use breakpoints calculated over just the NYSE stocks. Using all stocks increases the average excess returns, so our main results using only the NYSE universe are conservative. The 5–1 spreads in average excess returns increase substantially using all stocks. For the β^- (relative β^-) quintile portfolios, the 5–1 difference increases to 15.2% (8.6%) per annum, compared to 11.8% (6.6%) per annum using only NYSE stocks. By limiting our universe to NYSE stocks, we deliberately understate our results to avoid confounding influences of illiquidity and nonsynchronous trading.

In our last robustness check in Table 5, we use nonoverlapping observations. While the use of the overlapping 12-month horizon in Tables 1 and 2 is statistically efficient, we examine the effect of using nonoverlapping 12-month periods in the last column of Table 5. Our 12-month periods start at the beginning of January and end in December of each calendar year. With nonoverlapping samples, it is not necessary to control for the moving average errors with robust *t*-statistics,

¹⁰ We have also reproduced the Fama-MacBeth regressions using risk measures calculated at the weekly frequency and found virtually identical results to Table 2. In addition, when we examine realized betas and realized returns over a 60-month horizon using monthly frequency returns, we find the same qualitative patterns that are as statistically significant as using a 12-month horizon.

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.

In unreported results, we also conduct additional robustness checks to value weighting and using all stocks in a Fama-MacBeth regression setting. First, we run a set of value-weighted Fama-MacBeth regressions to make sure that small stocks are not driving our results. We do this by running a cross-sectional weighted least squares regression for each period, where the weights are the market capitalization of a firm at the beginning of each period. Using value-weighted regressions continues to produce a strong, statistically significant, positive relation between downside risk and contemporaneous returns with or without any additional controls. Similar to the results of using all stocks in the portfolio formations of Table 5, using all stocks in the Fama-MacBeth regressions only increases the magnitude of the downside risk premium, which remains overwhelmingly statistically significant.

2.5 Downside Risk in Size and Book-to-Market Portfolios

While we have demonstrated that exposures to high β^- or high relative β^- loadings are compensated by high average returns and this effect is consistent with investors placing greater weight on downside risk, we have not demonstrated that downside risk is useful in pricing portfolios sorted on other attributes. We now examine if portfolios of stocks sorted by other stock characteristics also exhibit contemporaneous exposure to downside risk. We focus on the Fama and French (1993) set of 25 portfolios sorted by size and book-to-market. To price these portfolios, Fama and French (1993) develop a linear asset pricing model that augments the excess market return factor, r_m , with size and value factors (SMB and HML, respectively). Table 6 examines if these portfolios exhibit exposure to downside risk, even after controlling for the standard Fama-French model.

In Table 6, we examine linear factor models of the form:

$$m = a + b_m \cdot r_m + b_{m^-} \cdot r_m^- + b_{SMB} \cdot SMB + b_{HML} \cdot HML, \quad (13)$$

where $r_m^- = \min(r_m, \mu_m)$ equals r_m if the excess market return is below its sample mean, or its sample mean otherwise. We estimate the coefficients b_m , b_{m^-} , b_{SMB} , and b_{HML} by GMM using the moment conditions:

$$E(mr) = 0, \quad (14)$$

Table 6
GMM Specification Tests

Linear factor model coefficients					
a	b_m	b_{m^-}	b_{SMB}	b_{HML}	$\Delta J \chi^2$ test
Specification I					
1.02	-5.25				
[57.6]	[3.87]				
Specification II					
1.35	-17.73	22.84			H_0 : Spec I
[8.70]	[3.03]	[2.16]			vs H_1 : Spec II
					(0.05)
Specification III					
1.08	-7.82		-0.95	-10.24	
[34.5]	[5.21]		[0.50]	[4.41]	
Specification IV					
1.60	-26.97	35.73	-3.63	-11.09	H_0 : Spec III
[7.70]	[3.42]	[2.56]	[1.52]	[4.11]	vs H_1 : Spec IV
					(0.00)

This table reports GMM estimates and GMM specification tests using the 25 Fama and French (1993) size and book-to-market sorted portfolios as base assets. Each linear factor model specification takes the form of Equation (13), with specifications I–III being special cases of Equation (13). In each case, we estimate the coefficients using the moment conditions in Equation (14). The column labeled “ ΔJ ” provides a χ^2 difference test of a restricted specification (H_0) versus an alternative specification (H_1). We report the χ^2 statistic with p -values below in parentheses. The sample period is from July 1963 to December 2001. In the table, robust t -statistics are reported in square brackets.

where r is a vector of the 25 Fama-French portfolios. The coefficient b_{m^-} reveals the exposure of the test portfolios to downside risk.¹¹ We conduct a χ^2 specification test to examine the fit with and fit without the downside risk exposure of various specifications of Equation (13). Specifically, we compute a $\Delta J \chi^2$ difference test of Newey and West (1987b) using an optimal weighting matrix of the moment conditions under the unrestricted model of the alternative hypothesis. In particular, if we reject the null hypothesis that $b_{m^-} = 0$, then we conclude that the restrictions imposed by the null hypothesis model that there is no downside risk exposure is too restrictive.

We consider two null models in Table 6, the null of the CAPM (Specification I) and the null of the Fama-French model (Specification III). In both alternatives (Specifications II and IV), Table 6 shows that the coefficient b_{m^-} is statistically significant at the 5% level. This indicates that, for pricing the size and book-to-market portfolios, the downside portion of market return plays a significant role, even in the presence of

¹¹ Note that the factor r_{m^-} is not traded, and thus the alpha of a time-series regression using the factors in Equation (13) does not represent the return of an investable strategy. Therefore, the alpha cannot be tested against the value of zero to examine possible mispricing. Similarly, we cannot compute a premium for a traded downside risk factor from Equation (13).

the standard market factor. This is true even when we allow for SMB and HML to be present in the model. This is a strong result because the SMB and HML factors are constructed specifically to explain the size and value premia of the 25 Fama-French portfolios. For both the CAPM and the Fama-French model, the ΔJ test strongly rejects both specifications in favor of allowing for downside market risk. For the Fama-French model, the p -value of the rejection is almost zero.

