

Option Implied Dependence

Carole Bernard,^{*} and Oleg Bondarenko^{†‡}

Preliminary draft: June 25, 2019

Abstract

Options prices provide information on the future anticipated probability distribution of asset returns. However, information on the dependence among assets is difficult to extract and interpret. We propose a new model-free method to obtain the full joint risk-neutral distribution among assets that is consistent with the prices of all individual assets and of index options. In the empirical application, we implement our approach using options on the S&P 500 index and its nine sectors. We compare it with parametric approaches based on multivariate Gaussian dependence and multivariate student dependence as well as skewed dependence. In particular, we find that the option-implied dependence is highly non-normal and time-varying. Using the estimated dependence, we then study the correlation risk conditional on the market going down or up. We find that the risk premium for the down correlation is strongly negative, whereas it is positive for the up correlation. These findings are consistent with the economic intuition that the investors are particularly concerned with the loss of diversification when financial markets fall. As a result, they are willing to pay a considerable premium to hedge against increases in correlation during turbulent times. However, the investors actually prefer high correlation when markets rally.

Key-words: Model-free dependence, implied correlations, forward-looking dependence, down and up correlation.

1 Introduction

There are well developed techniques to infer the forward-looking risk-neutral distribution of an asset return when a wide range of traded options on this asset is available. These techniques rely on the no-arbitrage relationship first examined by Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978) and can estimate risk-neutral distributions in a model-free way;

^{*}Carole Bernard, Department of Accounting, Law and Finance, Grenoble Ecole de Management, Univ. Grenoble Alpes ComUE, and Department of Economics and Political Sciences at Vrije Universiteit Brussel (VUB) (email: carole.bernard@grenoble-em.com).

[†]Oleg Bondarenko, Department of Finance, University of Illinois at Chicago (email: olegb@uic.edu).

[‡]The authors gratefully acknowledge funding from the Canadian Derivatives Institute (formerly called IFSID) and the Global Risk Institute (GRI) for the related project “Inferring tail correlations from option prices”. We are also thankful for the numerous discussions with Steven Vanduffel that contributed a lot to this paper. We thank the participants and discussants at the CDI September 2018 in Montreal, to the Research in Options 2018 Honoring Bruno Dupire’s 60th Birthday, to the 2019 conference on Risk Management and Financial Innovation Conference: In Memory of Peter Christoffersen, to the 2019 Finance ITAM conference, and to the EFMA 2019 meeting at the University of Azores.

see, for example, Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (2000), and Bondarenko (2003). In this paper, we generalize this approach to higher dimensions in that we use the prices of options written on individual assets as well as on the corresponding index to fully describe the dependence among assets in a forward-looking way, i.e., we derive a risk-neutral dependence.¹

Our method is new and builds on the Block Rearrangement Algorithm² (BRA) presented in Bernard, Bondarenko, and Vanduffel (2018). They show that finding a dependence among asset returns that is compatible with all available option prices on the assets and their weighted sum (the index), can be cast as a combinatorial problem in which a matrix needs to be arranged in a suitable manner. While there might be various compatible dependence structures, the BRA is shown to be the most natural one in the sense that one obtains the “most likely” implied dependence among asset returns given the information comprised in available option prices and given that no additional prior information is used. Our method makes it possible to obtain a complete spectrum for the implied dependence among asset returns and thus compares favorably with existing methods for estimating dependence, which primarily focus on the average correlation among asset returns. Moreover, our approach makes it possible to study dependence conditional on specific events, such as the market going down or up.

As a main application, we use our method to study correlation risk premia. Specifically, using options on the S&P 500 and its nine sectors, we find evidence of a slightly negative correlation premium and, as such, confirm results reported, among others, in Driessen, Maenhout, and Vilkov (2012), Buraschi, Trojani, and Vedolin (2014), Faria, Kosowski, and Wang (2016), and Buss, Schoenleber, and Vilkov (2017). These authors also use option data to assess implied dependence in a forward-looking way, but their methodology cannot be used for the study of possible state-dependent features of correlations. In particular, we study the down and up correlations (i.e., correlations conditional on the S&P 500 return being below or above its median value) and find that the risk premium is significantly negative for the down correlation, is significantly positive for the up correlation, and only marginally negative for the global correlation. Thus, the contribution of the global negative correlation premium that was found in Driessen, Maenhout, and Vilkov (2009, 2012) derives mainly from the difference between real-world and risk-neutral correlations when the market is going down. This feature confirms the economic intuition that the correlation premium is linked to the loss of diversification risk when the market falls (Driessen, Maenhout, and Vilkov (2009), Buraschi, Kosowski, and Trojani (2013)). Additional discussion of the correlation risk premium³ and its link to macroeconomic variables can be found in Faria, Kosowski, and Wang (2016), Engle and Figlewski (2014), Mueller, Stathopoulos, and Vedolin (2017), Pollet and Wilson (2010), Harvey, Liu, and Zhu (2016), and Cosemans (2011).

As already mentioned, we are not the first to estimate implied dependence and correlation risk premia using option data. A first attempt to measure implied dependence was proposed by the Chicago Board Options Exchange (CBOE), which started to disseminate S&P 500 Implied

¹Specifically, we derive in a first step a risk-neutral joint distribution of assets (X_1, \dots, X_d) . Clearly, this joint distribution is driven by the dependence among the assets, but also by their marginal distribution functions F_i . By applying in a second step the transformation $X_i \rightarrow F_i(X_i)$ we bring the assets on the same uniform scale and dependence becomes fully described by the joint distribution of $(F_1(X_1), \dots, F_d(X_d))$.

²The Block Rearrangement Algorithm generalizes the standard Rearrangement Algorithm introduced by Puccetti and Rüschendorf (2012). This algorithm has applications in various disciplines. Embrechts, Puccetti, and Rüschendorf (2013) use the RA in quantitative risk management to assess the impact of model uncertainty on Value-at-Risk estimates for portfolios. The algorithm has also important applications in operations research (fair allocation of goods) and engineering (image reconstruction).

³The correlation risk premium is also closely related to the variance risk premium (Bollerslev and Todorov (2011), Schneider and Trojani (2015)) and is positively related to the disagreement risk (Buraschi, Trojani, and Vedolin (2014)).

Correlation Indexes in July 2009.⁴ These are now well accepted dependence measures, based on individual implied volatilities and index implied volatility, and thus driven solely by option prices.

The CBOE approach has been significantly improved and applied to the estimation of correlation risk premia by Driessen, Maenhout, and Vilkov (2009, 2012). Specifically, Driessen, Maenhout, and Vilkov (2009) were the first to find evidence of priced correlation risk and to provide a risk-based explanation for it. Index options are expensive (i.e., correlation risk is priced) because they allow investors to insure against the risk of a loss in diversification benefits they experience during turbulent times when correlations are high. Driessen, Maenhout, and Vilkov (2012) propose a stochastic correlation model that is driven by ideas that underpin the CBOE approach and show it can be used to extract the implied average correlation (they assume equal correlations for each pair of stocks). Using data on index and individual options for all index constituents, they estimate the time-series of the implied average correlation for the S&P500 and DJ30 index and show these are significantly higher than the average realized correlations. The model of Driessen, Maenhout, and Vilkov (2012) is parsimonious in that at time t all pairwise correlations are identical and driven by a mean reverting process $\rho(t)$ (in the same spirit as in Cochrane, Longstaff, and Santa-Clara (2008)). Buss and Vilkov (2012) and Buss, Schoenleber, and Vilkov (2017) are able to relax the assumption that all pairwise correlations are equal, but this comes with the cost of imposing additional structure on dependence. We further discuss the features of these methods in Section 2.

The CBOE methodology and its improvements are confronted with two main challenges. First, they provide an incomplete measurement of dependence insofar as the obtained (average) pairwise correlations provide a limited perspective on dependence among assets. In contrast, our method yields a complete joint model for the asset returns and thus also depicts the full spectrum of dependence. Furthermore, it is documented in the literature that correlations among companies are time-dependent and tend to be stronger in market downturn phases; see for instance Bollerslev, Engle, and Wooldridge (1988) for empirical evidence under the historical probability measure. Finding such evidence in a forward-looking way is not possible using CBOE implied correlations, since these cannot be linked to specific market states. Using our methodology, however, we are able to measure correlation risk premia conditional on the market going down or up and to provide for the first time evidence that they are linked directly to the desire of investors to hedge against the loss of diversification in downturns.

The rest of the paper is organized as follows. In Section 2, we discuss the state of research on implied correlation and review, in particular, the methodology used by the CBOE to calculate its correlation indices. In Section 3, we present our method for inferring dependence among several assets. In Section 3, we implement the approach using options written on the S&P 500 index and on each of the nine sectors that comprise the index. In Section 5, we introduce new correlation indices and study the correlation risk-premium. We conclude in Section 6.

2 State of Research on Estimating Implied Correlation

To better understand their features and limitations, we review in this section existing approaches to extracting forward-looking estimates of dependence from option prices. We first discuss the so-called CBOE implied correlation index, followed by a discussion of alternative approaches that arise from recent literature on estimating the correlation risk premium and tail risk. These approaches are certainly useful, but they are also limited in that they make it neither possible to derive a complete dependence nor to assess its state-dependent features without making strong

⁴The methodology is outlined in a white paper from the Chicago Board Options Exchange (2009).

model assumptions. As far as we are aware only the approach that we outline in Section 3 overcomes these limitations.

2.1 CBOE Implied Correlation Index

The methodology used by CBOE to construct implied correlation indices is explained in a white paper on the CBOE S&P 500 implied correlation index (Chicago Board Options Exchange (2009)). Using ATM option prices, this index is an attempt to estimate the average pairwise correlation among the stocks that comprise the S&P 500 index. It thus displays some information on dependence under the risk-neutral probability. In what follows, we summarize the CBOE methodology and discuss its properties.

Let $S_t = \sum_i \omega_{i,t} X_{i,t}$ denote a weighted index value. For ease of presentation, we omit the time index t . Let us recall the following relationship between the variance of the index and the variances of its components:

$$\text{var}(S) = \sum_{i=1}^d \omega_i^2 \text{var}(X_i) + 2 \sum_{i=1}^{d-1} \sum_{j>i} \omega_i \omega_j \rho_{ij} \sqrt{\text{var}(X_i)} \sqrt{\text{var}(X_j)}, \quad (2.1)$$

in which ρ_{ij} is the correlation between X_i and X_j . The basic idea of the CBOE implied correlation index is to replace in equation (2.1) the different ρ_{ij} by a single number ρ , which thus can be interpreted as an average correlation among the index components. However, the variances are not estimated from past return data but are forward-looking option-implied measures. That is, they are the square of the implied ATM volatilities for the index and its components. Specifically, the implied volatility for the index, σ_S , solves the following equation:

$$C_S = BSCall(S_0, K, r, T, \sigma_S, q), \quad (2.2)$$

where C_S is the market price of an ATM call on the index with strike $K = S_0$, r is the continuously compounded risk-free rate, q is the dividend yield, S_0 the initial index price, and $BSCall$ denotes the Black-Scholes price of a call option on a single asset with maturity T and strike K . The same process is used on each of the components of the index to compute their respective implied volatilities. Implied volatilities are standard deviations of log-returns under the assumption that these are normally distributed. So, using equation (2.1), in which variances are substituted by the square of implied volatilities, would be justified when $\log(S) = \sum_i \omega_i \log(X_i)$, which is not the case (see Appendix A for further discussion). Hence, some bias is introduced, and we can only state that

$$\sigma_S^2 \approx \sum_{i=1}^d \omega_i^2 \sigma_i^2 + 2 \sum_{i=1}^{d-1} \sum_{j>i} \omega_i \omega_j \sigma_i \sigma_j \tilde{\rho}_{ij} \quad (2.3)$$

where σ_i and σ_j denote the individual implied volatility of the index components i and j , ω_i and ω_j are the weights of i and j in the index, and $\tilde{\rho}_{ij}$ is the pairwise correlation of index components $\log(X_i)$ and $\log(X_j)$. The CBOE correlation index is then obtained by replacing $\tilde{\rho}_{ij}$ in (2.3) with a single correlation number that we denote by ρ_{cboe} :

$$\rho_{cboe} = \frac{\sigma_S^2 - \sum_{i=1}^d \omega_i^2 \sigma_i^2}{2 \sum_{i=1}^{d-1} \sum_{j>i} \omega_i \omega_j \sigma_i \sigma_j}. \quad (2.4)$$

Here, $d = 50$ since the CBOE correlation index is defined using a subset of the 50 largest components of the index, measured by market capitalization. Furthermore, the weights are defined by

$$\omega_i = \frac{P_i Y_i}{\sum_{j=1}^d P_j Y_j},$$

where P_i is the price of the i th index component and Y_i is the “float-adjusted shares outstanding of the i th index component,” and, finally, the implied volatility σ_i of each component of the subset of 50 stocks is an ATM volatility (interpolated between strikes if no ATM options are available). The exact procedure is described in the CBOE white paper (Chicago Board Options Exchange (2009)).

Assuming that the approximation (2.3) is exact, we obtain that the expression of ρ_{cboe} in (2.4) can be written as an average pairwise correlation

$$\rho_{cboe} = \frac{\sum_{i < j} \omega_i \sigma_i \omega_j \sigma_j \tilde{\rho}_{ij}}{\sum_{i < j} \omega_i \sigma_i \omega_j \sigma_j}. \quad (2.5)$$

However, since the approximation (2.3) is not exact, the CBOE implied correlation index cannot be strictly interpreted as an average pairwise correlation coefficient, and it can potentially take values that are strictly larger than 1.⁵ A more detailed discussion on the limitations of the CBOE index and of cases in which this implied correlation index can be accurately interpreted as an average pairwise correlation is provided in Appendix A.

2.2 Improved Approaches to Implied Correlations

The CBOE methodology for estimating the implied correlation index ρ_{cboe} has spurred research on estimating correlation risk premia, which in turn has led to further improvements of the CBOE methodology. Most notably, Driessen, Maenhout, and Vilkov (2012) use the same expression as in the CBOE implied correlation index formula, but the variances are not driven by at-the-money options only, but rather based on all option prices from which risk-neutral densities and next variances are deduced (see Bakshi, Kapadia, and Madan (2003) and Britten-Jones and Neuberger (2000)). The model of Driessen, Maenhout, and Vilkov (2012) is relatively parsimonious in that at time t all pairwise correlations are assumed to be equal to $\rho(t)$, which is a mean reverting process (in the same spirit as Cochrane, Longstaff, and Santa-Clara (2008)). The model is illustrated using S&P 500 and DJ30 options; see also Skintzi and Refenes (2005). Clearly, the only restriction, which is that the index variance equates with the variance of the portfolio of index components, is insufficient to pin down the entire correlation structure, i.e., there are still many possible ways to construct the dependence among the components.

