A framework for robust measurement of implied correlation

Daniël Linders *, Wim Schoutens

KU Leuven, Leuven, Belgium

HIGHLIGHTS

- We determine correlation estimates by matching the observed index option price with its model price.
- Different basket price approximations lead to different implied correlation estimates.
- The industry standard for measuring implied correlation is a bad estimate when it is needed the most.
- We propose a new approach for determining accurately implied correlations.

ARTICLE INFO

ABSTRACT

In this paper we consider the problem of deriving correlation estimates from observed option data. An implied correlation estimate arises when we match the observed index option price with a corresponding model price. The underlying model assumes that stock prices can be described using a lognormal distribution, while a Gaussian copula describes the dependence structure. Within this multivariate stock price model, the index option price is not given in a closed form and has to be approximated. Different methods exist and each choice leads to another implied correlation estimate.

We show that the traditional approach for determining implied correlations is a member of our more general framework. It turns out that the traditional implied correlation underestimates the real correlation. This error is more pronounced when some stock volatilities are large compared to the other volatility levels. We propose a new approach to measure implied correlation which does not have this drawback. However, our numerical illustrations show that determining implied correlations with the traditional approach may be justified for strike prices which are close to the at-the-money strike price.

We also show that implied correlation estimates can be used to define an index, called the Implied Correlation Index (ICX), which reflects the market's perception about future (short-term) co-movement between stock prices. Using a volatility index together with the ICX gives an accurate description of the current level of market fear.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

The increased activity in multi-asset derivatives has resulted in an increased exposure to correlation risk for financial institutions. As a consequence, the correlation between different assets is an important input variable for determining portfolio risk measures like the VaR. In [1,2], the authors investigate the impact of a misspecification of the correlation on various risk measures. It is shown that even small errors in the correlation estimates can lead to serious errors in the value of the VaR.

Asset correlations are varying over time and a sudden change in the correlation levels may have an important impact on the risk profile of the asset portfolio; see e.g. [3] or [4]. Therefore, it is important to track these correlation levels over time, because it gives information about the level of diversification one can obtain by composing an asset portfolio. In this paper we build a framework for determining today's level of diversification possible in a basket of stocks. It is well-documented that in a market in distress, correlations are relatively high, indicating that the diversification benefit one can obtain by composing a portfolio is evaporating. In the extreme case there is no diversification possible and 'stock picking' does not make sense anymore; the portfolio return is not determined by the particular stocks composing the portfolio, but whether one is exposed to the stock market or not. The stock prices are moving in unison and the market behaves as one big asset.

Having an idea of today's diversification level is a challenging task. An estimate based on historical time series will result in a backward looking measure and can never contain information about an event which is not captured in the data set. Stock and index options are publicly traded instruments and these derivative prices contain information about the market's view about the future movements of the financial market. Backing out correlation estimates from these traded derivatives results in a *forward-looking* estimate of the future level of diversification, to which we refer as the *implied correlation*. A similar approach was followed in [5,6].

The traditional approach for determining implied correlations between stocks composing an index is based on a lognormal approximation for the distribution of the index. This methodology is widely adopted because it results in an easy and analytical expression for the implied correlation which is a combination of implied stock and index volatilities, provided that we change the weights of the index in an appropriate way. Knowledge about the implied level of correlation is specially needed in volatile times. The lognormal approximation tends to fail in these situations, which makes it a *bad estimate in times when it is needed the most*. Indeed, we show that the traditional approach results in a systematic underestimation of the true correlation parameter. The error becomes more pronounced when some of the stock volatilities become large. As a result, one has to be careful with the use of the traditional implied correlation.

We construct a framework for deriving implied correlation estimates. The implied correlation is determined by matching the observed index option price with a corresponding model price. This model assumes that the stock price dynamics can be described by a multivariate Black & Scholes model. Although this model suffers from some major drawbacks, it pays to consider this stock price model, because it is the most straightforward multivariate extension of the one-dimensional Black & Scholes model; different Brownian motions are connected by a Gaussian copula resulting in a parsimonious multivariate stock price model where each stock is described by one volatility parameter and the dependence is fully captured by the pairwise correlations. The multivariate Black & Scholes index option pricing formula can be considered as a benchmark pricing formula, similar to the one-dimensional Black & Scholes formula. In reality, implied stock volatilities will not be constant, which reflects the fact that a single volatility parameter is not sufficient to fully capture the stock dynamics. We account for this departure from the Black & Scholes setting by using the whole volatility surface. We also show that our new implied correlation can be used to construct an *Implied Correlation Index (ICX)* which is a more accurate co-movement measure than the current market standard for determining implied correlations, which is described in [7].

We show that the traditional approach for determining implied correlation estimates is a particular member of our general framework. However, within this framework, more accurate correlation measures can be investigated. For example, we consider a correlation estimate based on convex upper and lower bounds. These convex bounds have proven their efficiency in various actuarial and financial problems; see e.g. [8,9], among others. This new correlation index is more accurate and reliable than the traditional approach. In a regime where some volatilities are large, the new implied correlation still gives an accurate picture of the mean level of co-movement between the different stocks. Using implied correlation estimates for the period January 2000–October 2009, we find that the traditional approach always underestimates the real correlation levels, but the error remains small and is on average 1.5%. We conclude that the traditional approach is justified when determining at-the-money implied correlation levels. However, when deep out-of-the-money options are used, the traditional approach sometimes underestimates the level of correlation by 6%.

The paper is organized as follows. In Section 2, we introduce the financial market and the multivariate Black & Scholes model. In Section 3 the implied correlation smile and the Implied Correlation Index are introduced. After describing the most important numerical issues in Section 4, an illustration of the obtained results is given in Section 5. We show how the Implied Correlation Index behaves during the period January 2000 and October 2009. Market participants often use volatility indices as indicators for market fear. We argue that combining an estimate for the degree of co-movement with an estimate for the level of volatility results in a more accurate description of the concept market fear. Section 6 concludes the paper.

2. The financial market

We assume a financial market[1] where n different (dividend or non-dividend paying) stocks, labeled from 1 to n, are traded. The financial market is arbitrage-free and there exists a pricing measure $\mathbb{Q}$, equivalent to the physical probability measure $\mathbb{P}$, such that the current price of any pay-off at time T can be represented as the expectation of the discounted

---

[1] We use the common approach to describe the financial market via a filtered probability space $\left(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \le t \le T}, \mathbb{P}\right)$.

pay-off. Assume that the risk-neutral stock prices $X_i(t)$, $i = 1, 2, \ldots, n$, can be described by the following set of SDEs:

$$\frac{dX_i(t)}{X_i(t)} = (r - q_i)\,dt + \sigma_i dB_i(t) \quad \text{for } i = 1, 2, \ldots, n, \tag{1}$$

where $\underline{B}(t) = (B_1(t), B_2(t), \ldots, B_n(t))$, q_i is the continuously compounded dividend yield and r is the risk-free rate. Both r and q_i are assumed to be deterministic and constant over time. The process $\{\underline{B}(t) \mid t \geq 0\}$ is a standard n-dimensional Brownian motion. The correlation $\rho_{i,j}$ between the Brownian motions B_i and B_j is defined as

$$\rho_{i,j} = \text{Corr}\left[\sigma_i B_i(t), \sigma_j B_j(t+s)\right]. \tag{2}$$

The multivariate stock price model described above is called the *multivariate Black & Scholes model*. Under the risk-neutral pricing measure $\mathbb{Q}$, the time-T stock prices are lognormal distributed:

$$\ln \frac{X_i(T)}{X_i(0)} \overset{\mathbb{Q}}{=} N\left(\left(r - q_i - \frac{1}{2}\sigma_i^2\right)T, \sigma_i^2 T\right), \quad \text{for } i = 1, 2, \ldots, n, \tag{3}$$

where '$\overset{\mathbb{Q}}{=}$' denotes an 'equality in distribution under the $\mathbb{Q}$-measure'. For a detailed discussion about conditions for completeness and no-arbitrage in the multivariate Black & Scholes model, we refer to [10,11] and the references therein.

The time-0 prices of a vanilla call and put option on $X_i(T)$ with strike K and maturity T are denoted by $C_i[K, T]$ and $P_i[K, T]$, respectively. The price for an out-of-the-money vanilla option with strike price K and maturity T is denoted by $Q_i[K, T]$, where

$$Q_i[K, T] = \begin{cases} P_i[K, T], & \text{if } K < X_i(0), \\ \dfrac{C_i[K, T] + P_i[K, T]}{2}, & \text{if } K = X_i(0), \\ C_i[K, T], & \text{if } K > X_i(0). \end{cases}$$

If the risk-neutral dynamics of the stock price $X_i(T)$ can be described by the lognormal distribution (3), the option prices $C_i[K, T]$ and $P_i[K, T]$ can be determined using the well-known Black & Scholes option pricing formulae; see e.g. [12].

2.1. Index options and the multivariate Black & Scholes model

The market index is composed of a linear combination of the n underlying stocks. Denoting the price of the index at time t by $S(t)$, $0 \leq t \leq T$, we have that

$$S(t) = w_1 X_1(t) + w_2 X_2(t) + \cdots + w_n X_n(t), \tag{4}$$

where w_i, $i = 1, 2, \ldots, n$, are positive weights that are fixed up front. If the multivariate stock price dynamics can be described by (1), vanilla options can be priced using the Black & Scholes option pricing formula. Index options written on S, however, cannot be determined in a closed form and the cdf F_S of S is not given in an analytical attractive form. The most straightforward approach is to use Monte Carlo simulation to derive estimates for the index option prices. However, our search is for implied correlation estimates; hence an inversion of the index option pricing formula is required. Moreover, we want to construct a fast algorithm for determining implied correlations, which stresses the importance of a closed form (approximate) index option pricing formula which can be evaluated in an accurate and fast way.

