

The Price of Correlation Risk: Evidence from Equity Options

JOOST DRIESSEN, PASCAL J. MAENHOUT, and GRIGORY VILKOV*

ABSTRACT

We study whether exposure to marketwide correlation shocks affects expected option returns, using data on S&P100 index options, options on all components, and stock returns. We find evidence of priced correlation risk based on prices of index and individual variance risk. A trading strategy exploiting priced correlation risk generates a high alpha and is attractive for CRRA investors without frictions. Correlation risk exposure explains the cross-section of index and individual option returns well. The correlation risk premium cannot be exploited with realistic trading frictions, providing a limits-to-arbitrage interpretation of our finding of a high price of correlation risk.

CORRELATIONS PLAY A CENTRAL ROLE in financial markets. There is considerable evidence that correlations between asset returns change over time¹ and that stock return correlations increase when returns are low.² A marketwide increase in correlations negatively affects investor welfare by lowering diversification benefits and by increasing market volatility, so that states of nature with unusually high correlations may be expensive. It is therefore natural to ask whether marketwide correlation risk is priced in the sense that assets that pay off well when marketwide correlations are higher than expected (thereby providing a

*Driessen is at the University of Amsterdam. Maenhout and Vilkov are at INSEAD. We would like to thank Yacine Aït-Sahalia, David Bates, Jonathan Berk, Oleg Bondarenko, Michael Brandt, Menachem Brenner, John Campbell, Mike Chernov, Greg Duffee, Darrell Duffie, Rob Engle, Jan Ericsson, Gerard Genotte, Jens Jackwerth, Chris Jones, Frank de Jong, Hayne Leland, Toby Moskowitz, Anthony Neuberger, Josh Rosenberg, Mark Rubinstein, Pedro Santa-Clara, Ken Singleton, Otto van Hemert, Robert Whitelaw, Zhipeng Zhang, and especially Bernard Dumas for comments and stimulating discussions. We are particularly grateful for the detailed and constructive comments of an anonymous referee and the Editor. We received helpful comments from seminar participants at Berkeley Haas School of Business, BI Oslo, Cornell Johnson School, HEC Lausanne, INSEAD, LBS-LSE-Oxford Asset Pricing Workshop, MIT Sloan, NY Fed, NYU Stern, Stanford GSB, Tilburg, University of Amsterdam, University of Bonn, University of Frankfurt, University of Rotterdam, Warwick Business School, Yale SOM, CEPR Summer Symposium Gerzensee, EFA 2005, Duke-UNC Asset Pricing Conference, and WFA 2006. We gratefully acknowledge the financial support of INSEAD R&D.

¹ See Bollerslev, Engle, and Woolridge (1988) and Moskowitz (2003), among others. Brandt and Diebold (2006), and Engle and Sheppard (2005) present recent innovations in the estimation of dynamic correlations.

² Financial crises are often viewed as episodes of unusually high correlations. Roll (1988) analyzes the 1987 crash and Jorion (2000) studies the Russia/LTCM crisis. Longin and Solnik (2001) use extreme value theory to study whether international equity correlations increase in volatile times.

hedge against correlation risk) earn lower returns than can be justified by their exposure to other priced risk factors. Index options are an obvious example of such assets, as they will appear expensive when correlation risk is priced.

This is the first paper to analyze whether cross-sectional differences in exposure to marketwide correlation risk can account for cross-sectional differences in expected returns. Our first contribution is to provide evidence of a large correlation risk premium. We show that the differential pricing of index and individual stock options contains unique information on the price of correlation risk. In particular, our analysis of the cross-section of index and individual option returns, as well as the study of variance risk premia in index and individual options, highlights an important tension between index and individual option prices. Demonstrating this tension and offering a risk-based explanation for it forms the second contribution of this paper.

The bulk of recent work on empirical option pricing studies index options. Although there is growing evidence that individual option prices and returns behave differently empirically, most work focuses on Black and Scholes (1973) and Merton (1973a) implied volatility functions.³ We add formal evidence that individual options, unlike index options, do not embed a negative variance risk premium, nor earn economically significant returns in excess of a one-factor model. By considering individual options on all index components, our analysis emphasizes that a challenge in option pricing concerns explaining the difference between expected index and individual option returns. This is challenging since the index process is the weighted average of the individual processes. A risk-based explanation for the contrast between index and individual options requires that aggregated individual processes be exposed to a risk factor that is lacking from the individual processes. Priced correlation risk makes this possible. Intuitively, index options are expensive and earn low returns, unlike individual options, because they offer a valuable hedge against correlation increases and insure against the risk of a loss in diversification benefits.^{4,5} Our results thereby also offer a novel view on the source of the large volatility risk premium that recent work on index options has disclosed.

We use data on S&P100 index options and on individual options on all the S&P100 index components, combined with prices of the underlying stocks from January 1996 until the end of December 2003. We provide evidence for a correlation risk premium in three different ways.

³ See, for instance, Bakshi and Kapadia (2003b), Bakshi, Kapadia, and Madan (2003), Bollen and Whaley (2004), Branger and Schlag (2004), Dennis and Mayhew (2002), and Dennis, Mayhew, and Stivers (2006).

⁴ Rubinstein (2000) revisits the 1987 crash and lists correlation risk as a potential reason why stock market declines and increases in volatility coincide, noting that "Correlation increases in market declines, which increases volatility and reduces opportunities for diversification."

⁵ Garleanu, Pedersen, and Poteshman (2005) develop a model where risk-averse market makers cannot perfectly hedge a book of options, so that demand pressure increases the price of options. The authors document empirically that end-users are net long index options, which could explain their high prices, but the model is agnostic about the source of the exogenous demand by end-users. Our findings suggest that the demand for index options may well be driven by investors' desire to hedge against correlation risk.

First, we present a general decomposition of index variance risk. Index variance changes are due to changes in individual variances and changes in correlations, so that index variance risk is priced to the extent that individual variance risk and correlation risk are priced. We find a large negative index variance risk premium, in line with results in the recent literature.⁶ Unlike recent work, we also estimate variance risk premia in all individual options on all S&P100 components and find no evidence of a negative risk premium on individual variance risk.⁷ As the decomposition shows, these two findings are only consistent with each other in a risk-based model if exposure to correlation shocks is priced. Therefore, the stylized facts about index and individual variance risk provide model-free indirect evidence for priced correlation risk.

Second, we derive a simple option-based trading strategy aimed at exploiting priced correlation risk. The strategy sells index straddles and buys individual straddles and stocks in order to hedge individual variance risk and stock market risk, respectively. This trading strategy offers an attractive risk-return trade-off. Its Sharpe ratio is 77% higher than the one for bearing stock market risk in our sample. Correcting for standard risk factors, we find a large excess return of more than 10% per month. This is direct evidence of a large correlation risk premium. We demonstrate that this strategy has more attractive risk-return properties than the option-based trading strategies that have been suggested in the literature (like selling index puts or selling market variance), especially when considering higher moments of the return distributions. In a portfolio choice setting we find that the correlation strategy generates a utility gain for a CRRA investor that is substantially larger than what can be obtained with existing option-based strategies (selling market variance or selling index puts).

Finally, we estimate the correlation risk premium from the cross-section of index and individual option returns. Because of the large dispersion in their sensitivities to marketwide correlation shocks, these assets constitute a particularly well-chosen cross-section. Furthermore, recent work has shown that expected index option returns are very large in absolute value and extremely challenging to explain (e.g., Bondarenko (2003a, 2003b), Buraschi and Jackwerth (2001), Coval and Shumway (2001), and Jones (2006)). We show that differences in exposure to the correlation risk factor account for 70% of the cross-sectional variation in CAPM residuals of index and individual option returns. The estimated correlation risk premium is large and highly significant. Exposure to individual variance risk is not priced in this cross-section, in line with our other results.

⁶ The relevant literature includes Aït-Sahalia and Kimmel (2005), Andersen, Benzoni, and Lund (2002), Bakshi and Kapadia (2003a), Bollerslev, Gibson, and Zhou (2004), Bondarenko (2004), Broadie, Chernov, and Johannes (2007), Buraschi and Jackwerth (2001), Carr and Wu (2004), Coval and Shumway (2001), Eraker, Johannes, and Polson (2003), Eraker (2004), Jones (2006), and Pan (2002). Bates (2003) surveys earlier work.

⁷ In fact, we obtain weak evidence of a positive variance risk premium in individual options, which strengthens the evidence of a correlation risk premium.

In sum, our findings strongly suggest that correlation risk is priced. Merton's ICAPM (1973b) may provide a theoretical explanation for this finding to the extent that marketwide correlation levels have predictive power for market variance. As an alternative hypothesis, the large correlation risk premium we document may be interpreted as reflecting mispricing of index options due to investor irrationality and lack of arbitrage. For example, some investors may be overly cautious about correlation risk and this may lead to an irrationally high correlation risk premium. Simultaneously, rational arbitrageurs may face market frictions, which prevent them from exploiting the high correlation risk premium. To explore this limits-to-arbitrage hypothesis, we analyze the impact on the profitability and feasibility of our correlation trading strategy of market frictions in the form of transaction costs and margin requirements as in Santa-Clara and Saretto (2007). We show that transaction costs have an important impact on the profitability of the trading strategy. Its Sharpe ratio no longer exceeds the equity Sharpe ratio and the optimal portfolio weight for the correlation strategy becomes statistically insignificant. The impact of transaction costs on the correlation strategy is large because of the high bid-ask spreads for individual options. Furthermore, margin requirements make the correlation trading strategy infeasible for risk-tolerant investors, who stand to gain most from the strategy. Thus, if the large correlation risk premium reflects mispricing of index options, rational investors facing realistic market frictions cannot arbitrage the mispricing away and cannot exploit the correlation risk premium.