Buss and Vilkov (2012) allow pairwise correlations among the assets to differ by making the structural assumption that $\rho_{ij}^{\mathbb{Q}} - \rho_{ij}^{\mathbb{P}} = \alpha(1 - \rho_{ij}^{\mathbb{P}})$ with $\alpha \in (0, 1)$, \mathbb{P} as the real-world measure and \mathbb{Q} as the risk-neutral measure. The parameter α is computed from historical correlation values, weights, and implied variances. Knowing α , the correlation matrix under the risk-neutral probability can be derived. By construction, the implied correlation matrix is positive definite with all coefficients between -1 and 1, the correlation risk premium is negative ($\rho^{\mathbb{Q}} > \rho^{\mathbb{P}}$), and the correlation risk premium is higher in magnitude for low or negatively correlated stocks that are exposed to a higher risk of losing diversification benefits, which is consistent with the literature (Mueller, Stathopoulos, and Vedolin (2017)). Buss and Vilkov (2012) combine the option-implied correlations with option-implied volatilities to compute implied betas and show their better predictive quality with respect to future realized betas. Buss, Schoenleber, and Vilkov (2017) also impose additional structure and relax the equal pairwise correlation constraint by estimating a block diagonal heterogeneous correlation matrix. These authors apply their method to study implied correlations inferred from the S&P 500 ETF (economy) and its nine sector ETFs (economic sectors).

⁵For example, the KCJ index was 100.8 on November 6, 2008, 105.93 on November 13, 2008, and 103.4 on November 20, 2008.

A very different approach to capturing information on the dependence using option prices is proposed by Kelly, Lustig, and Van Nieuwerburgh (2016). Their idea is as follows: There is always a negative difference between the price of a basket option and the weighted sum of individual options written on the index's components with well chosen strikes, in which equality can only be obtained when the assets exhibit perfect positive dependence (comonotonicity). This finding motivates the use of this difference among prices as a measure of dependence; the more negative the difference, the smaller the correlation (see also Kelly and Jiang (2014)).

Deriving forward-looking information on the dependence among assets is an important but difficult issue. One common limitation of the existing literature is that only partial dependence information is obtained; either average correlation estimates are derived or several additional assumptions need to be made in order to obtain an (unconditional) correlation matrix under the risk-neutral probability measure. By contrast, our approach makes it possible to obtain a dependence structure without imposing any restrictions and to relate it to market states. One of the key features is that we fully use the information on all options with given maturity on an index and on all its components. Specifically, in the next section we outline a novel approach that enables us to extract a compatible dependence structure using these option prices. It allows us to derive the full dependence among assets and not only some specific aggregate statistics, such as the global correlation.

3 Methodology for Inferring Dependence

Our method for inferring dependence among assets is model-free in that it is driven solely by the available information on index options and options on its components. Moreover, the dependence structure obtained is consistent with maximum entropy and thus the most likely one given the available financial data and given that no further prior information is used other than that revealed by option prices.

To fix notation, let $C(K)$ and $P(K)$ denote the time- t price of the European-style call and put options with strike K and fixed maturity T on the underlying asset X . Under the standard assumptions, the option prices are equal to the expected value of their payoffs under a suitably chosen *risk-neutral* probability measure \mathbb{Q} :

$$\begin{aligned} C(K) &= e^{-r\tau} E^{\mathbb{Q}} [(X - K)^+] = e^{-r\tau} \int_0^{\infty} (x - K)^+ f(x) dx, \\ P(K) &= e^{-r\tau} E^{\mathbb{Q}} [(K - X)^+] = e^{-r\tau} \int_0^{\infty} (K - x)^+ f(x) dx, \end{aligned}$$

where $\tau = T - t$ is time to maturity and $f(x)$ denotes the *risk-neutral density* (RND). The RND satisfies the relationship first established by Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978):

$$f(x) = e^{r\tau} \left. \frac{\partial^2 C(K)}{\partial K^2} \right|_{K=x} = e^{r\tau} \left. \frac{\partial^2 P(K)}{\partial K^2} \right|_{K=x}. \quad (3.6)$$

Similarly, the *risk-neutral cumulative distribution* (RNCD) satisfies:

$$F(x) = e^{r\tau} \left. \frac{\partial C(K)}{\partial K} \right|_{K=x} = e^{r\tau} \left. \frac{\partial P(K)}{\partial K} \right|_{K=x}. \quad (3.7)$$

Although not directly observable, the RNCD can be recovered using the relationship in (3.7), provided that prices of options with continuum of strikes $K \in \mathbb{R}$ are available. In practice,

option prices are only available for a finite number of strikes. Nevertheless, a number of efficient nonparametric approaches have been proposed in the literature that make it possible to circumvent this shortcoming; see, for example, Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (2000), and Bondarenko (2003).

Consider now an index with d assets. For ease of exposition we assume that all weights are equal to one. Denote by $X_{j,t}$ the value of each component at some given date t so that the value of the index is $S_t = \sum_{j=1}^d X_{j,t}$. When there is no confusion, we simplify notation and drop the time index t . For example, we write X_j instead of $X_{j,t}$ and S instead of S_t . We assume that the options market offers sufficient strikes in that the risk-neutral distributions F_1, \dots, F_d for the assets X_1, \dots, X_d can be estimated accurately. In addition, assume that the risk-neutral distribution F_S of the index itself can be estimated from the index options. Option prices are thus used to estimate the marginal distributions of the assets, and information on their dependence appears through knowledge of the marginal distribution of the index.

Armed with the risk-neutral cumulative distributions for the d assets and for the index, we construct a multivariate distribution for the d assets that is consistent with this information. We start by approximating each asset's RNCD with a discrete distribution (which can be done to any degree of accuracy), and we then build a dependence among these discrete distributions. For each asset $j \in \{1, \dots, d\}$, the distribution of X_j is approximated as follows. There are n possible states and X_j takes on the values x_{ij} , $i = 1, \dots, n$ with a probability $1/n$. Specifically, these elements x_{ij} are defined as the realizations $x_{ij} := F_j^{-1}(\frac{i-0.5}{n})$ ($i = 1, \dots, n$). Using this discretization, we are able to represent the multivariate vector of assets (X_1, X_2, \dots, X_d) by a $n \times d$ matrix:

$$\begin{bmatrix} x_{11} & x_{12} & \dots & x_{1d} \\ x_{21} & x_{22} & \dots & x_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nd} \end{bmatrix},$$

where the j -th column corresponds to the j -th asset X_j and where each i -th row represents a state of the world in which a joint outcome $(X_1 = x_{i1}, \dots, X_d = x_{id})$ occurs with probability $1/n$. If one permutes the order of the values in the j -th column, then the marginal distribution of X_j remains unchanged because all realizations are equally likely. In contrast, the dependence of X_j with the other variables X_k for $k \neq j$ is affected, since permutations yield other joint events having probability $1/n$.

To infer the joint dependence among variables X_1, \dots, X_d with known marginal distributions and for which the distribution of the sum S is also known, it is convenient to add one column to the above matrix and to introduce the following $n \times (d+1)$ matrix \mathbf{M} :

$$\mathbf{M} = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1d} & -s_1 \\ x_{21} & x_{22} & \dots & x_{2d} & -s_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nd} & -s_n \end{bmatrix}. \quad (3.8)$$

In the last column, the elements $-s_i$ are all possible realizations of the sum $-S$; i.e., $s_i := -F_S^{-1}(\frac{i-0.5}{n})$, in which F_S is the distribution function of S .

3.1 Toy example

To explain the method for constructing a joint dependence, we begin with an oversimplified example. There are $d = 3$ assets and $n = 5$ discretization steps. Therefore, X_1 , X_2 , X_3 and $-S$

all take five values with probability $1/5$, which we report in a matrix:

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 0 & -19 \\ 2 & 2 & 2 & -13 \\ 3 & 3 & 3 & -10 \\ 5 & 5 & 4 & -8 \\ 6 & 7 & 9 & -6 \end{bmatrix}. \quad (3.9)$$

The first three columns of matrix \mathbf{M} thus depict the random vector (X_1, X_2, X_3) . Its joint outcomes are displayed in five rows, each row reflecting one of the five states of the world occurring with probability $1/5$. The random vector (X_1, X_2, X_3) is not yet admissible in that the five row sums (taken over the first three columns) do not match the realizations of S , i.e., we do not yet meet the constraint that $X_1 + X_2 + X_3 - S = 0$. However, permuting elements within a column is allowed, as it does not affect marginal distributions. We thus aim at permuting elements within columns such that $X_1 + X_2 + X_3 - S = 0$ (in other words, such that the row sums in \mathbf{M} are all equal to zero).

With only five rows, it might be feasible to try all possible permutations. However, in a realistic situation one could not try all possible permutations, as there would be too many possible configurations. In our empirical application, we will employ our method using at least $n = 1,000$ discretization points (states of the world), and we therefore need an efficient way to obtain a candidate solution. To achieve such efficiency, we observe that the condition $X_1 + X_2 + X_3 - S = 0$ is equivalent to the condition that $X_1 + X_2 + X_3 - S$ has zero variance. Clearly, in order to minimize the variance of $X_1 + X_2 + X_3 - S$, it must hold that X_1 is as negatively correlated as possible with $X_2 + X_3 - S$, which means that the elements of the first column of the matrix \mathbf{M} in (3.9) should appear in opposite order (are antimonotonic) to those that correspond to $X_2 + X_3 - S$. Since permuting (rearranging) values within columns does not affect marginal distributions, we thus rearrange the values in the first column to achieve this situation. Let us illustrate this statement on the first column of the matrix displayed in (3.9), which is rearranged such that its realizations are placed in opposite order to the realizations of $X_2 + X_3 - S$. We obtain the matrix $\mathbf{M}^{(1)}$:

$$X_2 + X_3 - S = \begin{bmatrix} -18 \\ -11 \\ -3 \\ 2 \\ 10 \end{bmatrix} \quad \mathbf{M}^{(1)} = \begin{bmatrix} 6 & 1 & 0 & -19 \\ 5 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 5 & 5 & -8 \\ 1 & 7 & 9 & -6 \end{bmatrix}. \quad (3.10)$$

For the starting configuration \mathbf{M} , we have that $\text{var}(X_1 + X_2 + X_3 - S) = 126$; after rearranging its first column, however, one obtains $\mathbf{M}^{(1)}$ and $\text{var}(X_1 + X_2 + X_3 - S) = 58$ is strictly decreased. We then repeat this process for each of the subsequent columns of the matrix. We can further improve this procedure by noting that in order to yield zero row sums we actually need that $X_1 + X_2$ to be antimonotonic to $X_3 - S$ and, likewise, $X_1 + X_3$ to be antimonotonic to $X_2 - S$, and $X_2 + X_3$ to $X_1 - S$. The application of this procedure to the toy example is fully described in Appendix B and after 12 steps leads to the following output:

$$\widetilde{\mathbf{M}} = \mathbf{M}^{(12)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 3 & 3 & 0 & -6 \\ 1 & 7 & 5 & -13 \\ 6 & 1 & 3 & -10 \\ 2 & 2 & 4 & -8 \end{bmatrix}. \quad (3.11)$$

We thus obtain the ideal situation that the row sums of the rearranged matrix are all equal to zero, i.e., we have found an admissible multivariate model for the assets (X_1, X_2, X_3) that is

consistent with the distribution function of their sum S and which can now be used to compute various statistics of interest. For instance, we can find that $P(X_2 > 3|X_1 > 3) = 0.5$. In order to specifically study the dependence among the three assets, we remove the effect of the marginal distributions by applying the transformation $X_i \rightarrow F_i(X_i)$ ($i=1,2,3$). We obtain the following discrete dependence structure:

$$\tilde{\mathbf{U}} = \begin{bmatrix} 0.8 & 0.8 & 1 \\ 0.6 & 0.6 & 0.2 \\ 0.2 & 1 & 0.8 \\ 1 & 0.2 & 0.4 \\ 0.4 & 0.4 & 0.6 \end{bmatrix}. \quad (3.12)$$

In practice, due to discretization errors and the fact that the algorithm is a heuristic, the variance of the row sums might not be exactly equal to zero. Indeed, finding a perfect rearrangement is an NP-complete problem (Haus (2015)), which means that, most likely, there exists no deterministic algorithm with polynomial complexity. Thus, the above procedure can only be a heuristic. In practice, however, the procedure performs extremely well (Bernard, Bondarenko, and Vanduffel (2018)) and the remaining noise can be ignored. A full exposition of the method is relegated to Appendix C.

3.2 Discussion

In Bernard, Bondarenko, and Vanduffel (2018) it is shown that the obtained multivariate model for (X_1, X_2, \dots, X_d) exhibits maximum entropy in that it yields the “most likely”⁶ configuration given the information present in available option prices and given that no additional prior information is used. To make clear what is meant by “most likely,” let us take a step back and assume for a moment that we only have information on the marginal distributions of the assets, i.e., we only agree on the values that appear in the first d columns but not on the order in which they appear, i.e., all permutations within columns are equally plausible and there is no reason to privilege one permutation over another. Hence, randomizing the assignment of realizations to the different states leads to marginal distribution functions (reflected by the columns) that are most likely to be independent, which corresponds precisely to the maximum entropy case. Using the additional known information, namely, the marginal distribution of the sum, merely implies that the set of admissible permutations reduces to those that yield row sums that are zero. Our method implements the idea of randomizing the assignment of realizations to the different states, but now under the additional constraint provided by distribution of the index. In the context of inference of marginal RNDs from observed option prices, the principle of maximum entropy has been explored in Rubinstein (1994) and in Jackwerth and Rubinstein (1996) and has been employed in Stutzer (1996).

4 Empirical Application

In this section, we present an empirical application of our algorithm for assessing dependence using daily prices for ETF options (with approximately one month maturity) written on the

⁶In this regard, it is also worth citing Jaynes (2003, p.370), who developed the principle of maximum entropy in its modern form and who stated “*In summary, the principle of maximum entropy is not an oracle telling which predictions must be right; it is a rule for inductive reasoning that tells us which predictions are most strongly indicated by our present information.*”

nine SPDR sector funds and on the SPDR S&P 500 ETF Trust. These ETFs seek to provide investment results that, before expenses, correspond generally to the price of the S&P 500 index and of the corresponding nine economic sectors. Abbreviated names for the ETFs are provided in Table 4.1. Our sample covers the time period April 1, 2007 until June 12, 2017. We first discuss the main characteristics of the sample over the nine sectors and the aggregate market.

