Option-Implied Correlations and the Price of Correlation Risk*

Joost Driessen[†] Pascal Maenhout[‡] Grigory Vilkov[§]

December 2006

Abstract

We study whether differences in exposure to market-wide correlation shocks can account for cross-sectional differences in expected returns. Consistent with correlation increases lowering diversification benefits and investor welfare, we present evidence of a large and significant correlation risk premium. We use data on S&P100 index options, options on all index components, and stock returns. First, we show that index variance risk is priced, while individual variance risk is not, which is indirect evidence of priced correlation risk. Second, a trading strategy exploiting priced correlation risk generates a high alpha and Sharpe ratio and is a very attractive investment for CRRA investors. Third, exposure to correlation risk explains the cross-section of index and individual option returns very well. Index options hedge correlation risk, offering a risk-based explanation for their high prices and low returns relative to individual options. Fourth, a parsimonious option pricing model with priced correlation risk shows that the substantial gap between average realized (28.7%) and option-implied (46.7%) stock correlations is direct evidence of a large correlation risk premium. Finally, we provide evidence suggesting that correlation risk may be priced in the cross-section of stock returns.

*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 Gennotte, Jens Jackwerth, Chris Jones, Frank de Jong, Hayne Leland, Toby Moskowitz, Anthony Neuberger, Pedro Santa-Clara, Josh Rosenberg, Mark Rubinstein, 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, the Duke-UNC Asset Pricing Conference and WFA 2006. We gratefully acknowledge the financial support of INSEAD R&D.

[†]University of Amsterdam, Finance Group, Faculty of Economics and Econometrics, Roetersstraat 11, 1018 WB Amsterdam, the Netherlands. Email: J.J.A.G.Driessen@uva.nl.

[‡]INSEAD, Finance Department, Boulevard de Constance, 77305 Fontainebleau Cedex, France. Email: pascal.maenhout@insead.edu.

[§]INSEAD, Finance Department, Boulevard de Constance, 77305 Fontainebleau Cedex, France. Email: grigory.vilkov@insead.edu.

Option-Implied Correlations and the Price of Correlation Risk

Abstract

We study whether differences in exposure to market-wide correlation shocks can account for cross-sectional differences in expected returns. Consistent with correlation increases lowering diversification benefits and investor welfare, we present evidence of a large and significant correlation risk premium. We use data on S&P100 index options, options on all index components, and stock returns. First, we show that index variance risk is priced, while individual variance risk is not, which is indirect evidence of priced correlation risk. Second, a trading strategy exploiting priced correlation risk generates a high alpha and Sharpe ratio and is a very attractive investment for CRRA investors. Third, exposure to correlation risk explains the cross-section of index and individual option returns very well. Index options hedge correlation risk, offering a risk-based explanation for their high prices and low returns relative to individual options. Fourth, a parsimonious option pricing model with priced correlation risk shows that the substantial gap between average realized (28.7%) and option-implied (46.7%) stock correlations is direct evidence of a large correlation risk premium. Finally, we provide evidence suggesting that correlation risk may be priced in the cross-section of stock returns.

Correlations play a central role in financial markets. There is considerable evidence that correlations between asset returns change over time[1] and that stock-return correlations increase when returns are low.[2] A market-wide 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 market-wide correlation risk is priced in the sense that assets that pay off well when market-wide correlations are higher than expected (thus providing a 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 and will appear expensive when correlation risk is priced.

This is the first paper to analyze whether cross-sectional differences in exposure to market-wide 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-Scholes (1973) and Merton (1973a) implied volatility functions.[3] We add formal evidence that individual options do not embed a variance risk premium, nor earn economically significant returns in excess of the one-factor Black-Scholes model. By considering individual options on all index components, our analysis emphasizes that a challenge in option pricing concerns reconciling the evidence of a second priced risk factor in index options (besides market risk) with the opposite evidence for individual options. 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

[1] See Bollerslev, Engle and Woolridge (1988) and Moskowitz (2003), among others. Brandt and Diebold (2003) 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.

[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 (2005).

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 five different ways.

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 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 risk premium on individual variance risk. As the decomposition shows, these two findings are only consistent with each other in a risk-based model if exposure to systematic 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 options and buys individual options and stocks in order to hedge individual variance risk and stock market risk, respectively. This trading strategy offers a remarkable risk-return trade-off. Its Sharpe ratio is 5 times larger than the one for bearing stock market risk in our sample. Correcting for standard risk factors, we find a highly significant excess return of roughly 20% per month. This is direct evidence of a very large correlation risk premium. We demonstrate that this strategy has more attractive risk-return properties than the option-based

[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.

[6] The relevant literature includes Ait-Sahalia and Kimmel (2005), Andersen, Benzoni and Lund (2002), Bakshi and Kapadia (2003a), Bollerslev, Gibson and Zhou (2004), Bondarenko (2004), Broadie, Chernov and Johannes (2005), 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.

trading strategies that have been suggested in the literature (like selling index puts or selling market variance), also when considering higher moments of the return distributions.

Third, 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 market-wide 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 and 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 83% of the cross-sectional variation in CAPM residuals of index and individual option returns. The estimated correlation risk premium is large (-15% per month) and significant. Exposure to individual variance risk is not priced in this cross-section, in line with our other results.

Fourth, motivated by these findings, we develop a simple option pricing model that incorporates stochastic correlations and a correlation risk premium. Each individual stock price follows a standard Ito process, and the correlation between stock returns is stochastic and mean-reverting. The model produces endogenous stochastic index-return volatility, even with constant individual volatilities. Moreover, we allow for a negative risk premium on correlation risk. Formally, the negative correlation risk premium generates higher expected correlation paths under the risk-neutral measure than under the actual measure. We prove that this divergence in expected correlations under the two measures results in option-implied correlations exceeding average realized correlations between equity returns. Another important implication of this model is that expected index option returns are much more negative than expected individual option returns.

To evaluate this model empirically, we calculate each day an option-implied correlation from 30-day index and individual options. The time-series for this implied correlation is intimately related to realized equity return correlations, both for levels and changes. Moreover, we show that implied correlations predict realized correlations, which supports the interpretation of the implied correlation as a measure of risk-neutral expected future correlations. Most importantly, we find a systematic difference between implied and realized correlations. The average implied correlation is 46.7%, while realized correlations are on average 28.7%. This provides further evidence for a large correlation risk premium and offers an intuitive perspective on the 'overpriced put puzzle': index puts are priced as if correlations between stocks in the index are on average 63% higher than seems historically the case.

We estimate the parameters of the model with GMM, matching the time-series properties of stock returns and of implied and realized correlations. Given the parameter estimates, we then analyze and formally test the implications of the model for index and individual option returns. The model gives a very good fit of expected index option returns, across the entire moneyness spectrum. In contrast, when we set the orthogonal correlation risk premium to zero, the model generates expected index option returns that are much less negative than the observed returns.

Fifth, we study the role played by exposure to correlation and individual variance risk in the cross-section of stock returns. While the analysis of index and individual options naturally forms the main focus of the paper, the cross-section of stock returns provides a useful robustness check of our results. Moreover, Ang et al. (2006) find that cross-sectional differences in expected stock returns are related to cross-sectional differences in exposures to market variance risk. Since we find strong evidence from options that market variance risk is priced because of priced correlation risk, not because of priced individual variance risk, we analyze whether the same holds for the cross-section of stock returns. Interestingly, we find that sensitivities to market variance are almost perfectly correlated with sensitivities to correlation risk, and essentially unrelated to loadings on individual variance risk. This strongly suggests that if market variance risk is priced in the cross-section of stock returns (as Ang et al. (2006) show in detail), it is driven by priced correlation risk, exactly as we found for options. For our short sample period, we do not obtain significant estimates of risk premia on either market variance risk or on market-wide correlation risk in stocks.

In sum, our findings strongly suggest that correlation risk is priced while individual variance risk is not. As a potential theoretical explanation for this, we explore the implications of Merton's ICAPM (1973b), in particular the extension of Chen (2003) that allows for time-varying market variance and shows that innovations in state variables that predict market variance carry a risk premium. Empirically, we find that both correlation and individual variance levels predict market variance. The ICAPM thus provides a justification for the existence of a correlation risk premium, but does not explain the lack of a risk premium on individual variance risk.

Several articles have investigated the correlation structure of interest rates of different maturities. Longstaff, Santa-Clara and Schwartz (2003), De Jong, Driessen and Pelsser (2004) and Han (2006) 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 the dynamics of the implied correlation measure with one 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.

Finally, it is interesting to note that practitioners have recognized the possibility of trading the difference between implied and historical correlations, 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 trading the difference between realized and implied correlations has been introduced, namely the correlation swap.

The paper is organized as follows. Section 1 presents the general decomposition of index variance risk. The data is described in section 2. Section 3 provides empirical evidence on variance and correlation risk premia, based on the framework of section 1. Section 4 develops, estimates and tests a model of priced correlation risk, focusing on its implications for option-implied versus realized correlations and expected returns on index versus individual options. In section 5 we study whether priced correlation risk can account for the empirical cross-section of stock returns and analyze the predictive power of average correlations and individual variances for future market returns and variance. We also discuss our findings in light of the ICAPM. Section 6 concludes.

1 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 similarly be decomposed. This section also briefly discusses the model-free implied variance estimator used in our empirical analysis in section 3.

1.1 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, unlike existing work, which does not entertain the possibility of priced correlation risk.

The stock market index is composed 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.[7] The instantaneous correlation between Wiener processes B_i and B_j that drive stocks i and j is

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

While the specific model in section 4 will impose more structure on the dynamics of $\rho_{ij}(t)$, we now 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.

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

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

It is clear from (2) that index variance changes are driven by shocks to individual variances $\phi_i^2(t)$ and to 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 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 $E_t^Q\left[d\phi_I^2\right] - E_t^P\left[d\phi_I^2\right]$.[8] Given constant index weights $\{w_i\}$ and defining $\iota_i \equiv w_i^2 + \frac{1}{2}\sum_{j\neq i} w_i w_j \frac{\phi_j}{\phi_i}\rho_{ij}$, applying Ito's lemma to (2) shows that

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

[8] This definition represents the total variance risk premium, i.e. 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.

$$E_t^Q\left[d\phi_I^2\right]-E_t^P\left[d\phi_I^2\right]=\sum_{i=1}^{N}\iota_i\left\{E_t^Q\left[d\phi_i^2\right]-E_t^P\left[d\phi_i^2\right]\right\}+\sum_{i=1}^{N}\sum_{j\neq i}w_iw_j\phi_i\phi_j\left\{E_t^Q\left[d\rho_{ij}\right]-E_t^P\left[d\rho_{ij}\right]\right\} \tag{3}$$

In words, the index variance risk premium reflects all individual variance risk premia $E_t^Q\left[d\phi_i^2\right]-E_t^P\left[d\phi_i^2\right]$, as well as correlation risk premia $E_t^Q\left[d\rho_{ij}\right]-E_t^P\left[d\rho_{ij}\right]$.[9] Below, we first present a detailed study of index and individual variance risk premia, i.e. 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 final sum in equation (3), i.e. on correlation risk premia.[10] We then 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.

1.2 Model-Free Implied Variances and Variance Risk Premia

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

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

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 the correct measure of the (total) risk-neutral expected integrated variance

[9] 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 $N=100$. Furthermore, we incorporate stochastic index weights in the calibration and estimation in section 4.

[10] Empirically, the time-series average of ι_i is positive for all stocks that make up the S&P100 index over our 8 year sample, with one exception. Moreover, the single stock for which the time-series average of ι_i is negative has an insignificant individual variance risk premium.

in a jump-diffusion setting. The measure is therefore considered 'model-free', and can be labeled the *model-free implied variance* ($MFIV$).

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

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

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, Gibson and Zhou (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 3.2.

2 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

[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.

due to changes in prices is small and can be neglected for our purposes (see also footnote 9).

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. We focus on short-maturity options, which are known to trade most liquidly, and take the discrete time interval τ over which we calculate the risk-neutral variances to be 30 calendar days. As is common in the literature, we focus on put options. OptionMetrics provides daily interpolated implied volatility surfaces for each security, calculated with a kernel smoothing algorithm from option prices, for a maturity of 30 calendar days, at deltas of -0.2 to -0.8 (in increments of -0.05). To calculate the option price from the data, we first obtain the option moneyness by inverting the Black-Scholes formula for delta and then use Black-Scholes to calculate the option price.[12] For the riskfree rate and dividend rate for the index, we utilize OptionMetrics data on forward prices, the zero-coupon curve and the index dividend yield. To infer the dividend rate for each stock, we combine the forward price with the associated spot rate. OptionMetrics uses a tree method with discrete dividends to price options, so before using the Black-Scholes formula, we calculate the mean continuously compounded dividend rate for each stock by averaging the implied OptionMetrics dividends. With this panel of option prices, we implement the model-free implied variance measure of section 1.2 following the procedure in Jiang and Tian (2005), suitably adjusted for put options.[13]

To calculate the realized variance, we use daily returns from CRSP for individual stocks and from OptionMetrics for the S&P100. We calculate the realized variance over 30-day windows, requiring that the stock has at least 15 non-zero return observations. Since the window spans 30 calendar days, this means that we require 15 observations out of approximately 22 trading days.

3 Evidence on Variance and Correlation Risk Premia

Based on the general framework of section 1, 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

[12] Jiang and Tian (2005) argue that using the interpolated volatility surface to calculate option prices, as we do here, is superior to working with interpolated option prices. Unreported results with raw (non-interpolated) option prices show that our main findings are unaffected.

[13] 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.

variance swap returns. These tests are conducted using the model-free implied variance of section 1.2. In light of the general decomposition of index variance risk above, this analysis provides indirect evidence on the importance of priced correlation risk. We then present direct evidence of a risk premium on correlation risk by developing and implementing a simple option-based trading strategy that exploits priced correlation risk. Finally, we test whether correlation and individual variance risk is priced in a well-chosen cross-section of assets (individual and index options).

