

Measuring Equity Risk with Option-Implied Correlations*

Adrian Buss

Grigory Vilkov

May 31, 2012

Abstract

We use forward-looking information from option prices to estimate option-implied correlations and to construct an option-implied predictor of factor betas. With our implied market betas, we find a monotonically increasing risk-return relation, not detectable with standard rolling-window betas, with the slope close to the market excess return. Our implied betas confirm a risk-return relation consistent with linear factor models, because, when compared to other beta approaches: (i) they are better predictors of realized betas, and (ii) they exhibit smaller and less systematic prediction errors. The predictive power of our betas is not related to known relations between option-implied characteristics and returns.

Keywords: option-implied, correlation, beta, risk-return relation, pairs trading, CAPM, factor model

JEL: G11, G12, G14, G17

*Previously distributed under the title “Option-Implied Correlation and Factor Betas Revisited.” We are grateful for comments and suggestions from the editor, Geert Bekaert, and three anonymous referees. We would also like to thank Yacine Aït-Sahalia, Peter Christoffersen, Joost Driessen, Eduard Dubin, Ralph Koijen, Pascal Maenhout, Christian Schlag, and Raman Uppal for comments and stimulating discussions. We received many helpful comments from the participants of the Adam Smith Asset Pricing Conference 2009, the Humboldt-Copenhagen Conference 2009, the 16th Annual Meeting of the German Finance Association (DGF) 2009, the 12th Conference of the Swiss Society for Financial Market Research 2009, the European Finance Association 2009, as well as the following seminars: ICEF-HSE 2011, Exeter Business School 2012, Tilburg University 2012, and the Brownbag seminar at the University of Frankfurt. We are also grateful to Alexandra Hansis and Yuliya Plyakha for their programming expertise. Send correspondence to Grigory Vilkov, Finance Department, Goethe University, Grüneburgplatz 1, 60323 Frankfurt am Main, Germany. E-mail: vilkov@vilkov.net. Telephone: +49 (69) 798 33732. Adrian Buss, Graduate Program “Finance and Monetary Economics,” Goethe University, buss@finance.uni-frankfurt.de.

Introduction

The linear form of the risk-return relation suggested by the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965) and the Arbitrage Pricing Theory of Ross (1976) has been heavily tested over the years, but it has often been rejected. For example, the early study by Black, Jensen, and Scholes (1972) puts the CAPM in doubt by finding that the expected return on an asset is not strictly proportional to its beta. Later, studies by Reinganum (1981), Lakonishok and Shapiro (1986), and Fama and French (1992) find no relation between market beta and average returns during the 1963–1990 period; recently, Baker, Bradley, and Wurgler (2011) show that high-beta stocks significantly underperform low-beta stocks, more so in the later decades.

All these studies rely solely on historical return information to estimate the most important input factor for testing the linear risk-return relation—the betas of the stocks. Instead of using only historical stock returns, we use *information extracted from current option prices* to measure market betas. Such *forward-looking* betas, reflecting the most recent market information, may be a better proxy for future realized betas than historical ones, especially when betas vary considerably over time.²

We proceed in three steps. First, we propose a new way to model option-implied correlations and demonstrate how to estimate them from option prices.³ Combining these option-implied correlations with option-implied volatilities, we compute forward-looking betas under the risk-neutral probability measure. Though we empirically deal only with option-implied *market* betas, one can, in general, use our methodology to compute betas for any asset pricing factors—even for factors without traded options available.⁴

Second, we use our implied betas to study the risk-return tradeoff implied by linear factor models. The well-documented anomalies that market betas and returns are only

²On the time-variation in betas, and its implications for tests of the CAPM, see, among others, Blume (1975), Bollerslev, Engle, and Wooldridge (1988), Harvey (1989), and Ghysels (1998).

³We are not the first to introduce the notion of implied correlations: Driessen, Maenhout, and Vilkov (2005) computed implied correlations, assuming all pairwise correlations are equal, to quantify the correlation risk premium. In our paper, all pairwise implied correlations are allowed to be different.

⁴In an earlier version of this paper, we also investigated the betas of the Fama and French (1993) size and value factors, and the Carhart (1997) momentum factor, as well as their principal components.

weakly related and that high-beta stocks do not outperform low-beta stocks—common reasons for rejecting the CAPM—do not show up with our implied betas in our 1996–2009 sample. In contrast, we reject a monotone risk-return relation by using historical rolling-window betas computed with five-year monthly and one-year daily returns. The implicit market excess return, measured by the slope of the cross-sectional regression of the portfolios’ mean excess returns on the portfolios’ mean expected betas, is also much closer to the historical equity risk premium for option-implied betas than for historical ones.

Third, we show that the ability of our implied betas to find a “correct” risk-return relation is due to their better predictive quality with respect to future realized betas. Most importantly, our implied betas exhibit a more uniform distribution of the prediction error across beta-sorted portfolios, while other beta methodologies have a large and systematically varying bias. For example, historical betas from daily returns on average overstate the realized betas of the beta-sorted top quintile stocks by 0.13 and understate the realized betas of the beta-sorted bottom quintile by 0.13, thus producing a much flatter risk-return relation.⁵ In contrast, our option-implied betas exhibit almost no bias across portfolios, with the prediction error ranging from -0.04 to 0.03 across all quintiles, resulting in a pronounced relation between market betas and returns.

Technically, we propose a parametric way to estimate implied correlations such that (i) our implied correlation matrix is positive definite with correlations not exceeding one, (ii) the correlation risk premium is negative, consistent with the literature, and (iii) the correlation risk premium is higher in magnitude for low or negatively correlated stocks that are exposed to a higher risk of losing diversification benefits, also consistent with the literature.

We compare our methodology to standard rolling-window historical beta estimators,⁶ and to two alternative approaches using option-implied information for estimating betas. French, Groth, and Kolari (1983) (hereafter, FGK) compute market betas using corre-

⁵We consistently observe this pattern for different numbers of beta-sorted portfolios.

⁶See Fama and MacBeth (1973), Keim and Stambaugh (1986), and Breen, Glosten, and Jagannathan (1989), among others.

lations from historical return data and the ratio of stock-to-market implied volatilities. Chang, Christoffersen, Jacobs, and Vainberg (2012) (hereafter, CCJV) show that in a one-factor model, and under the assumption of zero skewness of the market return residual, option-implied market betas are the product of a function of the stock-to-market implied skewness ratio, serving as a proxy for expected correlation, and the ratio of stock-to-market implied volatilities.

Thus, all considered option-implied betas rely on the same ratio of the stock-to-market implied volatilities in the computation. The only difference is the choice of the proxy for the expected future stock-to-market correlation, going from historical correlation (FGK), to a function of the implied skewness ratio (CCJV), to an explicit parametric form of the implied correlation (our model). We show that the improper choice of the expected correlation proxy can lead to a systematic bias in implied betas.

For example, in our empirical analysis the average weighted market beta, which theoretically must be equal to 1, is about 0.8 for FGK betas and 1.1 for CCJV betas. In contrast, our approach, after taking into account that risk-neutral moments might deviate from their objective counterparts, delivers an unbiased option-implied predictor of realized market betas.

Our empirical analysis of the risk-return relation and predictive qualities of the betas can be used to quantify the gain from using option-implied information in estimating betas. A comparison of the results for the FGK and the historical betas indicates that there is information in individual- and market-implied volatilities beneficial for predicting betas. An even stronger relation between beta and return in our approach suggests that there is additional useful information contained in implied correlations.

The paper is organized as follows. Section 1 introduces the stock market model and discusses the modeling and identification of implied volatilities and correlations. Section 2 provides an overview of the data. In Section 3, we discuss new evidence on the risk-return relation with our option-implied betas, contrast it to the risk-return relation with other beta estimation approaches, and analyze the sources of the differences in the results.

Section 4 concludes with a short summary.

1. Stock Market and Implied Correlation Model

In Section 1.1, we present the underlying stock market model—a linear factor model—and its asset pricing implications. In Section 1.2, we show how to use information from current option prices to obtain implied correlations.

1.1 The model

Our economy contains N traded assets, $i = 1, \dots, N$. The return-generating process for each asset i is given by a linear factor model:

$$r_{i,t+1} = \mu_{i,t} + \sum_{k=1}^K \beta_{ik,t} \delta_{k,t+1} + \varepsilon_{i,t+1},$$

where $\mu_{i,t}$ is the asset’s expected return, each $\delta_{k,t+1}$ denotes a mean zero systematic factor, $\beta_{ik,t}$ denotes the sensitivity of the return on asset i to the innovations in factor k , and $\varepsilon_{i,t+1}$ is the “nonsystematic” risk component with $E[\varepsilon_{i,t+1} | \delta_{k,t+1}] = 0, \forall k$.

For a one-factor model, or for the case of multiple non-correlated factors, the beta of stock i with respect to factor k , $\beta_{ik,t}$, can be directly estimated as the ratio of the stock-to-factor covariance $\sigma_{ik,t}$ to the factor variance $\sigma_{k,t}^2$, or from the pairwise stock correlations $\rho_{ij,t}$, stock volatilities $\sigma_{j,t}$, and the weights of the factor-mimicking portfolio w_j , as follows:⁷

$$\beta_{ik,t} = \frac{\sigma_{ik,t}}{\sigma_{k,t}^2} = \frac{\sigma_{i,t} \sum_{j=1}^N w_j \sigma_{j,t} \rho_{ij,t}}{\sigma_{k,t}^2}. \quad (1)$$

Under several alternative sets of assumptions, one can derive a pricing relation that

⁷For correlated factors, the betas are given by the coefficients of a multivariate regression of the asset’s return on the factors:

$$\beta_{i,t} = \Sigma_{\delta\delta,t}^{-1} \times \Upsilon_{i\delta,t},$$

where $\Sigma_{\delta\delta,t}$ is the factor variance-covariance matrix, and $\Upsilon_{i\delta,t}$ is the vector of stock-to-factor covariances. Both can be computed from stock correlations, volatilities, and the weights of the factor-mimicking portfolios.

links the factor exposures and factor risk premiums to a stock’s expected return, holding for each asset approximately or exactly, as shown in Ross (1976) and Chen and Ingersoll (1983), among others. This pricing, or risk-return, relation stipulates that each systematic factor δ_k bears a risk premium λ_k , and that the exposure of a stock to this factor, measured by the respective beta $\beta_{ik,t}$, is compensated in the expected return of a stock by $\beta_{ik,t} \times \lambda_k$, i.e.,

$$E_t[r_{i,t+1}] = \mu_{i,t} = \lambda_0 + \sum_{k=1}^K \beta_{ik,t} \times \lambda_k + v_i,$$

where v_i denotes the pricing error (that may be zero under certain assumptions).