Very few papers have studied trading strategies based on individual options. A notable exception is Goyal and Saretto (2007), who analyze trading strategies using the cross-section of individual options and obtain very high Sharpe ratios. Their paper is complementary to ours, since they study in detail the cross-sectional predictability of individual option returns, while we focus on the difference between index and individual option returns (without modeling the cross-sectional differences of individual stock options).

Our paper is also related to work on option-implied correlations. Several articles investigate the correlation structure of interest rates of different maturities. Longstaff, Santa-Clara, and Schwartz (2001), De Jong, Driessen, and Pelsser (2004), and Han (2007) provide evidence that interest rate correlations implied by cap and swaption prices differ from realized correlations. Collin-Dufresne and Goldstein (2001) propose a term structure model where bond return correlations are stochastic. Campa and Chang (1998) and Lopez and Walter (2000) study the predictive content of implied correlations obtained from foreign exchange options for future realized correlations between exchange rates. Skintzi and Refenes (2003) describe how index and individual stock options can be used to find implied equity correlations for the Dow Jones Industrial Average index. They study the statistical properties and dynamics of the implied correlation measure with 1 year of data, but do not analyze the key implications for index option pricing. In fact, none of these articles investigates or estimates a risk premium on correlation risk. The negative correlation risk premium we find implies higher expected correlation paths under the risk-neutral measure

than under the actual measure. This divergence in expected correlations under the two measures can explain why option-implied correlations exceed average realized correlations.

Finally, it is interesting to note that practitioners have recognized the possibility of trading priced correlation risk, by implementing a strategy known as “dispersion trading.” This strategy typically involves short positions in index options and long positions in individual options. Very recently, a new contract aimed at directly exploiting the correlation risk premium has been introduced, namely, the correlation swap.

The paper is organized as follows. Section I presents the general decomposition of index variance risk. The data are described in Section II. Section III provides empirical evidence on variance and correlation risk premia, based on the framework of Section I. Section IV develops and empirically implements a correlation trading strategy. In Section V, we study whether priced correlation risk can explain the empirical cross-section of option returns. Section VI discusses the impact of transaction costs and margin requirements on the feasibility and profitability of the correlation trading strategy. Section VII concludes.

I. Understanding Market Variance Risk

We show in a general framework how market variance risk can be decomposed into individual variance risk and correlation risk. The risk premium for bearing market variance risk can be similarly decomposed. This section also briefly discusses the model-free implied variance estimator used in our empirical analysis in Section III.

A. The Determinants of Market Variance Risk

We study (priced) market variance risk from a new perspective by explicitly acknowledging that market variance risk can be decomposed into individual variance risk and correlation risk. Existing work does not entertain the possibility of priced correlation risk.

The stock market consists of N stocks. The price of stock i , S_i , follows an Ito process with instantaneous variance ϕ_i^2 , which itself also follows an Ito process.⁸ The instantaneous correlation between Wiener processes B_i and B_j that drive stocks i and j is

$$E_t[dB_i dB_j] = \rho_{ij}(t)dt, \quad i \neq j. \quad (1)$$

While we impose more structure on the dynamics of $\rho_{ij}(t)$ in Section IV, for now we only assume that $\rho_{ij}(t)$ follows an Ito process and that the conditions on $\phi_i(t)$ and $\rho_{ij}(t)$ for the resulting variance–covariance matrix to be positive-definite are satisfied for all t .

⁸ We omit time as an argument for notational convenience throughout, except when placing particular emphasis.

Given a set of index weights $\{w_i\}$, the instantaneous index variance $\phi_I^2(t)$ at time t is

$$\phi_I^2(t) = \sum_{i=1}^N w_i^2 \phi_i^2(t) + \sum_{i=1}^N \sum_{j \neq i}^N w_i w_j \phi_i(t) \phi_j(t) \rho_{ij}(t). \quad (2)$$

It is clear from (2) that index variance changes are driven by shocks to both individual variances $\phi_i^2(t)$ and correlations $\rho_{ij}(t)$. We are interested in the extent to which exposure to these shocks is priced. If the price of correlation risk is negative (because states with higher-than-expected correlation are associated with a deterioration in investment opportunities and investor welfare), assets with payoffs that covary positively with correlation provide a hedge against unexpected correlation increases and earn negative excess returns relative to what is justified by their exposure to standard risk factors. An index option has by construction a large positive exposure to index-wide correlation risk and thus constitutes a prime example of such an asset. Formally, a negative correlation risk premium manifests itself in a higher drift for the instantaneous correlation under the risk-neutral measure Q than under the physical measure P , thus driving a wedge between expected correlations under the two distributions. Intuitively, an index option will then seem expensive relative to a benchmark without priced correlation risk like Black–Scholes. The concept of priced variance risk follows the same reasoning.

The total index variance risk premium is $E_t^Q[d\phi_I^2] - E_t^P[d\phi_I^2]$.⁹ Given constant index weights $\{w_i\}$ and defining $\iota_i \equiv w_i^2 + \sum_{j \neq i} w_i w_j \frac{\phi_i}{\phi_i} \rho_{ij}$, applying Ito's lemma to (2) shows that

$$\begin{aligned} E_t^Q[d\phi_I^2] - E_t^P[d\phi_I^2] &= \sum_{i=1}^N \iota_i \{E_t^Q[d\phi_i^2] - E_t^P[d\phi_i^2]\} \\ &\quad + \sum_{i=1}^N \sum_{j \neq i}^N w_i w_j \phi_i \phi_j \{E_t^Q[d\rho_{ij}] - E_t^P[d\rho_{ij}]\}. \end{aligned} \quad (3)$$

In words, the index variance risk premium reflects all individual variance risk premia $E_t^Q[d\phi_i^2] - E_t^P[d\phi_i^2]$, as well as correlation risk premia $E_t^Q[d\rho_{ij}] - E_t^P[d\rho_{ij}]$.¹⁰ The factor ι_i multiplying the individual variance risk premium represents the contribution of stock i 's return variance to the index variance, scaled by its own variance. This is intuitive since the ι_i 's are used as weights when summing the individual variance risk premia to obtain their importance for the index variance risk premium.

Below, we first present a detailed study of index and individual variance risk premia, that is, the left-hand side and the first sum on the right-hand side of equation (3). This analysis provides indirect evidence on the importance of the

⁹ This definition represents the total variance risk premium, that is, including compensation for market risk if variance shocks are correlated with market risk (the "leverage effect"). In the empirical analysis, we correct for this in order to obtain the risk premium for "pure" variance risk.

¹⁰ The simplifying assumption of constant index weights is innocuous. Simulations show that allowing for stochastic index weights has a negligible impact on the empirical results with $N = 100$.

final sum in equation (3), that is, on correlation risk premia. In particular, since the time-series average of ι_i is empirically positive for all stocks that make up the S&P100 index over our 8-year sample, any evidence of a negative index variance risk premium and of nonnegative individual variance risk premia implies a negative correlation risk premium. Subsequently, we test directly for a correlation risk premium by analyzing a correlation trading strategy and we investigate empirically whether a common correlation risk factor and a common individual variance risk factor can account for cross-sectional variation in option returns.

Before turning to the data description and the empirical results, we present the model-free methodology used to estimate variance risk premia.

B. Model-Free Implied Variances and Variance Risk Premia

Consider the risk-neutral expected integrated variance of the return on asset $a \in \{I, 1, \dots, i, \dots, N\}$ over a discrete interval of length τ starting at time t :

$$\sigma_a^2(t) = E_t^Q \left[\int_t^{t+\tau} \phi_a^2(s) ds \right]. \quad (4)$$

We follow the methodology of Britten-Jones and Neuberger (2000), Carr and Madan (1998), and Dumas (1995), who build on the work of Breeden and Litzenberger (1978), to estimate the risk-neutral expected integrated variance $\sigma_a^2(t)$ defined in (4) from index options for $a = I$ and from individual options for $a = i$. As derived in Britten-Jones and Neuberger, their procedure gives the correct estimate of the option-implied (i.e., risk-neutral) integrated variance over the life of the option contract when prices are continuous but volatility is stochastic, in contrast to the widely used, but incorrect, Black–Scholes implied volatility. Furthermore, Jiang and Tian (2005) show that the method also yields an accurate measure of the (total) risk-neutral expected integrated variance in a jump-diffusion setting. The measure is therefore considered “model-free,” and can be labeled the *model-free implied variance (MFIIV)*.

We denote the price of a τ -maturity call option on asset a with strike price K at time t by $C_a(K, t)$. The main result of Britten-Jones and Neuberger is that the risk-neutral expected integrated variance $\sigma_a^2(t)$ defined in (4) equals the model-free implied variance, which is defined as

$$\sigma_{MF,a}^2(t) \equiv 2 \int_0^\infty \frac{C_a(K, t) - \max(S(t) - K, 0)}{K^2} dK. \quad (5)$$

Jiang and Tian show that the integral over a continuum of strikes in (5) can be approximated accurately by a sum over a finite number of strikes. Finally, Bollerslev et al. (2004), Bondarenko (2004), and Carr and Wu (2004) establish that the difference between the model-free implied variance and the realized variance can be used to estimate the variance risk premium. In particular, the null of a zero total variance risk premium implies a zero difference between average realized and average model-free implied variance.

Finally, it is noteworthy that *MFIV* equals the no-arbitrage variance swap rate. Equation (5) can therefore be used to synthetically create variance swaps from options across strikes K . This interpretation will prove useful in the subsequent tests in Section III.B.