An extensive bibliography is dedicated to the problem of finding accurate approximations for the index option prices $C[K, T]$ and $P[K, T]$; see e.g. [13–18]. An analysis of the pricing of index options in the multivariate Black & Scholes model can be found in [19]. Throughout the paper, we use the notations $\overline{C}[K, T]$ and $\overline{P}[K, T]$ for the approximate values of the index option prices $C[K, T]$ and $P[K, T]$, respectively. Any approach can be used to determine $\overline{C}[K, T]$ and $\overline{P}[K, T]$, provided that they can be determined using the risk-free rate r, the dividend yields $q_1, q_2, \ldots, q_n$ and the variance–covariance matrix Σ. Of course, the approximation should be accurate:

$$\overline{C}[K, T] \approx C[K, T], \quad \text{for all } K \geq 0,$$
$$\overline{P}[K, T] \approx P[K, T], \quad \text{for all } K \geq 0.$$

Note that we use the notations $Q[K, T]$ and $\overline{Q}[K, T]$ to denote an out-of-the-money index option price and its approximation, respectively.

For our numerical illustrations, which are presented in Section 5, we use the approach proposed in [20] to derive $\overline{C}[K, T]$ and $\overline{P}[K, T]$. More precisely, approximate index option prices are derived by combining a convex upper and lower bound using a moment matching method. Furthermore, the approximate index curves $\overline{C}$ and $\overline{P}$ can be considered as index option curves written on a synthetic market index $\overline{S}$, which serves as an approximate index for the real index S. The cdf $F_{\overline{S}}$ of $\overline{S}$ is given in a closed form. Convex approximations for sums of dependent lognormal r.v.'s proved to be successful in the earlier literature; see e.g. [9,21–24]. They are all based on the theory of comonotonicity; see e.g. [25,26].

In order to avoid unnecessary overloading of the notations, hereafter we will omit the fixed time index T when no confusion is possible. Furthermore, we assume that the stocks do not pay any dividends. Note, however, that all results remain valid for the situation where the stocks do pay dividends.

2.2. Implied stock volatilities

Given the volatility parameter σ_i, the model price $Q_i[K]$ of an (out-of-the-money) vanilla option is given in a closed form by the Black & Scholes formula. To emphasize the dependence on the volatility parameter, we sometimes write $Q_i[K; \sigma_i]$. Options on the individual stocks are traded on an option exchange and their prices can be observed. Information about the marginal volatilities $\sigma_1, \sigma_2, \ldots, \sigma_n$ is contained in these traded prices. The distance between the observed out-of-the-money option price, denoted by $\widehat{Q}_i[K]$, and the theoretical Black & Scholes price $Q_i[K; \sigma_i]$ depends solely on σ_i. An implied estimate for the volatility parameter σ_i arises when we match the model price and the market price.

Definition 1 (*Implied Volatility*). The *implied volatility* of stock i with moneyness $\frac{K}{X_i(0)}$, denoted by $\widehat{\sigma}_i\left[\frac{K}{X_i(0)}\right]$, is defined by the following equation:

$$Q_i\left[K; \widehat{\sigma}_i\left[\frac{K}{X_i(0)}\right]\right] = \widehat{Q}_i[K]. \tag{5}$$

In a multivariate Black & Scholes model, the implied volatility should be constant, i.e., $\widehat{\sigma}_i\left[\frac{K}{X_i(0)}\right] \approx \sigma_i$, for every K. However, one typically observes that implied volatilities depend on strike and maturity. The presence of the volatility surface implies that at least one of the assumptions in the model is wrong. For example, market practitioners are well-aware of the fact that modeling the dynamics of a stock by a lognormal distribution is flawed. The implied volatility can be interpreted as a number which contains the information about the factors affecting the value of the option which are not included in the Black & Scholes model.

The *at-the-money Black & Scholes implied volatility* is the implied volatility $\widehat{\sigma}_i[\pi]$ where the moneyness π is equal to 1. If no confusion is possible, we will denote this at-the-money implied volatility by $\widehat{\sigma}_i$. The calculation of $\widehat{\sigma}_i$ requires the market price $\widehat{Q}_i[X_i(0)]$. In Section 4, we describe how to determine $\widehat{\sigma}_i$ if $\widehat{Q}_i[X_i(0)]$ is unknown. For a detailed overview of implied volatility we refer to [27–29].

3. A framework for measuring implied Black & Scholes correlation

For a given strike K and maturity T, the (model) price $Q[K]$ of an index option depends on marginal information and the information about the dependence structure. In the multivariate Black & Scholes model (1), the volatilities σ_i, $i = 1, 2 \ldots, n$ unambiguously specify the marginal risk-neutral distributions, whereas the dependence structure is modeled by the pair-wise correlations $\rho_{i,j}$, $i, j = 1, 2, \ldots, n$.

Volatilities and correlations are not observable, but the prices of index and vanilla options are. As a result, the pricing of these derivatives is not a real issue. If we take the observed prices for granted, implied estimates for the volatilities and the correlations can be obtained by matching the theoretical model prices with the market prices. In Section 2.2 we described how to extract implied volatilities from vanilla option curves; in this section we show how index options can be used to determine implied estimates for the level of correlation between the stocks. Similar multivariate dependence measures can be found in [30,5,31].

3.1. Implied correlation

Assume that the dynamics of the stocks are described by the multivariate Black & Scholes model. Information about the marginal volatilities $\underline{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ is not sufficient to determine the approximation $\overline{Q}[K]$ for $Q[K]$. Indeed, $n(n-1)$ pairwise correlations $\rho_{i,j}$ have to be specified. Assume that all the correlations are equal to ρ:

$$\underline{Assumption}: \rho_{i,j} = \rho, \quad \text{for } \begin{array}{l} i \neq j \\ i = 1, 2, \ldots, n. \end{array} \tag{6}$$

The parameter ρ can be interpreted as the average correlation level. In general, we have that $\rho \in \left[\frac{-1}{n-1}, 1\right]$.

If Condition (6) holds, we have that the marginal volatilities $\underline{\sigma}$ together with the correlation ρ completely specify the approximation $\overline{Q}[K]$ for $Q[K]$. To emphasize the dependence on $\underline{\sigma}$ and ρ, we sometimes write $\overline{Q}[K; \underline{\sigma}, \rho]$. If we replace the unknown marginal volatilities $\underline{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ required to determine $\overline{Q}[K; \underline{\sigma}, \rho]$ by the implied volatilities $\widehat{\underline{\sigma}}[\pi] = (\widehat{\sigma}_1[\pi], \widehat{\sigma}_2[\pi], \ldots, \widehat{\sigma}_n[\pi])$ with equal moneyness $\pi = \frac{K}{S(0)}$, the implied correlation arises as the value for ρ such that the observed (out-of-the-money) index option price $\widehat{Q}[K]$ equals the model price $\overline{Q}[K]$.

Definition 2 (*Implied Correlation*). The *Implied correlation* of the index S with moneyness $\pi = \frac{K}{S(0)}$, denoted by $\widehat{\rho}[\pi]$, is defined by the following equation:

$$\overline{Q}[K; \widehat{\underline{\sigma}}[\pi], \widehat{\rho}[\pi]] = \widehat{Q}[K], \tag{7}$$

where $\widehat{\underline{\sigma}}[\pi] = (\widehat{\sigma}_1[\pi], \widehat{\sigma}_2[\pi], \ldots, \widehat{\sigma}_n[\pi])$ are the marginal implied volatilities with moneyness π, defined in Definition 1.

The implied correlation $\widehat{\rho}\,[\pi]$ can be determined if the model price $Q\,[S\,(0)\,\pi]$ can be approximated. The problem of determining the model prices $Q\,[S\,(0)\,\pi]$ was discussed in Section 2.1. We conclude that different choices exist for the approximation $\overline{Q}\,[S\,(0)\,\pi]$ and each choice will result in a different estimate for $\widehat{\rho}\,[\pi]$. The closer $\overline{Q}\,[S\,(0)\,\pi]$ is to $Q\,[S\,(0)\,\pi]$, the more accurate the corresponding correlation estimate $\widehat{\rho}\,[\pi]$ will be. Note also that the implied correlation is given in an implicit form. In Section 3.2 we show that for a particular choice of the approximation $\overline{Q}\,[S\,(0)\,\pi]$, the implied correlation is given in a closed form.

In a multivariate Black & Scholes model, the marginals are assumed to follow a lognormal distribution and these marginals are connected by a Gaussian copula. In case condition (6) holds, both the implied stock volatility surface and the implied correlation surface should be flat. If (6) does not hold, the implied correlation can be interpreted as a mean level of correlation.

The presence of the volatility smile for stock options indicates that the marginals are not lognormal, i.e., vanilla options are not priced by the Black & Scholes formula. In Definition 2, this departure from the Black & Scholes formula for vanilla options is taken into account by using the implied stock volatilities when deriving the (approximate) index option value $\overline{Q}\,[K]$. Numerical illustrations show that the correlation estimates $\widehat{\rho}\,[\pi]$ will change in function of π. The corresponding curve is referred to as the *correlation smile*. This smile indicates that, even if we account for the marginal volatility smiles, our model is not able to match the market prices with a single correlation parameter.