	Description	Ticker	Abbreviation
	Consumer Discretionary Sector SPDR Fund	XLY	cdi
	Consumer Staples Sector SPDR Fund	XLP	cst
	Energy Sector SPDR Fund	XLE	ene
	Financial Sector SPDR Fund	XLF	fin
	Health Care Sector SPDR Fund	XLV	hea
	Industrial Sector SPDR Fund	XLI	ind
	Materials Sector SPDR Fund	XLB	mat
	Technology Sector SPDR Fund	XLK	tec
	Utilities Sector SPDR Fund	XLU	uti
	SPDR S&P 500 ETF Trust	SPY	spx

Table 4.1: **S&P 500 Sectors.** This table lists the underlying assets used in this study. Option prices for these assets are obtained from OptionMetrics. The sample period is from April 1, 2007 to June 12, 2017.

In the first panel of Figure 4.1, the cumulative returns of three specific sectors and of the S&P 500 are represented. The corresponding ATM implied volatilities are displayed in the second panel of Figure 4.1. Individual sectors typically display more variable performance than the S&P 500 cumulative return, with very different behavior over time. During the 2008–2009 financial crisis, the financial sector started to decline first and then the crisis spread to the other sectors (i.e., a contagion effect was particularly observed in the last quarter of 2008 onwards), whereas during the 2015 energy crisis, primarily the energy sector was affected. The ATM implied volatilities are highly correlated and, as expected, the implied volatility of the financial sector is much larger than those for the other sectors in 2008–2010, whereas the implied volatility of the energy sector dominates in 2014–2015.

Table 4.2 provides various summary statistics for the nine sectors and the S&P 500 index. In particular, from Table 4.2, we observe that the ATM implied volatilities are consistently smaller than the standard deviations of the estimated risk-neutral densities. This bias reflects non-normality of sector’s risk-neutral densities, which manifest itself in the volatility skew (in the Black-Scholes model ATM implied volatility and σ^{RND} have to be equal). Consistent with the literature on the variance risk premium (e.g., Bollerslev, Tauchen, and Zhou (2009), Todorov (2010), Carr and Wu (2009), Bondarenko (2014)), we also observe a significant positive variance risk premium, as the standard deviations in the risk-neutral world, σ^{RND} , are consistently larger (across all sectors and the index) than the ones under the real-world measure (the second column of Table 4.2).

4.1 Implied Dependence for Selective Days

We now illustrate our approach for estimating dependence on two selective dates, September 8, 2008, and November 20, 2008. These two dates are both in the midst of the financial crisis, but the former represents a relatively calm period and the latter represents an extremely turbulent one during the financial crisis. As we will show, they display vastly different implied dependences.

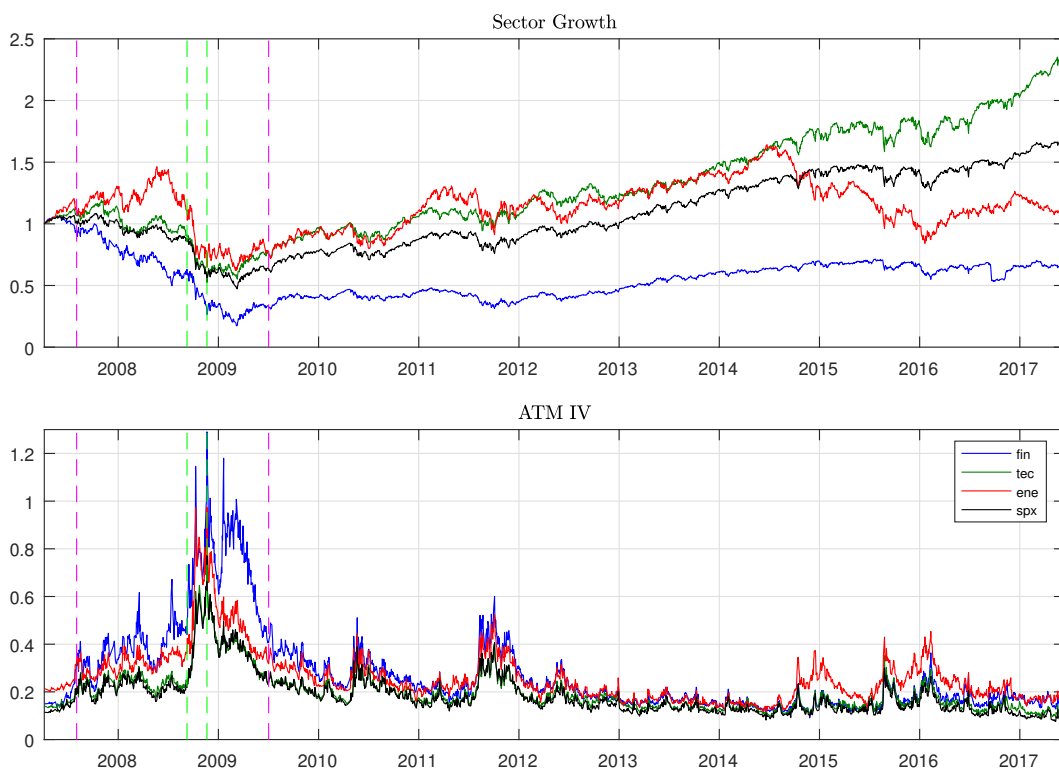


Figure 4.1: **Cumulative Returns and ATM Implied Volatilities.** The pink vertical lines indicate the period corresponding to the financial crisis (August 1, 2007 to July 1, 2008). The green vertical lines show two selective days: September 8, 2008, and November 20, 2008. Shown are three sectors (fin, tec, ene), as well as the S&P 500 index (the black line).

	Mean	StDev	Skew	Kurt	Corr	Beta	SR	Weight	IV	σ^{RND}
mat	0.056	0.216	-0.27	4.59	0.90	1.28	0.24	0.03	0.237	0.255
ene	0.033	0.220	-0.38	3.59	0.73	1.07	0.13	0.11	0.257	0.271
fin	0.010	0.247	-0.61	5.18	0.87	1.43	0.02	0.16	0.270	0.292
ind	0.082	0.192	-0.51	4.86	0.94	1.19	0.40	0.11	0.207	0.224
tec	0.102	0.171	-0.60	3.84	0.91	1.03	0.57	0.22	0.194	0.209
cst	0.081	0.116	-0.76	4.11	0.78	0.60	0.66	0.11	0.142	0.161
uti	0.040	0.140	-0.88	4.04	0.50	0.47	0.25	0.03	0.174	0.189
hea	0.090	0.140	-0.59	3.62	0.81	0.74	0.61	0.13	0.168	0.186
cdi	0.103	0.182	-0.26	4.43	0.92	1.10	0.54	0.11	0.208	0.226
spx	0.064	0.151	-0.69	4.42	1.00	1.00	0.39		0.183	0.197

Table 4.2: **Summary Statistics.** The table reports time-series averages over our sample for the nine sectors and for the aggregate market. The first six columns are computed from the monthly returns and include the mean, standard deviation, skewness, kurtosis, correlation with the market, beta with the market, and the Sharpe ratio (SR). The last three columns include the sector weight, option ATM implied volatility (IV), and standard deviation of the RND (σ^{RND}). The statistics are reported in annualized form and as decimals.

We first use options with one month maturity to obtain the risk-neutral marginal distributions F_j of the nine sectors X_j and F_S of the S&P 500 index S . We use $n = 1,000$ states for discretizing the distributions and apply the method provided in Section 3 to obtain an $n \times 10$ matrix, which

displays a joint model for the sectors X_i and the index S compatible with all marginal information.

Even though we thus fully obtain the joint distribution and thus also the dependence among these nine sectors and the index, we need to make choices in displaying it, as we cannot represent a dependence in ten dimensions. Hence, from the $n \times 10$ output matrix, we will extract triplets (x_i, y_i, z_i) for $i = 1, \dots, n$ to studying the dependence among the S&P 500 index (x_i), the financial sector (y_i), and the utilities sector (z_i). We present our findings in Figures 4.2 and 4.3:

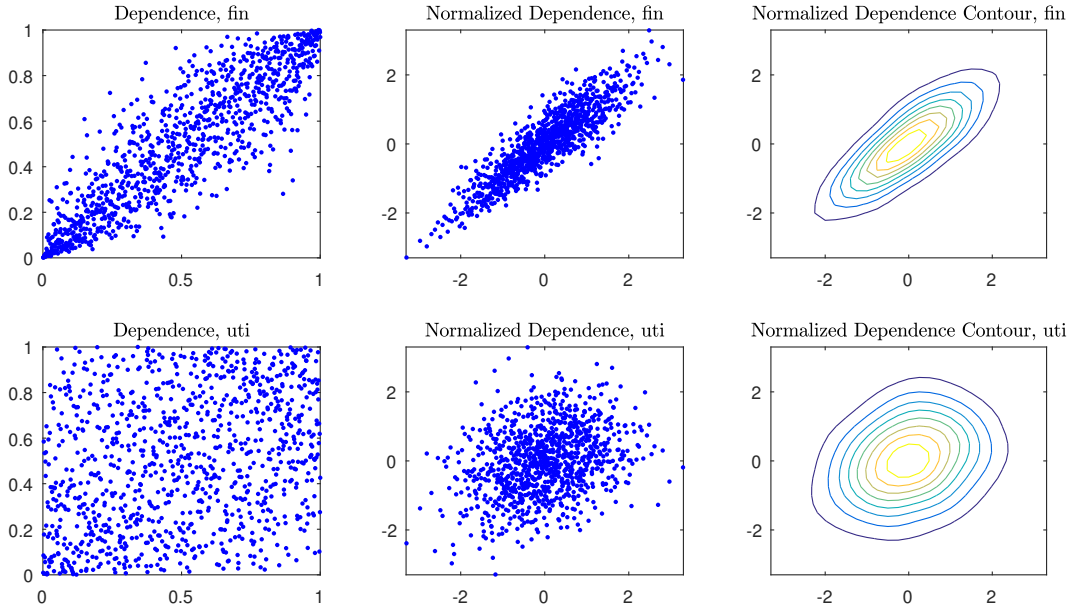


Figure 4.2: **Implied dependence on September 8, 2008.** The first column shows the dependence of the financial sector (top panels) and the utilities sector (bottom panels) relative to the S&P 500 index. The middle column shows the same dependence, but after transformation to normally distributed variables. The third column displays the corresponding contour plots.

We remove the effect of the marginal distributions on the joint distribution and display in the first column the scatterplots $(F_S(x_i), F_{fin}(y_i))$ for the top panel and $(F_S(x_i), F_{uti}(z_i))$ for the bottom panel. By doing so we bring all returns to the same (uniform) scale and obtain a visualization of true dependence. However, it is typically easier to interpret and visualize dependence between normally distributed variables instead of standard uniform distributions. Therefore, the graphs in the second column represent scatterplots of transformed variables that are now standard normal distributions by applying an additional transformation based on the quantile function Φ^{-1} of a standard normal distribution (normalized dependence). Let $G_S(x) := \Phi^{-1}(F_S(x))$ denote this transformation for the S&P 500 index, while $G_{fin}(x)$ and $G_{uti}(x)$ are defined similarly for the financial and utilities sectors, respectively. In the middle column, we then show scatterplots for the couples $(G_S(x_i), G_{fin}(y_i))$ in the top panel and for the couples $(G_S(x_i), G_{uti}(z_i))$ in the bottom panel. In the third column, we display the corresponding contour plots. These contour plots can easily be interpreted by looking at deviations from perfect ellipsoids. Indeed, when the dependence is Gaussian, these contours must be ellipsoids. From Figures 4.2 and 4.3, we find that for the financial sector the dependence on both days is strongly positive and appears to be rather symmetric, with left and right tails that are both pronounced. However, the left tail appears to be stronger than the right one. In contrast, the dependence for the utilities sector is much weaker on September 8, 2008, but becomes asymmetric with a stronger lower tail on November 20, 2008. The dependencies for both sectors are noticeably more pronounced on the second date.

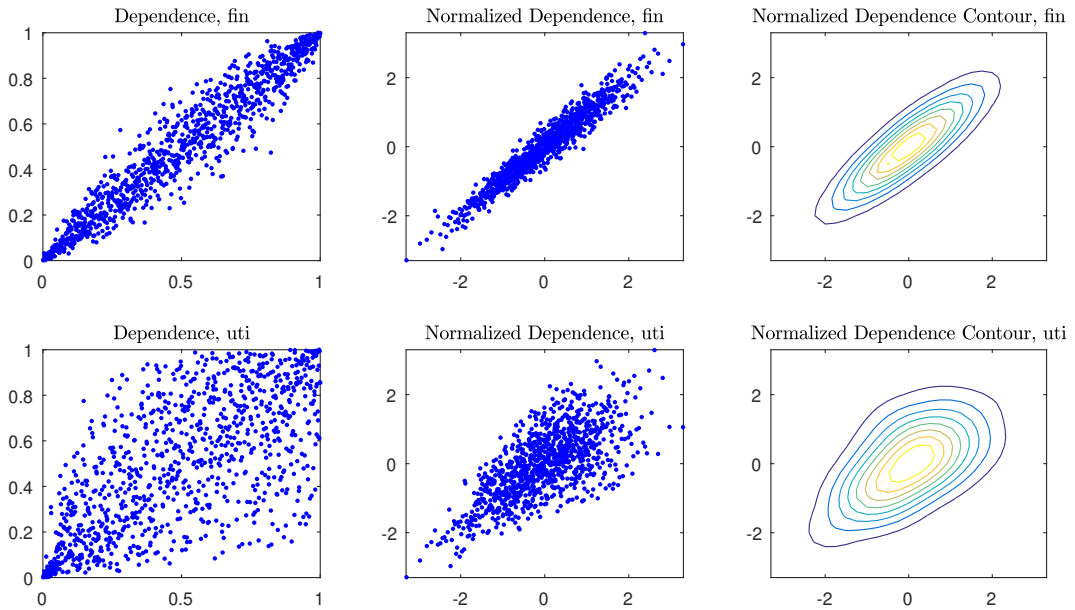


Figure 4.3: **Implied dependence on November 20, 2008.** The first column shows the dependence of the financial sector (top panels) and the utilities sector (bottom panels) relative to the S&P 500 index. The middle column shows the same dependence but for normally distributed variables. The third column displays the corresponding contour plots.

To further describe the features of the displayed dependence, we zoom in on some summary measures. From the output of the algorithm, we compute the pairwise correlations between sectors X_i and X_j . Specifically, we compute

$$\rho_{i,j}^{\mathbb{Q}} = \text{corr}^{\mathbb{Q}}(R_i, R_j), \quad (4.13)$$

in which

$$R_i := \frac{X_{i,t}}{X_{i,0}}, \quad R_j := \frac{X_{j,t}}{X_{j,0}}$$

are returns for sectors X_i and X_j , respectively, over the period $[0, t]$.