1.2 Implied volatility and implied correlation modeling

The exposure of a stock to the factors (the betas) is traditionally estimated from historical stock returns. However, as we are interested in the conditional betas, the use of historical data implies the assumption that the future is sufficiently similar to the past. Instead of relying solely on the historical time series of stock returns to compute the beta in equation (1), one can also infer the betas from current prices of traded options. The use of option information may improve the predictive quality of the beta, because option prices subsume current market expectations about future stock dynamics (e.g., Vanden 2008). For example, there is considerable evidence that the use of implied volatilities improves forecasts of realized volatilities [see Poon and Granger (2003) for a review].

Using option data to compute betas raises a number of issues. First, option-implied moments characterize the finite-period *risk-neutral* probability distribution of stock returns. Risk-neutral volatilities and correlations systematically differ from their objective counterparts due to the presence of a variance risk premium (e.g., Bakshi and Madan 2006; Carr and Wu 2009) and a correlation risk premium.⁸ Risk premiums make option-implied

⁸While the pairwise implied correlations for equity markets are not directly observable, Driessen, Maenhout, and Vilkov (2005) compute the homogenous (or equal for each pair of stocks) implied correlation for the S&P 100 and show that it is higher than the average realized correlation; Buraschi, Kosowski, and Trojani (2012) show that “the average monthly payoff of a correlation swap on the average correlation

moments biased predictors of realized moments, and this bias may be inherited by option-implied betas.

Second, in general options on the cross-moments of stock returns do not exist, e.g., options on a basket consisting of two stocks only. Hence, implied correlations are not directly observable from option prices, and one has to make a modeling choice for implied correlations.⁹ In order to formulate our model, we first list the technical conditions that the implied correlation matrix must satisfy, as well as the empirical stylized observations that the implied correlation matrix should satisfy. Then we specify a simple parametric form for implied correlations that is consistent with these technical conditions and empirical observations.

Specifically, for the identification of the implied correlations $\rho_{ij,t}^Q$, we have only one identifying restriction that equates the observed implied variance of the market index $(\sigma_{M,t}^Q)^2$ with the calculated implied variance of a portfolio of all market index constituents $i = 1 \dots N$:

$$(\sigma_{M,t}^Q)^2 = \sum_{i=1}^N \sum_{j=1}^N w_i w_j \sigma_{i,t}^Q \sigma_{j,t}^Q \rho_{ij,t}^Q, \quad (2)$$

where $\sigma_{i,t}^Q$ denotes the implied volatility of stock i in the index and w_i are the index weights. Implied correlations must satisfy two technical conditions: (i) all correlations $\rho_{ij,t}^Q$ do not exceed one, and (ii) the correlation matrix is positive definite. Moreover, we want our implied correlations to be consistent with two empirical observations. First, that the implied correlation $\rho_{ij,t}^Q$ is higher than the correlation under the true measure $\rho_{ij,t}^P$, and, second, that the correlation risk premium is larger in magnitude for pairs of stocks that provide higher diversification benefits (i.e., low or negatively correlated stocks), and, hence, are exposed to a higher risk of losing diversification in bad times characterized by increasing correlations. The second observation is supported by the negative correlation

of S&P 500 stocks is -10.96% "; and Driessen, Maenhout, and Vilkov (2009) derive a correlation factor and show that it bears a negative risk premium, again implying that the risk-neutral correlations are on average higher than the true correlations.

⁹Siegel (1995) created an exchange option that implicitly reveals the market beta of a stock. Yet these options are not traded. Foreign exchange literature uses currency options to infer the implied correlations between currency pairs; see, e.g., Siegel (1997).

between the correlation under the objective measure and the correlation risk premium in Mueller, Stathopoulos, and Vedolin (2012).¹⁰

To identify $N \times (N - 1)/2$ correlations that satisfy restriction (2), as well as the technical and empirical conditions stated above, we propose the following parametric form for implied correlations $\rho_{ij,t}^Q$:

$$\rho_{ij,t}^Q = \rho_{ij,t}^P - \alpha_t(1 - \rho_{ij,t}^P), \quad (3)$$

where $\rho_{ij,t}^P$ is the expected correlation under the objective measure, and α_t denotes the parameter to be identified. Substituting the implied correlations (3) into restriction (2), one can compute α_t in closed form:

$$\alpha_t = -\frac{(\sigma_{M,t}^Q)^2 - \sum_{i=1}^N \sum_{j=1}^N w_i w_j \sigma_{i,t}^Q \sigma_{j,t}^Q \rho_{ij,t}^P}{\sum_{i=1}^N \sum_{j=1}^N w_i w_j \sigma_{i,t}^Q \sigma_{j,t}^Q (1 - \rho_{ij,t}^P)}, \quad (4)$$

and then use equation (3) to identify the full implied correlation matrix Γ_t^Q , with elements $\rho_{ij,t}^Q$.

If $-1 < \alpha_t \leq 0$, then the model satisfies the technical conditions listed above and the model is also consistent with the empirical observations. We explain this below.

First, the identifying restriction (2) is satisfied by construction. Second, for a fixed α_t the correlation risk premium $\rho_{ij,t}^P - \rho_{ij,t}^Q = \alpha_t(1 - \rho_{ij,t}^P)$ is greater in magnitude for those pairs of stocks with low $\rho_{ij,t}^P$, i.e., for stocks with a low correlation under the true measure; and, if $\alpha_t \leq 0$, then $\rho_{ij,t}^Q \geq \rho_{ij,t}^P$. Hence, for $-1 < \alpha_t \leq 0$, both empirical conditions are satisfied. Moreover, if $|\alpha_t| \leq 1$, then $|\rho_{ij,t}^Q| \leq 1$, and, finally, in Theorem 1 below, we show that for $\alpha_t \in (-1, 0]$ the implied correlation matrix is positive definite if the correlation matrix under the true measure is positive definite. Thus for $-1 < \alpha_t \leq 0$, both technical conditions are also satisfied.

¹⁰The average correlation between the correlation under the objective measure and the correlation risk premium is -0.46 for the six pairs of currencies in their paper. We would like to thank the authors for generously providing us with this evidence.

Theorem 1. The implied correlation matrix Γ^Q with elements $\rho_{ij,t}^Q = \rho_{ij,t}^P - \alpha_t(1 - \rho_{ij,t}^P)$ is positive definite if and only if the correlation matrix under the true probability measure Γ^P with elements $\rho_{ij,t}^P$ is positive definite and $\alpha_t \in (-1, 0]$.¹¹

Empirically, all estimated values of the parameter α_t lie in $(-1, 0]$, and, hence, all requirements for the model listed above are satisfied in the empirical part of our paper.

2. Data Description

Our study is based on the major U.S. market proxy, the S&P 500 Index, and its constituents for a sample period from January 4, 1996, to December 31, 2009, a total of 3,524 trading days. In Section 2.1, we briefly describe the stock and option data. Then, in Section 2.2, we discuss the estimation of the realized and option-implied measures.

2.1 Stock and option data

The daily and monthly stock data consist of prices, returns, and number of shares outstanding and come with the S&P 500 Index returns from the CRSP database. Sorted by PERMNO, we have a total of 950 names in our data, which exceeds 500 because of index additions and deletions. To construct a proxy for index weights, we use Compustat. We first merge this database with the daily CRSP data, and then we compute on each day the weights using the closing market capitalization of all current index components from the previous day.

The data for the equity and index options are obtained from the IvyDB OptionMetrics Volatility Surface file that provides Black-Scholes implied volatilities for options with standard maturities and delta levels. We use options with approximately one year to maturity, as we want to estimate the betas for a horizon of one year, and because options with one year to maturity provide a good balance between the stability of the implied volatilities and the liquidity of the underlying market.¹² We select out-of-the-money

¹¹Please refer to the Appendix for the proof.

¹²Note that we use options not as instruments for trading, but as an information source only. If the

options (puts with deltas strictly larger than -0.5 and calls with deltas smaller than 0.5), and on average we have option data for 445 out of the 500 stocks in the S&P 500 Index at each point in time—with the number growing from 373 in 1996 to 483 in 2009, or from 76% to 98% of the index market capitalization.

2.2 Volatility, variance, and skewness estimation

We estimate realized (co)variances as central moments from daily and monthly returns using the rolling-window methodology with one-year and five-year window lengths, respectively.

Risk-neutral variances can be estimated parametrically as the Black and Scholes (1973) implied volatility (IV), or non-parametrically following the model-free implied variance (MFIV) approach of Britten-Jones and Neuberger (2000), Bakshi, Kapadia, and Madan (2003), and others. The advantages of the MFIV are that it (i) subsumes information from all options expiring on one date (e.g., Vanden 2008), (ii) does not rely on any one model, except for some mild restrictions on the underlying stock process (see Bakshi and Madan 2000), and (iii) is the price of a linear portfolio of options.¹³ Accordingly, we use MFIV as a proxy for the expected risk-neutral variance $(\sigma_{i,t}^Q)^2$ for the respective maturity. To compute the model-free moments [in addition to MFIV, we also need the model-free implied skewness (MFIS)] we follow the formulas in Bakshi, Kapadia, and Madan (2003).¹⁴

For analytical purposes, we also need the expected variance risk premium. As in Carr and Wu (2009), we use as proxy for it the realized variance risk premium, estimated for stock i on day t as the average difference between the realized and implied variances over

underlying option market is not liquid or inefficient, and we can still take advantage of it by improving the predictive abilities of the betas, then with a “good” market, we may even have better results.

