

Option-Implied Measures of Equity Risk*

BO-YOUNG CHANG¹, PETER CHRISTOFFERSEN^{1,2,3}, KRIS JACOBS^{4,5}, and GREGORY VAINBERG¹

¹*Desautels Faculty of Management, McGill University,* ²*Rotman School of Management, University of Toronto,* ³*Copenhagen Business School,* ⁴*University of Houston,* and ⁵*Tilburg University*

Abstract. Equity risk measured by beta is of great interest to both academics and practitioners. Existing estimates of beta use historical returns. Many studies have found option-implied volatility to be a strong predictor of future realized volatility. We find that option-implied volatility and skewness are also good predictors of future realized beta. Motivated by this finding, we establish a set of assumptions needed to construct a beta estimate from option-implied return moments using equity and index options. This beta can be computed using only option data on a single day. It is therefore potentially able to reflect sudden changes in the structure of the underlying company.

JEL Classification: G12

1. Introduction

Many studies have demonstrated that option-implied volatility is a strong predictor of future volatility in equity markets. Classic contributions include Day and Lewis (1992), Canina and Figlewski (1993), Lamoureux and Lastrapes (1993), Christensen and Prabhala (1998), Fleming (1998), and Blair, Poon, and Taylor (2001). The predictive power of option-implied equity volatility has been confirmed recently

* Christoffersen acknowledges financial support from the Center for Research in Econometric Analysis of Time Series, CREATES, funded by the Danish National Research Foundation. Christoffersen and Jacobs thank Fonds de Recherche sur la Société et la Culture, Institut de Finance Mathématique de Montréal, and Social Sciences and Humanities Research Council for financial support. We would like to thank Gurdip Bakshi, Federico Bandi, Morten Bennedsen, Bob Dittmar, Jin Duan, Edwin Elton, Rob Engle, Andras Fulop, Steve Heston, Gerard Hoberg, Nikolay Gospodinov, Shingo Goto, Bjarne Astrup Jensen, Søren Johansen, Madhu Kalimipalli, Lawrence Kryzanowski, David Lando, Christian Lundblad, Dilip Madan, Adolfo de Motta, Jesper Rangvid, Andrei Simonov, Carsten Sørensen, Peter Norman Sørensen, Alex Triantis, Bas Werker, and participants in seminars at Copenhagen Business School, Chicago Graduate School of Business, Copenhagen, CREATES, Ivey, Maryland, McGill, National University of Singapore, Schulich, Stockholm School of Economics, Tilburg, Wilfrid Laurier, and at the American Finance Association, National Bureau of Economic Research-National Science Foundation, European Finance Association, and Bendheim Center Conferences for helpful comments. Any remaining inadequacies are ours alone.

by Busch, Christensen, and Nielsen (2008), who compare option-implied forecasts with state-of-the-art realized volatility forecasts.¹

Volatility is clearly not the only risk measure of interest. The central equity risk concept is arguably market beta, which captures the covariation of the return on an individual security with the return on the market portfolio, as approximated by a broad market index. Accurate measurement of market betas is critical for important issues such as cost of capital estimation, performance measurement, and the detection of abnormal returns.

The importance of companies' market beta in corporate finance as well as in asset pricing raises the question whether option-implied information can be used to compute and predict these betas. Interestingly, we find that option-implied volatility and skewness are good predictors of future beta, just as previous authors have found option-implied volatility to be a strong predictor of future volatility. Motivated by this finding, we establish a set of assumptions under which a company beta can be estimated from option-implied volatility and skewness measures from equity and index options.

Existing techniques for beta estimation use historical returns data. These methods thus assume that the future will be sufficiently similar to the past to justify simple extrapolation of current or lagged betas. There is widespread agreement that betas are time varying, and historical methods can easily allow for this. One popular approach uses a rolling window of historical returns to capture time variation, while other approaches model the time variation in historical betas more explicitly and in a more sophisticated fashion. However, no matter how sophisticated the modeling of the time variation in the betas, a historical method may not perform well if historical patterns in the data are unstable. The appeal of our procedure is that option prices are inherently forward looking and therefore contain information on future betas as opposed to lagged betas.

A key strength of our approach is that betas can be computed using closing prices of options observed only on a single day. This may be an important advantage when a company experiences major changes in its operating environment or capital structure, in which case historical return data do not constitute a reliable source for estimating betas. Examples include firms involved in mergers or acquisitions, reorganized firms emerging from Chapter 11, firms undertaking initial public offerings or seasoned equity offerings, as well as firms undertaking large-scale expansions and/or major changes in the composition of debt and equity.

We are not the first to propose extracting market beta from options. French, Groth, and Kolari (1983) and Buss and Vilkov (2009) combine option-implied

¹ Papers focusing on other markets include Blair, Poon, and Taylor (2001), Ederington and Guan (2002), Figlewski (1997), Jorion (1995), and Pong *et al.* (2004). See Poon and Granger (2003) and Granger and Poon (2005) for surveys.

volatility with historical correlation to improve the measurement of betas. However, these approaches rely on conventional correlation estimates from historical returns. McNulty *et al.* (2002) emphasize the problems with historical beta when computing the cost of capital and propose as an alternative the forward-looking market-derived capital pricing model, which uses option data to assess equity risk. Siegel (1995) notes the advantage of a beta based on option data and proceeds to propose the creation of a new derivative, called an exchange option, which would allow for the computation of what he refers to as “implicit” betas. Unfortunately, the exchange options discussed by Siegel (1995) are not yet traded, and therefore, his method cannot be used in practice to compute betas. While others have thus suggested various forms of option-implied betas, we suggest a new measure that exclusively uses available options data, and furthermore, we are the first to conduct a large-scale empirical study of the properties of betas based on option information.

We show that market betas can be computed without using historical correlation estimates and without the creation of a new derivative, by using prices on existing equity options and index options. Our proposed beta is computed using option-implied estimates of variance and skewness, which can be computed using the methods proposed by Bakshi and Madan (2000), Bakshi, Kapadia, and Madan (2003) (henceforth BKM), Britten-Jones and Neuberger (2000), and Carr and Madan (2001) and developed further in Jiang and Tian (2005).² These methods allow us to retrieve the moments of the underlying distributions for index options and stock options from the cross section of option prices. We then use a traditional one-factor model and express the option-implied beta as a function of the variance and the skewness of the underlying distributions.

We implement this model using option contracts for the one hundred components of the S&P100 index as well as data on S&P500 options. Daily option data are obtained from the OptionMetrics database over the period 1996–2004. We compare the performance of the option-implied beta with that of historical betas. Although the option-implied betas are computed using only one day of data, we find that they perform well compared to traditional historical beta estimates. The option-implied betas outperform the historical betas in cross-sectional regressions. For the purpose of predicting future betas, we find that the option-implied betas contain information that is not contained in historical betas. Interestingly, we find that the option-implied betas have significant predictive power for future betas that extend beyond the options’ maturities.

The empirical findings in this paper are related to some of the results in the growing literature on the pricing of equity options and, more specifically, to results on the differential pricing of equity and index options. Dennis and Mayhew (2002) and Duan and Wei (2009) find that firms with high historical market betas have higher

² See also Jackwerth and Rubinstein (1996).

negatively skewed risk-neutral distributions. This paper demonstrates that their result is to be expected because option-implied betas are higher if the underlying risk-neutral distribution is more negatively skewed. We also demonstrate that a stock's option-implied beta is partly determined by the difference between the skew of the stock's risk-neutral distribution and the skew of the index's risk-neutral distribution. BKM document differences between the skew of stock and index options. Driessen, Maenhout, and Vilkov (2009) use equity options to obtain correlations between stocks in a parametric setup.

The remainder of this paper is organized as follows: Section 2 discusses the derivation of option-implied return moments and illustrates their ability to anticipate changes in future betas using the components of the S&P100. Section 3 establishes a set of sufficient conditions for deriving an option-implied beta using equity and index options data only. Section 4 presents empirical results comparing the option-implied beta with more traditional estimates computed using daily returns. Section 5 concludes. The appendix contains a brief description of the derivations of the option-based moments, as well as a Monte Carlo study of the moment estimators.

2. Option-Implied Return Moments and Equity Risk

In this section, we show how current option-implied variance and skewness estimates from individual stock and index options are informative about the future riskiness of a stock. Accurate measurement of equity risk is critically important, and practitioners as well as academics usually approach this issue by specifying factor models and estimating the required betas. Factor models and betas are used by practitioners for a number of reasons. First, they provide a benchmark for performance measurement because they indicate the return a portfolio manager ought to have made given the risk present in his portfolio. Second, factor models provide a benchmark for the detection of abnormal returns as is, for instance, done in event studies. Third, factor models can be used to determine the required cost of capital.

We first discuss how to compute return moments from option prices. Subsequently, we discuss the data and the choices we make with respect to the implementation of these return moments. Then we document the informativeness of the moments for future equity risk.

2.1 THE LINK BETWEEN OPTION PRICES AND RETURN MOMENTS

In this subsection, we explain how to compute return moments using option prices. We employ the methods of Carr and Madan (2001) as used in BKM.³ The key result

³ See also Bakshi and Madan (2000), Britten-Jones and Neuberger (2000), Derman and Kani (1998), Jiang and Tian (2009), and Rubinstein (1994).

is that any twice differentiable payoff function can be spanned by a position in bonds, stocks, and out-of-the-money options. A brief overview of this general result is presented in Appendix A.

Let q denote the probability distribution function under the risk-neutral measure. The variance and skewness under the risk-neutral measure are defined as

$$\begin{aligned}\text{SKEW} &\equiv \frac{E^q[(R - E^q[R])^3]}{\text{VAR}^{3/2}}, \\ \text{VAR} &\equiv E^q[(R - E^q[R])^2].\end{aligned}\quad (1)$$

Using the noncentral moments, we can rewrite them as

$$\begin{aligned}\text{SKEW} &= \frac{E^q[R^3] - 3E^q[R]E^q[R^2] + 2E^q[R]^3}{\text{VAR}^{3/2}}, \\ \text{VAR} &= E^q[R^2] - E^q[R]^2.\end{aligned}\quad (2)$$

Following BKM, we define the “Quad” and “Cubic” contracts as having a payoff function equal to the squared return and cubed return, respectively, for a given horizon τ . The fair values of these contracts are

$$\begin{aligned}\text{Quad} &= e^{-r\tau} E^q[R^2], \\ \text{Cubic} &= e^{-r\tau} E^q[R^3].\end{aligned}$$

Substituting these expressions into the variance and skewness formulas in Equation (2), we get the option-implied moments

$$\text{SKEW}^{\text{OI}} = \frac{e^{r\tau} \text{Cubic} - 3E^q[R]e^{r\tau} \text{Quad} + 2E^q[R]^3}{\text{VAR}^{3/2}}, \quad (3)$$

$$\text{VAR}^{\text{OI}} = e^{r\tau} \text{Quad} - E^q[R]^2. \quad (4)$$

BKM show that under any martingale pricing measure, the Quad and Cubic contract prices can be recovered from the market prices on portfolios of out-of-the-money European calls $C(\tau, K)$ and puts $P(\tau, K)$, where K denotes the strike price and τ denotes the time to maturity.

The price of the Quad contract is

$$\text{Quad} = \int_S^\infty \frac{2(1 - \ln[\frac{K}{S}])}{K^2} C(\tau, K) dK + \int_0^S \frac{2(1 + \ln[\frac{S}{K}])}{K^2} P(\tau, K) dK, \quad (5)$$

where S is the price of the underlying stock. The price of the Quad contract can be interpreted as the forward price of volatility.⁴

The price of the Cubic contract is

$$\text{Cubic} = \int_S^{\infty} \frac{6\ln(\frac{K}{S}) - 3\ln(\frac{K}{S})^2}{K^2} C(\tau, K) dK - \int_0^S \frac{6\ln(\frac{S}{K}) + 3\ln(\frac{S}{K})^2}{K^2} P(\tau, K) dK. \quad (6)$$

BKM further show that using a third-order approximation the risk-neutral log-return mean can be approximated by

$$E^q[R] = e^{r\tau} - 1 - \frac{e^{r\tau}}{2} \text{Quad} - \frac{e^{r\tau}}{6} \text{Cubic}. \quad (7)$$

This relationship provides the last element required to compute variance and skewness from Quad and Cubic contracts.

2.2. DATA AND MOMENT ESTIMATION

We obtain option data from OptionMetrics that is a comprehensive source of high-quality historical data for the US equity and index options markets. We extract the security ID, date, expiration date, call or put identifier, strike price, best bid, best offer, and implied volatility from the option price file. For European options, implied volatilities are calculated using mid-quotes and the Black–Scholes formula. For American options, a binomial tree approach that takes into account the early exercise premium is employed. In our empirical analysis, we focus on the quotes of the one hundred stocks in the S&P100 as of December 31 2004 and the S&P500 index, for the period January 1 1996 to December 31 2004.