Specifically, our base estimator for correlations is the usual sample Pearson correlation coefficient that is calculated on the basis of the last sixty observed daily sector returns.⁷ While this classic estimator has highest possible efficiency under normality, it is also well known that efficiency drops significantly if the distribution changes, and its value can be misleading if outliers are present. To deal with lack of robustness, we also assessed correlations using the sample rank correlation coefficient (Spearman correlation coefficient), which is a simple non-parametric estimator that is known to present a good compromise between efficiency and robustness (Croux and Dehon (2010)). This analysis confirms that that all conclusions made in this section are robust to the estimation method.

We represent all pairwise correlation coefficients on the left panels of Figure 4.4 for September 8, 2008 and of Figure 4.5 for November 20, 2008. We observe that the financial, energy, and technology sectors are highly correlated and diversify little with the other sectors. The best diversifiers are clearly the sectors of materials and of utilities. The average global correlation is 0.665 for September 8, 2008 and 0.88 for November 20, 2008.

⁷See Jackwerth and Vilkov (2015) for a discussion of the impact of frequency on the estimation of real-world correlations

We are also able to compute two additional correlation coefficients that are critical to the understanding of what drives dependence in relation to market states. We define the down and up correlation coefficients as correlations conditional on the S&P 500 having a low or high return, respectively. Specifically,

$$\rho_{i,j}^{d,\mathbb{Q}} = \text{corr}^{\mathbb{Q}}(R_i, R_j | R_S \leq R_S^M) \quad (4.14)$$

and

$$\rho_{i,j}^{u,\mathbb{Q}} = \text{corr}^{\mathbb{Q}}(R_i, R_j | R_S > R_S^M), \quad (4.15)$$

where R_S^M denotes the median index return of the S&P 500 index.⁸

In the above definitions, if we replace X_j with the S&P 500 index S , we obtain the global correlation $\rho_{i,S}^{\mathbb{Q}}$ as well as the down and up correlations of sector i with the index, that is, $\rho_{i,S}^{d,\mathbb{Q}}$ and $\rho_{i,S}^{u,\mathbb{Q}}$. We plot the former versus the latter on the right panel of Figure 4.4 for September 8, 2008 and of Figure 4.5 for November 20, 2008. There are nine points corresponding to the nine sectors. Also shown is the first bisectrix $\rho_{i,S}^{d,\mathbb{Q}} = \rho_{i,S}^{u,\mathbb{Q}}$.

From Figures 4.4 and 4.5, it is clear that the down correlations tend to be much higher than the up correlations, i.e., $\rho_{i,S}^{d,\mathbb{Q}} > \rho_{i,S}^{u,\mathbb{Q}}$. In fact, the respective average of down correlations is 0.584 for September 8, 2008 and 0.858 for November 20, 2008. The corresponding average of the up correlations is significantly smaller: 0.411 and 0.653, respectively. On both days, the correlation conditional on the market going down is thus considerably lower than the correlation conditional on the market going up. This feature also indicates that the forward-looking dependence that we infer is asymmetric and thus non-Gaussian. In Section 5.3, we formally assess the extent by which the dependence is Gaussian or not. Finally, note that on both days all conditional correlations are smaller than the global ones. This is not surprising, as this result holds in a theoretical multivariate normal model for the returns.⁹

5 Correlation Risk Premia

In this section, we use the full sample of options on the nine sectors and the index to study risk correlation risk premia.

5.1 Average Implied Correlations Indices

Since there are many sector pairs ($\frac{1}{2}d(d-1) = 36$), it is convenient to define weighted average global, down, and up correlation indices. Specifically, for some positive weights π_i , we define:

$$\rho^{\mathbb{Q}} = \frac{\sum_{i<j} \pi_i \pi_j \rho_{i,j}^{\mathbb{Q}}}{\sum_{i<j} \pi_i \pi_j}, \quad (5.16)$$

⁸We also considered alternative definitions for the down and up correlations, where we used the mean instead of the median R^M as the return cutoff. The empirical results were very similar. However, there are important theoretical and practical advantages to defining the conditional correlations with respect to a specific quantile of index return. Specifically, using the median, we guarantee that the calculation of up and down correlations is always based on the same amount of realizations, making it possible to better control the degree of robustness of the estimates we obtain.

⁹Assuming multivariate normality, one obtains the following relation between the up (or down) correlation $\rho_{i,S}^{u,\mathbb{Q}}$ ($\rho_{i,S}^{d,\mathbb{Q}}$) and the global correlation $\rho_{i,S}^{\mathbb{Q}}$, namely that (we omit for convenience the reference to the measure \mathbb{Q} and the assets): $\rho^u (= \rho^d) = \rho \sqrt{\frac{1-\frac{2}{\pi}}{1-\frac{2\rho^2}{\pi}}}$ and thus $|\rho^u| = |\rho^d| < |\rho|$. This formula is consistent with equation (3) in Campbell, Forbes, Koedijk, and Kofman (2008).

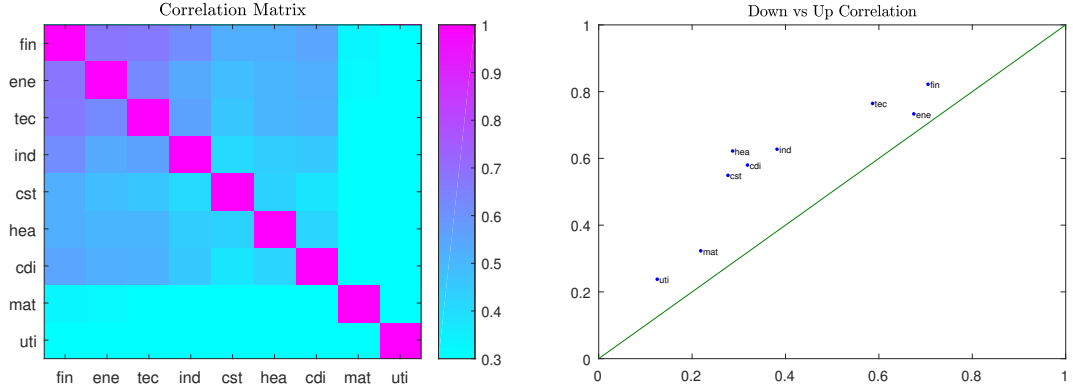


Figure 4.4: **Implied Correlations for the Nine Sectors on September 8, 2008.** The left panel shows the correlation matrix. The right panel shows the implied down correlation $\rho_{i,S}^{d,Q}$ versus up correlation $\rho_{i,S}^{u,Q}$. Also shown is the 45-degree line.

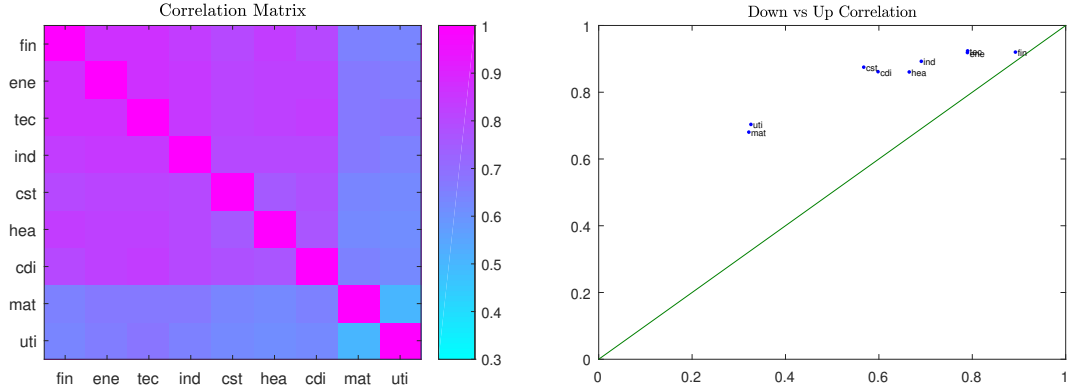


Figure 4.5: **Implied Correlations for the Nine Sectors on November 20, 2008.** The left panel shows the correlation matrix. The right panel shows the implied down correlation $\rho_{i,S}^{d,Q}$ versus up correlation $\rho_{i,S}^{u,Q}$. Also shown is the 45-degree line.

$$\rho^{d,Q} = \frac{\sum_{i<j} \pi_i \pi_j \rho_{i,j}^{d,Q}}{\sum_{i<j} \pi_i \pi_j}, \quad (5.17)$$

$$\rho^{u,Q} = \frac{\sum_{i<j} \pi_i \pi_j \rho_{i,j}^{u,Q}}{\sum_{i<j} \pi_i \pi_j}. \quad (5.18)$$

There could be several sensible choices for weights π_i , such as $\pi_i = 1/d$ (equally-weighted) or $\pi_i = \omega_i$ (value-weighted). Here, we focus on the case in which

$$\pi_i = \omega_i \sigma_i, \quad (5.19)$$

where σ_i is the model-free standard deviation estimated from the risk-neutral density. This case corresponds to “risk-weighted” averaging, and it has been used in the CBOE methodology and all other existing approaches. (Recall that the average global correlation ρ^Q with weight as in (5.19) can be computed from the marginal distributions and the distribution of the weighted sum without our BRA technique.) Thus, focusing on this case allows for direct comparison with the alternative approaches in the literature. However, the findings that we present in Section 5.2 on correlation risk premia and its state-dependent features are robust to the choice of weights.

These implied average correlations are computed for each trading day in our sample. The time series of $\rho^{\mathbb{Q}}$ is displayed in the first panel of Figure 5.6. In the second panel, we display both the down and up correlations, $\rho^{d,\mathbb{Q}}$ and $\rho^{u,\mathbb{Q}}$. The down correlations are consistently higher than the up correlations, which is consistent with the observations made in the previous section based on two specific dates.

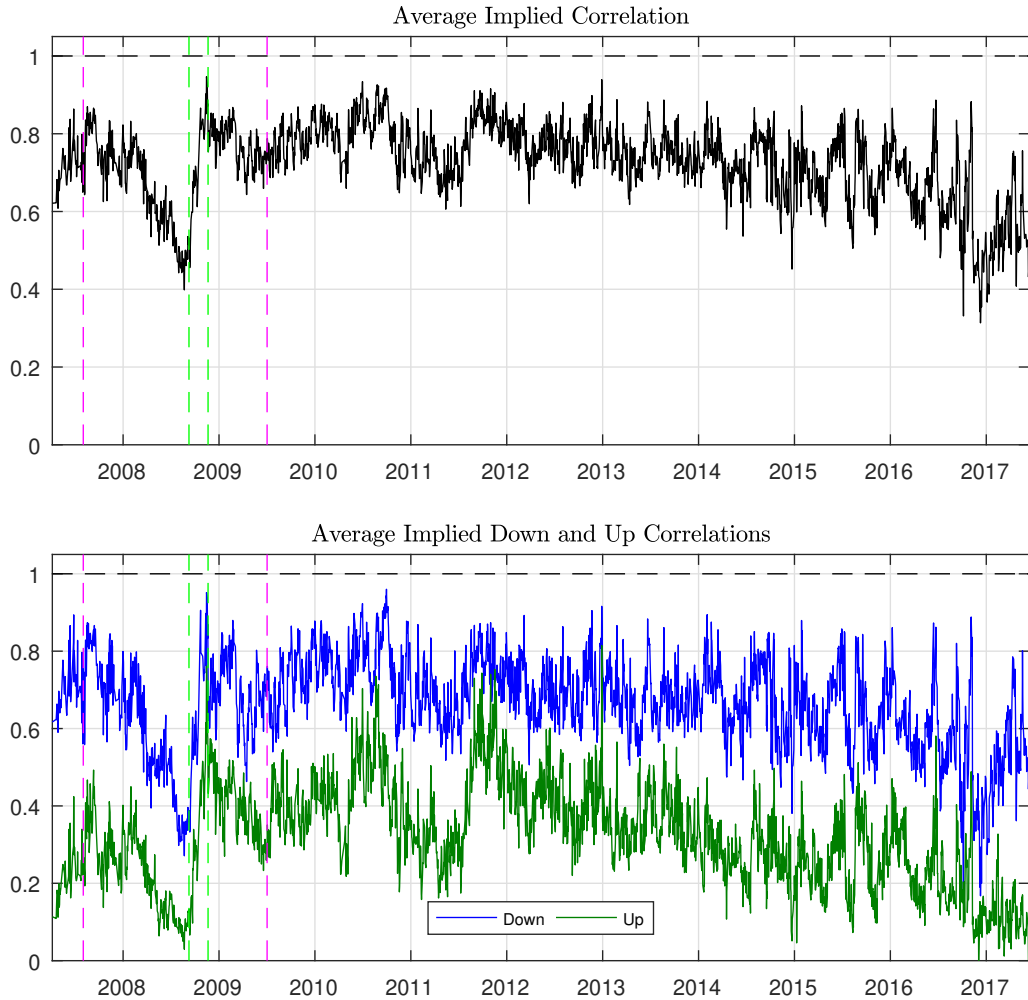


Figure 5.6: **Average Implied Correlations.** The top panel shows the average implied global correlation $\rho^{\mathbb{Q}}$. Bottom panel shows the average implied down (blue) and up (green) correlations $\rho^{d,\mathbb{Q}}$ and $\rho^{u,\mathbb{Q}}$. The pink vertical lines indicate the period corresponding to the financial crisis (August 1, 2007 to July 1, 2008). The green vertical lines show two selective days: September 8, 2008 and November 20, 2008.

We also observe from Figure 5.7 that that all correlations dropped significantly during the second and third quarter of 2008, which seems surprising, as this period fell in the midst of the financial crisis. To explain this feature and to put it in proper perspective, we first observe from Figure 5.7 that on a longer time horizon significant variations in the levels of correlations are not unusual and that the fall in correlations that we observe during 2008 is not a unique event. For example, between January 1999 and October 2000, realized correlations (red curve) dropped from 60 percent to 15 percent to reach two years later the unprecedented level of nearly 80 percent. Next, we aim at explaining the considerable drop in correlations between March 2008 and September 2008. To do so, we split the time period August 2007 - July 2009 into three subperiods: Period I covers the time period until March 2008 (the bailout of Bear Sterns),

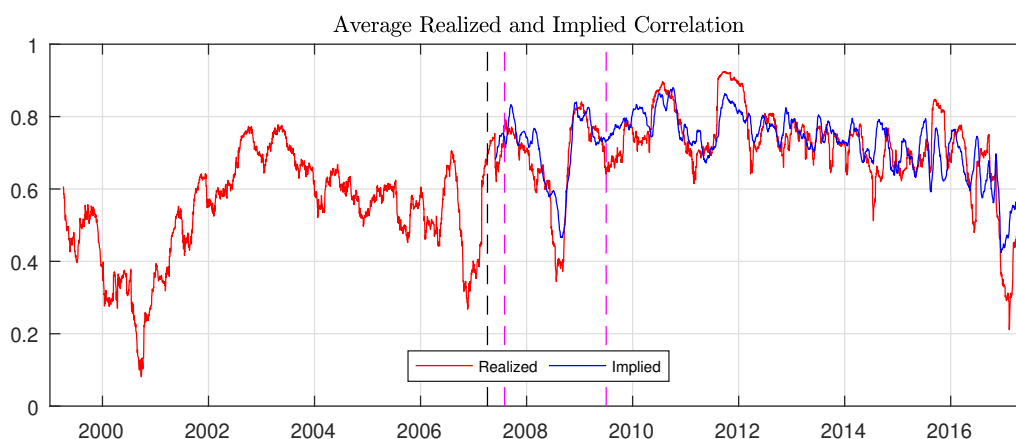


Figure 5.7: **Average Realized and Implied Correlation.** The realized correlation (red line) is shown for a longer period (starting from January 1999), while the implied correlation (blue line) is shown for a shorter period (starting from April 2007) and as a 1-month moving average.