¹³See Carr and Wu (2006) for a discussion of why the industry has chosen MFIV over IV by switching the methodology for the CBOE traded implied volatility index VIX.

¹⁴To calculate the MFIV and MFIS precisely, in principle, we need a continuum of option prices. We discretize the respective integrals and approximate them from the available options. We have thirteen implied volatilities for OTM options from the surface file at our disposal for each maturity. Using cubic splines, we interpolate these implied volatilities inside the available moneyness range and extrapolate using the last known value (boundary for each side) to fill in 1,000 grid points in the moneyness range from $1/3$ to 3 . We then calculate the option prices from the interpolated volatilities, and we use these prices to compute the implied moments.

the past year:

$$VRP_{i,t} = \frac{1}{252 - \tau} \sum_{\delta=t-252}^{t-\tau-1} \left[(\sigma_{i,\delta}^P)^2 - (\sigma_{i,\delta}^Q)^2 \right], \quad (5)$$

where $(\sigma_{i,\delta}^P)^2$ and $(\sigma_{i,\delta}^Q)^2$ denote the realized and implied variances for a period from δ to $\delta + \tau$. We compute the $VRP_{i,t}$ for a maturity τ of 21 (trading) days.

3. The Risk-Return Relation

In this section, we examine the relation between stock market betas and expected returns. In Section 3.1, we compare the risk-return relation generated by our option-implied beta predictors with that generated by traditional historical betas. Next, in Section 3.2, we discuss the risk-return relation for alternative option-implied beta approaches. Finally, in Section 3.3, we reconcile the observed differences in the risk-return relation by contrasting the predictive qualities of the different beta approaches.

3.1 Risk-return with option-implied vs. historical betas

Traditionally, tests of the beta-return relation use historical rolling-window betas. Most typically, authors use five years of trailing monthly returns (e.g., Black, Jensen, and Scholes 1972), or one year of daily returns (e.g., Baker, Bradley, and Wurgler 2011, among many others) for their computations. For our tests, we follow these accepted practices and compute historical betas using daily and monthly stock and index returns, with respective rolling-window lengths of one and five years. At the end of each month within our sample period, we compute the stock-to-market covariances, as well as the market variance, and use them in equation (1) to produce market beta predictors for each stock.

For the option-implied betas, we first estimate implied volatilities and correlations. The implied volatilities $\sigma_{i,t}^Q$ are obtained from the individual and index options, as discussed in Section 2.2. The implied correlations $\rho_{ij,t}^Q$ are computed from definition (3), after calibrating the only unknown parameter α_t from equation (4). We use as inputs the implied volatilities from options with one year to maturity, and historical correla-

tions from daily or monthly stock returns with respective rolling-window lengths of one or five years, as a proxy for the expected correlations under the objective measure. Then option-implied market beta $\beta_{iM,t}^Q$ of stock i can be computed from equation (1) as:

$$\beta_{iM,t}^Q = \frac{\sigma_{i,t}^Q \sum_{j=1}^N w_j \sigma_{j,t}^Q \rho_{ij,t}^Q}{(\sigma_{M,t}^Q)^2}.$$

To establish the risk-return relation, we perform a portfolio sorting exercise, where we sort, at the end of each month, and separately for each set of betas, the stocks into portfolios according to their expected beta; compute for each portfolio the value-weighted realized return over the next month; and analyze the panel of portfolio returns to infer the form of the resulting beta-return relation.¹⁵ To gauge the value of the option-implied information, we always analyze the performance of the betas pairwise, i.e., we contrast historical daily betas with option-implied daily betas, and historical monthly betas with option-implied monthly betas.

In Table 1, we provide a summary of the mean expected betas and the mean realized returns for the beta-sorted quintile portfolios. Figure 1, Panels A and B present the risk-return relations for the two pairs of betas in question. We observe that the implied betas show in both panels a less noisy relation across different quintiles and informally can be labeled “more linear.” In addition, the return difference between the extreme quintile portfolios is more pronounced for implied betas than for the historical ones. For example, for daily historical and implied betas, we have 0.60% p.a. vs. 5.57% p.a. return difference, respectively, and for monthly historical and implied betas, we have -0.74% p.a. vs. 7.63% p.a., respectively. This result indicates that the relation between market beta risk and returns is much stronger for our implied betas.

We also perform a formal monotonicity test of the risk-return relation, applying the

¹⁵This portfolio-sorting methodology is similar to the early study of Black, Jensen, and Scholes (1972), as well as that of Baker, Bradley, and Wurgler (2011). Differences exist in the sample period, data frequency, and stocks selected; for example, Baker, Bradley, and Wurgler use monthly data over a longer sample period and select the top 1,000 stocks by market capitalization, while we have daily and monthly data for all U.S. stocks, which are in the S&P 500 Index at a given point in time.

non-parametric technique of Patton and Timmermann (2010).¹⁶ Table 1 shows the results of the test for two null hypotheses: (1) there is a monotonically decreasing risk-return relation, and (2) there is a monotonically increasing risk-return relation, with p -values obtained from time-series block bootstrapping. To find confirmation of a monotonically increasing relation consistent with linear market models, one needs to reject the null of a decreasing relation and fail to reject the null of an increasing relation. For our sample period, we fail to reject a monotonically increasing relation for all historical as well as option-implied betas but reject the null of a monotonically decreasing relation for the option-implied betas only—at the 5% level for daily implied betas and at the 10% level for monthly implied betas. Most important, by using option-implied information in the beta computation we always improve the statistics in the monotonicity tests in favor of linear market models: for the daily betas, the p -value for the decreasing null hypothesis drops from 31% to 5%, and for the monthly betas from 67% to 10%; the p -value for the increasing null goes up from 40% to 71%, and from 60% to 81% for daily and monthly betas, respectively.

In addition to the visual inspection and formal monotonicity test for the quintile beta-sorted portfolios, we also implement cross-sectional regressions for a larger number of portfolios. Similar to the procedure above, we sort the stocks into 10, 25, or 50 portfolios, according to their expected beta, and compute the time-series of the realized portfolio excess returns. Next, we regress the mean realized excess returns of the portfolios on a constant and the mean expected betas of these portfolios, similar to Black, Jensen, and Scholes (1972). Table 2 shows the results of these regressions, together with a significance test on the coefficients.

For the historical daily betas, the slope is typically insignificant, and we cannot find an anticipated relation between market betas and returns. In contrast, for our implied daily betas, the slope of the regression is always significant, and the average (across regressions with a different number of portfolios) slope coefficient of 5.3% indicates that

¹⁶Following the recommendation of Romano and Wolf (2011), we use quintile portfolios in monotonicity tests.

the compensation for taking market risk is close to the historical market excess return. Nor can we find a relation between monthly historical betas and returns, while we observe a strong positive relation with our monthly implied betas.

Thus, compared to historical betas, our option-implied betas deliver better results in terms of the risk-return relation: (i) we reject a monotonically decreasing risk-return relation at 5% and 10% levels, consistent with the predictions of linear factor models; (ii) the slope in the cross-sectional regression, i.e., the compensation for market risk, is significant; and (iii) the estimate of the expected market excess return is in magnitude close to the realized market excess return.

3.2 Risk-return with alternative option-implied betas

There exist two alternative procedures for computing betas using option-implied information. French, Groth, and Kolar (1983) (FGK, henceforth) combine historical stock-to-market correlations with option-implied volatilities for a stock and the market in a conventional formula:

$$\beta_{iM,t}^{FGK} = \rho_{iM,t}^P \times \frac{\sigma_{i,t}^Q}{\sigma_{M,t}^Q}.$$

Chang, Christoffersen, Jacobs, and Vainberg (2012) (CCJV, henceforth) suggest using the risk-neutral model-free skewness (*MFIS*) implied by current option prices, combined with option-implied volatilities, as follows:

$$\beta_{iM,t}^{CCJV} = \left(\frac{MFIS_{i,t}}{MFIS_{M,t}} \right)^{\frac{1}{3}} \times \frac{\sigma_{i,t}^Q}{\sigma_{M,t}^Q}. \quad (6)$$

The first part of the above beta expression (6) serves as a proxy for the risk-neutral correlation.¹⁷

We compute two versions of the FGK betas: one using stock-to-market correlations $\rho_{iM,t}^P$ from daily returns, and the other from monthly returns, with one- and five-year

¹⁷Note that this correlation proxy is only defined for individual stocks with a negative skewness, given that the market skewness proxy is typically negative. As a result, one has to eliminate all stocks with a positive implied skewness from the sample used for the CCJV performance analysis.

rolling-window lengths, respectively. The implied volatility $\sigma_{i,t}^Q$ and the implied skewness $MFIS_{i,t}$ are computed from options with one year to maturity, as before.¹⁸ We then perform the same risk-return relation investigation for the alternative option-implied betas as we did in the previous section. That is, at the end of each month we sort the available stocks into portfolios based on the expected FGK or CCJV betas and compute the one-month holding period returns for each portfolio.

Table 1 provides a summary of the mean expected betas and the mean realized returns for the beta-sorted quintile portfolios. Figure 1, Panels C and D present the risk-return relation for both FGK betas and CCJV betas, respectively. We observe that for both daily and monthly historical correlations, the FGK betas show an increasing relation between betas and returns—though the return difference between the extreme quintile portfolios is rather small (about 2.6% and 1.3%, respectively), whereas the CCJV line in Panel D is noisy and almost flat.

Not surprisingly, these patterns are reflected in the results of the monotonicity tests in Table 1. While we always fail to reject the null of an increasing relation for the FGK betas, and reject a monotone decreasing relation for the daily FGK betas at the 7% level, we fail to reject the decreasing null for the monthly FGK betas. That is, the use of option-implied information in the way French, Groth, and Kolari (1983) propose improves the risk-return relation compared to the historical betas, but the results are always weaker than for our option-implied betas. For the CCJV betas, we are more compelled to reject the increasing relation with a p -value of 15% than to reject the decreasing relation—a consequence of the very flat and negative risk-return picture, as shown in Panel D of Figure 1.

