Pricing and hedging basket credit derivatives in the Gaussian copula

The static assumptions of the Gaussian copula model have long presented an obstacle to dynamic hedging of credit portfolio tranches. Here, ***Jean-David Fermanian*** *and* ***Olivier Vigneron*** *combine the copula with a spread diffusion to derive hedging error as proportional to the gamma times the difference between implied and realised spread correlations, analogous to the classic Black-Scholes result. Consequently any basket derivative has a unique price and can be replicated using credit default swaps*

Since the seminal paper of Li (2000), the Gaussian copula model has become the market standard of the structured credit derivatives world. By postulating a correlation structure for the default times of each issuer directly, this model allows the practitioner to price collateralised debt obligation (CDO) payouts and gives a simple solution for their risk management. However, the choice of the pairwise correlations appears largely arbitrary and lacks a clear link with a concept of realised correlation that could be estimated from credit default swap (CDS) prices. After the dramatic collapse in the value of numerous CDOs during the credit crunch, some people have even called Li's formula 'the formula that killed Wall Street'. A common argument against the Gaussian copula model is its lack of replication arguments, contrary to the Black-Scholes model. Option pricing theory relies heavily on the concepts of implied and realised volatilities. Indeed, in the idealised setting of the Black-Scholes framework, both these numbers should be equal and a continuous delta-hedging strategy marked at this volatility perfectly replicates the final payout.

By studying simple basket derivatives, we show how the implied copula correlation should be linked to the realised correlations of CDS spread variations as well as their realised levels of volatility. This is a new result and the main finding of our article. To do so, we apply results of Fermanian & Vigneron (2009) in the simple framework of first-p-to-default securities. Our starting point is to evaluate analytically the hedging error for a simple first-to-default basket on two names. We then exhibit the single spread dynamics that allow a perfect replication, under the assumption that no jump-to-default events occur. We define and solve for the breakeven correlation level, and link it to the volatility and correlation of the credit spreads. We then generalise these results to first-p-to-default and $n \geq 2$ names, and derive the shape of the base correlation skew. Finally, we give some examples of the generated correlation skews.

To simplify the picture and highlight the main results, we will only consider European-style payouts with a given maturity T, no running premiums (only upfront payments), identical recovery rates R and zero interest rates.

The hedging error of a first-to-default basket

In the classical Black-Scholes model, the hedging error of replicating a stock option between t and $t + dt$ is proportional to the gamma of the underlying option multiplied by the difference between the implied and the realised variances of stock returns (see El Karoui, Jeanblanc-Picqué & Shreve, 1998). With some obvious notation, it is:

$$d\pi_t = \frac{S_t^2\Gamma_t}{2}\left(\sigma^2_{implied} - \sigma^2_{real,t}\right) \tag{1}$$

Thus, to replicate the option payout, the implied volatility should be chosen as the average of the future realised instantaneous volatilities up to maturity.

With credit derivatives, no such result has been stated until now, and it is the purpose of this section to calculate the hedging error that results from replicating the payout of a basket credit derivatives with the Gaussian copula model. Here, we consider a basket of $n = 2$ risky names and the simplest structured credit derivative based on this basket, a first-to-default, whose payout is:

$$\psi_{FtD,T}\left(\vec{\delta}\right) = (1-R)\min\left(\delta_1 + \delta_2, 1\right)$$

where δ_j is the default indicator function of name j, that is, $\delta_j = \mathbf{1}(\tau_j \leq T)$. Its t-value will be denoted by $\psi_{FtD,t}(\vec{\delta})$ or even ψ_t.

For a given name j, a key quantity is its current (risk-neutral) survival probability up to T, denoted by $Q_{jt} := Q(\tau_j > T|\mathcal{F}_t)$. We assume it can be observed in the market through CDSs written on j. In other words, the quantities Q_{jt}, $j = 1, 2$ and $t \in [0, T]$ can be seen as our traded underlyings. At time t, the information set $\mathcal{F}_t$ contains all the available market information (that is, current and past credit spreads, past default events, etc).

The first-to-default will be priced under the usual Gaussian copula model. Conditional on a single common factor X, both default events are independent. The law of X is standard normal. Let ϕ and Φ be its density and its distribution function respectively. The (risk-neutral) conditional default probability of j at time $t \leq T$ is defined by:

$$p_{j|x} = Q\left(\tau_j \leq T | X = x, \mathcal{F}_t\right) = \Phi\left(\frac{\bar{\Phi}^{-1}\left(Q_{jt}\right) - \beta_j x}{\sqrt{1-\beta_j^2}}\right) \tag{2}$$

where $\bar{\Phi} = 1 - \Phi$ and β_j is the pricing 'beta' factor of j. As an exercise, it can be checked that $\psi_{FtD,t}$, the current price of this first-to-deafult, depends only on the pricing factors β_j, $j = 1, 2$, through their product $\rho_{12}^{model} = \beta_1\beta_2$. To lighten the notation, we set:

$$d_j := d_j(T, x) = \frac{\bar{\Phi}^{-1}(Q_{jt}) - \beta_j x}{\sqrt{1-\beta_j^2}}$$

which implies $p_{j|x} = \Phi(d_j)$ and $q_{j|x} = 1 - p_{j|x} = \bar{\Phi}(d_j)$.