Period II runs from March 2008 until September 2008 (the failure of Lehman Brothers), and Period III is from September 2008 until July 2009.

Back in 2006, the US financial industry was highly exposed to subprime mortgages. When house prices started to fall in July 2006, it had a direct impact on banks. Their stock prices tended to fall and their implied volatility increased on average (as confirmed by the pattern that we observe from the blue curves in the second resp. third panels of Figure 5.8 during Period I). By August 2007, correlations had also increased considerably, as confirmed by the pattern observed from the red curve in Figure 5.7. As common factors affected all financial institutions, correlations increased. However, the crisis had no spread yet to other sectors. There was no market crash and market implied volatility remained fairly stable (as can be seen in the pattern observed from the black curves in the third panel of Figure 5.8 during Period I). A plausible explanation is a widespread belief that by lowering interest rates the Federal Reserve could restore liquidity and enhance confidence.

In March 2008, Bear Stearns became the first of several financial institutions to be bailed out by the government and the market was then still confident that the government would rescue any other banks and thus that the crisis would not spill-over to other sectors. Correlations and volatilities temporarily decreased returning to pre-crisis levels (second and third panel of Figure 5.8 during Period II). But the implied volatility (and stress in the market) suddenly increased again after the bankruptcy of Fannie Mae and Freddie Mac on September 7 followed by Lehman Brothers a week after. As the government decided not to bail out Lehman, investors started to panic right after Lehman filed bankruptcy on September 15. A new crisis period started as investors seek safe investment with cash or government securities. During the last quarter of 2008, markets fell worldwide, volatilities and correlation peaked and remained at high levels afterwards (as confirmed by the patterns displayed by the curves in all three panels of Figure 5.8 during Period III).

5.2 Down and Up Correlation Risk Premia

We use the average correlation indices defined in the previous section to study down and up correlation risk premia. Formally, we define the global, down, and up correlation premium as

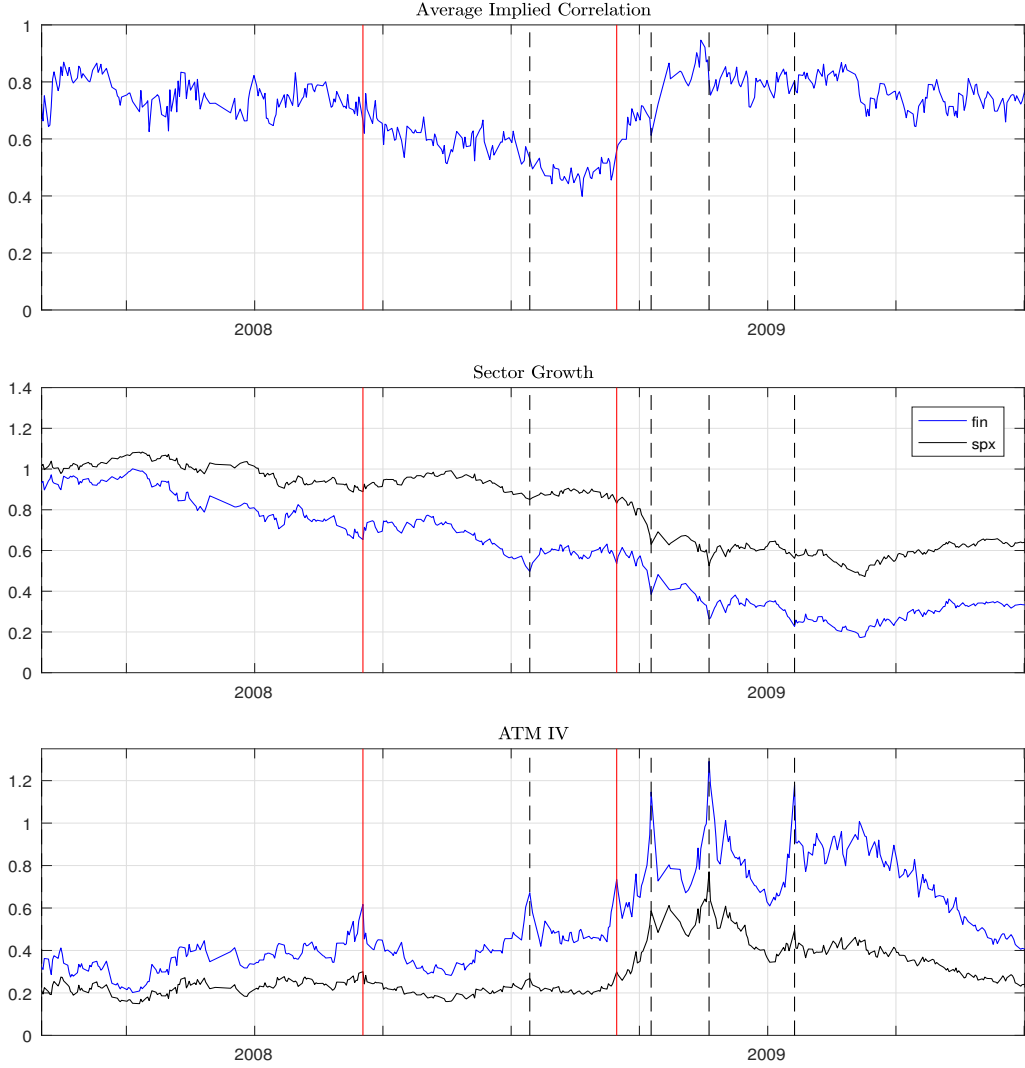


Figure 5.8: **Average Implied Correlation, Cumulative Returns, and ATM Implied Volatilities.** This figure focuses on the period of the financial crisis (August 1, 2007 to July 1, 2009). The red solid lines show the dates that separate Period I, Period II, and Period III of the financial crisis (March 17, 2008, and September 15, 2008). The black dashed lines indicate several extreme trading days (July 14, 2008, October 9, 2008, November 20, 2008, and January 20, 2009). The second and third panels show the financial sector (blue line) and the S&P 500 index (black line).

$$\theta = \rho^{\mathbb{P}} - \rho^{\mathbb{Q}}, \quad (5.20)$$

$$\theta^d = \rho^{d,\mathbb{P}} - \rho^{d,\mathbb{Q}}, \quad (5.21)$$

$$\theta^u = \rho^{u,\mathbb{P}} - \rho^{u,\mathbb{Q}}, \quad (5.22)$$

in which $\rho^{\mathbb{P}}$, $\rho^{d,\mathbb{P}}$ and $\rho^{u,\mathbb{P}}$ are the real-world counterparts to the risk-neutral correlations that we introduced in (5.16)-(5.18).

In the first panel of Figure 5.9, we display the average implied correlation $\rho^{\mathbb{Q}}$ (the blue line) and the average realized correlation $\rho^{\mathbb{P}}$ (the red line), across time. We observe that the correlation risk premium θ , which appears as the difference between the blue and the red curve, is small and mostly negative. On average, it is equal to -0.015 across our sample period (Table 5.3). This is generally with findings in the literature regarding the correlation risk premium.

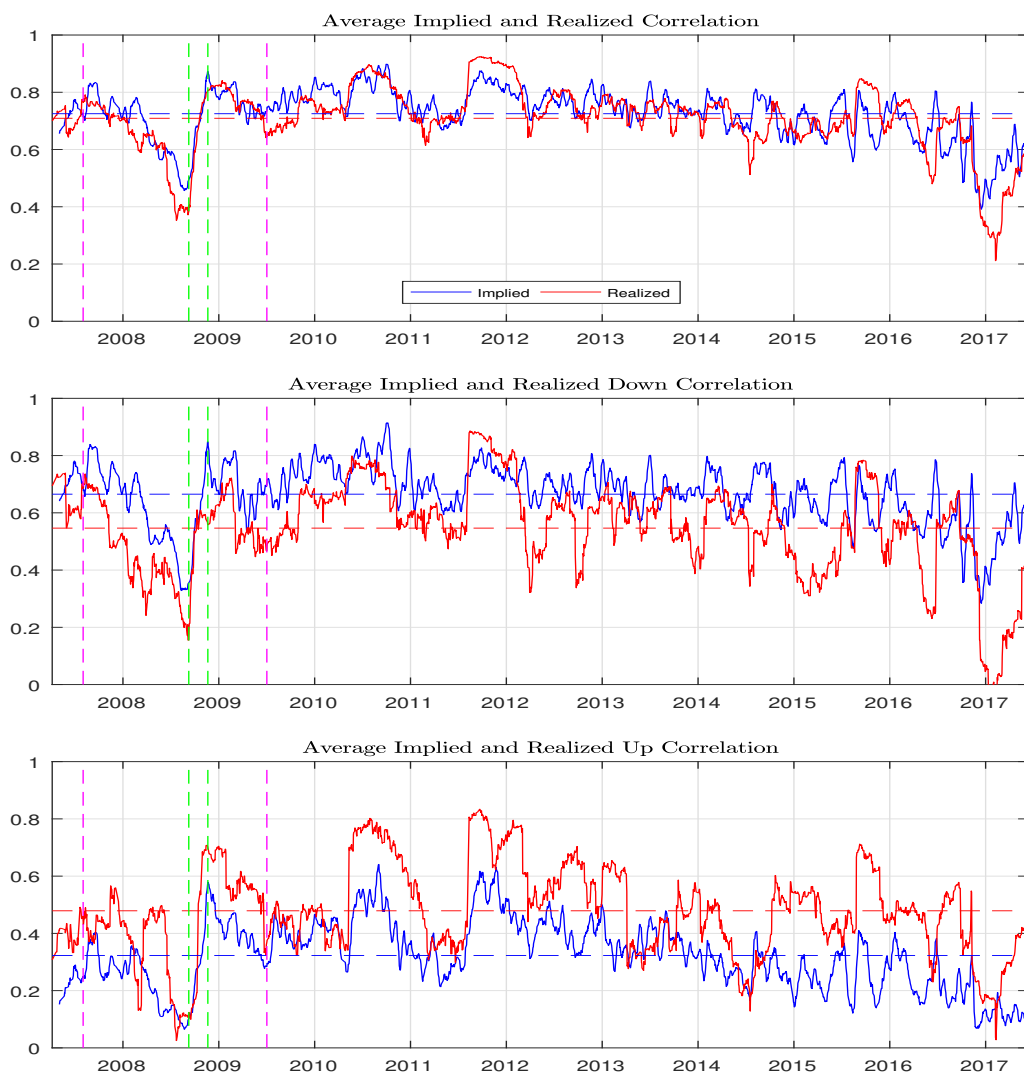


Figure 5.9: **Implied and Realized Correlations.** Implied correlations (blue) are computed from option inferred dependences; realized correlations (red) are computed from sector index returns. The corresponding means of the two series are shown with the horizontal dashed lines. The pink vertical lines indicate the period corresponding to the financial crisis (August 1, 2007 to July 1, 2008). The green vertical lines show two selective days: September 8, 2008 and November 20, 2008.

The second and third panels of Figure 5.9 and Table 5.3 provide our novel contributions to the literature on correlation premium. These panels visually demonstrate that the average realized up (down) correlations are systematically higher (lower) than their implied counterparts. Specifically, we observe from Table 5.3 that the up correlation risk premium θ^u is significant in magnitude and positive on average (0.157 in Table 5.3). In contrast, the down correlation risk premium θ^d is significant in magnitude but negative on average (-0.118 in Table 5.3). This observation allows us to conclude that the negative global correlation premium θ that we observe comes from the dependence behavior in the left tail, i.e., from down correlation coefficients, as the implied down correlations (under the risk-neutral probability measure) are much larger than the realized ones (under the real-world probability measure). This conclusion is consistent with the economic intuition that the correlation premium compensates the loss of diversification in times of crisis, i.e., when the market falls. As far as we know, this is the first time that a proposed methodology has been able to estimate conditional correlations and confirms this common belief, which is

discussed extensively in the literature on the correlation risk premium.

	N_{obs}	Under \mathbb{P}	Under \mathbb{Q}	RP	t -stat
Global	2418	0.709	0.725	-0.015	-1.9
Down	2418	0.547	0.665	-0.118	-7.8
Up	2418	0.479	0.322	0.157	10.7

Table 5.3: **Correlation Risk Premium.** The table reports statistics for the risk premia (RP) θ , θ^d and θ^u computed for the average global, down, and up correlations. The last column shows Newey-West t -statistics computed with 63 lags.

From Table 5.3, we also confirm that on average the up correlations are lower than the down correlations under both probability measures. This asymmetry is a strong indication that the dependence may not be Gaussian. This point is further explored in the next section.

5.3 Dependence or Marginal Distributions?

The down correlation coefficients that we compute are consistently larger than the up correlation coefficients (see e.g., the second panel of Figure 5.6 and Table 5.3). This feature does not comply with a multivariate normal (MVN) distribution for the asset returns, as in this case symmetry dictates that up and down correlations match with each other. Since correlation coefficients are affected jointly by the margins and by the dependence, the differences that we observe may be caused either by the deviation from normality of the margins or by the deviation from the Gaussian copula. In this section, we investigate this issue in more detail.

Our starting point is the output empirical multivariate distribution obtained using our model-free algorithm, consistent with all available option prices on the nine sectors and on the S&P 500 index. This case is denoted by EE to refer to the situation in which we use both the information on the empirical copula and on empirical margins. We denote by GG the case of Gaussian margins and Gaussian copula (i.e., a MVN distribution). Specifically, the standard deviations from the Gaussian margins comply with those of the empirical margins, and the Gaussian copula is calibrated to a constant correlation matrix in such a way that the model preserves the average pairwise correlation among the nine sectors. For both cases, we then compute the average pairwise global correlations (5.16), the down correlations (5.17), and the up correlations (5.18). We display them as a time series in the different panels of Figure 5.10: the black line corresponds to the case of EE and the blue line corresponds to GG (MVN distribution).