The *at-the-money Black & Scholes implied correlation* is the implied correlation $\widehat{\rho}\,[\pi]$ where we take $\pi = 1$. If no confusion is possible, the at-the-money implied correlation is denoted by $\widehat{\rho}$. In order to determine $\widehat{\rho}$, an index option with strike $S\,(0)$ is required. In Section 4, we give a detailed description about the calculation of $\widehat{\rho}$ when an index option with strike $S\,(0)$ is not traded. The implied correlation $\widehat{\rho}$ is closely related to the implied correlation index, defined by the CBOE; see e.g. [7]. In Section 3.2, we will show that the latter approach will outperform the former approach when some of the volatilities become large.

3.2. Traditional implied correlation

The implied correlation $\widehat{\rho}\,[\pi]$ defined in Definition 2 is an approximation for the correlation parameter which matches the observed index option price with the corresponding model price. This correlation estimate is reliable provided that the real index S can be replaced by an appropriate approximate index $\overline{S}$ and the corresponding approximate index option price $\overline{Q}\,[K]$ is given in a closed form and is close to the model price $Q\,[K]$. The implied correlation is, in general, only given by the implicit relation (7) and cannot be determined in an analytical way. In this subsection we consider a particular choice for the approximate index, which we will denote by $\widetilde{S}$. We show that if $\widetilde{S}$ has a lognormal distribution, the implied correlation is given in a closed form, provided that we change the weights in an appropriate way.

The index S defined in (4) is a weighted average of lognormal r.v.'s. Having knowledge about the volatilities $\sigma_1, \sigma_2, \ldots, \sigma_n$ together with the pairwise correlations $\rho_{i,j}$ does not enable us to determine the cdf F_S and the index option price $Q\,[K]$ in a closed form. However, we can determine the variance $\text{Var}[S]$ of the index S:

$$\text{Var}\,[S] = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j X_i\,(0)\,X_j\,(0)\,\mathrm{e}^{2rT}\left(\mathrm{e}^{\rho_{i,j}\sigma_i\sigma_j T} - 1\right). \tag{8}$$

Assume for the moment that the (unknown) risk-neutral dynamics of the index S can be approximated by the index $\widetilde{S}$, where

$$\frac{\mathrm{d}\widetilde{S}\,(t)}{\widetilde{S}\,(t)} = r\mathrm{d}t + \sigma_{\widetilde{S}}\mathrm{d}\widetilde{B}\,(t)\,. \tag{9}$$

Here, $\left\{\widetilde{B}\,(t) \mid t \geq 0\right\}$ is a standard Brownian motion and $\widetilde{S}\,(0) = S\,(0)$. In this model, the index $\widetilde{S}$ follows a lognormal distribution and its variance $\text{Var}\left[\widetilde{S}\right]$ is given by

$$\text{Var}\left[\widetilde{S}\right] = S\,(0)^2\,\mathrm{e}^{2rT}\left(\mathrm{e}^{\sigma_{\widetilde{S}}^2 T} - 1\right). \tag{10}$$

Index options written on $\widetilde{S}$ can be valued using the Black & Scholes option pricing formula. The price of an out-of-the-money index option on $\widetilde{S}$ with maturity T and strike K is denoted by $Q^{BLS}\left[K; \sigma_{\widetilde{S}}\right]$. For a given moneyness π, the implied volatility $\widehat{\sigma}_S\,[\pi]$ of S is determined such that the theoretical Black & Scholes option price $Q^{BLS}\,[K; \widehat{\sigma}_S\,[\pi]]$ matches the market price $\overline{Q}\,[K]$. The at-the-money implied index volatility $\widehat{\sigma}_S\,[1]$ is denoted by $\widehat{\sigma}_S$. Note that in this case, matching the model price with the market price results in a volatility, not in a correlation parameter. However, an implied correlation estimate arises when we link the implied stock volatilities $\widehat{\sigma}_1\,[\pi], \widehat{\sigma}_2\,[\pi], \ldots, \widehat{\sigma}_n\,[\pi]$ with the implied index volatility $\widehat{\sigma}_S\,[\pi]$.

We consider the multivariate Black & Scholes model (1) where (6) holds. For a given moneyness π, knowledge about the marginal implied volatilities $\widehat{\sigma}_1\,[\pi], \widehat{\sigma}_2\,[\pi], \ldots, \widehat{\sigma}_n\,[\pi]$ is not sufficient to determine $\text{Var}[S]$ using expression (8). However, if we take $\sigma_{\widetilde{S}} = \widehat{\sigma}_S\,[\pi]$, $\text{Var}\left[\widetilde{S}\right]$ follows from expression (10) and an implied correlation estimate, denoted by $\widehat{\theta}\,[\pi]$, is defined such that

$$\text{Var}\,[S] = \text{Var}\left[\widetilde{S}\right]\,. \tag{11}$$

This equation can be rewritten as

$$S(0)^2 \left(e^{\widehat{\sigma}_S^2 [\pi] T} - 1 \right) = \sum_{i=1}^n w_i^2 X_i(0)^2 \left(e^{\widehat{\sigma}_i^2 [\pi] T} - 1 \right) + \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n w_i w_j X_i(0) X_j(0) \left(e^{\widehat{\theta}[\pi] \widehat{\sigma}_i[\pi] \widehat{\sigma}_j[\pi] T} - 1 \right). \tag{12}$$

Approximating the exponential terms by a Taylor series results in

$$\widehat{\sigma}_S^2 [\pi] \approx \sum_{i=1}^n \widetilde{w}_i^2 \widehat{\sigma}_i^2 [\pi] + \sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \widetilde{w}_i \widetilde{w}_j \widehat{\theta}[\pi] \widehat{\sigma}_i[\pi] \widehat{\sigma}_j[\pi], \tag{13}$$

where

$$\widetilde{w}_i = \frac{w_i X_i(0)}{S(0)}.$$

We find the following analytical expression for the implied correlation:

$$\widehat{\theta}[\pi] \approx \frac{\widehat{\sigma}_S^2 [\pi] - \sum_{i=1}^n \widetilde{w}_i^2 \widehat{\sigma}_i^2 [\pi]}{\sum_{i=1}^n \sum_{\substack{j=1 \\ j \neq i}}^n \widetilde{w}_i \widetilde{w}_j \widehat{\sigma}_i[\pi] \widehat{\sigma}_j[\pi]}. \tag{14}$$

The correlation index $\widehat{\theta}[\pi]$ can be interpreted as the average level of correlation between the stocks, based on out-of-the-money option prices with moneyness π. The at-the-money implied correlation $\widehat{\theta}[1]$ is denoted by $\widehat{\theta}$. The correlation estimate $\widehat{\theta}$ is extensively investigated in [5] for the Dow Jones Industrial Average. In [7], an implied correlation index for the S&P500 is defined which is closely related to $\widehat{\theta}$.

Remark that for given implied index and stock volatilities, the use of formula (14) does not guarantee that the corresponding correlation estimate $\widehat{\theta}[\pi]$ is smaller than one. We discuss this issue in more detail in Section 4.

3.2.1. Fallacies of the traditional implied correlation

The at-the-money correlation coefficient $\widehat{\theta}$ defined in (12) depends on a lognormal approximation for a weighted sum of dependent lognormal r.v.'s. It is well-known that the sum S is not lognormal distributed, see e.g. [21,32], but one might hope that the approximation is reasonable and that the implied correlation $\widehat{\theta}[\pi]$ is a good approximation. In this example we illustrate that there exist situations where the lognormal approximation fails to provide an accurate description for the sum S, which results in a serious underestimation of the implied correlation; see also [33].

Consider a market with two stocks. The risk-neutral dynamics can be described by the following SDEs:

$$\begin{cases} \dfrac{dX_1(t)}{X_1(t)} = r dt + \sigma_1 dB_1(t), \\[2mm] \dfrac{dX_2(t)}{X_2(t)} = r dt + \sigma_2 dB_2(t). \end{cases} \tag{15}$$

For the initial stock prices we take $X_1(0) = X_2(0) = 100$. We assume that $r = 3\%$ and $\sigma_1 = 20\%$. The correlation coefficient ρ is equal to 80%. The maturity T is set to 1 year. The stock market index S is

$$S(t) = 0.5 X_1(t) + 0.5 X_2(t). \tag{16}$$

The only unknown parameter of the vector $(X_1(t), X_2(t))$ is the volatility σ_2, which we will fix later. Having values for all the parameters, the variance $\text{Var}[S]$ of S can be determined using expression (8) and is a function of σ_2.

Assume now that the risk-neutral dynamics of the index S can be approximated by the r.v. $\widetilde{S}$, which can be described by (9). If σ_2 is fixed, we know the exact value for $\text{Var}[S]$ and we can choose the volatility parameter $\sigma_{\widetilde{S}}$ such that $\text{Var}[S] = \text{Var}\left[\widetilde{S}\right]$. We approximate the distribution of S using the lognormal r.v. $\widetilde{S}$ such that the first two moments of this lognormal approximation match the exact first two moments; see e.g. [15].

Given $\sigma_2 > 0$, the price of an index call option with strike $K = 100$ and maturity $T = 1$ year is denoted by $C[100]$. This price can be approximated using Monte Carlo simulation and we denote this simulated value by $C^{sim}[100]$. Using a lognormal approximation $\widetilde{S}$ for the index S, this option price can be determined using the Black & Scholes option pricing formula. The Black & Scholes price is denoted by $C^{BLS}[100]$. Finally, we also include an alternative approximation $\overline{S}$ for S. Here, we use the approximation based on convex upper and lower bounds derived in [20]. This approximate index option price is denoted by $\overline{C}[100]$. The three different approximations are presented in Fig. 1 for various choices of σ_2.