II. Data Description

We use daily data from OptionMetrics for S&P100 index options and for individual options on all the stocks included in the S&P100 index from January 1996 until December 2003.¹¹

The S&P100 is a value-weighted index with quarterly rebalancing. During our sample period, the new index shares for the quarter are fixed (unless the number of floating shares changes during the quarter by more than 5%) based on the market values at the closing prices of the third Friday of the last month in the previous quarter. In addition, 47 changes in the list of constituent companies took place in our sample. These also occur on the rebalance dates. At each rebalance date, we construct index component weights using market values based on stock prices from CRSP. We keep these weights fixed until the next rebalance date. This introduces a small discrepancy between actual S&P100 daily weights and our fixed weights because the (actual) value-based weights fluctuate daily due to price changes. As we have 100 companies in the index, any such discrepancy due to changes in prices is small and can be neglected for our purposes (see also footnote 10).

From the OptionMetrics database, we select all put and call options on the index and on the index components. We work with best bid and ask closing quotes rather than the interpolated volatility surfaces constructed by OptionMetrics. In Sections III to V, we use the midquotes for these option data (average of bid and ask), and we assess the effect of bid-ask spreads in Section VI. We discard options with zero open interest, with zero bid prices, and with missing implied volatility or delta (which occurs for options with non-standard settlement or for options with intrinsic value above the current mid price). We focus on short-maturity options, which are known to trade most liquidly and consider all options with remaining maturity between 14 and 60 days. When multiple maturities are available within this interval, we select the maturity that generates the largest average number of call and put options with matching strike prices (to enable us to construct straddles). We also eliminate options of extreme moneyness (Black–Scholes delta below 0.15 for calls and above -0.05 for puts) as outliers, which filters out options with extremely high implied volatilities. From Section III onwards, when constructing straddle returns and a cross-section of option returns, we eliminate calls and puts without a matching option of the other type for the straddle construction. The options are American-style. However, for short-maturity options, the early exercise premium is typically negligible. Using a binomial tree, we find that,

¹¹ Interestingly, Standard and Poor's mentions on its website that a requirement for companies to be included in the S&P100 index is that they have listed options. This makes the S&P100 a natural index to consider for our study.

indeed, this premium is between 0.3% and 1.1% of the 1-month option price for puts (depending on volatility and moneyness), and thus has a small impact on option returns. For call options, the early exercise premium is zero with a continuous dividend yield (which is an appropriate assumption for index options). Santa-Clara and Saretto (2007) find that returns on American and European index options are very similar. Any early exercise premium in individual call options (due to discrete dividends) will actually bias against finding evidence for a correlation risk premium, since our correlation strategy buys individual options (see Section IV) and because *MFIV* would be too high in this case for individual options.

To construct the model-free implied variances, we require observations over time and across strikes of prices of S&P100 index options and individual stock options. Across the strike dimension, we use out-of-the-money (OTM) options, namely, calls with Black–Scholes delta below 0.5 and puts with delta above -0.5 . We implement the model-free implied variance measure of Section I.B following the procedure in Jiang and Tian (2005), suitably adjusted for put options when needed.¹² We calculate model-free implied variance on each day for each underlying that has at least three available options outstanding, with the restriction that at least one put and call be included.

We use daily returns from CRSP for individual stocks and from OptionMetrics for the S&P100 to estimate the realized variance. For each day, we calculate the realized variance over the same period as the one over which implied variance is calculated for that day, that is, ranging from 14 to 60 days, requiring that no more than 10 returns be missing from the sample. Since the window spans on average 1 month, this means that we require on average at least 12 observations out of 22 trading days.

We annualize both model-free and realized variance using 252 trading days in a calendar year. We use the T-bill rate of appropriate maturity (interpolated when necessary) from OptionMetrics as the risk-free rate.

III. Evidence on Variance and Correlation Risk Premia

Based on the general framework of Section I, we test for the presence of variance risk premia in index options, in individual options on all constituent stocks, and in the cross-section of individual variance swap returns. These tests are conducted using the model-free implied variance of Section I.B. In light of the general decomposition of index variance risk in equation (3), this analysis provides indirect evidence on the importance of priced correlation risk. Section IV presents direct evidence of a risk premium on correlation risk by developing

¹² One subtlety regarding the index weights emerges. If the expiration of the index option occurs after the next rebalance date, the index variance will reflect both the “old” and the “new” index weights. We calculate the projected weights of the index components using current market values. Moreover, in the period between rebalance dates there may be announcements of deletions from and additions to the index, which take effect at the next rebalance. We incorporate this migration in the projected weights. We weight the old fixed weights and the new projected weights using the relative time to maturity of the index option before the rebalance date and after the rebalance date.

and implementing a simple option-based trading strategy that exploits priced correlation risk. Finally, in Section V we test whether correlation and individual variance risk is priced in a well-chosen cross-section of assets (individual and index options).

A. Implied versus Realized Variances

The recent empirical literature on equity options primarily studies index options. Individual options have attracted much less attention. The majority of the recent work on individual options focuses on Black–Scholes implied volatility functions (Bakshi and Kapadia (2003b), Bakshi et al. (2003), Bollen and Whaley (2004), Branger and Schlag (2004), Dennis and Mayhew (2002), Dennis et al. (2006), and Garleanu et al. (2005)). A common finding is that implied volatility functions are flatter for individual options than for index options. While implied volatility functions provide very interesting information, they do not permit a formal test of the presence of variance risk premia. This section presents such a formal test, based on the model-free methodology described in Section I.B. Moreover, our OptionMetrics sample is more recent and spans 8 years (January 1996 up until December 2003) and includes options on all stocks that were included in the S&P100 over that period. Carr and Wu (2004) also use OptionMetrics and a related methodology, but focus on a subsample of 35 individual options.

We start with the index variance. Figure 1 plots the time series of (the square root of) the implied index variance and of (the square root of) the realized historical variance. The well-established finding that option-implied index variance is higher than realized index variance also holds for our recent sample. While all calculations are done for variances, we take square roots of the computed variances for interpretation purposes. Table I reports an average (annualized) realized index volatility of 20.80%, while the *MFIV* average is 24.69%. The null hypothesis that implied and realized index variance are on average equal is very strongly rejected, based on a *t*-test with Newey and West (1987) autocorrelation consistent standard errors for 22 lags (*t*-statistic of 6.81).

Turning to the equally weighted average of the individual variances in Figure 2, there is, quite remarkably, less systematic difference between the two volatility proxies. On average, the square root of realized variance (41.44%) actually exceeds the square root of implied variance (38.97%). The null hypothesis that, on average across all stocks in the index, the implied and realized variance are equal is rejected (*t*-statistic of 3.2), which suggests a significantly positive variance risk premium in individual options. However, when conducting the test in ratio form the null that $\frac{RV}{MFIV} = 1$ is only marginally rejected at the 5% confidence level. More importantly, carrying out the test for all stocks individually, the null of a zero variance risk premium ($RV = MFIV$) is not rejected at the 5% confidence level for 98 stocks out of the 127 stocks that are included in the sample for this analysis. Of the remaining 29 stocks, only seven exhibit a significant positive difference between implied and realized variance. We therefore find no evidence for the presence of a negative variance risk premium in




Figure 1. Implied versus realized volatility for index options. The figure presents the time series of the square root of the model-free implied index variance and of the square root of the realized variance over our 1/1996 to 12/2003 OptionMetrics sample. The model-free implied index variance is calculated from a cross-section (across strikes) of 1-month options on the S&P100, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in Section I.B. Realized variance is calculated from daily index returns over a 1-month window. Variances are expressed in annual terms.

individual stock options. If anything, there is weak evidence of a positive risk premium for variance risk in individual options.

This is quite surprising, given the well-known empirical regularity for index options. Bakshi and Kapadia (2003b) find a difference of 1% to 1.5% (depending on the treatment of dividends) between the average implied and the average historical volatility in their 1991 to 1995 sample of 25 individual stock options. They also stress that the difference is smaller than for index options. The discrepancy between our results and theirs may not only reflect the difference in sample, but also the difference in methodology to calculate the option-implied variance. Bollen and Whaley (2004) also report that the average deviation between (Black–Scholes) implied volatility and realized volatility is approximately zero for the 20 individual stocks in their sample. Finally, Carr and Wu (2004) use a similar methodology to ours and also report much smaller average variance risk premia for individual stocks than for S&P indices. The mean variance risk premia are insignificant for 32 out of the 35 individual stocks they study.¹³

¹³ Carr and Wu (2004) also report that estimates of mean log variance risk premia are significantly negative for 21 out of 35 individual stocks. However, mean log variance risk premia are expected to be negative (because of Jensen's inequality), even under the null of a zero variance risk premium, and thus lead to a biased test.

Table I
Variance Risk Premia in Index and Individual Options

The table reports the time-series averages of realized and model-free implied variances, for S&P100 options and for individual options on the stocks in the S&P100 index over the 1/1996 to 12/2003 sample period. For individual options the variances are equally weighted cross-sectional averages across all constituent stocks. Realized variance RV is calculated from daily returns over a 1-month window. The model-free implied variance $MFIV$ is calculated from a cross-section (across strikes) of 1-month options, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in Section I.B. The data on option prices are from OptionMetrics and variances are expressed in annual terms. The p -values, based on Newey and West (1987) autocorrelation consistent standard errors with 22 lags, are for the null hypothesis that implied and realized variance are on average equal ($RV = MFIV$ and $RV/MFIV = 1$).

	Index Options	Individual Options
Mean Realized Variance	0.2080 ²	0.4144 ²
Mean Model-Free Implied Variance	0.2469 ²	0.3897 ²
Difference $\sqrt{RV} - \sqrt{MFIV}$	-0.0389	0.0247
p value for $H_0 : RV - MFIV = 0$	0.0000	0.0014
p value for $H_0 : \frac{RV}{MFIV} - 1 = 0$	0.0000	0.0485
Individual Tests of Variance Risk Premia	# Stocks	
$H_0 : RV - MFIV = 0$ not rejected	98	
$H_0 : RV - MFIV \leq 0$ rejected	22	
$H_0 : RV - MFIV \geq 0$ rejected	7	

These findings provide indirect evidence of a negative correlation risk premium. As can be seen from equation (3), when individual variance risk is not priced (or carries a positive risk premium), index variance risk only carries a negative risk premium to the extent that the price of correlation risk is negative. Our results strongly suggest that this is the case.