	GG	EG	GE	EE
Global	0.725	0.695	0.692	0.725
Down	0.483	0.515	0.594	0.665
Up	0.483	0.475	0.323	0.322

Table 5.4: **Implied Correlations.** Average implied global, down, and up correlations are computed for the four cases (GG, EG, GE, EE), where the first letter denotes the type of margins (Gaussian or Empirical) and the second letter denotes the type of the copula (Gaussian or Empirical).

By construction, the average pairwise global correlations (5.16) match for EE and GG cases (first panel). However, it is clear that the average pairwise down (up) correlation from the model-free distribution (EE) is consistently larger (smaller) over time than the one from the case GG

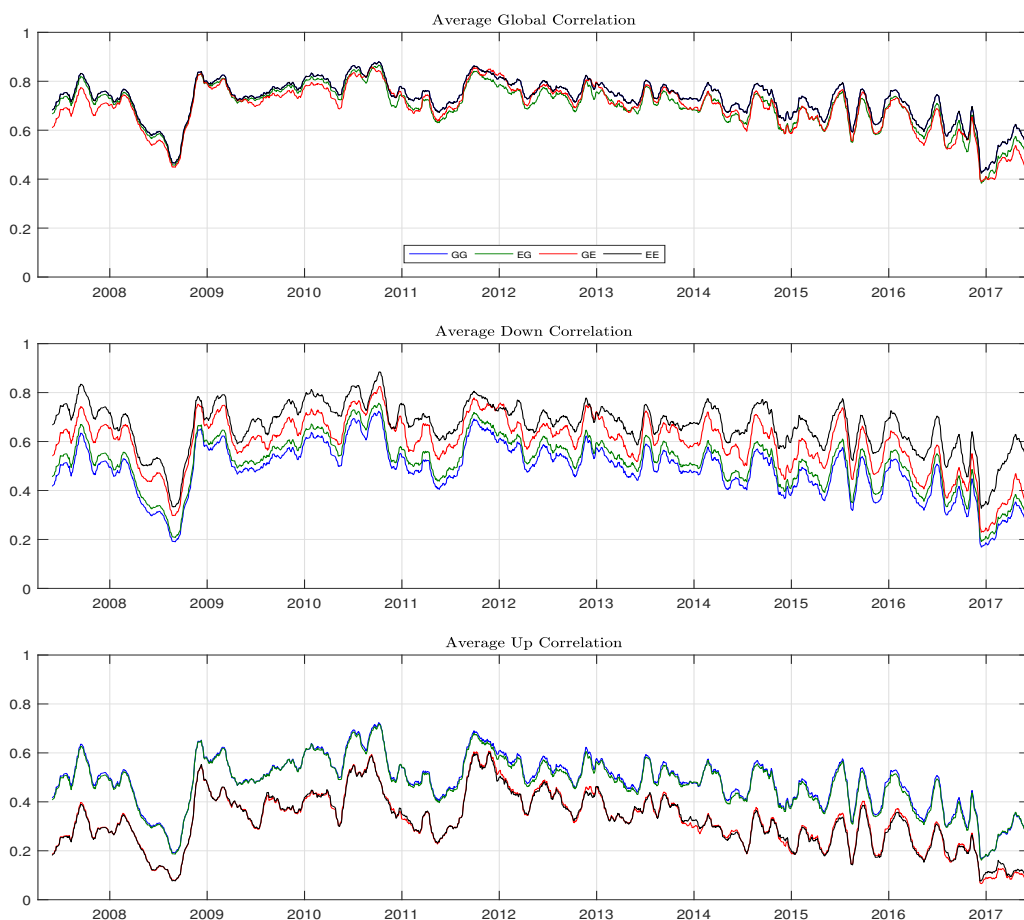


Figure 5.10: **Implied Correlations.** Average implied global, down, and up correlations are computed for the four cases (GG, EG, GE, EE), where the first letter denotes the type of margins (Gaussian or Empirical) and the second letter denotes the type of the copula (Gaussian or Empirical). Statistics are plotted as 1-month moving averages.

of the MVN fitted distribution as seen in the second (third) panel. This observation is also clear from the first and last columns of Table 5.4, in which we report the average of each quantity over the full sample.

In a next step we investigate the contributions of the marginal distributions and of the dependence to the differences that we observe between the average up and down correlations. To this end, we disentangle the information from the margin distribution and the dependence information (copula).¹⁰ We then implement two experiments. In the first, the nine sectors returns are modeled with normal margins and the dependence is the empirical copula that comes as the algorithm’s output. We denote this case by GE (misspecified margins and correct dependence). The second experiment consists in taking the algorithm’s empirical margins but using a Gaussian copula as the dependence structure. We denote this case by EG (correct margins and misspecified dependence). The first experiment is displayed with a red line and the second with a green line in all three panels of Figure 5.10. On the one hand, the green lines follow closely the blue lines, which represent up correlation and down correlation obtained using the MVN fitted distribution (in both the second and the third panels). On the other hand, the red lines are much closer to the

¹⁰Any joint distribution can be decomposed into the information on the margins and the information on the copula from Sklar (1959)’s theorem.

empirical results (black line). We thus conclude that the difference between the up correlation and down correlation coefficients is driven by the specificity of the empirical copula under the risk-neutral probability, and not by the potential asymmetry or fat tails of the margins. This conclusion is reinforced by comparing the average down and up correlations in the second and third columns of Table 5.4.

Finally, we formally test whether the option inferred dependence can be seen as Gaussian, i.e., consistent with the dependence of a multivariate normal distribution. The classical tests in this regard are Mardia’s tests of multinormality and variants thereof. For general multivariate data, Mardia (1970) constructed two statistics for measuring multivariate skewness and kurtosis, which can be used to test the hypothesis of normality (Mardia (1974), Mardia (1975) and Mardia, Kent, and Bibby (1980)). Usually, the test for whether skewness (MS) and kurtosis (MK) are consistent with a Gaussian model are performed separately; however, there also exist so-called omnibus tests that assess them simultaneously. In this paper, we perform these two tests on the normalized dependence defined as in the previous subsection. That is, we use option inferred dependence with attached normal margins and assess whether multivariate normality holds. The two tests are performed separately on each trading day in the sample. The results are shown in Table 5.5. Both tests provide strong evidence that the risk-neutral dependence among assets is not of a Gaussian nature, although the evidence for non-zero skewness is mostly pronounced. In the next section, we provide further evidence of non-Gaussian dependence in that it is shown that correlations in decreasing markets (i.e., down correlations) tend to be higher than in increasing markets (i.e., up correlations), a feature that does not comply with Gaussian dependence.

	N_{obs}	Mean	$p > 0.01$	$p > 0.05$	$p > 0.10$
MS	2230	0.0010	0.0067	0.0045	0.0027
MK	2230	0.1343	0.3946	0.3148	0.2655

Table 5.5: **Test for Gaussian Dependence.** The table reports the results of two multinormality tests, MS and MK, which were run for all trading days. Shown are mean p -values and the fraction of the days when p -value exceeds threshold of 0.01, 0.05, and 0.10 (that is, when the null hypothesis of the Gaussian dependence is not rejected).

6 Conclusions

We propose a novel methodology to estimate the risk-neutral dependence among assets in a manner that is consistent with option prices on the index and its components, without adding model assumptions. We obtain the full dependence among assets rather than merely, not just an average correlation coefficient (such as the CBOE implied correlation index). To do so, we use all the information available rather than merely relying upon ATM volatilities or variances of the components and the index.

We have described one application of our methodology, namely to deriving the correlation risk premia conditional on market conditions (e.g., down or up markets). This empirical application already demonstrates how the methodology can explore the risk-neutral dependence in a way that previous researchers have been unable to do without making specific model assumptions.

We anticipate that numerous additional applications of the methodology will emerge. Our method could provide a new way to revisit the phenomenon described in Kelly, Lustig, and Van Nieuwerburgh (2016). It could also make it possible to price a multivariate derivative on functions of the sector return in a way that is consistent with both index options and individual

options without introducing arbitrage. In addition, our method makes it possible to detect arbitrage when no dependence structure is compatible with the prices of options on the index and the prices of options on its components.

Finally, note that it is typically difficult to know whether comovements in the market are due to simultaneous changes in the volatility (changes in marginal distributions) or to changes in the dependence among assets. Our methodology, however, allows us to extract the dependence structure and disentangle the effects of marginal distribution and dependence. Moreover, it can also be useful to better measure the respective magnitudes of the variance risk and correlation risk premia.

References

- AÏT-SAHALIA, Y., AND A. W. LO (2000): “Nonparametric risk management and implied risk aversion,” *Journal of Econometrics*, 94(1), 9–51.
- BAKSHI, G., N. KAPADIA, AND D. MADAN (2003): “Stock return characteristics, skew laws, and the differential pricing of individual equity options,” *Review of Financial Studies*, 16(1), 101–143.
- BANZ, R., AND M. MILLER (1978): “Prices for State-Contingent Claims: Some Estimates and Applications,” *Journal of Business*, 51, 653–672.
- BARONE-ADESI, G., AND R. E. WHALEY (1987): “Efficient analytic approximation of American option values,” *The Journal of Finance*, 42(2), 301–320.
- BERNARD, C., O. BONDARENKO, AND S. VANDUFFEL (2018): “Rearrangement algorithm and maximum entropy,” *Annals of Operations Research*, 261(1-2), 107–134.
- BOLLERSLEV, T., R. F. ENGLE, AND J. M. WOOLDRIDGE (1988): “A capital asset pricing model with time-varying covariances,” *Journal of Political Economy*, 96(1), 116–131.
- BOLLERSLEV, T., G. TAUCHEN, AND H. ZHOU (2009): “Expected stock returns and variance risk premia,” *The Review of Financial Studies*, 22(11), 4463–4492.
- BOLLERSLEV, T., AND V. TODOROV (2011): “Tails, fears, and risk premia,” *Journal of Finance*, 66(6), 2165–2211.
- BONDARENKO, O. (2003): “Estimation of risk-neutral densities using positive convolution approximation,” *Journal of Econometrics*, 116(1), 85–112.
- (2014): “Why are put options so expensive?,” *The Quarterly Journal of Finance*, 4(03), 1450015.
- BREEDEN, D. T., AND R. H. LITZENBERGER (1978): “Prices of state-contingent claims implicit in option prices,” *Journal of Business*, 51(4), 621–651.
- BRITTEN-JONES, M., AND A. NEUBERGER (2000): “Option prices, implied price processes, and stochastic volatility,” *Journal of Finance*, 55(2), 839–866.
- BURASCHI, A., R. KOSOWSKI, AND F. TROJANI (2013): “When there is no place to hide: Correlation risk and the cross-section of hedge fund returns,” *Review of Financial Studies*, 27(2), 581–616.
- BURASCHI, A., F. TROJANI, AND A. VEDOLIN (2014): “When uncertainty blows in the orchard: Comovement and equilibrium volatility risk premia,” *Journal of Finance*, 69(1), 101–137.
- BUSS, A., L. SCHOENLEBER, AND G. VILKOV (2017): “Option-Implied Correlations, Factor Models, and Market Risk,” *INSEAD Working Paper No. 2017/20/FIN*.
- BUSS, A., AND G. VILKOV (2012): “Measuring Equity Risk with Option-implied Correlations,” *Review of Financial Studies*, 25(10), 3113–3140.
- CAMPBELL, R. A., C. S. FORBES, K. G. KOEDIJK, AND P. KOFMAN (2008): “Increasing correlations or just fat tails?,” *Journal of Empirical Finance*, 15(2), 287–309.
- CARR, P., AND L. WU (2009): “Stock options and credit default swaps: A joint framework for valuation and estimation,” *Journal of Financial Econometrics*, 8(4), 409–449.
- CHICAGO BOARD OPTIONS EXCHANGE (2009): “CBOE S&P500 Implied Correlation Index,” *White paper*.
- COCHRANE, J. H., F. A. LONGSTAFF, AND P. SANTA-CLARA (2008): “Two trees,” *Review of Financial Studies*, 21(1), 347–385.

- COOKE, R., AND C. KOUSKY (2010): “Climate Dependencies and Risk Management: Micro-correlations and Tail Dependence,” *Issue brief*, pp. 10–13.
- COSEMANS, M. M. (2011): “The pricing of long and short run variance and correlation risk in stock returns,” *Working paper of the University of Amsterdam Business School*.
- CROUX, C., AND C. DEHON (2010): “Influence functions of the Spearman and Kendall correlation measures,” *Statistical methods & applications*, 19(4), 497–515.
- DRIESSEN, J., P. J. MAENHOUT, AND G. VILKOV (2009): “The Price of Correlation Risk: Evidence from Equity Options,” *Journal of Finance*, 64(3), 1377–1406.
- DRIESSEN, J., P. J. MAENHOUT, AND G. VILKOV (2012): “Option-implied correlations and the price of correlation risk,” *Working paper, INSEAD*.
- EMBRECHTS, P., G. PUCETTI, AND L. RÜSCHENDORF (2013): “Model uncertainty and VaR aggregation,” *Journal of Banking & Finance*, 37(8), 2750–2764.
- ENGLE, R., AND S. FIGLEWSKI (2014): “Modeling the dynamics of correlations among implied volatilities,” *Review of Finance*, 19(3), 991–1018.
- FARIA, G., R. KOSOWSKI, AND T. WANG (2016): “The Correlation Risk Premium: International Evidence,” *Working Paper of Imperial College Business School*.
- HARVEY, C. R., Y. LIU, AND H. ZHU (2016): “... and the cross-section of expected returns,” *Review of Financial Studies*, 29(1), 5–68.
- HAUS, U.-U. (2015): “Bounding stochastic dependence, joint mixability of matrices, and multidimensional bottleneck assignment problems,” *Operations Research Letters*, 43(1), 74–79.
- JACKWERTH, J. C., AND M. RUBINSTEIN (1996): “Recovering probability distributions from option prices,” *Journal of Finance*, 51(5), 1611–1631.
- JACKWERTH, J. C., AND G. VILKOV (2015): “Asymmetric volatility risk: Evidence from option markets,” *SSRN working paper*.
- JAYNES, E. T. (2003): *Probability theory: the logic of science*. Cambridge university press.
- KELLY, B., AND H. JIANG (2014): “Tail risk and asset prices,” *Review of Financial Studies*, 27(10), 2841–2871.
- KELLY, B., H. LUSTIG, AND S. VAN NIEUWERBURGH (2016): “Too-Systemic-to-Fail: What Option Markets Imply about Sector-Wide Government Guarantees,” *American Economic Review*, 106(6), 1278–1319.
- MARDIA, K. (1975): “Assessment of multinormality and the robustness of Hotelling’s T² test,” *Applied Statistics*, pp. 163–171.
- MARDIA, K. V. (1970): “Measures of multivariate skewness and kurtosis with applications,” *Biometrika*, 57(3), 519–530.
- (1974): “Applications of some measures of multivariate skewness and kurtosis in testing normality and robustness studies,” *Sankhyā: The Indian Journal of Statistics, Series B*, pp. 115–128.
- MARDIA, K. V., J. T. KENT, AND J. M. BIBBY (1980): *Multivariate analysis*. Academic press.
- MUELLER, P., A. STATHOPOULOS, AND A. VEDOLIN (2017): “International correlation risk,” *Journal of Financial Economics*, 126(2), 270–299.
- POLLET, J. M., AND M. WILSON (2010): “Average correlation and stock market returns,” *Journal of Financial Economics*, 96(3), 364–380.
- PUCETTI, G., AND L. RÜSCHENDORF (2012): “Computation of sharp bounds on the distribution of a function of dependent risks,” *Journal of Computational and Applied Mathematics*, 236(7), 1833–1840.
- ROSS, S. (1976): “Options and Efficiency,” *Quarterly Journal of Economic*, 90, 75–89.
- RUBINSTEIN, M. (1994): “Implied binomial trees,” *Journal of Finance*, 49(3), 771–818.
- SCHNEIDER, P., AND F. TROJANI (2015): “Fear trading,” *Working paper of the Swiss Finance Institute, Research Paper 15-03*.
- SKINTZI, V. D., AND A.-P. N. REFENES (2005): “Implied correlation index: A new measure of diversification,” *Journal of Futures Markets*, 25(2), 171–197.
- SKLAR, M. (1959): “Fonctions de repartition à n dimensions et leurs marges,” *Publ. inst. statist. univ. Paris*, 8, 229–231.