Table 2 presents the results of the cross-sectional regression for the alternative beta approaches. For the FGK daily betas, we find a positive and significant slope coefficient supporting a relation between beta and return, though the mean slope of 3.6% is lower than for our implied betas. For the FGK monthly betas, the support of the risk-return

¹⁸CCJV use a 180-day maturity “... to some extent based on a trade-off between option liquidity which is largest for options with 30 to 90 days to maturity, and the relevant horizon for firm risk, which is arguably considerably longer.”

relation is weaker, with consistently higher p -values and smaller slope coefficients. Thus, the use of implied volatilities in the FGK approach improves the risk-return relation fit compared to historical betas, though not as much as the use of implied volatilities *and* implied correlations in our implied betas. For the CCJV betas, we find no evidence for a positive cross-sectional relation between betas and returns, and for 50 portfolios we even find a significantly negative slope coefficient, indicating a negative expected market excess return.

In summary, using FGK betas we find evidence that the risk-return relation is relatively monotone and that one is compensated for bearing market beta risk; however, the evidence is weaker than for our option-implied betas. In contrast, the CCJV betas deliver a risk-return picture that is too flat, resulting in an expected market excess return close to zero.

3.3 Why option-implied betas work: comparison with other methods

The results in the preceding sections indicate that our option-implied betas demonstrate a risk-return relation anticipated from linear market models—it is positive and monotone, and it delivers an expected market excess return estimate close to the historical equity risk premium. In the following section we relate the performance of our implied betas to their forecasting qualities, and study biases that may potentially affect the documented results.

3.3.1 Beta predictability

One can expect to observe a realistic risk-return relation only if one uses expected betas that have a good predictive power and are unbiased with respect to realized betas. For instance, even allowing the CAPM to be the true model, but using an empirical methodology to sort the stocks in the wrong order, one may not find any positive risk-return relation in the sample.

We start our analysis of the predictive qualities of the different beta approaches with the summary statistics in Table 3. First, we want to see if some methods deliver biased conditional betas with respect to the realized ones. We therefore compute for each month the value-weighted average of the individual stock's betas for all stocks in the S&P 500. Assuming constant market weights over the period for which the betas are computed,¹⁹ this quantity is theoretically equal to one. Our computations show that the weighted average FGK betas are 19% to 30% lower, and the CCJV betas are about 10% higher than one. In part, the positive bias of the CCJV betas can be explained by the fact that these betas are positive by construction, with the minimal beta being around 0.04, while all other methodologies give some very negative betas. In contrast, the value-weighted mean of the implied betas that we propose in this paper is unbiased by construction.

Note that a uniform bias does not affect the results from a cross-sectional beta-sorting exercise, nor does it lead to erroneous estimates of the market risk premium. For example, consider a beta approach that overestimates true betas uniformly by a certain amount. The portfolio sorting would then still generate portfolios in which the stocks are sorted according to their true betas. But if size and direction of the bias change across portfolios, the sorting procedure, and hence the form of the risk-return relation, will be strongly affected.

Figure 2 presents the mean prediction error for the portfolios sorted on market betas, computed for each quintile and beta methodology as the time-series mean of the difference between the realized and the expected portfolio betas. Panels A and B highlight that the prediction errors are smallest in magnitude and do not demonstrate any pronounced pattern across quintiles for our option-implied betas. In contrast, there is a strong monotone decreasing pattern for the historical rolling-window methodologies. For instance, for the daily historical approach, the realized beta of the bottom quintile portfolio is much higher than the expected beta, with an average prediction error of 0.13, and the realized beta of the top quintile portfolio is lower, with an average prediction error of -0.13 . For the

¹⁹Simulations, available upon request, show that the changing market weights have a marginal influence on the value-weighted sum of stock betas.

historical monthly betas, the prediction errors are even bigger, ranging from about 0.30 to -0.40 .²⁰ Such a misestimation of portfolio betas has a strong influence on the results of the beta-return test. Specifically, the underestimation of the low quintile portfolio betas yields too high returns for the low quintile portfolios, in case the relation holds for the “true” betas; similarly, the overestimation of the high quintile portfolio betas causes too low returns for the high quintile portfolios. In this case, detecting an existing linear beta-return relation would be extremely hard, given these prediction errors, especially if one uses the top-minus-bottom portfolio only. For example, prediction error differences between the top and bottom quintile portfolio of about 0.26 and 0.70—as is the case with the historical methods—would lead to an erroneous underestimation of the difference in the annualized portfolios’ returns of about 1.5% and 4.2%, assuming that the linear beta-return relation holds and the historical market excess return is about 6%.

For the FGK monthly and the CCJV betas, plotted in Figure 2, Panels C and D, we find patterns similar to the ones observed for the historical rolling-window methods (differing only in magnitude), which also causes problems in establishing a correct risk-return picture. The prediction error for the daily FGK betas is more uniform, which explains the better-looking beta-return relation found for these betas compared to the other two option-implied beta approaches.

In addition to the average prediction error analysis, we also look at scatterplots, shown in Figure 3, where we plot, for all quintile portfolios and all months in the sample period, the expected beta of a portfolio against the realized beta of the same portfolio for a forecast horizon of one year.

If a beta model fits the data well, we expect to observe a clustering of points around the 45 degree line. In Panels B and D of Figure 3 the implied beta forecasts are located close to the realized betas, hence providing a very good fit, with no clear asymmetries emerging. The fit of the historical methodologies, shown in Panels A and C, is worse, and, especially for monthly betas, we observe some strong asymmetries—high betas are

²⁰Blume (1975) used the term “regression tendency” to describe the “tendency for a portfolio with either an extremely low or high estimated beta in one period to have a less extreme beta as estimated in the next period.”

underestimated and low betas overestimated. A similar but even stronger pattern emerges for the FGK betas: about 90% of the betas are underestimated, due to the systematic bias discussed before. From the two FGK methods, the daily betas look slightly better. The graph in Panel G for the CCJV betas also shows strong asymmetries, and about 70% of the betas are overestimated.

Table 4 demonstrates the predictive abilities of the different beta approaches in the time-series for setups with 10, 25, and 50 beta-sorted portfolios. First, we show: the average mean-squared error (MSE) across all portfolios; a p -value for the test that the MSE for a given beta predictor is the same as for the historical counterpart; and the number of portfolios for which a specific approach yields the lowest MSE.²¹ Second, we compute for the bottom and top portfolios the mean prediction error, defined as the mean difference between the realized and the predicted portfolio betas, the top-bottom difference, and a p -value testing if the top-bottom difference is significantly different from zero. While the MSE yields a measure of the overall predictability performance, the difference in the prediction errors for extreme portfolios indicates if there is a systematic bias in the beta predictability.

We find that our option-implied betas always have the lowest mean-squared errors across all methodologies, i.e., they predict future betas better than the historical approaches, as well as the FGK and CCJV approaches. The difference between our betas and the historical as well as other approaches is always highly significant, and our implied betas have on average the lowest MSE for about 80% of the portfolios. Due to their systematic bias, the daily FGK betas always show a significantly higher MSE than that of the daily historical methodology. On the monthly frequency, the results are about the same. Finally, the MSE for the CCJV betas is consistently higher than the errors for the daily approaches, and about the same as the errors for the monthly approaches.

In addition, Table 4 presents evidence that our implied betas exhibit a significantly

²¹Note that we only compare approaches that use the same data frequency for the correlation computation—either daily or monthly. For example, at the daily frequency we only compare historical daily, our daily option-implied, and FGK daily beta approaches. We proceed similarly for the monthly frequency.

smaller difference in the prediction errors for extreme portfolios, i.e., they have a smaller systematic bias in the cross-section at the daily and monthly frequency. For example, the mean difference in the prediction errors for the extreme portfolios for the daily historical approach is about -0.40 , resulting from a positive bias of about 0.17 for the bottom portfolio and a negative bias of about -0.23 for the top portfolio, and it is always significant. In contrast, the difference for our option-implied daily betas is only -0.05 and not significantly different from zero. Similarly, the FGK daily betas always have a small and insignificant difference in prediction errors for the extreme portfolios. Compared to our implied betas, FGK betas demonstrate a smaller difference in the prediction errors for the extreme portfolios for all setups at the daily frequency, but they are worse for all setups at the monthly frequency. Similar to the MSE results, the difference in the prediction errors for the extreme portfolios for the CCJV betas lies between the differences for the other methods at daily and monthly frequencies. As expected, the difference in the prediction errors for extreme portfolios grows for all approaches with the number of portfolios, as the top and bottom portfolios get more extreme in terms of expected beta.

To see how the different beta methodologies can predict realized betas on the individual stock level, as opposed to the portfolio level, we perform one more exercise. It resembles the popular pairs-trading strategy, in which we aim at getting a zero market risk exposure of a two-stock portfolio by adding a position in the market proxy to it.

We first form a large number of portfolios, each consisting of two stocks from our sample, randomly going long one stock and short the other one.²² To study the ability of the different methods to make accurate market beta predictions, we add the market proxy to the portfolio in such a way that it becomes expected market neutral, i.e., we

²²This practice is typical for “stock picking,” in which a fund expects the first stock to outperform the second one. It is also typical for statistical arbitrage models, in which a fund trades two cointegrated stocks if the spread strongly deviates from its long-run mean and liquidates the positions when the spread returns to the equilibrium.

choose the weight w such that the expected market beta of the portfolio is zero:

$$(1 - w) \times (\beta_{1M,t} - \beta_{2M,t}) + w \times \beta_{MM,t} = 0,$$

where $\beta_{MM,t} = 1$ denotes the market beta of the market itself, and $\beta_{iM,t}$, $i = 1, 2$, denotes the beta of each stock with respect to the market. Given the optimal weight w , we compute the *realized* market beta of the portfolio over the next month, which also gives us the deviation from the target beta of zero.²³ We use the latter to compute the squared error for each portfolio and month—our main performance criterion in this exercise. We finally pool all observations, i.e., across time and portfolios, and compute the mean squared error, together with a p -value for the difference with the historical beta counterpart, based on standard errors clustered by time and pair, following Petersen (2009).

