XAIA INVESTMENT

Copula-regression for the risk-neutral distribution of two highly correlated assets

Amelie Hüttner
XAIA Investment GmbH
Sonnenstraße 19, 80331 München, Germany
amelie.huettner@xaia.com

Jan-Frederik Mai
XAIA Investment GmbH
Sonnenstraße 19, 80331 München, Germany
jan-frederik.mai@xaia.com

Date: June 3, 2016

Abstract We seek to model the multivariate distribution of two financial indices (X, Y) at some future time point, with the following partial information given. From derivatives prices on financial markets we can infer information about (the market's opinion of) the univariate distribution functions F_X of X and F_Y of Y. Empirical evidence and experts' opinions further justify the assumption that $Y \approx g(X)$ for some monotone function g. Consequently, a reasonable model for the multivariate law of (X, Y) is to join the retrieved marginal laws F_X and F_Y with a copula function C_ρ, where the model parameter $\rho < 1$ controls the level of deviation from the case $\rho = 1$ representing the base "regression" model $Y = g(X)$. The present article clarifies in which situations such a model can be helpful, discusses desirable properties of the applied copula function C_ρ, and presents a method to retrieve F_X and F_Y in such a way that their inverses are given in closed form in order to be able to simulate from the resulting model for (X, Y) efficiently via combining a copula sampling algorithm with the classical inversion method for the marginals.

1 Motivation

The motivation for the present research is to find a bivariate risk-neutral probability distribution for the pair of two financial indices that are highly correlated at a future point in time, say (X, Y). For example, economic reasoning suggests that an equity index is highly negatively correlated with a credit index associated with the same regional sector. That is because if stock markets perform well, this is likely to coincide with associated credit spreads becoming tighter, cf. Fung et al. (2008). When looking at historical data, regressing one index on the other is a resonable and widely applied method to study their functional relationship, and the deviations from such fundamental relationship. In contrast, a risk-neutral distribution is forward-looking by definition, since it represents the market's opinion about the values of the two indices in the future. The market provides partial information about such a risk-neutral distribution by quoting prices for financial contracts whose cashflows depend on the indices' future values. In particular, it is a classical task in Mathematical Finance to retrieve the univariate marginal risk-neutral distribution functions F_X and F_Y of X and Y from prices for standardized option contracts referring to either X or Y. Even though there is typically no information available to determine the (risk-neutral) dependence structure between the two indices, economic fundamentals justify the use of a very strong dependence between the two, not to say an almost functional relationship between the two. Consequently, it is a natural idea to try to regress one marginal probability distribution on the other to model the pair (X, Y) as the outcome of a deviation around a fundamental regression relationship $Y = g(X)$, like this is done on historical data. One important difference is that the univariate probability distributions of X and Y are known a priori (i.e. are retrieved from option data), only the dependence structure has to be modeled via an adequate copula function. Akin to the standard historical regression idea, this copula should intuitively represent a symmetric deviation around a fundamental regression relationship $Y = g(X)$. This is precisely the idea that is pursued in the present article. In order to make this copula-regression idea more plastic, the whole procedure is described along a concrete example involving an equity index for X (the EuroStoxx) and a highly negatively correlated credit default swap (CDS) index for Y (the ITRX EUR). The remainder of this article is organized as follows: In Section 2, we explain the extraction of the marginal distribution functions F_X and F_Y from quoted call option data, and indicate some necessary approximations for the application of the proposed method to options on Index CDS. Section 3 discusses the choice of an appropriate copula class for the joint modeling of (X, Y). A concrete application of the presented method is finally discussed in Section 4. Section 5 concludes.

2 Model-free extraction of the marginal distribution function from option prices

We denote by X the value of a finanical index at a future time point T. The market trades European call options with maturity T and strike prices $0 = K_0 < K_1 < K_2 < \ldots < K_n < \infty$ on this index, i.e. we can observe market prices for these options. Under the assumption that these market prices are free of arbitrage, there exists a distribution function F_X such that the price $C(K_i)$ of the call option with strike price K_i can be written as

$$C(K_i) = DF(T) \int_{K_i}^{\infty} (x - K_i) \,\mathrm{d}F_X(x), \quad i = 0, \ldots, n, \tag{1}$$

where $DF(T)$ denotes the price of a risk-free zero coupon bond with maturity T. The function F_X is called a *risk-neutral distribution function* for X and is by no means unique. Its interpretation is that if X has distribution function F_X (under some so-called risk-neutral probability measure $\mathbb{Q}$), we write $X \sim F_X$, then the call prices can be written as expected value of the terminal payoff, i.e.

$$C(K_i) = DF(T) \, \mathbb{E}^{\mathbb{Q}}\big[\max\{X - K_i, 0\}\big], \quad i = 0, \ldots, n.$$

In order to make the choice of F_X unambiguous, the articles Neri, Schneider (2012, 2013) postulate that $\mathrm{d}F_X(x) = f_X(x)\,\mathrm{d}x$ for a probability density function f_X, and demand that f_X maximizes entropy among all probability densities satisfying the constraint (1), i.e.

$$f_X := \text{argmax}_{f \in \mathcal{C}(K_1,\ldots,K_n)} \Big\{ -\int_0^\infty f(x)\,\log\big(f(x)\big)\,\mathrm{d}x \Big\},$$
$$\mathcal{C}(K_1,\ldots,K_n) := \big\{ f : (0,\infty) \to (0,\infty)\ : \\ \mathrm{d}F_X(x) = f(x)\,\mathrm{d}x \text{ satisfies } (1) \big\}.$$