Thus, not only do individual stocks sorted directly on β^- loadings reveal a large reward for stocks with high downside risk exposure, but other portfolios commonly used in asset pricing also exhibit exposure to downside risk. In particular, linear factor model tests using the Fama and French (1993) size and book-to-market portfolios reject the hypothesis that these portfolios do not have exposure to downside market risk.

3. Predicting Future Downside Risk

The previous section demonstrates a strong positive relationship between stocks that exhibit high downside risk and returns for holding such stocks over the same period. While this is the essence of the relationship implied by a risk-to-reward explanation, knowing this relationship may not be of practical value if we cannot predict downside risk prior to the holding period. Therefore, we now examine if we can predict downside risk in a future period using past information. If today's information can predict future downside risk, then we can form an investable trading strategy that has exposure to downside risk. In particular, one potential valuable predictor of future downside risk exposure may be the covariation of a stock with market down movements in the past if downside risk is a persistent risk characteristic of stocks.

In Section 3.1, we explore the determinants of downside risk using past information. Section 3.2 examines returns of portfolios of stocks sorted by past downside risk, as well as sorts by past coskewness. We find that forecasting future downside risk is difficult for stocks with high volatility and we explore why in Section 3.3. Nevertheless, we can forecast future downside risk, and predict high future returns for stocks with high past downside risk for a large portion of the market. We conduct additional robustness checks in Section 3.4.

3.1 Determinants of Future Downside Risk Exposure

In this section, we begin by exploring the cross-sectional determinants of downside risk. Since we have very little theoretical guidance as to what firm characteristics determine riskiness of a stock, our investigation is merely exploratory in nature. Our analysis complements Harvey and Siddique (1999), who characterize coskewness loadings, and Ang and Chen (2002), who examine how some stock characteristics are related

with upside and downside correlations. Neither one of these studies examine how downside beta is related to firm characteristics at the individual stock level. In Table 7, we explore how relative downside betas are correlated with other cross-sectional variables. We run Fama-MacBeth (1973) regressions of realized relative downside beta on various firm characteristics that are known *ex ante* and on other future risk characteristics measured *ex post*. In particular, finding some firm characteristics that are cross-sectionally correlated with future downside beta can help us develop some possible investable strategy that generates a future spread in downside beta, and hence, future returns.

In Table 7, we first consider regressions of future realized relative downside beta over the next 12 months on variables that are estimated over the past 12 months in the columns labeled “Past variables.” All the independent variables in these regressions are measured in a period prior to the realization of relative downside beta. The regressions are run at a monthly frequency, so we use 12 Newey-West (1987a) lags. Regressions I–X use one independent variable at a time, in addition to industry dummy variables, while regression XI uses all the past variables simultaneously.¹²

Regression I shows that past relative beta does predict future relative beta over the next 12 months. However, while highly statistically significant, the coefficient is only 0.077 and far from one. Hence, future relative downside beta is difficult to predict simply by using past relative downside risk. In fact, the average 12-month autocorrelation of relative downside risk across all NYSE stocks is only 0.082, which is significantly lower than the average autocorrelation of downside beta at 0.435 and regular beta at 0.675. Thus, although past downside risk can also be used to predict future downside risk, we should expect that using only past downside risk to predict future downside risk may be difficult.

In regressions II–X, we examine how past firm characteristics are related to future relative downside beta. Past standard volatility increases relative downside beta, while smaller stocks tend to exhibit greater relative downside risk. While regression IV suggests that value stocks exhibit greater relative downside risk, controlling for additional characteristics in regression XI indicates that it is growth stocks that tend to have more relative downside risk. This is consistent with the results in Table 6 that downside risk is present in the Fama-French size and book-to-market portfolios. Stocks that are past winners also exhibit greater relative downside risk, consistent with the interpretation in Rouwenhorst (1998).

¹² Industry classifications are based on groupings of two-digit Standard Industrial Classification (SIC) codes following Ferson and Harvey (1991). They are Miscellaneous, Petroleum, Finance, Durables, Basic Industry, Food & Tobacco, Construction, Capital Goods, Transportation, Utilities, Textile & Trade, Service, and Leisure. In unreported results, we find very little pattern of downside risk exposure across industries, except that utilities generally exhibit lower exposure to downside risk than other industries. This is consistent with the notion that utilities are traditionally defensive stocks that tend to hold their value relative to other industries during market downturns.

Table 7
Determinants of Relative Downside Beta

Regression	Past variables										Realized risk measures			Industry dummy?	
	Past rel. beta	Past std dev	Log size	Bk-mkt	Past ret	ROE	Asset growth	Sales growth	Leverage	Ind [Div]	Std dev	Coskew	Cokurt		Liquidity beta
I	0.007 [10.7]														Yes
II		0.251 [6.10]													Yes
III			-0.038 [10.8]												Yes
IV				0.007 [2.04]											Yes
V					0.052 [4.85]										Yes
VI						-0.020 [3.04]									Yes
VII							-0.003 [0.43]								Yes
VIII								-0.006 [1.04]							Yes
IX									0.000 [0.13]						Yes
X											-0.066 [6.82]				Yes

Table 7
(continued)