Interest rates are taken from the CRSP Zero Curve file, and underlying security prices are obtained through CRSP. As in Bakshi, Cao, and Chen (1997), BKM, and Jiang and Tian (2005), we use the average of the bid and ask quotes for each option contract and we filter out average quotes that are less than \$3/8. We also filter out quotes that do not satisfy standard no-arbitrage conditions. Finally, we eliminate in-the-money options because they are less liquid than out-of-the-money and at-the-money options. We eliminate put options with strike prices of more than 103% of the underlying asset price ($K/S > 1.03$), as well as call options with strike prices of less than 97% of the underlying asset price ($K/S < 0.97$).

⁴ Volatilities derived in this fashion have been studied among others by Britten-Jones and Neuberger (2000), Carr and Madan (2001), Carr and Wu (2009), and Jiang and Tian (2005). Alternatively, one could use at-the-money implied Black–Scholes volatility as an estimate of volatility.

Moments are computed by integrating over moneyness. In practice, we do not have a continuum of option prices across moneyness, and we therefore have to make a number of choices regarding implementation. We follow Carr and Wu (2009) and Jiang and Tian (2005) in imposing structure on implied volatilities. Our implementation is specifically designed to improve the quality of the integration procedure. First, as mentioned above, we limit our attention to options on the firms in the S&P100 in order to maximize the availability of strike prices and the size of the integration domain. Second, we only estimate the moments for days that have at least two out-of-the-money call prices and two out-of-the-money put prices available. Third, as in Carr and Wu (2009) and Jiang and Tian (2005), for each maturity, we interpolate implied volatilities using a cubic spline across moneyness levels (K/S) to obtain a continuum of implied volatilities. The cubic spline is only effective for interpolating between the maximum and minimum available strike price. For moneyness levels below (above) the available moneyness level in the market, we simply extrapolate the implied volatility of the lowest (highest) available strike price.

After implementing this interpolation–extrapolation technique, we are able to extract a fine grid of 1,000 implied volatilities for moneyness levels between 1% and 300%. We then convert these implied volatilities into call and put prices using the following rule: moneyness levels less than 100% ($K/S < 1$) are used to generate put prices and moneyness levels more than 100% ($K/S > 1$) are used to generate call prices. This fine grid of option prices is then used to compute the option-implied moments by approximating the Quad and Cubic contracts using trapezoidal numerical integration. It is important to note that this procedure does not assume that the Black–Scholes model correctly prices options. It merely provides a translation between option prices and implied volatilities. Also note that our implementation follows Jiang and Tian (2005) and is slightly different from the one in BKM and Dennis and Mayhew (2002). We document and discuss the benefits of our approach in more detail with the help of a Monte Carlo experiment in Appendix B.

We can in principle compute several option-implied moments for every underlying asset, one for each available option maturity. For each day, we linearly interpolate using the two contracts nearest to the 180-day maturity to get the 180-day VAR and SKEW contracts, always using one contract with maturity more than 180 days and one contract with maturity less than 180 days. The choice of a 180-day horizon is to some extent based on a trade-off between option liquidity that is largest for options with 30–90 days to maturity and the relevant horizon for firm risk, which is arguably considerably longer.

Panel A in Table I presents descriptive statistics of the option data. We report the average firm size (market capitalization in billion dollars) and the number of firms in each industry from our list of one hundred stocks. We also show the average

Table I. Descriptive statistics on option data and moment estimates

Panel A reports the average firm size and the number of firms in each industry from our list of one hundred stocks observed from 1996 through 2004. It also shows the average daily option volume, the average number of quotes each day, and the minimum number of quotes per day for call and put options separately. Panel B reports standard deviations (SD, annualized) and skewness implied from 180-day options averaged by industry. The average *ex post* realized betas are also reported for reference. Panel C reports the option-implied standard deviation and skewness as well as *ex post* realized beta by firm-size decile.

Panel A: Market capitalization and option data by industry								
Industry	Number of firms	Average firm size	Call options			Put options		
			Average daily volume	Average daily # quotes	Minimum daily # quotes	Average daily volume	Average daily # quotes	Minimum daily # quotes
Energy	6	48.33	549.73	6.44	5.09	250.17	6.43	4.97
Materials	6	21.67	231.26	5.95	4.43	135.78	6.24	4.73
Industrials	13	44.15	445.22	6.62	4.84	265.62	7.32	4.50
Consumer discretionary	17	22.18	304.50	6.26	4.82	195.54	6.39	4.98
Consumer staples	12	61.75	403.13	6.20	4.93	234.82	6.81	4.55
Health care	10	66.20	638.83	6.97	5.09	336.18	8.27	5.63
Financials	13	58.46	681.83	8.00	5.19	470.97	9.21	5.77
Information technology	14	81.21	1,849.48	9.54	6.49	903.35	10.50	5.72
Telecommunications services	4	68.25	894.44	6.60	5.75	544.98	6.72	5.78
Utilities	5	11.40	117.14	4.88	4.00	66.63	5.24	4.35
S&P500 index			4,652.29	27.08	10.00	8,731.36	35.92	9.00

Panel B: Moments sorted by industry				Panel C: Moments sorted by firm size			
Industry	SD 180-day	SKEW 180-day	Average <i>ex post</i> beta	Firm size	SD 180-day	SKEW 180-day	Average <i>ex-post</i> beta
Energy	0.4022	−0.4782	0.7997	1 (low)	0.4499	−0.3588	1.0976
Materials	0.3455	−0.4368	0.9380	2	0.3925	−0.4370	0.8827
Industrials	0.3343	−0.4711	0.9011	3	0.3398	−0.4641	0.7781
Consumer discretionary	0.3724	−0.4440	0.9505	4	0.3787	−0.4402	0.9303
Consumer staples	0.2887	−0.4932	0.5914	5	0.3196	−0.4595	0.7693
Health care	0.3534	−0.5142	0.8358	6	0.3447	−0.5082	0.9252
Financials	0.3590	−0.6116	1.2329	7	0.3961	−0.5870	1.2341
Information technology	0.4860	−0.5112	1.4601	8	0.3966	−0.5289	1.2009
Telecommunications services	0.4266	−0.5034	1.0839	9	0.3136	−0.5870	0.7665
Utilities	0.3010	−0.4964	0.5228	10 (high)	0.3576	−0.6245	1.1880
S&P500 index	0.2357	−1.4622					

daily option volume, the average number of quotes each day, and the minimum number of quotes per day for call and put options separately. The option trading volumes are much higher for index options than for individual equity options. It is also clear that there is substantial industry variation in option trading volumes, with volumes for IT and telecom much higher than those for materials and utilities.

Descriptive statistics of the option-implied moments are reported in Panels B and C of Table I. Panel B reports the average standard deviation (annualized) and skewness implied from 6-month options averaged by industry. The average *ex post* betas from returns are reported for reference. Panel C reports the option-implied standard deviation and skewness as well as *ex post* beta averaged by firm-size decile.

Figure 1 displays some properties of the moments over time. In Panel A, we plot the value of the volatility index (VIX) index during 1996–2004. Panel B plots




Figure 1. Option-implied volatility and skewness. (A) The value of the VIX index during 1996–2004. (B) The S&P500 option-implied volatility for the same period. (C) The average of the option-implied volatility for the S&P100 components. (D) S&P500 option-implied skewness and (E) the average of the option-implied skewness for the S&P100 components.

the S&P500 option-implied volatility for the same period. Panel C plots the average of the option-implied volatility across the one hundred firms in the S&P100. Panel D plots S&P500 option-implied skewness, and Panel E plots the average of the option-implied skewness for the S&P100 components.

It is reassuring that our measure of S&P500 option-implied volatility in Panel B is very highly correlated with the VIX in Panel A. The average of the option-implied volatilities for the S&P100 stocks in Panel C is also highly correlated with the VIX. From Panels D and E, we can draw some conclusions regarding skewness. First, the S&P500 option-implied skewness exceeds the average option-implied skewness of the S&P100 firms. This finding is consistent with the results in BKM. Note that here and in the rest of the paper, we refer to more negative skewness as higher skewness. Second, the correlation between option-implied skewness and option-implied volatility for the S&P500 is -0.06 , and the correlation between the average option-implied skewness and average option-implied volatility for the S&P100 components is 0.05 . Finally, while the option-implied volatility for the S&P500 is strongly correlated with the average option-implied volatility for the S&P100 components (0.70), the correlation between the two skewness time series is virtually zero.

2.3 OPTION-IMPLIED MOMENTS PREDICT FUTURE (*EX POST*) BETA

We are now ready to assess the informational content of the option-implied moments for future equity risk. There is an ongoing discussion on the appropriate choice of factor model and the empirical performance of alternative factor models and the appropriate measure of equity risk. We derive our results in the context of a particular factor model, the capital asset pricing model (CAPM). There is evidence in the literature that some multifactor models can improve on the performance of the CAPM, but this issue remains hotly debated.⁵ We merely note that the academic literature remains divided about the performance of the CAPM and that the CAPM is still the factor model most often used in financial practice. In our empirical application, the CAPM performs relatively well.

We measure the *ex post* market beta using the realized betas proposed by Andersen *et al.* (2006). While Andersen *et al.* use high-frequency data, we use daily data to compute the covariance between the market and the equity as well as the variance of the market. This implementation is inspired by Schwert's (1989) construction of a measure of realized volatility. In our main results, we

⁵ See Campbell, Lo, and MacKinlay (1997), Cochrane (2001), Jagannathan and McGrattan (1995), Fama (1991), Ferson (1995, 2004), and the references therein for an overview of this extensive debate. See also Ferson and Korajczyk (1995) and MacKinlay (1995).

use 6 months of daily returns on the S&P500 and on the individual equities to compute *ex post* beta for the 6-month or 180-day period. This construction is consistent with the choice of 180-day maturity options discussed above because presumably 180-day options are most informative for the distribution of returns over the next 180 days. We also conduct robustness exercises using 1-year and 2-year *ex post* betas.

In Table II, we report on forecasting regressions that use various permutations of the option-implied moments. In all cases, coefficients of multiple correlation are reported.⁶ Because of space constraints, we do not report results for individual companies, but instead we focus on industry results and firm-size deciles. Results for individual companies are available from the authors on request. In Panel A, we regress the industry-averaged *ex post* beta estimated from the subsequent 6 months on today's option-implied variance (VAR) and skewness (SKEW) averaged across industry.⁷

Table II. Multiple correlation from regressing *ex-post* betas on option-implied moments

We first create *ex-post* realized beta estimates using rolling 6-month regressions on daily data. We then regress the *ex-post* beta estimated on the subsequent 6 months on today's option-implied variance (VAR) and skewness (SKEW). Subscript "m" denotes moments implied from S&P500 options, and subscript "i" denotes option-implied moments from individual equity options.

Regressors	VAR _m	SKEW _m	VAR _i	SKEW _i	VAR _m , SKEW _m	VAR _i , SKEW _i	VAR _m , SKEW _m , VAR, SKEW _i
Panel A: By industry							
Energy	0.232	0.009	0.113	0.324	0.232	0.383	0.410
Materials	0.259	0.113	0.057	0.546	0.276	0.548	0.607
Industrials	0.324	0.263	0.177	0.505	0.404	0.522	0.643
Consumer discretionary	0.045	0.207	0.215	0.516	0.209	0.540	0.573
Consumer staples	0.175	0.164	0.559	0.108	0.249	0.672	0.754
Health care	0.057	0.319	0.474	0.071	0.329	0.482	0.578
Financials	0.404	0.115	0.315	0.542	0.413	0.597	0.644
Information technology	0.013	0.326	0.598	0.087	0.328	0.601	0.687
Telecommunications	0.194	0.104	0.623	0.098	0.227	0.643	0.679
Utilities	0.124	0.207	0.219	0.287	0.241	0.363	0.489
Panel B: By firm size							
1 (small)	0.065	0.229	0.317	0.464	0.234	0.557	0.631
2	0.175	0.051	0.086	0.422	0.179	0.427	0.475
3	0.193	0.298	0.438	0.287	0.367	0.494	0.551
4	0.278	0.157	0.200	0.264	0.310	0.340	0.516
5	0.115	0.075	0.469	0.427	0.142	0.612	0.685
6	0.109	0.026	0.317	0.266	0.110	0.427	0.522
7	0.069	0.207	0.552	0.015	0.223	0.553	0.630
8	0.029	0.362	0.439	0.267	0.366	0.467	0.583
9	0.070	0.199	0.493	0.063	0.216	0.494	0.603
10 (big)	0.104	0.207	0.187	0.120	0.239	0.227	0.277

⁶ The conclusions do not change when using the square roots of the adjusted R^2 .

⁷ Correlations are overall somewhat higher when we first regress on a firm-by-firm basis and subsequently average the correlation coefficients.