B. The Cross-section of Individual Variance Swap Returns

We find that the total individual variance risk premium in individual options is not significantly negative for almost all index components. To gain further insight into this important result, we study cross-sectional pricing of individual variance risk in stock options. This analysis complements the approach above and investigates explicitly whether exposure of individual variances to market risk or to a common variance factor is priced in individual options.

We consider the cross-section of returns on synthetic individual variance swaps, which are natural assets to consider for a study of priced variance risk. Each variance swap can be synthetically created from a cross-section of options on the underlying stock.

Denoting the realized return variance of asset a from t to $t + \tau$ by $RV_a(t) \equiv \int_t^{t+\tau} \phi_a^2(s) ds$ and using the model-free implied variance $MFIV_a(t) \equiv \sigma_{MF,a}^2(t)$ defined in Section I.B, the return on a variance swap from t to $t + \tau$ is $r_a(t) \equiv \frac{RV_a(t)}{MFIV_a(t)} - 1$ (Bondarenko (2004) and Carr and Wu (2004)). The variance swap




Figure 2. Cross-sectional average of implied versus realized volatility for individual options. The figure presents the time series of the (equally weighted) cross-sectional average of the square root of the model-free implied individual stock variance and of the square root of the realized individual stock variance over our 1/1996 to 12/2003 OptionMetrics sample. For each stock in the S&P100 index, the model-free implied variance is calculated from 1-month options, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in Section I.B. Realized variance is calculated from daily CRSP stock returns over a 1-month window. Variances are expressed in annual terms. Because of migrations, a total of 135 individual stocks are considered over the entire 8-year sample period.

return is driven by unexpected variance shocks and on average reflects the variance risk premium. We construct the time series of 1-month variance swap returns for all stocks with at least 18 monthly observations during the sample period, approximating $RV_a(t) \equiv \int_t^{t+\tau} \phi_a^2(s) ds$ using daily return observations. To enhance the practical feasibility of the variance swap as a (synthetically) tradable instrument, we focus on monthly observations, selecting options that have 1 month to maturity, and construct a time series of non-overlapping returns.

As factors to explain the cross-section we consider the excess return on the market (proxied by S&P100) and a common individual variance component. The latter is calculated as the cross-sectional weighted average (using index weights) of the returns on individual variance swaps $\sum_{i=1}^N w_i r_i(t)$. We follow the standard two-step procedure, estimating first the factor loadings and then regressing average returns on factor loadings to estimate factor risk premia.

We find that exposure to market risk does not explain the cross-section of individual variance swap returns, as almost all of the first-step betas are statistically insignificant. The common individual variance risk factor $\sum_{i=1}^N w_i r_i(t)$ is also not priced in individual options; the individual variance risk premium

is found to be insignificant. The point estimate is actually positive (0.077 with a t -statistic of 1.74), which is in line with the finding above that realized individual variance exceeds option-implied variance. Including both factors simultaneously leads to the same finding.

In sum, covariance of individual variance shocks with either market risk or with a systematic individual variance risk factor does not command a risk premium in individual options, consistent with the results in the previous subsection. Combined with the finding of priced market variance risk in index options, this provides more indirect evidence of priced correlation risk.

IV. Risk and Return of a Correlation Trading Strategy

Motivated by the results of a significantly negative risk premium for index variance risk, but not for individual variance risk, we now explicitly examine whether exposure to correlation risk is priced. We do this by constructing a trading strategy that only loads on correlation risk. Examining the risk-return properties of this trading strategy provides direct evidence on the extent to which correlation risk carries a risk premium that cannot be explained by exposure to standard risk factors.

A. Trading Correlation Risk

We derive a correlation trading strategy based on the general framework of Section I and make some additional simplifying assumptions allowing us to implement the trading strategy empirically.

First, we assume that a single state variable $\rho(t)$ drives all pairwise correlations in (1):

$$\rho_{ij}(t) = \overline{\rho_{ij}} \rho(t). \quad (6)$$

This is a natural assumption, given our interest in priced correlation risk. If the risk of correlation changes carries a risk premium, we expect this to be compensation for the risk of marketwide correlation changes. The specific process for ρ can be left unspecified, except that $d\rho - E[d\rho] = \sigma(\rho)dB_\rho$, where the Brownian motion dB_ρ may be correlated with other risk factors, and suitable conditions on $\sigma(\rho)$ and the drift of ρ such that the resulting variance-covariance matrix be positive-definite are assumed to be satisfied. As a second simplification, we assume an equally weighted stock market index, that is, $I = \frac{1}{N} \sum_{i=1}^N S_i$, when initiating the trading strategy.

Each stock's instantaneous variance $\phi_i^2(t)$ follows an Ito process, with diffusion term $\varsigma_i(\phi_i)dB_{\phi_i}$:

$$d\phi_i^2 - E[d\phi_i^2] = \varsigma_i(\phi_i)dB_{\phi_i}. \quad (7)$$

Denoting the price of an at-the-money (ATM) straddle on asset $a \in \{I, 1, \dots, i, \dots, N\}$ by O_a (i.e., the sum of the put price and call price), we focus on unexpected straddle returns:

$$\frac{dO_i}{O_i} - E \left[\frac{dO}{O_i} \right] = \frac{S_i}{O_i} \frac{\partial O_i}{\partial S_i} \phi_i dB_i + \frac{1}{O_i} \frac{\partial O_i}{\partial \phi_i^2} \varsigma_i(\phi_i) dB_{\phi_i} \quad (8)$$

and

$$\frac{dO_I}{O_I} - E \left[\frac{dO_I}{O_I} \right] = \sum_{i=1}^N \frac{S_i}{O_I} \frac{\partial O_I}{\partial S_i} \phi_i dB_i + \sum_{i=1}^N \frac{1}{O_I} \frac{\partial O_I}{\partial \phi_i^2} \varsigma_i(\phi_i) dB_{\phi_i} + \frac{1}{O_I} \frac{\partial O_I}{\partial \rho} \sigma(\rho) dB_\rho. \quad (9)$$

The correlation trading strategy aims to short the index straddle in order to exploit the correlation risk premium, while hedging the exposure to stock return shocks dB_i and to individual volatility shocks dB_{ϕ_i} .¹⁴ We first hedge the individual variance risk. When shorting index straddles worth 100% of initial wealth, the portfolio weight y_i in individual straddle i is then obtained by solving

$$-\frac{1}{O_I} \frac{\partial O_I}{\partial \phi_i^2} \varsigma_i(\phi_i) + y_i \frac{1}{O_i} \frac{\partial O_i}{\partial \phi_i^2} \varsigma_i(\phi_i) = 0. \quad (10)$$

These weights are the same across stocks if we assume that the parameters of the variance processes are common across stocks.¹⁵

Delta-hedging for each return shock dB_i , the portfolio weight z_i for each stock satisfies

$$-\frac{S_i}{O_I} \frac{\partial O_I}{\partial S_i} \phi_i + y_i \frac{S_i}{O_i} \frac{\partial O_i}{\partial S_i} \phi_i + z_i \phi_i = 0, \quad (11)$$

which will again be the same across all stocks so that delta-hedging can be implemented with the stock market index.

The resulting strategy thus shorts index straddles worth all initial wealth and invests a fraction y_i of initial wealth in each individual straddle and z_i of initial wealth in each individual stock, the remainder being invested in the risk-free asset so that the portfolio weights sum to 100%. This strategy only has (negative) exposure to correlation shocks and thus collects the correlation risk premium if correlation risk is priced. The simplest way to implement the strategy is to use Black–Scholes deltas and vegas for $\partial O_i / \partial S_i$ and $\partial O_i / \partial \phi_i^2$, evaluated at the implied volatility of the options in the straddle.

The trading strategy resembles a so-called “dispersion trade,” which sells index options and buys individual options. However, our strategy also takes positions in equity to hedge stock market risk. Furthermore, the portfolio weights y_i and z_i vary over time with $\rho(t)$, since the delta and vega of the index straddle depend on ρ . We calculate these numerically for different levels of ρ , using

¹⁴ Note that while dB_ρ does not enter equation (8) directly, dB_ρ may still impact individual straddle returns through correlation between dB_ρ and dB_i or dB_{ϕ_i} . However, as we hedge exposure to both dB_i and dB_{ϕ_i} , this has no effect on the analysis.

¹⁵ Our correlation strategy could be refined by incorporating the apparent mispricing of individual options as documented in Goyal and Saretto (2007) by overweighting and underweighting specific individual straddles.

Table II
Summary Statistics for the Correlation Trading Strategy, Underlying S&P100 Index and Other Option-Based Trading Strategies

The table reports summary statistics for the correlation trading strategy analyzed in Section IV.B. All results are based on non-overlapping monthly returns on options and the equity index for our 1/1996 to 12/2003 sample period. The correlation trading strategy is constructed using the equity index, index straddles, and individual straddles. The underlying index is the S&P100. The index put is an equally weighted portfolio of 1-month index put options with Black–Scholes deltas ranging from -0.8 to -0.2 . All statistics are monthly, except the Sharpe ratios, which are annualized.

Strategy	Corr. Strategy	S&P100 Index	Short Index Straddle	Short Index Put
Excess Return	0.1037	0.0068	0.1187	0.3178
Standard Dev.	0.4904	0.0574	0.6397	1.7419
Skewness	-0.2784	-0.0581	-1.0837	-3.2627
Kurtosis	3.1529	4.0644	4.4110	16.6672
Ann. Sharpe	0.7325	0.4134	0.6429	0.6371
CAPM α	0.1059		0.1107	0.1630
t -stat	(1.96)		(1.47)	(1.79)
CAPM β	0.0282		0.0557	15.8347
t -stat	(0.02)		(0.03)	(5.46)

1-month average lagged historical correlations as an estimate for ρ at each point in time. For each pair of stocks, we calculate the historical correlation at time t over a 1-month window, imposing the same requirements as for the calculation of realized variances. The historical pairwise correlations can then be aggregated into a cross-sectional weighted average across all pairs of stocks, using the appropriate weights from the S&P100 index.