Provided that the market prices are free of arbitrage, it is shown in Neri, Schneider (2012, 2013) that f_X is well-defined, unique, and piecewise exponential. Further, an iterative procedure – based on a Newton algorithm – is demonstrated to be reliable and efficient for the computation of (the piecewise exponential parameters of) f_X. Moreover, f_X is called the *Buchen-Kelly density* associated with the call option prices $K_1, \ldots, K_n$. Besides the ability to be well-defined and being reliably computable, the Buchen-Kelly density additionally has the following two key features: (a) its definition in terms of entropy-maximization is interpreted as no additional (model) information - thus no model mis-specification - being present in f_X, and (b) a Monte Carlo simulation of $X \sim f_X$ is highly efficient via the so-called inversion method, see, e.g. Mai, Scherer (2012, Chapter 6.3, p. 234). There is one practical problem when applying the Buchen-Kelly algorithm of Neri, Schneider (2012, 2013): the Buchen-Kelly density f_X depends critically on the observed input prices $C(K_i)$ and tends to be quite spiky if the latter are chosen thoughtlessly. In practice, for each strike price K_i the market quotes a bid price $\underline{C}(K_i)$ and an ask price $\overline{C}(K_i)$, and any choice of mid prices $C(K_i)$ within the respective intervals $[\underline{C}(K_i), \overline{C}(K_i)]$ is an admissible input for the Buchen-Kelly algorithm, provided obvious no-arbitrage conditions are satisfied[1]. However, in order to make sure that the associated Buchen-Kelly density f_X is smooth, a deliberate choice needs to be made. We recommend to generate the input mid prices $C(K_i)$ from a simple model with smooth parametric density f_θ, which is convenient to implement and flexible enough to provide an excellent fit to the observed bid and ask prices. More precisely, we consider mid prices of the form

$$C(K_i;\theta) = DF(T)\int_{K_i}^\infty (x-K_i)\,f_\theta(x)\,\mathrm{d}x, \quad i = 0,\ldots,n,$$

where the parameter(s) θ are chosen as the minimizer of a non-negative penalization function $P(\theta)$ which is strictly positive if and only if θ is such that at least one $C(K_i;\theta)$ lies outside its admissible interval $[\underline{C}(K_i), \overline{C}(K_i)]$. In the present article, we apply the three-parametric submodel of Gatheral's five-parametric SVI-model which is shown to be free of butterfly arbitrage in Gatheral, Jacquier (2014). It is a simple parametric model for the implied volatility smile, which can be transformed into a smooth density function f_θ because it is guaranteed to be free of butterfly arbitrage. The model's fitting capacity to observed call prices (or, equivalently, to the observed volatility smile) is incredibly good and the model is easy to implement because the volatility smile is by construction given in closed form. Hence, the penalization

[1] Essentially, the function $K_i \mapsto C(K_i)$ needs to be non-increasing and convex.

function $P(\theta)$ can directly be implemented on implied volatilities instead of call prices. Having retrieved the mid prices $C(K_i;\theta)$, we stick them into the Buchen-Kelly algorithm of Neri, Schneider (2012, 2013) to retrieve f_X. By construction, f_X is a piecewise exponential approximation of the smooth density f_θ, hence itself quite smooth, at least satisfactorily smooth for practical applications.

We provide an example for X denoting the value of the EuroStoxx index on March 18, 2016. The distribution F_X is extracted from option prices quoted on November 26, 2015, with strikes in the range $[2600, 4000]$. The price quote for the EuroStoxx was 3498.62. Figure 1 displays model-generated implied volatilities for the SVI and Buchen-Kelly model compared to market quotes, as well as the corresponding implied densities. Both models fit the market quotes very well.

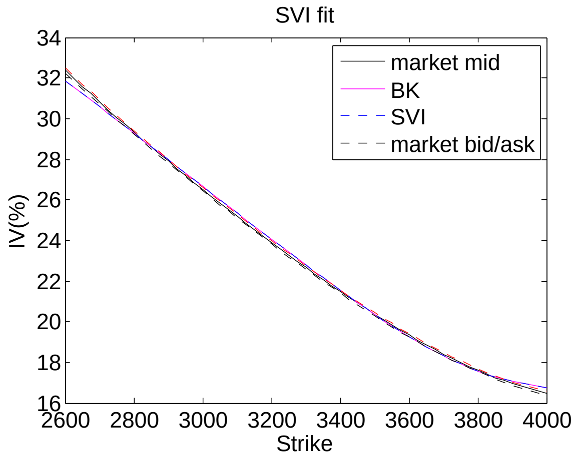

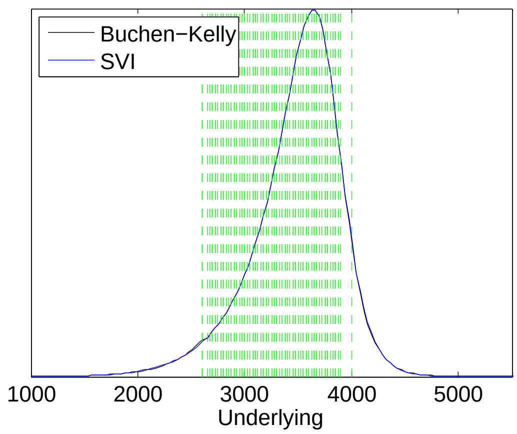

Fig. 1: SVI and Buchen-Kelly fit to EuroStoxx data. Top: Model-generated implied volatilities compared to market quotes. Bottom: Implied SVI and Buchen-Kelly densities.

2.1 Application to index CDS options

As indicated in the introduction, we seek to apply the Buchen-Kelly method of the current Section also to X denoting a credit default swap (CDS) index, because intuitively there is a high correlation between certain stock indices and CDS indices due to the fact that there is a huge overlap in the underlying companies of which the index is composed in both market segments. Unlike a stock index, unfortunately, a CDS index is not directly a tradeable asset, and the market does not directly quote call options on the index. However, the market does quote so-called Payer options, which can be approximated pretty well by call options on the CDS index, so that the Buchen-Kelly algorithm may still be applied. These CDS market-specific technicalities are explained in the sequel.