Subscript “m” denotes moments implied from S&P500 options, and subscript “i” denotes option-implied moments from individual equity options.

The results in Panel A are of substantial interest and allow for several conclusions. When using all four option-implied moments in forecasting beta (in the last column), the resulting coefficient of multiple correlation is quite high in all sectors and particularly in the consumer staples, IT, and telecom sectors. Second, when considering the four individual moments as predictors, the individual skewness, $SKEW_i$, typically has the highest correlation with *ex post* beta. Third, when comparing the correlation from the two market-based moments with the correlation from the two firm-specific moments, the latter are always larger.

Consider now Panel B where we regress the average *ex post* beta in each firm-size decile on the option-implied moments averaged across firms in the relevant decile as well. The rightmost column shows that the correlation when using all four moments is smallest for the biggest firms and quite large in all the other size categories. Only $SKEW_i$ seems to show a clear pattern across firm size: correlations tend to be larger for smaller firms and smaller for larger firms.

The findings in Table II are obtained using 6-month (180-day) *ex post* betas. A natural question is if the options with 180-day maturities have explanatory power for *ex post* betas computed over longer time horizons. We repeated the results in Table II using 1-year and 2-year *ex post* betas. We did not investigate *ex post* betas computed over horizons more than 2 years because this excessively reduces the sample period in view of the fact that the option data are available only for 1996–2004. Results for 1-year and 2-year *ex post* betas are not reported but available on request. The results are very surprising. Not only do the moments constructed from options with 180-day maturities have substantial predictive power for 1- and 2-year *ex post* betas but also the multiple correlation coefficients are almost uniformly higher than the ones for 180-day *ex post* betas in Table II. We discuss the implications of this finding in more detail below.

The main conclusion from Table II is that option-implied moments seem to contain a lot of information about future equity risk as captured by the *ex post* realized beta. This finding begs the question of how to sensibly combine the information in the moments. We suggest ways to do so in the next section.

3. Option-Implied Betas

In this section, we show how the univariate variance and skewness from individual equity and market index options can be aggregated into a single stock-specific risk measure. We show that under certain assumptions, the risk measure can be interpreted as a beta. We also analyze the potential biases in the option-implied risk measure.

3.1 FACTOR MODELS AND BETA ESTIMATION

The computation of betas for the CAPM is an issue that is hotly debated. In several classic applications of the CAPM, betas were computed by running a regression of stock returns on market returns, using returns for the past 60 months.⁸ This technique is still used in many academic and practitioner approaches to beta estimation. The choice of sixty lagged returns reflects the tension implicit in using historical returns to compute betas: on the one hand, one would like to use as long a history as possible to obtain more precise estimates, but, on the other hand, one does not want to use older returns data because it is likely that the beta changes over time.

Over the last two decades, the modeling of the time variation in the market beta has taken center stage in the CAPM literature. There is by now widespread consensus that beta is time varying, although it is less clear how much this time variation helps the model's empirical performance out of sample. A number of different econometric techniques are available to model the time variation in the beta.⁹ While modeling the time variation in the beta is helpful, it does not address the criticism that this type of measurement is backward looking, much like the simple regressions that use 60 months of lagged returns. But many applications of factor models, such as computing the cost of capital, are inherently forward looking. We now show that it is possible to compute forward-looking betas using the information embedded in option prices.

3.2. ESTIMATING MARKET BETAS USING OPTION DATA: EXISTING APPROACHES

The standard estimates of beta derived from historical market returns and individual equity returns can be written as

$$\beta_i = \frac{\text{COV}_{i,m}}{\text{VAR}_m} = \text{CORR}_{i,m} \left(\frac{\text{VAR}_i}{\text{VAR}_m} \right)^{1/2},$$

where the moments are computed from times series of historical returns. This is also the method used when computing *ex post* beta in Section 2.3 above.

⁸ See, for instance, Black, Jensen, and Scholes (1972), Fama and MacBeth (1973), and Fama and French (1992, 1993, 1996).

⁹ On modeling time variation in betas, see Bollerslev, Engle, and Wooldridge (1988), Blume (1971, 1975), Bos and Newbold (1984), Cochrane (2001), Ferson (1995, 2004), Ferson and Harvey (1999), Ferson, Kandel, and Stambaugh (1987), Ghysels (1998), Harvey (1989), and Jagannathan and Wang (1996). See Bauer *et al.* (2008) for a Bayesian approach. Lewellen and Nagel (2006) advocate estimating conditional betas using daily or weekly returns and short data windows. Companies such as BARRA provide investors and risk managers with time-varying estimates of market betas.

French, Groth, and Kolari (1983) suggest introducing option information in a hybrid estimation method, where the correlation between equity and market return, $\text{CORR}_{i,m}$, is estimated from historical returns but where the variances are implied from options. We can write

$$\beta_i^{\text{HYBR}} = \text{CORR}_{i,m} \left(\frac{\text{VAR}_i^{\text{OI}}}{\text{VAR}_m^{\text{OI}}} \right)^{1/2},$$

where VAR_i^{OI} and VAR_m^{OI} are the option-implied variances implied from option prices on individual equity options and on index options. We will refer to this as the hybrid beta estimate below. French, Groth, and Kolari (1983) empirically investigate the performance of the hybrid beta using a single day of data, and index option data were not available for their sample. Below we provide a more extensive empirical investigation of their approach that has also been investigated and extended by Buss and Vilkov (2009).

Note that the hybrid estimator still requires historical data and so cannot be used for firms where little or no history is available. It is also problematic to use it when the nature of the firm has changed dramatically, for example, following a merger or a major reorganization or expansion when the correlation between firm and market is likely to change. McNulty *et al.* (2002) cite the sensitivity of the correlation estimates from historical data as a key detractor for corporate practitioners when using the CAPM for cost of capital computations.

Siegel (1995) circumvents the computation of historical correlations by assuming the existence of an option to exchange shares in firm i for units of the market index. From the accounting identity dollars per share = (dollars per index unit) $\times$ (index units per share), one can derive the beta of firm i as follows:

$$\beta_i = \frac{\text{VAR}_i^{\text{OI}} + \text{VAR}_m^{\text{OI}} - \text{VAR}_X^{\text{OI}}}{2\text{VAR}_m^{\text{OI}}},$$

where VAR_X^{OI} is the option-implied variance implied from the option to exchange index units for equity shares. Note that this is the method used to imply covariance estimates from range-based volatility estimates in currency markets in Brandt and Diebold (2006). Unfortunately, since exchange options are not yet traded, the implicit beta cannot be computed in practice.

Below we argue that betas similar in spirit to Siegel's (1995) "implicit" betas can be computed without the creation of a new derivative by using prices on existing equity and index options. We will refer to these betas as "option-implied" betas. First, we derive an estimator of beta that relies on univariate moments only. Second, we implement the moment-based beta with the option moments derived above.

3.3 COMPUTING BETAS FROM THE MOMENTS OF STOCK AND INDEX RETURNS

Following BKM, we assume that the log return on stock i follows a single-factor model of the form

$$R_i = \alpha_i + \beta_i R_m + \epsilon_i, \quad (8)$$

where the market return R_m has mean μ_m . The idiosyncratic shock ϵ_i has zero mean and is assumed to be independent of the market return R_m .

Consider now the moments of the return distribution. To simplify notation, we do not provide time subscripts for the moments, but empirically there will be a different estimate for the moments at each point in time. We use the conventional definition of skewness as the third central moment divided by the standard deviation to the third power. Using the single-factor return structure with an independent idiosyncratic term, we can write the skewness of R_i as

$$\begin{aligned} \text{SKEW}_i &= \frac{E[\alpha_i + \beta_i R_m + \epsilon_i - E[\alpha_i + \beta_i R_m + \epsilon_i]]^3}{\text{VAR}_i^{3/2}} \\ &= \frac{\beta_i^3 E[R_m - \mu_m]^3 + E[\epsilon_i]^3}{\text{VAR}_i^{3/2}} \\ &= \frac{\beta_i^3 \text{SKEW}_m \text{VAR}_m^{3/2} + \text{SKEW}_{\epsilon,i} \text{VAR}_{\epsilon,i}^{3/2}}{\text{VAR}_i^{3/2}}. \end{aligned} \quad (9)$$

From this, we can write beta as

$$\beta_i = \left[\left(\frac{\text{SKEW}_i}{\text{SKEW}_m} \right) \left(\frac{\text{VAR}_i}{\text{VAR}_m} \right)^{3/2} - \left(\frac{\text{SKEW}_{\epsilon,i}}{\text{SKEW}_m} \right) \left(\frac{\text{VAR}_{\epsilon,i}}{\text{VAR}_m} \right)^{3/2} \right]^{1/3}. \quad (10)$$

In order to be able to implement an estimator using only moments of R_i and R_m , we now make the identifying assumption that the skewness of the idiosyncratic shock is zero, $\text{SKEW}_{\epsilon,i} = 0$. Note that Duan and Wei (2009) make a similar assumption. We can then solve for the market beta of stock i to be simply

$$\beta_i^{\text{MM}} = \left(\frac{\text{SKEW}_i}{\text{SKEW}_m} \right)^{1/3} \left(\frac{\text{VAR}_i}{\text{VAR}_m} \right)^{1/2}, \quad (11)$$

where the “MM” superscript indicates that the market beta is computed using the moments of the return distribution.¹⁰

¹⁰ Fouque and Kollman (2009) suggest an implied volatility regression approach to implementing Equation (11).

Note that expression (11) indicates that the skewness of the market return has to be nonzero for the market beta to be well defined in this setup. Recall from Figure 1 that the market skewness is always negative and large in magnitude empirically. However, this requirement implies that multivariate normality cannot be used as a sufficient condition for the existence of the CAPM. Instead, we have to resort to quadratic utility, see, for instance, the discussion in Huang and Litzenberger (1988) and Cochrane (2001). Alternatively, Kraus and Litzenberger (1976, 1983) demonstrate that we obtain the CAPM if all risk assets have linear characteristic lines, even if the market return is characterized by an asymmetric distribution.

Note also that setting $\text{SKEW}_{\epsilon,i} = 0$ is a sufficient but not necessary condition for using option information in beta estimation. Thus, even if the $\text{SKEW}_{\epsilon,i} = 0$ assumption is violated, return moments may be useful for forecasting, as we saw in Table II. Moreover, if prior information is available about the value of $\text{SKEW}_{\epsilon,i}$, then this can be used to estimate beta from Equation (10) rather than Equation (11).

Figure 2 quantifies the bias $\beta_i^{\text{MM}} - \beta_i$ as a function of $\text{SKEW}_{\epsilon,i}$. On the horizontal axis, we let $\text{SKEW}_{\epsilon,i}$ take values on a grid from -0.5 to $+0.5$. Motivated by the average of the empirical values in Table I above, we set $\beta_i = 1$, $\text{VAR}_m = 0.24^2$, and $\text{SKEW}_m = -1.44$. $\text{VAR}_{\epsilon,i}$ is set equal to the difference between the average firm variance and the market variance which is roughly 0.24^2 . VAR_i and SKEW_i can




Figure 2. Bias from idiosyncratic skewness

We plot the bias $\beta_i^{\text{MM}} - \beta_i$ as a function of $\text{SKEW}_{\epsilon,i}$. We set $\beta_i = 1$, $\text{VAR}_m = 0.24^2$, and $\text{SKEW}_m = -1.44$, based on sample averages. $\text{VAR}_{\epsilon,i}$ is set equal to the difference between the average firm variance and the market variance which is roughly 0.24^2 . VAR_i and SKEW_i can then be computed from $\text{VAR}_i = \beta_i^2 \text{VAR}_m + \text{VAR}_{\epsilon,i}$ and Equation (9), which gives the inputs needed for computing β_i^{MM} in Equation (11).

then be computed from $\text{VAR}_i = \beta_i^2 \text{VAR}_m + \text{VAR}_{\epsilon,i}$ and Equation (9) that gives the inputs needed for computing β_i^{MM} in Equation (11). Figure 2 shows that the bias amounts to 10% of the true beta when the idiosyncratic skewness is close to -0.5 or $+0.5$. The figure also shows that the bias is roughly linear in the idiosyncratic skewness. Assuming that SKEW_m is negative, which is the case empirically, then β_i^{MM} will be biased downward if $\text{SKEW}_{\epsilon,i}$ is positive and vice versa.

It is not straightforward to further verify the appropriateness of the $\text{SKEW}_{\epsilon,i} = 0$ assumption. We note that if there is a systematic asymmetry pattern across the idiosyncratic distribution of firms, we would expect the moment-based betas to be systematically higher or lower than the historical betas. This is not the case as we will see in Figure 4 below. However, we do observe that moment-based betas are slightly more centered around one than historical betas. This finding may be due to the assumption of a symmetric idiosyncratic return component, if the idiosyncratic skewness is negative for low beta stocks and positive for high beta stocks.