Fig. 1 clearly illustrates the fallacies of the lognormal approximation. If the volatility σ_2 becomes large (compared to σ_1), the call option $C^{BLS}[100]$ diverges from the simulated option price $C^{sim}[100]$, whereas the approximate price $\overline{C}[100]$ is always an accurate approximation.

Assume for the moment that the marginal volatilities σ_1 and σ_2 are known and the price $C[100]$ of the index option can be observed, while the volatility $\sigma_{\widetilde{S}}$ and the correlation ρ are unknown and have to be estimated from the observed index
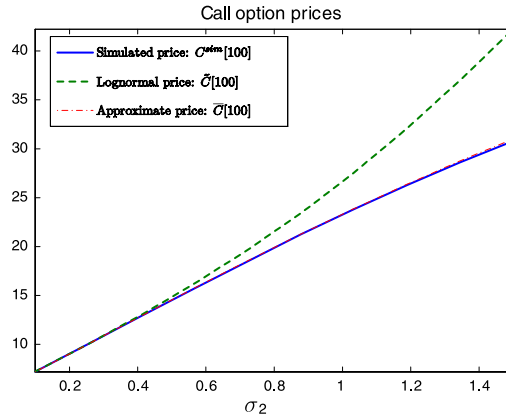
Fig. 1. Approximations for the price of an index call option with strike $K = 100$ and maturity $T = 1$ year.

option price. Fig. 1 shows that the implied index volatility $\widehat{\sigma}$ will be lower than the volatility parameter $\sigma_{\widetilde{S}}$, for which (11) holds. As a result, the implied correlation $\widehat{\theta}$ determined using Eq. (12) will also be too low. This effect is more pronounced when the volatility σ_2 is large. The accurate performance of the approximation $\overline{C}[100]$ will result in a better estimate $\widehat{\rho}$ for the correlation. This example can be generalized to the situation where the index S consists of more than two stocks. When the volatilities of some stocks are large compared to the rest of the volatilities, a similar dysfunctioning of the implied correlation is observed.

The situation where some volatilities become large occurs in a market where there is a high degree of market stress. In this example, we showed that during these periods of increased market fear, the traditional implied correlation gives misleading information. We can state that *the traditional implied correlation is a bad indicator when it is needed the most.*

4. Numerical issues

Before we give some numerical illustrations of the implied correlation for the Dow Jones Industrial Average in Section 5, we first consider some practical considerations.

Bid–ask spread. In practice, we will not observe a single index option price $\widehat{Q}[K]$ for each traded strike K. Instead, we will observe a bid price and a larger ask price. In order to cope with this bid/ask spread, we propose to use midquote prices as an approximation for the theoretical option prices. Similar conventions are made for vanilla option prices.

Implied volatilities. In reality, only a finite number of vanilla options are traded. For stock i, the traded strikes below the current stock price $X_i(0)$ are denoted by $K_{i,-l_i} < K_{i,-l_i+1} < \cdots < K_{i,0}$. The strike prices which exceed $X_i(0)$ are given by $K_{i,1} < K_{i,2} < \cdots < K_{i,h_i}$. The vanilla option prices $\widehat{Q}_i[K_j]$ for $j = -l_i, -l_{i+1}, \ldots, h_i$ can be observed in the market. We assume that $l_i > 0$ and $h_i > 1$.

The implied volatility $\widehat{\sigma}_i[\pi]$ of stock i is determined such that the corresponding Black & Scholes price with moneyness π matches the observed option price with the same moneyness. If the underlying options are of the European type, $\widehat{\sigma}_i[\pi]$ follows from an inversion of the Black & Scholes formula. We determine the implied volatility $\widehat{\sigma}_i[\pi]$, defined in Definition 1, such that the distance between the theoretical price and the observed price becomes minimal. To be more precise, $\widehat{\sigma}_i[\pi]$ follows from:

$$\widehat{\sigma}_i[\pi] = \arg\min_{\sigma \geq 0} \frac{\left| Q_i[K; \sigma] - \widehat{Q}_i[K] \right|}{\widehat{Q}_i[K]}, \tag{17}$$

where $K = \pi X_i(0)$. Note that another choice for the objective function leads to a different implied volatility estimate; see e.g. [34]. In general, K will not be a traded strike, so $\widehat{Q}_i[K]$ cannot be observed. Assume for the moment that K lies in between the traded strikes $K_{i,j}$ and $K_{i,j+1}$, so $K_{i,j} < K < K_{i,j+1}$, where $j \in \mathbb{Z}$ and $j \in [-l_i, h_i)$. The option prices $\widehat{Q}_i[K_{i,j}]$ and $\widehat{Q}_i[K_{i,j+1}]$ can be observed in the market, so $\widehat{\sigma}_i[\pi_{i,j}]$ and $\widehat{\sigma}_i[\pi_{i,j+1}]$ can be determined via (17), where we have put $\pi_{i,j} = \frac{K_{i,j}}{X_i(0)}$ and $\pi_{i,j+1} = \frac{K_{i,j+1}}{X_i(0)}$. The implied volatility with moneyness π is determined as an interpolation between $\widehat{\sigma}_i[\pi_{i,j}]$ and $\widehat{\sigma}_i[\pi_{i,j+1}]$:

$$\widehat{\sigma}_i[\pi] = \widehat{\sigma}_i[\pi_{i,j}] \frac{K_{i,j+1} - K}{K_{i,j+1} - K_{i,j}} + \widehat{\sigma}_i[\pi_{i,j+1}] \frac{K - K_{i,j}}{K_{i,j+1} - K_{i,j}}. \tag{18}$$

If $K < K_{i,-l_i}$, we put $\pi_{i,j} = \frac{K_{i,-l_i}}{X_i(0)}$ and $\pi_{i,j+1} = \frac{K_{i,-l_i+1}}{X_i(0)}$, whereas for $K > K_{i,h_i}$, we choose $\pi_{i,j} = \frac{K_{i,h_i-1}}{X_i(0)}$ and $\pi_{i,j+1} = \frac{K_{i,h_i}}{X_i(0)}$. In this situation, $\widehat{\sigma}_i[\pi]$ follows from an extrapolation between $\widehat{\sigma}_i[\pi_{i,j}]$ and $\widehat{\sigma}_i[\pi_{i,j+1}]$.

The options written on the stocks composing the Dow Jones index are of the American type. The industry standard for pricing American type options is the Cox–Ross–Rubinstein binomial tree model; see [35]. This discrete stock price model

requires, besides the risk-free rate, the strike price and the time to maturity, only the volatility as an input variable. The implied volatility $\widehat{\sigma}_i[\pi]$ follows from Eq. (17), where $Q_i[K;\sigma]$ has to be understood as the Cox–Ross–Rubinstein option price with volatility parameter σ.

The implied volatility of the index with moneyness π is denoted by $\widehat{\sigma}[\pi]$. The *at-the-money index volatility* $\widehat{\sigma}[1]$ is closely related to the Black & Scholes implied volatility index (Ticker: VXO); see e.g. [36,37]. In the sequel of the paper, we define VXO[T] as follows:

$$VXO[T] = \widehat{\sigma}[1].$$

Until here, we always assumed that vanilla and index options with maturity T are traded. In practice, the closest traded maturities are T_1 and T_2, satisfying $T_1 < T < T_2$. Options with maturity T_1 are called near term options, while next term options have a maturity equal to T_2. If the near term maturity T_1 is less than 7 days, we define T_1 to be the first maturity exceeding T while T_2 is the first maturity exceeding T_1. Using near term and next term options, VXO[T_1] and VXO[T_2] can be determined using (18). We then define VXO[T] as follows:

$$VXO[T] = \sqrt{\frac{1}{T}\left(T_1 VXO^2[T_1]\frac{T_2 - T}{T_2 - T_1} + T_2 VXO^2[T_2]\frac{T - T_1}{T_2 - T_1}\right)}. \tag{19}$$

The at-the-money Black & Scholes volatility index VXO[T] represents the perception of the market about future volatility; see e.g. [37]. VXO can be considered as the model-based counterpart of the volatility index VIX, which is discussed in [36,38].

Implied correlation. The current price level of the index is $S(0)$. Index options are traded for a finite number of strike prices. The market prices of these index options are denoted by $\widehat{C}[K]$ and $\widehat{P}[K]$. The traded strike prices below $S(0)$ are $K_{-l} < K_{-l+1} < \cdots < K_{-1} < K_0$ whereas $K_1 < K_2 < \cdots < K_h$ are the strike prices which exceed $S(0)$. We assume that $l > 0$ and $h > 1$.

For a given moneyness π, the implied correlation $\widehat{\rho}[\pi]$ is determined such that the approximate index option price coincides with the corresponding observed index option price. In this paper, the approximate index option price is determined by combining convex upper and lower bounds for the index S, see [20], but other approximations may be used, as was discussed in Section 2.1. In general, it is not possible to find an explicit formula for $\widehat{\rho}[\pi]$ such that (7) holds. Via (18), we can determine the vector $\underline{\widehat{\sigma}}[\pi]$ containing the marginal implied volatilities with moneyness π. We then search for $\widehat{\rho}[\pi]$ such that the distance between the theoretical index option price and the market price is minimized:

$$\widehat{\rho}[\pi] = \arg\min_{\rho\in(0,1)}\frac{\left|\overline{Q}\left[K;\underline{\widehat{\sigma}}[\pi],\rho\right] - \widehat{Q}[K]\right|}{\widehat{Q}[K]}, \tag{20}$$

where $K = \pi S(0)$.