B. Empirical Results

Implementing the correlation trading strategy empirically, we find that its portfolio weight in individual straddles according to equation (10), aggregated across all index components, equals on average 101.12% of initial wealth. By construction the strategy also sells 100% of initial wealth worth of index straddles, while the fraction invested in the stock index (from equation (11)) is -32.54% of initial wealth. The remaining 131.42% is invested in the riskless asset. Shorting 100% of index straddles and buying 101.12% of individual straddles corresponds on average to buying 0.58 individual straddles (aggregated across all stocks) per shorted index straddle, when normalizing the initial value of all underlying assets (the index as well as its components) to one.

Table II reports the first four moments for the trading strategy return in excess of the risk-free rate, estimated from monthly non-overlapping returns of holding options to their maturity date. The excess return is 10.37% per month and the annualized Sharpe ratio is 0.73. Compared to the annualized Sharpe ratio for the S&P100 index itself (0.41 over the sample period), the Sharpe ratio on the trading strategy is 77% higher. Although theoretically speaking the trading strategy is hedged against return and volatility shocks and only

exposed to correlation shocks, the hedge is expected to be imperfect, since the trading strategy is model-based and the parameters are not chosen to minimize in-sample hedging errors. Moreover, the theoretical hedge requires (costly) continuous rebalancing. It is therefore important to analyze the excess returns in more detail, as Sharpe ratios may not fully reveal the risk of a strategy based on derivatives (Ingersoll et al. (2007)). We do this in several different ways.

First, we estimate the CAPM beta and alpha of the strategy. The CAPM beta is 0.028 (t -statistic of 0.02) and the CAPM alpha is 10.59% per month (t -statistic of 1.96). The beta estimate shows that the trading strategy successfully hedges market risk, even though rebalancing is done only at a monthly frequency. The zero beta is especially noteworthy in light of the extreme betas of alternative derivatives-based trading strategies. For example, an equally weighted portfolio of 1-month index put options with Black–Scholes deltas ranging from -0.8 to -0.2 has a CAPM beta of -15.83 . The CAPM alpha generated by the correlation trading strategy is highly significant economically and indicates that the high return on the trading strategy cannot be justified by exposure to stock market risk. The t -statistic of 1.96 reflects the relatively short sample.

In the Fama–French model (Fama and French (1993)) and the four-factor model that adds momentum (Jegadeesh and Titman (1993)), the loadings on any of the factors are insignificant and the strategy has similar alphas as for the CAPM (10.59% and 10.63% with t -statistics of 1.93 and 1.82, respectively). Finally, we also control for systematic liquidity risk. Pastor and Stambaugh (2003) construct equity portfolios on the basis of exposure to a systematic liquidity risk measure. We use the return difference between high and low liquidity-risk portfolios as a liquidity risk factor. These data are available at a monthly frequency covering calendar months, while our analysis uses monthly returns between option expiration dates. Maximizing the overlap between these staggered time series, we find that the exposure of the trading strategy to liquidity risk is positive but insignificant. Correcting the trading strategy return for market risk, the Fama–French factors, momentum, and the liquidity risk factor gives an alpha of 10.83% per month with a t -statistic of 1.91.

Second, to put the higher moments into perspective, we compare with the summary statistics for two strategies that have been analyzed extensively in the recent literature on index options, namely, writing 1-month index straddles and writing 1-month index puts. The index straddle allows the investor to exploit the market variance risk premium. Recent work has argued that this risk premium is very large and that investors can benefit by selling index variance (e.g., Coval and Shumway (2001) and Bondarenko (2004)). The attractiveness of selling index puts has also been widely documented (e.g., Bondarenko (2003b)). We find that the Sharpe ratio of our correlation trading strategy exceeds those for the alternative strategies by roughly 15%. More importantly, selling correlation risk involves substantially less negative skewness and kurtosis than selling market variance. The difference with the short index put position is even more pronounced. While the correlation strategy has skewness and kurtosis of -0.28 and 3.15 , the third and fourth moments for the short put strategy are -3.26 and 16.67 , respectively. Finally, it can be noted that the

correlation trading strategy return has less kurtosis in our sample than the underlying index return itself.

The risk-adjusted excess returns for the alternative strategies are also interesting. While the CAPM alpha for the index straddle is similar in magnitude to the alpha for the correlation trading strategy and is somewhat higher for the index put, the estimates are noisier and the alphas are insignificant with a t -statistic of 1.47 for the straddle and 1.79 for the index put. This finding is robust to correcting for additional risk factors (size, value, momentum, and/or liquidity risk). Intuitively, it is straightforward to understand the difference between the results for the index straddle and the correlation strategy. By adding individual straddles to the index straddle the correlation strategy hedges out the individual variance risk. Since individual variance risk has a negligible risk premium, both strategies have similar alphas, but the correlation strategy has lower risk leading to a higher t -statistic for its alpha. Further, by shorting index straddles and buying individual straddles, part of the gamma of the option position is neutralized, which explains why the correlation strategy has lower skewness and kurtosis than the short index straddle.

The analysis of alphas and Sharpe ratios above neglects the considerable degree of skewness and kurtosis often exhibited by option strategies. We therefore consider the portfolio choice problem of a CRRA investor, because CRRA preferences penalize for negative skewness and high kurtosis (in contrast to mean-variance preferences). Specifically, we estimate the optimal portfolio weights in the derivatives-based trading strategies for a CRRA investor with a 1-month horizon, who can also invest in the underlying equity index and in the risk-free asset. Based on these portfolio weights, we also report the certainty equivalent wealth that this investor, when already investing in the market index and the riskless asset, is willing to pay in order to gain access to the correlation trading strategy. Table III reports the results, starting with the certainty equivalents for the two alternative strategies, namely, the short index straddle and the short index put. The alternative strategies generate certainty equivalent wealth gains of 1.29% and 0.77% per month, respectively, for an investor with $\gamma = 1$, illustrating the attractiveness of these strategies. For this coefficient of risk aversion, the certainty equivalent for the correlation trading strategy is 36% higher than for the index straddle and 129% higher than for the index put. For any level of risk aversion considered in Table III, we find certainty equivalents for the correlation strategy that exceed the certainty equivalents for the short index straddle by at least 26% and for the short index put by at least 49%. While the derivatives weights are statistically insignificant for both alternative strategies (t -statistics between 1.28 and 1.68), the portfolio weights for the correlation strategy are substantially larger and (marginally) statistically significant, with t -statistics between 1.88 and 1.99 (depending on risk aversion). For example, an investor with $\gamma = 2$ invests 18.85% of his wealth in the correlation strategy. Given that the S&P 100 index options are defined on \$100 times the index value, an investor with \$100,000 of financial wealth would sell 5.89 index straddles to implement the correlation strategy (on average over the 1996 to 2003 sample period).

Table III
Correlation Trading Strategy: Portfolio Weights and Certainty
Equivalents for a CRRA Investor

The table reports the optimal empirical portfolio weights (and t -statistics) of a CRRA investor in derivatives-based trading strategies obtained by maximizing in-sample expected utility. The investor also invests in the underlying equity index and in the risk-free asset (these weights are not reported in the table). The monthly certainty equivalent is the percentage of initial wealth that a CRRA investor demands as compensation for not being able to invest in a particular derivatives-based strategy and instead only investing in the equity index and the risk-free asset. Three derivative strategies are considered: a short position in the index straddle, a short position in the index put, and the correlation trading strategy. The certainty equivalents are estimated using the optimal CRRA portfolio weights over the 1/1996 to 12/2003 sample period, for different levels of risk aversion γ . The value of $\gamma = 1.8$ generates (approximately) a 100% equity index weight when derivatives are not available.

Risk Aversion γ	1	2	5	10	20	1.8
Short Index Straddle						
Portfolio weight	0.2113	0.1182	0.0502	0.0255	0.0129	0.1298
t -stat	(1.68)	(1.62)	(1.58)	(1.55)	(1.55)	(1.63)
Cert. Equiv.	1.2910%	0.7240%	0.3079%	0.1569%	0.0792%	0.7948%
Short Index Put						
Portfolio weight	0.1208	0.0772	0.0353	0.0184	0.0094	0.0836
t -stat	(1.28)	(1.31)	(1.29)	(1.29)	(1.29)	(1.31)
Cert. Equiv.	0.7659%	0.5249%	0.2483%	0.1308%	0.0671%	0.5644%
Correlation Strategy						
Portfolio weight	0.3475	0.1885	0.0782	0.0395	0.0199	0.2078
t -stat	(1.99)	(1.92)	(1.89)	(1.88)	(1.88)	(1.93)
Cert. Equiv.	1.7558%	0.9477%	0.3934%	0.1988%	0.0999%	1.0451%

The final column ($\gamma = 1.8$) of Table III is of particular interest. It considers the investor who optimally holds the market (equity weight of approximately 100%) when derivatives are not available. This investor's optimal portfolio weight in the correlation strategy and the corresponding gain in certainty equivalent wealth provide quantitative measures of the extent to which the observed risk-return trade-off of the strategy could arise in an equilibrium with a CRRA representative investor (if the optimal weight is zero), or instead, of whether the correlation risk premium represents a "good deal" (Santa-Clara and Saretto (2007)). The investor stands to gain 1.05% of initial wealth per month from the correlation strategy (31% and 85% higher than for the index straddle and put, respectively), based on an optimal correlation-strategy weight of 21% of initial wealth (t -statistic of 1.93) and a positive (but insignificant) equity weight. It is clear that the risk-return trade-off of the correlation trading strategy is a good deal for a $\gamma = 1.8$ investor and could not arise in a simple no-trade equilibrium with this CRRA investor as representative agent.