At this point, Equation (11) is simply a method to compute a market beta using moments. The moments can be computed using either historical data or option-implied information. We next analyze the use of option-implied information to compute Equation (11).

3.4 COMPUTING BETAS FROM OPTION-IMPLIED MOMENTS

Using the option-implied moments from Equations (3) and (4) as well as the moment-based beta estimate in Equation (11), we can compute the option-implied beta of stock i with data on individual stocks and the index. We get

$$\beta_i^{\text{OI}} = \left(\frac{\text{SKEW}_i^{\text{OI}}}{\text{SKEW}_m^{\text{OI}}} \right)^{1/3} \left(\frac{\text{VAR}_i^{\text{OI}}}{\text{VAR}_m^{\text{OI}}} \right)^{1/2}. \quad (12)$$

Recall from the conventional beta computations above that

$$\beta_i^{\text{HYBR}} = \text{CORR}_{i,m} \left(\frac{\text{VAR}_i}{\text{VAR}_m} \right)^{1/2}.$$

Therefore, in Equation (12), the term $\left(\frac{\text{SKEW}_i^{\text{OI}}}{\text{SKEW}_m^{\text{OI}}} \right)^{1/3}$ captures the option-implied correlation. As discussed above, French, Groth, and Kolari (1983) construct betas that combine option-implied volatilities with historical correlations in an effort to improve beta measurement. They have

$$\beta_i = \text{CORR}_{i,m} \left(\frac{\text{VAR}_i^{\text{OI}}}{\text{VAR}_m^{\text{OI}}} \right)^{1/2}.$$

Our option-implied beta can be thought of as a logical extension of their idea.

The moments computed from options data are risk-neutral moments. The question arises as to how a beta computed from these moments is related to that computed from physical moments? Fortunately, BKM provide us with useful guidance on the potential bias arising from our use of option-implied moments as proxies for physical moments. Consider first the relationship between physical and risk-neutral variance. BKM show that for the market index,

$$\text{VAR}_m^q \approx \text{VAR}_m^p (1 - \gamma \text{SKEW}_m^p \sqrt{\text{VAR}_m^p}),$$

where γ is the coefficient of relative risk aversion and the superscripts p and q denote the physical and risk-neutral measure, respectively. If we assume a similar relationship for the individual equity, we can form the ratio used in our beta computation

$$\text{VAR}_i^q / \text{VAR}_i^p \approx \frac{\text{VAR}_i^p / \text{VAR}_m^p (1 - \gamma \text{SKEW}_i^p \sqrt{\text{VAR}_i^p})}{(1 - \gamma \text{SKEW}_m^p \sqrt{\text{VAR}_m^p})}.$$

Figure 3 quantifies the bias $\beta_i^{\text{OI}} - \beta_i$ as a function of γ for three different values of the idiosyncratic skewness $\text{SKEW}_{\epsilon,i}$. On the horizontal axis, we let γ take values on a grid from 1 to 4. We again set $\beta_i = 1$, $\text{VAR}_m^p = 0.24^2$, and $\text{SKEW}_m^p = -1.44$. $\text{VAR}_{\epsilon,i}$ is set equal to the difference between the average firm variance and the market variance which is roughly 0.24^2 , and SKEW_i can then be computed from $\text{VAR}_i = \beta_i^2 \text{VAR}_m + \text{VAR}_{\epsilon,i}$ and Equation (11) that gives the inputs needed for computing β_i^{OI} . Figure 3 provides a number of insights. First, the bias is of course zero when the idiosyncratic skewness is zero (dot-dashed line) and the coefficient of relative risk aversion is zero. Second, the bias is decreasing with the coefficient of relative risk aversion. Third, recall from Figure 2 that a negative idiosyncratic skewness creates a positive bias in beta. Thus, when the idiosyncratic skewness is negative, the two biases to some extent cancel each other out (solid line). Fourth, when the idiosyncratic skewness is positive the bias is the largest. Fortunately, a negative idiosyncratic skewness due to bankruptcy risk may be the most empirically relevant.

Consider next the relationship between physical and risk-neutral skewness. BKM show that for the market index,

$$\text{SKEW}_m^q \approx \text{SKEW}_m^p - \gamma (\text{KURT}_m^p - 3) \sqrt{\text{VAR}_m^p}.$$

Assuming again a similar relationship for the individual equity, we get

$$\text{SKEW}_i^q / \text{SKEW}_i^p \approx \frac{[\text{SKEW}_i^p - \gamma (\text{KURT}_i^p - 3) \sqrt{\text{VAR}_i^p}]}{[\text{SKEW}_m^p - \gamma (\text{KURT}_m^p - 3) \sqrt{\text{VAR}_m^p}]}.$$




Figure 3. Bias from risk-neutral variance

We plot the bias $\beta_i^{OI} - \beta_i$ as a function of the coefficient of relative risk aversion for three different levels of idiosyncratic skewness, $SKEW_{\epsilon,i}$. We set $\beta_i = 1$, $VAR_m = 0.24^2$, and $SKEW_m = -1.44$, based on sample averages. $VAR_{\epsilon,i}$ is set equal to the difference between the average firm variance and the market variance which is roughly 0.24^2 . VAR_i and $SKEW_i$ can then be computed from $VAR_i = \beta_i^2 VAR_m + VAR_{\epsilon,i}$ and Equation (9), which gives the inputs needed for computing β_i^{OI} .

If the excess kurtosis is close to zero in individual stocks and in the market returns, this is approximately equal to $SKEW_i^p / SKEW_m^p$. Kurtosis is difficult to estimate reliably from returns as well as options, and we therefore do not attempt to assess the bias arising from nonzero excess kurtosis here.

Even if the moments are different under the two measures, the option-implied beta estimates may be useful for making assessments about beta. This issue also arises in the literature on volatility forecasting. Even though the implied volatility from options is a risk-neutral volatility measure, empirical researchers continue to use implied volatility for volatility estimation because its forecasting performance is very good. We refer the reader to Britten-Jones and Neuberger (2000) for further theoretical discussion. As an excellent recent example of this approach, Busch, Christensen, and Nielsen (2008) find that the implied volatility from options works better for forecasting future volatility than the state-of-the-art model based on lagged realized volatility (from intraday data) only. This is true in foreign exchange, equity, and bond markets. This strong empirical finding obtains in spite of the fact that the presence of a stochastic volatility risk premium theoretically ought to reduce the forecasting power of implied volatility.

We thus argue that the usefulness of the option-implied beta is largely an empirical issue. The (unobserved) differences between the risk-neutral and the

physical moments will translate into empirically measurable estimation and forecasting errors, and the size of these forecasting errors will determine the method's usefulness.

3.5 OPTION-IMPLIED BETAS AND OPTION-IMPLIED STATE PRICE DENSITIES

Several papers have used index option prices to construct option-implied state price densities. See, for example, Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (2000), Jackwerth (2000), and Bliss and Panigirtzoglou (2004).

The absence of arbitrage guarantees the existence of a stochastic discount factor, M , with the property that

$$E[M(1 + R_i)] = 1.$$

Assuming the existence of a risk-free rate, R_f , we can write the expected excess return on asset i as

$$E[R_i - R_f] = -(1 + R_f)\text{Cov}(M, R_i).$$

So note that even if we have the stochastic discount factor implied from index options, we still need to construct an estimate of the covariance between the stochastic discount factor and the return on each individual stock in order to get the desired estimate of the expected return on the stock. The stochastic discount factor alone does not provide all the information required to compute the expected returns on individual stocks.

It is not obvious how to obtain alternative estimates of the required covariance using only index and individual equity option prices. Using only this information is crucial because our goal is to construct estimates of expected returns using only one day of index and individual equity option prices. Our approach essentially assumes that the stochastic discount factor is linear in the market return so that the single-factor CAPM model obtains. In that case, the estimation of the expected return on the individual stock requires an estimate of its β with the market return that we obtain using options on the index and on individual stocks as outlined above.

4. Empirical Results on Option-Implied Betas

In this section, we first discuss, for all firms in the S&P100 index, the estimates of option-implied betas averaged over time. We then give a detailed discussion of the conditional performance of the option-implied betas using three firms, namely Walt Disney Corporation, Exxon Mobil Corporation, and Verizon Communications. Next we investigate the cross section of asset returns using option-implied and

traditional beta estimates. Finally, we compare the ability of option-implied and traditional historical beta estimates to anticipate changes in future beta. Throughout we also analyze the empirical performance of the hybrid beta proposed by French, Groth, and Kolari (1983).

4.1 OPTION-IMPLIED AND *EX POST* BETA: UNCONDITIONAL EVIDENCE

As mentioned above, we can in principle compute several option-implied betas for every underlying asset, one for each available option maturity. However, we choose to focus on a fixed maturity of 180 days. This choice reflects a trade-off between option liquidity, which is largest for options with 30–90 days to maturity, and the relevant horizon for firm risk, which is arguably considerably longer.

For each day, we linearly interpolate using the two contracts nearest to the 180-day maturity to get the 180-day VAR and SKEW contracts, always using one contract with maturity more than 180 days and one contract with maturity less than 180 days. Using these contracts, we obtain estimates of the 180-day option-implied beta β^{OI} using Equation (12) for each day t .

As a first assessment of the option-implied betas, Figure 4 graphically compares the time-averaged option-implied beta with the time-averaged *ex post* beta for each firm. Each “×” marks the average option-implied beta and average *ex post* beta for one of the S&P100 components. The dashed line is the 45-degree line, and the solid line represents the regression line of *ex post* beta on option-implied beta. Note that the regression fit is remarkably good at 67.5%. The solid line has a slope of 1.14 and an intercept of -0.212 . On average, option-implied betas are slightly more centered around one than the *ex post* betas. However, recall that raw beta estimates are adjusted toward one in many applications of historical beta. See, for instance, Damodaran (1999, chapter 4), Bodie, Kane, and Marcus (2007, p. 284), and the references therein for more on the motivation for these types of adjustments. If we apply this sort of adjustment to the realized *ex post* betas in Figure 4, the resulting regression line is very close to the 45-degree line.

We conclude that the correspondence across firms between option-implied betas and *ex post* betas is high. This is quite remarkable as the option-implied beta uses only one day of option data, and the *ex post* beta is computed using daily returns over a 6-month period. Encouraged by this finding, we proceed by analyzing the cross-sectional and forecasting performance of the option-implied betas. We start by providing a detailed discussion of the option-implied beta dynamics for three of the firms in our S&P100 sample. These results are very different from the unconditional evidence discussed here and emphasize the conditional performance of the beta estimates.




Figure 4. Mean *ex post* beta versus option-implied beta

We plot time averages of *ex post* beta for each of the S&P100 components versus the time averages of option-implied beta. The averages are computed for the period January 1 1996 to December 31 2004. The solid line is the regression line of mean *ex post* beta on mean option-implied beta. The slope of the regression is 1.14, and the intercept is -0.212 .

4.2 OTHER BENCHMARK BETAS

When studying the conditional performance of the option-implied betas, we have to formulate historical benchmark betas. In implementing historical benchmark betas, a number of choices have to be made. For an overview, see, for instance, Damodaran (1999). For our purpose, two important choices are the length of the estimation period and the frequency of the return data. With respect to the choice of data frequency, we are constrained by the fact that we do not have an extensive time series of option data available. We thus estimate historical betas using daily returns to construct the most powerful test possible of the option-implied beta.¹¹

With respect to the length of the estimation period, we implement historical beta in a number of different ways. Because we use 180-day option-implied beta β^{OI} , the use of 180 daily returns to compute historical beta is an obvious alternative. On the other hand, many academic studies of the CAPM have followed the approach of

¹¹ The use of daily betas has some disadvantages, see, for instance, Scholes and Williams (1977). However, in the event study literature, daily data and short estimation windows have been used extensively. See MacKinlay (1997) for a review. Several recent cross-sectional studies have also started using daily betas, typically constructed using estimation windows less than 5 years. See, for example, Ang, Chen, and Xing (2006), Barberis, Shleifer, and Wurgler (2005), and Lewellen and Nagel (2006).

Black, Jensen, and Scholes (1972) and Fama and MacBeth (1973) and computed historical betas using returns for the past 60 months. Industry providers of betas typically use estimation periods between 2 and 5 years (see Damodaran, 1999, p. 88; Bodie, Kane, and Marcus, 2007). When investigating the robustness of our results with respect to the estimation window, we find that using historical returns estimated using 1–5 years yields much worse results than using 6 months (180 days) of daily returns. If we interpret the 180-day historical betas as estimates of the conditional beta, as in Lewellen and Nagel (2006), this indicates that the conditional historical betas perform better than the unconditional ones obtained using longer historical windows. Please note that we do not always report historical beta results for all estimation windows because of space constraints.