3.1 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, Kapadia and Madan (2003), Bollen and Whaley (2004), Branger and Schlag (2004), Dennis and Mayhew (2002), Dennis, Mayhew and Stivers (2005) and Garleanu, Pedersen and Poteshman (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 1.2. 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.

Figure 1A 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 variance is higher than realized 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 1 reports an average (annualized) realized index volatility of 21.02%, while the $MFIV$ average is 24.79%. The null hypothesis that implied and realized variance are on average equal is very strongly rejected, based on a t-test with Newey-West (1987) autocorrelation consistent standard errors for 22 lags.

Turning to the equally-weighted average of the individual options in Figure 1B, there is, quite remarkably, little systematic difference between the two volatility proxies. On average, the square root of realized variance (39.91%) actually exceeds the square root of implied variance (38.56%).

The null hypothesis that, on average across all stocks in the index, the implied and realized variance are equal 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 is not rejected at the 5% confidence level for 108 stocks out of the 135 stocks that are included in the sample for this analysis. Of the remaining 27 stocks, 16 exhibit a significant positive difference between implied and realized variance. There is therefore very little evidence for the presence of a (negative) variance risk premium in individual stock options.

This is quite surprising, given the well-known empirical regularity for index options. Bakshi and Kapadia (2003b) found a difference of 1% to 1.5% (depending on the treatment of dividends) between the average implied and the average historical volatility in their 1991-1995 sample of 25 individual stock options. They also stressed 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.[14]

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, 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.

3.2 The Cross-Section of Individual Variance Swap Returns

We found that the total individual variance risk premium in individual options is not significantly different from zero. To gain further insight into this important result, we now 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.

[14] 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.

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 $MFIV_a(t) \equiv \sigma^2_{MF,a}(t)$ defined in section 1.2, 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 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 at daily frequency for all stocks with at least 600 observations during the sample period. As before, we use Newey-West (1987) standard errors for 22 lags, since we work with overlapping monthly 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 2-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 the majority of first-step betas is statistically insignificant. The common individual variance risk factor $\sum_{i=1}^{N} w_i r_i(t)$ is not priced either in individual options. The individual variance risk premium is found to be insignificant. Surprisingly, the point estimate is actually positive (0.0432 with a t-statistic of 1.14). Including both factors simultaneously leads to the same finding. 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 section 3.1 of a zero total individual variance risk premium. Combined with the finding of priced market variance risk in index options, this provides more indirect evidence of priced correlation risk.

3.3 Risk and Return of a Correlation Trading Strategy

Motivated by the finding that index variance risk is priced, in contrast to individual variance risk, we now explicitly examine whether exposure to correlation risk is priced. We do this by constructing a trading strategy that 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.

3.3.1 Trading Correlation Risk

We derive a correlation trading strategy based on the general framework of section 1 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):

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

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 market-wide correlation changes. The specific process for ρ can be left unspecified for now, 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 that 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, i.e. $I = \frac{1}{N}\sum_{i=1}^{N} S_i$, when initiating the trading strategy.

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

$$d\phi_i^2 - E\left[d\phi_i^2\right] = \varsigma_i(\phi_i)\,dB_{\phi_i} \tag{7}$$

Denoting the price of a put on asset $a \in \{I, 1, ..., i, ..., N\}$ by P_a, we focus on unexpected put returns:

$$\frac{dP_i}{P_i} - E\left[\frac{dP_i}{P_i}\right] = \frac{S_i}{P_i}\frac{\partial P_i}{\partial S_i}\phi_i dB_i + \frac{1}{P_i}\frac{\partial P_i}{\partial \phi_i^2}\varsigma_i(\phi_i)\,dB_{\phi_i} \tag{8}$$

and

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

The correlation trading strategy aims to short the index put 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}.[15] We first hedge the individual variance risk. When shorting index puts worth 100%

[15] Note that while dB_ρ does not enter equation (8) directly, dB_ρ may still impact individual put returns through 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.

of initial wealth, the portfolio weight y_i in individual put i is then obtained by solving

$$-\frac{1}{P_I}\frac{\partial P_I}{\partial \phi_i^2}\varsigma_i\left(\phi_i\right)+y_i\frac{1}{P_i}\frac{\partial P_i}{\partial \phi_i^2}\varsigma_i\left(\phi_i\right)=0 \tag{10}$$

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

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

$$-\frac{S_i}{P_I}\frac{\partial P_I}{\partial S_i}\phi_i+y_i\frac{S_i}{P_i}\frac{\partial P_i}{\partial S_i}\phi_i+z_i\phi_i=0 \tag{11}$$

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 puts worth all initial wealth and invests a fraction y_i of initial wealth in each individual put and z_i of initial wealth in each individual stock, the remainder being invested in the riskfree 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 $\partial P_i/\partial S_i$ and $\partial P_i/\partial \phi_i^2$, evaluated at the implied volatility of the option. For the 'correlation-vega' of the index put, we combine the Black-Scholes vega with equation (2).[16]

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 price, delta and vega of the index put depend on ρ. We calculate these numerically for different levels of ρ, using 30-day average lagged historical correlations as an estimate for ρ at each point in time.

3.3.2 Empirical Results

While the trading strategy requires, theoretically speaking, only one strike level, we implement the strategy empirically for different moneyness levels. However, to ensure that the individual and index puts that are used are comparable in terms of moneyness, we implement the strategy across a range of Black-Scholes deltas rather than strike-to-spot ratios.[17] We report results for each delta,

[16] We obtain very similar results when using the specific correlation model of section 4 (with endogenous stochastic market variance due to correlation risk) to calculate $\partial P_I/\partial\rho$.

[17] 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 vastly different volatilities for example.

as well as for an equally-weighted portfolio across all deltas.

Table 2A reports the first 4 moments for the trading strategy return in excess of the riskfree rate, estimated from monthly overlapping returns at a daily frequency. The excess returns range from 10% to 43% per month. The annualized Sharpe ratio is extremely high and always exceeds 1. Compared to the annualized Sharpe ratio for the S&P100 index itself (0.286 over the sample period), the Sharpe ratio on the trading strategy is for most deltas actually 4 to 5 times higher. Although the trading strategy is theoretically speaking 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 (Goetzmann et al. (2002)). We do this in several different ways.

First, we estimate the CAPM beta and alpha of the strategy. For the equally-weighted portfolio, the CAPM beta is 3.13 (t-statistic of 3.16) and the CAPM alpha is 18.85% per month (t-stat of 5.00). While the beta estimate reflects that the trading strategy is model-based and imperfectly hedged, the beta is far from extreme for a strategy based on derivatives.[18] More importantly, the CAPM alpha is highly significant (for all individual deltas), both statistically and economically, and indicates that the high return on the trading strategy cannot be justified by exposure to stock market risk. In the Fama-French model and the 4-factor model that adds momentum, the loadings on any of the additional factors are insignificant and the strategy has even higher alphas (20% and 21.5% with t-statistics of 6.1 and 6.7, 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, while our analysis so far used overlapping monthly returns. At the monthly frequency, correcting the trading strategy return for market risk, Fama-French factors, momentum and the liquidity risk factor gives an alpha of 24.4% per month with a t-statistic of 5.44. The exposure of the trading strategy to liquidity risk is positive, but insignificant.

Secondly, to put the higher moments into perspective, we compare with the summary statistics for a strategy that consists of writing 1-month index puts (Table 2B). For each delta, the correlation

[18] For example, an equally-weighted portfolio of 1-month index puts with deltas ranging from -0.8 to -0.2 has a CAPM beta of -17.85.

trading strategy has a substantially higher Sharpe ratio and exhibits remarkably less negative skewness. For deep OTM puts, the difference in Sharpe ratio between our strategy and the short index put is somewhat smaller, but the negative skewness and kurtosis is actually more than 50% higher for the put-only strategy.

The final column of Table 2B reports the Sharpe ratio obtained by the in-sample optimal mean-variance portfolio, i.e. with constant weights in the index put portfolio, individual put portfolio and stock market index chosen to maximize the Sharpe ratio. The maximum in-sample Sharpe ratio that can be achieved in this way is always lower than the Sharpe ratio of our correlation strategy. For the equally-weighted portfolio across all deltas, the correlation strategy improves the maximal in-sample Sharpe ratio by 21%. Since the improvement over the tangency portfolio stems from having time-varying portfolio weights in the correlation strategy, this indicates that our approach is able to capture relevant time-variation.

Finally, to further gauge the economic significance of these results, we report the certainty equivalent wealth that a CRRA investor, who is already investing in the market index, is willing to pay in order to gain access to the correlation trading strategy. Doing this for a CRRA investor rather than for mean-variance preferences penalizes for negative skewness and high kurtosis. Panel C reports the results, starting with the certainty equivalents for two alternative strategies, namely a (synthetic) index variance swap and an equally-weighted portfolio of index puts (as in the last row of Table 2B). The index variance swap 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 in the literature (e.g. Bondarenko (2003b)). This is also the case for our sample period: a CRRA investor with risk-aversion coefficient $\gamma = 5$ and a 1-month horizon is willing to pay 0.22% of initial wealth per month to gain access to the index variance swap and 0.56% in the case of the index put portfolio. The certainty equivalent for the correlation trading strategy is 5 to 6 times larger than for the index variance swap and always more than double the index put certainty equivalent. For example, the $\gamma = 5$ investor gains 1.26% of initial wealth per month when having access to the correlation trading strategy. The optimal weight for this investor is 11.96% with a t-statistic of 4.82. The optimal portfolio weight in the correlation strategy is always highly significant, for all values of γ, with t-statistics between 7.5 ($\gamma = 1$) and 4.4 ($\gamma = 20$).

The final column ($\gamma = 1.75$) of Table 2C is of particular interest. It considers the investor who optimally holds the market (equity weight of 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 tradeoff of the strategy could arise in an equilibrium with a CRRA representative investor (if the optimal weight is zero) or instead whether the correlation risk premium represents a 'good deal' (Santa-Clara and Saretto (2005)). The investor stands to gain 3.22% of initial wealth per month from the correlation strategy, based on an optimal correlation-strategy weight of 29% of initial wealth (t-statistic of 6.7) and a negative (but insignificant) equity weight. It is clear that the risk-return tradeoff of the correlation trading strategy is a 'very good deal' for a $\gamma = 1.75$ 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 highly significant. The risk-return tradeoff is actually substantially more generous than what can be obtained with short positions in index puts or in market variance.

3.4 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 CAPM excess returns on individual and index options.

Studying the cross-section of expected index and individual option returns presents several additional benefits. First, our test procedure is identical to standard procedures used in asset pricing to test for the presence of priced risk factors. Applying these techniques to option returns avoids the need for specific parametric modeling assumptions that are otherwise needed when testing option pricing models. Therefore, before developing a specific model of priced correlation risk in section 4, we first 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 this section.

Our cross-section contains the same 13 index puts as in Table 2B (1-month puts with deltas ranging from -0.8 to -0.2), as well as 13 individual puts with these characteristics. We calculate option returns as holding-period returns, i.e. the return at time $t+\tau$ on a put written at t is given by the option payoff at maturity $(t+\tau)$ divided by the option price at t. For each delta, we average the individual option returns cross-sectionally using the index weights for each day. To account for serial dependence in the overlapping monthly returns, t-statistics are based on Newey-West (1987) standard errors with 22 lags.

We use the standard 2-step procedure for cross-sectional asset pricing, estimating first 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 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 puts we study.

3.4.1 CAPM Results

For index options, the CAPM betas range from -13 to -20 and are all highly significant. The betas for individual options are somewhat smaller (ranging from -6 to -13), 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 -43% per month (t-stat of -3.96). All alphas are negative for index puts and average -13% per month.

For individual puts, the results are completely different. No time-series alpha is significantly negative and the average alpha is much smaller (1.39% per month). Even though all individual alphas are economically small, for 4 individual puts we obtain a statistically significant positive alpha. This is due to the averaging of individual put returns across more than 100 stocks, thus leading to much more precise estimates than for index puts. Economically speaking, the contrast between index and individual puts is clear: the average of the significant alphas is -30% for index puts, but only 3% for individual puts.

In summary, unlike for index options, the CAPM does quite well for individual options. The evidence against it is economically small and points to a positive rather than negative individual variance risk premium, consistent with average realized variance slightly exceeding average risk-neutral variance in section 3.1 and a positive (but insignificant) risk premium in section 3.2. Section 4 shows that all individual-option alphas become insignificant when using theoretical (based on Black-Scholes) rather than estimated betas.[19]

3.4.2 Results for Correlation and Individual Variance Risk Factors

We now study whether the CAPM excess returns on index and individual puts can be explained by their exposure to a correlation risk factor and an individual variance risk factor. The return on the correlation trading strategy (equally-weighted across deltas) is taken as the correlation risk factor. For the individual variance risk factor, we use an equally-weighted portfolio (across deltas) of index-weighted individual put returns. In a first step we regress CAPM residuals on both factors. Given the estimated factor loadings, we obtain the correlation and individual variance risk premium from the cross-section of average CAPM excess returns.

CAPM excess returns on index puts exhibit large and highly significant loadings on the correlation risk factor. The correlation betas range from -0.26 (ITM index put) to -1.83 (OTM index put). All index puts have correlation betas with t-statistics above 8.9 in absolute value. Individual puts have much smaller correlation loadings, ranging from a minimum of -0.03 to a maximum of 0.035. In contrast, CAPM residuals of index and individual puts exhibit remarkably similar sensitivities to individual variance risk, with loadings between 0.4 and 1.35 for index puts and between 0.41 and 1.75 for individual puts. All are statistically significant.