Overall, the results for the correlation trading strategy indicate that the compensation for bearing correlation risk is substantial. The risk-return trade-off is considerably more generous than what can be obtained with short positions in index puts or in market variance.

V. The Cross-section of Individual and Index Option Returns

We now examine whether a correlation risk factor can account for cross-sectional variation in index and individual option returns. A cross-section of index and individual options is an ideal testing ground for this hypothesis, since returns on index options are driven by index variance shocks and thus by correlation shocks, while individual option returns are likely to be much less dependent on correlation shocks. We use the return on the trading strategy developed above as a correlation risk factor to explain the cross-section of expected index and individual option returns. Our test procedure is identical to standard procedures used in asset pricing to test for the presence of priced risk factors, and avoids the need for specific parametric modeling assumptions that are otherwise needed when testing option pricing models. Rather than developing a specific model of priced correlation risk, we test a generic prediction shared by all option pricing models with priced correlation risk, namely, that differences in exposure to correlation risk justify differences in expected returns. Simultaneously, we also test whether individual variance risk is priced, complementing earlier analyses in the paper.

Our cross-section contains 24 short-maturity options and is constructed as follows. We include both calls and puts and consider three different moneyness ranges, with deltas ranging from -0.8 to -0.2 for puts and from 0.2 to 0.8 for calls.¹⁶ This results in six index options and six (portfolios of) individual options. To obtain a larger cross-section, we further divide each individual option portfolio into three volatility categories by sorting options on the implied volatility of their underlying asset, resulting in 18 portfolios of stock options. Sorting individual options on volatility to construct a cross-section of option returns is natural since the volatility risk premium is in many option-pricing models a function of the volatility level. We calculate non-overlapping monthly option returns as holding-period returns, that is, the return at time $t + \tau$ on an option written at t is given by the option payoff at maturity ($t + \tau$) divided by the option price at t . Within each delta-volatility portfolio, we average the individual option returns cross-sectionally using the index weights for each day.

We use the standard two-step procedure for cross-sectional asset pricing, estimating first the factor loadings for all assets and then regressing average returns cross-sectionally on these loadings to obtain factor risk premia. The standard errors for the cross-sectional regression are calculated with the methodology of Shanken (1992) to correct for the estimation error in the first-step betas. We start by testing the CAPM, with the excess return on the market (proxied by the S&P100) as a factor. The S&P100 is arguably a narrow definition of the market, but natural for our setting as it is the underlying asset for the index options we study.

¹⁶ We categorize options according to Black–Scholes deltas rather than strike-to-spot ratios to ensure that the individual and index options are comparable in terms of economic moneyness. The strike-to-spot ratio of an index option cannot easily be compared with the one of individual options, as the underlying assets obviously have very different volatilities, for example.

A. CAPM Results

For index options, the CAPM betas range from -17 to 22 and are all highly significant. The betas for individual options are somewhat smaller (ranging from -16 to 15), but also very significant. Consistent with existing empirical work, a one-factor pricing model like the CAPM generates very large mispricing for index options, with time-series alphas of up to -31% per month and cross-sectional alphas of up to -30% per month. All alphas are negative for index options and average -17% per month. While the alphas are all economically significant, options returns are quite noisy resulting in only two significant time-series alphas (out of six) and three significant cross-sectional alphas.

For individual options, the results are quite different. No time-series or cross-sectional alpha is significant, even though the formation of portfolios (with averaging of individual option returns across more than 30 stocks) would be expected to lead to more precise estimates for individual options than for index options. Economically speaking, the contrast between index and individual options is clear: While the average index-option alpha is -17% per month, the average individual-option alpha is much smaller (-3.55%).

In summary, unlike for index options, the CAPM does quite well for individual options and we find no statistical evidence against it. This is consistent with our earlier findings of an insignificant difference between average realized and average risk-neutral variance for 98 out of the 127 stocks (Section III.A) and of an insignificant variance risk premium in the cross-section of individual variance swap returns (Section III.B).¹⁷

B. Results for Correlation and Individual Variance Risk Factors

We now add the correlation and individual variance risk factors to the CAPM and estimate a three-factor model. The return on the correlation trading strategy is taken as the correlation risk factor. For the individual variance risk factor, we use the return on the index-weighted portfolio of individual straddles.

In the first-step time-series regressions of option returns, we find very similar market betas as for the CAPM. All index option returns exhibit large and significantly negative loadings on the correlation risk factor. The correlation betas range from -0.36 (ITM index call) to -1.61 (OTM index call), with an average of -0.96 . All index options have correlation betas with t -statistics above 5.35 in absolute value. Individual options have smaller correlation loadings, with an average of -0.24 . Only 6 out of 18 individual-option portfolio returns have significant correlation betas. In contrast, index and individual options exhibit similar sensitivities to individual variance risk, with loadings between 0.42 and 1.89 (average of 1.07) for index options and between 0.07 and 1.63 for individual options (average of 0.73). All but three are statistically significant.

¹⁷ Bollen and Whaley (2004) present simulated returns of a delta-hedged trading strategy that shorts options (on the S&P500 and on 20 individual stocks). Unlike for index options, they find small abnormal returns for stock options, in line with our results for a larger sample (all stocks in the index) and using a different methodology.

Table IV
The Cross-section of Index and Individual Option Returns

The top panel (Three-Factor Model) of the table reports estimates for the risk premia on market risk, correlation risk, and individual variance risk, obtained from a cross-sectional regression of the average monthly excess returns on 24 index and individual options on their exposures to market risk, correlation risk, and individual variance risk. The exposures are estimated in a first step, regressing the time series of each excess option return on the market (S&P100) excess return, on the correlation trading strategy excess return (Table II), and on the excess return on a portfolio of individual straddles (mimicking individual variance risk). The bottom panel (Two-Factor Model Applied to CAPM Residuals) of the table reports estimates for the risk premia on correlation risk and individual variance risk, obtained from a cross-sectional regression of the average monthly excess returns on 24 index and individual options on their exposures to correlation risk and individual variance risk. The exposures are estimated in a first step, regressing the time series of each excess option return on the excess return on the correlation trading strategy and on the excess return on a portfolio of individual straddles. The option returns in the Two-Factor Model are all in excess of the CAPM-predicted return, using the S&P100 index return as market factor. The table reports *t*-statistics as in Shanken (1992) and the cross-sectional R^2 .

Three-Factor Model	
Market Risk Premium	0.0120
(<i>t</i> -stat)	(1.87)
Correlation Risk Premium	0.1751
(<i>t</i> -stat)	(2.56)
Individual Variance Risk Premium	0.0078
(<i>t</i> -stat)	(0.16)
Cross-sectional R^2	89.2%
Two-Factor Model Applied to CAPM Residuals	
Correlation Risk Premium	0.1726
(<i>t</i> -stat)	(2.56)
Individual Variance Risk Premium	0.0072
(<i>t</i> -stat)	(0.15)
Cross-sectional R^2	70.4%

For index options, the three-factor model generates time-series alphas that are all statistically insignificant and that have a mean of -1.78% and a mean absolute value of 4.43% . The CAPM generates negative alphas for all index options with an average of -17.15% . Accounting for exposure to correlation risk and individual variance risk leads to a notable reduction in mispricing and index options no longer seem significantly “overpriced.” The improvement for individual options is small, as the CAPM already performs quite well (the mean absolute alpha goes from 5.97% to 5.03%).

Table IV presents the results for the cross-sectional regression of average index and individual option returns on their factor loadings. The risk premium for the correlation factor is estimated to be 17.5% per month (*t*-statistic of 2.56). While this is higher than the average return on the trading strategy (10.37%), the difference between the two estimates of the correlation risk premium is not statistically significant. In contrast, the price of individual variance risk is small and statistically insignificant. Note that the positive risk premium for the correlation factor corresponds to a negative price of correlation risk,

since the trading strategy sells correlation and pays off well when correlations are low. Given that the correlation factor betas for index options are always negative, the positive estimate in Table IV leads to negative excess returns for index options relative to the CAPM, that is, consistent with the definition in Section I, the price of correlation risk is negative in the sense that assets with payoffs that covary positively with correlation (e.g., index options) earn negative excess returns.

As a robustness check, we now apply the correlation and individual variance risk factors to CAPM residuals, rather than regressing option returns simultaneously on the market, correlation, and individual variance risk factors. Given the estimated factor loadings, we obtain the correlation and individual variance risk premia from the cross-section of average CAPM excess returns. This alternative analysis of priced correlation risk is conservative in the sense that any correlation between the market return and the correlation risk factor is now automatically attributed to the market return (and its associated risk premium). We find identical results for the factor loadings on correlation and individual variance risk. The point estimates and t -statistics for the factor risk premia in Table IV are also very similar to the results for the three-factor model. The high cross-sectional R^2 of 70.4% is remarkable, given that the cross-section concerns CAPM residuals and that the model imposes linearity.

In conclusion, we find that individual variance risk is not priced in the cross-section of index and individual options, consistent with the results in Sections III.A and III.B. We also obtain strong evidence that exposure to correlation risk accounts for a substantial part of the cross-sectional variation in average excess returns that cannot be explained by standard market risk.