Technically, the quoted value of the CDS index equals the so-called running spread of an associated Index CDS contract. An Index CDS is an insurance contract between two parties that offers protection on a basket of credit-risky assets. In case of default of one or more of the referenced assets during the lifetime of the contract, the protection seller compensates the protection buyer for the losses suffered. In return, he receives a standardized premium payment c from the protection buyer, which is proportional to the remaining nominal in the reference basket of assets. The value of the Index CDS at the valuation date $t = 0$ (from the point of view of the protection buyer) equals the expected discounted value of protection payments, denoted by $\mathbb{E}[DDL(0, \hat{T})]$, less the expected discounted value of premium payments $c \cdot \mathbb{E}[DPL(0, \hat{T})]$, which are linear in the premium c, where $\hat{T}$ is the maturity date of the contract. The market price of the contract is not directly quoted in the market, however, has to be computed from the observed CDS index value, the so-called *running spread* $s_{0,\hat{T}}$ of the Index CDS contract, which is defined by:

$$s_{0,\hat{T}} = \frac{\mathbb{E}[DDL(0, \hat{T})]}{\mathbb{E}[DPL(0, \hat{T})]}.$$

The value of an Index CDS in terms of the observed running spread is thus computed as

$$\begin{aligned} ICDS(0, \hat{T}) &= \mathbb{E}[DDL(0, \hat{T})] - c \cdot \mathbb{E}[DPL(0, \hat{T})] \\ &= (s_{0,\hat{T}} - c)\mathbb{E}[DPL(0, \hat{T})]. \end{aligned}$$

The quantity $\mathbb{E}[DPL(0, \hat{T})]$ is model-dependent. By market convention it is computed from a standard conversion formula based on simple modeling assumptions as a function of the quoted running spread $s_{0,\hat{T}}$, i.e. the quoted CDS index. Under the additional assumption that the CDS premium is paid continuously and the assumption of a constant interest rate r, it is very well approximated by

$$\mathbb{E}[DPL(0, \hat{T})] \approx f(s_{0,\hat{T}}) := \frac{1 - e^{-(r + s_{0,\hat{T}}/(1-R))\hat{T}}}{r + s_{0,\hat{T}}/(1 - R)},$$

where $R \in [0, 1]$ denotes a market standard recovery rate assumption, e.g. $R = 40\%$.

Index CDS options come in the form of payer and receiver options. The payer (receiver) option gives its holder the right to enter as protection buyer (seller) into an Index CDS at option expiry T. Both are usually traded European-style, and involve compensation payments for the proportion of the reference basket that defaults prior to the option maturity. For ease of calculations, we

assume, however, that the considered options are traded without this so-called front-end protection.
The strike s^K of an index CDS option is typically quoted in terms of running spread, whereas its price is given in terms of the value of the underlying Index CDS at T. The price of a Payer option with expiry T and maturity of the underlying Index CDS $\hat{T}$ is therefore:

$$\begin{aligned} Payer(0,T,\hat{T}) &= \mathbb{E}\big[DF(T)(ICDS(T,\hat{T}) - (s^K - c)\mathbb{E}[DPL(T,\hat{T})])^+\big] \\ &= \mathbb{E}\big[DF(T)\big((s_{T,\hat{T}} - c)f(s_{T,\hat{T}}) - (s^K - c)f(s^K)\big)^+\big], \end{aligned}$$

where $s_{T,\hat{T}}$ denotes the CDS index at the future time T. For a detailed introduction to the valuation of Index CDS and Index CDS options, see e.g. Martin (2012). For approaches to Index CDS option valuation involving up-front default compensation, see e.g. Armstrong, Rutkowski (2009) and Brigo, Morini (2011).

The payoff of the payer option as a function of the random variable $s_{T,\hat{T}}$ is

$$g(s_{T,\hat{T}}) := \big((s_{T,\hat{T}} - c)f(s_{T,\hat{T}}) - (s^K - c)f(s^K)\big)^+,$$

which, unfortunately, is not precisely a call payoff, as desired for the Buchen-Kelly method described earlier. However, it can be approximated satisfactorily well by a call payoff $g(s) \approx x^K(s - s^K)^+$, where x^K is chosen such that $g(s_*) = x^K\,(s_* - s^K)$ for some cutoff level $s_* > s^K$, i.e.

$$x^K = \frac{(s^* - c)f(s^*) - (s^K - c)f(s^K)}{s^* - s^K}.$$

For the remainder of this article, we set $s^* = 2s_{0,\hat{T}}$, twice the currently quoted CDS index at valuation date. Usually, quoted strikes for Index CDS options lie in a range of $[0.5s_{0,\hat{T}}, 1.5s_{0,\hat{T}}]$, therefore $2s_{0,\hat{T}}$ represents a sensible cutoff point in most cases. The presented call approximation only differs significantly from the true payoff for terminal running spreads $s_{T,\hat{T}} \gg s^K$, which, as we will see from the market-extracted density functions, are highly unlikely under the resulting risk-neutral distribution. Figure 2 displays the exact payoff in comparison to the approximating call-like payoff.
Summing up, we apply the described Buchen-Kelly algorithm to the following call approximations on the underlying $X := s_{T,\hat{T}}$:

$$\frac{1}{x^K} Payer(0,T,\hat{T}) \approx DF(T)\mathbb{E}[(X - s^K)^+].$$

Consider exemplarily the value X of the iTraxx Europe 5Y index on March 18, 2016. The presented method is applied to option quotes of November 26, 2015, with strike spreads in the range $[40, 105]$ bps[2] and a price quote of 69.762 bps for the iTraxx index itself. Figure 3 shows that both the SVI and Buchen-Kelly model fit the quoted implied volatilities very well. Further the implied densities for both models are displayed.