Regression	Past variables										Realized risk measures				
	Past rel. beta	Past std dev	Log size	Bk-mkt	Past ret	ROE	Asset growth	Sales growth	Leverage	Ind [Div]	Std dev	Coskew	Cokurt	Liquidity beta	Industry dummy?
XI	0.040 [10.7]	0.094 [2.58]	-0.035 [12.8]	-0.010 [3.56]	0.063 [8.12]	-0.011 [1.10]	0.032 [2.62]	-0.001 [0.09]	0.001 [1.47]	-0.005 [1.08]					Yes
XII	0.022 [9.19]	-0.150 [6.39]	-0.042 [18.8]	-0.006 [3.08]	0.003 [0.78]	0.007 [1.21]	0.016 [2.02]	-0.009 [1.34]	0.000 [0.63]	0.004 [1.16]	0.273 [7.08]	-2.441 [23.5]	0.092 [3.21]		Yes
XIII	0.017 [5.72]	-0.145 [5.88]	-0.041 [18.5]	-0.005 [2.64]	0.001 [0.33]	0.008 [1.33]	0.011 [1.40]	-0.001 [1.46]	-0.001 [1.46]	0.004 [1.10]	0.275 [6.79]	-2.484 [22.5]	0.106 [3.48]	0.006 [1.54]	Yes

This table reports the results of Fama-MacBeth (1973) regressions of realized relative downside risk ($\beta^- - \beta$) over a 12-month period on various firm characteristics and risk measures. The independent variables include both past variables in the information set observable at time t ("Past variables") and also risk measures with contemporaneous realization as the dependent variables ("Realized risk measures"). For all specifications except regression XIII, 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 XIII, the sample period is from January 1967 to December 2001 (397 months). The t -statistics in square brackets are computed using Newey-West (1987a) heteroskedastic-robust standard errors with 12 lags. All regressions include industry dummy variables using the industry classification codes of Ferson and Harvey (1991). The past variables include realized relative downside risk beta ("Past relative β^- ") and realized volatility ("Past std dev") over the previous 12-month period. The firm characteristics are log of market capitalization ("Log-size"), the firm book-to-market ratio ("Bk-mkt"), and past 12-month excess returns ("Past ret"), all computed at the beginning of each period. We also include firm growth measured over the most recently available four-quarter period—return on equity ("ROE"), the growth rate of assets ("Asset growth"), and the growth rate of sales ("Sales growth")—as well as book leverage ("Leverage"), and an indicator which equals one if the firm pays dividends ("Ind[Div]"). The realized risk characteristics measured contemporaneously as the realized relative downside beta dependent variable are the standard deviation of excess returns ("Std dev"), coskewness, and cokurtosis. All of the realized characteristics are computed over the same 12-month period as the relative downside beta using daily continuously compounded returns. We also include the realized Pastor-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.

Regressions VI–VIII relate accounting measures of performance (return on equity, ROE; asset growth; and sales growth) with downside risk. There is strong evidence that stocks with high past ROE tend to have high future exposure to downside risk. In contrast, the evidence for asset or sales growth as predictors of future downside relative beta is weak, as either one of these variables are statistically insignificant on their own, or statistically insignificant in the full regression XI. Regression IX shows that there is little leverage effect in relative downside risk. In regression X, we use a dummy variable that equals one if the firm has paid a dividend in the past 12 months. Dividend-paying firms exhibit less relative downside risk, though this relation disappears when we include all of these variables simultaneously in regression XI.¹³

In regressions XII and XIII of Table 7, we include other contemporaneous risk measures on the right-hand side. The motivation behind these regressions is that we want to remove any possible confounding effects of other risk factors that are correlated with downside risk. We include as independent variables the stock's standard deviation, coskewness, cokurtosis, and Pástor-Stambaugh (2003) liquidity betas, all of which are measured over the same 12-month interval as the realized relative downside beta. We find that high volatility, negative coskewness, and high cokurtosis are all related to greater future relative downside risk. Note that realized high relative downside beta is contemporaneously correlated with negative coskewness and high cokurtosis, both of which are associated with average high returns [see Harvey and Siddique (2000) and Dittmar (2002)]. The earlier Tables 2 and Table 3 show that the downside beta is different from coskewness and cokurtosis risk. However, the effect of standard deviation is just the opposite—high downside beta is contemporaneously associated with high average returns, but high standard deviation is associated with lower average returns. This is consistent with Ang et al. (2006). Therefore, regression XII suggests that, in order to find an investable strategy that provides a spread in downside beta, controlling for volatility may be necessary. Finally, regression XIII shows that the Pástor-Stambaugh liquidity effect is not related to relative downside risk.

3.2 Past Downside Risk and Future Returns

The results of Section 2 suggest that downside betas are contemporaneously correlated with high average returns. However, a trading strategy to exploit this pattern of expected returns must successfully predict future downside beta loadings from past information. We now examine if stocks

¹³ In an unreported table, we do find that firms with low ROE, high growth, and high leverage have high downside beta—but, these characteristics do not add more than what is already captured by regular beta. We also find that dividend-paying firms exhibit less downside risk, without accounting for the regular beta.

simply sorted on past downside risk provide enough variation in future downside risk to provide large spreads in future returns. However, given the results of Table 7, we expect that using only past downside beta to predict future downside beta might be difficult, since past downside beta is a poor predictor of future downside beta. We confirm that this is indeed the case using a portfolio trading rule. Fortunately, the Fama-MacBeth regressions in Table 7 also suggest a way of refining the simple trading strategies, which we implement in Section 3.3.

To construct these investable portfolios, we sort stocks into portfolios at time t based on preformation characteristics and then examine monthly holding period returns from t to $t + 1$. Our main preformation criterion is downside beta computed using daily returns over the past 12 months. At the beginning of each month t , we sort stocks into five quintiles based on their past β^- and coskewness. In the column labeled “Return” in Table 8, 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 12-month risk measurement period ending in June 1963 for the portfolio formation in July 1963.

