

Option-Implied Correlations and the Price of Correlation Risk*

Joost Driessen[†]

Pascal Maenhout[‡]

Grigory Vilkov[§]

September 2005

Abstract

Motivated by ample evidence that stock-return correlations are stochastic, we study the economic idea that the risk of correlation changes (affecting diversification benefits and investment opportunities) may be priced. We propose a powerful test by comparing option-implied correlations between stock returns (obtained by combining S&P100 option prices with prices of individual options on all index components) with realized correlations. Our parsimonious model shows that the substantial gap between average implied (46.7%) and realized (28.7%) correlations is direct evidence of a large negative correlation risk premium, since individual variance risk is not priced in our 1996-2003 sample. Empirical implementation of our model also indicates that the entire index variance risk premium can be attributed to the high price of correlation risk. Finally, the model offers a quantitatively accurate risk-based explanation of the empirically observed discrepancy between expected returns on index and on individual options. Index options are expensive, in contrast to individual options, because they hedge correlation risk. Standard models for individual equity returns with priced jump or volatility risk, but without priced correlation risk, fail to explain this discrepancy.

Keywords: Correlation risk; Dispersion trading; Index volatility; Stochastic volatility; Expected option returns.

*We would like to thank Yacine Ait-Sahalia, Oleg Bondarenko, Michael Brandt, Menachem Brenner, John Campbell, Mike Chernov, Darrell Duffie, Rob Engle, Jan Ericsson, Gerard Genotte, Jens Jackwerth, Frank de Jong, Hayne Leland, Anthony Neuberger, Josh Rosenberg, Otto van Hemert, Robert Whitelaw and especially Bernard Dumas for comments and stimulating discussions. We received helpful comments from seminar participants at the Cornell Johnson School, HEC Lausanne, INSEAD, LBS-LSE-Oxford Asset Pricing Workshop, NY Fed, NYU Stern, Tilburg, University of Amsterdam, University of Rotterdam, Warwick Business School, Yale SOM, CEPR Summer Symposium Gerzensee and EFA 2005. 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

Motivated by ample evidence that stock-return correlations are stochastic, we study the economic idea that the risk of correlation changes (affecting diversification benefits and investment opportunities) may be priced. We propose a powerful test by comparing option-implied correlations between stock returns (obtained by combining S&P100 option prices with prices of individual options on all index components) with realized correlations. Our parsimonious model shows that the substantial gap between average implied (46.7%) and realized (28.7%) correlations is direct evidence of a large negative correlation risk premium, since individual variance risk is not priced in our 1996-2003 sample. Empirical implementation of our model also indicates that the entire index variance risk premium can be attributed to the high price of correlation risk. Finally, the model offers a quantitatively accurate risk-based explanation of the empirically observed discrepancy between expected returns on index and on individual options. Index options are expensive, in contrast to individual options, because they hedge correlation risk. Standard models for individual equity returns with priced jump or volatility risk, but without priced correlation risk, fail to explain this discrepancy.

Correlations play a central role in financial markets. There is considerable evidence that correlations between asset returns change over time (see Bollerslev, Engle and Woolridge (1988), and Moskowitz (2003), among others¹) and that stock-return correlations may increase when returns are low.² Despite this evidence and despite the importance to investors (who risk vanishing diversification benefits precisely in down markets), the pricing of correlation risk in equity-derivatives markets has not received any attention in the literature.

This is surprising because index and individual stock options contain unique information on the pricing of correlation risk and constitute an ideal testing ground for the model of priced correlation risk we develop. The main idea is particularly intuitive and builds on two observations. First, option prices encode the risk-neutral expected variance of the return on the underlying asset over the option's life. Second, the index variance is a weighted average of individual variances and covariance terms. Therefore, we can combine index option prices with prices of individual options on all index components to infer the option-implied average correlation between the index components over the option's life. This forward-looking correlation measure reflects the risk-neutral distribution and a comparison with realized correlations, which mirror the actual measure, directly reveals the correlation risk premium. Our first contribution is to show that this risk premium is very large and almost entirely a compensation for orthogonal correlation risk.

As a second contribution, we shed new light on recent findings in the option-pricing literature and offer a novel view on the source of the large volatility risk premium that recent work on index options has disclosed.³ We demonstrate that a large fraction of changes in market volatility stem from correlation changes. Remarkably, a model where the entire market variance risk premium is due to a correlation risk premium explains the data extremely well. In other words, while changes in individual volatility empirically contribute to changes in market volatility, these changes are not priced and cannot explain the market volatility risk premium. Priced correlation risk constitutes the missing link.

Our third contribution is a risk-based explanation for why index puts seem expensive relative

¹See Brandt and Diebold (2003) and Engle and Sheppard (2005) for recent innovations in the estimation of dynamic correlations.

²Financial crises, with multiple assets and markets suffering simultaneous declines in asset values, 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 model multivariate distribution tails in international equity markets and to shed light on the hypothesis that correlations increase in volatile times.

³The 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 Liu (2004), Coval and Shumway (2001), Eraker, Johannes and Polson (2003), Eraker (2004), Jones (2005) and Pan (2002). Bates (2003) surveys earlier work.

to the Black-Scholes model, while individual options do not. Although index options have received most attention in the recent literature, there is growing evidence that individual options exhibit rather different stylized facts, especially as regards (Black-Scholes) implied volatility functions.⁴ We add formal evidence that individual options do not embed a variance or jump risk premium, nor earn returns in excess of the one-factor Black-Scholes model. Without priced correlation risk it is challenging to reconcile these findings with the opposite evidence for index options, because the index process is just the weighted average of the individual processes; risk premia present in index options should also appear in individual options. Indeed, we show that a standard market model with priced variance risk cannot explain the discrepancy between index and individual options and is strongly rejected by the data. The same goes for the one-factor stochastic volatility model and jump-diffusion models without correlation risk. A risk-based explanation for the contrast between index and individual options requires that aggregated individual processes be exposed to a risk factor that is lacking from the individual processes. Priced correlation risk makes this possible. Intuitively, index options are expensive and earn low returns, unlike individual options, because they offer a valuable hedge against correlation increases and insure against the risk of a loss in diversification benefits.^{5,6}

We propose a simple model to understand the role of stochastic correlations and the price of correlation risk. Each individual stock price follows a geometric Brownian motion (with potentially stochastic volatility), but the correlation between stock returns is stochastic and mean-reverting. This model produces endogenous stochastic index-return volatility, even with constant individual volatilities. Moreover, we allow for a negative risk premium on correlation risk, which generates higher expected correlation paths under the risk-neutral measure than under the actual measure. The negative correlation risk premium has at least three important implications that we test. First, the option-implied correlation exceeds the average realized correlation between equity returns if individual variance risk is not priced. Second, the risk-neutral expected integrated variance implied by index options is higher than the realized equity index variance. Third and most importantly,

⁴See 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).

⁵In a very insightful note revisiting the 1987 crash, Rubinstein (2000) 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."

⁶Garleanu, 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.

expected index-put returns are much more negative than expected individual-put returns.

To test these implications, 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. Our sample runs from January 1996 until the end of December 2003. We find that the option-implied variance systematically exceeds realized variance for index options, in line with previous findings, but not for individual options. Because implied variances are obtained with the model-free approach of Britten-Jones and Neuberger (2000), Carr and Madan (1998), Dumas (1995) and Jiang and Tian (2005), this is formal evidence that individual option prices do not embed a variance risk premium. These findings are further supported by the size of expected excess put returns. Index options have much more negative excess returns than individual options. The difference in excess returns varies from -7% per month for in-the-money (ITM) puts to -39% for out-of-the-money (OTM) puts.

We then analyze the three main implications of the correlation risk premium. First, we calculate each day an option-implied correlation from 30-day index and individual options, using various assumptions on the correlation structure. 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 the first 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 further show that it is crucial to incorporate changes in (implied) correlations to understand changes in option-implied index variance. If one only incorporates changes in individual variances, more than half of the variation in the index variance remains unexplained.

To assess the second and third implication, we calibrate the correlation risk model discussed above. While the model generates no difference between individual implied and realized variances by construction, it also closely matches the large observed difference between implied and realized index variance. Interestingly, when the correlation risk premium is calibrated to the average observed difference between implied and realized correlations, the model gives a very good fit of expected index-option returns, across the entire moneyness spectrum. For example, for an ATM index put

option, the expected return equals -24% per month. The model-implied value is -27%, of which -16% is due to the correlation risk premium and -11% to the regular equity premium effect. In contrast, when we set the orthogonal correlation risk to zero, the model generates expected index option returns that are much less negative than the observed returns. Finally, we show that commonly used models with priced jumps or stochastic volatility for individual stock prices, including an intuitive market model, fail to generate the observed discrepancy between expected individual- and index-option returns. The main intuition for the failure of these models is that the risk of jumps and stochastic volatility will affect both individual and index options, so that premia on these risks should be reflected in both index and individual options. Instead, our correlation model is able to generate a large difference between premia for index versus individual options, since the correlation risk premium only affects index options.

Several articles have investigated the correlation structure of interest rates of different maturities using cap and swaption prices. Longstaff, Santa-Clara, and Schwartz (2003) and De Jong, Driessen, and Pelsser (2004) both 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 and driven by a separate factor. 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, and apply this to the case of 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 rest of this paper is organized as follows. Section 1 presents the theoretical model of priced correlation risk. This model is implemented in Section 2, where we analyze option-implied correlations empirically and present evidence on implied versus realized variances, for both index

and individual options. Section 3 presents the empirical results of a simple decomposition of changes in index variance, highlighting the relevance of changes in the correlation structure. New stylized facts about the differences between expected returns on index options and individual options are provided in Section 4. Section 5 shows how the model of priced correlation risk can explain these empirical regularities, while Section 6 demonstrates that standard models with stochastic volatility and jumps fail to generate sufficiently large differences between index and individual options. Finally, Section 7 concludes.

1 A Model of Priced Correlation Risk

This section describes the model of priced correlation risk. We start by specifying the stochastic process for each stock that is included in the stock market index. The correlations between different stock returns are stochastic and the risk of changes in the correlation structure is priced. We then derive the process for the return variance of the stock market index and show how our model endogenously generates stochastic index volatility. The final subsection introduces the concept of the risk-neutral expected average correlation. The risk-neutral expected average correlation can be extracted from option prices and is shown to carry information about the price of correlation risk when contrasted with actual historical correlations estimated from stock returns.

1.1 Individual Stock Price Processes

The stock market index is composed of N stocks. Under the physical probability measure P , the price of stock i , S_i , is assumed to follow a geometric Brownian motion with expected return μ_i and possibly stochastic diffusion $\phi_i(t)$:

$$dS_i = \mu_i S_i dt + \phi_i S_i dB_i \tag{1}$$

where B_i is a standard scalar Wiener process. We omit time as an argument for notational convenience throughout, except when placing particular emphasis. The special case where $\phi_i(t)$ is constant simplifies (1) to the standard Black-Scholes set-up. More generally, the instantaneous volatility $\phi_i(t)$ is taken to be a strictly positive Ito process, driven by a standard scalar Wiener process B_{ϕ_i} , which is taken to be uncorrelated with B_j for all j .⁷ For example, the process for $\phi_i(t)$

⁷This means that there is no asymmetric volatility or ‘leverage effect’ (Black (1976)) at the level of individual stock returns. However, our correlation-risk model nonetheless generates an endogenous leverage effect for index

could be specified as the special case of the well-known Heston (1993) stochastic volatility (SV) model with non-priced volatility risk. An important premise will be that $\phi_i(t)$ follows the same process under the physical probability measure P as under the risk-neutral probability measure Q . In other words, we claim that the volatility risk in individual stock returns is not priced. We will present strong evidence supporting this premise.

Under both the physical measure P and the risk-neutral measure Q , the instantaneous correlation between the Wiener processes B_i and B_j for $i \neq j$ is modeled as:

$$E_t^P [dB_i dB_j] = E_t^Q [dB_i dB_j] = \rho_{ij}(t) dt. \quad (2)$$

The stochastic nature of this instantaneous correlation $\rho_{ij}(t)$ and especially the associated risk premium constitute the key innovations of this paper. We assume that a single state variable $\rho(t)$ drives all pairwise correlations:

$$\rho_{ij}(t) = \overline{\rho}_{ij} \rho(t). \quad (3)$$

Furthermore, most of the analysis is conducted under the following simplifying homogeneity assumption:

$$\overline{\rho}_{ij} = 1, \quad \forall i \text{ and } j. \quad (4)$$

While this assumption is admittedly restrictive, it has important advantages as explained below, while allowing us to capture correlation risk in the most parsimonious way. In addition, we carry out robustness checks to demonstrate that our results are not sensitive to this homogeneity restriction.

We impose the initial condition that $\rho(0) \in (0, 1)$. Under measure P , the correlation state variable $\rho(t)$, which is also the instantaneous correlation because of (4), is assumed to follow a mean-reverting process with long-run mean $\bar{\rho}$, mean-reversion parameter λ and diffusion parameter σ_ρ .⁸

$$d\rho = \lambda(\bar{\rho} - \rho) dt + \sigma_\rho \sqrt{\rho(1 - \rho)} dB_\rho. \quad (5)$$

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

returns. This 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.

⁸Collin-Dufresne and Goldstein (2001) discuss a similar model of correlation dynamics.

certain parameter restrictions $\rho(t)$ remains between zero and one with probability 1, as is shown by the following lemma.⁹

Lemma 1: *The correlation state variable $\rho(t)$ following (5) with initial condition $\rho(0) \in (0, 1)$ remains within interval $(0, 1)$ with probability 1 if $\lambda\bar{\rho} > \sigma_\rho^2/2$ and $\lambda(1 - \bar{\rho}) > \sigma_\rho^2/2$, and remains within interval $[0, 1)$ with probability 1 if $\lambda(1 - \bar{\rho}) > \sigma_\rho^2/2$.*

Lemma 1 is important to guarantee the positive definiteness of the resulting correlation and variance-covariance matrices, as is shown in the following proposition.

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

1.2 Priced Correlation Risk

Priced correlation risk is introduced into the model by having the instantaneous correlation follow a similar mean-reverting process under the risk-neutral measure Q as in (5), but with a correlation risk premium proportional to κ subtracted from its drift:

$$d\rho = \left[\lambda(\bar{\rho} - \rho) - \kappa\sigma_\rho\sqrt{\rho(1 - \rho)} \right] dt + \sigma_\rho\sqrt{\rho(1 - \rho)}dB_\rho^Q \quad (6)$$

A strictly negative value for κ implies a negative correlation risk premium, so that the expected path (or integral) of future correlations under Q exceeds the expected correlation path under P . Below we develop a methodology for extracting information about the price of correlation risk embedded in the risk-neutral correlation process. This is achieved by using option prices to back out the risk-neutral correlation process, while historical stock returns are used to estimate the actual correlation process.¹⁰

The correlation risk premium may have multiple sources. First, we let the Brownian motion B_ρ driving $\rho(t)$ in (5) be correlated with the market risk factor in the pricing kernel. We therefore

⁹We could have accommodated negative values for ρ by constraining it to the $(-1, 1)$ interval, but since ρ is the average correlation between the stocks in the index, this generalization is not needed empirically.