For index options, the two-factor model generates time-series alphas that have a mean of 6% and a mean absolute value of 7.8%. The CAPM generated -12.9% and 12.9%, respectively. Accounting for exposure to correlation risk and individual variance risk leads to a notable reduction in mispricing and index options no longer seem 'overpriced' (except for deep OTM puts). The improvement for individual options is very small, as the CAPM already performs quite well (the mean absolute α goes from 2.6% to 2.2%).

Table 3 presents the results for the cross-sectional regression of average CAPM residuals of index and individual puts on their factor loadings. The risk premium for the correlation factor is estimated to be 15% per month (t-statistic of 4.31). In contrast, the price of individual variance risk is small and statistically insignificant. Note that the positive risk premium for the correlation

[19] 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.

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 puts are always negative, the positive estimate in Table 3 leads to negative excess returns for index put options relative to the CAPM, i.e. consistent with the definition in section 1, the price of correlation risk is negative in the sense that assets with payoffs that covary positively with correlation (e.g. index puts) earn negative excess returns.

The estimated risk premia are close to the time-series averages for the factor CAPM-excess returns (18.85% and 1.39%). The empirical success of the model is also apparent from the high cross-sectional R^2 of 82.6%. This is particularly high, given that the cross-section concerns CAPM residuals and that the model imposes linearity. Section 4 studies priced correlation risk in a model that incorporates inherent nonlinearities.

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 3.1 and 3.2. We 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.

4 A Model of Priced Correlation Risk

Motivated by the empirical findings of the previous section, we develop a detailed model of priced correlation risk, showing how individual and index options can be combined to obtain option-implied correlations and then test the predictions of the model for the cross-section of expected index and individual option returns.

4.1 Individual Stock Price Processes

Under the physical probability measure P, the price of stock i, S_i, is assumed to follow an Ito process with expected return μ_i and possibly stochastic diffusion $\phi_i(t)$:

$$dS_i = \mu_i S_i dt + \phi_i S_i dB_i \tag{12}$$

where B_i is a standard scalar Wiener process. The special case where $\phi_i(t)$ is constant simplifies (12) to the standard Black-Scholes set-up. More generally, the instantaneous variance $\phi_i^2(t)$ is an Ito process, driven by a standard scalar Wiener process B_{ϕ_i}, which is taken to be uncorrelated with

B_j for all j.[20] Importantly, $\phi_i^2(t)$ follows the same process under the physical probability measure P as under the risk-neutral probability measure Q, so that individual volatility risk is not priced. This assumption is motivated by the empirical results in section 3.

We continue to focus on a single-factor model of correlation risk, i.e. we assume that equation (6) holds. Furthermore, we make the following simplifying homogeneity assumption:

$$\overline{\rho_{ij}} = 1, \ \forall i \text{ and } j. \tag{13}$$

While this assumption is admittedly restrictive, it allows us to capture pervasive correlation risk in the most parsimonious way. Moreover, the homogeneity assumption is instrumental in guaranteeing a positive-definite correlation matrix. Also, we are interested in market-wide average correlation dynamics rather than in cross-firm differences in correlations. Finally, despite this simplifying assumption, the model is shown to perform very well empirically and option-implied correlations that are estimated using (13) have predictive power for future realized correlations, although these are estimated without restriction (13).

We impose as initial condition $\rho(0) \in (0,1)$. Under measure P, the correlation state variable $\rho(t)$, which is also the instantaneous correlation because of (13), is assumed to follow a mean-reverting process with long-run mean $\overline{\rho}$, mean-reversion parameter λ and diffusion parameter σ_ρ:[21]

$$d\rho = \lambda(\overline{\rho} - \rho)\,dt + \sigma_\rho \sqrt{\rho(1-\rho)}dB_\rho. \tag{14}$$

Because of the $\sqrt{\rho(1-\rho)}$ factor, the process is of the Wright-Fisher type, used extensively in genetics (see e.g. Karlin and Taylor (1981)), and also in financial economics (Cochrane, Longstaff and Santa-Clara (2004)). We multiply the diffusion parameter by $\sqrt{\rho(1-\rho)}$ to ensure that under certain parameter restrictions $\rho(t)$ remains between zero and one with probability 1, as is shown by lemma 1 in appendix B.[22] This guarantees the positive definiteness of the resulting correlation and variance-covariance matrices (proposition 1, appendix B).

[20] This means that there is no asymmetric volatility or 'leverage effect' (Black (1976)) at the level of individual stock returns, in line with the results in section 3.2. However, our correlation-risk model nonetheless generates an endogenous leverage effect for index returns. This also is consistent with the empirical findings of Dennis, Mayhew and Stivers (2005), who obtain a much smaller leverage effect for individual stock returns than for the index.

[21] Collin-Dufresne and Goldstein (2001) discuss a similar model of correlation dynamics.

[22] We could have accommodated negative values for ρ, but since ρ is the average correlation between the stocks in the index, this generalization is not needed empirically.

4.2 Priced Correlation Risk

Under the risk-neutral measure Q, the instantaneous correlation ρ follows the same process as in (14), but with a correlation risk premium proportional to κ subtracted from its drift:

$$d\rho = \left[\lambda\left(\overline{\rho} - \rho\right) - \kappa\sigma_\rho\sqrt{\rho\left(1-\rho\right)}\right]dt + \sigma_\rho\sqrt{\rho\left(1-\rho\right)}dB_\rho^Q \tag{15}$$

A negative value for κ implies a negative correlation risk premium, so that the expected path of future correlations under Q exceeds the expected correlation path under P.[23] Assets that load positively on correlation risk earn a negative instantaneous excess return κ per unit of (standardized) correlation risk exposure. Intuitively, the negative excess return represents the insurance premium paid for assets that hedge against unexpected correlation increases.

If correlation risk is correlated with market risk, the total correlation risk premium κ includes compensation for market risk. A first component of κ therefore equals $\psi\eta$, where η denotes the Sharpe ratio for market risk and ψ is the correlation between the Brownian motion B_ρ driving $\rho\left(t\right)$ in (14) and the market risk factor in the pricing kernel. We will later estimate ψ. A negative value would be in line with empirical work that finds that correlations increase when prices decline.

Importantly, we found in sections 3.3 and 3.4 that positive exposure to correlation risk is associated with negative excess returns relative to what is justified by exposure to market risk. Put differently, we find empirically that the total correlation risk premium exceeds $\psi\eta$ or that the part of market-wide correlation risk that is orthogonal to stock market risk is also priced. In that case, correlation risk constitutes a second priced risk factor in the economy. Denoting the compensation for orthogonal correlation risk by κ_o, the total correlation risk premium κ can then be decomposed (through an orthogonal decomposition of the Brownian motion dB_ρ) in a part driven by the equity risk premium and an orthogonal part as follows

$$\kappa = \psi\eta + \sqrt{1-\psi^2}\kappa_o. \tag{16}$$

Our empirical estimates will show that the orthogonal part constitutes by far the biggest component of the price of correlation risk. Section 5 discusses in detail how a nonzero orthogonal correlation risk premium κ_o might arise in equilibrium.

[23] Lemma 2 (appendix B) states the conditions under which the process $\rho(t)$ remains within interval $(0, 1)$ under measure Q. These conditions nest the ones listed in Lemma 1, so that, if satisfied (as we assume and as is valid for our empirical parameter estimates), measures P and Q are equivalent.

Both the existence of correlation risk and of a correlation risk premium have fundamental implications for the pricing of index options. Even in a simplified version of our model with standard Black-Scholes dynamics for all individual stocks, we endogenously generate stochastic index variance. Furthermore, because individual variance risk is not priced in our model, the index variance risk premium is entirely due to priced correlation risk.

4.3 Risk-Neutral Expected Average Correlation

We now introduce the concept of the risk-neutral expected average correlation. The idea is to infer the expected average future correlation between N stocks from the expected future variances of these N stocks and from the expected future variance of the stock index.

From the definition of the market index, (4) can be written out for $a = I$ as

$$\sigma_I^2(t) = E_t^Q\left[\int_t^{t+\tau}\sum_{i=1}^N w_i^2\phi_i^2(s)\,ds\right] + E_t^Q\left[\int_t^{t+\tau}\sum_{i=1}^N\sum_{j\neq i} w_i w_j \phi_i(s)\,\phi_j(s)\,\rho(s)\,ds\right]. \tag{17}$$

Rather than attempting to extract the entire path of future correlations $\rho(s)$ from (17), we aim to obtain a 'certainty equivalent' of the future stochastic correlations, which we call the *risk-neutral expected average correlation* ($RNEAC$) and which is constant over the $[t, t+\tau]$ time interval. $RNEAC$ is defined as

$$RNEAC(t) \equiv \frac{E_t^Q\left[\int_t^{t+\tau}\phi_I^2(s)\,ds\right] - \sum_{i=1}^N w_i^2 E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}{\sum_{i=1}^N\sum_{j\neq i} w_i w_j E_t^Q\left[\int_t^{t+\tau}\phi_i(s)\,\phi_j(s)\,ds\right]}. \tag{18}$$

$RNEAC$ is a certainty equivalent to the correlation process $\rho(s)$ under the risk-neutral measure Q over the $[t, t+\tau]$ time interval in the sense that it yields the same risk-neutral expected integrated index variance as when the entire stochastic process for $\rho(s)$ is used.

Index and individual option prices can be used to estimate the risk-neutral integrated variances, as explained in 1.2. The risk-neutral expectation in the denominator, however, is not observed, since it requires the instantaneous volatilities $\phi_i(s)$ and $\phi_j(s)$. As an alternative to $RNEAC(t)$, we calculate instead the *implied correlation* $IC\left(t\right)$, defined as

$$IC\left(t\right)\equiv\frac{E_t^Q\left[\int_t^{t+\tau}\phi_I^2\left(s\right)ds\right]-\sum\limits_{i=1}^N w_i^2E_t^Q\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}{\sum\limits_{i=1}^N\sum_{j\neq i}w_iw_j\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}}\tag{19}$$

This measure of the option-implied correlation is readily estimated and is useful for the following reasons. First, IC is closely related to $RNEAC$. Lemma 3 (Appendix B) establishes that $IC(t)\leq RNEAC\left(t\right)$. Secondly, the first main contribution of this paper is to provide direct evidence on the importance of a negative correlation risk premium. This is achieved by comparing estimates of $IC(t)$ with estimates of realized correlations. More precisely, we can take expectations under measure P rather than under Q in equation (19) so as to capture the actual expected average correlation. Since this will later be estimated from the (cross-sectionally weighted average of) realized correlations, we call this in short the realized (or historical) correlation $RC\left(t\right)$, defined as:

$$RC\left(t\right)\equiv\frac{E_t^P\left[\int_t^{t+\tau}\phi_I^2\left(s\right)ds\right]-\sum\limits_{i=1}^N w_i^2E_t^P\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}{\sum\limits_{i=1}^N\sum_{j\neq i}w_iw_j\sqrt{E_t^P\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}\sqrt{E_t^P\left[\int_t^{t+\tau}\phi_i^2\left(s\right)ds\right]}}\tag{20}$$

Proposition 2 states that the difference between IC and RC is crucially linked to the price of correlation risk.

Proposition 2: *Given a set of fixed index weights* $\{w_i\}$*, individual stock price processes given by (12), (1), (6) and (13) with non-priced individual volatility risk, and a correlation process given by (14) under* P *and by (15) under* Q*, the difference between IC and RC can be written as*

$$IC\left(t\right)-RC\left(t\right)=\Xi\left(t\right)\left[\int_t^{t+\tau}\Psi\left(s\right)\left\{E_t^Q\left[\rho\left(s\right)\right]-E_t^P\left[\rho\left(s\right)\right]\right\}ds\right]\tag{21}$$

where $\Xi\left(t\right)>0$ *and* $\Psi\left(s\right)>0,\forall s$. $IC\left(t\right)-RC\left(t\right)=0$ *if and only if* $\kappa=0$ *and is strictly increasing in* -κ.

Testing whether correlation risk is priced can therefore be done by testing whether $IC=RC$. Note that the proposition requires that individual stock return variance risk is not priced. As can easily be seen from equations (19) and (20), $IC-RC$ is then proportional to $E_t^Q\left[\int_t^{t+\tau}\phi_I^2\left(s\right)ds\right]-E_t^P\left[\int_t^{t+\tau}\phi_I^2\left(s\right)ds\right]$, so that correlation risk is priced if and only if index variance risk is priced, in line of course with the general result in equation (3).

It follows from proposition 2 that when $\phi_i(t) = \phi_i$, $\forall i$, $IC(t) - RC(t)$ simplifies to

$$IC(t) - RC(t) = \frac{1}{\tau}\int_t^{t+\tau} \left\{E_t^Q[\rho(s)] - E_t^P[\rho(s)]\right\} ds \tag{22}$$

i.e. the difference in expected integrated correlation paths under Q and P. Interestingly, $IC(t)$ then gives the no-arbitrage correlation swap rate, fixed at t and to be paid at $t+\tau$ in exchange for the realized integrated correlation path.

4.4 Implied versus Realized Correlations

We estimate the implied correlation from option prices and show there is substantial correlation risk. We also investigate whether the implied correlation measure indeed captures the dynamics of historically observed correlations. We then present evidence that correlation risk is priced.

Given the time-series of model-free implied variances from index and individual options, the implied correlation $IC(t)$ is calculated for each day t as

$$IC(t) = \frac{\sigma^2_{MF,I}(t) - \sum\limits_{i=1}^{N} w_i^2 \sigma^2_{MF,i}(t)}{\sum\limits_{i=1}^{N} \sum_{j\neq i} w_i w_j \sqrt{\sigma^2_{MF,i}(t)}\sqrt{\sigma^2_{MF,j}(t)}}. \tag{23}$$

We obtain historical correlations between the stocks in the S&P100 index from CRSP returns. Because standard historical correlations are more commonly used and because the difference between historical correlations calculated from (20) and from the standard definition is very small and not economically significant, we calculate standard historical correlations. For each pair of stocks, we calculate the historical correlation at time t over a 30-day window, imposing the same requirements as for the calculation of realized variances. Repeating the procedure each day gives a moving estimate of the historical pairwise correlations, which can then be aggregated into a cross-sectional weighted average across all pairs of stocks, using the appropriate weights from the S&P100 index.