Panel A of Table 8 reports the differences in the excess returns between the highest and the lowest past β^- quintile portfolios in the row labeled “High-low.” We also report the cross-sectional realized β , β^- , and β^+ of each quintile portfolio. These realized factor loadings are computed over the following 12 months. The last column of Table 8 shows the cross-sectional realized coskewness of each quintile portfolios, also computed over the next 12 months.¹⁴

Panel A reports a monotonically increasing pattern in the realized downside 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^- between portfolios 5 and 1 sorted on past β^- is 0.80. Hence, past β^- seems to predict future β^- , but this range is far less than the range of β^- loadings for portfolios sorted on realized downside beta. In contrast, Table 1 shows that the range of realized β^- loadings is much wider at 1.72. Although past downside beta is persistent, it is far from a perfect predictor of future downside risk exposure. This is also illustrated by the 12-month autocorrelation of β^- among stocks listed on the NYSE, which is only 0.43. The difference in realized β , at 0.77, is similar to the difference in realized β^- . Panel A also shows that the past variation of β^- produces little variation in realized coskewness. Hence, the past β^- sorts are not reflecting coskewness risk.

¹⁴ We find that past β^+ has no predictive ability for future returns, even excluding the most volatile stocks as in Section 3.3. This is consistent with the results in Section 2, which fails to find a consistent contemporaneous pattern in expected returns for realized β^+ risk. These results are available on request.

Table 8
Returns of Stocks Sorted by Past Asymmetry Measures

Portfolio	Average next month return	Realized statistics			
		β	β^-	β^+	Coskew
Panel A: Stocks sorted by past β^-					
1 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
5 High β^-	0.70%	1.31	1.41	1.25	-0.18
High-low t -stat	0.11% [0.60]	0.77	0.80	0.77	-0.05
Panel B: Stocks sorted by past coskewness					
1 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
5 High coskew	0.57%	0.85	0.91	0.81	-0.13
High-low t -stat	-0.28% [2.76]	-0.06	-0.10	-0.04	0.05

This table reports the equal-weighted average returns and risk characteristics of stocks sorted by past β^- (panel A) and past coskewness (panel B). For each month, we compute β^- and coskewness with respect to the market of all stocks listed on the 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 labeled “Average next month return” reports the average return in excess of the one-month Treasury bill rate over the next month. The row labeled “High-low” reports the difference between the returns of portfolio 5 and portfolio 1. The entry labeled “ t -stat” is the simple OLS t -statistic in square brackets. The columns labeled “ 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 labeled “Coskew” reports the time-series averages of equal-weighted cross-sectional averages of individual stock coskewness over the next 12-month period of the High-low difference.

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 do 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. This difference is statistically significant with a t -statistic of 2.06. However, the highest β^- quintile portfolio has a low average excess return of 0.70% per month. Thus, the predictive relation is not as strong among stocks with the highest past β^- .

In contrast, panel B of Table 8 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 preranking spread in coskewness between

portfolios 5 and 1 is 0.50, while the average postranking 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% per month, whereas the portfolio of the least negative coskewness stocks (quintile 5) has an average excess return of 0.57% per month. 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 downside risk.

3.3 Investable Portfolios with a Spread in Downside Risk

The Fama-MacBeth regressions in Table 7 suggest that a simple trading strategy of forming portfolios on past β^- to predict future exposure to downside risk can be refined by considering subsets of portfolios focusing on volatility, book-to-market, and momentum interactions with downside risk. In this section, we consider one possible refinement of predicting future β^- using past β^- , which the Fama-MacBeth regressions of the determinants of β^- in Table 7 suggest is an important confounding influence. Specifically, we examine the role of stock return volatility and β^- and show that only for stocks with the highest volatility levels does past high β^- fail to predict future high average returns.

There are two reasons why high stock volatility causes past β^- to be a poor predictor of future β^- . First, in order for the strong contemporaneous pattern between β^- loadings and average returns to be exploited in an investable portfolio strategy, we must accurately forecast future β^- exposure. When return volatility is very high, the past β^- estimates contain substantial measurement error. Thus, high volatility causes the persistence of β^- to be lower, and thus future β^- loadings to be less predictable, because of sampling error.

Second, Ang et al. (2006) identify a puzzling anomaly that stocks with very high total or idiosyncratic volatility have extremely low returns. We suspect that this effect confounds the relationship between high β^- and high return. High-volatility stocks also 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 . To clearly see the confounding interaction between downside beta and volatility, we rewrite Equation (5) as:

$$\beta^- = \rho^- \cdot \frac{\sigma_i^-}{\sigma_m^-}, \quad (15)$$

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 and market volatilities 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.¹⁵ Therefore the Ang et al. (2006) volatility effect works in the opposite way as the expected return pattern for β^- , making it hard to predict downside risk for stocks with very high volatility.

Table 9 shows that past β^- poorly predicts future β^- only for high volatility stocks. In panel A, we report selected summary statistics of portfolios sorted by volatility, measured using daily continuously compounded returns over the previous 12 months. For the overall sample and for each of the highest volatility groups, we report the average market capitalization, past β^- , and average returns adjusted for size and book-to-market using a characteristic control similar to Daniel, Grinblatt, Titman, and Wermers (1997). While the average stock has an annualized volatility of 36%, stocks in the highest volatility quintile have a considerably higher average volatility of 61% per annum. Stocks in the highest demi-decile (5%-tile) have an average volatility of over 85% per annum.