In case we take $\rho > 1$, it is still possible to determine $\overline{Q}\left[K;\underline{\widehat{\sigma}}[\pi],\rho\right]$, in the sense that it will not lead to a breakdown of the calculations. So we can search for a minimum in the set of all non-negative values ρ and allow correlations bigger than one. We have witnessed that an implied correlation strictly bigger than one typically occurs when the moneyness is low. Of course, $\rho > 1$ does not make any sense. Therefore, we only consider implied correlation estimates where the moneyness is sufficiently large (e.g. $\pi > 0.75$) and we do not recommend the proposed methodology below this threshold. The observation that implied correlations bigger than one can occur, may be interpreted as the price we have to pay for working in the simple multivariate stock price model (1). In order to cope with this problem, one has to model the marginals with a more realistic distribution than the lognormal distribution. However, determining $\overline{Q}\left[K;\underline{\widehat{\sigma}}[\pi],\rho\right]$ in this new multivariate model is in general a hard problem and, to the best of our knowledge, this problem is not solved yet. Related results can be found in e.g. [39–41], where implied correlation estimates are determined from CDO prices and in [42] for implied correlation estimates based on spread options.

If strike K is not a traded strike price, $\widehat{Q}[K]$ cannot be observed in the market. Assume for the moment that K lies in between the two traded strike prices K_j and K_{j+1}, so $K_j < K < K_{j+1}$ and $j \in \mathbb{Z}$ and $j \in [-l,h)$. The index option prices $\widehat{Q}[K_j]$ and $\widehat{Q}[K_{j+1}]$ can be observed in the market and the implied correlations $\widehat{\rho}[\pi_j]$ and $\widehat{\rho}[\pi_{j+1}]$ can be determined using (20), where $\pi_j = K_j/S(0)$ and $\pi_{j+1} = K_{j+1}/S(0)$. The implied correlation with moneyness π follows from an interpolation between $\widehat{\rho}[\pi_j]$ and $\widehat{\rho}[\pi_{j+1}]$:

$$\widehat{\rho}[\pi] = \widehat{\rho}[\pi_j]\frac{K_{j+1} - K}{K_{j+1} - K_j} + \widehat{\rho}[\pi_{j+1}]\frac{K - K_j}{K_{j+1} - K_j}. \tag{21}$$

The *at-the-money implied correlation index* $\widehat{\rho}$ with maturity T is denoted by ICX[T]:

$$ICX[T] = \widehat{\rho}[1].$$

Having vanilla and index options for the near term and next term maturities T_1 and T_2, the implied at-the-money correlation indices ICX[T_1] and ICX[T_2] can be determined.[2] The implied correlation index ICX[T] can now be determined as follows:

$$ICX[T] = \frac{T_2 - T}{T_2 - T_1}ICX[T_1] + \frac{T - T_1}{T_2 - T_1}ICX[T_2]. \tag{22}$$

---

[2] Note that we silently assume that the near term and next term maturities for stock i are given by T_1 and T_2, respectively.
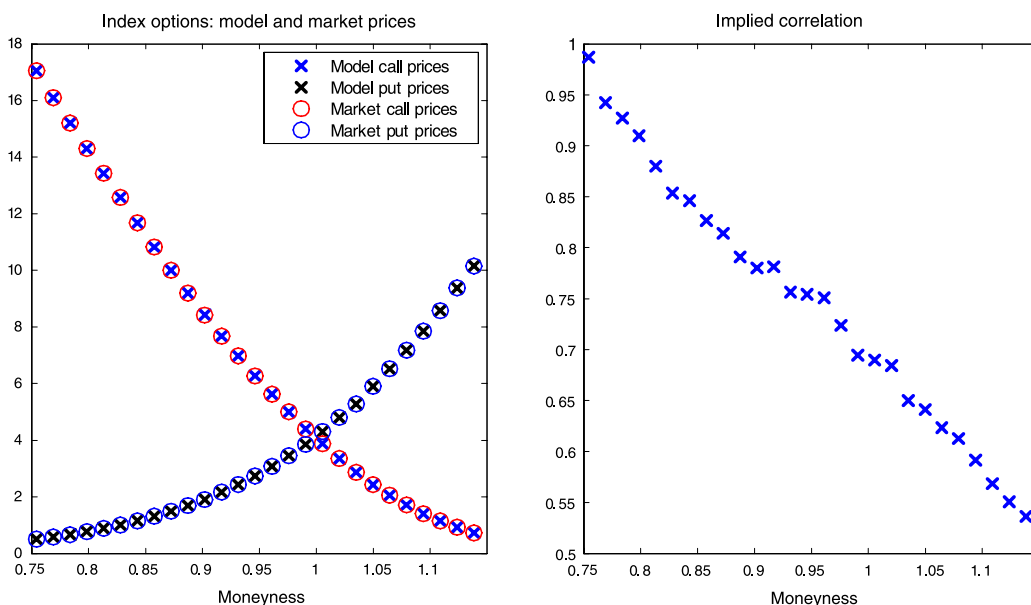
Fig. 2. March 2, 2009, Dow Jones, time to maturity 47 days. Implied correlation smile (right) $\rho\,[\pi]$ for the traded strike prices. The left panel shows the market prices $\widehat{C}\,[K]$ and $\widehat{P}\,[K]$ (circles) of index call and put options and the corresponding approximations (dots) $\overline{C}\left[K, \widehat{\rho}\left[\frac{K}{S(0)}\right]\right]$ and $\overline{P}\left[K, \widehat{\rho}\left[\frac{K}{S(0)}\right]\right]$.

Similar to the implied volatility defined by (19), the implied at-the-money correlation index with maturity T is determined as an inter- or extrapolation between the traded maturities T_1 and T_2, where we 'roll' to the next maturities if T_1 is smaller than 7 days. The implied correlation index ICX[T] can be considered as a measure for the perception of the market about the degree of co-movement in the period $[0, T]$.

Note that the main difference between our implied correlation index ICX and the CBOE implied correlation index is that we keep the time to maturity fixed at 30 days, whereas the CBOE implied correlation index fixes a maturity date (e.g. 1 year).

5. Numerical illustration

Dow Jones. The Dow Jones Industrial Average, established 1896, is a price-weighted index composed of the 30 largest, most liquid NYSE and NASDAQ listed stocks. Options with the DJ index as underlying are called DJX options. DJX options are based on 1/100th of the current value of the DJ. Therefore, hereafter $S(t)$ has to be interpreted as 1/100th of the value of the DJ at time t.

Implied correlation smile. For a given maturity T, the curve $\widehat{\rho}\,[\pi]$ is called the implied correlation smile. We find that the implied correlation *decreases* in function of the moneyness: low strike prices require a higher correlation parameter than high strike prices. The implied correlation for low strike prices is determined using out-of-the-money index put options, whereas out-of-the-money index call options are used to determine the implied correlation for high strike prices. Out-of-the money index put options are a bet on a decline of the market, whereas an out-of-the-money call option will end in the money when the market rises. With this in mind, the decreasing implied correlation curve may imply that the market expects stocks to go down simultaneously, while they are expected to go up more independently.

The right panel of Fig. 2 shows the correlation smile for the Dow Jones on March 2, 2009, where the time to maturity is equal to 47 days. For each traded strike $K_j, j = -l, -l+1, \ldots, h$, the implied correlation $\widehat{\rho}\,[\pi_j]$ with moneyness $\pi_j = \frac{K_j}{S(0)}$ is shown in the graph. The left panel shows the index option prices $\widehat{C}\,[K_j]$ and $\widehat{P}\,[K_j]$ (circles) together with the model prices $\overline{C}\left[K_j; \widehat{\sigma}\,[\pi_j], \widehat{\rho}\,[\pi_j]\right]$ and $\overline{P}\left[K_j; \widehat{\sigma}\,[\pi_j], \widehat{\rho}\,[\pi_j]\right]$ (crosses). We find that for the majority of the trading days, the implied correlation is a decreasing function of π. (See also Fig. 3.)

Implied correlation over time. We choose T to be 30 days. The implied correlation index ICX[T] is a number which represents the level of diversification possible by investing in the stocks $X_1, X_2, \ldots, X_n$. A high value for ICX[T] indicates that market participants expect that stocks will move more strongly together in the near future. Indeed, ICX[T] is based on short term traded index and vanilla options, which contain the view of the market about the near future. We say that the ICX is a forward-looking index for the degree of co-movement between the stock prices composing the index. Having daily (or even more frequent) quotes for the ICX is important for shareholders, investors and regulators, because one can anticipate highly correlated markets.