Figure 2 plots the time-series of the 1-week moving averages of the implied correlation $IC(t)$ defined in (23) and of the weighted-average historical correlation. The implied correlation is very volatile and peaks during financial crises. For instance the Asian crisis of 1997, the Russian crisis of 1998, or the event of September 11 2001 clearly show up as periods over which index options reflect high risk-neutral cross-stock correlations.[24] The fact that the 1-week moving average of

[24] The exchanges were closed from 11 to 16 September 2001, leading to missing data for that period.

the implied correlation fluctuates between 0.2 and 0.9 over our 8 year sample suggests that there is substantial correlation risk. It is noteworthy that the two series comove very strongly. Figure 2 shows that our implied correlation measure indeed captures the dynamics of the cross-sectional average of correlations between all stocks in the S&P100 index. The time-series correlation between the daily level of $IC(t)$ and of the historical correlation measure in Figure 2 is 68.78%. Table 4 presents further information about the two series. Of particular interest is the extent to which the implied correlation $IC(t)$ exceeds the historical correlation, since the difference between $IC(t)$ and $RC(t)$ reflects the price of correlation risk κ. Table 4 shows that the time-series average for the historical correlation is 28.68%, while the time-series average for $IC(t)$ is 46.68%, which suggests a large risk premium for correlation risk. In other words, index put options are on average priced as if the correlations between stocks in the index are 62.8% higher than seems historically the case.[25]

Table 4 also shows that the high time-series correlation between the implied correlation $IC(t)$ and the historical correlation (68.78%) is robust to whether we look at levels or quarterly changes. To further investigate the relation between implied and realized correlations, we study whether IC can predict future realized correlations. Table 4 reports the results from predictive regressions of cross-sectionally averaged realized correlations on lagged values of implied and of realized correlations. The predictive power of $IC(t)$ for future realized correlation is quite high, especially when taking into account that realized correlations are themselves rather noisy estimates: $IC(t)$ explains 35% of the variation in future realized correlations. Computing t-statistics based on Newey-West (1987) standard errors, the slope coefficients are statistically significant. Table 4 shows that realized correlations also predict future realized correlations. In the bivariate predictive regressions, the implied correlation continues to be a significant predictor.

4.5 Estimates and Tests of the Correlation Risk Model

We estimate the parameters of the correlation risk model using GMM and test its implications for expected index versus individual option returns.

[25] Because estimates of realized correlations based on daily returns may be affected by cross-autocorrelations (due to e.g. microstructure effects), we also use two-day and three-day returns for the estimation. This leads to almost identical (and in fact slightly lower) estimates of average realized correlations, namely 27.87% and 27.19%, respectively.

4.5.1 Estimation Methodology

For individual options, the results discussed in 3.1 show that there is no evidence for a systematic difference between implied and realized variance. In addition, there is little evidence that expected individual option returns deviate from the prediction of a one-factor Black-Scholes type model. This motivates us to assume that each individual stock price follows the standard geometric Brownian motion, as in (12) with constant diffusion coefficients. It follows immediately that individual implied variance equals on average realized variance in the model, in line with empirical observations.

The model contains 7 parameters: $\overline{\rho}$, λ, σ_ρ, κ, ψ, μ and ϕ. For simplicity, we assume that all 100 stocks are equally weighted in the index[26] and that each stock has the same volatility ϕ and expected return μ, thus extending the homogeneity assumption in (13) to all individual parameters. The objective of the estimation is to match the time-series properties of stock returns and of implied and realized correlations. We use 7 moment restrictions on these time-series properties and estimate the parameters using the special case of GMM where the number of parameters and moment conditions coincide. The moment conditions are the following: average equity index return (10.47% over our 1996-2003 sample period), individual stock volatility (38.56%, the square root of the mean individual $MFIV$ in Table 1), the average monthly realized correlation RC (28.7% in Table 4), the average monthly implied correlation IC (46.7% in Table 4), the time-series standard deviation of daily sampled $IC - RC$ (12.47%), the covariance between monthly changes in RC and its lagged level (-0.0055) and the covariance between monthly equity index returns and monthly changes in RC (-0.0011). It is intuitive to see how the parameters of interest are directly related to these moment conditions: μ and ϕ drive the average index return and individual stock volatility, respectively; $\overline{\rho}$ and κ are closely related to the average RC and IC; σ_ρ is related to the time-series standard deviation of $IC - RC$; λ reflects the mean-reversion tendency of the ρ process and the last moment condition is driven by ψ.[27]

The moment conditions ensure that our model matches the Sharpe ratio on the equity index as well as the empirical Sharpe ratio on a hypothetical correlation swap. Given the simplifying assumptions in this section, IC equals the no-arbitrage correlation swap rate, as discussed in section 4.3. The excess return on (a short position in) a correlation swap is then $IC - RC$, so that

[26] The index is equally-weighted at $t = 0$, but in the simulation these index weights change over time according to the simulated relative stock price movements.

[27] It is straightforward to show that the correlation between index return shocks and dB_ρ equals ψ (as $N \to \infty$) if the correlation between individual return shocks and dB_ρ is $\psi\sqrt{\rho}$, as we impose in the calibration.

its Sharpe ratio equals $(IC - RC) / StdDev(RC - IC)$.

We simulate the model under the P and Q measures to obtain these moment conditions and iterate until the model-implied moment restrictions exactly match the observed ones. Simulations under Q are based on the sample average for the 30-day riskfree rate of 4.46% per year. We simulate the model at the same daily frequency as the data are sampled. This is particularly important for the realized correlation, since some of the variation in RC is due to sampling error, i.e., due to the lack of continuous observations of RC. The estimation strategy thus corrects for such sampling error. GMM-style standard errors for the parameter estimates are readily obtained, and we use Newey-West (1987) to correct for autocorrelation with 12 monthly lags.

4.5.2 Parameter Estimates and Model Implications

The parameter estimates in Table 5A satisfy the parameter restrictions identified in lemmas 1 and 2 to guarantee that the correlation process remains within the interval $(0, 1)$. The correlation risk premium κ can be interpreted as the annualized continuous-time Sharpe ratio for correlation risk and is estimated at -6.6. This estimate is extremely high and ensures that the model matches the empirical monthly Sharpe ratio for a hypothetical correlation swap ($(0.287 - 0.467) / 0.1247 = -1.44$). The estimation procedure incorporates that the realized correlation is estimated using daily rather than continuous observations, which increases the variance of $IC - RC$ relative to the continuous-time estimate.

It is useful to compare with the Sharpe ratio we found in section 3.3 for the correlation trading strategy. The continuous-time Sharpe ratio estimate κ is substantially larger in absolute value (6.6 versus 1.44). The discrepancy has multiple sources. First, the model we consider here is stylized. For simplicity, the excess return and Sharpe ratio of the hypothetical correlation swap are obtained under the assumption of constant individual variances. In addition, the trading strategy in section 3.3 was implemented with monthly rebalancing. Higher-frequency rebalancing would be expected to further lower the risk of the strategy and to increase its Sharpe ratio.

The correlation process has considerable mean reversion, in line with the observed correlation behavior in Figure 2. The correlation between the stock index return and correlation changes is about -25%. The estimates clearly demonstrate that the correlation risk premium is almost entirely a reward for orthogonal correlation risk. Plugging $\psi = -0.253$ and $\eta = 0.286$ into the decomposition in (16), the part of the correlation risk premium that can be attributed to ρ being correlated with the market risk factor is negligible and represents about 1.1% of the total κ ($\frac{-0.0724}{-6.597} = 0.011$).[28]

An important and unique feature of our model is that correlation risk affects expected index option returns, but not individual option returns. Comparing the expected index versus individual option returns that are generated by the model with empirically observed expected returns is therefore an interesting test of the model. Figure 3 reports the results of this comparison and summarizes the main effects of our correlation model. Empirical estimates for unconditional expected option returns are obtained from the time-series averages of the holding-period returns that we introduced before (section 3.4). The graph first of all shows the implications for expected individual option returns. Given that individual stocks each follow a Black-Scholes process in the model, the generated expected individual option returns fully come from the negative exposure (or 'beta') of put options to diffusive shocks in the equity prices, combined with a positive annual equity premium of $10.47\% - 4.46\% = 6.01\%$. In line with the results in section 3.4, this beta-effect explains most of the negative expected returns on individual options in Figure 3, i.e. there is not much evidence that individual options embed risk premia on top of the standard equity risk premium.

Turning to the implications for index options, we first consider a case where the risk premium for orthogonal correlation risk (κ_o in decomposition (16)) is set to zero, so that κ only reflects the (small) 'ψ-part' of priced correlation risk. Even when this premium κ_o is zero, expected index option returns are more negative than individual option returns, since index options have a beta that is more negative. However, Figure 3 shows that a zero orthogonal correlation risk premium generates expected returns that are much closer to zero than the observed index option returns. Most importantly, Figure 3 presents expected index option returns for $\kappa = -6.6$ as estimated. The graph shows that this has a dramatic impact on the expected index option returns: the returns are much more negative in this case, and quite close to the observed average returns. Only for deep OTM put options is there some difference between the model-implied and observed values.

4.5.3 Statistical Model Tests

We test whether our correlation risk model can fit the difference between expected index and individual option returns across 13 deltas ($\Delta = -0.8$ to -0.2 in steps of 0.05). We focus on individual t-statistics per delta, because a Wald test suffers from the common problem that the 13 return differences are highly cross-correlated (the covariance matrix of option return differences

[28] The finding that κ is mainly compensation for orthogonal correlation risk is not dependent on the estimates of η and ψ. Even with $\eta = 0.5$ and $\psi = -0.999$ would orthogonal correlation risk represent more than 92% of κ.

across deltas is very close to singular), which inflates the test-statistic (see e.g. Cochrane (2001)). The test on expected-return differences incorporates that the risk premia on stock-market risk and on the correlation risk factor are estimated and corrects for serial correlation.

The results in Table 5B are clear: the correlation risk model offers an excellent fit and is never rejected according to the tests. We conduct the same tests also for the restricted version of the model where orthogonal correlation risk is not priced ($\kappa_o = 0$). Interestingly, in this case the individual t-statistics always reject the model, for all Δ's.

Finally, we turn to the fit of individual option returns. The individual t-tests validate our earlier findings that individual options do not embed additional risk premia. Unlike in section 3.4 where the betas are estimated, we now use the theoretical betas and find that all individual options have insignificant alphas. The predictions of the model for individual option returns are therefore not rejected either.

5 Correlation Risk and Expected Stock Returns

While our main focus is the pricing of correlation risk in index and individual options (given their obvious differences in exposure to correlation risk), we now study the role of market-wide correlation shocks in the empirical cross-section of stock returns. We also analyze whether market-wide correlations have predictive power for market returns and market variance and discuss how these and earlier findings are consistent with the ICAPM.

5.1 Correlation Risk and the Cross-Section of Stock Returns

Most work on priced variance risk studies index options. Ang et al. (2006), however, provide evidence that cross-sectional variation in expected stock returns is related to differences in sensitivities to market variance risk, while controlling for other well-documented sources of risk. Since our model and empirical findings clearly attribute the market variance risk premium to priced correlation risk (and not to priced individual variance risk), it is interesting to conduct a basic analysis of the ability of both factors to account for the pricing of market variance risk in empirical stock returns. This analysis leaves the confines of the parametric model of section 4, where the homogeneity assumption precludes any differences across stocks in their correlation sensitivities.

We have focused on options since their exposures to correlation risk are very clear (zero for individual options and of first-order importance for index options), so that the risk premia could be identified relatively easily, even with only 8 years of (high-frequency) data. For the cross-section of stocks, it seems unlikely that such clear differences in correlation exposures exist and we do not expect to find strong significance for the risk premia on variance and correlation risk based on 8 years of stock return data. We therefore perform the following analysis. We decompose variance dynamics into a correlation component and an individual variance component. We then analyze to what extent cross-sectional differences in exposure to market variance risk, which Ang et al. find can account for differences in expected returns, are driven by cross-sectional differences in exposure to correlation risk or by differences in exposure to individual variance risk. This analysis exploits the fact that exposures (betas) can be estimated quite accurately with daily data.

We decompose market variance changes as follows. Identically, changes in index variance must be due to changes in the variances of the returns of the constituents and to changes in the correlations between their returns. We decompose daily changes in index variance into what we call a pure variance effect and a pure correlation effect. The variance effect is the change in index variance one would expect if individual variances changed as they do in the data, while correlations remained constant. Similarly, the pure correlation effect is the change in index variance stemming purely from the observed change in correlations, but with constant individual variances.

The two components $\Delta Var(t)$ and $\Delta Corr(t)$ are constructed as follows:

$$\begin{aligned}\Delta Var(t) \equiv & \sum\nolimits_i \left[w_i^2 \sigma_{MF,i}^2(t) + \sum\nolimits_{j \neq i} w_i w_j \sigma_{MF,i}(t) \sigma_{MF,j}(t) IC(t-1)\right] \\ & - \sum\nolimits_i \left[w_i^2 \sigma_{MF,i}^2(t-1) + \sum\nolimits_{j \neq i} w_i w_j \sigma_{MF,i}(t-1) \sigma_{MF,j}(t-1) IC(t-1)\right] \end{aligned} \tag{24}$$

and

$$\begin{aligned}\Delta Corr(t) \equiv & \sum\nolimits_i \left[w_i^2 \sigma_{MF,i}^2(t-1) + \sum\nolimits_{j \neq i} w_i w_j \sigma_{MF,i}(t-1) \sigma_{MF,j}(t-1) IC(t)\right] \\ & - \sum\nolimits_i \left[w_i^2 \sigma_{MF,i}^2(t-1) + \sum\nolimits_{j \neq i} w_i w_j \sigma_{MF,i}(t-1) \sigma_{MF,j}(t-1) IC(t-1)\right]. \end{aligned} \tag{25}$$

In words, the pure variance effect $\Delta Var(t)$ is the difference between hypothetical index variance today and yesterday, where we hold the implied correlations fixed at yesterday's level ($IC(t-1)$), we fix the index weights at today's level ($w_i = w_i(t)$) and we only update the individual variances from $\sigma_{MF,i}^2(t-1)$ to $\sigma_{MF,i}^2(t)$. Instead, for the pure correlation effect $\Delta Corr(t)$ we fix individual variances at the level of the previous trading day ($\sigma_{MF,i}^2(t-1)$), but update the implied correlation from $IC(t-1)$ to $IC(t)$.