¹⁰Lemma 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 valid for our empirical parameter estimates), measures P and Q are equivalent.

let each Brownian B_i affecting individual stock price S_i have a nonzero correlation with B_ρ :

$$E^P [dB_i dB_\rho] = \psi_i \neq 0 \tag{7}$$

and are particularly interested in the correlation between dB_ρ and shocks to the market return. We will later estimate this correlation. A negative value might be expected and would be in line with empirical work that finds that correlations increase when prices decline. This negative correlation between returns shocks and changes in correlation would be an obvious reason for having a (negative) risk premium on correlation risk. In addition, the part of correlation risk that is orthogonal to stock market risk may also be priced. In that case, correlation risk constitutes a second priced risk factor in the economy. Formally, denoting the Sharpe ratio for (equity) market risk by η and the cross-sectional weighted average of ψ_i by ψ , the correlation risk premium κ can be decomposed in a systematic part and an orthogonal part as follows

$$\kappa = \kappa_1 + \kappa_2 = \psi\eta + \sqrt{1 - \psi^2} \kappa_o. \tag{8}$$

While we remain agnostic about the source of the risk premium on correlation risk for now, our empirical results will in fact show that the latter (an orthogonal second priced risk factor) constitutes by far the biggest component of the price of correlation risk.

An equilibrium model with correlation risk for the endowment processes, extending the equilibrium models of priced volatility risk of Bansal and Yaron (2004) and Tauchen (2005), would show that κ_o is zero in a representative-agent model with i.i.d. consumption growth and CRRA preferences. Heteroskedastic consumption growth combined with Epstein-Zin preferences on the other hand could support a nonzero risk premium for orthogonal correlation risk (κ_o).

1.3 Endogenous Stochastic Index Volatility

Both the existence of correlation risk and of a correlation risk premium have fundamental implications for the pricing of index options, as they are driven by the volatility of the stock market index. Even in a simplified version of our model with standard Black-Scholes dynamics for all individual stocks (the version of (1) with constant diffusion coefficients $\phi_i(t) = \phi_i$), we endogenously generate stochastic index volatility. More generally, the volatilities of individual stock returns are stochastic and as described above. Denoting the drift of $\phi_i^2(t)$ by γ_i , and the diffusion scaling parameter

(which will determine the ‘vol of vol’) by ς_i , the instantaneous variance process of the individual stock return, under both P and Q , is

$$d\phi_i^2 = \gamma_i dt + \varsigma_i \phi_i dB_{\phi_i} \quad (9)$$

It is then possible to obtain the process for the variance of the stock market index endogenously generated by our model with correlation risk as presented in Proposition 2.

Proposition 2: Define $\nu_\phi \equiv \sum_{i=1}^N \sum_{i \neq j} w_i w_j \phi_i \phi_j$, $\delta_\phi \equiv \frac{1}{2} \sum_{i=1}^N \left[w_i^2 + 2 \sum_{i \neq j} w_i w_j \phi_j \rho \right] \phi_i^{-1} \left[\gamma_i - \frac{1}{2} \varsigma_i^2 \right]$ and $\iota_i \equiv \frac{1}{2} \sum_{i=1}^N \left[w_i^2 + 2 \sum_{i \neq j} w_i w_j \phi_j \rho \right] \varsigma_i$. Given a set of fixed index weights $\{w_i\}$, individual stock price processes given by (1), (2), (3), (4) and (9), and a correlation process given by (5) under P and by (6) under Q , the instantaneous variance ϕ_I^2 of the stock market index under P solves

$$d\phi_I^2 = \nu_\phi \left[\lambda (\bar{\rho} - \rho) dt + \sigma_\rho \sqrt{\rho(1-\rho)} dB_\rho \right] + \delta_\phi dt + \sum_{i=1}^N \iota_i dB_{\phi_i}. \quad (10)$$

Under Q , the instantaneous variance ϕ_I^2 of the stock market index solves

$$d\phi_I^2 = \nu_\phi \left[\left(\lambda (\bar{\rho} - \rho) - \kappa \sigma_\rho \sqrt{\rho(1-\rho)} \right) dt + \sigma_\rho \sqrt{\rho(1-\rho)} dB_\rho^Q \right] + \delta_\phi dt + \sum_{i=1}^N \iota_i dB_{\phi_i}. \quad (11)$$

The process for the index variance implied by our model inherits the $\sqrt{\rho(1-\rho)}$ factor in its diffusion term from the bounds imposed on the correlation process and is therefore also of the Wright-Fisher type. The terms δ_ϕ and ι_i stem from the non-priced individual volatility risk and vanish in the special case of constant ϕ_i 's as in the Black-Scholes model. Interestingly, because of our premise of non-priced individual volatility risk, the sole difference between (10) and (11) is precisely the correlation risk premium. In other words, the model attributes the entire index variance risk premium to priced correlation risk. In section 3 we show empirically that correlation risk is indeed a crucial driver of stochastic index variance.

1.4 Risk-Neutral Expected Average Correlation

We are now ready to 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.

Consider the risk-neutral expected integrated variance of the return on asset a (where $a \in$

$\{I, 1, \dots, i, \dots, N\}$) over a discrete interval of length τ starting at time t :

$$\sigma_a^2(t) = E_t^Q \left[\int_t^{t+\tau} \phi_a^2(s) ds \right] \quad (12)$$

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

$$\sigma_I^2(t) = E_t^Q \left[\int_t^{t+\tau} \sum_{i=1}^N w_i^2 \phi_i^2(s) ds \right] + E_t^Q \left[\int_t^{t+\tau} \sum_{i=1}^N \sum_{i \neq j} w_i w_j \phi_i(s) \phi_j(s) \rho(s) ds \right] \quad (13)$$

Rather than attempting to extract the entire path of future correlations $\rho(s)$ from (13), 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

$$RNEAC(t) \equiv \frac{E_t^Q \left[\int_t^{t+\tau} \phi_I^2(s) ds \right] - \sum_{i=1}^N w_i^2 E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]}{\sum_{i=1}^N \sum_{i \neq j} w_i w_j E_t^Q \left[\int_t^{t+\tau} \phi_i(s) \phi_j(s) ds \right]} \quad (14)$$

RNEAC is a certainty equivalent to the postulated 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 contain information about the risk-neutral expected future variances $E_t^Q \left[\int_t^{t+\tau} \phi_I^2(s) ds \right]$ and $E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]$, respectively. 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(t)*, defined as

$$IC(t) \equiv \frac{E_t^Q \left[\int_t^{t+\tau} \phi_I^2(s) ds \right] - \sum_{i=1}^N w_i^2 E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]}{\sum_{i=1}^N \sum_{i \neq j} w_i w_j \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]} \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_j^2(s) ds \right]}} \quad (15)$$

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 $IC(t) \leq RNEAC(t)$. 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 (15) so as to capture the actual (as opposed to risk-neutral) expected average correlation. Since this will later be estimated from the time-series average of the (cross-sectionally weighted-average of) realized correlations, we call this in short the realized (or historical) correlation $RC(t)$, defined as:

$$RC(t) \equiv \frac{E_t^P \left[\int_t^{t+\tau} \phi_I^2(s) ds \right] - \sum_{i=1}^N w_i^2 E_t^P \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]}{\sum_{i=1}^N \sum_{i \neq j} w_i w_j \sqrt{E_t^P \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]} \sqrt{E_t^P \left[\int_t^{t+\tau} \phi_j^2(s) ds \right]}} \quad (16)$$

Proposition 3 states that the difference between IC and RC is crucially linked to the price of correlation risk.

Proposition 3: *Given a set of fixed index weights $\{w_i\}$, individual stock price processes given by (1), (2), (3) and (4) with non-priced individual volatility risk, and a correlation process given by (5) under P and by (6) under Q , the difference between IC and RC is zero if and only if $\kappa = 0$ and is monotonically increasing in $-\kappa$.*

Testing whether correlation risk is priced can therefore be done by testing whether $IC = RC$. Note that the proposition requires that individual stock return volatility risk is not priced. As can easily be seen from equations (15) and (16), $IC - RC$ is then proportional to $E_t^Q \left[\int_t^{t+\tau} \phi_I^2(s) ds \right] - E_t^P \left[\int_t^{t+\tau} \phi_I^2(s) ds \right]$, so that correlation risk is priced if and only if index volatility risk is priced, in line of course with the result of proposition 2.

Another crucial implication of the model with priced correlation risk concerns expected index versus individual option returns. We pursue this in detail in sections 4 and 5, after testing the other implications of the model in sections 2 and 3.

2 Empirical Features of Variances and Correlations

Before turning to the data description and the empirical results, we present the methodology we use for estimating the risk-neutral expected integrated variances. Not only are these integrated variances needed to estimate the time series for IC , they will also be used to obtain estimates of the variance risk premia in index and individual options. In particular, this will allow us to provide evidence for our premise that variance risk is not priced in individual options.

2.1 Model-Free Implied Variances and Variance Risk Premia

We use the methodology of Britten-Jones and Neuberger (2000), Carr and Madan (1998) and Dumas (1995), who build on the seminal work of Breeden and Litzenberger (1978), to estimate the risk-neutral expected integrated variance $\sigma_a^2(t)$ defined in (12) 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 in a jump diffusion setting. The measure is therefore considered ‘model-free’, and can be labeled the *model-free implied variance* (or *MFIV*).

We denote the price of a call option on asset a with strike price K and maturing at time t by $C_a(K, t)$, and the model-free implied variance of the return on asset a by $\sigma_{MF,a}^2(t)$. The main result of Britten-Jones and Neuberger is that the risk-neutral expected integrated variance ($\sigma_a^2(t)$) defined in (12) equals the model-free implied variance, which is defined as

$$\sigma_{MF,a}^2(t) \equiv 2 \int_0^\infty \frac{C_a(K, t + \tau) - \max(S(t) - K, 0)}{K^2} dK \quad (17)$$

Jiang and Tian show that, for implementation purposes, the integral over a continuum of strikes in (17) can be approximated accurately by a discrete sum over a finite number of strikes. Finally, Bollerslev, Gibson and Zhou (2004), Bondarenko (2004) and Carr and Liu (2004) establish that the difference between the model-free implied variance and the realized variance is an estimate of the variance risk premium. As a first empirical contribution we apply this methodology to show that while index options embed large variance risk premia, there is no such evidence for individual options. This provides empirical support for the premise in Section 1 that the individual volatility process is the same under the physical measure P as under the risk-neutral measure Q .

2.2 Data Description and Empirical Implementation

We use data from OptionMetrics for S&P100 index options and for individual options on all the stocks included in the S&P100 index.¹¹ OptionMetrics provides daily data on individual and index

¹¹Interestingly, 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.

options starting in January 1996. Our sample runs until the end of December 2003.

The S&P100 is a value-weighted index with rebalancing taking place each quarter. 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 value of the constituent companies at the closing price of the third Friday of the last month in the previous quarter. For instance, in 2004 the Fall quarterly rebalance took place after the close on Friday September 17 2004 (going into effect Monday September 20). The list of constituent companies may also change over time whenever a company is deleted from the index in favor of another. During our sample period, 47 such additions and deletions took place. These also occur on the rebalance dates. We reconstruct the index components and corresponding index weights for the entire sample period, using stock prices from CRSP as follows. At each rebalance date for the S&P100, we calculate the weight for stock i as the market value (from CRSP) of company i divided by the total market value of all companies that are present in the index after the rebalancing. We keep these weights fixed for the entire period 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, even though the number of shares stays fixed throughout the period. As we have 100 companies in the index, any such discrepancy due to changes in prices is extremely small and can be neglected for our purposes.

To construct the model-free implied variances, we require a panel (observations over time and across strikes) of option prices of both S&P100 index options and individual stock options for all the S&P100 constituents. 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 expected integrated variances (as well as the implied correlation IC) to be 30 calendar days. To obtain the panel of prices for options with 30-day maturity, we use the Volatility Surface File from OptionMetrics, in combination with the Standardized Options File and Security Prices File. As is common in the literature, we focus on put options. The Volatility Surface File in OptionMetrics contains daily interpolated volatility surfaces for each security, calculated with a kernel smoothing algorithm. The volatility surfaces are available for a maturity of 30 calendar days, at deltas of -0.2 to -0.8 (in increments of -0.05) and are calculated from option prices. For each option and each day, OptionMetrics only provides the delta and implied volatility. To calculate the option price from the data, we first need to obtain the option moneyness by inverting the Black-Scholes formula for delta

and then use Black-Scholes to calculate the option price.¹² Several inputs for these formulae were missing and obtained from other files. For the riskfree rate and dividend rate for the index, we utilize the daily data provided by OptionMetrics, namely the zero-coupon curve and the index dividend yield. To infer the continuously compounded dividend rate for each stock, we combine the forward price in the Standardized Options File with the spot rate used for the forward price calculations (in the Security Prices File). OptionMetrics uses a tree method with discrete dividends to price options, so before we can use the Black-Scholes approximation, we calculate the mean continuously compounded dividend rate for each stock by averaging the implied OptionMetrics dividends. Once we have the panel of option prices, we implement the model-free implied variance measure described in Section 2.1 following the procedure in Jiang and Tian (2005), suitably adjusted for put options.

To calculate the realized variance, we use daily returns from CRSP for individual stocks and from OptionMetrics for the S&P100. On each date and for each stock, we require at least 15 return observations during the 30-day window over which the realized variance is calculated. Since the window spans 30 calendar days, this means that we require 15 observations out of approximately 22 trading days. The variances are expressed in annual terms.

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

$$IC(t) = \frac{\sigma_{MF,I}^2(t) - \sum_{i=1}^N w_i^2 \sigma_{MF,i}^2(t)}{\sum_{i=1}^N \sum_{i \neq j} w_i w_j \sqrt{\sigma_{MF,i}^2(t)} \sqrt{\sigma_{MF,j}^2(t)}}. \quad (18)$$

One subtlety regarding the index weights emerges however: if the expiration of the index option occurs after the next rebalance date, it is clear that 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. If the index option expires after the rebalancing, the index option price will reflect the public information that the old company will be replaced by a new one over the life of the option. We therefore incorporate the migration in the projected weights. We then weight the old fixed weights and the new projected

¹²Jiang 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. However, unreported results with raw (non-interpolated) option prices show that our main findings are unaffected.

weights using the relative time to maturity of the index option before the rebalance date and after the rebalance date, that is we linearly interpolate the weights in order to apply (18) accurately.

To obtain historical correlations between the stocks in the S&P100 index, we use CRSP returns. Because standard historical correlations are more commonly used and intuitive and because the difference between historical correlations calculated from (16) 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. We require that both stocks have a 30-day history of returns, remain in the index until time t , and have at least 15 non-zero returns over the 30-day window (as for realized variances). For robustness checks, we also calculate historical correlations over a 180-day window, with a requirement of 90 non-zero returns (out of approximately 128 trading days). Repeating the procedure at different points in time gives a moving average estimate of the historical correlation between each pair of stocks, 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.