A graph of historical values of the ICX is shown in Fig. 6. The third graph is a smoothed version of the second graph, where at each day we have taken the average over the last 7 consecutive days. The average implied correlation level over the whole
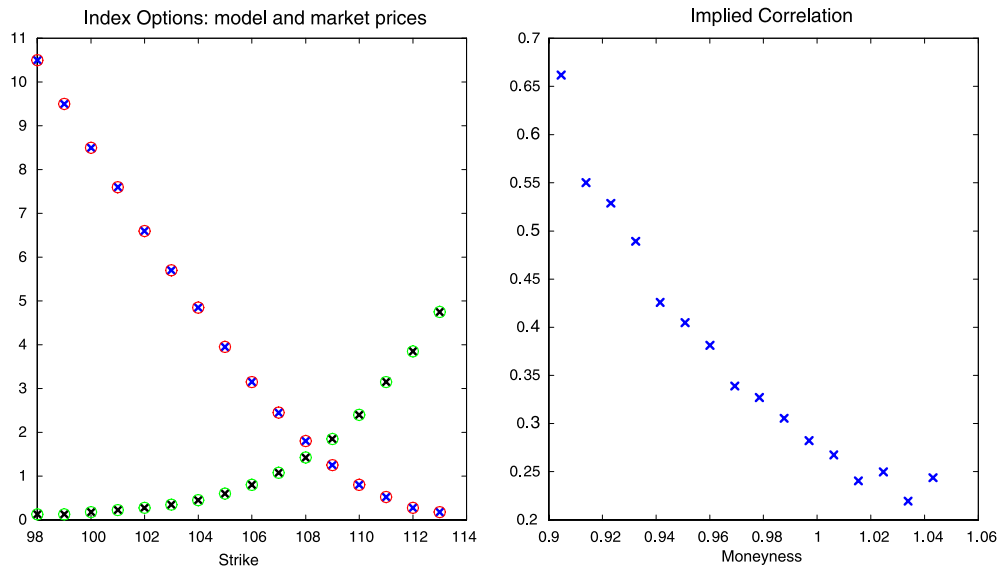
Fig. 3. March 5, 2005, Dow Jones, time to maturity 44 days. Implied correlation smile (right) $\rho\,[\pi]$ for the traded strike prices. The left panel shows the market prices $\widehat{C}\,[K]$ and $\widehat{P}\,[K]$ (circles) of index call and put options and the corresponding approximations (dots) $\overline{C}\left[K, \widehat{\rho}\left[\frac{K}{S(0)}\right]\right]$ and $\overline{P}\left[K, \widehat{\rho}\left[\frac{K}{S(0)}\right]\right]$.

period is 36%. Begin 2000, implied levels of correlation for the Dow Jones are around 20%. The degree of co-movement starts increasing during the burst of the Dot-Com bubble. On September 17, 2001, the ICX peaked around 51%. On August 2, 2002, the ICX attained a new maximum of 70% and the highest value of 82% was recorded on February 22, 2003. From 2007 on, correlations are again increasing, with a peak of 85% in October, 2008, after the Lehmann Brothers collapse. During 2009, the average level of correlation remains at a historical high level (51% on average). From Fig. 6 we find that the implied correlation is behaving in a stochastic way. In order to have an idea of the future direction of the correlation, one has to calibrate an appropriate stochastic correlation process to this data set. A possible approach for modeling correlation, and more general co-movement, measures in proposed in [43].

During periods of strong co-movement, markets are in most of the situations going down. The peaks in Fig. 6 can be related in time with major events like September 11, 2001, the default of Lehmann Brothers,. . .. In these 'headline driven' markets, there is a high degree of uncertainty about the future and stocks tend to move more strongly together than during periods where there are no major news items dominating the trading floor.

The ICX is a model-based index for the degree of co-movement, or herd behavior. We already showed that a price has to be paid for working in the simple multivariate Black & Scholes model (1), in the sense that implied correlations may be bigger than one. Alternatively, one can choose to use a model-free index for the degree of herd behavior; see e.g. [44–48]. However, in order to determine these co-movement measures a lot of data is needed. Indeed, one needs a vanilla option curve for each stock and an index option curve in order to determine a model-free co-movement estimate. An implied correlation estimate, on the other hand, only requires one vanilla option per stock and a single index option.

Traditional implied correlation index. The only quoted index which reflects the degree of co-movement between stock prices is the CBOE S&P 500 implied correlation index; see e.g. [7]. This index is the industry standard and is widely used to set up so-called dispersion trades, which are strategies designed to trade co-movement. The methodology for calculating this implied correlation index can easily be adapted to other indices, as long as there are traded options on the index and the underlyings. The methodology for the CBOE implied correlation is based on expression (14). In Section 3 we introduced a framework for determining implied correlation estimates. A particular member of this framework is the traditional implied correlation $\widehat{\theta}\,[\pi]$, which is given by the easy expression (14). By using a more sophisticated approximation for the index option prices, a more reliable and accurate implied correlation estimate $\widehat{\rho}\,[\pi]$ can be determined; see Section 3.2.1.

Fig. 4 shows the difference between the at-the-money implied correlation estimates $\widehat{\rho}\,[1]$ and $\widehat{\theta}\,[1]$, for a maturity $T = 30$ days. We see that this difference is always positive, which shows that the traditional implied correlation systematically underestimates the correlation levels. However, the average difference is 1.5%, which shows that for determining at-the-money implied correlation levels, the traditional approach may be justified.

In the following example we show that the difference $\widehat{\rho}\,[\pi] - \widehat{\theta}\,[\pi]$ can become large when the moneyness π is small. Fig. 5 shows the estimates $\widehat{\rho}\,[\pi]$ and $\widehat{\theta}\,[\pi]$ for the DJ implied correlation on July 21, 2008, and October 20, 2008. In both cases, the time to maturity is 89 days. In Section 3.2.1, we showed that the traditional approach underestimates the real level of correlation. Furthermore, this error increases when some volatilities are large compared to the other volatilities. The difference $\widehat{\rho}\,[\pi] - \widehat{\theta}\,[\pi]$ is depicted in the right panels of Fig. 5. We find that the difference is positive and increases when the strike becomes smaller. On October 20, 2008, the error was bigger compared to July 21, 2008. The difference
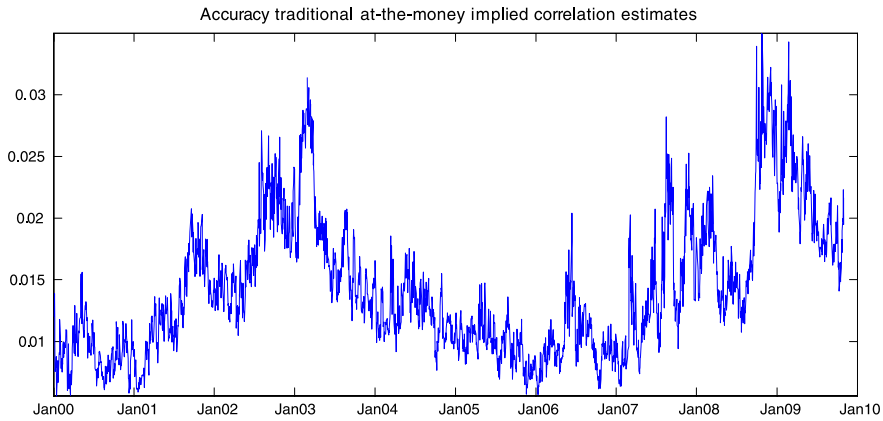
Accuracy traditional at-the-money implied correlation estimates

Fig. 4. We test the accuracy of the traditional implied correlation estimate $\widehat{\theta}[1]$, by showing the difference $\widehat{\rho}[1] - \widehat{\theta}[1]$ for the period January 2000–October 2009.
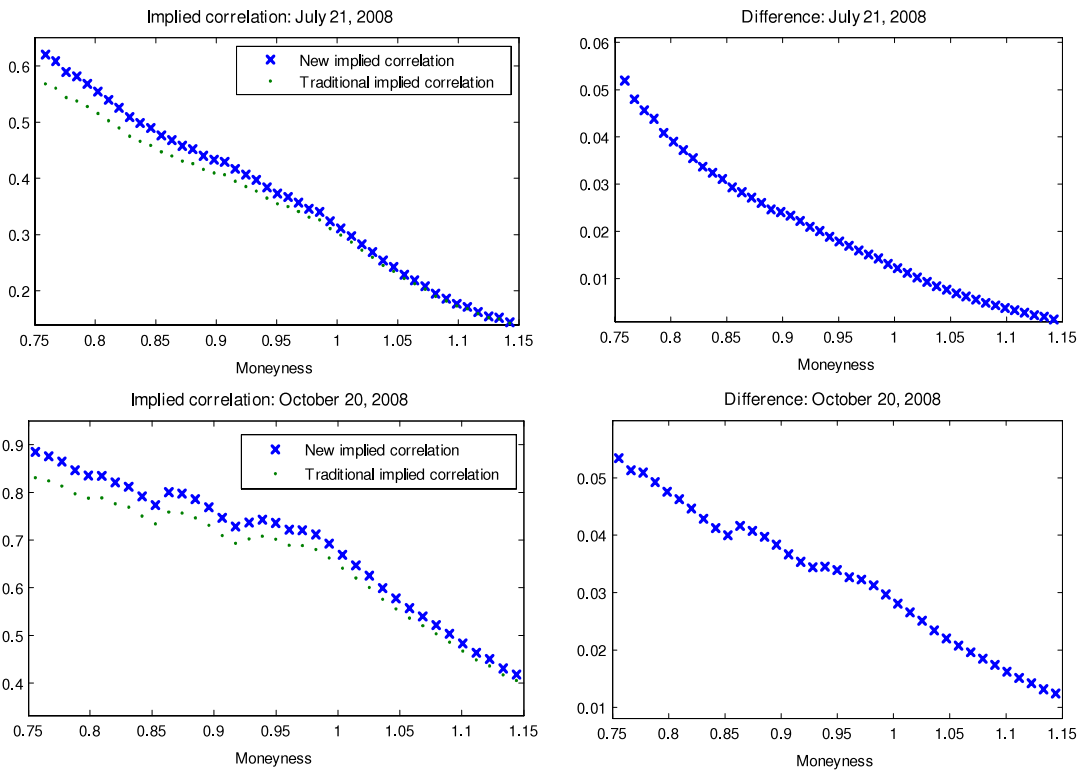

Fig. 5. The Dow Jones implied correlation $\widehat{\rho}[\pi]$ and the traditional implied correlation $\widehat{\theta}[\pi]$ for July 21, 2008 (upper left) and October 20, 2008 (lower left), with time to maturity equal to 89 days. The difference $\widehat{\theta}[\pi] - \widehat{\rho}[\pi]$ for these two trading days is shown on the right.

between the two trading days may be explained if we have a look at the at-the-money implied volatilities listed in Table 1. On October 20, 2008, we were in the middle of the financial crisis and some volatilities exceed 100%. July 21, 2008, on the contrary, was much calmer and volatility levels were in a more normal range.