The results of this pairs-trading application are presented in Table 5. The performance of the option-implied betas, with a mean squared error of 0.0329 and 0.0344 for daily and monthly frequencies, respectively, is significantly better than 0.0347 and 0.0675 of the corresponding historical methods, with p -values below 0.01. The FGK betas perform worse than our implied betas at the daily frequency and monthly frequency due to their bias.

In summary, while both our implied betas and the FGK betas exhibit by far the lowest systematic prediction error across portfolios, only our option-implied betas exhibit the lowest prediction error in the time-series, i.e., they also fit the level of the betas best. This evidence demonstrates the reason why our implied betas and the FGK betas outperform the other methods in showing a correct risk-return relation, and why our betas perform better than FGK betas, especially in estimating the slope coefficient in the cross-sectional regression that should resemble the historical equity risk premium.

²³Because pairs trading positions typically are rebalanced frequently, we consider a short portfolio holding period of only four weeks.

3.3.2 Can other option-implied return patterns explain the results?

Given the fact that our option-implied betas use risk-neutral volatilities and correlations, and that these implied estimates differ from the volatilities and correlations under the objective probability measure due to the presence of risk premiums, option-implied betas may subsume some other option-related characteristics for predicting stock returns, thus demonstrating a spurious beta-return relation.

For example, both the variance risk premium and implied skewness are good predictors of stock returns: Goyal and Saretto (2009)²⁴ and Bali and Hovakimian (2009) show that the variance risk premium is negatively linked to returns; and Cremers and Weinbaum (2010), Rehman and Vilkov (2010), Xing, Zhang, and Zhao (2010), and Conrad, Dittmar, and Ghysels (2012), among others, find a significant relation between implied skewness and future returns. If there is a difference in the distribution of these characteristics across beta-sorted portfolios for the different beta approaches, then the risk-return relation found in the previous section may be spurious and not generated by the exposure to the market factor.

We specify three conditions that are necessary for implied betas to capture a spurious return pattern in such a way that the risk-return relation looks correct with our betas, though it does not look correct with other betas. First, there must be a sizable difference and clear pattern in the average characteristic value of the beta-sorted portfolios. Second, there must be strong differences in the average characteristic value of the portfolios across beta methodologies. Third, the return pattern arising from the predictive power of a given characteristic and the beta-return relation for our betas must work in the same direction. Importantly, *all* three conditions must be satisfied to give rise to a spurious risk-return relation.

To see if these three conditions are satisfied for the aforementioned characteristics, and we indeed documented a correct-looking but spurious risk-return relation with our option-implied betas, we compute for each beta methodology and each quintile portfolio

²⁴While Goyal and Saretto (2009) concentrate on the analysis of option returns, in their working paper version they also provide the results for stock portfolios.

the value-weighted average of the variance risk premium and the implied skewness of the individual stocks.

In Table 6, Panel A shows that the average variance risk premium for the option-implied betas is only slightly increasing across quintile portfolios. Moreover, comparing the average variance risk premium across quintile portfolios for the historical and the option-implied methods, one can see that the difference is marginal for all portfolios. Similar results hold if compared to the FGK methods. To understand if this difference is economically significant, we compute the excess return associated with a long-short portfolio strategy based on the variance risk premiums of individual stocks.²⁵ Similar to Goyal and Saretto (2009), we find that stocks with a high variance risk premium significantly underperform stocks with a low variance risk premium by -5.54% p.a. The difference in the average variance risk premium for the top and bottom quintiles in the long-short portfolio is about ten times higher than the difference for our beta-sorted portfolios, implying that the return difference between the quintile portfolios based on the historical betas may be biased downwards by about 0.5% p.a. By correcting the risk-return relation with historical betas for a difference of 0.5% , we would make it look better but still not as good as the risk-return picture with our implied betas.

Consequently, these results do not provide any evidence for a spurious risk-return relation due to return patterns induced by the variance risk premium of the individual stocks in the portfolios. Looking at the alternative option-implied methods, the results suggest that the FGK betas may be biased downwards, but in very weak a fashion, and that for the CCJV betas the bias (if any) should go in the opposite direction.

In Panel B of Table 6, we observe for all methods, except the CCJV betas, no clear pattern in the average portfolio skewness. The increasing pattern for the CCJV betas can be explained with the help of equation (6): the correlation proxy $\left(\frac{MFIS_{i,t}}{MFIS_{M,t}}\right)^{\frac{1}{3}}$ is increasing in the individual stock skewness $MFIS_{i,t}$, and, hence, the CCJV betas are higher for the stocks with more negative skewness values.

²⁵We sort the available stocks at the end of each month into quintile portfolios based on their individual volatility risk premium and compute the value-weighted return on the long-short portfolio of the top and bottom quintiles in the next month.

Comparing the option-implied methods to their historical counterparts, we find that the difference in the average skewness between the two extreme quintiles is slightly higher for our option-implied methods. To see if this difference can explain the observed differences in the risk-return relation, we compute the excess return associated with a long-short portfolio strategy based on the skewness of the individual stocks.²⁶ For our sample, the value-weighted return on such a portfolio is slightly negative and insignificant. Thus, the higher positive skewness difference between the extreme beta-sorted quintiles for our implied betas would work even against an increasing risk-return relation, and, hence, excessive exposure to option-implied skewness cannot explain the better risk-return relation for our option-implied betas.

4. Conclusion

Linear factor models, and the resulting linear risk-return tradeoff, play an important role in asset pricing and portfolio management. It is clear that for these models to perform well, an accurate prediction of the stock factor exposures, i.e., the factor betas, is crucial. In this paper, we show how to use forward-looking information contained in current option prices to construct implied correlations and, accordingly, to build option-implied beta predictors.

With respect to the market factor, our option-implied betas demonstrate the correct positive and monotone risk-return relation, while standard historical rolling-window betas fail to detect a monotonically increasing relation between beta and return.

We show that the main reason for the stronger performance of our implied betas is their predictive quality with respect to the realized betas. Implied betas demonstrate a smaller and more uniformly distributed bias (average prediction error) than the historical and other option-implied betas, which is especially important for creating a correct ranking by factor exposure in the stock sample. Controlling for the effect of the option-implied

²⁶Similar to the analysis of the variance risk premium, we create the top-minus-bottom quintile portfolio from an option-implied skewness sorting and compute its value-weighted return.

characteristics (variance risk premium and implied skewness), which are able to predict stock returns and hence may affect the sorting in the case of our implied betas, we show that the performance of our betas is not linked to the return predictive qualities of these characteristics.

Our method of modeling option-implied correlations, and the application of these option-implied correlations for computing forward-looking betas, have numerous empirical applications in asset pricing and risk management and can be used to address various problems in research and practice. For example, the better predictive quality of our betas should lead to better cost-of-capital estimates.

Appendix: Proofs

Theorem 1. The implied correlation matrix Γ^Q with elements $\rho_{ij,t}^Q = \rho_{ij,t}^P - \alpha_t(1 - \rho_{ij,t}^P)$ is positive definite if and only if the correlation matrix under true probability measure Γ^P with elements $\rho_{ij,t}^P$ is positive definite and $\alpha_t \in (-1, 0]$.

Proof. The implied correlation matrix Γ^Q can be rewritten as:

$$\Gamma^Q = \Gamma^P \times (1 + \alpha_t) - \alpha_t \times \iota \cdot \iota',$$

where Γ^P denotes the correlation matrix under the true measure P , and ι denotes an $N \times 1$ vector of ones. For $\alpha_t \in (-1, 0]$, the first part in the expression above $\Gamma^P(1 + \alpha_t)$ is positive definite, and the second part $-\alpha_t \times \iota \cdot \iota'$ is positive semidefinite, because the matrix $\iota \cdot \iota'$ is positive semidefinite. The sum of positive definite and positive semidefinite matrices is positive definite, and it implies the positive definiteness of Γ^Q for $\alpha_t \in (-1, 0]$. For other values of α_t we cannot establish the positive definiteness of the implied correlation matrix, because for $\alpha_t \leq -1$ the first part of the expression above will be negative definite, and for $\alpha_t > 0$ the second part will be negative semidefinite. ■

References

- Baker, M. P., B. Bradley, and J. A. Wurgler. 2011. Benchmarks as Limits to Arbitrage: Understanding the Low Volatility Anomaly. *Financial Analysts Journal* 67:1–15.
- Bakshi, G. S., N. Kapadia, and D. B. Madan. 2003. Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies* 16:101–43.
- Bakshi, G. S., and D. B. Madan. 2000. Spanning and Derivative Security Valuation. *Journal of Financial Economics* 55:205–38.
- Bakshi, G. S., and D. B. Madan. 2006. A Theory of Volatility Spreads. *Management Science* 52:1945–56.
- Bali, T. G., and A. Hovakimian. 2009. Volatility Spreads and Expected Stock Returns. *Management Science* 55:1797–812.
- Black, F., M. C. Jensen, and M. S. Scholes. 1972. The Capital Asset Pricing Model: Some Empirical Tests. In M. C. Jensen (ed.), *Studies in the Theory of Capital Markets*. New York: Praeger Publishers Inc.
- Black, F., and M. S. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81:637–54.
- Blume, M. E. 1975. Betas and Their Regression Tendencies. *Journal of Finance* 30:785–95.
- Bollerslev, T., R. F. Engle, and J. M. Wooldridge. 1988. A Capital Asset Pricing Model with Time-Varying Covariances. *Journal of Political Economy* 96:116–31.
- Breen, W., L. R. Glosten, and R. Jagannathan. 1989. Economic Significance of Predictable Variations in Stock Index Returns. *Journal of Finance* 44:1177–89.
- Britten-Jones, M., and A. Neuberger. 2000. Option Prices, Implied Price Processes, and Stochastic Volatility. *Journal of Finance* 55:839–66.
- Buraschi, A., R. Kosowski, and F. Trojani. 2012. When There is No Place to Hide — Correlation Risk and the Cross-Section of Hedge Fund Returns. Working Paper.
- Carhart, M. M. 1997. On Persistence in Mutual Fund Performance. *Journal of Finance* 52:57–82.