We then run for each stock in CRSP the following two regressions. Like Ang et al. (2006), we regress daily stock returns on the market return and on changes in implied market variance $\Delta\sigma^2_{MF,I}(t) \equiv \sigma^2_{MF,I}(t) - \sigma^2_{MF,I}(t-1)$ to obtain the CAPM beta and the 'market variance beta', respectively. In a second regression, we regress stock returns on the market return, on $\Delta Var(t)$ and on $\Delta Corr(t)$. We follow Ang et al. (2006) and estimate the betas over a rolling 1-month window to account for time-variation in factor loadings. Because of the decomposition that $\Delta\sigma^2_{MF,I}(t) = \Delta Var(t) + \Delta Corr(t)$, stocks may load on market variance risk because they are exposed to correlation risk, or because they load on individual variance risk, or both. To analyze this empirically, we calculate the cross-sectional correlation between the market variance betas and the correlation betas, as well as the cross-sectional correlation between the market variance betas and the individual variance betas.

The time-series average of the cross-sectional correlation between the market variance loadings of a stock and its correlation beta is 97.23% (Table 6). In contrast, the average cross-sectional correlation between the market variance beta and the individual variance beta is only 15%. This strongly suggests that if market variance risk is priced (as Ang et al. (2006) find), it must be due to priced correlation risk, exactly as we found for options. Put differently, the fact that cross-sectional differences in exposure to market variance risk account for differences in expected returns cannot be due to priced individual variance risk, since the market variance beta of a stock is essentially unrelated to its individual variance loading. On the contrary, cross-sectional differences in market variance betas correspond almost perfectly to cross-sectional differences in correlation sensitivities.

We analyze the robustness of this result by controlling for other risk factors that have previously been shown to account for cross-sectional variation in expected stock returns. We consider size and value (Fama and French (1993)) and momentum (Jegadeesh and Titman (1993)). The inclusion of these factors in the estimation of the loadings on market variance risk, correlation risk and individual variance risk has little impact on the beta patterns. The cross-sectional correlation between the market variance beta and the correlation beta is still very high (93.21%) and the cross-sectional correlation between betas for market variance and individual variance remains small (12.84%).

As discussed above, we do not expect statistically precise estimates of the risk premia from the stock cross-section. Still, for completeness we now turn to cross-sectional regressions to estimate the risk premia on correlation risk and individual variance risk. Because our sample is half the length of the sample of Ang et al. (2006), we also estimate the market variance risk premium for our sample.

We use Fama-MacBeth (1973) and each month perform a cross-sectional regression of individual stock returns on market variance betas, where we correct for the market index return (and size, book-to-market and momentum) by correcting the individual stock return for the factor beta times the factor return. As before, the betas are estimated using daily observations and monthly windows. Similarly, we regress individual stock returns on correlation betas and individual variance betas, again correcting for standard risk factors. The bottom panel of Table 6 reports the time-series average of the factor risk premia estimates and their t-statistics. As expected, we find that the market variance risk premium is not statistically significant. The same holds for the correlation and individual variance risk premia.

In sum, the beta patterns we documented above indicate that if market variance risk is priced in the cross-section of stock returns (as suggested by Ang et al. (2006)), this is attributable to priced correlation risk, exactly as we found for options. These results are also in line with those of Krishnan, Petkova and Ritchken (2006), who take an alternative approach to study the price of correlation risk in the cross-section of stocks. They do not use high-frequency option data to construct correlation and variance risk factors, but instead use monthly changes in historical equity correlations and variances as risk factors. Clearly, these factors are more prone to estimation error than option-implied estimates of correlations and variances, but they can be constructed for much longer sample periods. Krishnan, Petkova and Ritchken (2006) study 40 years of US equity data and find evidence for a risk premium on correlation risk, in line with our results.

5.2 An ICAPM Interpretation of Priced Correlation Risk

Table 4 already showed that option-implied correlations predict future realized correlations and hence future investment opportunities. We now study directly whether correlations can be used to predict the market return and market variance. According to Merton's ICAPM, market-wide correlations may then be a priced state variable. In particular, Chen (2003) extends the ICAPM to allow for time-varying market volatility and shows that an asset's expected return depends on the covariation between the asset's return and innovations in a state variable that predicts future market volatility. Market-wide correlations are a potential candidate for being such a state variable and Chen's extension of the ICAPM may then provide an explanation for our finding of a negative correlation risk premium. The intuition is that higher future market volatility represents a deterioration in investment opportunities causing risk-averse investors to increase precautionary savings and to reduce current consumption. Assets that pay off well when future market variance is expected to increase provide a hedge against the deterioration in investment opportunities and therefore earn lower expected returns.

We also analyze the predictive power of average individual variances for future investment opportunities, since our results so far indicate that individual variance risk is not priced. Analyzing the predictive power of average individual variances is also interesting in light of recent work on the relation between idiosyncratic risk and market returns (Goyal and Santa-Clara (2003)).

We start with predictive regressions of the monthly CRSP (value-weighted or equally-weighted) index return over the 1960-2006 sample on lagged values of (value-weighted or equally-weighted) average correlation and average individual variance, calculated from all CRSP stocks that have non-missing returns for all trading days in a given month.[29] We find no evidence of any predictability, whether using average correlations, average individual variance or market variance.[30] Numerous recent papers starting with the seminal work of French, Schwert and Stambaugh (1987) have examined the empirical time-series relation between the conditional mean and variance of the market return, testing the fundamental tradeoff between risk and return (as in for instance Merton's ICAPM). It is fair to say that this large literature has obtained mixed evidence. Recent work, however, has obtained clear evidence of a significant risk-return tradeoff by developing sophisticated techniques for estimating either conditional variances (e.g. MIDAS of Ghysels, Santa-Clara and Valkanov (2004 and 2005)) or conditional expected returns (e.g. the implied cost of capital of Pastor, Sinha and Swaminathan (2006) or the predictive systems of Pastor and Stambaugh (2006)).[31] Extending these methods (e.g. the MIDAS approach) to construct conditional correlation estimates is a promising area of future research and may lead to different findings than in our simple exploratory analysis.

Table 7 reports predictive regressions for market variance, using both average correlations and average individual variances as predictors. We consider both the variance of the value-weighted CRSP index and of the equally-weighted index, where average correlations and variances are also value-weighted or equally-weighted across stocks. For the regressions we take square roots of market variance and of the predictive variables, since the time-series of variance exhibit strong positive

[29] Although daily CRSP data are available from 1926, before 1960 very few stocks have non-missing returns for all trading days in each month.

[30] This is robust to using standard deviations instead of variances, to the exclusion of the 1987 crash from the sample, to nonlinear transformations of correlation and to considering sample splits.

[31] Brandt and Wang (2006) avoid the need to estimate expected returns and conditional variance by using the cross-section of stocks returns.

skewness. In both the equally-weighted and value-weighted cases we find that the correlation level has predictive power for future market variance. According to Chen's extension of the ICAPM, correlation risk should therefore be priced. In particular, Chen shows that when market volatility is stochastic, covariation with factors that provide information on future market variances commands a risk premium. Intriguingly, average individual variances also have predictive power for future market variances and contribute in fact more. Table 7 reports the implied λ for each factor (defined in Chen's equation (16)), which summarizes the factor's predictive power and persistence and to which the factor risk premium is proportional. We find that λ is actually larger for the individual variance factor than for the correlation factor. While the ICAPM therefore provides an explanation for why correlation risk is priced, it simultaneously predicts that individual variance risk should also carry a risk premium and in fact a larger one. In contrast, we have presented extensive and robust evidence of a large correlation risk premium, but we find no evidence of priced individual variance risk.

These findings leave us with the task of constructing equilibrium models where correlation risk is priced, but not individual variance risk. This is a challenging task. As discussed above, the ICAPM would generate risk premia for covariation with both state variables, as they both predict future market variance. While the ICAPM is based on a number of restrictive assumptions (e.g. a rational long-term representative investor), it is clear that developing a rational equilibrium model with fundamentally different predictions is non-trivial.

6 Conclusion

We show empirically that correlation risk is priced in the sense that assets that pay off well when market-wide correlations are higher than expected earn negative excess returns. This result is consistent with increases in market-wide 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 variance risk premium, we find no significant 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 20% per month and has a very large Sharpe ratio. This strategy has more attractive risk-return properties (including higher moments) than other option-based strategies. Third, the return on this correlation trading strategy explains 83% of the cross-sectional variation in index and individual option returns that is not accounted for by market risk. Fourth, we prove that a divergence in expected correlations between the risk-neutral and physical distributions indicates that correlation risk is priced. Empirically, we find a substantial gap between option-implied correlations (46.7% on average) and realized correlations (28.7% on average). Finally, we provide evidence suggesting that correlation risk may also be priced in the cross-section of stock returns.

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 market-wide correlation shocks and the ensuing loss in diversification benefits.

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.

Appendix A

This appendix contains the proof of Proposition 2.

Proof of Proposition 2: As individual variance risk is not priced, we have $E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right] = E_t^P\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right], \forall i$. To compare $IC(t)$ with $RC(t)$ it suffices to compare the part of the numerators of the definition:

$$
\begin{aligned}
IC(t) - RC(t) &= \Xi(t)\left(E_t^Q\left[\int_t^{t+\tau}\phi_I^2(s)\,ds\right] - E_t^P\left[\int_t^{t+\tau}\phi_I^2(s)\,ds\right]\right) \\
&= \Xi(t)\left[\int_t^{t+\tau}\sum_{i=1}^N\sum_{j\neq i} w_i w_j \left\{E_t^Q\left[\phi_i(s)\,\phi_j(s)\,\rho(s)\right] - E_t^P\left[\phi_i(s)\,\phi_j(s)\,\rho(s)\right]\right\} ds\right] \\
&= \Xi(t)\left[\int_t^{t+\tau}\Psi(s)\left\{E_t^Q\left[\rho(s)\right] - E_t^P\left[\rho(s)\right]\right\} ds\right] \\
\textit{where}\ \ \Xi(t) &\equiv \frac{1}{\sum_{i=1}^N\sum_{j\neq i} w_i w_j \sqrt{E_t^P\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}\sqrt{E_t^P\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}} > 0 \\
\Psi(s) &\equiv \sum_{i=1}^N\sum_{j\neq i} w_i w_j E_t^P\left[\phi_i(s)\,\phi_j(s)\right] > 0,\ \forall s.
\end{aligned}
$$

We use the Fubini Theorem to change the order of integration for the second equality and exploit that $\phi_i(s)\,\phi_j(s)$ is independent of $\rho(s)$ for the third equality.

If correlation risk is not priced ($\kappa = 0$), then $E_t^Q[\rho(s)] = E_t^P[\rho(s)]$, so that $IC - RC = 0$. To prove that $IC - RC$ is increasing in -κ, consider IC_{κ_1} and IC_{κ_2}, which are defined as in (19) for different equivalent martingale measures Q_1 and Q_2, each constructed from κ_1 and κ_2, respectively, where $\kappa_1 < \kappa_2$. It suffices then to prove that $IC_{\kappa_1} > IC_{\kappa_2}$. Using the arguments above, we find:

$$
IC_{\kappa_1}(t) - IC_{\kappa_2}(t) = \Xi(t)\left[\int_t^{t+\tau}\Psi(s)\left\{E_t^{Q_1}[\rho(s)] - E_t^{Q_2}[\rho(s)]\right\} ds\right].
$$

The (normalized) kernel is $\xi_i(t) = e^{-\int_0^t(\eta_i dB_M + \kappa_{i,o} dB_{\rho,o}) - \frac{1}{2}\int_0^t(\eta_i^2 + \kappa_{i,o}^2)ds}$, $i = 1, 2$, where dB_M and $dB_{\rho,o}$ are the independent Wiener processes driving stock market risk and orthogonal correlation risk, respectively. Define $f(t) \equiv \frac{\xi_2(t)}{\xi_1(t)}$. Under Q_1, f is a martingale and it is straightforward to show that $\frac{df}{f} = (\eta_1 - \eta_2)\,dB_M^{Q_1} + (\kappa_{1,o} - \kappa_{2,o})\,dB_{\rho,o}^{Q_1}$. Then, using $\kappa_i = \psi\eta_i + \sqrt{1-\psi^2}\kappa_{i,o}$,

$$
\begin{aligned}
E_t^{Q_1}\left[\rho\left(s\right)\right]-E_t^{Q_2}\left[\rho\left(s\right)\right] &= E_t^{Q_1}\left[\rho\left(s\right)\right]-f\left(t\right)^{-1}E_t^{Q_1}\left[f\left(s\right)\rho\left(s\right)\right] \\
&= E_t^{Q_1}\left[\rho\left(s\right)\right]-f\left(t\right)^{-1}E_t^{Q_1}\left[f\left(s\right)\right]E_t^{Q_1}\left[\rho\left(s\right)\right]-f\left(t\right)^{-1}cov_t^{Q_1}\left(f\left(s\right),\rho\left(s\right)\right) \\
&= -f\left(t\right)^{-1}\int_t^s E_t^{Q_1}\left[df\left(u\right)d\rho\left(u\right)\right]du \\
&= -f\left(t\right)^{-1}\left(\kappa_1-\kappa_2\right)\sigma_\rho\int_t^s E_t^{Q_1}\left[f\left(u\right)\sqrt{\rho(u)\left(1-\rho(u)\right)}\right]du \\
&= \left(\kappa_2-\kappa_1\right)\sigma_\rho\int_t^s E_t^{Q_2}\left[\sqrt{\rho(u)\left(1-\rho(u)\right)}\right]du>0
\end{aligned}
$$