Implied volatility. The implied at-the-money volatility of a broad index is often used as an estimate for the future volatility of the market. The third graph of Fig. 7 shows the implied Dow Jones volatility VXO[30 days] for the period January 2000–October 2009. The VXO and the ICX index move strongly together, indicating that periods of increased co-movement are going hand in hand with periods of increased volatility. It is well-documented that periods of high market stress are characterized by high levels of volatility and a strong co-movement between stock prices. Today, market participants often use indices for the future level of volatility as a fear gauge.

Fig. 8 shows the implied correlation as a function of the implied volatility. The regime around January 2009 (circles) was characterized by high implied volatilities and high implied correlations. However, watching only volatility can be too narrow
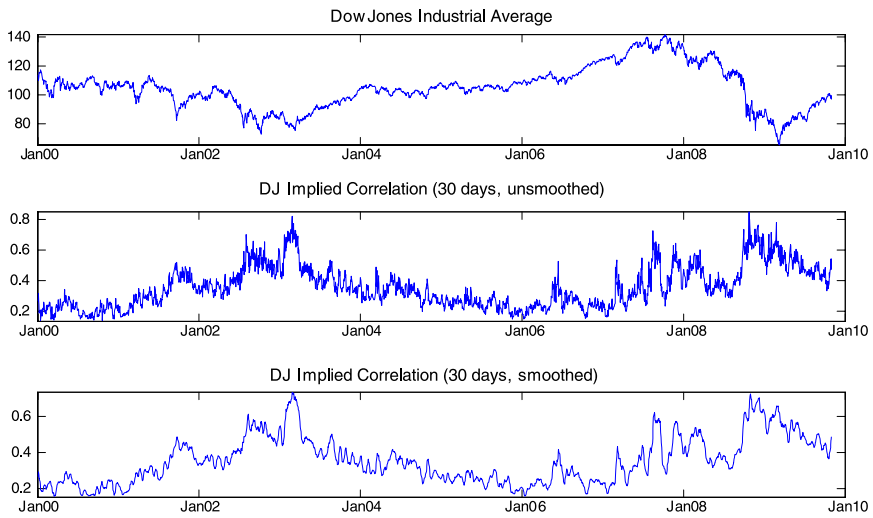
Fig. 6. The Dow Jones price levels between January 2000 and October 2009 (first graph). The second plot shows the 30 days DJ-ICX. The third graph shows a smoothed version of the second plot.
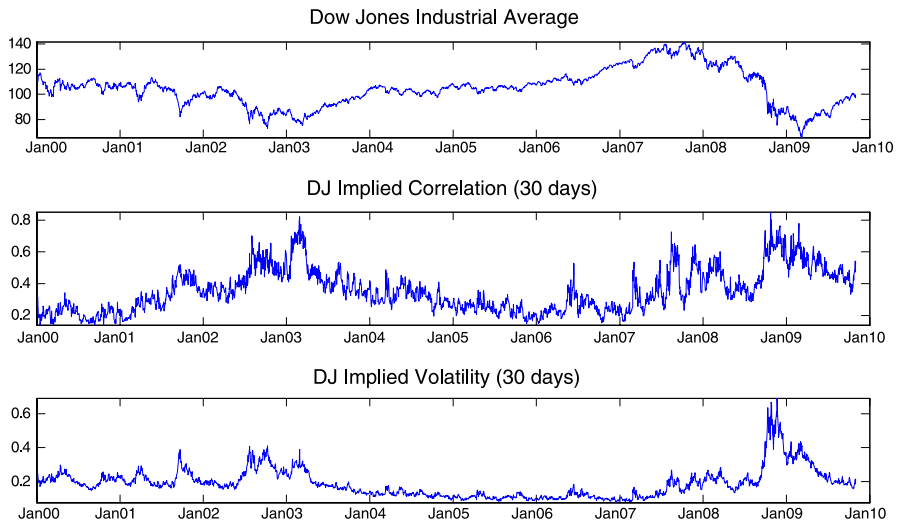

Fig. 7. The Dow Jones price levels between January 2000 and October 2009 (first graph). The second plot shows the 30 days DJ-ICX. The third graph shows the 30 days implied at-the-money index volatility.
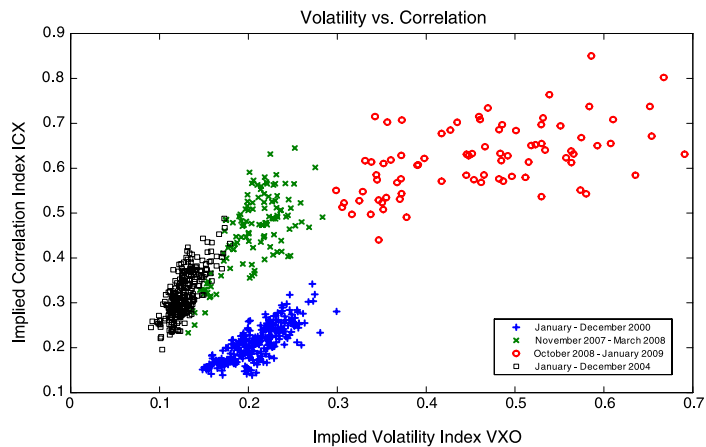

Fig. 8. Dow Jones implied correlation as a function of the implied volatility.

Table 1
At-the-money implied volatilities.

	July 21, 2008	October 20, 2008
Alcoa Incorporated	52.79%	96.30%
American Express Company	49.17%	111.59%
Bank of America	66.26%	91.05%
Boeing Corporation	37.10%	70.56%
Caterpillar	35.28%	75.69%
JP Morgan	48.97%	80.19%
Chevron	32.03%	66.50%
Citigroup	54.89%	101.28%
Coca Cola Company	26.64%	47.05%
Walt Disney Company	29.86%	56.31%
DuPont	28.08%	69.23%
Exxon Mobile	30.03%	68.44%
General Electric	31.30%	76.45%
General Motors	109.46%	215.81%
Hewlet-Packard	32.20%	49.07%
Home Depot	49.67%	68.50%
Intel	37.65%	73.30%
IBM	27.09%	64.86%
Johnson & Johnson	19.67%	46.32%
McDonald's	29.02%	48.29%
Merck & Company	36.41%	58.04%
Microsoft	32.10%	75.62%
3M	31.59%	47.20%
Pfizer	31.69%	53.32%
Proctor & Gamble	21.71%	48.15%
AT&T	33.29%	66.65%
United Technologies	31.54%	54.79%
Verizon	28.51%	55.08%
Wal-Mart Stores	29.66%	62.23%
Kraft Foods		44.89%
American International group	67.52%	

to estimate the sentiment of the market. For example, there is no big difference between the period January–December 2000 (plusses) and the period November 2007–March 2008 (crosses) when we only look at the volatility levels. But the two periods do differ when we look at the corresponding implied correlation levels. Fig. 8 also shows that during 2004 (squares), volatility levels were relatively stable, while the changes in implied correlation were more pronounced. Interesting to notice is that in the period January 2000–October 2009, there are no situations where volatility is relatively big, but correlation is low. For example, there are no trading days where $VXO[T] > 0.3$ but $ICX[T] < 0.4$.

Combining an estimate of future volatility with an estimate for the future level of co-movement between stocks gives a broader view of the fear present in the market. Instead of looking only at volatility, the couple $(VXO[T], ICX[T])$ gives a more accurate picture of today's level of market fear and may be used to detect periods of increased market fear in an early stage; see also [47].

6. Conclusion

In this paper it is assumed that options on an index and the stocks composing this index are traded on an option exchange and their prices can be observed. Therefore, the pricing of these derivatives is not the real issue. Given that the marginals are lognormal distributed, we can extract the marginal volatility parameters implied by the observed vanilla options. Typically, out-of-the-money put options require a higher volatility parameter than out-of-the-money call options. Provided that the index option prices are given, we now use the volatility surface for each stock and extract the single correlation parameter such that the model price matches the market price. The traditional approach for determining implied correlation levels is a particular member of our framework. Furthermore, we can improve the traditional approach by considering more accurate approximations for the index option price. We show that this new implied correlation index outperforms the traditional approach. It turns out that the traditional implied correlation underestimates the real correlation. This error is more pronounced when some stock volatilities are large compared to the other volatility levels, which may occur in times of market stress. We conclude that in a situation of increased market fear, the traditional implied correlation is a dangerous measure, because it may give misleading signals. However, our numerical illustrations show that determining implied correlations with the traditional approach may be justified for strike prices which are close to the at-the-money strike price.