2.3 Implied versus Realized Variances

The recent empirical literature on equity options primarily studies index options.¹³ Individual options have attracted much less attention, in part because recent and comprehensive data was lacking until now.¹⁴ 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 and intuitive information, they do not permit a formal test of the presence of variance risk premia as we need for our analysis. This section presents such a formal test, based on the model-free methodology described in Section 2.1. Moreover, our OptionMetrics sample is 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

¹³See Bates (2003) for a survey. More recent contributions subsequent to this survey include Ait-Sahalia and Kimmel (2005), Andersen, Benzoni and Lund (2002), Bondarenko (2003 and 2004), Broadie, Chernov and Johannes (2005), Eraker, Johannes and Polson (2003), Eraker (2004) and Jones (2005).

¹⁴An early analysis (1976-1978) of the empirical performance of alternative option pricing models for individual options can be found in Rubinstein (1985). He shows that while the Black-Scholes model generates moderate mispricing, there is no single alternative model that improves substantially.

(2004) also use OptionMetrics and a related methodology, but focus on a subsample of just 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. It is immediately obvious that index variance is far from constant over the span of our sample. The well-established finding that option-implied variance is higher than realized variance also holds for our recent sample. Notice that all calculations and estimations are done for variances, not for volatilities, but for interpretation purposes we take square roots of the computed variances when plotting or presenting numbers. 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, it seems first of all that both implied and realized volatility are smoother and less volatile over time. Quite remarkably, there is little systematic difference between the two volatility proxies. On average, it is in fact the case that the square root of realized variance (39.91%) 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. 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 (again using Newey-West consistent standard errors with 22 lags) for 108 stocks out of the 135 stocks that are included in the sample for this analysis. Of the remaining 27 stocks, only 16 exhibit a statistically significantly 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 1.07% to 1.5% (depending on the treatment of dividends) between the average implied volatility and the average historical volatility in their 1991-1995 sample of 25 individual stock options. While they also stressed that the difference is smaller than for index options and therefore that the evidence for risk premia is less firm, it is clear that our results suggest that there may not be a variance risk premium in individual options at all. 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. Dennis and Mayhew (2002) argue that

variation in individual skews is more closely linked to firm-specific than to systematic factors, which is consistent with the lack of priced individual variance risk we find. 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.

The results above are very important since they provide indirect evidence that correlation risk must be priced. As was demonstrated in Section 1.3, when individual variance risk is not priced, index variance risk only carries a risk premium to the extent that correlation risk is priced. Our results strongly suggest that this is the case. The next subsection provides direct evidence of priced correlation risk (according to Proposition 3) in the form of a large difference between option-implied and realized correlations.

2.4 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. Finally, we examine the robustness of our results to the homogeneity assumption of our model (equation (4)), by considering an alternative formulation with heterogeneity in pairwise correlations.

Figure 2A plots the time-series of the 1-week moving averages of the implied correlation $IC(t)$ defined in (18) and of the historical correlation, which is a cross-sectional weighted average of all historical pairwise correlations each calculated over a 30-day window as explained above. The implied correlation is very volatile and peaks not surprisingly 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.¹⁵ The fact that the 1-week moving average of 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. At least informally, Figure 2A shows that our implied correlation measure indeed captures the dynamics of the cross-sectional average of correlations between all stocks in the

¹⁵The exchanges were closed from 11 to 16 September, leading to missing data for that period.

S&P100 index. The time-series correlation between $IC(t)$ and the historical correlation measure in Figure 2A is in fact 68.78%. Table 2 presents further information about the two series plotted in Figure 2A. Of particular interest is the extent to which the implied correlation $IC(t)$ exceeds the historical correlation. Keeping in mind that the implied correlation $IC(t)$ reflects the dynamics in (6), while the historical measure estimates the correlation under the physical measure (equation (5)), the difference between $IC(t)$ and $RC(t)$ reflects the price of correlation risk κ . Table 2A shows that the time-series average for the historical correlation is 28.68%, while the time-series average for $IC(t)$ is 46.68%. The implied correlation is therefore 62.8% higher, 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.¹⁶

Table 2 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, or to whether we calculate the historical correlation over a 30-day window (as in Figure 2A) or instead with a longer 180-day window. To further investigate the relation between the implied correlation and historical correlations, we study whether the implied correlation can predict future realized correlations. Table 1 reports the results from these predictive regressions of cross-sectionally averaged realized correlations on lagged values of the implied correlation and on lagged values of the 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, for both time windows. Table 2A shows that realized correlations also predict future realized correlations. In the bivariate predictive regressions, the implied correlation continues to be a significant predictor, for both time windows.

The results presented so far are obtained under the homogeneity assumption that $\overline{\rho_{ij}} = 1, \forall i$ and j . As a robustness check for this assumption, we now modify the procedure to estimate option-implied correlation dynamics so as to accommodate heterogeneity in $\overline{\rho_{ij}}$. We incorporate $\overline{\rho_{ij}} \neq 1$ in equation (18) and estimate $\overline{\rho_{ij}}$ by its rolling 180-day historical estimate for each pair of stocks. This gives rise to an alternative empirical strategy for identifying the option-implied correlation, which we now call $\widehat{IC}(t)$, the *correlation scaling factor*. However, this approach can be problematic since

¹⁶One may worry that estimates of realized correlations are biased downwards if there is any cross-autocorrelation. Unreported results using two-day returns rather than daily returns show that this is not the case.

it no longer imposes sufficient structure to guarantee that derived pairwise correlations $\overline{\rho_{ij}}\widehat{IC}(t)$ remain below 1. We therefore bound all derived pairwise correlations $\overline{\rho_{ij}}\widehat{IC}(t)$ below 1. As an indication of the severity of this problem, 2.17% of all correlations need to be bounded or ‘truncated’ in this way.

The resulting time-series of the 1-week moving average of the cross-sectional average of $\overline{\rho_{ij}}\widehat{IC}(t)$ is plotted in Figure 2B, along with the 1-week moving average of the cross-sectional average of the historical correlations. The time-series in Figure 2B are extremely close to the ones in Figure 2A. This is comforting, since it shows that the homogeneity assumption is not driving our results. In fact, the time-series correlation between the implied correlation of Figure 2A and the average derived correlation in Figure 2B (based on the correlation scaling factor $\widehat{IC}(t)$) is 0.9895. All other results obtained for $IC(t)$ carry over to the heterogeneous case. Figure 2B and Table 2 contain the same evidence of a large correlation risk premium. The time-series mean of the correlation scaling factor is 1.626, virtually identical to the 1.628 ratio of the mean implied correlation to the average historical correlation. The correlation scaling factor also exhibits a lot of time-series variation (time-series standard deviation of 0.4819). Finally, the time-series average and standard deviation of the cross-sectional average of $\overline{\rho_{ij}}\widehat{IC}(t)$ are extremely similar to the statistics for the implied correlation $IC(t)$ obtained under the homogeneity assumption: 0.4478 and 0.1417 versus 0.4668 and 0.1394 for $IC(t)$.

3 Decomposing Index Variance Changes: The Importance of Correlation Risk

A basic but powerful idea we would like to emphasize in this paper, is that stochastic index volatility may be driven by changes in the correlation structure between the components, i.e. by correlation risk. Of course, an alternative driver of index volatility may be stochastic volatility at the level of individual stocks, as can also be seen from equations (10) and (11). To show that correlation risk is relevant for understanding the dynamics of the volatility of the stock market index, we construct a simple decomposition of changes in index variance. Identically, changes in the variance of the index must be due to changes in the variances of the returns of the constituents and to changes in the correlations between their returns. In a bivariate regression of changes in index variance on changes in both determinants, the R^2 must be unity and the coefficients identically equal to

one. Univariately however, it is an empirical question which determinant is more important for explaining the changes in the index variance. To implement this idea empirically, we regress weekly changes (i.e. over 5 trading days) in index variance on 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 for constant individual variances.

The explanatory variables $\Delta Var(t)$ and $\Delta Corr(t)$ in our regressions are constructed as follows:

$$\begin{aligned} \Delta Var(t) &\equiv \sum_{i=j} w_i^2 \sigma_{MF,i}^2(t) + \sum_{i \neq j} w_i w_j \sigma_{MF,i}(t) \sigma_{MF,j}(t) IC(t-5) \\ &- \sum_{i=j} w_i^2 \sigma_{MF,i}^2(t-5) + \sum_{i \neq j} w_i w_j \sigma_{MF,i}(t-5) \sigma_{MF,j}(t-5) IC(t-5). \end{aligned} \quad (19)$$

and

$$\begin{aligned} \Delta Corr(t) &\equiv \sum_{i=j} w_i^2 \sigma_{MF,i}^2(t-5) + \sum_{i \neq j} w_i w_j \sigma_{MF,i}(t-5) \sigma_{MF,j}(t-5) IC(t) \\ &- \sum_{i=j} w_i^2 \sigma_{MF,i}^2(t-5) + \sum_{i \neq j} w_i w_j \sigma_{MF,i}(t-5) \sigma_{MF,j}(t-5) IC(t-5). \end{aligned} \quad (20)$$

In words, the pure variance effect $\Delta Var(t)$ is the difference between hypothetical index variance today and 5 trading days ago, where we hold the implied correlations fixed at the level of 5 days ago ($IC(t-5)$), we fix the index weights at today's level ($w_i = w_i(t)$) and we only update the individual variances from $\sigma_{MF,i}^2(t-5)$ to $\sigma_{MF,i}^2(t)$. Instead, for the pure correlation effect $\Delta Corr(t)$ we fix the individual variances at the level of last week ($\sigma_{MF,i}^2(t-5)$), impose the index weights of today (w_i), but update the implied correlation from $IC(t-5)$ to $IC(t)$.

Defining the dependent variable as $\Delta \sigma_{MF,I}^2(t) \equiv \sigma_{MF,I}^2(t) - \sigma_{MF,I}^2(t-5)$, Table 3 presents OLS estimates for the following regression:

$$\Delta \sigma_{MF,I}^2(t) = \alpha + \beta_1 \Delta Var(t) + \beta_2 \Delta Corr(t) + \varepsilon(t) \quad (21)$$

Table 3 shows that weekly changes in individual variances alone explain less than half the variation in weekly changes in index variance. The slope coefficient is very significant, but the estimate itself is 1.20, which is larger than the value of 1 that would be expected if individual variance risk were the sole driver of index variance changes. The slope coefficient for the pure

correlation effect is closer to one (1.1) and also very statistically significant. Most important for our study however is the R^2 : changes in correlations *alone* explain 67% of the changes in index variance. This is evidence that correlation risk is quantitatively and empirically relevant for understanding the dynamics of index variance. It should also be noted that this result obtains even when imposing the homogeneity restriction that $\overline{\rho_{ij}} = 1, \forall i$ and j , as we continue to do here. For completeness, we also report the results for the bivariate regression, which is approximately an identity.¹⁷

One may worry at this point that $IC(t)$ absorbs all the noise in the individual and index implied variances (from which it is constructed), thus increasing its explanatory power for changes in index variance by what is lacking from individual variances due to measurement error. To avoid this problem we now conservatively use $IC(t-1)$ as a proxy for $IC(t)$: yesterday’s estimation of implied correlation does not have today’s noise in it. The second regressor, $\widehat{\Delta Corr}(t)$, is then defined as in (20) except that $IC(t)$ in the second term is replaced by $IC(t-1)$.

The result of this conservative procedure is obviously to lower the explanatory power of the pure correlation effect in the second column of Table 3. However, the R^2 is still 33%. Most noteworthy is that even after accounting for (contemporaneous) changes in individual variances, which alone explain 45% of the variation in index variance changes, adding the ‘lagged’ correlation effect increases the R^2 by a further 22 percentage points. Therefore, even when only relying on ‘stale’ correlation changes, our understanding of the dynamics of index variance is vastly improved by incorporating correlation risk.

This shows that a pure SV model where the correlation structure in the economy is assumed fixed will have a hard time fully explaining changes in index variance empirically. The new channel of correlation risk we propose here is empirically relevant as well. Section 5 will show that the index variance risk premium is almost entirely due to a correlation risk premium.

4 Individual versus Index Option Returns: Stylized Facts

We now turn to an important and unique feature of our model, namely that correlation risk affects expected index option returns, but not individual option returns. First we present new stylized facts about the differential risk premia in index versus individual options, which complement the analysis of Section 2.3. Then we show in the next section how our correlation-risk model performs very well

¹⁷Strictly speaking, the bivariate version is in fact no longer an identity. This is because we impose fresh index weights, lest we commingle the (mechanical) effect of a change in index weights with the (economic) effect of changes in individual variances or in correlations.

in addressing the empirical findings. Alternative models on the other hand, where stochastic index volatility stems from the presence of jumps in individual stock returns or from stochastic volatility in individual stock returns cannot explain these stylized facts adequately.

4.1 Calculating Expected Option Returns

Recent work has shown that expected index option returns are very large in absolute value and extremely challenging for the one-factor Black-Scholes model (Bondarenko (2003a and 2003b), Buraschi and Jackwerth (2001), Coval and Shumway (2001) and Jones (2005) among others). Much less is known about individual options, however, and we therefore present a detailed comparison of index options and individual options on all the stocks in the index.¹⁸

As for any asset, the expected return on an option reflects the difference between its expected payoff and price. Because the physical measure P pins down the actual dynamics of the underlying and thus the expected option payoff, while the risk-neutral measure Q is encoded in the option price, expected option returns summarize the discrepancy between both measures, i.e. the presence of priced risk factors. As stressed by Coval and Shumway (2001), while the prediction of a pricing model for the difference between P and Q could be tested in many ways, the major attraction of using option returns for this purpose is its easily interpretable economic content.

The unconditional expected put return is estimated as the sample average of the time-series of holding-period returns. The holding-period return at time $t + \tau$ on a put written at t can be computed by dividing the option payoff at maturity ($t + \tau$) by the option price at t . Scaling both the option payoff at $t + \tau$ and the option price at t by the stock price at t lets us calculate the holding-period return from the scaled put price and the return on the underlying. As before, we use the Volatility Surface File from OptionMetrics for 30-day puts, complemented with the Standardized Options and Security Prices files to obtain a panel of option prices. For the 30-day return on the underlying, we use CRSP data for individual stocks and OptionMetrics data for the S&P100.

In addition to ‘raw’ expected returns, we can also obtain expected excess returns relative to a one-factor option pricing model (Black-Scholes and single-factor extensions thereof). To do this, it suffices to de-mean the returns on the underlying, so that its expected return equals the riskfree rate minus the dividend yield, before calculating each option payoff and holding-period return.

¹⁸As mentioned in Section 2.3, most recent work on individual options focuses on implied volatility smiles. Bollen and Whaley (2004) form an exception and present simulated returns of a delta-hedged trading strategy that shorts options (on the S&P 500 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.