Finally, since $\Xi\left(t\right)>0$ and $\Psi\left(s\right)>0$, $\forall s$, it follows that $IC_{\kappa_1}-IC_{\kappa_2}>0$ if and only if $\kappa_1<\kappa_2$, so that $IC-RC$ is strictly increasing in -κ. ■

Appendix B

This appendix elaborates some aspects of the analysis in section 1: lemmas 1 through 3 and proposition 1.[32]

Lemma 1: *The correlation state variable $\rho\left(t\right)$ following (14) with initial condition $\rho\left(0\right)\in\left(0,1\right)$ remains within interval $\left(0,1\right)$ with probability 1 if $\lambda\overline{\rho}>\sigma_\rho^2/2$ and $\lambda\left(1-\overline{\rho}\right)>\sigma_\rho^2/2$, and remains within interval $\left[0,1\right)$ with probability 1 if $\lambda\left(1-\overline{\rho}\right)>\sigma_\rho^2/2$.*

Proof of Lemma 1: Starting with the lower boundary, the solution to the Feller diffusion

$$
d\rho=\lambda\left(\overline{\rho}-\rho\right)dt+\sigma_\rho\sqrt{\rho}dB \tag{26}
$$

never reaches 0 from a strictly positive initial condition $\rho\left(0\right)$ with probability 1 under the Feller condition $\lambda\overline{\rho}>\sigma_\rho^2/2$. As the drift term $\lambda\left(\overline{\rho}-\rho\right)dt$ of (26) is strictly positive when $\rho\to 0$, it dominates the diffusion term $\sigma_\rho\sqrt{\rho}dB$ (in the sense that $\left|\lambda\left(\overline{\rho}-\rho\right)dt\right|\geq\left|\sigma_\rho\sqrt{\rho}dB\right|$ with probability 1) when $\rho\to 0$ and the Feller condition is satisfied. As $\left|\sigma_\rho\sqrt{\rho}\sqrt{1-\rho}dB\right|\leq\left|\sigma_\rho\sqrt{\rho}dB\right|$ for $\rho\to 0$, the diffusion term of (14) is also dominated by the drift, and hence the solution to (14) never reaches 0 from a strictly positive initial condition $\rho\left(0\right)$ with probability one if $\lambda\overline{\rho}>\sigma_\rho^2/2$.

To show that the solution to (14) never reaches the upper boundary we show that 0 is never

[32] We thank Zhipeng Zhang for suggestions on the proofs in Appendix B.

reached by $1-\rho$, which follows

$$d\left(1-\rho\right)=\lambda\left(\left(1-\overline{\rho}\right)-\left(1-\rho\right)\right)dt-\sigma_{\rho}\sqrt{\rho}\sqrt{1-\rho}dB \tag{27}$$

For this SDE we can also look at the Feller diffusion for $1-\rho$:

$$d\left(1-\rho\right)=\lambda\left(\left(1-\overline{\rho}\right)-\left(1-\rho\right)\right)dt-\sigma_{\rho}\sqrt{1-\rho}dB \tag{28}$$

This diffusion never reaches 0 from a strictly positive initial condition $1-\rho\left(0\right)$ with probability 1 under the Feller condition $\lambda\left(1-\overline{\rho}\right)>\sigma_{\rho}^{2}/2$. Following the logic for the lower boundary case, we show that the solution to (27) is bounded away from 0 with probability one if $\lambda\left(1-\overline{\rho}\right)>\sigma_{\rho}^{2}/2$. ■

Proposition 1: *The stochastic correlation matrix $\Omega\left(t\right)=\left[\rho_{ij}\left(t\right)\right]$, where $i,j=1,...,N$, with elements $\rho_{ij}\left(t\right)=1$ for $i=j$, and $\rho_{ij}\left(t\right)=\rho\left(t\right)$ for $i\neq j$, is positive definite for all t if $\rho\left(t\right)\in\left[0,1\right)$. The instantaneous variance-covariance matrix $\Sigma\left(t\right)$ of the $N-$dimensional stock price process with positive definite correlation matrix $\Omega\left(t\right)$ and instantaneous variance for any stock $E\left[d\left\langle S_{i}\right\rangle\right]=\phi_{i}^{2}\left(t\right)dt$ with $\phi_{i}\left(t\right)\neq 0$, $\forall i$, is positive definite for all t.*

Proof of Proposition 1: Take the correlation matrix $\Omega\left(t\right)=\left[\rho_{ij}\left(t\right)\right]$ at time t and omit the time argument. We can rewrite $\Omega=\left(1-\rho\right)\mathbf{I}+\rho\mathbf{11'}$, where $\mathbf{I}$ is the identity matrix and $\mathbf{1}$ is a column vector of ones. Let $X=\left(x_{i}\right)$ be a column vector with at least one element different from 0. Then the quadratic form created from this vector and the correlation matrix is always positive and the matrix Ω is positive definite by definition:

$$X'\Omega X=\left(1-\rho\right)X'\mathbf{I}X+\rho X'\mathbf{11'}X=\left(1-\rho\right)\sum_{i=1}^{n}x_{i}^{2}+\rho\left(\sum_{i=1}^{n}x_{i}\right)^{2}>0 \tag{29}$$

The variance matrix $\Phi=\left[\phi_{i}^{2}\right]$, where $i=1,...,N$, has positive diagonal elements $\phi_{i}^{2}\left(t\right)$ and zero off-diagonal elements. Next define A as the symmetric positive definite matrix with zero off-diagonal elements and ϕ_{i} on the diagonal, so that $\Phi=AA^{T}$. Then the instantaneous variance-covariance matrix $\Sigma=A\Omega A^{T}$. Pre- and post-multiplying this matrix by any $N\times 1$ vector $w\neq\mathbf{0}$ gives $w^{T}A\Omega A^{T}w=w_{1}^{T}\Omega w_{1}>0$, where $w_{1}^{T}=w^{T}A=\left(A^{T}w\right)^{T}$, as Ω is positive definite. ■

Lemma 2: *For $\kappa<0$, the correlation state variable $\rho\left(t\right)$ following (15) with initial condition $\rho\left(0\right)\in\left(0,1\right)$ remains within interval $\left(0,1\right)$ with probability 1 if $\lambda\overline{\rho}>\sigma_{\rho}^{2}/2$ and $\lambda\left(1-\overline{\rho}+\frac{\kappa\sigma_{\rho}\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}\right)>$*

$\sigma_\rho^2/2$ *for some arbitrarily small positive* ε*, and remains within interval* $[0,1)$ *with probability* 1 *if* $\lambda\left(1-\overline{\rho}+\frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}\right)>\sigma_\rho^2/2$ *for some arbitrarily small positive* ε*.*

Proof of Lemma 2: The instantaneous correlation process under Q can be written as:

$$d\rho=\lambda\left(\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\rho\left(1-\rho\right)}}{\lambda}-\rho\right)dt+\sigma_\rho\sqrt{\rho\left(1-\rho\right)}dB_\rho^Q \tag{30}$$

We see that $\overline{\rho}^*=\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda}$ approaches the initial long-run mean $\overline{\rho}$ from above when ρ approaches 0 or 1, and the initial process (14) approaches the boundary 0 faster (informally) than the process under Q. Then 0 is a natural boundary for the process under Q if $\lambda\overline{\rho}>\sigma_\rho^2/2$.

For the upper boundary we modify Lemma 1's proof slightly. Approximate process (30) as

$$d\rho = \begin{cases}\lambda\left(\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda}-\rho\right)dt+\sigma_\rho\sqrt{\rho\left(1-\rho\right)}dB_\rho^Q \text{ if } 1-\rho>\varepsilon\\ \lambda\left(\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}-\rho\right)dt+\sigma_\rho\sqrt{\rho\left(1-\rho\right)}dB_\rho^Q \text{ if } 1-\rho\leq\varepsilon\end{cases} \tag{31}$$

where ε is a very small positive number

The continuous process ρ, while approaching its boundary 1, passes through the point where the distance to the boundary is ε. At this point the nonlinearly increasing (for this range of ρ) part of the drift due to the change of measure $\frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda}$ becomes a constant $\frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}$. This new process remains continuous.

For any ρ closer to the upper boundary 1 than ε the following is true:

$$\frac{\kappa\sigma_\rho\sqrt{\varepsilon\left(1-\varepsilon\right)}}{\lambda}\leq\frac{\kappa\sigma_\rho\sqrt{\rho\left(1-\rho\right)}}{\lambda}\text{ for }\frac{\kappa\sigma_\rho}{\lambda}<0$$

Hence for $\rho\in[1-\varepsilon,1)$ the 'long-run mean' of the approximating correlation process will be higher than any value of the long-run mean of the original process:

$$\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\varepsilon\left(1-\varepsilon\right)}}{\lambda}\geq\overline{\rho}-\frac{\kappa\sigma_\rho\sqrt{\rho\left(1-\rho\right)}}{\lambda}$$

It follows that the approximating process is attracted more strongly (formally measured by the speed of convergence) to the upper boundary 1 than the original process.

The sufficient conditions under which the value 1 is a natural boundary[33] for the approximating

[33] Here we deviate slightly from the definition of natural (Feller) boundary given in Karlin and Taylor (1981). We do not show that the process cannot be started from the boundary, instead, we assume that.

process (31) are also sufficient for 1 being an unattainable boundary for the original process (14).

The sufficient conditions for the approximating process follow from the derivation of those for the original correlation process under P. We just need to replace $\overline{\rho}$ by a new long-run mean $\overline{\rho} - \frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}$ for a small positive number ε. Then the boundary 1 is unattainable if, for some arbitrarily small positive ε, $\lambda\left(1-\overline{\rho}+\frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}\right) > \sigma_\rho^2/2$. ■

Lemma 3: *Given a set of fixed index weights $\{w_i\}$ and strictly positive volatility processes $\phi_i(t)$, $\forall i$, the relationship between $RNEAC(t)$ defined in (18) and $IC(t)$ defined in (19) is given by:*

$$IC(t) = RNEAC(t)\frac{\sum_{i=1}^{N}\sum_{j\neq i} w_i w_j E_t^Q\left[\int_t^{t+\tau}\phi_i(s)\,\phi_j(s)\,ds\right]}{\sum_{i=1}^{N}\sum_{j\neq i} w_i w_j\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}} \leq RNEAC(t)$$

Proof of Lemma 3: From the definitions,

$$IC(t) = RNEAC(t)\frac{\sum_{i=1}^{N}\sum_{j\neq i} w_i w_j E_t^Q\left[\int_t^{t+\tau}\phi_i(s)\,\phi_j(s)\,ds\right]}{\sum_{i=1}^{N}\sum_{j\neq i} w_i w_j\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}}. \tag{32}$$

Therefore, it suffices to determine the following inequality:

$$E_t^Q\left[\int_t^{t+\tau}\phi_i(s)\,\phi_j(s)\,ds\right] \overset{?}{\leq} \sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_j^2(s)\,ds\right]}$$

We apply the Cauchy-Bunyakovsky-Schwarz inequality twice to obtain:

$$E_t^Q\left[\int_t^{t+\tau}\phi_i(s)\,\phi_j(s)\,ds\right] \leq E_t^Q\left[\sqrt{\int_t^{t+\tau}\phi_i^2(s)\,ds}\sqrt{\int_t^{t+\tau}\phi_j^2(s)\,ds}\right] \tag{33}$$

and

$$E_t^Q\left[\sqrt{\int_t^{t+\tau}\phi_i^2(s)\,ds}\sqrt{\int_t^{t+\tau}\phi_j^2(s)\,ds}\right] \leq \sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_i^2(s)\,ds\right]}\sqrt{E_t^Q\left[\int_t^{t+\tau}\phi_j^2(s)\,ds\right]} \tag{34}$$

from which the result follows directly. ■

References

[1] Ait-Sahalia, Y. and R. Kimmel, 2005, "Maximum Likelihood Estimation of Stochastic Volatility Models," working paper, Princeton University.

[2] Andersen, T., L. Benzoni and J. Lund, 2002, "An Empirical Investigation of Continuous-Time Models for Equity Returns," *Journal of Finance* 57, 1239-1284.

[3] Ang, A., R. Hodrick, Y. Xing and X. Zhang, 2006, "The Cross-Section of Volatility and Expected Returns," *Journal of Finance* 61, 259-299.

[4] Bakshi, G. and N. Kapadia, 2003a, "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies* 16, 527-566.

[5] Bakshi, G. and N. Kapadia, 2003b, "Volatility Risk Premiums Embedded in Individual Equity Options: Some New Insights," *Journal of Derivatives*, 45-54.

[6] Bakshi, G., N. Kapadia and D. Madan, 2003, "Stock Return Characteristics, Skew Laws, and Differential Pricing of Individual Equity Options," *Review of Financial Studies* 16, 101-143.

[7] Bates, D., 2003, "Empirical Option Pricing: A Retrospection," *Journal of Econometrics* 116, 387-404.

[8] Black, F., 1976, "Studies of Stock Price Volatility Changes," Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section, 177-181.

[9] Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 637-654.

[10] Bollen, N. and R. Whaley, 2004, "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?," *Journal of Finance* 59, 711-754.

[11] Bollerslev, T., R. Engle and J. Woolridge, 1988, "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy* 96, 116-131.