The second row of panel A reports that the quintile of stocks with the highest volatility constitutes, on average, only 3.9% of the overall market capitalization. Hence, by excluding the highest quintile of volatile stocks, we exclude stocks that represent only a small fraction of the market. In fact, the highest demi-decile (5%-tile) of the most volatile stocks constitutes only 0.4% of the total market capitalization. Hence, not surprisingly, the highest volatility stocks tend to be small. In the third row, we report the size and book-to-market adjusted returns.¹⁶ Stocks in the highest volatility quintile (demi-decile) underperform their benchmark portfolios by, on average, a large 0.38% (0.67%) per month. This is the puzzling Ang et al. (2006) effect—stocks with very high total or idiosyncratic volatility have low average returns.

The last two rows of panel A explore the interaction between volatility and β^- . While an average stock has a past β^- loading of almost one (0.99), high-volatility stocks tend to have high past β^- . In particular, the average past β^- is 1.44 for the stocks in the highest volatility quintile. We

¹⁵ 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 the tenth and first decile portfolios sorted on past ρ^- is approximately 5% per annum.

¹⁶ Since we use all stocks listed on the NYSE, AMEX, and NASDAQ in constructing size and book-to-market benchmark portfolios, the average adjusted returns of all stocks listed on the NYSE does not sum up to 0.0%.

Table 9
Stocks Sorted by Past β^- Excluding the Most Volatile Stocks

	Panel A: Selected characteristics of volatility portfolios				
	All stocks	Stocks of the highest volatility			
		Quintile	Octile	Decile	5%-tile
Proportion of market capitalization	100.0%	3.9%	1.9%	1.3%	0.4%
Annualized volatility	35.9%	61.0%	68.4%	72.5%	85.6%
Size/book-to-market adjusted returns	-0.08%	-0.38%	-0.50%	-0.56%	-0.67%
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%

Panel B: Average excess returns						
β^- Quintiles excluding most volatile stocks						
1 Low	2	3	4	5 High	High-low	Q4-low
0.58%	0.69%	0.82%	0.82%	0.92%	0.34% [2.31]	0.25% [2.28]

In panel A, each month, we calculate individual stock volatility of all stocks listed on the NYSE using daily continuously compounded returns over the previous 12 months. We sort stocks according to volatility into quintiles, octiles, deciles, and demi-deciles (5%-tiles). Panel A reports selected average characteristics of stocks in each volatility group. The first column reports the characteristics over the entire sample. The other columns report the characteristics within the highest volatility groups. The row labeled "Annualized volatility" reports the average stock volatility over the past 12 months, while the row labeled "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 past β^- over the previous 12 months, and the 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. In panel B, we report the average return in excess of the 1-month Treasury bill rate over the next one month of portfolios sorted on past β^- that exclude the highest volatility quintile of stocks. Each month, we first sort stocks into quintiles according to volatility measured using daily continuously compounded returns over the previous 12 months. Then, we exclude the stocks that fall into the highest volatility quintile and rank the remaining stocks into equal-weighted quintiles (1-5) according to past β^- measured using continuously compounded returns over the previous 12 months. We report the average excess return over the next month. The row labeled "High-low" ("Q4-low") reports the difference between the 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. The sample period is from July 1962 to January 2001.

compare this with stocks in the highest realized β^- quintile of Table 1, which have an average β^- of 1.92. Thus, on the one hand, stocks with high volatility tend to have high β^- . On the other hand, high β^- stocks are not completely stocks with high volatilities.

Moreover, the fact that high-volatility stocks have high past β^- loadings does not imply that they continue to exhibit high β^- loadings in the following period. For stocks in the highest volatility quintile, the average β^- loading over the next 12 months is only 1.25, compared to the β^- over the past 12 months of 1.44. This accounts for the fairly low persistence of the 12-month downside betas of 26%, compared to the average 12-month β^- autocorrelation of 43.5%. If we narrow our focus to stocks on the

highest volatility octile, decile or demi-decile, we find that these stocks exhibit even lower persistence of β^- across the formation period and future holding periods. We conjecture that this lower correspondence of past β^- to future β^- arises, in part, from the fact that high volatility increases the sampling error of the past β^- as a predictor of future downside risk. Large measurement error makes the preformation β^- loadings a more unreliable predictor of actual, postformation, downside betas. If we focus on stocks with lower volatility, our past β^- loadings are more accurate ex ante predictors of future ex post downside risk loadings.

In panel B of Table 9, we examine the average return patterns to ranking on past β^- when we exclude the most volatile stocks. Fortunately, panel A shows that the segment of stocks where past β^- provides a poor predictor of future β^- constitutes only a small fraction of total market capitalization (less than 4%). We first sort stocks based on their past total volatility over the past year. Then, we exclude stocks in the highest volatility quintile and sort the remaining stocks into quintiles according to past β^- . We report the difference in average (unadjusted) returns between the highest quintile and the lowest quintile β^- portfolios, as well as the difference in returns between the second highest and the lowest quintile portfolios. Panel B shows that, by excluding stocks in the highest volatility quintile, the average difference between the highest and the lowest quintile β^- portfolios of 0.34% per month is statistically significant with a t -statistic of 2.31. The difference between the second highest and the lowest quintile β^- portfolios is slightly smaller at 0.25% per month, but has approximately the same statistical significance.¹⁷

3.4 Robustness Checks

In Table 10, we check that the predictive relation between past β^- and future returns, excluding the most volatile stocks, is not due to size, book-to-market, momentum, coskewness, or liquidity effects. The first column of Table 10 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 panel B of Table 9. This difference is highly statistically significant, with a t -statistic of 3.36.

¹⁷ 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. However, the return difference between quintiles 1 and 4 is unaffected for all the volatility exclusions at 0.25% per month and is always statistically significant at the 5% level. This is due to the fact that the volatility effect is strongest among high-volatility stocks which tend to have the highest past β^- .