De-meaning stock returns in this way amounts to a full shift in the probability distribution, while preserving its shape. Notice that in a single-factor complete-market model, the de-meaning exactly represents the change of measure from P to Q . If the option is indeed priced according to such a model, the mean of the holding-period returns computed from these de-meaned stock returns equals zero. Therefore, a simple, but formal test of the presence of other priced risk factors consists of testing whether these one-factor excess returns are statistically significant. We do this for both index and individual options.

When contrasting index with individual options, we need to ensure that the individual and index options are comparable in terms of ‘economic moneyness’. Since 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), we calculate the expected returns across a range of Black-Scholes delta’s rather than strike-to-spot ratios.

The calculations produce a double panel (time-series for each option and each delta) of raw returns and of excess returns. For each delta, we average the individual option returns cross-sectionally using the weights for each day calculated by the same scheme as before. To account for serial dependence in the daily-frequency one-month returns, t -statistics are based on Newey-West (1987) standard errors with 22 lags.

4.2 Empirical Results

Figure 3A presents the one-factor excess returns for index options and for individual options (the cross-sectionally weighted average). In line with existing empirical work, the excess returns for one-month index puts are extremely negative and grow to -50% for deep OTM puts (delta of -0.20). Even the 97.5% confidence bound (the upper bound of a 95% confidence interval) shows a -25% monthly excess return for the -0.2 delta index put. Importantly, significantly negative excess returns also obtain for less OTM puts. Puts with a delta of -0.4 still have a negative excess return of -25% per month. Only as we move towards ITM index puts, do the one-factor excess returns become less negative and do they cease to be statistically significant at the 97.5% confidence level. Overall, Figure 3A presents strong evidence against the one-factor complete-market option pricing model for index puts.

For individual options, the results are completely different however. While the excess return is also estimated to be negative for all puts, the 97.5% confidence bound is always positive, even

for the most OTM option (-0.2 delta). The one-factor complete-market model can therefore not be rejected for individual options. While the surprisingly negative excess returns for index puts have been documented elsewhere, it is quite remarkable that the results for individual options are so different. We directly calculate the difference between the index and average individual excess returns and plot the 95% confidence interval in Figure 3B. The difference is between -7% and -39% per month and very statistically significant for OTM and ATM options. Even for ITM puts is the 97.5% confidence bound just borderline zero.

We retain as an important stylized fact therefore that while index option returns are vastly different from what would be expected if the single-factor complete-market model generated the data, there is no evidence for mispricing of individual options by that same model. Note that we are looking at excess returns relative to a one-factor model, in order to detect the presence of additional priced risk factors. If we were to look at raw returns instead, i.e. where expected option payoffs reflect the non-demeaned physical distribution, then these put returns would be slightly more negative even than what Figure 3A reports, because of a positive equity risk premium. Explaining these raw index put returns is the challenge we present our model with in the next section.

5 Implications of the Priced Correlation Risk Model

In this section we assess to what extent the model with priced correlation risk introduced in section 1 can explain the empirical features of individual versus index options discussed above. Importantly, we calibrate the correlation risk premium to the observed difference between the implied and historical correlation, and then analyze whether this correlation risk premium can explain the expected option returns and the difference between implied and realized volatility.

For individual options, the results discussed in 2.3 show that there is no evidence for a systematic difference between implied and realized variance. In addition, there is no 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 (1) with constant diffusion coefficients. This premise directly implies that implied variance equals on average realized variance, and that expected one-factor excess returns are zero for individual options. These two implications are thus in line with empirical observations.

The remainder of this section focuses on the implications for S&P 100 index options. For simplicity, we assume that all 100 stocks are equally weighted in the index and that each stock

has a constant volatility of 38.56% (the square root of the mean individual *MFIV* in Table 1). To incorporate the equity diffusion risk premium, we set the expected stock return equal to the average equity index return over our sample period, which equals 10.47% per year. We also make the homogeneity assumption in equation (4) and use the process for the correlation specified in equations (5) and (6) under the actual and risk-neutral measure, respectively. We Euler-discretize equations (1), (5) and (6) to simulate the actual and risk-neutral paths of individual and index prices. Simulations under the risk-neutral measure, based on the sample average for the 30-day riskfree rate of 4.46% per year, are used to obtain prices of one-month index options.¹⁹

The parameters in (5) and (6) are calibrated using the information in the implied and historical correlation time series. We first discretize equation (5) into the following mean-reverting heteroskedastic process:

$$\rho_{t+\Delta t} - \rho_t = (1 - e^{-\lambda\Delta t})(\bar{\rho} - \rho_t)\Delta t + \varepsilon_{t+\Delta t} \quad (22)$$

where the error term $\varepsilon_{t+\Delta t}$ has mean zero and conditional variance $\sigma_\rho^2 \rho_t(1 - \rho_t)\Delta t$. The instantaneous correlation is not observed. To proxy for this latent variable, we can either use the implied correlation or the historical correlation. Using the implied correlation series is problematic given the presence of the risk premium in its drift. We therefore use the 30-day historical correlation as a proxy for the instantaneous correlation to estimate (22). This is in line with the standard term structure literature, where the 1-month interest rate is typically taken as a proxy for the instantaneous interest rate. We acknowledge that the historical correlation is a proxy by specifying a measurement equation for the observed 1-month correlation

$$RC_{1month} = \rho_t + \chi_t \quad (23)$$

where χ_t is a mean-zero measurement error. The state-space model in (22) and (23) can be estimated with a Quasi Maximum Likelihood Kalman Filter approach. To avoid using overlapping observations (due to the 1-month period used to estimate the correlation), we take monthly observations of the historical correlation. The procedure also renders a time series of filtered values for the latent factor ρ . This time series can be used to estimate the correlation between stock index returns and correlation changes. Assuming that the factor value is an estimate for the correlation

¹⁹We divide the 1-month maturity interval into 40 time steps and simulate 50000 outcomes for each time step.

in the middle of the 1-month period, we can calculate this correlation, and find a correlation of -22.1% , which corresponds to the cross-sectional average of ψ_i in (7).²⁰

Finally, we estimate the correlation risk premium parameter κ that determines the change from the actual to the risk-neutral measure (equation (5) versus (6)) as follows. From the simulated index option prices and the Black-Scholes individual option prices, we can back out the implied correlation that our calibrated model generates. Given that this implied correlation crucially depends on the correlation risk premium, we fit the risk premium parameter κ to the average implied correlation level of 46.68% (see Table 2). In sum, our estimates imply the following risk-neutral correlation process

$$d\rho = 1.276(0.287 - \rho) dt + 0.346\sqrt{\rho(1 - \rho)}(dB + 26.21dt) \quad (24)$$

These estimates satisfy the parameter restrictions identified in lemma 1 and 2 to guarantee that the correlation process remains within the interval $[0, 1)$.

The estimates clearly demonstrate that the correlation risk premium is almost entirely a reward for orthogonal correlation risk. Plugging $\psi = -0.221$ and $\eta = 0.286$ into the decomposition in (8), the part of the correlation risk premium due to systematic correlation risk is negligible and only represents less than 0.25% of the total κ ($-0.06321 / -26.21 = 0.00242$).

We first calculate the difference between the model-free implied index variance and the realized index variance from the simulation results. The square root of the implied variance equals 23.64% , while the square root of the realized variance is 19.09% . The difference of 4.55% is quite close to the observed difference of 3.77% . Next we analyze the expected option returns that are generated by our model. Figure 4 reports the results of this comparison and summarizes the main effects of our correlation model. The graph first of all shows the implications for expected individual option returns. Given that individual stocks each follow a Black-Scholes process, the generated expected individual option returns fully come from the negative exposure (or ‘beta’) of put options to diffusion 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 4.2, this beta-effect explains most of the negative expected returns on individual options in Figure 4, 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 (8)) is set to zero, so that κ only reflects the

²⁰If we use the first or last day of the 30-day period, this estimate goes down to about -10% .

(small) ‘ ψ -part’ of priced correlation risk. Even when this premium 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 4 shows that a zero orthogonal correlation risk premium generates expected returns that are much closer to zero than the observed index option returns. Finally, and most importantly, Figure 4 presents expected index option returns when we choose the correlation risk premium such that we match the difference between implied and historical correlations. 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 empirical expected returns. Only for deep OTM put options is there some difference between the model-implied and observed values.

To summarize, the results in this section provide strong evidence that the same single value for the correlation risk premium can explain the empirically observed difference between (i) implied and historical correlations, (ii) implied and realized index variance, and (iii) expected returns on index versus individual options across a very wide moneyness spectrum.

6 Comparison with Jump and Stochastic Volatility Models

We now analyze to what extent commonly used models with jumps or stochastic volatility can explain the difference between expected returns on index and individual options. Our strategy is to calibrate jump and volatility risk premia to expected returns on individual options, and then to analyze the model implications for expected index-option returns. We first derive the index process from individual stock price processes, and the expected returns on individual and index options.

Consider stock i ($i = 1, \dots, 100$) with process

$$dS_i = \mu(t) S_i dt + \phi(t) S_i dB_{i1} + \tilde{\phi}_i(t) S_i dB_{i2} + \theta S_i (dJ - \lambda(t) dt) + \tilde{\theta}_i S_i (d\tilde{J}_i - \tilde{\lambda}_i(t) dt). \quad (25)$$

Each stock is exposed to a common ($\phi(t)$) and an idiosyncratic ($\tilde{\phi}_i(t)$) stochastic volatility process, and to common and idiosyncratic jumps, with fixed jump sizes θ and $\tilde{\theta}_i$, respectively, and with jump intensities $\lambda(t)$ and $\tilde{\lambda}_i(t)$, respectively. The diffusion shocks dB_{i1} have fixed correlation parameter ρ across stocks, whereas the dB_{i2} shocks are idiosyncratic. Furthermore, each dB_{i1} has a correlation with the shock to the common SV process $\phi(t)$, thus incorporating a leverage effect.²¹ Obviously, only the common SV and common jump processes carry risk premia. Jump intensities are allowed

²¹This makes our model consistent with the finding of Dennis, Mayhew and Stivers (2005) that the leverage effect is mainly driven by systematic factors.

to vary over time, as long as the variation is not priced.

This general model nests several special cases:

1. Pure one-factor SV model: $\tilde{\phi}_i(t) = 0$, $\theta = 0$, $\tilde{\theta}_i = 0$.
2. Pure jump-diffusion model: $\tilde{\phi}_i(t) = 0$, $\phi(t) = \phi$.
3. Market model: $\rho = 1$ and no jumps ($\theta = 0$, $\tilde{\theta} = 0$). This model extends the standard market model and generates stochastic volatility and stochastic correlations, both driven by the same factor, and may intuitively be a potential candidate for explaining the disparity between expected index and individual option returns.

Define as before the stock index as $I = \sum_{i=1}^N S_i$. The process for I is:²²

$$\frac{dI}{I} = \mu(t) dt + \phi(t) \sum_{i=1}^{100} \frac{S_i}{I} dB_{i1} + \sum_{i=1}^{100} \tilde{\phi}_i(t) \frac{S_i}{I} dB_{i2} + \theta(dJ - \lambda(t) dt) + \sum_{i=1}^{100} \frac{S_i}{I} \tilde{\theta}(d\tilde{J}_i - \tilde{\lambda}_i(t) dt). \quad (26)$$

Assuming S_i/I is constant over the option's life, $\forall i$, this process can be approximated as follows. Define $w_i \equiv S_i/I$, $W \equiv \sum_{i=1}^{100} w_i^2 + 2\rho \sum_{i=1}^{100} \sum_{j=1, i \neq j}^{100} w_i w_j$, and a new Brownian motion $d\tilde{B}_1$ to get

$$\frac{dI}{I} \simeq \mu(t) dt + \phi(t) \sqrt{W} d\tilde{B}_1 + \theta(dJ - \lambda(t) dt) + \sum_{i=1}^{100} w_i \tilde{\phi}_i(t) dB_{i2} + \sum_{i=1}^{100} w_i \tilde{\theta}(d\tilde{J}_i - \tilde{\lambda}_i(t) dt). \quad (27)$$

The process for the index inherits the structure from the individual components, the only difference being a scaling parameter W for the volatility, reflecting the imperfect correlation between the stocks. In the market model, where all stocks are driven by a common market factor and perfectly correlated, $W = 1$ so that the index and individual processes have the same structure. This will be useful below. The last two terms are averages of idiosyncratic components and are typically small.

We use a Heston model for the common variance $V \equiv \phi^2$:

$$dV = \alpha(\nu - V) dt + \bar{\sigma} \sqrt{V} \left(\omega d\tilde{B}_1 + \sqrt{1 - \omega^2} dZ \right) \quad (28)$$

under the actual measure. Note that the correlation between dB_{i1} and dZ is caused by the systematic part $d\tilde{B}_1$ (the market return), otherwise it will disappear in the index due to averaging. This is useful for the calculations below.

Under the risk-neutral measure Q , with ξ denoting the risk premium for orthogonal variance risk and the total volatility risk premium $\xi_{tot} = \omega\eta/\sqrt{W} + \sqrt{1 - \omega^2}\xi$ (η is the stock market Sharpe

²²Formally, dividing by $I(t-)$ because of the presence of jumps.

ratio, as before),

$$dV = \alpha(\nu - V)dt + \bar{\sigma}\sqrt{V} \left[\omega \left(d\tilde{B}_1^Q - \eta\sqrt{V}/\sqrt{W} \right) + \sqrt{1 - \omega^2} \left(dZ^Q - \xi\sqrt{V} \right) \right] \quad (29)$$

$$= (\alpha + \xi_{tot}\bar{\sigma}) \left(\frac{\alpha}{\alpha + \xi_{tot}\bar{\sigma}} \nu - V \right) dt + \bar{\sigma}\sqrt{V} \left(\omega d\tilde{B}_1^Q + \sqrt{1 - \omega^2} dZ^Q \right) \quad (30)$$

The term $1/\sqrt{W}$ enters because $dB_{i1}^Q = dB_{i1} + \eta\sqrt{V}dt$, so that $d\tilde{B}_1^Q = d\tilde{B}_1 + (\eta\sqrt{V}/\sqrt{W})dt$.