VI. The Impact of Transaction Costs and Margins

Our evidence points to a correlation risk premium that is both economically and statistically significant. Understanding the source and size of this risk premium is important. While unreported results indicate that marketwide correlations predict market variance, so that these correlations¹⁸ may be a priced state variable in Merton's ICAPM (in particular the extension in Chen (2003)), a general equilibrium model with priced correlation risk is needed to shed more light on the size of the correlation risk premium. Since developing such a general equilibrium model is beyond the scope of this paper, we instead analyze the impact of realistic trading frictions on the feasibility and profitability of the correlation trading strategy to explore whether limits to arbitrage may prevent investors from exploiting this correlation risk premium fully. In a recent paper, Santa-Clara and Saretto (2007) study the impact of transaction costs and margin requirements on the execution and profitability of index-option trading strategies and find that limits to arbitrage in the form of realistic trading frictions severely impact the risk-return trade-off of these strategies.

¹⁸ An interesting related question is why individual variance risk is not priced, which is also related to the pricing of idiosyncratic risk, analyzed by Ang et al. (2006) in the cross-section of stock returns and in the time-series sense in Goyal and Santa-Clara (2003).

It is interesting to apply their analysis to our setting for the following reasons. First, we found in Section IV that our correlation trading strategy outperforms two standard trading strategies based on index options when ignoring trading frictions. Second, the quantitative impact of trading frictions may be different, because our strategy involves not only index options, but also individual options. For comparison reasons, we also report the impact of frictions on the two alternative index-option trading strategies we consider, since the sample period as well as the type of index options (S&P500 versus S&P100) are different from the analysis in Santa-Clara and Saretto (2007).

We first account for transaction costs in the form of bid-ask spreads by using closing bid and ask quotes rather than mid quotes. As the correlation trading strategy sells index options and buys individual options, we calculate bid-to-maturity returns for index options and ask-to-maturity returns for individual options. Bid-ask spreads lower the excess return on the trading strategy by roughly 50%, specifically, from 10.4% (raw monthly excess return) and 10.6% (CAPM alpha) to 5.3% and 5.5%, respectively. The CAPM alpha is no longer statistically significant, with a t -statistic of 0.77. The annualized Sharpe ratio is also substantially lower when accounting for bid-ask spreads and drops from 0.73 in Table II to 0.41, which is very similar to the Sharpe ratio of the equity index in the absence of trading frictions. The impact of transaction costs on the alternative index-option trading strategies is less pronounced. For example, the raw excess return on the index straddle shrinks by only 2%, consistent with the findings of Santa-Clara and Saretto (2007) (taking into account that we hold options to maturity, thus avoiding roundtrip transaction costs on options). After transaction costs, the Sharpe ratios of the index straddle and put now exceed the Sharpe ratio of the correlation strategy (0.52 and 0.58 versus 0.41). The impact of transaction costs on our strategy is larger than for the two alternative index-option strategies because of the larger bid-ask spreads for individual options.

Table V reports the optimal portfolio allocation to the correlation trading strategy and to the two alternative index-option trading strategies, as well as the associated certainty equivalent wealth gains, for a CRRA investor facing transaction costs in the form of bid-ask spreads. As in Table III, the investor can also invest in the risk-free asset and in the underlying equity index. The transaction costs for the riskless asset and the equity index are expected to be an order of magnitude smaller than for the options strategies and for simplicity are assumed to be zero.

Not surprisingly, transaction costs have a major impact on the optimal allocation by CRRA investors in the correlation strategy. The portfolio shares are roughly 57% of the optimal weights without frictions (Table III) and are now statistically insignificant. However, they remain economically quite large for low γ . The point estimate for the certainty equivalent shows that the log investor still gains 0.47% of wealth per month from having access to the trading strategy, but there is no statistical evidence that the gain is significantly different from zero. Furthermore, since this certainty equivalent was 1.76% without frictions, it is clear that the economic impact of bid-ask spreads is substantial.

Table V
Correlation Trading Strategy with Transaction Costs: Portfolio
Weights and Certainty Equivalents for a CRRA Investor

The table reports the optimal empirical portfolio weights (and t -statistics) of a CRRA investor in derivatives-based trading strategies, accounting for transaction costs in the form of bid-ask spreads and obtained by maximizing in-sample expected utility. The investment opportunity set also includes the underlying equity index and the risk-free asset (these weights are not reported in the table). The monthly certainty equivalent is the percentage of initial wealth that a CRRA investor demands as compensation for not being able to invest in a particular derivatives-based strategy and instead only investing in the equity index and the risk-free asset. Three derivative strategies are considered: a short position in the index straddle, a short position in the index put, and the correlation trading strategy. The certainty equivalents are estimated using the optimal CRRA portfolio weights over the 1/1996 to 12/2003 sample period, for different levels of risk aversion γ . The value of $\gamma = 1.8$ generates (approximately) a 100% equity index weight when derivatives are not available. To account for transaction costs, we use closing bid and ask quotes rather than mid quotes. For the short index straddle and short index put, we use bid-to-maturity returns. For the correlation trading strategy, which sells index options and buys individual options, we use bid-to-maturity returns for index options and ask-to-maturity returns for individual options.

Risk Aversion γ	1	2	5	10	20	1.8
Short Index Straddle						
Portfolio weight	0.1612	0.0901	0.0383	0.0195	0.0098	0.0990
t -stat	(1.26)	(1.24)	(1.22)	(1.21)	(1.21)	(1.24)
Cert. Equiv.	0.7545%	0.4281%	0.1834%	0.0937%	0.0473%	0.4694%
Short Index Put						
Portfolio weight	0.0896	0.0581	0.0267	0.0139	0.0071	0.0628
t -stat	(0.92)	(0.99)	(1.00)	(1.00)	(1.00)	(0.99)
Cert. Equiv.	0.4115%	0.2971%	0.1441%	0.0765%	0.0394%	0.3178%
Correlation Strategy						
Portfolio weight	0.1976	0.1068	0.0445	0.0225	0.0113	0.1178
t -stat	(0.97)	(0.98)	(0.98)	(0.98)	(0.98)	(0.98)
Cert. Equiv.	0.4698%	0.2597%	0.1096%	0.0557%	0.0281%	0.2856%

As discussed above, the impact of transaction costs on the alternative index-option strategies is smaller. The optimal portfolio weights are only 25% smaller and the certainty equivalents are reduced by roughly 40% for index straddles and by approximately 60% for the index put portfolio. The certainty equivalent for the index straddle is now larger than for the correlation trade. However, for the index straddle and put strategies the associated portfolio weights are statistically insignificant (as before).

Another potentially important trading friction for options concerns margin requirements, as shown by Santa-Clara and Saretto (2007) for strategies based on index options. Our trading strategy shorts index options as well, but is simultaneously long individual options and also has a large component invested in the riskless asset, justifying an additional analysis of the effect of margins on the strategy's feasibility. We first calculate the required initial margin for the strategy based on its short position in index options and using the (stringent) CBOE margin rules for options described in Santa-Clara and Saretto (2007,

Table VI
Margin Requirements

The table reports the initial margin requirement, the maximum margin requirement, and the total holdings of the riskless asset for the portfolio choice problem of a CRRA investor having access to the riskless asset, the underlying equity index, and the correlation trading strategy. For each level of risk aversion, the margin requirements and riskless asset holdings reflect the optimal portfolio weights reported in Table III (without transaction costs) and in Table V (with transaction costs due to bid-ask spreads). The initial margin is based on the correlation strategy's short position in index options using the CBOE margin rules for options described in Santa-Clara and Saretto (2007, p. 12). We also report the maximum margin over the sample period, obtained by calculating for each 1-month period in our sample the largest margin requirement that occurs over the life of the trading strategy, using the description of margin updates in Santa-Clara and Saretto. The risk-free weight is the (total) portfolio weight allocated by the CRRA investor to the riskless asset, obtained as the sum of the direct holdings of the risk-free asset and of the indirect holdings through the correlation trading strategy (as reported in Section IV.B, the correlation strategy itself holds 131.42% in the riskless asset).

Risk Aversion γ	1	2	5	10	20	1.8
Correlation Strategy without Transaction Costs						
Initial margin	2.3757	1.2887	0.5346	0.2700	0.1360	1.4207
Max. margin	2.8215	1.5305	0.6349	0.3207	0.1616	1.6872
Riskfree weight	0.2134	0.4857	0.7661	0.8785	0.9382	0.4434
Correlation Strategy with Transaction Costs						
Initial margin	1.3509	0.7302	0.3042	0.1538	0.0773	0.7999
Max. margin	1.6044	0.8672	0.3613	0.1827	0.0918	0.9500
Riskfree weight	-0.2825	0.2963	0.7043	0.8499	0.9244	0.2256

p. 12).¹⁹ The initial margin can then be compared to the total risk-free investment held by the investor to see whether the initiation of the strategy is feasible. We also calculate for each 1-month period in our sample the largest margin requirement that occurs over the life of the trading strategy, using the description of margin updates in Santa-Clara and Saretto (2007). This gives the maximum margin over the sample period, which again can be compared to the holdings of the risk-free asset to check whether the position can be maintained.

Finally, we calculate the (total) portfolio weight allocated by the CRRA investor to the riskless asset, obtained as the sum of the direct holdings of the risk-free asset and of the indirect holdings through the correlation trading strategy (as reported in Section IV.B, the correlation strategy holds 131.42% in the riskless asset in case of a 100% portfolio weight for the strategy). We conduct the analysis with and without transaction costs in order to isolate the effect of both types of trading frictions.

The main finding in Table VI is that the optimal position in the correlation trading strategy is feasible for highly risk-averse investors, but not for investors

¹⁹ We conservatively do not allow for netting out of short positions in index options against long positions in individual options for margin purposes, as this would only be possible for market-maker accounts.

with $\gamma \leq 2$. Importantly, it is precisely this relatively risk-tolerant investor who stands to gain most from the strategy and for whom the optimal portfolio weight is economically significant when ignoring margin requirements. This conclusion obtains whether we incorporate transaction costs or not, indicating the relevance of margin requirements. The trading strategy is feasible for the $\gamma = 5$ investor precisely because her optimal weight in the strategy is small.