- STUTZER, M. (1996): “A simple nonparametric approach to derivative security valuation,” *Journal of Finance*, 51(5), 1633–1652.
- TODOROV, V. (2010): “Variance risk-premium dynamics: The role of jumps,” *The Review of Financial Studies*, 23(1), 345–383.

A CBOE Implied Correlation Index

In this appendix, we discuss in greater detail the basic features and limitations of the CBOE implied correlation index and compare it with our approach.

A.1 CBOE index is not a genuine correlation

An important limitation of the CBOE implied correlation index is that its output value ρ_{cboe} may not always be interpreted as an average pairwise correlation coefficient. For instance, it has already been observed that ρ_{cboe} may take a value that is strictly larger than 1. To understand the reasons for this issue, recall from Section 2 that the CBOE correlation index builds on equation (2.3), which is only correct when $\log(S) = \sum_i \omega_i \log(X_i)$, i.e., it is in principle required that the index S is a geometric average of its components. If S denotes a constantly rebalanced portfolio of the assets X_i with weights ω_i , one gets such an equality up to a constant term. The CBOE implied correlation index is thus an average correlation in a Black-Scholes framework in which the index S denotes the portfolio obtained by constantly rebalancing a weighted portfolio of the components X_i . Clearly, this theoretical requirement does not readily hold in real markets, which prevents ρ_{cboe} from being interpreted as a genuine average pairwise correlation. Moreover, there is an additional issue of a more practical nature that hinders such interpretation. In the calculation of ρ_{cboe} , one does not consider in equation (2.4) all 500 stocks of the index when computing the weighted sums, but only a subset thereof. Specifically, the 50 biggest stocks in terms of market capitalization are considered; however, since these stocks tend to have lower volatility, the value of ρ_{cboe} tends to be overestimated. In contrast, our methodology is model-free and yields a complete dependence structure that is inherently compatible with all option prices observed.

A.2 How wrong can the CBOE index be?

A question arises concerning the impact on the computed CBOE implied correlation index when one deviates from the theoretically required assumptions. In this regard, the mere fact that the index is an arithmetic average and not a geometric one is not a major concern. Indeed, if asset prices are lognormally distributed with a Gaussian dependence that is homogeneous (i.e., for all pairwise correlations $\tilde{\rho}_{ij}$ among the assets' logreturns it holds that $\tilde{\rho}_{ij} = \rho$), then the value of ρ_{cboe} still corresponds closely to the value of the constant pairwise correlation ρ . However, under some alternative distributional assumptions regarding the asset prices, ρ_{cboe} may fall short in depicting the average correlation among the assets' logreturns.

Specifically, in a regime switching market model, ρ_{cboe} might become negative even when assets are positively correlated. To show this, define the following lognormal variables:

$$Y_1 = 100e^{r - \frac{v_1^2}{2} + v_1 W_1}, \quad Y_2 = 100e^{r - \frac{v_2^2}{2} + v_2 W_2}, \quad \text{and} \quad Z = 100e^{r - \frac{\sigma_Z^2}{2} + \sigma_Z W_Z},$$

where W_1 , W_2 , and W_Z are standard normally distributed variables in which W_Z is independent of W_1 and W_2 , and W_1 and W_2 have correlation λ_{12} (under \mathbb{Q}). Define a regime-switching model for the assets X_1 and X_2 by

$$X_1 = (1 - \mathbb{I})Y_1 + \mathbb{I}Z, \quad X_2 = (1 - \mathbb{I})Y_2 + \mathbb{I}Z,$$

in which \mathbb{I} is the variable indicating the regime. Thus, when $\mathbb{I} = 1$, the two assets coincide (extreme regime) and we assume that this occurs when Z is small enough, i.e., $\mathbb{I} = \mathbb{1}_{Z < z_q}$ in which z_q is the Value-at-Risk at level q of Z . Furthermore, define the index S as the arithmetic

average of the two assets, i.e., $S = \frac{X_1}{2} + \frac{X_2}{2}$. We then proceed as follows: First, we obtain, by Monte Carlo simulations, prices for ATM calls on the index S and on the individual stocks X_i for stocks behaving as in the above mixture model. We then estimate the corresponding implied volatilities σ_i and σ_S and compute the CBOE implied correlation index ρ_{cboe} from the implied volatilities (using equation (2.4)). In Figure A.11 we display the true correlation $\text{corr}(X_1, X_2)$ (green line), obtained by Monte Carlo simulations, ρ_{cboe} (red line) as a function of q .

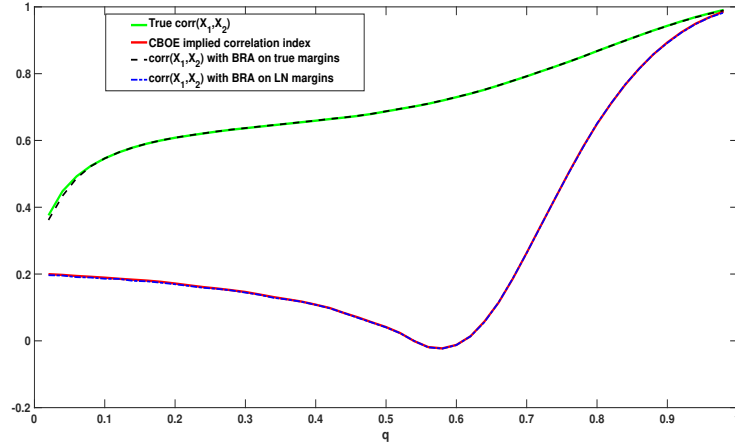


Figure A.11: This figure displays the CBOE implied correlation ρ_{cboe} , the Pearson correlation between the returns in a regime switching model, i.e., $\text{corr}(X_1, X_2)$, and the Pearson correlation after having run the BRA presented in Section 3): first when correct margins and next when margins are modelled with lognormal distributions with log-mean $(r - v_i^2/2)$ and log-variance v_i^2 . These quantities are plotted against q , which is the probability that $\mathbb{I} = 1$. The parameter values used are $v_1 = 0.15$ $v_2 = 0.25$, $\lambda_{12} = 0.2$, $r = 2\%$, $\omega_1 = \omega_2 = 1/2$ and $\sigma_Z = 0.4$.

We observe that the CBOE implied correlation index is unable to measure the true correlation between (the logs of) X_1 and X_2 (i.e., the green line does not match the red line). To better understand this feature, we have also run our algorithm (presented in Section 3) to infer the dependence, and have computed the correlation between the logreturns when we use the true marginal distributions of X_1 , X_2 , and S or when we assume that X_1 , X_2 , and the index S are all lognormally distributed (with log mean $r - \sigma_i^2/2$ and log variance σ_i^2 for $i = 1$ and $i = 2$). It is clear from Figure A.11 that the inferred correlation coefficient, which appears as the dashed black line, is now roughly equal to the true Pearson correlation (green line), whereas the blue dashed line roughly matches the CBOE. This observation shows that the CBOE implied correlation may fail to provide a good estimate of the true Pearson correlation when deviating away from a lognormal model.

Finally, we assess how important deviations from theoretical requirements are in practice. After all, the mixture model outlined above does not necessarily comply with real option markets and the question remains as to how much the CBOE correlations really differ from the ones that we compute using our approach. To ease the comparison between the CBOE index and our approach, we display the various CBOE correlation indices in the first panel of Figure A.12 (they overlap as there have been regular changes in the manner in which this index is calculated). In the second panel, we display the CBOE constant maturity (1-year) correlation index and our proposed average implied correlation index (1-month). We observe that these two indices behave in a similar manner, but that our computed correlation tends to be larger than the CBOE index over time. This phenomenon is natural, as the CBOE implied correlation measures the

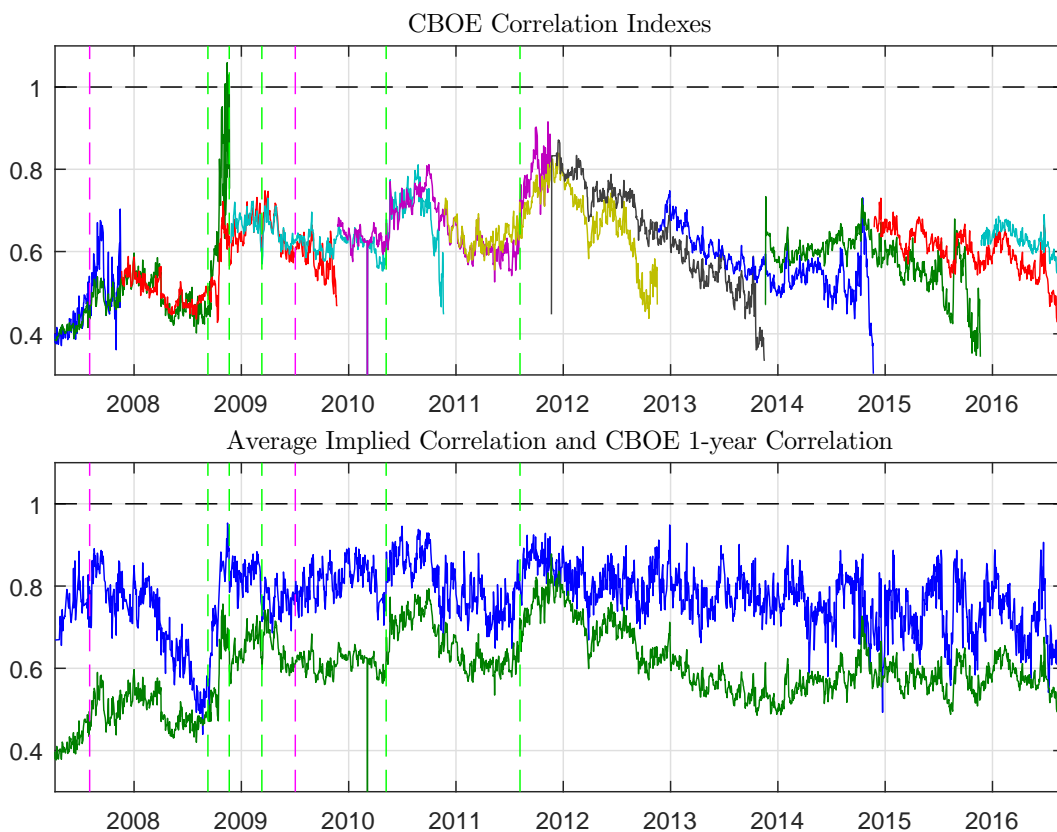


Figure A.12: **CBOE and Average Implied Correlations.** Top panel: CBOE implied correlations. Bottom panel: average Global implied correlation ρ^Q and CBOE constant maturity (1-year) correlation. The pink vertical lines indicate the period corresponding to the financial crisis (August 1st, 2007 to July 1st, 2008). The green vertical lines show two selective days: September 8, 2008 and November 20, 2008.

average pairwise correlation among 500 assets, whereas our index measures the average pairwise correlation among nine sectors. By aggregating 500 variables into nine sectors, a lot of the idiosyncratic risk disappears and the correlation increases.¹¹ All in all, we conclude that the CBOE implied correlation index provides a reasonable proxy for global average correlation in real markets. However, at times deviations can be important (particularly in stress scenarios), and our methodology is able to assess the full dependence, i.e., a joint distribution, whereas the CBOE implied correlation index merely yields a single number. This feature is further discussed in the

¹¹This effect is mostly pronounced in a fully homogeneous market in which all correlations ρ_{ij} among stock returns X_i and X_j are equal, i.e., $\rho_{ij} = \rho$. In a Gaussian framework, homogeneity is consistent with a stochastic model for asset returns X_i that are all driven in the same manner by a systematic component (in addition to their idiosyncratic component). Correlating sector indices then means that we correlate two sums, which, due to the effect of the law of large numbers, tend to become nearly linear functions of the systematic component, i.e., when the number of assets within the sectors grows, the correlation between the corresponding sector indices converges to 1. Next, we allow for the feature that for assets belonging to the same sector their pairwise correlation, say ρ_{intra} , differs from the correlation, say ρ_{betw} , that exists among assets belonging to different sectors. This correlation structure is compatible with a stochastic model in which normalized asset returns $X_{i,j}$ belonging to the j^{th} sector ($j = 1, \dots, 9$) write as $X_{i,j} = aM + bM_j + \sqrt{1 - a^2 - b^2}\epsilon_{ij}$, in which M is a systematic component, M_j is a sector specific component, and the ϵ_{ij} are idiosyncratic components (independent standard normals). In this instance, correlations among sector indices are still higher than pairwise asset correlations, but the effect will be smaller than in the case of full homogeneity. For instance, when $\rho_{intra} = 0.8$ and $\rho_{betw} = 0.6$, the average correlation among the nine sector indices is approximately equal to 0.75, whereas the average pairwise correlation among the assets is only equal to 0.62, which conforms well with the pattern observed in the second panel of Figure A.12; see Cooke and Kousky (2010) for further discussion.

following section.