[2]Running spreads are quoted in basis points (bps): 1 bps = $1/10000$.

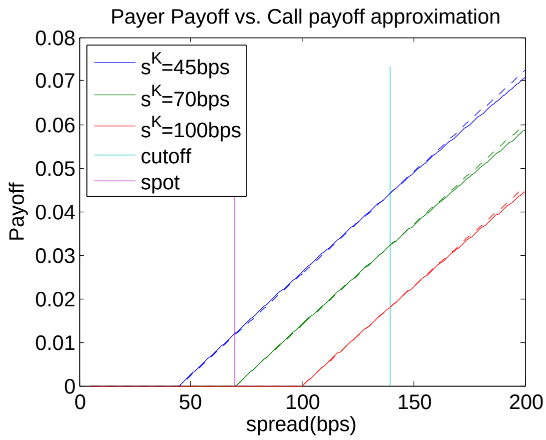

Fig. 2: Approximation (dotted lines) of the exact (solid lines) ITRX Payer option payoff for three different strikes with call payoffs $x^K \cdot (s_t - s^K)^+$.

3 Jointly modeling two highly correlated indices

Furnished with a robust method for inferring the marginal distributions from quoted option prices, we proceed to the joint modeling of two index prices (X, Y). Often, a monotonic relationship between the indices is apparent in historical data, and also a highly reasonable model from an expert's point of view. Consequently, we postulate $Y = g(X)$, with $g(\cdot)$ some monotonic function. We assume for now g increasing. The presented approach can easily be transferred to decreasing g, cf. Section 3.1.
Knowing the risk-neutral marginal laws F_X and F_Y, the function g relating the two indices is already determined:

$$\begin{aligned} F_Y(y) = \mathbb{P}(g(X) \leq y) = \mathbb{P}(X \leq g^{-1}(y)) = F_X(g^{-1}(y)) \\ \Leftrightarrow g(x) = F_Y^{-1}(F_X(x)). \end{aligned} \quad (2)$$

In reality, historical observations give rise to take into account the possibility that the pair (X, Y) deviates from this explicit relation $Y = g(X)$, but that we may interpret these deviations as noise blurring the postulated relation in a symmetric manner. Resorting to a traditional regression way of thinking, this variation would be taken into account by introducing a symmetrically distributed error term ϵ with $\mathbb{E}[\epsilon] = 0$, say $Y = g(X) + \epsilon$. In the presented case, however, the marginal distributions are known, so this additive inclusion of an error term is not appropriate, as it obviously affects the known marginal distribution function of Y. Instead, the deviation from the functional relation g has to be modeled in a way that leaves the margins unaltered. Taking this into consideration, we propose to model the joint distribution of (X, Y) by means of the extracted margins F_X and F_Y, and a bivariate copula C. A copula is a multivariate distribution function with standard uniform margins. Given C and the two margins, Sklar's Theorem states that $F_{X,Y}(x, y) = C(F_X(x), F_Y(y))$ defines a joint distribution function for (X, Y) with the given margins. Sample pairs (U_1, U_2) of uniform random variables with given dependence structure from C can be transformed to sample pairs of (X, Y) according to the inversion method as

$$(X, Y) = \big(F_X^{-1}(U_1), F_Y^{-1}(U_2)\big).$$

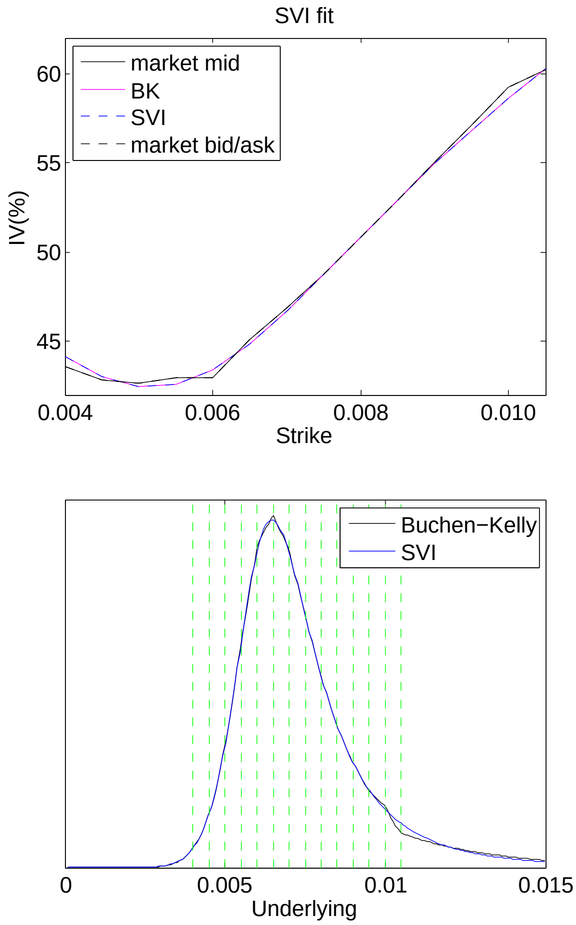

Fig. 3: SVI and Buchen-Kelly fit to iTraxx data. Top: Model-generated implied volatilities compared to market quotes. Bottom: Implied SVI and Buchen-Kelly densities.

The assumption of a perfect positive dependence between X and Y, i.e. $Y = g(X)$, corresponds to the copula $C(u_1, u_2) = \min\{u_1, u_2\}$, the so-called comonotonicity copula, see Mai, Scherer (2012, Example 1.2, p. 5). Therefore, when choosing an appropriate family of parametric copulas for modeling purposes, one should make sure that the chosen class contains the comonotonicity case. We further want to model the deviations from the explicit functional relation $Y = g(X)$ in a symmetric way, analogous to the symmetric error term modeling in a standard regression. To do so, the bivariate copula C desirably exhibits two types of symmetry, namely exchangeability and radial symmetry, whose definitions are recalled for the convenience of the reader.