4.3 THREE CASE STUDIES

We cannot discuss the results for all hundred S&P100 components in detail because of space constraints, but we provide a detailed discussion of the forecasting performance for three of the components of the index: Disney, Exxon, and Verizon. In Section 4.5, we discuss results for all hundred stocks and provide more details about the analysis.

Figure 5 plots the option-implied skewness and volatility for these three companies. Comparing with Figure 1, it is clear that the volatility for all three companies is positively correlated with S&P500 volatility. In fact, the correlations are 0.71, 0.78, and 0.48, respectively. Skewness patterns are more complex and harder to relate to S&P500 skewness. Verizon skewness substantially increases in the second half of 1998, at a time when S&P500 skewness is also large, but in fact the overall correlation between the two series is -0.01 . Disney skewness increases at the end of 1997 when S&P500 skewness increases, but the overall correlation is 0.01 . It is difficult to visually identify covariation patterns between Exxon and S&P500 skewness; the overall correlation is 0.10 .

The left column of panels in Figure 6 presents the time path of betas for Disney. We compare the option-implied beta with our proxy for the *ex post* “realized” beta, which is computed as the covariance divided by variance for *ex post* daily returns over the 180-day period, depicted in the top row of Figure 6. There is substantial variation in the *ex post* betas, and both the option-implied beta and the 180-day historical beta do an excellent job at capturing this variation. The 180-day historical beta does slightly better overall. The 5-year historical beta is much too smooth to capture the variation in the 180-day *ex post* beta.

The middle column of Figure 6 presents a similar analysis for Exxon. Once again the 5-year historical beta does not perform well. The 180-day historical beta substantially outperforms the option-implied beta. However, closer inspection of the




Figure 5. Option-implied moments for Disney, Exxon, and Verizon

We plot the time series of the option-implied moments needed for computing the option-implied beta for Walt Disney Corporation (DIS), Exxon Mobil Corporation (XOM), and Verizon Communications (VZ). The other moments needed are the option-implied volatility and skewness for the S&P500 in Figure 1. We use moments for 180-day returns in all cases. The moments are calculated at a daily frequency for the period January 1 1996 to December 31 2004.

figures indicates that this result is largely due to the patterns in *ex post* beta during 2000 and 2001, when the *ex post* beta becomes very small and even negative. A negative beta is widely regarded as unrealistic and therefore indicative of a problem with our proxy for the *ex post* beta. In our opinion, this suggests that it is important to look beyond the simple correlation between the *ex post* beta and the forecasts. In the case of Exxon, the 180-day historical beta substantially outperforms the option-implied beta when judged by correlation but that is only because the 180-day historical beta is equivalent to the lagged *ex post* beta and therefore ends up being unrealistically small or negative in the period 2000–01. The time path for the option-implied betas is much more plausible. See also Lewellen and Nagel (2006) for negative estimates of historical betas obtained using short data windows.

The right column of Figure 6 presents our findings for Verizon. In this case, the option-implied beta substantially outperforms both historical betas, as well




Figure 6. Realized, option-implied, and historical betas: Disney, Exxon, and Verizon

We plot the time series of different types of betas for Disney, Exxon, and Verizon for the period January 1, 1996 to December 31, 2004. The top row plots the time series of *ex post* 180-day realized betas (β^{REAL}). On each day, a 180-day *ex post* beta is computed using the returns for the following 180 days. The second row plots the time series of option-implied betas. On each day, a 180-day option-implied beta (β^{OI}) is computed. The third row plots the time series of 180-day historical betas, computed using the previous 180 days of returns. The bottom row plots the time series of 5-year historical betas, computed using the previous 1,825 days of returns.

as the 1-year and 2-year historical beta forecasts that are not included because of space constraints. Note how the option-implied beta nicely captures the increase in the *ex post* beta toward the end of the sample, even though it overshoots in 2002.

The results from the three case studies encourage a full-scale investigation using the one hundred S&P100 components. However, they also illustrate that the *ex post* realized beta computed over a relatively short period such as 180 days can produce questionable estimates of the true but unknown beta. This caveat is important when interpreting the results below. The three case studies also illustrate that the option-implied betas contain a certain amount of high-frequency noise when estimated from only one day of options data. We could reduce this noise by estimating

the option-implied betas using several days of options. However, in order to be as transparent as possible, all the option-implied betas below are estimated using just one day of options data.

4.4 THE CROSS-SECTION OF RETURNS

Perhaps, the most important implications of factor models are cross-sectional. We therefore investigate the cross-sectional performance of the option-implied betas, even though our sample of one hundred firms observed over 9 years is far from ideal to study the cross section. We first regress average 180-day and then average 365-day excess returns on average beta in the spirit of Black, Jensen, and Scholes (1972).

Panel A of Table III presents the first set of results, with year-by-year cross-sectional regressions and 180-day excess returns. We show results for three different beta estimates: option-implied beta, historical beta using 180 days of returns, and hybrid beta using option-implied volatilities but correlations from 180 days of returns.

Consider first the regressions using historical betas. Note that the slope coefficient is positive for the 1996–99 years, negative for the 2000–02 years, and then again positive for the 2003–04 years. The market movements surrounding the bursting of the dot-com bubble in 2000–02 are clearly difficult to reconcile with the single-factor model. This pattern is also observed when using option-implied beta. We boldface the R^2 for option-implied beta and hybrid beta when it is larger than the R^2 from historical beta. Notice that the option-implied beta almost always has a larger R^2 than the historical beta. This is not the case for the hybrid beta.¹²

Panel B of Table III reports results for the overall 1996–2004 period. Given the negative slopes on beta during 2000–02, we also report results for three subperiods. For the overall sample, the slope is positive but small and the fit is better when using historical beta. Figure 7(A) shows the mean excess 180-day return scattered against beta when using option-implied and historical betas, respectively. The two beta estimates perform very similarly on the first subsample, whereas the option-implied beta performs best on the most recent subsample, perhaps because the options markets are getting more liquid over time. The 2000–02 period appears as an anomaly with a negative relationship between returns and beta for both beta estimates.

In Panels C and D of Table III, we repeat the analysis using excess returns over 365 days instead of 180 days. The results are very similar. Option-implied beta again has an R^2 that is higher than that from historical beta in most years. The magnitudes of the R^2 from the option-implied betas are similar as well.

¹² We recognize that having a high R^2 when the slope estimate is negative is not necessarily desirable.

Table III. Cross-sectional regression of average realized excess return on average predicted beta

For the one hundred firms in the S&P100, we cross-sectionally regress average excess return on the average predicted beta. The excess return is defined as the actual equity return less the 3-month *T*-bill rate. In Panel A, cross-sectional regressions are run for each year of data. In Panel B, we report the full-sample result as well as three subperiods corresponding to before, during, and after the 3 years with negative slope coefficients on historical beta in Panel A. Boldface font indicates that the R^2 is larger than that of the corresponding regression with historical beta.

	Option-implied beta			Historical beta			Hybrid beta		
	Slope	<i>t</i> -statistics	R^2	Slope	<i>t</i> -statistics	R^2	Slope	<i>t</i> -statistics	R^2
Panel A: 180-Day excess return, year-by-year regressions									
1996	0.07	1.25	0.02	0.13	2.72	0.07	0.14	1.86	0.04
1997	0.26	4.05	0.14	0.11	2.17	0.05	0.18	2.32	0.05
1998	0.35	4.24	0.16	0.17	3.36	0.10	0.27	3.60	0.12
1999	0.71	7.14	0.34	0.29	5.76	0.26	0.43	5.41	0.23
2000	-0.21	-3.64	0.12	-0.17	-4.88	0.20	-0.20	-4.73	0.19
2001	-0.16	-5.06	0.21	-0.08	-4.40	0.17	-0.10	-4.27	0.16
2002	-0.11	-2.87	0.08	-0.05	-1.21	0.02	-0.05	-0.99	0.01
2003	0.12	4.32	0.17	0.12	3.70	0.13	0.12	3.20	0.10
2004	0.03	0.72	0.01	0.01	0.45	0.00	-0.01	-0.29	0.00
Panel B: 180-Day excess return, subperiod regressions									
1996-99	0.41	9.75	0.49	0.27	9.88	0.50	0.42	9.67	0.49
2000-02	-0.16	-5.93	0.27	-0.09	-4.75	0.19	-0.10	-4.42	0.17
2003-04	0.13	4.66	0.19	0.08	2.99	0.09	0.10	2.85	0.08
1996-2004	0.05	2.36	0.05	0.05	3.95	0.14	0.07	3.86	0.13
Panel C: 365-Day excess return, year-by-year regressions									
1996	0.39	3.11	0.09	0.48	4.28	0.16	0.57	3.08	0.09
1997	0.67	4.39	0.17	0.22	1.80	0.03	0.39	2.05	0.04
1998	1.08	6.22	0.29	0.51	4.73	0.19	0.80	4.88	0.20
1999	1.08	5.23	0.22	0.53	5.43	0.23	0.75	4.90	0.20
2000	-0.44	-5.10	0.21	-0.31	-5.77	0.26	-0.37	-5.73	0.26
2001	-0.26	-5.02	0.21	-0.13	-4.08	0.15	-0.16	-3.98	0.14
2002	0.04	0.69	0.01	0.13	2.33	0.05	0.14	2.15	0.05
2003	0.13	3.04	0.09	0.08	1.50	0.02	0.06	1.07	0.01
2004	0.14	1.77	0.03	0.07	1.28	0.02	0.05	0.62	0.00
Panel D: 365-Day excess return, subperiod regressions									
1996-99	0.91	8.94	0.45	0.61	8.96	0.45	0.92	8.78	0.44
2000-02	-0.25	-5.53	0.24	-0.12	-3.62	0.12	-0.13	-3.39	0.11
2003-04	0.18	3.55	0.12	0.11	2.33	0.05	0.12	1.81	0.03
1996-2004	0.15	3.35	0.10	0.14	4.69	0.18	0.18	4.58	0.18

Figure 7(B) reports the mean excess 365-day return scattered against the average option-implied beta (left column) and against the average historical beta (right column). Figure 7(B) clearly looks very similar to Figure 7(A) where 180-day returns were used. Option-implied and historical beta estimates perform very similarly on the first subsample, they both perform poorly on the second subsample, and the option-implied beta performs best on the third subsample.

It is not straightforward to relate our results to the available literature for a number of reasons. First, we use a short sample because of the limited availability of option




Figure 7. (A) Mean excess return versus mean beta, 180-day returns. We plot the average excess 180-day return for each of the S&P100 components against the average option-implied beta (left column) and the average historical 180-day beta (right column). The averages are computed for three subperiods: 1996–99 (top row), 2000–02 (middle row), and 2003–04 (bottom row). The solid lines represent the regression of excess return on beta.

data. Second, we use individual stock returns rather than portfolios.¹³ Third, we use 180- and 365-day returns, while most of the literature uses monthly returns. Fourth, we investigate a relatively small cross-sectional sample. Fifth, our sample is limited to large cap stocks, and therefore the cross-sectional application has limited power to test a given model. For the purpose of this study, we merely conclude that the option-implied beta has significant explanatory power in the cross section. This is

¹³ It is well known that assessing the performance of the CAPM using individual stocks is subject to testing problems related to measurement error in the betas. See Black, Jensen, and Scholes (1972) for an early discussion of this problem. We limit ourselves to one hundred stocks because we want to use liquid option contracts.




Figure 7. (B) Mean excess return versus mean beta, 365-day returns. We plot the average excess 365-day return for each of the S&P100 components against the average option-implied beta (left column) and the average historical 180-day beta (right column). The averages are computed for three subperiods: 1996–99 (top row), 2000–02 (middle row), and 2003–04 (bottom row). The solid lines represent the regression of excess return on beta.

quite remarkable given that the option-implied betas are estimated using only one day of option data.

4.5 ANTICIPATING *EX POST* REALIZED BETA

We started the empirical investigation in Section 2.3 by assessing the ability of option-implied moments to track changes in future beta. We then derived a new beta estimator using the option-implied moments. In this section, we close the loop and assess the ability of option-implied beta to anticipate changes in future beta.

Consider Table IV. We run multivariate regressions of *ex post* beta on historical beta and either option-implied or hybrid beta, for each firm in the S&P100.¹⁴ We report for each industry the proportion of stocks where the estimated coefficient for OI or HYBR, respectively, is statistically significantly estimated. OI refers to option-implied beta, and HYBR refers to using option-implied volatilities but historical return correlations. Panel A reports results when the 180-day historical beta is used as the benchmark. Panels B and C repeat the exercise using the 1- and 5-year historical betas, respectively, as a benchmark. In each panel, we report on three different forecasting horizons: 180 days, 1 year, and 2 years. This is inspired by the somewhat surprising empirical results discussed in the context of Table II. While the performance of the option-implied moments for 180-day *ex post* betas in Table II is impressive, the performance improves even more when using 1- and 2-year *ex post* betas.