A.3 CBOE index is a number - our method yields a multivariate distribution

The CBOE implied correlation index yields only a global Pearson correlation (a single number), whereas our methodology measures the entire dependence (a multivariate distribution). As a result, the CBOE implied correlation index cannot capture tail dependence (i.e., higher correlation in tails of the distribution), even though there is evidence of “local” correlation increasing in times of crisis. Correlation is then state-dependent: it may depend, for instance, on the index returns but the methodology adopted by the CBOE cannot be adjusted to find such evidence based on option prices. By contrast, our methodology is able to assess these state-dependent features and to estimate conditional correlations.

There are additional limitations of the CBOE implied correlation index. For instance, by definition, the CBOE index only uses the information from ATM option prices: any other implied volatility could have been used instead. Our approach is model-free and integrates all the option prices available.

B Toy Example of the Algorithm Used to Infer Dependence

For ease of presentation, we illustrate the algorithm with an oversimplified example. Assume that $d = 3$; we use a very small discretization step, $n = 5$, so that X_1 , X_2 , X_3 , and $-S$ all take five values with probability $1/5$, which we report in a matrix as in (3.8):

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 0 & -19 \\ 2 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 5 & 5 & 5 & -8 \\ 6 & 7 & 9 & -6 \end{bmatrix}, \quad (\text{B.23})$$

where $V := \text{var}(X_1 + X_2 + X_3 - S) = 126$. In what follows we denote $X_4 = -S$.

The objective is to find a rearrangement of the values inside each column such that the row sums of the four columns of \mathbf{M} are all equal to zero. This is equivalent to rearranging the columns of the matrix \mathbf{M} such that $\text{var}(X_1 + X_2 + X_3 + X_4) = 0$. Note that we allow for rearrangements within columns, as doing so affects the dependence among the X_j , $j = 1, 2, 3$, but not their marginal distributions. By contrast, swapping values among columns will affect the marginal distributions and is not allowed. We illustrate the method on the matrix (B.23).

Step 1: Rearranging X_1 ,

$$X_2 + X_3 + X_4 = \begin{bmatrix} -18 \\ -8 \\ -3 \\ 2 \\ 10 \end{bmatrix} \quad \mathbf{M}^{(1)} = \begin{bmatrix} 6 & 1 & 0 & -19 \\ 5 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 5 & 5 & -8 \\ 1 & 7 & 9 & -6 \end{bmatrix} \quad V = 58.$$

Step 2: Rearranging X_2 ,

$$X_1 + X_3 + X_4 = \begin{bmatrix} -13 \\ -5 \\ -3 \\ -1 \\ 4 \end{bmatrix} \quad \mathbf{M}^{(2)} = \begin{bmatrix} 6 & 7 & 0 & -19 \\ 5 & 5 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 2 & 5 & -8 \\ 1 & 1 & 9 & -6 \end{bmatrix} \quad V = 12.4.$$

Step 3: Rearranging X_3 ,

$$X_1 + X_2 + X_4 = \begin{bmatrix} -6 \\ -3 \\ -6 \\ -4 \\ -4 \end{bmatrix} \quad \mathbf{M}^{(3)} = \begin{bmatrix} 6 & 7 & 9 & -19 \\ 5 & 5 & 0 & -13 \\ 3 & 3 & 5 & -10 \\ 2 & 2 & 3 & -8 \\ 1 & 1 & 4 & -6 \end{bmatrix} \quad V = 4.$$

In this case, the order of the fourth and fifth rows for X_3 is arbitrary and the rearrangement is not unique. But both lead to a new variance equal to 4.

Step 4: Rearranging X_4 ,

$$X_1 + X_2 + X_3 = \begin{bmatrix} 22 \\ 10 \\ 11 \\ 7 \\ 6 \end{bmatrix} \quad \mathbf{M}^{(4)} = \begin{bmatrix} 6 & 7 & 9 & -19 \\ 5 & 5 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 3 & -8 \\ 1 & 1 & 4 & -6 \end{bmatrix} \quad V = 2.8.$$

Step 5: Rearranging the block $[X_1 \ X_2]$; it is possible to keep the variance unchanged or to switch the first and second row. Neither results in a change in variance. For instance, we switch the two rows to illustrate the algorithm:

$$X_1 + X_2 = \begin{bmatrix} 13 \\ 10 \\ 6 \\ 4 \\ 2 \end{bmatrix} \quad X_3 + X_4 = \begin{bmatrix} -10 \\ -10 \\ -8 \\ -5 \\ -2 \end{bmatrix} \quad \mathbf{M}^{(5)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 6 & 7 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 4 & -8 \\ 1 & 1 & 3 & -6 \end{bmatrix} \quad V = 2.8.$$

Step 6: Rearranging the block $[X_1 \ X_3]$ may not help either, as it is also antimonotonic:

$$X_1 + X_3 = \begin{bmatrix} 14 \\ 6 \\ 8 \\ 6 \\ 4 \end{bmatrix} \quad X_2 + X_4 = \begin{bmatrix} -14 \\ -3 \\ -10 \\ -6 \\ -5 \end{bmatrix} \quad \mathbf{M}^{(6)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 6 & 7 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 4 & -8 \\ 1 & 1 & 3 & -6 \end{bmatrix} \quad V = 2.8.$$

Step 7: Rearranging the block $[X_1 \ X_4]$, we need to interchange the second, fourth, and fifth rows:

$$X_1 + X_4 = \begin{bmatrix} -14 \\ -4 \\ -10 \\ -6 \\ -5 \end{bmatrix} \quad X_2 + X_3 = \begin{bmatrix} 14 \\ 7 \\ 8 \\ 6 \\ 4 \end{bmatrix} \quad \mathbf{M}^{(7)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 1 & 7 & 0 & -6 \\ 3 & 3 & 5 & -13 \\ 6 & 2 & 4 & -10 \\ 2 & 1 & 3 & -8 \end{bmatrix} \quad V = 2.$$

Step 8-11: Now we go back to Step 1 and look again at the respective columns X_1 (which is already antimonotonic), X_2 (in matrix $\mathbf{M}^{(11)}$ which decreases the variance to 1.6), X_3 (which is already antimonotonic in $\mathbf{M}^{(11)}$), and X_4 (which is already antimonotonic in $\mathbf{M}^{(11)}$). After rearranging these four columns sequentially, we obtain

$$\mathbf{M}^{(11)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 1 & 7 & 0 & -6 \\ 3 & 3 & 5 & -13 \\ 6 & 1 & 4 & -10 \\ 2 & 2 & 3 & -8 \end{bmatrix} \quad V = 1.6.$$

Step 12: Next we apply again the rearrangement on block $[X_1 \ X_2]$. We switch rows 2 and 3 and find that the variance is equal to 0. The algorithm stops. The final matrix is

$$\mathbf{M}^{(12)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 3 & 3 & 0 & -6 \\ 1 & 7 & 5 & -13 \\ 6 & 1 & 4 & -10 \\ 2 & 2 & 3 & -8 \end{bmatrix} \quad V = 0.$$

The algorithm has converged, and all row sums are equal to zero.

C Algorithm to Infer Dependence

The method we propose for inferring dependence is inspired by the so-called Rearrangement Algorithm (RA) of Puccetti and Rüschendorf (2012) and of Embrechts, Puccetti, and Rüschendorf (2013), which was originally introduced to deal with the assessment of model risk and which was adjusted by Bernard, Bondarenko, and Vanduffel (2018) to make it suitable for inferring dependence, and labeled as Block Rearrangement Algorithm (BRA).

The objective is to find a rearrangement of the first d columns such that the row sums of the $d+1$ columns of \mathbf{M} are all equal to zero. In other words, the opposite of the last column is the sum of the previous ones, i.e., $S = \sum_{j=1}^d X_j$. Denote this last column by $X_{d+1} = -S$. This procedure is equivalent to finding a rearrangement of the matrix \mathbf{M} such that $X_1 + \dots + X_{d+1}$ is identically equal to zero, and thus such that $\text{var}(X_1 + \dots + X_{d+1}) = 0$. We allow for rearrangements within columns, as doing so affects the dependence among X_j , $j = 1, \dots, d$, but not their respective marginal distributions. By contrast, swapping values among columns will affect the marginal distributions and is not allowed. Clearly, in order for $X_1 + \dots + X_{d+1}$ to have the smallest possible variance, it must hold that for all $\ell = 1, \dots, d+1$, X_ℓ is as negatively correlated as possible with $\sum_{j=1, j \neq \ell}^{d+1} X_j$ (Puccetti and Rüschendorf 2012, Theorem 2.1), i.e., is anti-monotonic. This observation lies at the core of this rearrangement method.

In fact, it must actually hold that for any decomposition of $\{1, \dots, d+1\} = I_1 \cup I_2$ into two disjoint sets I_1 and I_2 , the sums $S_1 := \sum_{k \in I_1} X_k$ and $S_2 := \sum_{k \in I_2} X_k$ are anti-monotonic and not only for sets of the form $I_1 = \{j\}$ and $I_2 = \{1, \dots, d+1\} \setminus \{j\}$. This observation makes it possible to generalize the standard RA by rearranging “blocks of columns” instead of one column at a time: The columns in the first set I_1 are stacked into a matrix (block) \mathbf{X}_1 and one rearranges its rows (i.e., one swaps entire rows) such that the row sums of \mathbf{X}_1 (reflecting S_1) are in increasing order. As for the matrix \mathbf{X}_2 that is formed by stacking the remaining columns, the rows are rearranged such that the row sums (reflecting S_2) are in decreasing order.

Algorithm for inferring dependence

1. Select a random sample of n_{sim} possible partitions of the columns $\{1, \dots, d+1\}$ into two non-empty subsets $\{I_1, I_2\}$. When $d \leq 9$, we take $n_{sim} = 2^d - 1$ so that all non-trivial partitions are considered.
2. For each of the n_{sim} partitions, create the matrices (blocks) \mathbf{X}_1 and \mathbf{X}_2 with corresponding row sums S_1 and S_2 and rearrange rows of \mathbf{X}_2 so that S_2 is anti-monotonic to S_1 .
3. If there is no improvement in $\text{var}\left(\sum_{j=1}^{d+1} X_j\right)$, output the current matrix \mathbf{M} ; otherwise, return to step 1.

At each step of this algorithm, we ensure that the variance decreases or remains the same: the columns, say X_j before rearranging and \tilde{X}_j after rearranging, verify¹²

$$\text{var}\left(\sum_{j=1}^{d+1} X_j\right) \geq \text{var}\left(\sum_{j=1}^{d+1} \tilde{X}_j\right).$$

Each time, we randomize each of the $d+1$ columns of \mathbf{M} given in (3.8), and next we apply the algorithm. For each run k of the algorithm we obtain a candidate dependence among the d first components of the index. Note that the vectors $[X_1^{(k)}, \dots, X_d^{(k)}]$ differ only with respect to their interdependence. As for each run k , we randomize the initial matrix (i.e., we randomize the initial condition for the algorithm), and it appears reasonable to assert that the procedure provides a way in which to describe the set of *all* possible dependence structures (copulas) that are consistent with the given information. Nevertheless, we obtain that application of the procedure leads to “a very thin set” of dependence structures, which display (nearly) maximum entropy (Bernard, Bondarenko, and Vanduffel (2018)).

D Data

D.1 CBOE Options

We use CBOE options on nine Select Sector SPDR ETFs. These options are American-style, and their underlying assets (SPDR ETFs) pay quarterly dividends. The option contract size is 100 shares of the corresponding ETF. The minimum price movement is 0.05. The strikes are multiples of \$1. Sector options all expire on the same day. On any given trading day, we focus on the shortest available maturity with at least 30 days remaining.

D.2 Dataset Construction

1. We use closing quotes to compute mid-point prices. In the dataset, we match all puts and calls by trading date t , maturity T , and strike K . For each pair (t, T) , we drop very low (high) strikes with zero bids. We approximate the risk-free rate r over $[t, T]$ by the rate of Treasury bills.

3. Because sector options are American style, their prices $P_t^A(K)$ and $C_t^A(K)$ could be slightly higher than the prices of the corresponding European options $P_t(K)$ and $C_t(K)$. The difference,

¹²Indeed, $\text{var}\left(\sum_{k=1}^{d+1} X_k\right) = \text{var}\left(X_j + \sum_{k \neq j} X_k\right)$, and a necessary condition for $\text{var}\left(\sum_{k=1}^{d+1} X_k\right)$ to become minimum is that each X_j is anti-monotonic with $\sum_{k \neq j} X_k$.

however, is small for the short maturities, on which we focus. This is particularly true for OTM and ATM options.

To infer the prices of European options $P_t(K)$ and $C_t(K)$ on a given underlying X_t and maturity $\tau = T - t$, we proceed as follows: First, we discard all ITM options. That is, we use put prices for $K/Z_t \leq 1.00$ and call prices for $K/Z_t \geq 1.00$, where $Z_t := X_t e^{(r-\delta)\tau}$ is the forward price. Prices of OTM and ATM options are both more reliable and less affected by the early exercise feature. Second, we correct American option prices $P_t^A(K)$ and $C_t^A(K)$ for the value of the early exercise feature by using Barone-Adesi and Whaley (1987) approximation. Third, we compute the prices of ITM options through the put-call parity relationship: $C_t(K) - P_t(K) = (Z_t - K)e^{-r\tau}$.

4. We check option prices for violations of the no-arbitrage restrictions. To preclude arbitrage opportunities, call and put prices must be monotonic and convex functions of the strike. In particular, the call pricing function $C_t(K)$ must satisfy

$$(a) \quad C_t(K) \geq (F_t - K)^+ e^{-r\tau}, \quad (b) \quad -e^{-r\tau} \leq C_t'(K) \leq 0, \quad (c) \quad C_t''(K) \geq 0.$$

When restrictions (a)-(c) are violated, we enforce them by running the so-called *Constrained Convex Regression* (CCR) introduced in Bondarenko (2003). Intuitively, CCR searches for the smallest (in the sense of least squares) perturbation of option prices that restores the no-arbitrage restrictions. The procedure can also identify possible recording errors or typos.

5. For each pair (t, T) , we estimate RND using the method of *Positive Convolution Approximation* (PCA) developed in Bondarenko (2003). The method allows one to infer the RND $f_t(x)$ and RNCD $F_t(x)$ through the relationships in (3.6) and (3.7). The method addresses directly the important limitations of option data that (a) options are only traded for a discrete set of strikes, as opposed to a continuum of strikes, (b) very low and very high strikes are unavailable, and (c) option prices are recorded with substantial measurement errors, which arise from non-synchronous trading, price discreteness, and the bid-ask bounce. The PCA method is fully nonparametric, always produces arbitrage-free estimators, and controls against overfitting while allowing for small samples.