Definition 3.1 (Exchangeability and radial symmetry)

(a) A random vector (X, Y) is called *exchangeable* if the probability distribution of (X, Y) equals that of (Y, X).

(b) A random vector (X, Y) is called *radially symmetric* about a point (a_1, a_2) if the probability distribution of $(X - a_1, Y - a_2)$ equals that of $(a_1 - X, a_2 - Y)$.

A copula C is *exchangeable* (*radially symmetric*) if samples $(U_1, U_2) \sim C$ are exchangeable (radially symmetric about (0.5,0.5)). It is well-known that C is exchangeable if and only if $C(u_1, u_2) = C(u_2, u_1)$, and C is radially symmetric if and only if

$$C(u_1, u_2) = C(1 - u_1, 1 - u_2) + u_1 + u_2 - 1.$$

Exchangeabilty means that samples from C are distributed symmetrically about the diagonal, and radial symmetry means that samples are distributed point symmetrically about (0.5,0.5). Combining these two symmetries yields a third symmetry about the counterdiagonal $U_2 = 1 - U_1$. When translating samples $(U_1, U_2) \sim C$ to sample pairs $(X, Y) = (F_X^{-1}(U_1), F_Y^{-1}(U_2))$, these symmetries about the diagonal and about (0.5,0.5) translate to the following symmetries:

$$\begin{aligned} C \text{ exchangeable } \Rightarrow \quad & (X, Y) \overset{d}{=} (g^{-1}(Y), g(X)), \\ C \text{ radially symmetric } \Rightarrow \quad & \\ & (X, Y) \overset{d}{=} (F_X^{-1}(1 - F_X(X)), F_Y^{-1}(1 - F_Y(Y))). \end{aligned}$$

Both exchangeability and the additional symmetry about the counterdiagonal yield, in terms of (X, Y), symmetries on certain 'contour lines', whose shapes depend on g. Radial symmetry yields a point symmetry about $(F_X^{-1}(0.5), F_Y^{-1}(0.5))$. See Figure 4 for an illustration how points in $[0, 1]^2$ are translated from copula samples to samples of (X, Y) using the margins F_X, F_Y from our example.

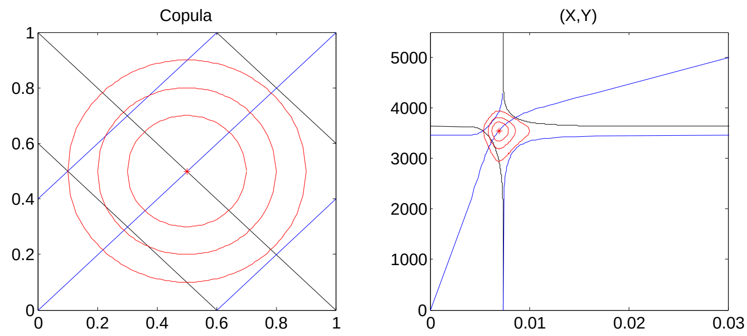

Fig. 4: Translation of copula samples in $[0, 1]^2$ lying on certain contour lines to samples of (X, Y).

In applications, we consider parametric copula families $(C_\rho)_{\rho \in [0,1]}$ with $C_1(u_1, u_2) = \min\{u_1, u_2\}$. Intuitively, we are interested in modeling the dependence between the highly correlated variables (X, Y) via a parameter $\rho \approx 1$ close to one, representing a small disturbance from the base case $Y = g(X)$, which is obtained for $\rho = 1$. The underlying idea of this approach stems from a traditional regresson way of thinking, with the sole difference that the error term is included via a copula in order to keep the known marginal distributions unchanged. Since the error term in a traditional regression analysis is modeled symmetrically, we also wish to use a copula that satisfies an appropriate meaning of "error symmetry", and we believe that copula families with C_ρ being both exchangeable and radially symmetric satisfy this demand. The following two copula families both satisfy the aforementioned demands but exhibit antipodal stochastic behavior:

1. A *Gaussian copula* to model noise, see Meyer (2009) for background. Intuitively, Gaussian copulas generate a random noise symmetric about the base case relation $Y = g(X)$. The magnitude of the noise is determined by the parameter ρ. Loosely speaking, almost all samples will violate the relation $Y = g(X)$ but all violations remain within a region around this relation whose size is controled by ρ.
2. A *Dirichlet copula* to model outliers, given by
$$C_\rho(u_1, u_2) = \min\{u_1, u_2\}\,\frac{(1-\rho)\max\{u_1, u_2\} + 1}{2-\rho},$$
see Mai et al. (2015). Compared to Gaussian copulas, Dirichlet copulas do not generate noise, but somehow the complete opposite, namely generate absolutely erratic and wild deviations from the base case relation $Y = g(X)$, whose rate of occurence is controled by ρ. Loosely speaking, most samples will satisfy the base case relation $Y = g(X)$ but a small set of samples, whose size is controled by ρ, violates this relation dramatically.

As a side remark, it is also possible to use a convex combination between Gaussian and Dirichlet copulas in order to model both noise and outliers. Such a convex combination also satisfies the desired symmetry properties.
For an illustration of sample pairs from these copula classes see Figure 5. Figure 6 shows the corresponding pairs (X, Y) in our example, where we have a strong negative implied relation between ITRX and EuroStoxx. The necessary modifications to treat this case are presented in the following.