- Carr, P., and L. Wu. 2006. A Tale of Two Indices. *Journal of Derivatives* 13:13–29.
- Carr, P., and L. Wu. 2009. Variance Risk Premiums. *Review of Financial Studies* 22:1311–41.
- Chang, B.-Y., P. Christoffersen, K. Jacobs, and G. Vainberg. 2012. Option-Implied Measures of Equity Risk. *Review of Finance* 16:385–428.
- Chen, N.-F., and J. E. Ingersoll. 1983. Exact Pricing in Linear Factor Models with Finitely Many Assets: A Note. *Journal of Finance* 38:985–8.
- Conrad, J. S., R. F. Dittmar, and E. Ghysels. 2012. Ex Ante Skewness and Expected Stock Returns. *Journal of Finance*.
- Cremers, M., and D. Weinbaum. 2010. Deviations from Put-Call Parity and Stock Return Predictability. *Journal of Financial and Quantitative Analysis* 45:335–67.
- Driessen, J., P. Maenhout, and G. Vilkov. 2005. Option-Implied Correlations and the Price of Correlation Risk. Working Paper, INSEAD.
- Driessen, J., P. Maenhout, and G. Vilkov. 2009. The Price of Correlation Risk: Evidence from Equity Options. *Journal of Finance* 64:1375–404.
- Fama, E. F., and K. R. French. 1992. The Cross-Section of Expected Stock Returns. *Journal of Finance* 47:427–65.
- Fama, E. F., and K. R. French. 1993. Common Risk Factors in the Returns on Stock and Bonds. *Journal of Financial Economics* 33:3–56.
- Fama, E. F., and J. D. MacBeth. 1973. Risk Return and Equilibrium: Empirical Tests. *Journal of Political Economy* 71:607–36.
- French, D. W., J. C. Groth, and J. W. Kolari. 1983. Current Investor Expectations and Better Betas. *Journal of Portfolio Management* 10:12–8.
- Ghysels, E. 1998. On Stable Factor Structures in the Pricing of Risk: Do Time-Varying Betas Help or Hurt? *Journal of Finance* 53:549–73.
- Goyal, A., and A. Saretto. 2009. Cross-Section of Option Returns and Volatility. *Journal of Financial Economics* 94:310–26.
- Harvey, C. R. 1989. Time-varying Conditional Covariances in Tests of Asset Pricing Models. *Journal of Financial Economics* 24:289–317.

- Keim, D. B., and R. F. Stambaugh. 1986. Predicting Returns in the Stock and Bond Markets. *Journal of Financial Economics* 17:357–90.
- Lakonishok, J., and A. C. Shapiro. 1986. Systematic Risk, Total Risk and Size as Determinants of Stock Market Returns. *Journal of Banking and Finance* 10:115–32.
- Lintner, J. 1965. The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *Review of Economics and Statistics* 47:13–37.
- Mueller, P., A. Stathopoulos, and A. Vedolin. 2012. International Correlation Risk. Working Paper, LSE.
- Patton, A. J., and A. Timmermann. 2010. Monotonicity in Asset Returns: New Tests with Applications to the Term Structure, the CAPM and Portfolio Sorts. *Journal of Financial Economics* 98:605–25.
- Petersen, M. A. 2009. Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. *Review of Financial Studies* 22:435–80.
- Poon, S.-H., and C. W. Granger. 2003. Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature* 41:478–539.
- Rehman, Z., and G. Vilkov. 2010. Risk-Neutral Skewness: Return Predictability and Its Sources. Working Paper.
- Reinganum, M. R. 1981. A New Empirical Perspective on the CAPM. *Journal of Financial and Quantitative Analysis* 16:439–62.
- Romano, J. P., and M. Wolf. 2011. Alternative Tests for Monotonicity in Expected Asset Returns. Working Paper, Stanford University.
- Ross, S. 1976. The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* 13:341–60.
- Sharpe, W. F. 1964. Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *Journal of Finance* 19:425–42.
- Siegel, A. F. 1995. Measuring Systematic Risk Using Implicit Beta. *Management Science* 41:124–8.
- Siegel, A. F. 1997. International Currency Relationship Information Revealed by Cross-Option Prices. *Journal of Futures Markets* 17:369–84.

- Vanden, J. M. 2008. Information Quality and Options. *Review of Financial Studies* 21:2635–76.
- Xing, Y., X. Zhang, and R. Zhao. 2010. What Does the Individual Option Volatility Smirk Tell Us About Future Equity Returns? *Journal of Financial and Quantitative Analysis* 45:641–62.

Figure 1 Risk-return relation

The figure shows the annualized mean realized return of the five quintile portfolios sorted by expected market beta, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into quintiles based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. We then compute the value-weighted realized return over the next month for each quintile portfolio. Exact numerical values for the realized returns and expected betas of each portfolio can be found in Table 1. The panels present the results for the different beta approaches.

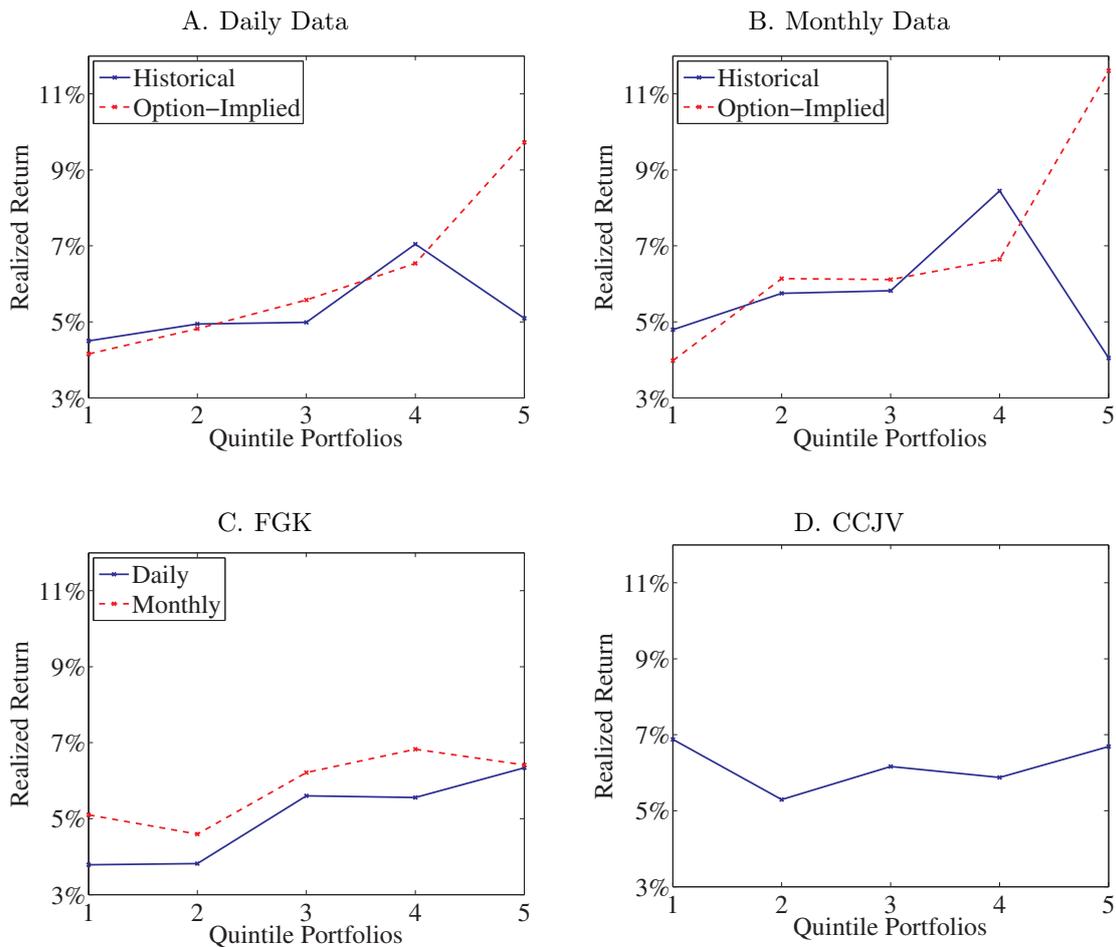


Figure 2 Portfolio beta predictions: time-series averages

The figure shows the mean prediction error in quintile portfolio betas, i.e., the difference between realized betas and expected betas of the quintile portfolios, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into quintiles based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each quintile portfolio, month, and methodology, we then compute the value-weighted realized portfolio beta as well as the value-weighted expected portfolio beta over the next year, and their difference is the prediction error. The figure presents the time-series mean of the prediction errors for each methodology. The panels present the results for the different beta approaches.