We now work out the expected option returns. First, the expected equity return $\mu(t)$ equals $r + \eta V(t) + \theta(\lambda(t) - \lambda^Q(t))$. Let $P(S_i, V)$ be the price of a put on stock i . From Liu and Pan (2003) we can derive the expected option return for this model:

$$\begin{aligned} \frac{E_t(dP)}{P} - rdt &= \frac{S}{P} [\eta V(t) \Delta_S + V(t) \bar{\sigma} \xi_{tot} \Delta_V + (\lambda(t) - \lambda^Q(t)) \Delta_{J,S}] dt \\ &= \frac{S}{P} \left[\eta V(t) \left(\Delta_S + \bar{\sigma} \omega \Delta_V W^{-1/2} \right) + \xi V(t) \bar{\sigma} \sqrt{1 - \omega^2} \Delta_V + (\lambda(t) - \lambda^Q(t)) \Delta_{J,S} \right] dt \end{aligned} \quad (31)$$

Δ_S represents the partial derivative with respect to S . The volatility-sensitivity Δ_V is defined as $\frac{\partial P(S,V)}{\partial V}/S = \frac{\partial P(1,V)}{\partial V}$. $\Delta_{J,S}$ stands for the jump delta, defined as

$$\Delta_{J,S} = \frac{P((1 + \theta)S, V) - P(S, V)}{S} = P((1 + \theta), V) - P(1, V) \quad (32)$$

The first term in equation (31) represents the diffusion risk premium (incorporating the leverage effect), the second and third term the volatility and jump risk premium, respectively. Given equation (27), a similar expression for the expected index-option return can be obtained:

$$\begin{aligned} \frac{E_t(dP_I)}{P_I} - rdt &= \frac{I}{P_I} [\eta V(t) \Delta_I + WV(t) \bar{\sigma} \xi_{tot} \Delta_{\tilde{V}} + (\lambda(t) - \lambda^Q(t)) \Delta_{J,I}] dt \\ &= \frac{I}{P_I} \left[\eta V(t) \left(\Delta_I + \bar{\sigma} \omega \Delta_{\tilde{V}} \sqrt{W} \right) + W \xi V(t) \bar{\sigma} \sqrt{1 - \omega^2} \Delta_{\tilde{V}} + (\lambda(t) - \lambda^Q(t)) \Delta_{J,I} \right] dt \end{aligned} \quad (33)$$

with P_I being the price of the index put and $\tilde{V}(t) = WV(t)$.

Comparing (31) and (33) reveals that the difference in expected returns between index and individual options is due to differences in delta's, volatility-sensitivities, the scaling parameter W (whenever volatility risk matters) and the ratio of underlying price to option price. The latter converts absolute delta's and sensitivities into relative ones so as to get percentage returns. The fact that the structure of the expected options returns in (31) and (33) is so similar is not surprising, given the similarity in price processes for the underlying assets (equations (25) and (27)). After

all, the index is the weighted average of the individual components.

We now present the quantitative calibration results for the pure SV model and the market model (section 6.1) and a pure jump-diffusion model (section 6.2). The details of the implementation are relegated to Appendix C. The methodology consists of substituting some moments directly by sample analogues from the data and focusing on expected return implications. Under the null hypothesis that the model holds, the substitution is valid and other moment restrictions can then be tested (section 6.3 presents formal model tests). The main idea underlying the approach is to focus on unconditional expected returns, so that the unknowns in (31) and (33) can be estimated from long-run sample averages and cross-sectional regressions thereof. For instance, the mean of $(\lambda(t) - \lambda^Q(t))$ is estimated from long-run sample means and the jump delta's are estimated from the cross-section of (scaled) option prices.²³

6.1 Market Model and One-Factor SV Model

Figure 5 presents the results for the market model and the SV model simultaneously, because as far as expected option returns are concerned, the market model is simply the special case where $W = 1$. The graph shows that the SV model generates a very small difference between index and individual option returns. In fact, underlying these results is the finding that the volatility risk premium is smaller for index options than for individual options in this model. This can be explained by noting the presence of two opposing effects. The factor W in equation (33) is smaller than one and thus brings expected index returns closer to zero. On the other hand, index options have lower prices, and equation (33) shows that this leads to more negative expected returns.²⁴ This latter effect is quantitatively dominated by the effect of $W < 1$. Expected index option returns are slightly more negative than individual option returns, because the effect of the equity risk premium is larger for index options due to their higher beta's. Since the market model imposes $W = 1$, both the volatility risk premium and the equity risk premium are larger for index options. This results in slightly more negative index option returns than for the SV model, but much less negative than in the data.

²³We also used a Fama-MacBeth methodology to allow for time-varying volatility-sensitivities and jump delta's, stemming from time-variation in vega's, I/P_I ratios and jump intensities. This leads to an even worse performance of the SV and jump models.

²⁴Furthermore, because index options have lower prices, they also have lower implied volatilities, which increases the volatility-sensitivity (as can be seen from equation (46) in Appendix C) and therefore also contributes to making the expected index-option return more negative.

6.2 Jump-Diffusion Model

The results for the jump-diffusion model can be understood as follows. Equations (31) and (33) indicate that expected individual and index option returns can differ because of differences between $\Delta_{J,S}$ and $\Delta_{J,I}$ and between the ratios S/P and I/P_I . Quantitatively, the difference in equity-to-option price ratios matters most. Given that index option prices are smaller, the index option risk premia are ‘leveraged’ up. Across all strikes, individual options are about 56% more expensive than index options, so that we expect a similar difference, all else equal, for the risk premia. Figure 6, which shows the empirical expected option returns and the predictions of the jump-diffusion model, confirms this intuition. The graph shows that the jump model is able to generate some difference between index and individual option returns, but this difference is clearly much smaller than the observed return difference. Even though the jump risk premium cannot explain the discrepancy between index and individual options, the estimates do show some evidence for the existence of a jump risk premium. The contribution of the jump risk premium to the total equity premium of 6.01% varies from 0.77% to 1.00% for the different jump sizes.

Bakshi, Kapadia and Madan (2003) and Branger and Schlag (2004) consider jump-diffusion market models to explain observed differences in implied volatility functions between index and individual options. Our analysis of a pure jump-diffusion model is consistent with their findings that such a model can generate some discrepancy between index and individual option results, but as far as expected option returns are concerned, the discrepancy is much smaller than is empirically observed.

6.3 Statistical Tests of the Models

Finally, we conduct formal statistical tests of the models considered in this section and of the model with priced correlation risk proposed in this paper.

We test whether each model can fit the difference between the expected index and individual option returns across 13 delta’s ($\Delta = -0.8$ to -0.2 in steps of 0.05). The Wald test on expected-return differences incorporates that the risk premia on stock-market risk and on the second risk factor (depending on the model, stochastic volatility risk, jump risk, or correlation risk) are estimated. The test also corrects for serial correlation using Newey-West (1987) standard errors. For the SV and jump models, the Wald test has 13 degrees of freedom (13 Δ ’s), since the risk premia are not fitted to the difference between the expected option returns. For the correlation model, we

conservatively use 12 degrees of freedom, since fitting to $IC - RC$ could be interpreted as fitting to expected index versus individual option returns for one particular Δ .

Table 4 shows that the SV, market and jump model are very strongly rejected by the Wald test. The priced correlation-risk model yields a much smaller test statistic and is only marginally rejected at the 1% confidence level.

The Wald test results may suffer from the common problem that the 13 return differences are highly cross-correlated, thus inflating the test-statistic (see e.g. Cochrane (2001)). Looking at t -statistics per option is then more informative. The results in the second panel of Table 4 are clear: the correlation-risk model is almost never rejected, while the alternative models considered in section 6 (SV, market model and pure-jump model) are always rejected, for all Δ 's.

In sum, this section shows that standard jump and SV models without priced correlation risk, as well as a market model, generate some difference between returns on index versus individual options, but this difference is much smaller than the empirically observed difference. Formal statistical tests show that these models are always rejected, in strong contrast to the correlation-risk model.

7 Conclusion

We develop a parsimonious model of priced correlation risk for stock returns and implement it empirically using data on S&P100 index options and options on all index components. We demonstrate that correlation risk carries a large negative risk premium in our 1996-2003 sample, as evidenced by the substantial gap between option-implied correlations (46.7% on average) and realized correlations (28.7% on average). As a second contribution, we show that the entire index variance risk premium can be attributed to this correlation risk premium. While individual variance risk contributes to market variance changes, this risk is empirically not priced and cannot explain the well-documented variance risk premium. Priced correlation risk constitutes the missing link between unpriced individual variance risk and priced market variance risk. Finally, priced correlation risk enables us to offer a risk-based explanation for the discrepancy between index-option returns and individual-option returns. Index puts are expensive because they allow investors to hedge against correlation increases and the ensuing loss in diversification benefits. This forms the third contribution of the paper.

Why does correlation risk carry such a high risk premium? Remarkably, our analysis shows that it is not because of the intuition that correlations are countercyclical. Indeed, the systematic

part of correlation risk accounts only for a small fraction of the correlation risk premium, with more than 99% being compensation for orthogonal correlation risk. Examining which equilibrium model would generate such a large price of orthogonal correlation risk is an exciting area of research, which could extend recent work of Bansal and Yaron (2004) and Tauchen (2005).

In future research we also plan to study the relevance and implications of correlation risk premia in other settings. Obvious applications include portfolio choice problems, modeling equity risk premia using time-varying beta's, and the pricing of basket credit derivatives, such as collateralized debt obligations.

Appendix A

This appendix contains the proofs that were omitted from the main text.

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

$$d\rho = \lambda(\bar{\rho} - \rho) dt + \sigma_\rho \sqrt{\rho} dB \quad (34)$$

never reaches 0 from a strictly positive initial condition $\rho(0)$ with probability 1 under the Feller condition $\lambda\bar{\rho} > \sigma_\rho^2/2$. As the drift term $\lambda(\bar{\rho} - \rho) dt$ of (34) is strictly positive when $\rho \rightarrow 0$, it dominates the diffusion term $\sigma_\rho \sqrt{\rho} dB$ (in the sense that $|\lambda(\bar{\rho} - \rho) dt| \geq |\sigma_\rho \sqrt{\rho} dB|$ with probability 1) when $\rho \rightarrow 0$ and the Feller condition is satisfied. As $|\sigma_\rho \sqrt{\rho} \sqrt{1 - \rho} dB| \leq |\sigma_\rho \sqrt{\rho} dB|$ for $\rho \rightarrow 0$, the diffusion term of (5) is also dominated by the drift, and hence the solution to (5) never reaches 0 from a strictly positive initial condition $\rho(0)$ with probability one if $\lambda\bar{\rho} > \sigma_\rho^2/2$.

To show that the solution to (5) never reaches the upper boundary we show that 0 is never reached by $1 - \rho$, which follows

$$d(1 - \rho) = \lambda((1 - \bar{\rho}) - (1 - \rho)) dt - \sigma_\rho \sqrt{\rho} \sqrt{1 - \rho} dB \quad (35)$$

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

$$d(1 - \rho) = \lambda((1 - \bar{\rho}) - (1 - \rho)) dt - \sigma_\rho \sqrt{1 - \rho} dB \quad (36)$$

This diffusion never reaches 0 from a strictly positive initial condition $1 - \rho(0)$ with probability 1

under the Feller condition $\lambda(1 - \bar{\rho}) > \sigma_\rho^2/2$. Following the logic for the lower boundary case, we show that the solution to (35) is bounded away from 0 with probability one if $\lambda(1 - \bar{\rho}) > \sigma_\rho^2/2$. ■

Proof of Proposition 1: An $N \times N$ symmetric matrix is positive definite if and only if all its N leading principal minors are strictly positive. Take the correlation matrix $\Omega(t) = [\rho_{ij}(t)]$ at time t and omit the time argument. For $\rho = 0$, $\Omega = I_N$ and hence is positive definite. For $\rho \in (0, 1)$ first look at the second order principal leading minor:

$$|M_2| = \begin{vmatrix} 1 & \rho \\ \rho & 1 \end{vmatrix} = 1 - \rho^2 > 0 \text{ for } \rho(t) \in (0, 1) \quad (37)$$

The third order principal leading minor can be transformed into the determinant of the 2×2 matrix as follows. Subtract from all rows $i > 1$ the first row multiplied by a positive number ρ :

$$|M_3| = \begin{vmatrix} 1 & \rho & \rho \\ \rho & 1 & \rho \\ \rho & \rho & 1 \end{vmatrix} = \begin{vmatrix} 1 & \rho & \rho \\ 0 & 1 - \rho^2 & \rho - \rho^2 \\ 0 & \rho - \rho^2 & 1 - \rho^2 \end{vmatrix} = (1 - \rho^2)^2 \begin{vmatrix} 1 & \frac{\rho - \rho^2}{1 - \rho^2} \\ \frac{\rho - \rho^2}{1 - \rho^2} & 1 \end{vmatrix} \quad (38)$$

For $\rho \in (0, 1)$ we get then $|M_3| > 0$ as $\frac{\rho - \rho^2}{1 - \rho^2} \in (0, 1)$ and $(1 - \rho^2) > 0$.

In general, using this scheme, the k^{th} order principal leading minor equals a positive scale factor times the determinant of the $(k - 1) \times (k - 1)$ symmetric matrix with diagonal elements all one, and equal off-diagonal elements in the range $(0, 1)$. Thus all leading principal minors of the given $N \times N$ matrix are positive, and the matrix $\Omega(t)$ is positive definite, for all t .

The volatility matrix $\Phi = [\phi_i]$, where $i = 1, \dots, N$ has positive diagonal elements $\phi_i(t)$ and zero off-diagonal elements. Hence all its principal leading minors are positive and the matrix is positive definite. Next define A as the symmetric positive definite matrix with zero off-diagonal elements and $\sqrt{\phi_i}$ 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 $w_1^T = w^T A = (A^T w)^T$, as Ω is positive definite. ■

Proof of Proposition 2: Given the model, the index process follows

$$\begin{aligned} \frac{dI}{I} &= \mathbf{w}^T \boldsymbol{\mu} dt + \mathbf{w}^T \Phi^{1/2} d\mathbf{B} \\ &= \mathbf{w}^T \boldsymbol{\mu} dt + \mathbf{w}^T \Phi^{1/2} \Omega^{1/2} d\widehat{\mathbf{B}} \end{aligned}$$

where $\Phi = [\phi_i]$ has positive diagonal elements $\phi_i(t)$ following (9) and zero off-diagonal elements, and Ω is the stochastic correlation matrix defined in Proposition 1. The variance of the value-weighted index return is

$$\phi_I^2 = \mathbf{w}^T \Phi^{1/2} \Omega \Phi^{1/2} \mathbf{w} = \sum_{i=1}^N \left[w_i^2 \phi_i^2 + \sum_{i \neq j} w_i w_j \phi_i \phi_j \rho \right].$$

Using Ito's Lemma, derive the process for variance:

$$\begin{aligned} d\phi_I^2 &= \left[\sum_{i=1}^N \sum_{i \neq j} w_i w_j \phi_i \phi_j \right] d\rho + \frac{1}{2} \sum_{i=1}^N \left[w_i^2 + 2 \sum_{i \neq j} w_i w_j \phi_j \rho \right] \phi_i^{-1} \left[\gamma_i - \frac{1}{2} \varsigma_i^2 \right] dt \\ &\quad + \frac{1}{2} \sum_{i=1}^N \left[w_i^2 + 2 \sum_{i \neq j} w_i w_j \phi_j \rho \right] \varsigma_i dB_{\phi_i}. \end{aligned}$$