3.1 Modifications for decreasing relations

For functional relations between $Y = g(X)$ with decreasing g, the presented approach is still applicable, with the following minor modifications:
Inferring the decreasing g from the marginal distribution yields
$$\begin{aligned} F_Y(y) = \mathbb{P}(g(X) \le y) = \mathbb{P}(X \ge g^{-1}(y)) = 1 - F_X(g^{-1}(y)) \\ \Leftrightarrow g(x) = F_Y^{-1}(1 - F_X(x)). \end{aligned}$$

When generating samples of (X, Y), one can still use the same copula families as in the increasing case, but (X, Y) are generated from $(U_1, 1-U_2)$:
$$(X, Y) := (F_X^{-1}(U_1), F_Y^{-1}(1-U_2))$$
For $(U_1, U_2) \sim C_\rho$, the vector $(U_1, 1-U_2)$ is distributed according to the copula[3]
$$\tilde{C}_\rho(u_1, u_2) = u_1 - C_\rho(u_1, 1-u_2).$$

One also retains the desired symmetries: For C_ρ exchangeable, samples from $\tilde{C}_\rho$ are distributed symmetrically about the counterdiagonal. A radially symmetric copula C_ρ yields a $\tilde{C}_\rho$ which is

[3] For Gaussian copulas, one has $\tilde{C}_\rho = C_{-\rho}$, i.e. again a Gaussian copula with parameter $-\rho$.

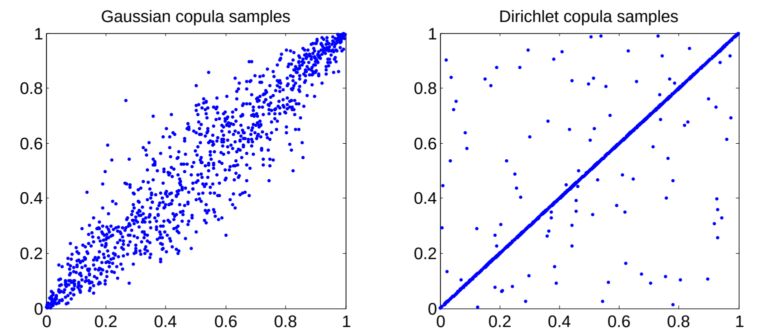

Fig. 5: Samples from a Gaussian copula $(\rho = 0.95)$ and a Dirichlet copula $(\rho = 0.9)$.

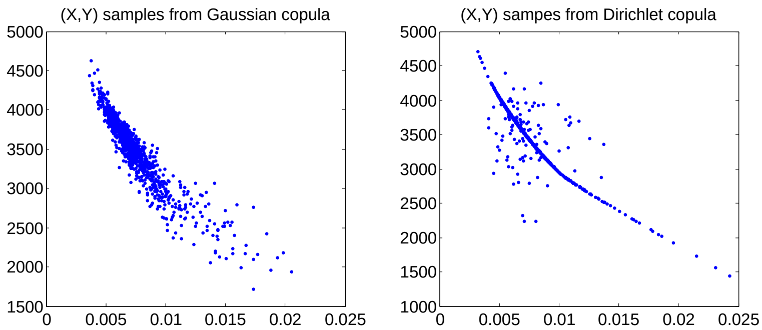

Fig. 6: Translation of these samples to pairs of (X, Y).

again radially symmetric. And finally, for C_ρ exchangeable and radially symmetric, $\tilde{C}_\rho$ also exhibits both kinds of symmetry. Furthermore, for $\rho = 1$ one obtains $\tilde{C}_1 = W$, the countermonotonicity copula that corresponds to $Y = g(X)$ for decreasing g, as desired.

4 Assessing trading opportunities

Often one might be tempted to enter into a trade solely based on historical information, forgetting to take into account that the market view on the considered assets might already have changed. While historical information, alongside with expert opinions, might give a good initial idea of the dependence structure between the considered assets, designing a trade based purely on the historical view neglects the current market opinion which is reflected in quoted option prices. The method presented in this paper therefore offers a more complete way for assessing potential trading opportunities, respecting also current market views.

In the following we consider a showcase trade on November 26, 2015. It consists of selling out-of-the-money (OTM) March 2016 Payer options on the iTraxx Eur 5Y Index (ITRX) with strike $s^K = 100$ bps and buying OTM March 2016 Put options with strike $K = 2850$ on the EuroStoxx 50 Index (SX5E), such that the initial PnL of the trade is zero. The current levels of ITRX and SX5E are 69.762 bps and 3498.62, respectively. Fixing the CDS nominal at x, the number of Puts to buy is then determined as $y = x \cdot Payer(0, T, \hat{T})/Put(0, T)$. Denoting the values of ITRX and SX5E respectively at option maturity T by X and Y respectively, where $\hat{T}$ is the corresponding CDS maturity, the PnL of this trade at T is given as follows:

$$\begin{aligned}\mathsf{PnL}(X,Y,x,y,s^K,K) &= y\max\{K-Y,0\} \\ &\quad - x\max\{(X-c)f(X)-(s^K-c)f(s^K),0\}.\end{aligned} \tag{3}$$

When looking at historical data, the trade looks quite attractive, see Figure 7. However, the implied relation obtained from quoted option prices (assuming a perfect negative dependence between the assets as seems appropriate here), which reflects current market views, is quite different from the implied relation obtained from regressing the historical time series.