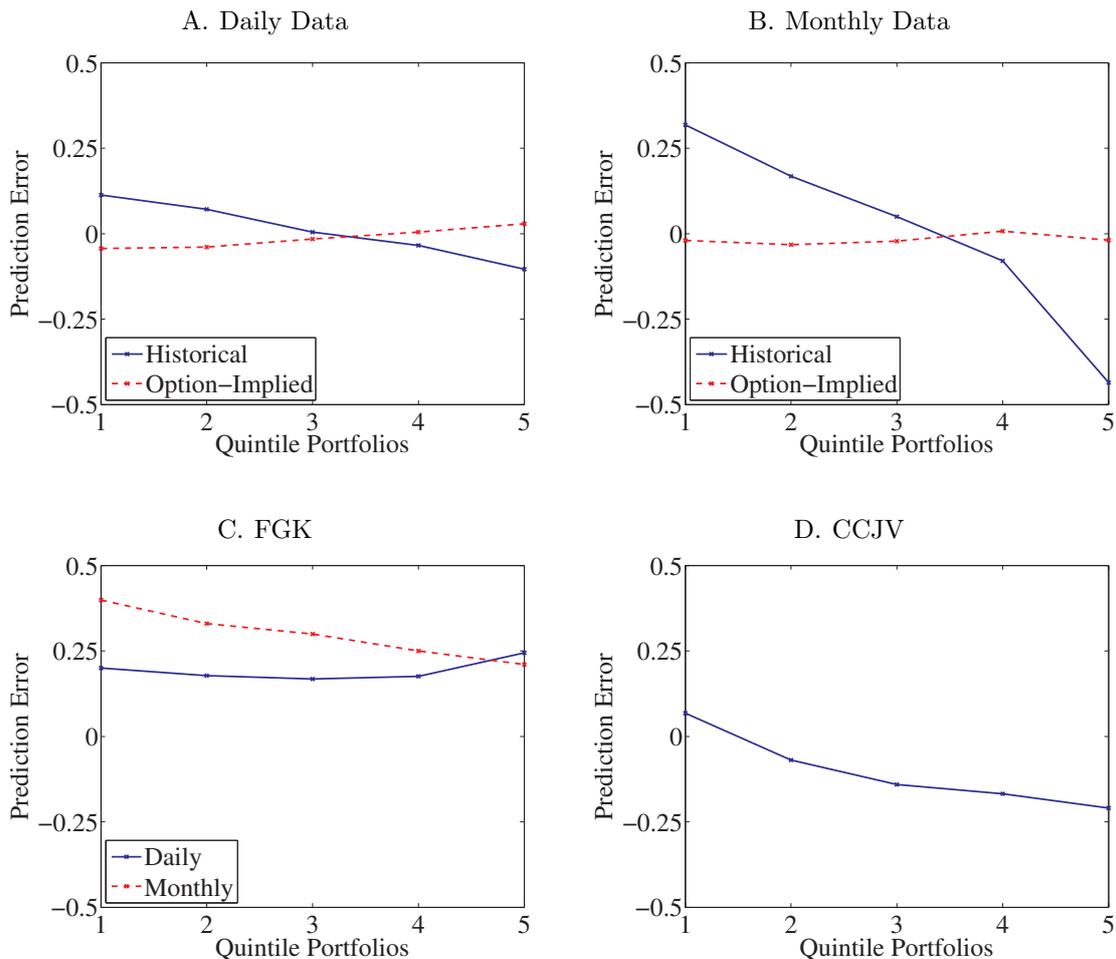


Figure 3 Portfolio beta predictions: scatterplots

The figure shows the scatterplots of expected and realized quintile portfolio betas, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into quintiles based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each quintile portfolio, month, and methodology, we then compute the value-weighted realized portfolio beta and the value-weighted expected portfolio beta over the next year. The figure plots the realized quintile portfolio betas against the expected quintile portfolio betas for all five quintiles, separately for each methodology. The panels present the results for the different beta approaches.

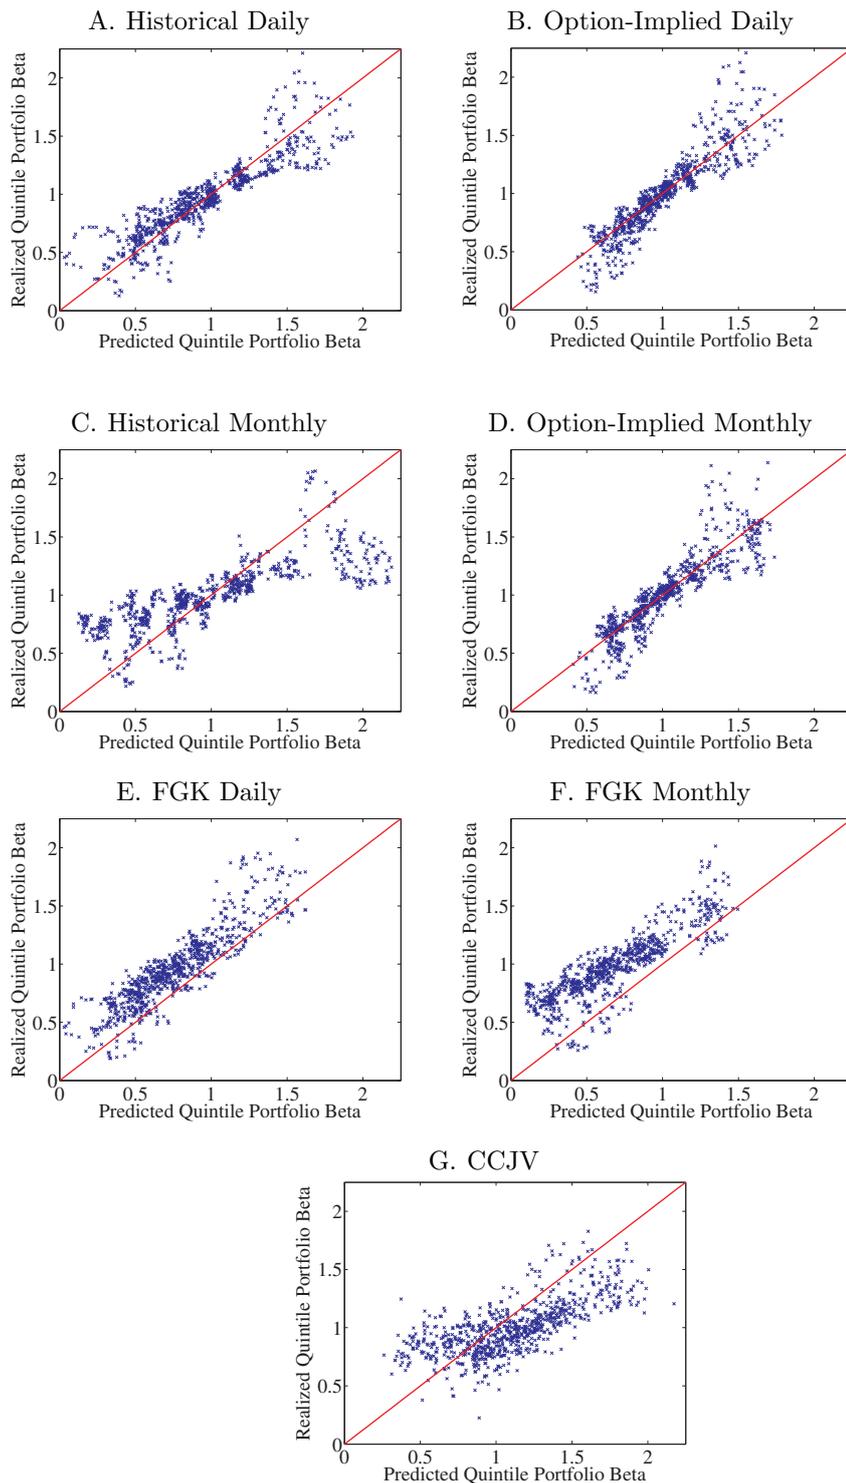


Table 1 Quintile portfolio betas and returns

	1	2	3	4	5	5-1	Monotonicity Tests	
							H_0 : increasing	H_0 : decreasing
Historical Daily								
Expected Beta	0.48	0.71	0.88	1.09	1.52	1.04	-	-
Realized Return	4.50	4.94	4.99	7.05	5.09	0.60	0.40	0.31
Option-Implied Daily								
Expected Beta	0.66	0.83	0.96	1.12	1.45	0.79	-	-
Realized Return	4.15	4.82	5.58	6.54	9.72	5.57	0.71	0.05
FGK Daily								
Expected Beta	0.40	0.59	0.72	0.88	1.19	0.78	-	-
Realized Return	3.79	3.82	5.60	5.56	6.34	2.55	0.40	0.07
Historical Monthly								
Expected Beta	0.36	0.66	0.90	1.20	1.79	1.43	-	-
Realized Return	4.79	5.75	5.82	8.45	4.05	-0.74	0.60	0.67
Option-Implied Monthly								
Expected Beta	0.65	0.83	0.98	1.15	1.49	0.84	-	-
Realized Return	3.98	6.14	6.11	6.65	11.61	7.63	0.81	0.10
FGK Monthly								
Expected Beta	0.25	0.48	0.64	0.84	1.16	0.91	-	-
Realized Return	5.10	4.59	6.21	6.83	6.42	1.32	0.34	0.12
CCJV								
Expected Beta	0.73	0.93	1.08	1.23	1.52	0.80	-	-
Realized Return	6.88	5.29	6.17	5.88	6.69	-0.19	0.15	0.18

Table 1 provides the mean expected beta and the annualized mean realized return for the five quintile portfolios sorted on expected market beta, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into quintiles based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each quintile portfolio, month, and methodology, we then compute the value-weighted expected portfolio market beta and the annualized value-weighted realized return over the next month. We only include a stock in our sorting procedure if its expected beta is available for all approaches within a certain group (Daily, Monthly, CCJV). The table reports the time-series means of the expected betas and the realized returns for each methodology. In addition, the table provides p -values, obtained from time-series block bootstrapping, for the Patton and Timmermann (2010) monotonicity test of the hypotheses for monotonically increasing and monotonically decreasing relations between expected betas and returns.