Table 10
Characteristic Controls on Stocks Sorted by Past β^-

	Size/bk-mkt adjusted	Including additional controls for		
		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%
<i>t</i> -stat	[3.36]	[2.71]	[2.69]	[2.15]

The table reports robustness checks of the results in Table 9. For each month, we compute individual stock volatility and β^- with respect to the market of all NYSE stocks using daily continuously compounded returns over the previous 12 months. We first sort stocks according to volatility into quintiles and exclude stocks that fall within the highest volatility quintile. 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 labeled “Size/bk-mkt adjusted,” we report the average returns in excess of size and book-to-market matched benchmark portfolios similar to Daniel et al. (1997). In the next three columns, we include additional controls for momentum (as measured by past 12-month returns), past coskewness, and past 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 size and book-to-market matched returns within each β^- quintile. The number of stocks in each portfolio varies across time from 177 to 277 stocks. The row labeled “High-low” reports the difference between the returns of quintile portfolios 1 and 5. The entry labeled “*t*-stat” in square brackets is the simple *t*-statistic for the High-low difference. 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.

In the next three columns, we control for additional past return characteristics: momentum, coskewness, and liquidity. 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 size and book-to-market matched returns within each β^- quintile. To control for momentum, we use past 12-month returns. Liquidity is measured using the historical liquidity betas of Pástor and Stambaugh (2003). Table 10 clearly shows that the spreads in size and book-to-market adjusted returns between the β^- quintiles 5 and 1 remain significant at the 5% level after controlling for momentum, coskewness, and liquidity. In each case, the point estimates of the differences average over 0.30% per month. Therefore, our predictive pattern of returns for past β^- are not due to size, book-to-market, past return, coskewness, or liquidity effects.

Finally, as an alternative control for volatility, we consider orthogonalizing the β^- portfolios with respect to volatility in Table 11. We first rank stocks according to past 12-month volatility into quintiles. Then, we

Table 11
Returns of Stocks Sorted by Past Volatility and Past Realized Downside Beta

Portfolio	Past volatility quintiles					Average
	1 Low	2	3	4	5 High	
1 Low β^-	-0.26%	-0.07%	-0.21%	-0.27%	-0.43%	-0.25%
2	-0.23%	0.01%	-0.01%	0.08%	-0.37%	-0.13%
3	-0.15%	0.07%	0.03%	0.00%	-0.31%	-0.07%
4	0.01%	0.11%	0.23%	0.17%	-0.47%	0.01%
5 High β^-	0.08%	0.17%	0.22%	0.14%	-0.31%	0.06%
High-low	0.35%	0.25%	0.43%	0.41%	0.12%	0.31%
<i>t</i> -stat	[3.25]	[2.35]	[3.59]	[2.84]	[0.55]	[3.14]

The table examines the relation between past realized volatility and past realized β^- . For each month, we compute past volatility and past β^- with respect to the market of all stocks listed on NYSE using daily continuously compounded returns over the past 12 months. We first rank stocks into quintiles (1–5) at the beginning of each month based on volatility calculated over the previous 12 months. Then, we rank stocks within each first-sort quintile into additional quintiles according to realized β^- , computed over the previous 12 months. For each 5×5 grouping, we form an equal-weighted portfolio. In each cell, we report the average returns in excess of size and book-to-market matched benchmark portfolios similar to Daniel et al. (1997). In the column labeled “Average”, we report the average size and book-to-market adjusted return of stocks in each second sort quintile, which controls for realized volatility. The row labeled “High-low” reports the difference between the returns of quintile portfolios 5 and portfolio 1. The entry labeled “*t*-stat” in square brackets is the simple *t*-statistic. The sample period is from July 1963 to December 2001.

perform a second ranking within each volatility quintile into quintiles according to β^- over the past 12 months. We then form equal-weighted portfolios within each doubly sorted group, hold the portfolio over the next one month, and report average size and book-to-market adjusted returns.

Table 11 shows a strong positive relation between past downside beta and the next one-month size and book-to-market adjusted return within each volatility quintile. In fact, the differences between the highest and the lowest downside beta quintile returns are significant in all quintiles, except in the highest volatility quintile. When we average the returns of β^- portfolios across volatility quintiles (in the last column), we find that there is a statistically significant difference in returns of 0.31% per month across the highest quintile downside beta stocks and the lowest quintile downside beta stocks. Hence, controlling for volatility, there is a strong predictive relation between past β^- and future returns.

In summary, although there is a strong contemporaneous relation between downside risk and average returns, this relation does not hold predictively for stocks with very high volatility. For the vast majority of stocks, we can use past downside beta to predict future returns. For stocks with high volatility, factor loadings contain large measurement error, making past downside beta a less reliable predictor of future downside risk, and Ang et al. (2006) report that stocks with high volatility have

low returns. Fortunately, these very volatile stocks constitute a small fraction of the total market capitalizations of only 4%. If we focus on stocks not in the highest volatility quintile, we find a strong pattern between past β^- and future holding period returns. However, by not examining the predictive ability of downside risk among stocks with very high volatility, we completely abstract from the puzzling Ang et al. (2006) volatility effect, which we cannot resolve in this article.

4. 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. This risk–return relation is consistent with an economy where agents place 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 list of cross-sectional effects, including size and book-to-market effects, coskewness risk, liquidity risk, and momentum effect. The effect is also different from regular 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. The downside premium is robust across a battery of robustness tests. 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 unconditional covariation of a stock with extreme downside movements of the market. In contrast, we find that the premium for upside risk is weak in the data and often changes signs depending on the set of cross-sectional risk controls.