The only priced source of risk in the index variance is the correlation process Wiener dB_ρ . Moving from P to Q , the process changes to

$$\begin{aligned} d\phi_I^2 &= \nu_\phi d\rho + \delta_\phi dt + \sum_{i=1}^N \nu_i dB_{\phi_i} \\ &= \nu_\phi \left[\lambda(\bar{\rho} - \rho) dt - \kappa \sigma_\rho \sqrt{\rho(1-\rho)} + \sigma_\rho \sqrt{\bar{\rho}} \sqrt{1-\rho} dB_\rho^Q \right] + \delta_\phi dt + \sum_{i=1}^N \nu_i dB_{\phi_i} \end{aligned}$$

As claimed, the change in drift in the index variance is proportional (scaled by ν_ϕ) to the adjustment in the correlation state variable process itself. ■

Proof of Proposition 3: From equations (15) and (16) it is clear that the denominators of their right-hand sides are equal up to a change of measure: in $IC(t)$ the denominator contains an expectation of integrated product of individual volatilities under Q , while in $RC(t)$ the expectation of the same expression is taken under P . As the individual variance risk is not priced, the individual variance process in (9) does not change when we move from P to Q . Applying Ito's lemma shows that the process for the product of individual volatilities is not affected when changing from P to Q . This directly implies

$$\begin{aligned} \sum_{i=1}^N \sum_{i \neq j} w_i w_j E_t^Q \left[\int_t^{t+\tau} \phi_i(s) \phi_j(s) ds \right] &= \sum_{i=1}^N \sum_{i \neq j} w_i w_j E_t^P \left[\int_t^{t+\tau} \phi_i(s) \phi_j(s) ds \right] \\ \sum_{i=1}^N w_i^2 E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right] &= \sum_{i=1}^N w_i^2 E_t^P \left[\int_t^{t+\tau} \phi_i^2(s) ds \right] \end{aligned}$$

To compare $IC(t)$ with $RC(t)$ it suffices to compare the part of the nominators of the definition:

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

The second equality follows from Fubini Theorem. From Proposition 2 and applying the expectation operator:

$$\begin{aligned}
E_{t_0}^P [\phi_I^2(t_1)] &= \phi_I^2(t_0) + E_{t_0}^P \left[\int_{t_0}^{t_1} \nu_\phi \lambda (\bar{\rho} - \rho) d\tau + \int_{t_0}^{t_1} \delta_\phi d\tau \right] \\
E_{t_0}^Q [\phi_I^2(t_1)] &= \phi_I^2(t_0) + E_{t_0}^Q \left[\int_{t_0}^{t_1} \nu_\phi \lambda \left(\bar{\rho} - \frac{\kappa \sigma_\rho \sqrt{\rho(1-\rho)}}{\nu_\phi \lambda} - \rho \right) d\tau + \int_{t_0}^{t_1} \delta_\phi d\tau \right]
\end{aligned}$$

Hence

$$\begin{aligned}
IC - RC &= -\Xi * \left(E_{t_0}^Q \left[\int_{t_0}^{t_1} \nu_\phi \lambda (\bar{\rho}^* - \rho) d\tau \right] - E_{t_0}^P \left[\int_{t_0}^{t_1} \nu_\phi \lambda (\bar{\rho} - \rho) d\tau \right] \right) \\
\text{where } \bar{\rho}^* &= \bar{\rho} - \frac{\kappa \sigma_\rho \sqrt{\rho(1-\rho)}}{\nu_\phi \lambda}
\end{aligned}$$

As $\nu_\phi \lambda > 0$ *a.s.* and $\Xi > 0$ *a.s.* and both do not depend on the correlation risk premium κ , κ (or, more generally, its dynamics) unambiguously defines the dynamics of the relation $IC - RC$. The long-run mean of the index variance process under Q is decreasing in κ , and hence the difference $IC - RC$ is monotonically decreasing in the correlation risk premium κ .

If correlation risk is not priced ($\kappa = 0$), we have $IC - RC = 0$ and because of the monotonic relation we also have $IC - RC < 0$ for $\kappa > 0$, and $IC - RC > 0$ for $\kappa < 0$. ■

Appendix B

This appendix elaborates some aspects of the analysis in Section 1: lemma 2 on the parameter restrictions for the process for $\rho(t)$ to remain within the $(0, 1)$ interval under Q , and lemma 3 on the relationship between $RNEAC(t)$ defined in (14) and $IC(t)$ defined in (15).

Lemma 2: For $\kappa < 0$, the correlation state variable $\rho(t)$ following (6) with initial condition $\rho(0) \in$

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

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

$$d\rho = \lambda\left(\bar{\rho} - \frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda} - \rho\right)dt + \sigma_\rho\sqrt{\rho(1-\rho)}dB_\rho^Q \quad (39)$$

For $\kappa < 0$, the long-run mean of the latent correlation process is higher under Q than under P . To prove that the state space for the process is of the form $(0, 1)$, i.e. that 0 and 1 are natural (Feller) boundaries for some parameter values, we have to show that the boundaries are unattainable for these parameters.

We see that $\bar{\rho}^* = \bar{\rho} - \frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda}$ approaches the initial long-run mean $\bar{\rho}$ from above when ρ approaches 0 or 1, and the initial process (5) approaches the boundary 0 faster (informally) than the process under Q . Then the boundary 0 is a natural boundary for the process under Q if $\lambda\bar{\rho} > \sigma_\rho^2/2$.

For the upper boundary we modify Lemma 1's proof slightly. Approximate process (39) as

$$d\rho_a = \begin{cases} \lambda\left(\bar{\rho} - \frac{\kappa\sigma_\rho\sqrt{\rho(1-\rho)}}{\lambda} - \rho\right)dt + \sigma_\rho\sqrt{\rho(1-\rho)}dB_\rho^Q & \text{if } 1 - \rho > \varepsilon \\ \lambda\left(\bar{\rho} - \frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda} - \rho\right)dt + \sigma_\rho\sqrt{\rho(1-\rho)}dB_\rho^Q & \text{if } 1 - \rho \leq \varepsilon \end{cases} \quad (40)$$

where ε is a very small positive number

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 $\frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}$. The process remains continuous.

It can be shown that for any ρ closer to the upper boundary 1 than ε the following is true:

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

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:

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

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 boundary²⁵ for the approximating process (40) are also sufficient for 1 being an unattainable boundary for the original process (5).

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 $\bar{\rho}$ by a new long-run mean $\bar{\rho} - \frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}$ for a small positive number ε . Then the boundary 1 is unattainable if, for some arbitrarily small positive ε , $\lambda \left(1 - \bar{\rho} + \frac{\kappa\sigma_\rho\sqrt{\varepsilon(1-\varepsilon)}}{\lambda}\right) > \sigma_\rho^2/2$. ■

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 (14) and $IC(t)$ defined in (15) is given by:*

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

Proof of Lemma 3: From the definitions,

$$IC(t) = RNEAC(t) \frac{\sum_{i=1}^N \sum_{i \neq j} w_i w_j E_t^Q \left[\int_t^{t+\tau} \phi_i(s) \phi_j(s) ds \right]}{\sum_{i=1}^N \sum_{i \neq j} w_i w_j \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]} \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_j^2(s) ds \right]}}. \quad (41)$$

Therefore, it suffices to determine the following inequality:

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

From its definition $\phi_i(t)$ is a strictly positive stochastic process. We can then apply the Cauchy-Bunyakovsky-Schwarz inequality twice to obtain:

$$E_t^Q \left[\int_t^{t+\tau} \phi_i(s) \phi_j(s) ds \right] \leq E_t^Q \left[\sqrt{\int_t^{t+\tau} \phi_i^2(s) ds} \sqrt{\int_t^{t+\tau} \phi_j^2(s) ds} \right] \quad (42)$$

²⁵Here 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.

and

$$E_t^Q \left[\sqrt{\int_t^{t+\tau} \phi_i^2(s) ds} \sqrt{\int_t^{t+\tau} \phi_j^2(s) ds} \right] \leq \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_i^2(s) ds \right]} \sqrt{E_t^Q \left[\int_t^{t+\tau} \phi_j^2(s) ds \right]} \quad (43)$$

from which the result follows directly. ■

Appendix C

This appendix contains the implementation details of Section 6, in particular the methodology for identifying the unknowns in equations (31) and (33).

1) SV Model and Market Model

For the SV model (including the market model, for which $W = 1$), we need to calculate Δ_V and $\Delta_{\tilde{V}}$. For simplicity, we first assume that the option price is given by $E_t^Q[BS(1, \sqrt{\frac{1}{\tau} \int_t^{t+\tau} (V(s) + \tilde{\phi}_i^2) ds})]$ (Hull and White (1987)), where τ is the option maturity and $BS(S, \sigma)$ denotes the Black-Scholes option price formula as a function of the underlying price S and volatility σ . Because we have Heston processes, rather than Hull-White processes, this simplification ignores the leverage effect. However, the simplification is only needed to calculate the volatility-sensitivity (not to calculate option prices, since we directly use sample moments for this) and extensive simulations (reported below) show that the approximation is very accurate. Note that the leverage effect is fully accounted for in (31) and (33), as is clear from the presence of $\eta V(t) \bar{\sigma} \omega \Delta_V W^{-1/2}$ and $\eta V(t) \bar{\sigma} \omega \Delta_{\tilde{V}} \sqrt{W}$.

To calculate partial derivatives with respect to $V(t)$, we approximate the option price by $BS(1, \sqrt{E_t^Q[\frac{1}{\tau} \int_t^{t+\tau} (V(s) + \tilde{\phi}_i^2) ds]})$. The expected variance paths are given by

$$\begin{aligned} \sigma_{BS,P}^2 &= E_t \left[\frac{1}{\tau} \int_t^{t+\tau} V(s) ds \right] = \nu + \frac{1 - e^{-(\alpha + \xi \bar{\sigma})\tau}}{(\alpha + \xi \bar{\sigma})\tau} (V(s) - \nu) \\ \sigma_{BS,I}^2 &= E_t \left[\frac{1}{\tau} \int_t^{t+\tau} \tilde{V}(s) ds \right] = W\nu + \frac{1 - e^{-(\alpha + \xi \bar{\sigma})\tau}}{(\alpha + \xi \bar{\sigma})\tau} (\tilde{V}(s) - W\nu) \end{aligned} \quad (44)$$

We can then calculate the Δ_V 's for individual options

$$\begin{aligned} \Delta_V &\simeq \frac{\partial BS(1, \sigma_{BS,P})}{\partial V} \Big|_{V=V(t)} = \frac{\partial BS(1, \sigma_{BS,P})}{\partial \sigma_{BS,P}} \frac{\partial \sigma_{BS,P}}{\partial \sigma_{BS,P}^2} \frac{\partial \sigma_{BS,P}^2}{\partial V} \Big|_{V=V(t)} \\ &= \frac{\partial BS(1, \sigma_{BS,P})}{\partial \sigma_{BS,P}} \frac{1}{2\sigma_{BS,P}} \frac{1 - e^{-(\alpha + \xi \bar{\sigma})\tau}}{(\alpha + \xi \bar{\sigma})\tau} \end{aligned} \quad (45)$$

and similarly for index options

$$\begin{aligned}\Delta_{\tilde{V}} &\simeq \frac{\partial BS(1, \sigma_{BS,I})}{\partial \tilde{V}} \Big|_{\tilde{V}=\tilde{V}(t)} = \frac{\partial BS(1, \sigma_{BS,I})}{\partial \sigma_{BS,I}} \frac{\partial \sigma_{BS,I}}{\partial \sigma_{BS,I}^2} \frac{\partial \sigma_{BS,I}^2}{\partial \tilde{V}} \Big|_{\tilde{V}=\tilde{V}(t)} \\ &= \frac{\partial BS(1, \sigma_{BS,I})}{\partial \sigma_{BS,I}} \frac{1}{2\sigma_{BS,I}} \frac{1 - e^{-(\alpha + \xi\bar{\sigma})\tau}}{(\alpha + \xi\bar{\sigma})\tau}\end{aligned}\quad (46)$$

The first term is the regular Black-Scholes vega, the second term corrects for the fact that we need a partial derivative with respect to the variance, and the third term corrects for mean-reversion in variance over the option maturity τ .

We estimate the unconditional volatility risk premium as follows. For each option we replace $\sigma_{BS,I}^2$ and $\sigma_{BS,P}^2$ by the time-series averages of the implied volatility and evaluate the Black-Scholes vega (the first term in (45) and (46)) at this sample average. The ratio of equity-to-option price S/P (or I/P_I) is also replaced by its sample mean. For individual options, the remaining unknown term is $\bar{\sigma}\xi V(t) \frac{1 - e^{-(\alpha + \xi\bar{\sigma})\tau}}{(\alpha + \xi\bar{\sigma})\tau}$, which is constant across all strikes. The fact that only this term is needed to obtain expected returns is attractive, since we circumvent making specific assumptions about the mean reversion of volatility, or the volatility of volatility. We estimate the unconditional expectation of this term from a cross-sectional regression of the (annualized) average individual excess option returns on the option-specific term $\frac{S}{P} \frac{\partial BS(1, \sigma_{BS,P})}{\partial \sigma_{BS,P}} \frac{1}{2\sigma_{BS,P}}$. Here we correct for the presence of an equity diffusion risk premium η by first subtracting $\frac{S}{P}\eta\Delta_S$ from the average option return for each option. Given that the volatility risk premium ξ does not enter the expected equity return directly, we can directly choose η to match the observed equity premium of 6.01%. Finally, given the estimate for $E[\bar{\sigma}\xi V(t) \frac{1 - e^{-(\alpha + \xi\bar{\sigma})\tau}}{(\alpha + \xi\bar{\sigma})\tau}]$, the last ingredient required to calculate expected index option returns is W . For the market model, $W = 1$ by definition. For the SV model, we calculate W using the actual weights of the S&P100 index and the implied correlation $IC(t)$ at each date, and subsequently take the time-series average, resulting in $W = 0.4755$. Using IC is highly conservative, since historical correlations would lead to a lower value for W and less negative expected index option returns.

THE EFFECT OF LEVERAGE ON THE VOLATILITY-SENSITIVITY

To examine the accuracy of the approximations for $\Delta_{\tilde{V}}$ and Δ_V in (45) and (46), we now simulate the Heston model of section 6.2 with parameter estimates from the recent empirical option-pricing literature²⁶ and calculate the volatility-sensitivity of individual and index options. Starting with the

²⁶The literature presents estimates for the parameters of the Heston model for index options. We obtain parameters

estimates of Pan (2002), the average (across all strikes) absolute percentage difference between the simulated volatility-sensitivities in the Heston model and the approximated volatility-sensitivities $\Delta_{\tilde{V}}$ and Δ_V is 0.77% for index options and 0.81% for individual options. The approximations are therefore very accurate, even with the -0.57 leverage effect estimated by Pan. When considering many alternative parameter estimates, the largest average absolute percentage differences between the ‘true’ and approximated vol-sensitivities arise for the parameter estimates of Ait-Sahalia and Kimmel (2005), who report a leverage effect of -0.75. Even then, the average absolute percentage differences are just 1.33% and 1.41%. It is clear that these small inaccuracies could not possibly alter our conclusion that the SV and market model are unable to simultaneously explain expected index- and individual-option returns.