We have shown that implied correlation estimates can be used to define an index, called the Implied Correlation Index (ICX), which reflects the market's perception about future (short term) co-movement between stock prices. A value close to 1 indicates that the market expects stocks to move almost perfectly together, which implies that one cannot benefit from a diversification effect when composing a portfolio of stocks. It is well-known that high levels of co-movement are often associated with an increased market fear. Hence the diversification benefit one hopes for is evaporating when it is needed the most.

Therefore, it is important to construct indicators which are able to signal periods of increased market fear. Nowadays, implied volatility levels serve as fear indicators. We have shown that using an estimate for the future (short term) volatility together with an estimate for the future level of co-movement gives a more accurate description of the future level of market fear.

Acknowledgments

References

[1] V.D. Skintzi, G. Skiadopoulos, A.-P.N. Refenes, The effect of misestimating correlation on value at risk, J. Alternat. Invest. 7 (4) (2005) 66–82.
[2] A.S. Wong, 'Quantitative impact of correlation errors on basket options with time-varying correlations', J. Futures Markets 32 (2) (2012) 152–165.
[3] T. Moskowitz, 'An analysis of covariance risk and pricing anomalies', Rev. Financ. Stud. 16 (2) (2003) 417–457.
[4] J. James, K. Kasikov, K.-A. Edwards, 'The end of diversification', Quant. Finance 12 (11) (2012) 1629–1636.
[5] V.D. Skintzi, A.N. Refenes, 'Implied correlation index: A new measure of diversification', J. Futures Markets 25 (2005) 171–197. http://dx.doi.org/10.1002/fut.20137.
[6] R. Cont, R. Deguest, 'Equity correlations implied by index options: Estimation and model uncertainty analysis', Math. Finance 23 (3) (2013) 496–530.
[7] Chicago Board Options Exchange 2009, CBOE S&P 500 implied correlation index, Working Paper.
[8] R. Kaas, J. Dhaene, M.J. Goovaerts, 'Upper and lower bounds for sums of random variables', Insurance Math. Econom. 27 (2) (2000) 151–168.
[9] J. Dhaene, M. Denuit, M. Goovaerts, R. Kaas, D. Vyncke, 'The concept of comonotonicity in actuarial science and finance: applications', Insurance: Math. Econom. 31 (2) (2002) 133–161.
[10] J. Dhaene, A. Kukush, D. Linders, 'The multivariate black & scholes market: conditions for completeness and no-arbitrage', Theory Probab. Math. Statist. (88) (2013) 1–14.
[11] T. Björk, Arbitrage Theory in Continuous Time, Oxford University Press, 1998, p. 311.
[12] F. Black, M. Scholes, 'The pricing of options and corporate liabilities', J. Polit. Econ. 81 (3) (1973) 637–654.
[13] R. Brooks, J. Corson, J.D. Wales, 'The pricing of index options when the underlying assets all follow a lognormal diffusion', Adv. Futures Options Res. 7 (1994).
[14] M. Milevsky, S. Posner, 'A closed-form approximation for valuing basket options', J. Deriv. 5 (4) (1998) 54–61.
[15] D. Brigo, F. Mercurio, F. Rapisarda, R. Scotti, 'Approximated moment-matching dynamics for basket-options pricing', Quant. Finance 4 (1) (2004) 1–16.
[16] G. Deelstra, J. Liinev, M. Vanmaele, 'Pricing of arithmetic basket options by conditioning', Insurance Math. Econom. 34 (1) (2004) 55–77.
[17] R. Carmona, V. Durrleman, 'Generalizing the black-scholes formula to multivariate contingent claims', J. Comput. Finance 9 (2006) 43–67.
[18] R. Korn, S. Zeytun, 'Efficient basket monte carlo option pricing via a simple analytical approximation', J. Comput. Appl. Math. 243 (1) (2013) 48–59.
[19] M. Krekel, J. de Kock, R. Korn, T.-K. Man, 'An analysis of pricing methods for baskets options', Wilmott Mag. (April) (2007) 82–89.
[20] D. Linders, Pricing index options in a multivariate Black & Scholes model, in: Research report afi-1383 feb, KU Leuven - Faculty of Business and Economics, Leuven, 2013.
[21] S. Vanduffel, T. Hoedemakers, J. Dhaene, 'Comparing approximations for risk measures of sums of non-independent lognormal random variables', N. Am. Actuar. J. 9 (4) (2005) 71–82.
[22] J. Dhaene, S. Vanduffel, M.J. Goovaerts, R. Kaas, D. Vyncke, 'Comonotonic approximations for optimal portfolio selection problems', J. Risk Insur. 72 (2005) 253–301.
[23] K. Van Weert, J. Dhaene, M.J. Goovaerts, 'Comonotonic approximations for a generalized provisioning problem with application to optimal portfolio selection', J. Comput. Appl. Math. 235 (2011) 3245–3256.
[24] K. Van Weert, J. Dhaene, M. Goovaerts, 'Optimal portfolio selection for general provisioning and terminal wealth problems', Insurance Math. Econom. 47 (1) (2010) 90–97.
[25] J. Dhaene, M. Denuit, M. Goovaerts, R. Kaas, D. Vyncke, 'The concept of comonotonicity in actuarial science and finance: theory', Insurance: Math. Econom. 31 (1) (2002) 3–33.
[26] G. Deelstra, J. Dhaene, M. Vanmaele, An overview of comonotonicity and its applications in finance and insurance, in: B. Oksendal, G. Nunno (Eds.), Advanced Mathematical Methods for Finance, Springer, Berlin, Heidelberg, 2011, pp. 155–179.
[27] M. Rubinstein, 'Implied binomial trees', J. Finance 49 (3) (1994) 771–818.
[28] N. Chriss, Black Scholes and Beyond: Option Pricing Models, McGraw-Hill Companies, Incorporated, 1996.
[29] J. Gatheral, The Volatility Surface: A Practitioner's Guide, Wiley Finance, Wiley, 2006.
[30] F. Schmid, R. Schmidt, 'Multivariate extensions of Spearman's rho and related statistics', Statist. Probab. Lett. 77 (4) (2007) 407–416.
[31] J. Dhaene, D. Linders, W. Schoutens, D. Vyncke, 'A multivariate dependence measure for aggregating risks', J. Comput. Appl. Math. 263 (2014) 78–87.
[32] H. Huang, M. Milevsky, J. Wang, Ruined moments in your life: How good are the approximations? Insurance: Math. Econom. 3 (34) (2004) 421–447.
[33] P. Embrechts, A. McNeil, D. Straumann, Correlation and dependence in risk management: properties and pitfalls, in: 'Risk Management: Value at Risk and Beyond', Cambridge University Press, 1999, pp. 176–223.
[34] F. Guillaume, W. Schoutens, 'Calibration risk: Illustrating the impact of calibration risk under the heston model', Rev. Deriv. Res. 15 (1) (2012) 57–79.
[35] J.C. Cox, S.A. Ross, M. Rubinstein, 'Option pricing: A simplified approach', J. Financ. Econom. 7 (3) (1979) 229–263.
[36] P. Carr, L. Wu, 'A tale of two indices', J. Deriv. 13 (3) (2006) 13–29.
[37] R. Whaley, 'The investor fear gauge', J. Portfolio Manag. 26 (2000) 12–17.
[38] Chicago Board Options Exchange 2003, The CBOE volatility index - VIX, White Paper.
[39] T. Moosbrucker, 'Explaining the correlation smile using Variance Gamma distributions', J. Fixed Income 16 (1) (2006) 71–87.
[40] H. Albrecher, S. Ladoucette, W. Schoutens, A generic one-factor lévy model for pricing synthetic cdos, in: M. Fu, R. Jarrow, J.-Y. Yen, R. Elliott (Eds.), Advances in Mathematical Finance', Applied and Numerical Harmonic Analysis, Birkhäuser, Boston, 2007, pp. 259–277.
[41] J. Garcia, S. Goossens, V. Masol, W. Schoutens, 'Lévy base correlation', Wilmott J. 1 (2009) 95–100.
[42] B. Tavin, Hedging dependence risk with spread options via the power frank and power student t copulas, in: Technical Report, Université Paris I Panthéon-Sorbonne, 2013, Available at SSRN: http://ssrn.com/abstract=2192430.
[43] F. Guillaume, D. Linders, Stochastic modeling of herd behavior indices, Research report AFI-1490 FEB, KU Leuven, 2014.
[44] J. Dhaene, D. Linders, W. Schoutens, D. Vyncke, 'The herd behavior index: A new measure for the implied degree of co-movement in stock markets', Insurance Math. Econom. 50 (3) (2012) 357–370.
[45] P. Laurence, A new tool for correlation risk management: the market implied comonotonicity gap, Global Derivatives, Paris, Invited Talk, May 2008.
[46] P. Laurence, T.-H. Wang, 'Distribution-free upper bounds for spread options and market-implied antimonotonicity gap', Eur. J. Finance 14 (8) (2008) 717–734.
[47] J. Dhaene, J. Dony, M.B. Forys, D. Linders, W. Schoutens, Fix - the fear index: Measuring market fear, in: M. Cummins, et al. (Eds.), 'Topics in Numerical Methods for Finance, in: Springer Proceedings in Mathematics and Statistics', 2012.
[48] C. Kim, Y. Choi, W. Lee, J.Y. Ahn, 2014, Analyzing herd behavior in global stock markets: an interncontinental comparison, Working paper, Department of statistics, Ewha Womans University, Seoul.