First consider Panel A. Of the one hundred stocks we investigate, 32% have statistically significant loadings on the option-implied beta at the 180-day forecast horizon, even though the historical beta is included in the regression. When the horizon of interest increases from 180 to 365 and 730 days, this percentage improves to 38% and 56%, respectively. Panel A also reports results for hybrid beta, where the correlation used for computing the hybrid beta is computed using either 180 days, 1 year, or 5 years of daily returns. Clearly the hybrid beta performs well. When forecasting 180-day *ex post* betas, the percentage of significant estimates is 32% when computing correlation using 180 days of data. The percentages are 26% and 39% when correlations are computed using 1- and 5-year windows, respectively.

The performance of the option-implied as well as hybrid betas improves when forecasting 1- and 2-year *ex post* betas. These results are consistent with the results for forecasting using return moments reported in Table II, and they are somewhat surprising. One potential explanation is the potential problems with *ex post* betas evident from Figure 6. As discussed before, *ex post* betas often take on implausible values when constructed using a short horizon. This is less often the case when using longer horizons, and this may explain the improved forecasting performance of the option-implied betas.

Panels B and C use 1- and 5-year historical betas, respectively, as a benchmark, instead of 180-day historical betas as in Panel A. The historical betas computed using longer windows have less forecasting power, as is evident from Figure 6, and as a result the option-implied and hybrid betas perform relatively better overall.

¹⁴ When implementing forecasting regressions, one potential problem is overlapping data (see Christensen and Prabhala, 1998). Two potential solutions are to use instrumental variable techniques and to use nonoverlapping data. Instrumental variable regressions were noninformative because of the poor quality of the instruments. Regression results were very similar when using a sample of non-overlapping data based on 30-day betas.

Table IV. Proportion of stocks for which option-implied or hybrid betas predict realized beta

We run a time series regression of realized beta (180-, 365-, and 730-day horizons) on the historical beta and OI beta for each firm in the S&P100. We then report for each industry the proportion of stocks for which the coefficient for OI is significant at 95% confidence level based on Newey–West *t*-statistics with lag 180. We repeat the procedure by replacing the OI beta with HYBR betas computed over horizons of 180, 365, and 1,825 days. OI refers to option-implied beta and HYBR refers to using option-implied volatilities but historical return correlations. We report the results when the historical betas are computed over 180-, 365-, and 1825-day horizons in Panels A, B, and C, respectively.

Industry	180-Day realized beta				365-Day realized beta				730-Day realized beta			
	OI	HYBR180	HYBR365	HYBR1825	OI	HYBR180	HYBR365	HYBR1825	OI	HYBR180	HYBR365	HYBR1825
Panel A: 180-Day historical beta as benchmark												
Energy	0.33	0.17	0.17	0.80	0.33	0.33	0.17	0.80	0.50	0.17	0.17	0.80
Materials	0.67	0.67	0.67	0.50	0.83	0.67	0.83	0.33	1.00	1.00	1.00	0.50
Industrials	0.38	0.31	0.38	0.31	0.46	0.54	0.31	0.54	0.77	0.62	0.46	0.77
Consumer discretionary	0.41	0.59	0.41	0.38	0.41	0.53	0.53	0.38	0.65	0.65	0.53	0.63
Consumer staples	0.17	0.08	0.00	0.33	0.17	0.08	0.00	0.50	0.50	0.33	0.08	0.58
Health care	0.20	0.30	0.00	0.56	0.20	0.30	0.30	0.44	0.30	0.20	0.20	0.56
Financials	0.31	0.23	0.15	0.11	0.38	0.23	0.15	0.11	0.62	0.46	0.38	0.22
Information technology	0.29	0.43	0.36	0.31	0.43	0.36	0.50	0.31	0.36	0.36	0.21	0.23
Telecommunications	0.25	0.00	0.50	0.00	0.25	0.25	0.25	0.00	0.75	0.50	0.25	0.33
Utilities	0.20	0.00	0.00	1.00	0.40	0.00	0.20	1.00	0.20	0.00	0.00	1.00
All hundred stocks	0.32	0.32	0.26	0.39	0.38	0.35	0.33	0.42	0.56	0.45	0.34	0.54
Panel B: 365-Day historical beta as benchmark												
Energy	0.33	0.50	0.17	0.80	0.33	0.67	0.17	0.80	0.50	0.67	0.17	0.80
Materials	0.67	0.50	0.67	0.17	0.67	0.83	0.67	0.00	1.00	0.83	0.67	0.33
Industrials	0.38	0.23	0.23	0.31	0.62	0.38	0.23	0.46	0.85	0.69	0.46	0.85
Consumer discretionary	0.41	0.65	0.59	0.44	0.53	0.65	0.76	0.44	0.65	0.65	0.71	0.69
Consumer staples	0.17	0.17	0.08	0.58	0.25	0.33	0.08	0.67	0.42	0.42	0.17	0.58
Health care	0.40	0.20	0.30	0.56	0.30	0.30	0.40	0.56	0.40	0.20	0.30	0.56
Financials	0.31	0.08	0.38	0.22	0.38	0.38	0.38	0.33	0.62	0.38	0.46	0.44
Information technology	0.29	0.43	0.36	0.31	0.36	0.64	0.57	0.23	0.50	0.29	0.43	0.23
Telecommunications	0.25	0.25	0.50	0.00	0.25	0.50	0.50	0.00	0.75	0.50	0.50	0.33
Utilities	0.60	0.60	0.00	0.75	0.60	0.60	0.00	0.75	0.20	0.20	0.20	0.75
All hundred stocks	0.36	0.35	0.34	0.41	0.43	0.51	0.41	0.43	0.59	0.48	0.43	0.57

(Continued)

Table IV. (Continued)

Industry	180-Day realized beta				365-Day realized beta				730-Day realized beta			
	OI	HYBR180	HYBR365	HYBR1825	OI	HYBR180	HYBR365	HYBR1825	OI	HYBR180	HYBR365	HYBR1825
Panel C: 1825-Day historical beta as benchmark												
Energy	0.80	0.60	0.40	0.60	0.40	0.40	0.40	0.60	0.80	0.40	0.40	0.80
Materials	0.83	1.00	1.00	0.83	1.00	1.00	1.00	0.67	1.00	1.00	1.00	0.83
Industrials	0.54	0.54	0.46	0.23	0.54	0.54	0.54	0.54	0.69	0.69	0.77	0.77
Consumer discretionary	0.63	0.69	0.50	0.56	0.63	0.63	0.69	0.75	0.75	0.69	0.75	0.75
Consumer staples	0.25	0.75	0.67	0.75	0.25	0.67	0.58	0.75	0.50	0.75	0.58	0.75
Health care	0.44	0.67	0.44	0.78	0.44	0.56	0.00	0.78	0.11	0.00	0.00	0.78
Financials	0.44	0.33	0.44	0.33	0.44	0.56	0.56	0.56	0.56	0.67	0.56	0.56
Information technology	0.62	0.62	0.54	0.62	0.62	0.69	0.62	0.62	0.77	0.77	0.77	0.69
Telecommunications	0.33	0.33	0.67	0.33	0.67	0.67	0.67	0.33	0.67	0.67	0.67	0.67
Utilities	0.75	1.00	1.00	0.25	0.75	1.00	1.00	0.25	0.25	1.00	0.75	0.25
All hundred stocks	0.54	0.64	0.57	0.54	0.54	0.64	0.58	0.63	0.62	0.66	0.63	0.71

The improvements provided by option-implied and hybrid betas differ across sectors. For the industrials, consumer discretionary, materials, and IT sectors, they perform well. Their performance is generally weaker for the consumer staples and health care industries.

In Table V, we cross-sectionally regress the improvement in correlation offered by option-implied and hybrid betas on various firm characteristics. We only report results obtained using the 180-day historical return as the benchmark as this is the best performing historical beta. Each regressor is normalized by dividing by its standard deviation. Note that for all horizons and for both option-implied and hybrid betas, the distance of average *ex post* beta from one is negative and significant. Thus, option-implied betas have relative difficulty when the *ex post* beta is far from unity. This effect is particularly strong for option-implied beta and was evident already in Figure 4 where on average the option-implied beta tended to be closer to one than *ex post* beta. Note again that the approach commonly used for historical betas that shrinks the betas toward unity may lessen this effect.

The mean *ex post* beta variable is significant in two of three option-implied beta cases, suggesting that option-implied beta performs better for higher beta stocks. Firm size is significantly negative in two of three option-implied columns, indicating that option-implied beta performs better for the smaller firms in the sample. Recall, however, that all the firms we consider are large caps. The option volume is positive in two of the three option-implied columns. This may suggest that estimates of the option-implied moments will improve as option markets become more liquid. We therefore expect that as option market liquidity continues to improve, the added value of using option-implied information to compute beta will increase.

Table V. Improvement in correlation regressed on firm characteristics

We cross-sectionally regress the improvement in correlation on various firm characteristics. HYBR refers to the 180-day hybrid beta. Each regressor is normalized by dividing by the standard deviation. Bold face font indicates a t-statistic significant at the 5% level.

Regressor	180-Day realized beta		365-Day realized beta		730-Day realized beta	
	OI	HYBR	OI	HYBR	OI	HYBR
Constant	-0.336	0.002	-0.249	0.020	0.102	0.142
Firm size	-0.197	-0.067	-0.159	-0.050	-0.116	-0.043
Option volume	0.169	0.075	0.128	0.044	0.087	0.026
Mean (REALB)	0.127	-0.006	0.112	0.011	0.011	-0.017
Mean (REALB) - 1	-0.163	-0.043	-0.156	-0.045	-0.153	-0.060
SD (REALB)	0.002	0.024	-0.005	0.012	0.027	0.017
R ²	0.35	0.13	0.26	0.09	0.15	0.10

4.6 COMBINING THE INFORMATION IN HISTORICAL AND OPTION-IMPLIED BETAS

So far we have focused on comparing betas computed from option prices with those computed from returns. In this section, we instead ask if there is scope for combining the two types of beta. They are computed using very different information sets, and so it is likely that the option-based betas complement the information in traditional return-based betas.

In order to assess the potential benefits of combining betas, we investigate multivariate regressions of *ex post* realized beta on the two kinds of betas. The multivariate forecasting regression is

$$\beta_{t,t+H}^{\text{REAL}} = \theta_1 + \theta_2 \beta_t^{\text{OI}} + \theta_3 \beta_t^{\text{HIST}} + u_{t,t+H}, \quad (13)$$

where β_t^{REAL} is *ex post* realized beta, β_t^{HIST} is the (best performing) 180-day historical beta, and β_t^{OI} is the option-implied beta computed using option prices on day t only.

The results using a horizon H of 180, 365, and 730 days are reported in Table VI. We run pooled panel regressions of the *ex post* beta on the forecasted betas, either by industry or by size decile. For the industry-based results in Panel A, the coefficient on historical beta is positive and significant virtually everywhere. When using a horizon of 180 days for *ex post* beta, the option-implied beta coefficient is positive and significant in five cases. Point estimates are large but usually smaller than the loadings on the historical betas, indicating that the information in the historical beta is more important. Looking next at the 365-day horizon, the option-implied beta is positive and significant in six cases. The results for the 730-day horizon are similar: seven significant option-implied betas. The results across industries largely confirm the findings from Table IV: the option-implied betas perform best for the materials, industrials, and consumer discretionary sectors.

These results suggest that the option-implied beta is complementary to the information in historical beta for predicting future beta. Once again, we obtain the interesting result that options contain important information for forecasting *ex post* betas computed over horizons that exceed the options' maturities.

In Panel B, both the historical beta and the option-implied betas are positive and significant in virtually every case. Unlike the case of the industry results, few patterns emerge as a function of firm size.

In summary, option-implied betas complement historical betas. Option-implied betas are computed using a very different information set than the one used for historical betas. This makes them useful in combination with the historical beta. We emphasize again that the option-implied beta is computed using only one day of data, and its performance is therefore quite remarkable.

Table VI. Multivariate beta prediction panel regressions

We run multivariate regressions of *ex post* beta on forecasted beta. We consider three horizons for the *ex post* beta: 180, 365, and 730 days. We include two betas in each regression: option-implied beta (OI) and historical beta estimated on the past 180 daily returns (HIST). Newey–West *t*-statistics with lag 180 are in parentheses. We compute *t*-statistics by running pooled panel regressions. In Panel A, we pool by industry and in Panel B by firm-size decile. Statistical significance at the 5% level is indicated in bold.