2) Jump-Diffusion Model

The procedure to estimate the right-hand side of equations (31) and (33) for the jump-diffusion model contains the following steps. First, the ratio of the equity-to-option price S/P (or I/P_I) is replaced by its sample average. Second, the jump delta (both for index and individual options) is estimated from the cross-section of option prices (again relative to the stock price level). For example, given a jump size of $\theta = -10\%$, the jump delta for an ATM put is calculated by subtracting its price from the price of a 10% OTM put (interpolating implied volatilities linearly to obtain prices that are not directly observed). We use the average cross-section over the full sample period to estimate the jump delta’s.²⁷ Finally, we perform a cross-sectional regression of the (annualized) cross-section of average individual option returns on the Δ_S ’s and $\Delta_{J,S}$ ’s to estimate η and $(\lambda - \lambda^Q)$, imposing the restriction that the average equity return during the sample period is matched, i.e. $\eta + \theta(\lambda - \lambda^Q) = 0.0601$. We repeat this analysis for a range of downward jump sizes θ between -1% and -30%. We thus fit the jump risk premium to the average individual option returns, and then analyze the implications for index options.

for individual options by applying the W factor in the model.

²⁷We also used a Fama-MacBeth methodology to allow for time-varying volatility-sensitivities (section 6.1) and jump delta’s (section 6.2), stemming from time-variation in vega’s, I/P_I ratios and jump intensities. This leads to an even worse performance of the SV and jump models.

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] Bakshi, G. and N. Kapadia, 2003a, "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies* 16, 527-566.
- [4] Bakshi, G. and N. Kapadia, 2003b, "Volatility Risk Premium Embedded in Individual Equity Options: Some New Insights," *Journal of Derivatives*, 45-54.
- [5] 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.
- [6] Bansal, R. and A. Yaron, 2004, "Risks For The Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59, 1481-1509.
- [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] Branger, N. and C. Schlag, 2004, "Why is the Index Smile So Steep?," *Review of Finance* 8, 109-127.
- [18] Breeden, D. and R. Litzenberger, 1978, "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business* 51, 621-652.
- [19] Britten-Jones, M. and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance* 55, 839-866.
- [20] Broadie, M., M. Chernov and M. Johannes, 2005, "Model Specification and Risk Premiums: Evidence from Futures Options," working paper, Columbia Business School.
- [21] Buraschi, A. and J. Jackwerth, 2001, "The Price of a Smile: Hedging and Spanning in Option Markets," *Review of Financial Studies* 14, 495-527.
- [22] 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.
- [23] 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.
- [24] Carr, P. and L. Wu, 2004, "Variance Risk Premia," working paper, Bloomberg L.P.
- [25] Cochrane, J., 2001, "*Asset Pricing*," Princeton University Press.
- [26] Cochrane, J., F. Longstaff and P. Santa-Clara, 2004, "Two Trees," working paper, University of Chicago.

- [27] 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.
- [28] Coval, J. and T. Shumway, 2001, "Expected Option Returns," *Journal of Finance* 56, 983-1009.
- [29] 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.
- [30] Dennis, P. and S. Mayhew, 2002, "Risk-Neutral Skewness: Evidence from Stock Options," *Journal of Financial and Quantitative Analysis* 37, 471-493.
- [31] 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*.
- [32] Dumas, B., 1995, "The Meaning of the Implicit Volatility Function in Case of Stochastic Volatility," unpublished note, HEC.
- [33] Engle, R. and K. Sheppard, 2005, "Evaluating the Specification of Covariance Models for Large Portfolios," working paper, NYU.
- [34] Eraker, B., M. Johannes and N. Polson, 2003, "The Impact of Jumps in Volatility and Returns," *Journal of Finance* 58, 1269-1300.
- [35] Eraker, B., 2004, "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices," *Journal of Finance* 59, 1367-1403.
- [36] Garleanu, N., L. Pedersen and A. Poteshman, 2005, "Demand-Based Option Pricing," working paper, Wharton School of Business.
- [37] Heston, S., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies* 6, 327-343.
- [38] Hull, J. and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance* 42, 281-300.
- [39] Jackwerth, J. and M. Rubinstein, 1996, "Recovering Probability Distributions from Contemporaneous Security Prices," *Journal of Finance* 51, 1611-1631.

- [40] Jiang, G. and Y. Tian, 2005, "Model-Free Implied Volatility and Its Information Content," forthcoming in *Review of Financial Studies*.
- [41] Jones, C., 2005, "A Nonlinear Factor Analysis of S&P 500 Index Option Returns," forthcoming in *Journal of Finance*.
- [42] Jorion, P., 2000, "Risk Management Lessons from Long-Term Capital Management," *European Financial Management* 6, 277-300.
- [43] Karlin, S. and H. Taylor, 1981, *A Second Course in Stochastic Processes*, Academic Press, London UK.
- [44] Liu, J. and J. Pan, 2003, "Dynamic Derivative Strategies," *Journal of Financial Economics* 69, 401-430.
- [45] 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.
- [46] 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.
- [47] Merton, R., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics* 4, 141-183.
- [48] Moskowitz, T., 2003, "An Analysis of Covariance Risk and Pricing Anomalies," *Review of Financial Studies* 16, 417-457.
- [49] Newey, W. and K. West, 1987, "A Simple, Positive-Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55, 703-708.
- [50] Pan, J., 2002, "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics* 63, 3-50.
- [51] Roll, R., 1988, "The International Crash of October 1987," *Financial Analysts Journal* 44, 19-35.
- [52] Rubinstein, M., 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978," *Journal of Finance* 40, 455-480.

- [53] 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.
- [54] Skintzi, V. and A. Refenes, 2003, “Implied Correlation Index: A New Measure of Diversification,” working paper, Athens University.
- [55] Tauchen, G., 2005, “Stochastic Volatility in General Equilibrium,” working paper, Duke University.

Table 1: Index versus Individual Options: Variances

The table 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 Jan. 1996 - Dec. 2003 sample period. For individual options the variances are equally-weighted cross-sectional averages across all constituent stocks. Realized variance is calculated from daily returns over a 30-day window. The model-free implied variance 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 2.1. The data on option prices is from OptionMetrics and variances are expressed in annual terms. The p-value for the null hypothesis that implied and realized variance are on average equal is based on Newey-West (1987) autocorrelation consistent standard errors with 22 lags.

	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
<i>p</i> value for $H_0 : RV = MFIV$	0.0000	0.046

Table 2: Implied and Realized Correlations: Summary Statistics

The table reports summary statistics (time-series mean and standard deviation) for the implied correlation (IC), realized correlation (RC), correlation scaling factor and the derived correlation. IC is calculated from daily observations on model-free implied variances for the S&P100 index and for all index components, using (18). $RC(t)$ 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 (180-day) window of daily stock returns. For each pair of stocks, both stocks need to have at least 15 (90) non-zero returns over the 30-day (180-day) window. The correlation scaling factor is estimated by the extension of (18) that allows for heterogeneity. \widehat{IC} is also calculated from daily observations on model-free implied variances. The average derived correlation is the cross-sectional average of $\overline{\rho_{ij}}\widehat{IC}$, where $\overline{\rho_{ij}}$ is estimated by the 180-day historical correlation. The middle panel reports the time-series correlation between IC and RC, both in levels and changes, and for both estimates of RC (using a 30-day or 180-day window of daily returns). The bottom panel gives the results of predictive regressions of RC on lagged IC and on lagged RC, both univariately and bivariate. t-statistics are based on Newey-West (1987) autocorrelation consistent standard errors with 22 lags (90 lags for 180-day RC).

Series	Mean	Std. Dev.
Implied Correlation $IC(t)$	0.4668	0.1394
Average Realized Corr. $RC(t)$ (30 day)	0.2868	0.1247
Average Realized Corr. $RC(t)$ (180 day)	0.2888	0.0996
Correlation Scaling Factor $\widehat{IC}(t)$	1.6260	0.4819
Average Derived Correlation $\overline{\rho_{ij}}\widehat{IC}(t)$	0.4478	0.1417
Time-Series Corr. with $IC(t)$	$RC(t)$ (30d)	$RC(t)$ (180d)
Levels	0.6878	0.6490
Quarterly Changes	0.5650	0.2220
Predictive Regressions of $RC(t)$	$RC(t)$ (30d)	$RC(t)$ (180d)
Univariate: Slope for $IC(t)$	0.5122	0.4123
(t -stat)	(8.34)	(4.68)
R^2	0.3495	0.3531
Univariate: Slope for $RC(t-1)$	0.6295	0.5779
(t -stat)	(8.90)	(3.84)
R^2	0.3988	0.3375
Bivariate: Slope for $IC(t)$	0.2575	0.2599
(t -stat)	(4.48)	(3.60)
Bivariate: Slope for $RC(t-1)$	0.4255	0.3344
(t -stat)	(5.93)	(2.27)
R^2	0.4453	0.4178

Table 3: Decomposition of Index Variance Changes

We report the results of the decomposition of index variance changes described in section 3. OLS estimates of (21) are presented for the univariate case (regressing only on ΔVar and only on ΔCorr) and for the bivariate case (both regressors). The construction of the regressors ΔVar and ΔCorr are explained in section 3. The last column uses $\text{IC}(t-1)$ as a proxy for $\text{IC}(t)$. The second regressor, $\widehat{\Delta\text{Corr}}$, is then defined as in (20) except that $\text{IC}(t)$ in the second term is replaced by $\text{IC}(t-1)$.

Regression	Original $IC(t)$	Proxied $IC(t)$
Univariate: Var Slope	1.2023	1.2023
(t -stat)	(40.88)	(40.88)
R^2	0.4547	0.4547
Univariate: Corr Slope	1.0979	0.7975
(t -stat)	(63.73)	(31.37)
R^2	0.6696	0.3293
Bivariate: Var Slope	1.0199	1.0592
(t -stat)	(260)	(45.75)
Bivariate: Corr Slope	0.9916	0.6575
(t -stat)	(336)	(36.43)
R^2	0.9905	0.6720

Table 4: Statistical Tests of the Models

Table 4 reports formal statistical tests of the model with priced correlation risk, a pure SV model, a market model and a pure jump-diffusion model. The tests are conducted on the difference between the model prediction for the discrepancy between expected returns on index versus individual options, and the empirically observed discrepancy, for 13 different put options (indexed by Δ ranging from -0.8 to -0.2 in increments of 0.05). The Wald test on expected-return differences considers all 13 Δ 's, while the second panel reports t-statistics for each Δ . The test incorporates that the risk premia on both risk factors are estimated and also corrects for serial dependence. The Wald test has 12 degrees of freedom for the correlation model and 13 for the other models.

Model Test	Correlation Model	SV Model	Market Model	Jump Model
Wald test	26.5643	82.6745	55.2116	57.1185
<i>p</i> value	0.0089	3.45×10^{-12}	3.71×10^{-7}	1.71×10^{-7}

Single Option Test: <i>t</i> -stat per Δ for observed difference index-individual versus fitted				
Option	Correlation Model	SV Model	Market Model	Jump Model
$\Delta = -0.80$	-1.7868	-3.8584	-4.2458	-3.2268
$\Delta = -0.75$	-1.1566	-3.7153	-3.7272	-3.0642
$\Delta = -0.70$	-0.5480	-3.5040	-3.1189	-2.8553
$\Delta = -0.65$	-0.0847	-3.3812	-2.7591	-2.7678
$\Delta = -0.60$	0.2408	-3.3577	-2.5521	-2.7907
$\Delta = -0.55$	0.5020	-3.3485	-2.4225	-2.8317
$\Delta = -0.50$	0.7110	-3.3542	-2.3541	-2.8852
$\Delta = -0.45$	0.8539	-3.4305	-2.3565	-3.0038
$\Delta = -0.40$	0.8658	-3.6170	-2.4021	-3.1952
$\Delta = -0.35$	0.7813	-3.8763	-2.5123	-3.4443
$\Delta = -0.30$	0.6683	-4.1918	-2.6207	-3.7188
$\Delta = -0.25$	0.4819	-4.6569	-2.7281	-4.0832
$\Delta = -0.20$	0.1556	-5.3968	-2.9110	-4.6022
Average	0.1296	-3.8222	-2.8239	-3.2669

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 Jan. 1996 - Dec. 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 2.1. Realized variance is calculated from daily index returns over a 30-day window. Variances are expressed in annual terms.

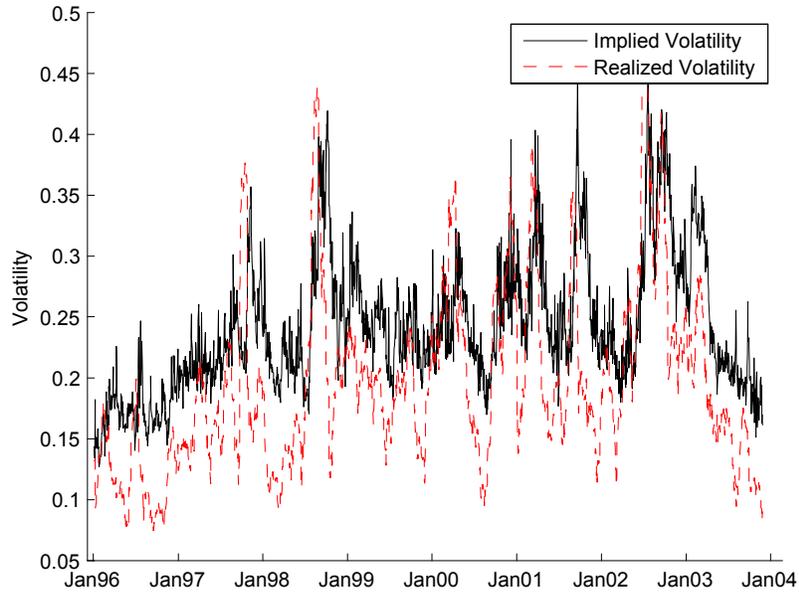


Figure 1B: Cross-Sectional Average of Implied versus Realized Volatility for Individual Options

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 Jan. 1996 - Dec. 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 2.1. 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.