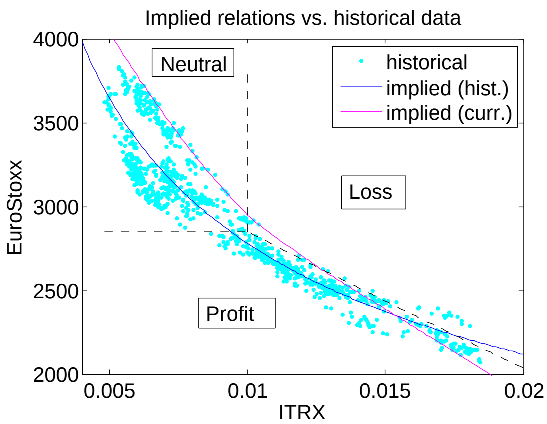

Fig. 7: PnL areas for selling ITRX Payer options ($s^K = 100$ bps) vs. buying SX5E Put options ($K = 2850$), with initial PnL equal to zero, compared to historical observations. The implied relations obtained using the presented method and obtained via regression of historical data are displayed as well.

Historical data suggests that a strong negative relation holds for (X,Y), hence we generate sample pairs (X,Y) from a Gaussian copula C_ρ with parameter[4] $\rho = 0.99$ and the extracted margins F_X and F_Y. Alternatively, we simulate $X \sim F_X$ and assume Y adheres approximately to a relation that is found from historical data via regression, namely

$$Y = \tilde{g}(X) + \epsilon, \qquad \tilde{g}(X) = \lambda \cdot X^a.$$

The parameters λ, a are estimated in a classical regression and $\epsilon \sim \mathcal{N}(0,\sigma^2)$, $\sigma = 1/3\sqrt{\mathbb{V}[\tilde{g}(X)]}$ are normally distributed residuals[5]. Comparing the empirical PnL distributions, one finds that both are asymmetric for the considered trade, although skewed in different directions, cf. Figure 8. Our proposed method implies that losses occur with a higher probability, but are less pronounced, and gains occur with a smaller probability but are potentially larger, whereas the simulation from F_X and the historical relation states exactly the opposite. Analogously, the PnL profiles one obtains from these two methods for the considered trade example differ significantly.

[4]A parameter $\rho \in [0.95, 0.99]$ seems appropriate regarding historical data.

[5]The variance of the residual terms was chosen such that a similar deviation from the implied relation was obtained as in our method with $C_{0.99}$.

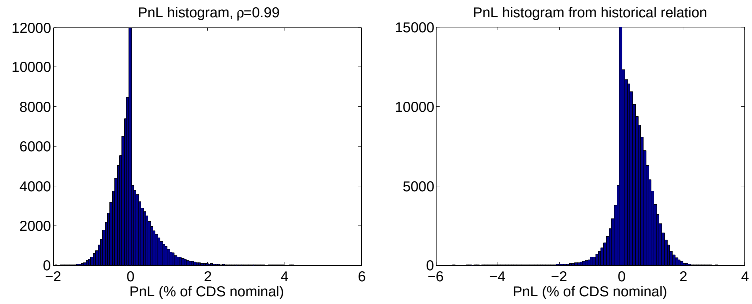

Fig. 8: Left: Histogram of PnL distribution for considered trade; (X, Y) simulated from Gaussian copula $C_{0.99}$ (y-axis cut off). Right: Histogram of PnL distribution for considered trade; (X, Y) simulated from F_X and historical relation (y-axis cut off).

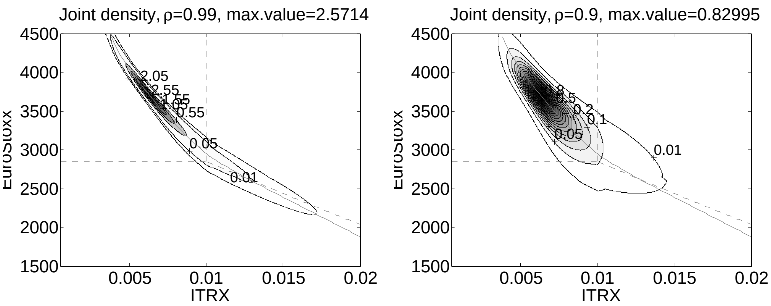

Fig. 9: Joint density of (X, Y) for $\rho = 0.99$ and $\rho = 0.9$ as contour plots, with maximal density value given in title. For comparison, the perfect negative relation and the borders of the PnL regions are also displayed. The major part of the probability mass is assigned to the PnL-neutral region, centered around the current levels, and spreading out with decreasing ρ.

Remark 4.1 (Implicit tail approximations)
As illustrated in Figure 9, the joint distribution assigns most of the probability mass on the explicit relation and inside the PnL-neutral region close to the current levels. Therefore the overall probability of ending up outside of the PnL-neutral region is quite low, and most of the simulated pairs (X, Y) generate no profit or loss. However, for an analysis of the considered trade it is crucial to focus on the probability mass in the tails, where one expects a non-zero PnL. The found relation $g(.)$ between both indices is retrieved completely from option data on both indices. In particular, the tails of F_X and F_Y, are extrapolated by the Buchen-Kelly method outside the observed moneyness range of the given option data. This implicit extrapolation carries over to a prediction of $g(X)$, when X is very large (or low). Consequently, the degree of model uncertainty within the present analysis is significant and has to be kept in mind.

4.1 Sensitivity analysis

Looking again at (3), the PnL at T, one finds that this is heavily dependent on the choice of the strike prices K and s^K. Figure 10 illustrates that many different PnL profiles can be achieved with the available strike ranges, assuming perfect negative dependence between (X, Y). Note that a special case is $K \approx g(s^K)$, where the PnL outcome is almost flat.

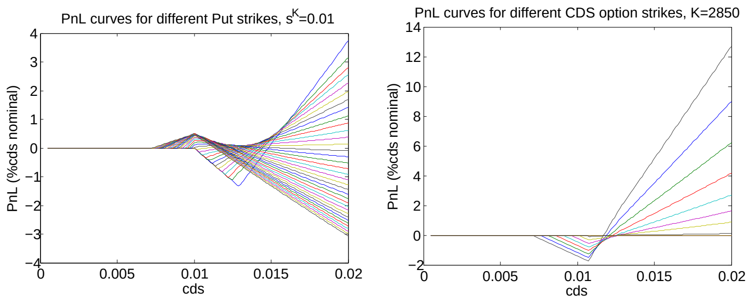

Fig. 10: Possible PnL curves as functions of ITRX spread, assuming a perfect negative relation. Left: ITRX strike s^K fixed. Right: EuroStoxx strike K fixed.