Table 2 Cross-sectional regressions

	Daily			Monthly			CCJV
	Historical	Option-Implied	FGK	Historical	Option-Implied	FGK	
10 Portfolios							
Constant	0.0097	-0.0402	-0.0033	0.0232	-0.0342	0.0125	0.0077
<i>p</i> -value Significance	0.40	0.05	0.71	0.06	0.06	0.29	0.61
Slope	0.0119	0.0686	0.0239	0.0012	0.0654	0.0171	0.0181
<i>p</i> -value Significance	0.29	0.00	0.02	0.92	0.00	0.27	0.18
25 Portfolios							
Constant	0.0041	-0.0284	-0.0180	0.0190	-0.0278	0.0102	0.0354
<i>p</i> -value Significance	0.80	0.14	0.15	0.14	0.13	0.34	0.08
Slope	0.0184	0.0554	0.0432	0.0090	0.0608	0.0232	-0.0056
<i>p</i> -value Significance	0.24	0.00	0.00	0.42	0.00	0.10	0.75
50 Portfolios							
Constant	0.0039	-0.0068	-0.0124	0.0085	-0.0135	0.0025	0.0595
<i>p</i> -value Significance	0.75	0.71	0.31	0.50	0.48	0.81	0.00
Slope	0.0205	0.0337	0.0410	0.0196	0.0460	0.0344	-0.0299
<i>p</i> -value Significance	0.08	0.05	0.01	0.08	0.01	0.01	0.04

Table 2 provides the results of the cross-sectional regressions, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into 10, 25, or 50 portfolios based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each portfolio, month, and methodology, we then compute the value-weighted expected portfolio market beta and the value-weighted realized excess return over the next month. We only include a stock into our sorting procedure if its expected beta is available for all approaches within a certain group (Daily, Monthly, CCJV). We then regress the mean realized excess return of the portfolios on a constant and the portfolios' mean expected betas, separately for each methodology. The table reports the regression coefficients, together with their *p*-values for significance using OLS standard errors.

Table 3 Market beta: summary statistics

	Daily			Monthly			CCJV
	Historical	Option-Implied	FGK	Historical	Option-Implied	FGK	
Number of Observations	79864	70845	70845	80453	71376	71376	64869
Mean Weighted	0.9967	1.0000	0.8106	0.9948	1.0000	0.6996	1.0958
Mean	0.9665	1.0317	0.7993	1.0181	1.0600	0.7023	1.1498
Standard Deviation	0.4625	0.3563	0.3807	0.6259	0.3976	0.4149	0.3926
Minimum	-0.6754	-0.4158	-0.8178	-0.8546	-0.1411	-0.6174	0.0371
Median	0.9008	0.9741	0.7468	0.9240	0.9913	0.6447	1.1258
Maximum	4.6592	5.2705	4.3565	5.1996	5.8827	4.7640	4.6247

Table 3 provides the results of the cross-sectional regressions, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into 10, 25, or 50 portfolios based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each portfolio, month, and methodology, we then compute the value-weighted expected portfolio market beta and the value-weighted realized excess return over the next month. We only include a stock into our sorting procedure if its expected beta is available for all approaches within a certain group (Daily, Monthly, CCJV). We then regress the mean realized excess return of the portfolios on a constant and the portfolios' mean expected betas, separately for each methodology. The table reports the regression coefficients, together with their p -values for significance using OLS standard errors.

Table 4 Market beta: predictability analysis

	Daily			Monthly			CCJV
	Historical	Option-Implied	FGK	Historical	Option-Implied	FGK	
10 Portfolios							
Average Mean-Squared Error	0.1733	0.1502	0.2488	0.3285	0.1579	0.3493	0.3001
<i>p</i> -value Difference to Historical	-	0.00	0.00	-	0.00	0.77	-
Percentage of Lowest MSE	10.00	90.00	0.00	0.00	90.00	10.00	-
Bottom Portfolio Prediction Error	0.1373	-0.0488	0.2110	0.3917	-0.0252	0.4485	0.1634
Top Portfolio Prediction Error	-0.1616	0.0018	0.2538	-0.7240	-0.1060	0.1792	-0.2682
Top-Bottom Portfolio Prediction Error	-0.2989	0.0506	0.0427	-1.1157	-0.0807	-0.2693	-0.4316
<i>p</i> -value for Significance	0.00	0.64	0.61	0.00	0.52	0.04	0.00
25 Portfolios							
Average Mean-Squared Error	0.1947	0.1741	0.2625	0.3555	0.1873	0.3663	0.3249
<i>p</i> -value Difference to Historical	-	0.00	0.00	-	0.00	0.82	-
Percentage of Lowest MSE	16.00	80.00	4.00	0.00	88.00	12.00	-
Bottom Portfolio Prediction Error	0.1697	-0.0418	0.2487	0.5115	-0.0013	0.5329	0.2837
Top Portfolio Prediction Error	-0.2316	-0.0901	0.2430	-1.0008	-0.3185	0.0948	-0.3856
Top-Bottom Portfolio Prediction Error	-0.4013	-0.0483	-0.0057	-1.5124	-0.3171	-0.4381	-0.6693
<i>p</i> -value for Significance	0.00	0.69	0.95	0.00	0.02	0.02	0.00
50 Portfolios							
Average Mean-Squared Error	0.2132	0.2002	0.2752	0.3795	0.2168	0.3835	0.3475
<i>p</i> -value Difference to Historical	-	0.00	0.00	-	0.00	0.90	-
Percentage of Lowest MSE	28.00	70.00	2.00	8.00	78.00	14.00	-
Bottom Portfolio Prediction Error	0.1879	-0.0206	0.2652	0.5834	0.0267	0.5888	0.3730
Top Portfolio Prediction Error	-0.3247	-0.2142	0.2118	-1.1588	-0.5348	-0.0274	-0.5708
Top-Bottom Portfolio Prediction Error	-0.5125	-0.1937	-0.0533	-1.7422	-0.5615	-0.6162	-0.9439
<i>p</i> -value for Significance	0.00	0.13	0.66	0.00	0.00	0.00	0.00

Table 4 provides the summary of the market beta predictability, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into 10, 25, or 50 portfolios based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each quintile portfolio, month, and methodology, we then compute the value-weighted expected portfolio market beta and the value-weighted realized portfolio beta over the next month. We only include a stock in our sorting procedure if its expected beta is available for all approaches within a certain group (Daily, Monthly, CCJV). First, we compute the mean squared error (MSE) between the portfolio's expected market betas and the portfolio's realized market betas over all months. In addition, we report a *p*-value, indicating if the MSE is significantly different from the MSE of the historical counterpart. Moreover, within each group (Daily, Monthly, CCJV) we compute the percentage of portfolios for which a specific beta methodology yields the lowest MSE. Second, we compute for the extreme portfolios (bottom and top) the time-series mean prediction error in betas (realized minus expected beta), the difference between the top and the bottom portfolio's prediction error, and a *p*-value, indicating if the difference is significantly different from zero, based on Newey-West standard errors with 24 lags.

Table 5 Market-neutral pairs trading

	Daily			Monthly			CCJV
	Historical	Option-Implied	FGK	Historical	Option-Implied	FGK	
Number of Observations	10,915,787	10,915,787	10,915,787	10,915,787	10,915,787	10,915,787	10,915,787
Mean Squared Error	0.0347	0.0329	0.0535	0.0675	0.0344	0.1005	0.0496
p -value Difference to Historical	-	0.01	0.00	-	0.00	0.00	-
0.25 Quantile Squared Error	0.0024	0.0024	0.0038	0.0050	0.0026	0.0095	0.0051
Median Squared Error	0.0114	0.0112	0.0178	0.0231	0.0121	0.0414	0.0222
0.75 Quantile Squared Error	0.0366	0.0357	0.0580	0.0724	0.0381	0.1190	0.0616

Table 5 provides the results of the market-neutral pairs trading application, over the sample period from January 1996 to December 2009. At the end of each month and for each available pair, we form a portfolio consisting of: a long position in one stock of the pair, a short position in the other stock, and a long or a short position in the market such that the expected market beta of the portfolio is zero, i.e., a market-neutral portfolio. Then we compute the realized market beta of this specific pair over the holding period (21 trading days) as well as the prediction error, i.e., the deviation from the zero beta expectation. We finally pool all observations, i.e., across time and portfolios, and compute the mean squared error together with a p -value for the difference to the historical beta counterpart based on standard errors clustered by time and pair following Petersen (2009).

Table 6 Quintile portfolios: option-implied characteristics

A. Variance Risk Premium						
	1	2	3	4	5	5-1
Daily						
Historical	-0.0183	-0.0172	-0.0187	-0.0202	-0.0243	-0.0059
Option-Implied	-0.0143	-0.0153	-0.0182	-0.0211	-0.0290	-0.0147
FGK	-0.0162	-0.0160	-0.0185	-0.0195	-0.0265	-0.0103
Monthly						
Historical	-0.0165	-0.0160	-0.0174	-0.0194	-0.0272	-0.0107
Option-Implied	-0.0155	-0.0153	-0.0180	-0.0213	-0.0315	-0.0160
FGK	-0.0183	-0.0154	-0.0175	-0.0196	-0.0259	-0.0076
CCJV	-0.0145	-0.0152	-0.0171	-0.0220	-0.0336	-0.0190
B. Option-Implied Skewness						
	1	2	3	4	5	5-1
Daily						
Historical	-0.3373	-0.3296	-0.3507	-0.3405	-0.3442	-0.0069
Option-Implied	-0.3479	-0.3457	-0.3545	-0.3438	-0.3279	0.0200
FGK	-0.3461	-0.3242	-0.3425	-0.3452	-0.3553	-0.0092
Monthly						
Historical	-0.3438	-0.3464	-0.3555	-0.3487	-0.3144	0.0295
Option-Implied	-0.3508	-0.3585	-0.3466	-0.3355	-0.3152	0.0356
FGK	-0.3423	-0.3477	-0.3511	-0.3444	-0.3389	0.0034
CCJV	-0.2877	-0.3324	-0.3498	-0.3688	-0.4108	-0.1231

Table 6 provides the average characteristics related to option-implied quantities for quintile portfolios sorted on the expected market beta, over the sample period from January 1996 to December 2009. At the end of each month, we sort, for each beta methodology, the stocks into quintiles based on their expected market beta. The first portfolio thereby contains the stocks with the lowest expected market betas, and the last portfolio contains the stocks with the highest expected market betas. For each quintile portfolio, month, and methodology, we then compute the value-weighted average of the individual stock variance risk premiums for all stocks in the quintile at this time (Panel A), using individual stock variance risk premiums computed according to equation (5). Similarly, we compute for each quintile portfolio, month, and methodology, the value-weighted average of the individual stock skewness for all stocks in the quintile at this time (Panel B). The table reports the time-series means of these two statistics for all quintile portfolios and the 5-1 portfolio, separately for each methodology.