In summary, the analysis of trading frictions reveals that the correlation risk premium cannot be captured by investors who are subject to realistic transaction costs and margin requirements. Limits to arbitrage could therefore explain the economic presence of a correlation risk premium that may otherwise, that is, when ignoring frictions, seem very large. While studying the equilibrium price of correlation risk in a general equilibrium model is an interesting topic for future research, the results in this section suggest that our finding of a large correlation risk premium may also be consistent with a hypothesis of index-option mispricing. According to this hypothesis, the correlation risk premium need not (only) be the equilibrium compensation for correlation risk, but may also reflect inefficiencies in the market for index options leading to “overpriced” index options, which cannot be arbitrated away in the presence of realistic market frictions.

VII. Conclusion

We show empirically that correlation risk is priced in the sense that assets that pay off well when marketwide correlations are higher than expected earn negative excess returns. This result is consistent with increases in marketwide correlations leading to a deterioration of investment opportunities in the form of smaller diversification benefits. The negative excess return on correlation-sensitive assets can therefore be interpreted as an insurance premium.

We provide evidence of a large correlation risk premium in a number of different ways. First, while index options reflect a large negative variance risk premium, we find no significant negative premium on variance risk in individual options on all index components. Second, a trading strategy that sells correlation risk by selling index options and buying individual options earns excess returns of 10% per month and has a large Sharpe ratio. This strategy has more attractive risk-return properties (especially higher moments) than other option-based strategies. Third, the return on this correlation trading strategy explains 70% of the cross-sectional variation in index and individual option returns that is not accounted for by market risk.

As a second contribution, we demonstrate that priced correlation risk constitutes the missing link between unpriced individual variance risk and priced market variance risk, and enables us to offer a risk-based explanation for the discrepancy between index and individual option returns. Index options are expensive, unlike individual options, because they allow investors to hedge against positive marketwide correlation shocks and the ensuing loss in diversification benefits.

When introducing realistic market frictions in the form of transaction costs and margin requirements, we find that the correlation trading strategy cannot be exploited by investors facing these frictions. This provides a potential limits-to-arbitrage interpretation for our finding of a large correlation risk premium. Simultaneously, it should be noted that the market makers who are active in markets for both index and individual options can be expected to earn the correlation risk premium, since end-users of options have been shown to be net long index options and net short individual options (Garleanu et al. (2005)) and since market-makers are only margined on their net positions.

Correlation risk is relevant in many areas of financial economics. Subsequent to our work, Buraschi, Porchia, and Trojani (2006) study the effect of correlation risk on dynamic portfolio choice and Krishnan, Petkova, and Ritchken (2006) show that correlation risk is priced in the cross-section of stock returns. Another interesting application concerns the pricing of basket credit derivatives, such as collateralized debt obligations.

REFERENCES

- Aït-Sahalia, Yacine, and Robert L. Kimmel, 2005, Maximum likelihood estimation of stochastic volatility models, Working paper, Princeton University.
- Andersen, Torben, Luca Benzoni, and Jesper Lund, 2002, An empirical investigation of continuous-time models for equity returns, *Journal of Finance* 57, 1239–1284.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Bakshi, Gurdip S., and Nikunj Kapadia, 2003a, Delta-hedged gains and the negative market volatility risk premium, *Review of Financial Studies* 16, 527–566.
- Bakshi, Gurdip S., and Nikunj Kapadia, 2003b, Volatility risk premiums embedded in individual equity options: Some new insights, *Journal of Derivatives* 11, 45–54.
- Bakshi, Gurdip S., Nikunj Kapadia, and Dilip B. Madan, 2003, Stock return characteristics, skew laws, and differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.
- Bates, David S., 2003, Empirical option pricing: A retrospection, *Journal of Econometrics* 116, 387–404.
- Black, Fischer S., and Myron S. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.
- Bollen, Nicolas P. B., and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–754.
- Bollerslev, Tim P., Robert F. Engle, and Jeffrey M. Woolridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116–131.
- Bollerslev, Tim P., Michael Gibson, and Hao Zhou, 2004, Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities, Working paper, Duke University.
- Bondarenko, Oleg, 2003a, Statistical arbitrage and securities prices, *Review of Financial Studies* 16, 875–919.
- Bondarenko, Oleg, 2003b, Why are puts so expensive? Working paper, University of Illinois, Chicago.
- Bondarenko, Oleg, 2004, Market price of variance risk and performance of hedge funds, Working paper, University of Illinois, Chicago.
- Brandt, Michael W., and Francis X. Diebold, 2006, A no-arbitrage approach to range-based estimation of return covariances and correlations, *Journal of Business* 79, 61–73.

- Branger, Nicole, and Christian Schlag, 2004, Why is the index smile so steep? *Review of Finance* 8, 109–127.
- Breedon, Douglas T., and Robert H. Litzenberger, 1978, Prices of state contingent claims implicit in option prices, *Journal of Business* 51, 621–652.
- Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied price processes, and stochastic volatility, *Journal of Finance* 55, 839–866.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premiums: Evidence from futures options, *Journal of Finance* 62, 1453–1490.
- Buraschi, Andrea, and Jens C. Jackwerth, 2001, The price of a smile: Hedging and spanning in option markets, *Review of Financial Studies* 14, 495–527.
- Buraschi, Andrea, Paolo Porchia, and Fabio Trojani, 2006, Correlation hedging, Working paper, Imperial College.
- Campa, Jose M., and Kevin Chang, 1998, The forecasting ability of correlations implied in foreign exchange options, *Journal of International Money and Finance* 17, 855–880.
- Carr, Peter P., and Dilip B. Madan, 1998, Towards a theory of volatility trading, in Robert A. Jarrow, ed.: *Volatility: New Estimation Techniques for Pricing Derivatives* (RISK Publications, London).
- Carr, Peter P., and Liuren Wu, 2004, Variance risk premia, Working paper, Bloomberg L.P.
- Chen, Joseph S., 2003, Intertemporal CAPM and the cross-section of stock returns, Working paper, University of Southern California.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2001, Stochastic correlation and the relative pricing of caps and swaptions in a generalized-affine framework, Working paper, Carnegie Mellon University.
- Coval, Joshua D., and Tyler Shumway, 2001, Expected option returns, *Journal of Finance* 56, 983–1009.
- Dennis, Patrick J., and Stewart Mayhew, 2002, Risk-neutral skewness: Evidence from stock options, *Journal of Financial and Quantitative Analysis* 37, 471–493.
- Dennis, Patrick J., Stewart Mayhew, and Chris Stivers, 2006, Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381–406.
- Dumas, Bernard, 1995, The meaning of the implicit volatility function in case of stochastic volatility, Working paper, HEC.
- Engle, Robert F., and Kevin Sheppard, 2005, Evaluating the specification of covariance models for large portfolios, Working paper, NYU.
- Eraker, Bjorn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1403.
- Eraker, Bjorn, Michael Johannes, and Nick Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Garleanu, Nicolae B., Lasse H. Pedersen, and Allen M. Poteshman, 2005, Demand-based option pricing, Working paper, Wharton School of Business.
- Goyal, Amit, and Pedro Santa-Clara, 2003, Idiosyncratic risk matters, *Journal of Finance* 58, 975–1007.
- Goyal, Amit, and Alessio Saretto, 2007, Option returns and the cross-sectional predictability of implied volatility, Working paper, Emory University.
- Han, Bing, 2007, Stochastic volatilities and correlations of bond yields, *Journal of Finance* 62, 1491–1524.
- Ingersoll, Jonathan E., Matthew I. Spiegel, William N. Goetzmann, and Ivo Welch, 2007, Portfolio performance manipulation and manipulation-proof performance measures, *Review of Financial Studies* 20, 1503–1546.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

- Jones, Christopher S., 2006, A nonlinear factor analysis of S&P 500 index option returns, *Journal of Finance* 61, 2325–2363.
- de Jong, Frank, Joost Driessen, and Antoon Pelsser, 2004, On the information in the interest rate term structure and option prices, *Review of Derivatives Research* 7, 99–127.
- Jorion, Philippe, 2000, Risk management lessons from long-term capital management, *European Financial Management* 6, 277–300.
- Krishnan, CNV, Ralitsa Petkova, and Peter Ritchken, 2006, The price of correlation risk, Working paper, Case Western Reserve University.
- Longin, François, and Bruno Solnik, 2001, Extreme correlation of international equity markets, *Journal of Finance* 56, 651–678.
- Longstaff, Francis A., Pedro Santa-Clara, and Eduardo S. Schwartz, 2001, The relative valuation of caps and swaptions: Theory and empirical evidence, *Journal of Finance* 56, 2067–2109.
- Lopez, Jose A., and Christian Walter, 2000, Is implied correlation worth calculating? Evidence from foreign exchange options and historical data, *Journal of Derivatives* 7, 65–82.
- Merton, Robert C., 1973a, Theory of rational option pricing, *Bell Journal of Economics* 4, 141–183.
- Merton, Robert C., 1973b, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Moskowitz, Tobias J., 2003, An analysis of covariance risk and pricing anomalies, *Review of Financial Studies* 16, 417–457.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Roll, Richard, 1988, The international crash of October 1987, *Financial Analysts Journal* 44, 19–35.
- Rubinstein, Mark, 2000, Comments on the 1987 stock market crash: 11 years later, in *Symposium Proceedings of the Actuarial Foundation* (Actuarial Foundation, Schaumburg, IL).
- Santa-Clara, Pedro, and Alessio Saretto, 2007, Option strategies: Good deals and margin calls, Working paper, UCLA.
- Shanken, Jay A., 1992, On the estimation of beta pricing models, *Review of Financial Studies* 5, 1–34.
- Skintzi, Vasiliki D., and Apostolos N. Refenes, 2003, Implied correlation index: A new measure of diversification, Working paper, Athens University.