[12] Bollerslev, T., M. Gibson and H. Zhou, 2004, "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities," working paper, Duke University.

[13] Bondarenko, O., 2003a, "Statistical Arbitrage and Securities Prices," *Review of Financial Studies* 16, 875-919.

[14] Bondarenko, O., 2003b, "Why are Puts So Expensive?," working paper, University of Illinois, Chicago.

[15] Bondarenko, O., 2004, "Market Price of Variance Risk and Performance of Hedge Funds," working paper, University of Illinois, Chicago.

[16] Brandt, M. and F. Diebold, 2003, "A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations," forthcoming in *Journal of Business*.

[17] Brandt, M. and L. Wang, 2006, "Measuring the Time-Varying Risk-Return Relation from the Cross-Section of Equity Returns," working paper, Duke University.

[18] Branger, N. and C. Schlag, 2004, "Why is the Index Smile So Steep?," *Review of Finance* 8, 109-127.

[19] Breeden, D. and R. Litzenberger, 1978, "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business* 51, 621-652.

[20] Britten-Jones, M. and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance* 55, 839-866.

[21] Broadie, M., M. Chernov and M. Johannes, 2005, "Model Specification and Risk Premiums: Evidence from Futures Options," forthcoming in *Journal of Finance*.

[22] Buraschi, A. and J. Jackwerth, 2001, "The Price of a Smile: Hedging and Spanning in Option Markets," *Review of Financial Studies* 14, 495-527.

[23] Buraschi, A., P. Porchia and F. Trojani, 2006, "Correlation Hedging," working paper, Imperial College.

[24] Campa, J. and P. Chang, 1998, "The forecasting ability of correlations implied in foreign exchange options," *Journal of International Money and Finance* 17, 855-880.

[25] Carr, P. and D. Madan, 1998, "Towards a Theory of Volatility Trading," in Jarrow, R. (ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*, RISK Publications, London.

[26] Carr, P. and L. Wu, 2004, "Variance Risk Premia," working paper, Bloomberg L.P.

[27] Chen, J., 2003, "Intertemporal CAPM and the Cross-Section of Stock Returns," working paper, University of Southern California.

[28] Cochrane, J., 2001, "*Asset Pricing*," Princeton University Press.

[29] Cochrane, J., F. Longstaff and P. Santa-Clara, 2004, "Two Trees," working paper, University of Chicago.

[30] Collin-Dufresne, P. and R. Goldstein, 2001, "Stochastic Correlation and the Relative Pricing of Caps and Swaptions in a Generalized-Affine Framework," working paper, Carnegie Mellon University.

[31] Coval, J. and T. Shumway, 2001, "Expected Option Returns," *Journal of Finance* 56, 983-1009.

[32] De Jong, F., J. Driessen and A. Pelsser, 2004, "On the Information in the Interest Rate Term Structure and Option Prices," *Review of Derivatives Research* 7, 99-127.

[33] Dennis, P. and S. Mayhew, 2002, "Risk-Neutral Skewness: Evidence from Stock Options," *Journal of Financial and Quantitative Analysis* 37, 471-493.

[34] Dennis, P., S. Mayhew and C. Stivers, 2005, "Stock Returns, Implied Volatility Innovations, and the Asymmetric Volatility Phenomenon," forthcoming in the *Journal of Financial and Quantitative Analysis*.

[35] Dumas, B., 1995, "The Meaning of the Implicit Volatility Function in Case of Stochastic Volatility," unpublished note, HEC.

[36] Engle, R. and K. Sheppard, 2005, "Evaluating the Specification of Covariance Models for Large Portfolios," working paper, NYU.

[37] Eraker, B., M. Johannes and N. Polson, 2003, "The Impact of Jumps in Volatility and Returns," *Journal of Finance* 58, 1269-1300.

[38] Eraker, B., 2004, "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," *Journal of Finance* 59, 1367-1403.

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

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

[41] French, K., W. Schwert and R. Stambaugh, 1987, "Expected Stock Returns and Variance," *Journal of Financial Economics* 19, 3-29.

[42] Garleanu, N., L. Pedersen and A. Poteshman, 2005, "Demand-Based Option Pricing," working paper, Wharton School of Business.

[43] Ghysels, E., P. Santa-Clara and R. Valkanov, 2004, "The MIDAS Touch: Mixed Data Sampling Regression Models," working paper, UCLA.

[44] Ghysels, E., P. Santa-Clara and R. Valkanov, 2005, "There is a Risk-Return Tradeoff After All," *Journal of Financial Economics* 76, 509-548.

[45] Goetzmann, W., J. Ingersoll Jr., M. Spiegel and I. Welch, 2002, "Sharpening Sharpe Ratios," working paper, Yale School of Management.

[46] Goyal, A. and P. Santa-Clara, 2003, "Idiosyncratic Risk Matters!," *Journal of Finance* 58, 975-1007.

[47] Han, B., 2006, "Stochastic Volatilities and Correlations of Bond Yields," forthcoming in *Journal of Finance*.

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

[49] Jiang, G. and Y. Tian, 2005, "The Model-Free Implied Volatility and Its Information Content," *Review of Financial Studies* 18, 1305-1342.

[50] Jones, C., 2006, "A Nonlinear Factor Analysis of S&P 500 Index Option Returns," *Journal of Finance* 61, 2325-2363.

[51] Jorion, P., 2000, "Risk Management Lessons from Long-Term Capital Management," *European Financial Management* 6, 277-300.

[52] Karlin, S. and H. Taylor, 1981, "*A Second Course in Stochastic Processes*," Academic Press, London UK.

[53] Krishnan, C., R Petkova and P. Ritchken, 2006, "The Price of Correlation Risk," working paper, Case Western Reserve University.

[54] Longstaff, F., P. Santa-Clara and E. Schwartz, 2001, "The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence," *Journal of Finance* 56, 2067-2109.

[55] Lopez, J. and C. Walter, 2000, "Is Implied Correlation Worth Calculating? Evidence from Foreign Exchange Options and Historical Data," *Journal of Derivatives* 7, 65-82.

[56] Merton, R., 1973a, "Theory of Rational Option Pricing," *Bell Journal of Economics* 4, 141-183.

[57] Merton, R., 1973b, "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41, 867-887.

[58] Moskowitz, T., 2003, "An Analysis of Covariance Risk and Pricing Anomalies," *Review of Financial Studies* 16, 417-457.

[59] Newey, W. and K. West, 1987, "A Simple, Positive-Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55, 703-708.

[60] Pan, J., 2002, "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics* 63, 3-50.

[61] Pastor, L., M. Sinha and B. Swaminathan, 2006, "Estimating the Intertemporal Risk-Return Tradeoff Using the Implied Cost of Capital," working paper, Chicago GSB.

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

[63] Pastor, L. and R. Stambaugh, 2006, "Predictive systems: Living with imperfect predictors," working paper, Chicago GSB.

[64] Roll, R., 1988, "The International Crash of October 1987," *Financial Analysts Journal* 44, 19-35.

[65] Rubinstein, M., 2000, "Comments on the 1987 Stock Market Crash: 11 Years Later," in *Risks in Investment Accumulation Products of Financial Institutions*, published by The Society of Actuaries.

[66] Santa-Clara, P. and A. Saretto, 2005, "Option Strategies: Good Deals and Margin Calls," working paper, UCLA.

[67] Shanken, J., 1992, "On the Estimation of Beta Pricing Models," *Review of Financial Studies* 5, 1-34.

[68] Skintzi, V. and A. Refenes, 2003, "Implied Correlation Index: A New Measure of Diversification," working paper, Athens University.

Table 1: Variance Risk Premia in Index and Individual Options

Table 1 reports the time-series averages of realized and model-free implied variances, for S&P100 puts and for individual puts on the stocks in the S&P100 index over the 1/1996 - 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 30-day window. The model-free implied variance MFIV is calculated from a cross-section (across strikes) of 30-day put options, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in section 1.2. The data on option prices are from OptionMetrics and variances are expressed in annual terms. The p-value is for the null hypothesis that implied and realized variance are on average equal.

	Index Options	Individual Options
Mean Realized Variance	0.2102^2	0.3991^2
Mean Model-Free Implied Variance	0.2479^2	0.3856^2
Difference (1)-(2)	-0.0377	0.0135
p value for $H_0 : RV = MFIV$	0.0000	0.046

Table 2: Correlation Trading Strategy: Summary Statistics

The table reports summary statistics for the correlation trading strategy derived in 3.3.1. All results are based on daily data of overlapping monthly returns on put options and equity, for our 1/1996 - 12/2003 sample period. Panel A reports summary statistics of the correlation trading strategy return, constructed using the equity index, index puts and individual puts with strike prices corresponding to a given Black-Scholes delta. For individual options, the index-weighted cross-sectional average across firms is used. The last row reports the results for an equally-weighted portfolio of the correlation trading strategies across Black-Scholes deltas. Panel B reports summary statistics for index put option returns, for different Black-Scholes deltas and for an equally-weighted portfolio across deltas. The column 'Optimal' contains the in-sample optimal Sharpe ratio, when investing statically (with fixed portfolio weights) in index puts and individual puts (with a given delta, or equally-weighted) and the equity index. All statistics are monthly, except the Sharpe ratios, which are annualized. Panel C reports the monthly certainty equivalent (as a percentage of initial wealth) that a CRRA investor demands as compensation for not being able to invest in a particular derivative strategy and instead only investing in the equity index and the riskfree asset. Three derivative strategies are considered: the correlation trading strategy and index put strategy (both equally-weighted across deltas), and a strategy that mimicks the return of an index variance swap contract. As discussed in section 3.3.2, these certainty equivalents are estimated using the optimal CRRA portfolio weights over the sample period, for different levels of risk aversion γ. The value of $\gamma = 1.75$ generates a 100% equity index weight when derivatives are not available.

Panel A: Return Statistics for Correlation Strategy

Option	Excess Return	Std. Dev.	Skewness	Kurtosis	Ann. Sharpe
$\Delta = -0.80$	0.1026	0.2419	0.0283	5.0081	1.4698
$\Delta = -0.75$	0.1137	0.2690	-0.4999	3.8270	1.4641
$\Delta = -0.70$	0.1227	0.3216	-0.7014	3.7837	1.3211
$\Delta = -0.65$	0.1336	0.3815	-0.8428	3.9385	1.2131
$\Delta = -0.60$	0.1475	0.4423	-0.9745	4.2709	1.1551
$\Delta = -0.55$	0.1627	0.5055	-1.1182	4.7365	1.1152
$\Delta = -0.50$	0.1798	0.5702	-1.2781	5.3640	1.0921
$\Delta = -0.45$	0.2006	0.6364	-1.4676	6.2044	1.0920
$\Delta = -0.40$	0.2291	0.7078	-1.6455	7.2296	1.1213
$\Delta = -0.35$	0.2646	0.7834	-1.8085	8.4832	1.1700
$\Delta = -0.30$	0.3056	0.8557	-2.0170	10.3280	1.2370
$\Delta = -0.25$	0.3575	0.9257	-2.3073	13.2221	1.3377
$\Delta = -0.20$	0.4268	1.0032	-2.6792	17.9442	1.4738
EW portfolio	0.2113	0.5098	-2.0292	9.4604	1.4357

Panel B: Return Statistics for Short Index Put and In-Sample Optimal Portfolio

	Short Index Put					Optimal
Option	Excess Return	Std. Dev.	Skewness	Kurtosis	Sharpe	Sharpe
$\Delta = -0.80$	0.1368	0.7767	-0.7900	2.9724	0.6102	1.0850
$\Delta = -0.75$	0.1497	0.8800	-0.9826	3.2820	0.5892	1.0482
$\Delta = -0.70$	0.1619	0.9711	-1.1620	3.6810	0.5774	0.9810
$\Delta = -0.65$	0.1773	1.0519	-1.3315	4.1411	0.5840	0.9383
$\Delta = -0.60$	0.1964	1.1274	-1.5013	4.6757	0.6035	0.9256
$\Delta = -0.55$	0.2182	1.2004	-1.6818	5.3260	0.6298	0.9152
$\Delta = -0.50$	0.2433	1.2714	-1.8820	6.1492	0.6628	0.9112
$\Delta = -0.45$	0.2726	1.3396	-2.1135	7.2282	0.7049	0.9272
$\Delta = -0.40$	0.3107	1.4056	-2.3860	8.6671	0.7657	0.9710
$\Delta = -0.35$	0.3580	1.4671	-2.7151	10.6586	0.8453	1.0351
$\Delta = -0.30$	0.4130	1.5179	-3.1394	13.6335	0.9426	1.1179
$\Delta = -0.25$	0.4822	1.5500	-3.7428	18.5664	1.0777	1.2390
$\Delta = -0.20$	0.5767	1.5511	-4.6883	27.7130	1.2880	1.4237
EW portfolio	0.2844	1.1824	-2.4236	9.3551	0.8332	1.1855

Panel C: Certainty Equivalent for a CRRA investor

γ	1	2	5	10	20	1.75
Index Variance Swap	0.91	0.52	0.22	0.11	0.06	0.58
Index Put	2.10	1.28	0.56	0.29	0.15	1.43
Correlation Strategy	4.66	2.89	1.26	0.65	0.33	3.22

Table 3: The Cross-Section of Index and Individual Option CAPM-Excess Returns

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 return on index and individual put options across 13 Black-Scholes deltas 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 return on the correlation trading strategy (Table 2A) and a portfolio of individual puts (mimicking individual variance risk). The option returns 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.