Considering the characteristics of the PnL distribution obtained from samples (X, Y) generated from C_ρ, $\rho < 1$, we find the following: Regardless of the chosen strikes, the expectation of the PnL is zero. For ρ increasing from 0 to 1, the standard deviation and the upper tail percentiles of the PnL distribution (corresponding to gains) decrease, whereas both VaR and CVaR of the respective trade increase. This is in line with the "spreading out" of probability mass observable for decreasing copula parameters ρ, cf. Figure 9. Figure 11 illustrates the behavior of VaR/CVar and the 99-/95-/90-percentiles of the PnL distribution of our considered trade with respect to changes in the copula parameter ρ: For $\rho = 0.9$ for example, in 1% of all cases one loses $\approx 1.3\%$ of the CDS nominal or more. On the gains side respectively, in 1% of all cases one earns $\approx 1.7\%$ of the CDS nominal or more. We found that for the considered trade, the empirical PnL dis-

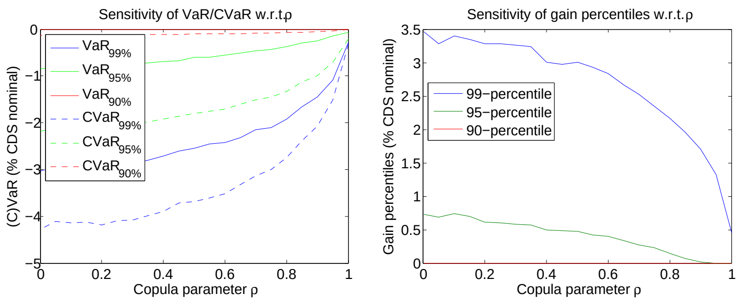

Fig. 11: Sensitivity of VaR/CVaR (left) and 99-/95-/90-percentiles of PnL distribution (right) with respect to copula parameter ρ for our trade example ($K = 2850$, $s^K = 100$ bps).

tribution is asymmetric. This generalizes in the following way: When selling[6] ITRX options, the PnL distribution is asymmetric for $K \neq g(s^K)$, where the asymmetry gets more pronounced for larger values of $|s^K - g^{-1}(K)|$. For $g^{-1}(K) > s^K$, as in our trade example, one has a higher probability for losses than gains, but possible losses are mostly moderate, whereas the probability for gains is smaller, but gains are potentially larger (expressed in % of CDS nominal). For $g^{-1}(K) < s^K$, the opposite holds.

[6] For the case of buying ITRX options against SX5E Puts, the opposite statements hold.

In the light of these characteristics of the PnL distribution, each trader can choose the strikes K, s^K to reflect his or her preferred PnL profile and set up the trade accordingly.

5 Conclusion

We presented an approach for modeling the joint distribution of two correlated indices (X, Y) using market information. Margins are extracted from quoted option prices using a robust method that combines the benefits of having a smooth density with closed-form evaluation of the inverse distribution function. The joint distribution is then obtained by combining the extracted margins with a bivariate parametric copula model C_ρ, which desirably exhibits exchangeability and radial symmetry, and contains the comonotonicity copula as a limiting case. With a quick inversion method for the marginal distributions at hand, one can generate a large number of sample pairs (X, Y) from copula samples $(U_1, U_2) \sim C_\rho$ using the quantile transform, which can be used, for example, to identify and analyze trading opportunities involving options on two closely correlated financial indices.

References

A. Armstrong, M. Rutkowski, Valuation of credit default swaptions and credit default index swaptions, *International Journal of Theoretical and Applied Finance* **12** (2009) pp. 1027–1053.

D. Brigo, M. Morini, No-armageddon measure for arbitrage-free pricing of index options in a credit crisis, *Mathematical Finance* **21:4** (2011) pp. 573–593.

P.W. Buchen, M. Kelly, The maximum entropy distribution of an asset inferred from option prices, *Journal of Financial and Quantitative Analysis* **31** (1996) pp. 143–159.

P. Carr, V. Linetsky, A jump to default extended CEV model: An application of Bessel processes, *Finance and Stochastics* **10** (2006) pp. 303–330.

H.-G. Fung, G.E. Sierra, J. Yau, G. Zhang, Are the U.S. stock market and credit default market related? Evidence from CDX indices, *Journal of Alternative Investments* **11:1** (2008) pp. 43–61.

J. Gatheral, A. Jacquier, Arbitrage-free SVI volatility surfaces, *Quantitative Finance* **14** (2014) pp. 59–71.

Y. Liu, P. Jäckel, Options on credit default index swaps, *Wilmott Magazine Technical Article* (2005)

J.-F. Mai, M. Scherer, Simulating copulas, *Imperial College Press* (2012).

J.-F. Mai, S. Schenk, M. Scherer, Analyzing model uncertainty via distortion of the stochastic root: a Dirichlet prior approach, *Working paper* (2015).

R. Martin, A CDS option miscellany, *Working paper* (2012).

C. Meyer, The bivariate normal copula, *Working paper* (2009).

R.B. Nelsen, An introduction to copulas, 2^{nd} edition, *Springer* (2006).

C. Neri, L. Schneider, Maximum entropy distribution of an asset inferred from option prices, *Finance and Stochastics* **6** (2012) pp. 293–318.

C. Neri, L. Schneider, A family of maximum entropy densities matching call option prices, *Applied Mathematical Finance* **20** (2013) pp. 548–577.