	REALB (180)				REALB (365)				REALB (730)			
	Constant	OI	HIST	<i>R</i> ²	Constant	OI	HIST	<i>R</i> ²	Constant	OI	HIST	<i>R</i> ²
Panel A: Industry												
Energy	0.179	0.219	0.475	0.28	0.294	0.242	0.347	0.23	0.468	0.305	0.114	0.19
Materials	0.195	0.286	0.491	0.54	0.230	0.330	0.430	0.63	0.263	0.401	0.341	0.65
Industrials	0.319	0.230	0.424	0.30	0.331	0.258	0.393	0.39	0.371	0.266	0.350	0.40
Consumer discretionary	0.438	0.188	0.344	0.25	0.488	0.190	0.297	0.28	0.591	0.176	0.222	0.26
Consumer staples	0.217	−0.004	0.634	0.43	0.229	0.040	0.560	0.42	0.218	0.138	0.422	0.34
Health care	0.378	−0.019	0.563	0.33	0.398	0.026	0.479	0.33	0.430	0.129	0.294	0.23
Financials	0.293	0.033	0.724	0.56	0.372	0.022	0.663	0.52	0.466	0.006	0.596	0.49
Information technology	0.546	0.184	0.447	0.33	0.650	0.143	0.404	0.31	0.859	0.055	0.334	0.22
Telecommunications	0.112	0.046	0.859	0.58	0.192	0.026	0.791	0.60	0.381	0.009	0.628	0.43
Utilities	0.117	0.176	0.573	0.43	0.194	0.233	0.411	0.36	0.292	0.308	0.202	0.32
Panel B: Firm size												
1 (small)	0.324	0.277	0.426	0.43	0.369	0.269	0.399	0.50	0.448	0.252	0.360	0.48
2	0.293	0.197	0.453	0.30	0.394	0.195	0.376	0.28	0.538	0.224	0.221	0.23
3	0.153	0.117	0.683	0.55	0.195	0.137	0.623	0.56	0.258	0.163	0.545	0.51
4	0.098	0.162	0.726	0.60	0.166	0.153	0.654	0.60	0.263	0.175	0.528	0.51
5	0.228	0.104	0.605	0.43	0.273	0.119	0.549	0.45	0.306	0.178	0.444	0.42
6	0.116	0.140	0.714	0.62	0.155	0.153	0.663	0.64	0.220	0.182	0.572	0.61
7	0.310	0.112	0.645	0.50	0.358	0.112	0.601	0.52	0.468	0.048	0.567	0.48
8	0.293	0.189	0.566	0.47	0.362	0.172	0.515	0.46	0.484	0.128	0.445	0.40
9	0.331	− 0.042	0.614	0.38	0.347	0.015	0.519	0.36	0.380	0.105	0.344	0.25
10 (big)	0.239	0.167	0.630	0.56	0.319	0.158	0.565	0.53	0.439	0.151	0.463	0.44

5. Concluding Remarks

Market betas are one of the most important concepts in the practice and theory of finance, and for many interesting applications of market betas, such as computing the cost of capital, out-of-sample performance is key. Currently, market betas are obtained by using regression techniques on historical data. Many historical implementations use a simple rolling regression approach, while other approaches allow more explicitly for time-varying betas. However, no matter how sophisticated the approach, historical betas implicitly assume that the past offers a good guide to the future.

This paper presents a radically different approach that extracts betas from option data. Because option data contain information about the future, this approach is inherently forward looking. The approach is inspired by the literature on volatility forecasting, where a number of authors have compared the forecasting performance of implied volatility with that of more traditional historical methods. The strength of our approach is that betas can be computed using a single cross section of option data, which may be an important advantage when a company experiences (potential) changes in its operating environment or capital structure.

We test our method using a very conservative approach by investigating how it compares with historical methods on average, including stable periods. We find that the option-implied estimates perform relatively well. Option-implied betas that were extracted for equities with liquid options often outperform the historical market beta in predicting the future beta in the following period. In some cases, combining historical betas with option-implied betas further improves results. We also provide a detailed study of the hybrid beta method proposed by French, Groth, and Kolar (1983). Hybrid betas that combine option-implied volatilities with historical correlations perform very well. Option-implied and hybrid betas also explain a sizeable amount of the cross-sectional variation in expected returns.

Much remains to be done. The computation of the option-implied beta uses moments extracted from options data. While we have made full use of recent innovations in the implementation of these procedures, and added some innovations ourselves, it may be that the out-of-sample performance of the option-implied betas can be improved through more efficient estimation of moments. Alternatively, it may be possible to better estimate option-implied correlation using the methods of Carr and Madan (2000). Also, we have not provided an explanation for the relatively poor performance of option-implied betas for stocks with very small *ex post* betas, and the differences in performance across industries also merit further investigation. Another remaining question is if option-implied beta performs even better in situations and time periods that can *a priori* be labeled as inherently

unstable. It may also prove interesting to investigate the optimal combination of option-implied and historical betas, which is of course related to the previous question. Finally, we note that it may be possible to compute option-implied betas that are estimated for a particular purpose using the relevant statistical loss function, for instance, in a portfolio model (Granger, 1969).

Our technique can be used for a variety of applications, such as the detection of abnormal returns in event studies, and the uncovering of abnormal returns for portfolio management. The very definition of the word “event” indicates that some aspects of the firm or the firm’s environment change, and therefore it will prove interesting to contrast the results obtained using option-implied betas with those obtained using historical betas. One particularly promising application is the computation of the cost of capital for newly merged companies.

The main focus in this paper has been on forecasting 180-day *ex post* betas, which are relevant for certain applications such as abnormal returns. For other applications, such as cost of capital calculations, longer horizon betas may be needed. We plan to investigate the performance of option-implied betas in this context by using long-term equity anticipation securities as well as option contracts with longer maturities traded on non-US markets. Indeed, our option-implied beta approach allows for the computation of a complete term structure of beta for each company so long as the options data are available.

We compute option-implied moments and betas using option prices on a given day. While this is the most obvious and transparent initial approach to investigating the method’s merits, the performance of the option-implied betas may be improved by adjusting these betas using a pre-determined rule or by smoothing betas and/or moments using information extracted from option prices on other days. The optimal use and optimal smoothing of information contained in option prices is certainly worthy of further study.

Finally, we have compared and combined the option-implied betas with simple historical betas from daily data. Option-implied betas could alternatively be compared and combined with other methods. Many alternatives exist, including realized betas from intraday data as in Andersen *et al.* (2006), the GARCH betas in Bollerslev, Engle, and Wooldridge (1988), the stochastic Bayesian beta in Jostova and Philipov (2005), simple Bayesian adjustments to ordinary least squares such as Vasicek (1973), or commercial beta estimates such as those from BARRA, Bloomberg, and Ibbotson.

Appendix A: Replicating Payoff Functions

Carr and Madan (2001) show that any twice continuously differentiable payoff function $f(S)$ can be replicated with bonds, the underlying stock, and the cross section of out-of-the-money options. For convenience, we replicate their argument here. The fundamental theorem of calculus implies that for any fixed F ,

$$\begin{aligned}
 f(S) &= f(F) + 1_{S>F} \int_F^S f'(u) du - 1_{S<F} \int_S^F f'(u) du \\
 &= f(F) + 1_{S>F} \int_F^S [f'(F) + \int_F^u f''(v) dv] du \\
 &\quad - 1_{S<F} \int_S^F [f'(F) + \int_u^F f''(v) dv] du.
 \end{aligned}$$

Because $f'(F)$ is not a function of u , we are able to apply Fubini's theorem

$$f(S) = f(F) + f'(F)(S - F) + 1_{S>F} \int_F^S \int_v^S f''(v) du dv + 1_{S<F} \int_S^F \int_S^v f''(v) du dv.$$

Now integrate over u

$$\begin{aligned}
 f(S) &= f(F) + f'(F)(S - F) + 1_{S>F} \int_F^S f''(v)(S - v) dv + 1_{S<F} \int_S^F f''(v)(v - S) dv \\
 &= f(F) + f'(F)(S - F) + \int_F^\infty f''(v)(S - v)^+ dv + \int_0^F f''(v)(v - S)^+ dv.
 \end{aligned}$$

If we set F equal to the initial stock price, $F = S_0$, and integrate over K instead of v , where K is interpreted as the strike, we are left with the spanning equation

$$\begin{aligned}
 f(S) &= [f(S_0) - f'(S_0)S_0] + f'(S_0)S + \int_{S_0}^\infty f''(K)(S - K)^+ dK \\
 &\quad + \int_0^{S_0} f''(K)(K - S)^+ dK.
 \end{aligned}$$

From this equation, we see that the payoff $f(S)$ is spanned by a $[f(S_0) - f'(S_0)S_0]$ position in bonds, $f'(S_0)$ position in shares of the stock, and a $f''(K)dK$ position in out-of-the-money options.

Appendix B: The Accuracy of the Moment Computations

As we do not have a continuum of option prices available to compute the moments, it is inevitable that certain biases will be induced in the estimation. We conducted a number of Monte Carlo experiments to determine the importance of these biases

on the estimation of volatility and skewness, and we used the results of these Monte Carlo experiments to guide our empirical implementation. There are three types of biases that are of particular concern in the estimation of volatility and skewness: discretization of strike prices, truncation of the integration domain, and asymmetry of the integration domain. Jiang and Tian (2005) examine the effects of discretization and truncation on the computation of option-implied volatility. We closely follow their setup and extend their analysis to the computation of option-implied skewness.

To examine the size of the approximation errors induced by these biases, we generate option prices using Heston's (1993) stochastic volatility (HSV) model with standard parameterization based on the empirical results of Bakshi, Cao, and Chen (1997): $\theta = 0.04$, $\kappa = 2$, $\sigma_v = 0.225$, and $\rho = -0.5$. We set the initial instantaneous variance (V) to be equal to the long-run variance $V = \theta = 0.04$, the initial stock price equal to $S_0 = 100$, and the risk-free rate equal to $r = 0.05$. Because we are interested in a 180-day beta, we set the time to maturity equal to 180 days, $T = 180/365$. To measure the variance and the skewness of the distribution implied by the HSV parameterization, we use the following approach. To ensure that we do not encounter negative stock prices or negative variances, we use Ito's lemma to convert the risk-neutral stochastic volatility dynamics in HSV to $d\ln S$ and $d\ln V$. Once we discretize the equations, we get the following risk-neutral dynamics:

$$\begin{aligned} \ln(S_{t+\Delta t}) &= \ln(S_t) + \left(r - \frac{1}{2}V_t\right)\Delta t + \sqrt{V_t}\sqrt{\Delta t}\epsilon_{t+\Delta t}^S, \\ \ln(V_{t+\Delta t}) &= \ln(V_t) + \frac{1}{V_t}\left(\left(\kappa + \lambda\right)\left(\frac{\kappa\theta}{\kappa + \lambda} - V\right) - \frac{1}{2}\sigma_v^2\right)\Delta t \\ &\quad + \frac{1}{\sqrt{V_t}}\sigma\sqrt{\Delta t}\epsilon_{t+\Delta t}^V, \\ \epsilon_{t+\Delta t}^V &= \rho\epsilon_{t+\Delta t}^S + \sqrt{1 - \rho^2}\epsilon_{t+\Delta t}, \end{aligned} \quad (14)$$

where $\Delta t = 1/252$, assuming that there are 124 trading days in 180 calendar days. By iterating through these equations 124 times, we can generate a 180-day stock price path. We repeat this exercise 250,000 times, calculate the log returns $\ln(\frac{S_T}{S_0})$ of each path, and compute volatility, which is equal to 0.2037, and skewness, which is equal to -0.4610 . To verify the accuracy of this simulation-based approach, we compare the call price with strike $K = 100$ obtained using simulation, C_{SIM} , with the closed-form solution C_{HEST} . We find $C_{\text{HEST}} = 6.5905$ and $C_{\text{SIM}} = 6.5780$.

We now have benchmarks to evaluate the accuracy of our estimation procedure. The first bias we investigate results from the discreteness of the strike prices. To compute the moments arbitrarily precisely, we need a continuum of option prices

from 0 to plus infinity, while in reality we only have prices at fixed strike price levels. In the first part of this experiment, we generate HSV call and put prices with different discrete strike price intervals d ranging from $d = 0.1$ to $d = 5$ using the integration domain $[Su^{-1}, Su]$ with $u = 2$. The percentage approximation errors for the volatility and skewness estimates are plotted in the top row of panels in Figure 8 for both simple trapezoidal integration using only observed prices (dotted line) and the interpolation–extrapolation technique of Jiang and Tian (2005) described in Section 2 (solid line). We conclude that when using the interpolation–extrapolation technique of Jiang and Tian, the bias is negligible for realistic discreteness.