Past downside beta is a good predictor of future covariation with down market movements, except for stocks that are extremely volatile. For the vast majority of stocks, past downside beta cross-sectionally predicts future returns. However, for stocks with very high volatility, consisting of less than 4% of market capitalization, past downside beta provides a poor predictor of future downside risk. While high-volatility stocks constitute only a small fraction of the total market, so a predictive downside beta relationship holds for the vast majority of stocks, it remains to be explored why the cross-sectional predictive relation for downside risk does not hold for stocks with very high levels of volatility.

Appendix

A Solution of the Disappointment Aversion Asset Allocation Problem

Ang, Bekaert, and Liu (2005) develop an algorithm for solving the portfolio allocation problem for DA utility that transforms the DA asset allocation 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 setup 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 (1990, 2001) show that the First-Order Conditions (FOC) for Equation (3) are given by:

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

where $\mathbf{1}_{\{\cdot\}}$ 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:

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

where wealth in state s is given by:

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

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

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

Ang, Bekaert, and Liu (2005) 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:

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

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:

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

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 $\mu_{W_i}^* = \mu_{W_i}^*$.

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 Condition (A-6) is satisfied. This provides 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 baseline 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 six 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 $w_y^* = 0.4457$, which sum to 1 and which represents the market. 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 alphas, 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 alphas increase with increasing β^- but also increasing coskewness. For this case, $\gamma = 6$, $A = 0.7$, $R_f = 1.05$, $x_u = 0.50 + \mu$, $x_m = 0.20 + \mu$, $x_d = -0.30 + \mu$, $y_u = 0.35$, 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 3, 1962 to December 31, 2001, including NASDAQ data, which is only available after 1972. We use the one-month Treasury bill rate from Ibbotson Associates as 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 12-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:

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

where T is the number of trading days in a period. We estimate downside beta and upside beta by conditioning the observations for which the market realization is below or above the sample mean, $\hat{\mu}_m = \frac{1}{T} \sum r_{mt}$. We calculate demeaned excess return of asset i and demeaned market excess returns conditional on market excess return being below the sample means, denoted \tilde{r}_{it}^- and \tilde{r}_{mt}^- , respectively. We also calculate demeaned excess return of asset i and demeaned market excess returns conditional on market excess return being above the sample means, denoted \tilde{r}_{it}^+ and \tilde{r}_{mt}^+ , respectively. We then calculate $\hat{\beta}^-$ and $\hat{\beta}^+$ as:

$$\hat{\beta}^- = \frac{\sum_{\{r_{mt} < \hat{\mu}_m\}} \tilde{r}_{it}^- \tilde{r}_{mt}^-}{\sum_{\{r_{mt} < \hat{\mu}_m\}} \tilde{r}_{mt}^{-2}}, \text{ and } \hat{\beta}^+ = \frac{\sum_{\{r_{mt} > \hat{\mu}_m\}} \tilde{r}_{it}^+ \tilde{r}_{mt}^+}{\sum_{\{r_{mt} > \hat{\mu}_m\}} \tilde{r}_{mt}^{+2}}. \tag{B-8}$$

We calculate higher-order moments of stock returns using continuously compounded daily returns over each 12-month period. Coskewness and cokurtosis are estimated as:

$$\widehat{\text{coskew}} = \frac{\frac{1}{T} \sum \tilde{r}_i \tilde{r}_i^2}{\sqrt{\frac{1}{T} \sum \tilde{r}_i^2 (\frac{1}{T} \sum \tilde{r}_i^2)}} \text{ and } \widehat{\text{cokurt}} = \frac{\frac{1}{T} \sum \tilde{r}_i \tilde{r}_i^3}{\sqrt{\frac{1}{T} \sum \tilde{r}_i^2 (\frac{1}{T} \sum \tilde{r}_i^2)^{3/2}}}. \tag{B-9}$$

We also collect market capitalizations, book-to-market ratio, and past 12-month returns at the beginning of each 12-month period for each stock.

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:

$$r_{i,t} = \beta_i^0 + \beta_i^L L_t + \beta_i^M r_{m,t} + \beta_i^S SMB_t + \beta_i^H HML_t + \epsilon_{i,t}, \tag{B-10}$$

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 each 12-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 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

- Ang, A., G. Bekaert, and J. Liu, 2005, "Why Stocks May Disappoint," *Journal of Financial Economics*, 76, 471–508.
- Ang, A., and J. Chen, 2002, "Asymmetric Correlations of Equity Returns," *Journal of Financial Economics*, 63, 443–494.
- Ang, A., and J. Chen, 2005, "CAPM Over the Long-Run: 1926–2001," forthcoming in *Journal of Empirical Finance*.
- Ang, A., J. Chen, and Y. Xing, 2002, "Downside Correlation and Expected Stock Returns," NBER Working Paper 8643.
- Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang, 2006, "The Cross-Section of Volatility and Expected Returns," *Journal of Finance*, 51, 259–299.
- Bansal, R., R. Dittmar, and C. Lundblad, 2005, "Consumption, Dividends, and the Cross-Section of Equity Returns," *Journal of Finance*, 60, 1639–1672.
- Bansal, R., D. A. Hsieh, and S. Viswanathan, 1993, "A New Approach to International Arbitrage Pricing," *Journal of Finance*, 48, 1719–1747.
- Banz, R., 1981. "The Relation between Return and Market Value of Common Stocks," *Journal of Financial Economics*, 9, 3–18.
- Barberis, N., and M. Huang, 2001, "Mental Accounting, Loss Aversion, and Individual Stock Returns," *Journal of Finance*, 56, 1247–1292.
- 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.
- 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.
- 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.
- 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.
- Chen, J., Hong, H., and J. Stein, 2001, "Forecasting Crashes: Trading Volume, Past Returns and Conditional Skewness in Stock Prices," *Journal of Financial Economics*, 61, 345–381.
- Daniel, K., M. Grinblatt, S. Titman, and R. Wermers, 1997, "Measuring Mutual Fund Performance with Characteristic Based Benchmarks," *Journal of Finance*, 52, 1035–1058.
- Daniel, K., and S. Titman, 1997, "Evidence on the Characteristics of Cross-Sectional Variation in Stock Returns," *Journal of Finance*, 52, 1–33.
- Dittmar, R., 2002, "Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross-Section of Equity Returns," *Journal of Finance*, 57, 369–403.
- Epstein, L. G., and S. E. Zin, 1990, "First-Order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics*, 26, 387–407.
- Epstein, L. G., and S. E. Zin, 2001, "The Independence Axiom and Asset Returns," *Journal of Empirical Finance*, 8, 537–572.
- Fama, E. F., and K. R. French, 1992, "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, 427–465.