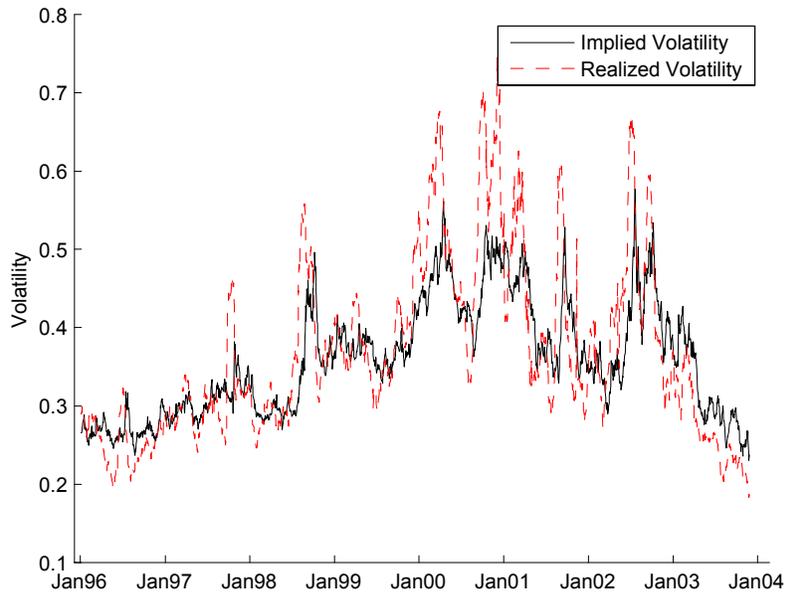


Figure 2A: Implied versus Realized Correlation

Figure 2A 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 (18). 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.

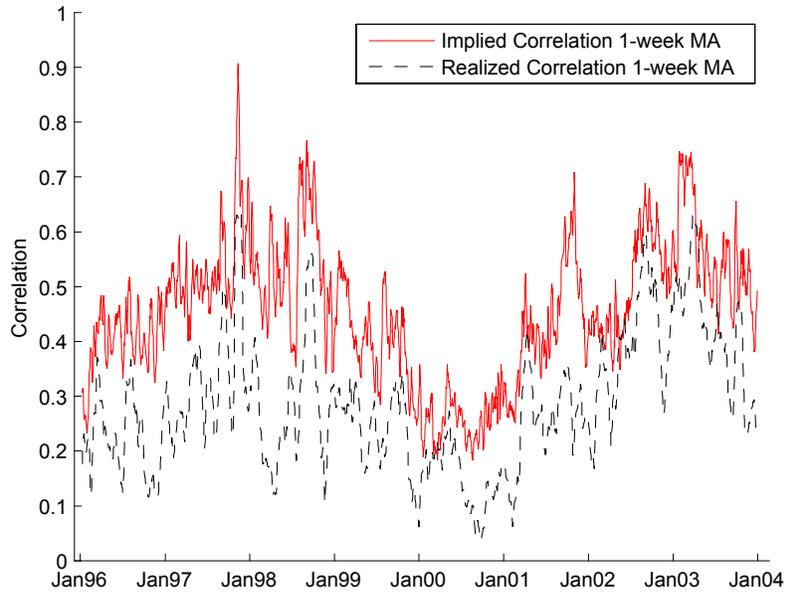


Figure 2B: Average Derived Correlation based on $\widehat{IC}(t)$ and Realized Correlation

Figure 2B shows 1-week moving averages of the average derived correlation and the realized correlation. The average derived correlation is the cross-sectional average of $\overline{\rho_{ij}}\widehat{IC}$, where $\overline{\rho_{ij}}$ is estimated by the 180-day historical correlation and \widehat{IC} by the extension of (18) that allows for heterogeneity. Like IC , \widehat{IC} is calculated from daily observations on model-free implied variances for the S&P100 index and for all index components. 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.

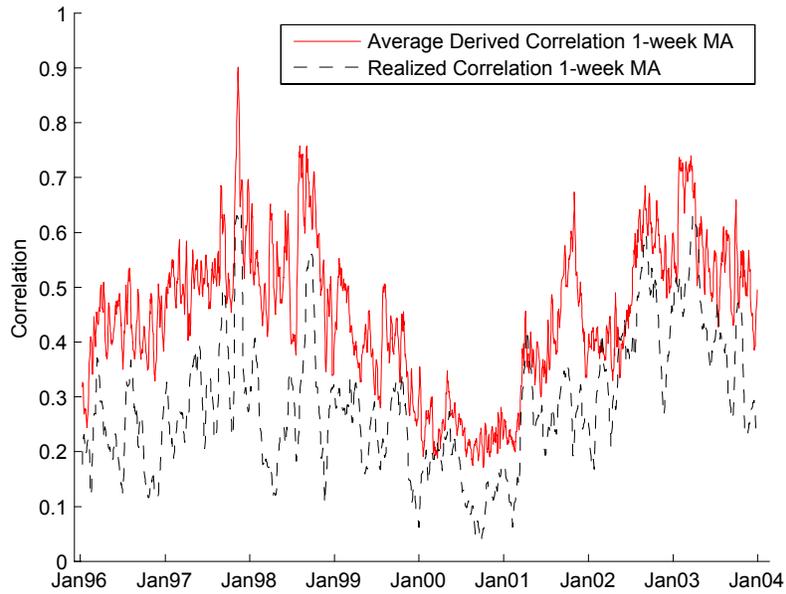


Figure 3A: One-Factor Excess Returns: Index and Individual Options

Figure 3A shows unconditional expected one-factor excess returns on 30-day S&P100 puts and individual puts over our Jan. 1996 - Dec. 2003 OptionMetrics sample, across a range of Black-Scholes delta's. Individual-option excess returns are index-weighted cross-sectional averages across individual puts on all the stocks in the S&P100. The unconditional expected one-factor excess returns are calculated as sample means of the time-series of 30-day holding-period returns, where each option payoff is calculated from demeaned distributions for the underlying asset. The de-meaning sets the expected return on the underlying asset equal to the riskfree rate minus the dividend yield. Confidence bounds are based on Newey-West (1987) standard errors with 22 lags.

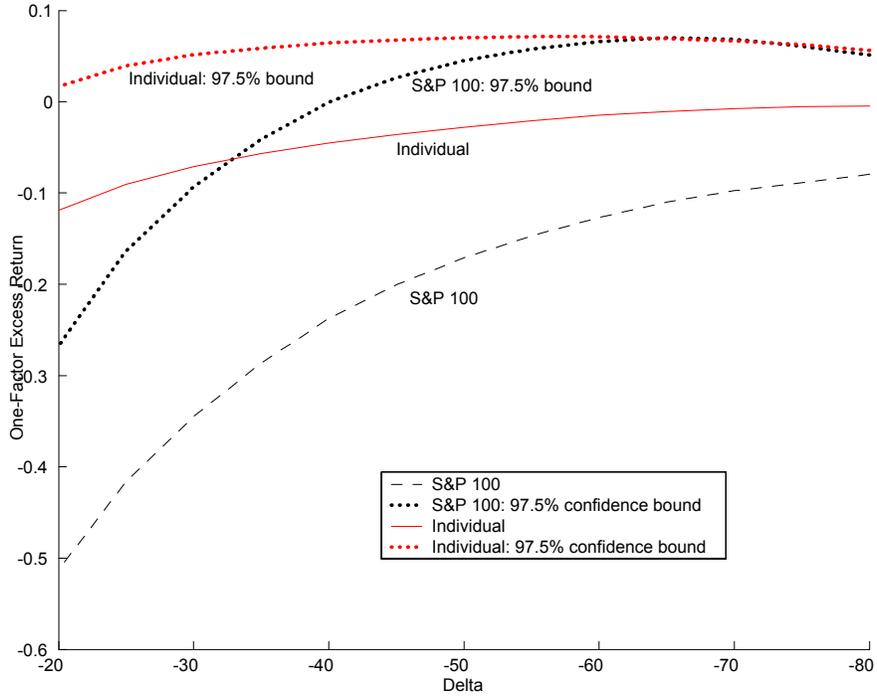


Figure 3B: Difference in One-Factor Excess Returns between Index and Individual Options

Figure 3B shows the difference in unconditional expected one-factor excess returns between 30-day S&P100 puts and 30-day individual puts over our Jan. 1996 - Dec. 2003 OptionMetrics sample, across a range of Black-Scholes delta's. Individual-option excess returns are index-weighted cross-sectional averages across individual puts on all the stocks in the S&P100. The unconditional expected one-factor excess returns and confidence bounds are calculated as in Figure 3A.

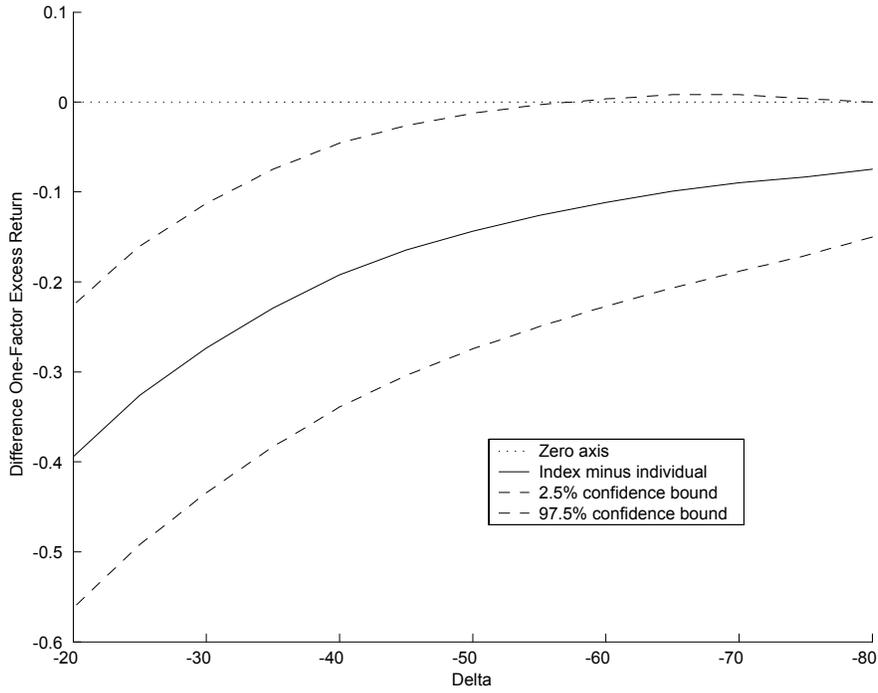


Figure 4: Expected Option Returns in the Correlation Risk Model

Figure 4 compares empirical unconditional expected option returns on 30-day index puts and (the index-weighted cross-sectional average of) individual puts with theoretical unconditional expected option returns predicted by the correlation risk model. The empirical unconditional expected option returns are calculated as sample means of the time-series of 30-day holding-period returns, where each option payoff is calculated from the distribution for the underlying asset (not demeaned like in Figure 3A). 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 a correlation risk premium calibrated to the observed difference between IC and RC (or set to zero in 1 case), as described in Section 5.

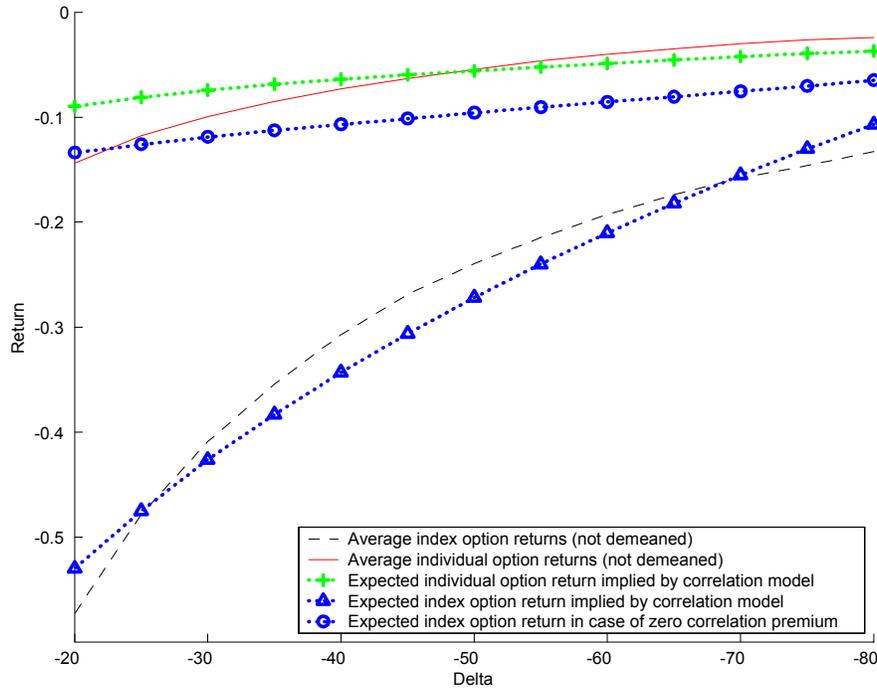


Figure 5: Expected Option Returns in the Market Model and SV Model

Figure 5 compares empirical unconditional expected option returns on 30-day index puts and (the index-weighted cross-sectional average of) individual puts with theoretical unconditional expected option returns predicted by the one-factor SV model and by the market model. The empirical unconditional expected option returns are calculated as in Figure 4. The model predictions for expected index-option returns are endogenously generated by the model given a volatility risk premium calibrated to expected individual-option returns, as described in Section 6 and Appendix C.

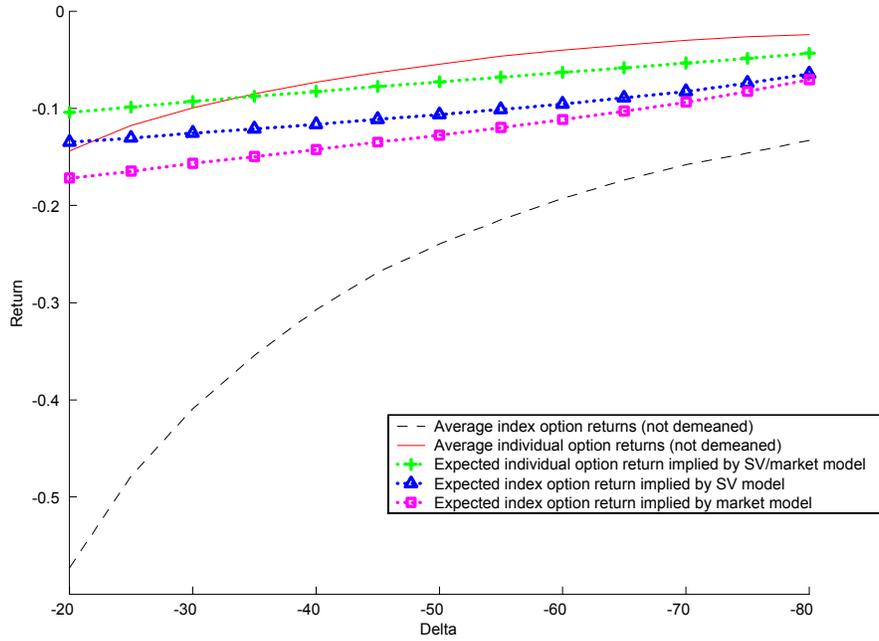


Figure 6: Expected Option Returns in the Jump-Diffusion Model

Figure 6 compares empirical unconditional expected option returns on 30-day index puts and (the index-weighted cross-sectional average of) individual puts with theoretical unconditional expected option returns predicted by the jump-diffusion model. The empirical unconditional expected option returns are calculated as in Figure 4. The model predictions for expected index-option returns are endogenously generated by the model given a jump risk premium calibrated to expected individual-option returns, as described in Section 6 and Appendix C.