We next investigate the bias resulting from the truncation of the integration domain. We keep the strike price interval constant at $d = 0.1$, and we vary the width of the integration domain $[Su^{-1}, Su]$ by changing u between 1.1 and 2. The middle




Figure 8. Volatility and skewness approximation error

We plot the percentage approximation errors for the simple integration approach (dotted line) and the interpolation–extrapolation technique (solid line) of Jiang and Tian (2005). HSV call and put prices are generated with parameters $\theta = 0.04$, $\kappa = 2$, $\sigma_v = 0.225$, $\rho = -0.5$, and $V = 0.04$. The approximation error is calculated as the percent difference between the estimated moments and the actual moments. The actual moments are obtained by generating 250,000 HSV 180-day log returns and computing the sample standard deviation (0.2037) and skewness (-0.4610).

row of panels in Figure 8 plots the percentage approximation errors in the volatility and skewness estimation for simple trapezoidal integration using only observed prices (dotted line) and for the interpolation–extrapolation technique (solid line). We see that it is difficult to estimate skewness accurately when the width of the integration domain is small. Therefore, as mentioned above, we choose a sample of stocks with liquid option data.

Finally, we investigate the bias induced by the asymmetry of the integration domain, which has not been investigated in the previous studies. This asymmetry presumably has little effect on the measurement of volatility, but it may be important for the measurement of skewness because skewness in effect compares the two sides of the distribution. If the integration domain is asymmetric, there is more information about one side of the log return distribution than about the other. To illustrate the resulting bias, we generate HSV call and put prices with a constant strike price interval $d = 0.1$ and we vary the integration domain from $[Su_L^{-1}, Su_H]$, where $u = 1.6$, $u_L = u - du$, $u_H = u + du$ and we vary du between -0.3 and 0.3 . The bottom row of Figure 8 plots the percentage approximation errors in the volatility and skewness estimation for simple trapezoidal integration using only observed prices (dotted line) and for the interpolation–extrapolation technique (solid line). For the interpolation–extrapolation technique used, the errors are very small for the relevant strike prices.

We conclude from these three experiments that the integration technique introduced by Jiang and Tian (2005) does a good job of mitigating potential biases. We therefore use this approach in our empirical implementation.

References

- Ait-Sahalia, Y. and Lo, A. (2000) Nonparametric risk management and implied risk aversion, *Journal of Econometrics* **94**, 9–51.
- Andersen, T., Bollerslev, T., Diebold, F., and Wu, J. (2006) Realized beta: Persistence and predictability, in: T. Fomby and D. Terrell, (eds.), *Advances in Econometrics: Econometric Analysis of Economic and Financial Time Series in Honor of R.F. Engle and C.W.J. Granger, Vol. 20, Part B*, Elsevier, Oxford, UK, 1–40.
- Ang, A., Chen, J., and Xing, Y. (2006) Downside risk, *Review of Financial Studies* **19**, 1191–1239.
- Bakshi, G., Cao, C., and Chen, Z. (1997) Empirical performance of alternative option pricing models, *Journal of Finance* **52**, 2003–2049.
- Bakshi, G., Kapadia, N., and Madan, D. (2003) Stock return characteristics, skew laws, and differential pricing of individual equity options, *Review of Financial Studies* **10**, 101–143.
- Bakshi, G. and Madan, D. (2000) Spanning and derivative security valuation, *Journal of Financial Economics* **55**, 205–238.
- Barberis, N., Shleifer, A., and Wurgler, J. (2005) Comovement, *Journal of Financial Economics* **75**, 283–317.
- Bauer, R., Cosemans, M., Frehen, R., and Schotman, P. (2008) A Bayesian panel data approach to explaining market beta dynamics, Working paper, University of Maastricht.

- Black, F., Jensen, M., and Scholes, M. (1972) The capital asset pricing model: Some empirical tests, in: M. Jensen, (ed.), *Studies in the Theory of Capital Markets*, Praeger, New York, New York.
- Blair, B., Poon, S.-H., and Taylor, S. (2001) Forecasting S&P100 volatility: The incremental information content of implied volatility and high frequency index returns, *Journal of Econometrics* **105**, 5–26.
- Bliss, R. and Panigirtzoglou, N. (2004) Option-implied risk aversion estimates, *Journal of Finance* **59**, 407–446.
- Blume, M. (1971) On the assessment of risk, *Journal of Finance* **20**, 1–10.
- Blume, M. (1975) Betas and their regression tendencies, *Journal of Finance* **30**, 785–795.
- Bodie, Z., Kane, A., and Marcus, A. (2007) *Investments*, 7th edition, McGraw-Hill, New York.
- Bollerslev, T., Engle, R., and Wooldridge, J. (1988) A capital asset pricing model with time varying covariances, *Journal of Political Economy* **96**, 116–131.
- Bos, T. and Newbold, P. (1984) An empirical investigation of the possibility of stochastic systematic risk in the market model, *Journal of Business* **57**, 35–41.
- Brandt, M. and Diebold, F. (2006) A no-arbitrage approach to range-based estimation of return covariances and correlations, *Journal of Business* **79**, 61–74.
- Britten-Jones, M. and Neuberger, A. (2000) Option prices, implied price processes, and stochastic volatility, *Journal of Finance* **55**, 839–866.
- Busch, T., Christensen, B. J., and Nielsen, M. (2011) The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets, *Journal of Econometrics* **160**, 48–57.
- Buss, A. and Vilkov, G. (2009) Option-implied correlation and factor betas revisited. Working paper, Goethe University.
- Campbell, J., Lo, A., and MacKinlay, A. C. (1997) *The Econometrics of Financial Markets*. Princeton University Press, Princeton, New Jersey.
- Canina, L. and Figlewski, S. (1993) The informational content of implied volatility, *Review of Financial Studies* **6**, 659–681.
- Carr, P. and Madan, D. (2000) Factor models for option pricing, Working paper, University of Maryland.
- Carr, P. and Madan, D. (2001) Optimal positioning in derivative securities, *Quantitative Finance* **1**, 19–37.
- Carr, P. and Wu, L. (2009) Variance risk premia, *Review of Financial Studies* **22**, 1311–1341.
- Christensen, B. and Prabhala, N. (1998) The relation between implied and realized volatility, *Journal of Financial Economics* **50**, 125–150.
- Cochrane, J. (2001) *Asset Pricing*, Princeton University Press, Princeton, New Jersey.
- Damodaran, A. (1999) *Applied Corporate Finance: A User's Manual*, Hoboken, John Wiley & Sons, New Jersey.
- Day, T. and Lewis, C. (1992) Stock market volatility and the information content of stock index options, *Journal of Econometrics* **52**, 267–287.
- Dennis, P. and Mayhew, S. (2002) Risk-neutral skewness: Evidence from stock options, *Journal of Financial and Quantitative Analysis* **37**, 471–493.
- Derman, E. and Kani, I. (1998) Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility, *International Journal of Theoretical and Applied Finance* **1**, 61–110.
- Driessen, J., Maenhout, P., and Vilkov, G. (2009) The price of correlation risk: Evidence from equity options, *Journal of Finance* **64**, 1377–1406.
- Duan, J-C and Wei, J. (2009) Systematic risk and the price structure of individual equity options, *Review of Financial Studies* **22**, 1981–2006.

- Ederington, L. and Guan, W. (2002) Is implied volatility an informationally efficient and effective predictor of future volatility?, *Journal of Risk* **4**, 29–46.
- Fama, E. (1991) Efficient capital markets: II, *Journal of Finance* **46**, 1575–1618.
- Fama, E. and French, K. (1992) The cross-section of expected stock returns, *Journal of Finance* **47**, 427–465.
- Fama, E. and French, K. (1993) Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**, 3–56.
- Fama, E. and French, K. (1996) Multifactor explanations of asset pricing anomalies, *Journal of Finance* **53**, 55–84.
- Fama, E. and MacBeth, J. (1973) Risk, return, and equilibrium: Empirical tests, *Journal of Finance* **81**, 607–636.
- Ferson, W. (1995) Theory and empirical testing of asset pricing models, in: R. Jarrow, V. Maksimovic, and W. Ziemba, (eds.), *Handbooks in Operations Research and Management Science*, Elsevier Science B.V., Amsterdam, The Netherlands.
- Ferson, W. (2004) Tests of multifactor pricing models, volatility bounds and portfolio performance, in: G. Constantinides, M. Harris, and R. Stultz, (eds.), *Handbook of the Economics of Finance*, Amsterdam, The Netherlands: Elsevier Science Publishers.
- Ferson, W. and Harvey, C. (1999) Conditioning variables and the cross-section of stock returns, *Journal of Finance* **54**, 1325–1360.
- Ferson, W., Kandel, S., and Stambaugh, R. (1987) Tests of asset pricing with time-varying expected risk premiums and market betas, *Journal of Finance* **42**, 201–220.
- Ferson, W. and Korajczyk, R. (1995) Do arbitrage pricing models explain the predictability of stock returns?, *Journal of Business* **68**, 309–349.
- Figlewski, S. (1997) Forecasting volatility, financial markets, *Institutions and Instruments* **6**, 2–87.
- Fleming, J. (1998) The quality of market volatility forecasts implied by S&P100 index option prices, *Journal of Empirical Finance* **5**, 317–345.
- Fouque, J.-P. and Kollman, E. (2009) Calibration of stock betas from skews of implied volatilities, *Applied Mathematical Finance*, forthcoming.
- French, D., Groth, J., and Kolari, J. (1983) Current investor expectations and better betas, *Journal of Portfolio Management* **10**, 12–17.
- Ghysels, E. (1998) On stable factor structures in the pricing of risk: Do time-varying betas help or hurt?, *Journal of Finance* **51**, 425–461.
- Granger, C. (1969) Prediction with a generalized cost of error function, *Operations Research Quarterly* **20**, 199–207.
- Granger, C. and Poon, S.-H. (2005) Practical issues in forecasting volatility, *Financial Analysts Journal* **61**, 45–56.
- Harvey, C. (1989) Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* **24**, 289–317.
- Heston, S. (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* **6**, 327–343.
- Huang, C.-F and Litzenberger, R. (1988) *Foundations for Financial Economics*, Prentice Hall, Englewood Cliffs, New Jersey.
- Jackwerth, J. (2000) Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* **13**, 433–451.
- Jackwerth, J. and Rubinstein, M. (1996) Recovering probability distributions from contemporaneous security prices, *Journal of Finance* **51**, 1611–1631.
- Jagannathan, R. and McGrattan, E. (1995) The CAPM debate, *Federal Reserve Bank of Minneapolis Quarterly Review* **19**, 2–17.

- Jagannathan, R. and Wang, Z. (1996) The conditional CAPM and the cross-section of expected returns, *Journal of Finance* **51**, 3–53.
- Jiang, G. and Tian, Y. (2005) The model-free implied volatility and its information content, *Review of Financial Studies* **18**, 1305–1342.
- Jiang, G. and Tian, Y. (2009) A random walk down the options market, Working paper, Schulich School of Business, York University.
- Jorion, P. (1995) Predicting volatility in the foreign exchange market, *Journal of Finance* **50**, 507–528.
- Jostova, G. and Philipov, A. (2005) Bayesian analysis of stochastic betas, *Journal of Financial and Quantitative Analysis* **40**, 747–778.
- Kraus, A. and Litzenberger, R. (1976) Skewness preference and the valuation of risky assets, *Journal of Finance* **31**, 1085–1100.
- Kraus, A. and Litzenberger, R. (1983) On the distributional conditions for a consumption-oriented three moment CAPM, *Journal of Finance* **38**, 1381–1391.
- Lamoureux, G. and Lastrapes, W. D. (1993) Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities, *Review of Financial Studies* **6**, 293–326.
- Lewellen, J. and Nagel, S. (2006) The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* **82**, 289–314.
- MacKinlay, A. C. (1995) Multifactor models do not explain deviations from the CAPM, *Journal of Financial Economics* **38**, 3–28.
- MacKinlay, A. C. (1997) Event studies in economics and finance, *Journal of Economic Literature* **35**, 13–39.
- McNulty, J., Yeh, T., Schulze, W., and Lubatin, M. (2002) What's your real cost of capital?, *Harvard Business Review* **80**, 114–121.
- Poon, S-H and Granger, C. (2003) Forecasting volatility in financial markets: A review, *Journal of Economic Literature* **26**, 478–539.
- Pong, S., Shackleton, M., Taylor, S., and Xu, X. (2004) Forecasting sterling/dollar volatility: A comparison of implied volatility and AR(FI)MA models, *Journal of Banking and Finance* **28**, 2541–2563.
- Rubinstein, M. (1994) Implied binomial trees, *Journal of Finance* **49**, 771–818.
- Scholes, M. and Williams, J. (1977) Estimating betas from nonsynchronous data, *Journal of Financial Economics* **5**, 309–327.
- Schwert, W. (1989) Why does stock market volatility change over time?, *Journal of Finance* **44**, 1115–1153.
- Siegel, A. (1995) Measuring systematic risk using implicit beta, *Management Science* **41**, 124–128.
- Vasicek, O. (1973) A note on using cross-sectional information in Bayesian estimation on security betas, *Journal of Finance* **28**, 1233–1239.