- 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.
- Fama, E. F., and K. R. French, 1997, "Industry Costs of Equity," *Journal of Financial Economics*, 43, 153–193.
- Fama, E. F., and K. R. French, 2005, "The Value Premium and the CAPM," forthcoming *Journal of Finance*.
- Fama, E. F., and J. D. MacBeth, 1973, "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 71, 607–636.
- Ferson, W. E., and C. R. Harvey, 1991, "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99, 385–415.
- Friend, I., and R. Westerfield, 1980, "Co-Skewness and Capital Asset Pricing," *Journal of Finance*, 35, 897–913.
- Gibbons, M. R., 1982, "Multivariate Tests of Financial Models," *Journal of Financial Economics*, 10, 3–27.
- Guidolin, M., and A. Timmermann, 2002, "Optimal Portfolio Choice Under Regime Switching, Skew and Kurtosis Preferences," working paper, University of California, San Diego.
- Gul, F., 1991, "A Theory of Disappointment Aversion," *Econometrica*, 59, 667–686.
- 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.
- Harvey, C. R., J. C. Liechty, M. W. Liechty, and P. Müller, 2003, "Portfolio Selection with Higher Moments," working paper, Duke University.
- Harvey, C. R., and A. Siddique, 1999, "Autoregressive Conditional Skewness," *Journal of Financial and Quantitative Analysis*, 34, 465–477.
- Harvey, C. R., and A. Siddique, 2000, "Conditional Skewness in Asset Pricing Tests," *Journal of Finance*, 55, 1263–1295.
- Hong, H., and J. Stein, 2003, "Differences of Opinion, Short-Sales Constraints and Market Crashes," *Review of Financial Studies*, 16, 487–525.
- Isakov, D., 1999, "Is Beta Still Alive? Conclusive Evidence from the Swiss Stock Market," *European Journal of Finance*, 5, 202–212.
- Jagannathan, R., and Z. Wang, 1996, "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance*, 51, 3–53.
- Jahankhani, A., 1976, "E-V and E-S Capital Asset Pricing Models: Some Empirical Tests," *Journal of Financial and Quantitative Analysis*, 11, 513–528.
- Jegadeesh, N., and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48, 65–91.
- Jondeau, E., and M. Rockinger, 2006, "How Higher Moments Affect the Allocation of Assets," forthcoming in *European Financial Management*.
- Kahneman, D., and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.
- 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, 1355–1382.
- Kothari, S. P., J. Shanken, and R. G. Sloan, 1995, "Another Look at the Cross Section of Expected Stock Returns," *Journal of Finance*, 50, 185–224.

- Kraus, A., and R. Litzenberger, 1976, "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance*, 31, 1085–1100.
- Kraus, A., and R. Litzenberger, 1983, "On the Distributional Conditions for a Consumption-Oriented Three Moment CAPM," *Journal of Finance*, 38, 1381–1391.
- Kyle, A. W., and W. Xiong, 2001, "Contagion as a Wealth Effect of Financial Intermediaries," *Journal of Finance*, 56, 1401–1440.
- Lettau, M., and S. Ludvigson, 2001, "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *Journal of Political Economy*, 109, 1238–1287
- Lewellen, J., and S. Nagel, 2005, "The Conditional CAPM Does Not Explain Asset-Pricing Anomalies," forthcoming in *Journal of Financial Economics*.
- Markowitz, H., 1959, *Portfolio Selection*. Yale University Press, New Haven, CT.
- Newey, W. K., and K. D. West, 1987a, "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–8.
- Newey, W. K., and K. D. West, 1987b, "Hypothesis Testing with Efficient Method of Moments," *International Economic Review*, 28, 777–787.
- Pástor, L., and R. F. Stambaugh, 2003, "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, 642–685.
- Pettengill, G. N., S. Sundaram, and I. Mathur, 1995, "The Conditional Relation between Beta and Returns," *Journal of Financial and Quantitative Analysis*, 30, 101–116.
- 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, 843–855.
- Rouwenhorst, K. G., 1998, "International Momentum Strategies," *Journal of Finance* 53, 267–284.
- Routledge, B. R., and S. E. Zin, 2003, "Generalized Disappointment Aversion and Asset Prices," NBER Working Paper 10107.
- Roy, A. D., 1952, "Safety First and the Holding of Assets," *Econometrica*, 20, 431–449.
- Rubinstein, M., 1973, "The Fundamental Theory of Parameter-Preference Security Valuation," *Journal of Financial and Quantitative Analysis*, 8, 61–69.
- Scholes, M., and J. T. Williams, 1977, "Estimating Betas from Nonsynchronous Data," *Journal of Financial Economics*, 5, 309–27.
- 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, 915–919.
- Shanken, J., 1992, "On the Estimation of Beta-Pricing Models," *Review of Financial Studies*, 5, 1–33.
- Shumway, T., 1997, "Explaining Returns with Loss Aversion," working paper, University of Michigan.