Correlation Risk Premium	0.1495
(t-stat)	(4.31)
Individual Variance Risk Premium	0.0137
(t-stat)	(0.62)
Cross-sectional R^2	82.6%

Table 4: Implied and Realized Correlations: Summary Statistics

The table reports summary statistics (time-series mean and standard deviation) for the implied correlation (IC) and realized correlation (RC). IC is calculated from daily observations on model-free implied variances for the S&P100 index and for all index components, using (23). RC(t) is a cross-sectional weighted average (using the appropriate weights from the S&P100 index) of all historical pairwise correlations at time t, each calculated over a 30-day window of daily stock returns. For each pair of stocks, both stocks need to have at least 15 non-zero returns over the 30-day window. The middle panel reports the time-series correlation between IC and RC, both in levels and changes. The bottom panel gives the results of predictive regressions of RC on lagged IC and on lagged RC, both univariately and bivariately. The t-statistics are based on Newey-West (1987) autocorrelation consistent standard errors with 22 lags.

Series	Mean	Std. Dev.
Implied Correlation $IC(t)$	0.4668	0.1394
Average Realized Corr. $RC(t)$	0.2868	0.1247

Time-Series Correlation between $IC(t)$ and $RC(t)$	
Levels	0.6878
Quarterly Changes	0.5650

Predictive Regressions of $RC(t)$	
Univariate: Slope for $IC(t)$	0.5122
(t-stat)	(8.34)
R^2	0.3495
Univariate: Slope for $RC(t-1)$	0.6295
(t-stat)	(8.90)
R^2	0.3988
Bivariate: Slope for $IC(t)$	0.2575
(t-stat)	(4.48)
Bivariate: Slope for $RC(t-1)$	0.4255
(t-stat)	(5.93)
R^2	0.4453

Table 5: Parameters Estimates and Statistical Tests

Table 5 reports GMM parameter estimates and standard errors (Panel A) and statistical tests (Panel B) of the model with priced correlation risk. The estimation is based on 7 moment restrictions as explained in section 4.5.1 and standard errors are obtained using the delta method. The tests are conducted on the discrepancy between the model prediction for the difference between expected returns on index and individual options, and the empirically observed difference, for 13 different put options (indexed by (Black-Scholes) delta ranging from -0.8 to -0.2 in increments of 0.05). The table reports t-statistics for each delta. The test incorporates that the risk premia on both risk factors are estimated and also corrects for serial dependence. The restricted model imposes that the price of orthogonal correlation risk is zero. The final column reports t-statistics for the difference between the expected individual put return as predicted by the model and as empirically observed.

Panel A: Parameter Estimates

	Estimate	Std. Error
$\overline{\rho}$	0.287	0.026
λ	6.02	3.258
σ_ρ	1.823	0.245
κ	-6.597	0.544
ψ	-0.253	0.075
μ	0.1047	0.0866
ϕ	0.386	0.012

Panel B: Statistical Tests: Individual t-tests

	Index minus Individual Put Return		Indiv. Put Return
Option	Full Model	Restricted Model ($\kappa_o = 0$)	Full Model
$\Delta = -0.80$	-1.3826	-3.5438	1.2515
$\Delta = -0.75$	-0.6848	-3.4278	1.0773
$\Delta = -0.70$	-0.0387	-3.2514	0.9224
$\Delta = -0.65$	0.4591	-3.1445	0.7242
$\Delta = -0.60$	0.8042	-3.1394	0.5285
$\Delta = -0.55$	1.0671	-3.1433	0.3029
$\Delta = -0.50$	1.2565	-3.1539	0.0626
$\Delta = -0.45$	1.3520	-3.2287	-0.1459
$\Delta = -0.40$	1.2940	-3.4099	-0.3351
$\Delta = -0.35$	1.1404	-3.6587	-0.5230
$\Delta = -0.30$	0.9830	-3.9370	-0.7007
$\Delta = -0.25$	0.7153	-4.3206	-0.8875
$\Delta = -0.20$	0.2865	-4.8714	-1.1561
Average abs. value	0.8819	3.5562	0.6629

Table 6: Correlation Risk and the Cross-Section of Stock Returns

Table 6 reports cross-sectional correlations (across stocks in the CRSP universe) of market variance betas with correlation betas and individual variance betas. All betas are estimated using daily return data and a monthly rolling window. In a first set of regressions, the market variance betas are estimated by regressing a stock's return on a market variance factor, controlling for other risk factors (either the equity market return or the Fama-French factors and momentum). In a second set of regressions, the correlation betas and individual variance betas are estimated by regressing stock returns on correlation and individual variance factors, again controlling for other risk factors. Each month, a stock is included in the cross-section if it has at least 17 daily return observations. The market variance factor is the daily change in model-free implied index variance, and the correlation factor and individual variance factor are constructed from index and individual option data (equations (24) and (25)). Since all betas are re-estimated each month, the table reports the average cross-sectional correlation of betas. The table also reports estimates for risk premia on the different risk factors, estimated using Fama-MacBeth from the average coefficients of cross-sectional regressions of monthly stock returns on lagged betas. The risk premia are reported at a monthly frequency.

Beta Patterns: Cross-sectional Correlation of Market Var Beta with		
	Corr. Beta	Indiv. Var. Beta
Controlling for Market Risk	0.9723	0.1502
Controlling for FF and Mom.	0.9321	0.1284

Cross-sectional Risk Premia Estimates			
	Mkt Var.	Correlation	Indiv. Var.
Controlling for Market Risk	−0.0042%	−0.0081%	0.0143%
(t-stat)	(0.80)	(0.52)	(0.22)
Controlling for FF and Mom.	−0.0022%	−0.0064%	0.0123%
(t-stat)	(0.54)	(0.43)	(0.20)

Table 7: Predicting Market Variance

Table 7 reports results of a monthly regression of the level of the market index variance on the one-month lagged levels of equity correlations and individual variances over the 1960-2006 sample. The monthly market variance is estimated using daily data of equity market index returns (either the CRSP value-weighted or equally-weighted return). Each month, using all stocks in the CRSP universe with no missing returns in that month, we also construct the average pairwise correlation across all stocks and the average individual variance across all stocks using daily data (where the average is value-weighted or equally-weighted). We also report results of AR(1) regressions for the correlation and individual variance levels. Together, these regression results determine the λ-parameter in Chen's (2003) ICAPM. In this model, λ determines the factor risk premium (together with the risk aversion parameter).

Predicting Market Variance	Value-Weighted		Equally-Weighted	
	Correlation	Indiv. Var.	Correlation	Indiv. Var.
Slope	0.0099	0.0719	0.0073	0.1338
(t-stat)	(6.27)	(5.67)	(4.43)	(6.86)
R^2	0.146		0.167	
λ	0.0015	0.0025	0.0010	0.0036

AR(1) for Predict. Variables	Value-Weighted		Equally-Weighted	
	Correlation	Indiv. Var.	Correlation	Indiv. Var.
Constant	0.0875	0.0103	0.0922	0.0063
(t-stat)	(13.72)	(6.92)	(14.38)	(4.23)
Slope	0.4310	0.7398	0.4098	0.7973
(t-stat)	(10.42)	(18.51)	(10.37)	(15.76)
R^2	0.186	0.549	0.168	0.637

Figure 1A: Implied versus Realized Volatility for Index Options

Figure 1A 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 - 12/2003 OptionMetrics sample. The model-free implied index variance is calculated from a cross-section (across strikes) of 30-day put options on the S&P100, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in section 1.2. Realized variance is calculated from daily index returns over a 30-day window. Variances are expressed in annual terms.

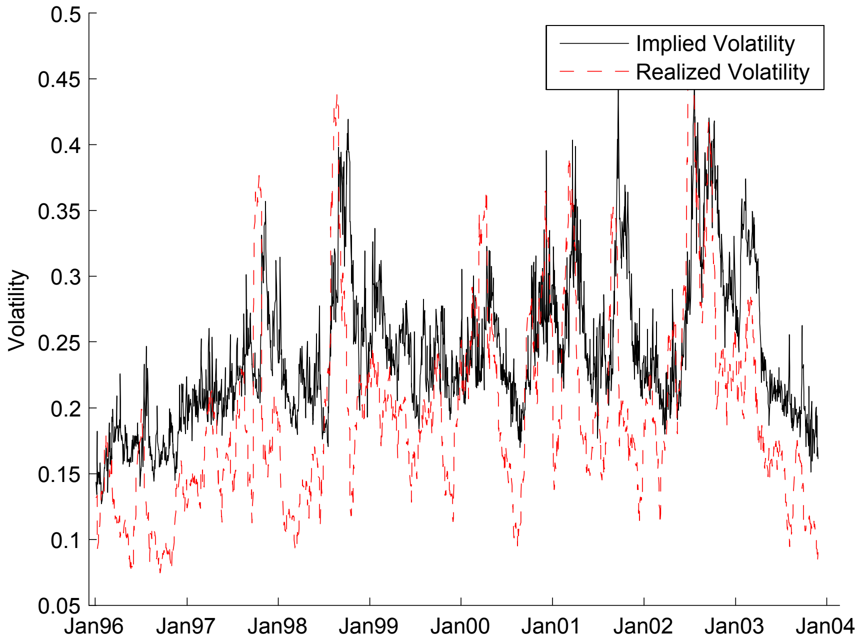

Figure 1B: Cross-Sectional Average of Implied versus Realized Volatility for Individual Options

Figure 1B 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 - 12/2003 OptionMetrics sample. For each stock in the S&P100 index, the model-free implied variance is calculated from 30-day put options, using the methodology of Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) described in section 1.2. Realized variance is calculated from daily CRSP stock returns over a 30-day 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.

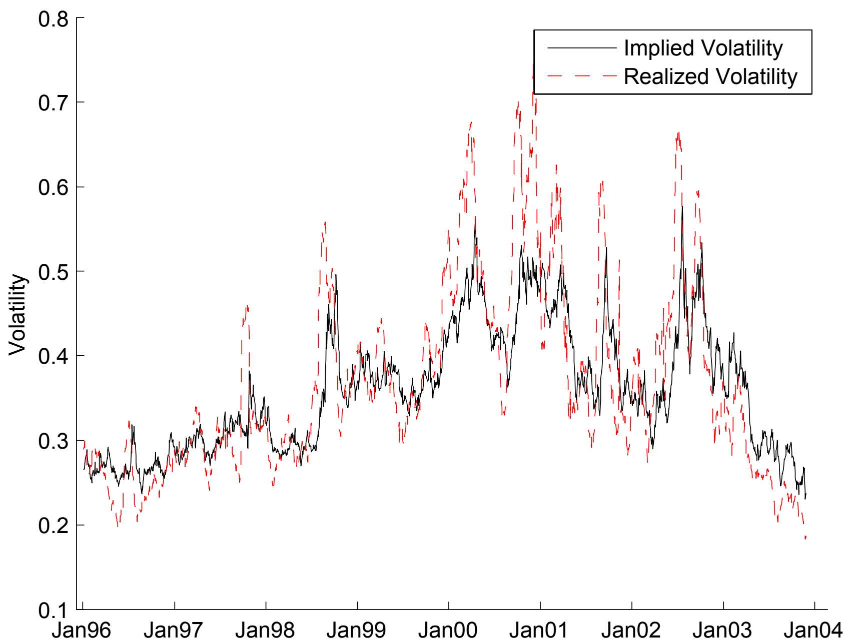

Figure 2: Implied versus Realized Correlation

Figure 2 shows 1-week moving averages of the implied correlation and the realized correlation. The implied correlation is calculated from daily observations on model-free implied variances for the S&P100 index and for all index components, using (23). Each model-free implied variance is calculated from 30-day put options. The realized correlation at time t is a cross-sectional weighted average (using the appropriate weights from the S&P100 index) of all historical pairwise correlations at time t, each calculated over a 30-day window of daily stock returns. For each pair of stocks, both stocks need to have at least 15 non-zero returns over the 30-day window.

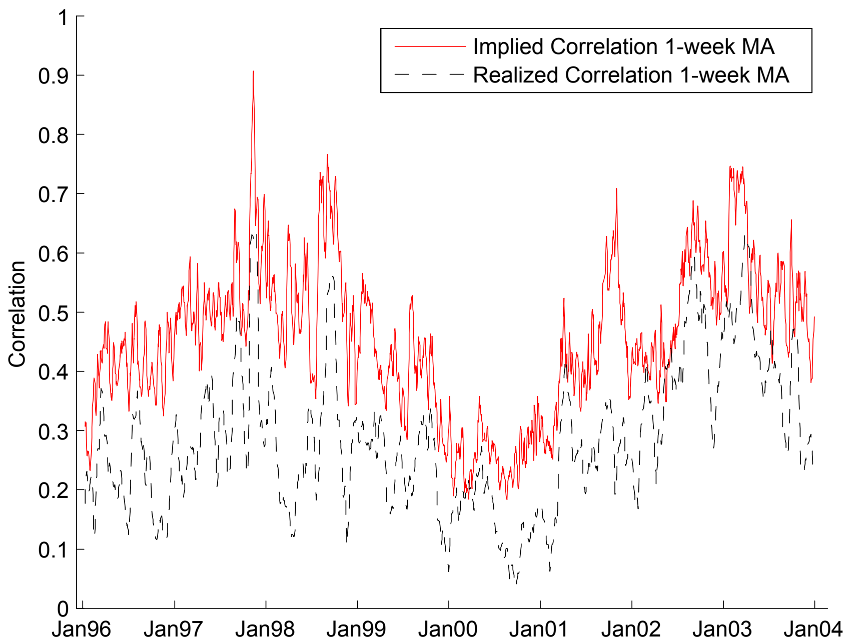

Figure 3: Expected Option Returns in the Correlation Risk Model

Figure 3 compares average option returns on 30-day index puts and individual puts with theoretical unconditional expected option returns predicted by the correlation risk model. The average individual put returns are averaged over time and across index components (using the index weights). Average option returns are calculated as sample means of the time-series of monthly holding-period returns. Model-implied expected individual-option returns are generated by the Black-Scholes model. The model predictions for expected index-option returns are endogenously obtained in the model given the GMM estimates of the parameters in Table 5A (except in one case where κ_o is set to zero), as described in section 4.5.