The standard delta-hedging strategy consists of neutralising the value of the structured product against spread moves using individual CDSs as our hedging instruments. Since we have assumed they are fully upfront and European-style, the t-price of a T-maturity CDS on j is simply $(1 - R)(1 - Q_{jt})$. The hedging ratio (or 'delta') of the first-to-default is simply the derivative of the first-to-default price with respect to the quantities Q_{jt} divided by minus the loss given default $-(1 - R)$. They are denoted by Δ_{jt}. Between two successive times, the hedging error is then:

$$d\pi_t = d\psi_{FtD,t} + \Delta_{1t}(1-R)dQ_{1t} + \Delta_{2t}(1-R)dQ_{2t}$$

Next, we assume dynamics for the survival probabilities (which we will refer to as 'spread dynamics') under the real-world probability measure: $dQ_{jt} = \mu_{jt}dt + \sigma_{jt}dW_{jt}$ for $j = 1, 2$, with $E[dW_{1t}dW_{2t}] =: \rho_{12}^S dt$, a constant level of spread correlation.

It is apparent that there is no effect of time in this delta-hedging strategy, except through the updates of the survival probabilities. In other words, the first-to-default 'theta' is zero.[1] By Itô's lemma, we deduce the hedging error:

$$d\pi_t = \frac{1}{2}\frac{\partial^2\psi_t}{\partial^2 Q_{1t}} d\langle Q_{1t}, Q_{1t}\rangle + \frac{1}{2}\frac{\partial^2\psi_t}{\partial^2 Q_{2t}} d\langle Q_{2t}, Q_{2t}\rangle + \frac{\partial^2\psi_t}{\partial Q_{1t}\partial Q_{2t}} d\langle Q_{1t}, Q_{2t}\rangle$$

To be perfectly hedged, all these 'gamma-type' terms have to cancel each other out, for every realisation of the underlying spreads. Under the Gaussian copula model, this can be achieved with some particular spread dynamics only (see below).

By an integration by parts, this can be rewritten:

$$d\pi_t = \frac{(1-R)A_{12,t}}{2}\left(\gamma_{1t}^2 + \gamma_{2t}^2\right)\left[\rho_{12}^{model} - \frac{2\gamma_{1t}\gamma_{2t}}{\gamma_{1t}^2 + \gamma_{2t}^2}\rho_{12}^S\right]dt \quad (3)$$

where $\gamma_{it} := \sigma_{it}/\phi(\Phi^{-1}(Q_{it}))$ is a 'volatility-type' quantity, and where we have set:

$$A_{12,t} := \int \frac{\phi(d_1)\phi(d_2)\phi(x)}{\sqrt{1-\beta_1^2}\sqrt{1-\beta_2^2}} dx$$

Like the first-to-default price itself, the latter expression depends only on $\rho_{12}^{model} = \beta_1\beta_2$, and not on the way we split this product into the two factors β_1 and β_2. The proof of (3) can be found in Fermanian & Vigneron (2009). It is based crucially on an analytical property of ϕ, that is, for all x:

$$\phi(d_j)\phi(x) = \phi\left(\frac{x - \beta_j\bar{\Phi}^{-1}(Q_{jt})}{\sqrt{1-\beta_j^2}}\right)\phi\left(\bar{\Phi}^{-1}(Q_{jt})\right) \quad (4)$$

The parallel of the formula (3) with (1) is striking. The cross-gamma term $A_{12,t}$ above has the same status as the gamma factor Γ_t of a vanilla option. In brackets, we observe differences between the pricing correlations ρ_{12}^{model} and the realised spread correlation ρ_{12}^S, but rescaled by a term that captures the dispersion of spread volatilities. In other words, the right correlation level that should be used in the Gaussian copula model is a mix of spread correlations and (instantaneous) spread volatilities.

The previous result is available when no sudden default event occurs between t and $t + dt$. In other words, we exclude the occurrence of jump-to-default that would not be anticipated in the market. It may appear a strong assumption, but most default events occurred after significant widenings of the credit spreads of the defaulted entity historically.

Breakeven correlations of a first-to-default

From these hedging errors, we can deduce the flat pricing correlation that would neutralise the portfolio mark-to-market variations between two successive dates. It is called the instantaneous breakeven correlation. It depends on the current spread levels, the instantaneous volatilities and correlation of the credit spread

[1] This result would be slightly different if we reasoned with hazard rates instead of CDS prices. In this case, the theta of the unhedged first-to-default would be non-zero while the delta-hedged first-to-default theta would still be zero

moves. With the previous notations, the breakeven correlation of a two-name first- to-default between t and $t + dt$ is:

$$\rho^2_{BE,t} = \frac{2\gamma_{1t}\gamma_{2t}}{\gamma^2_{1t} + \gamma^2_{2t}}\rho^S_{12}$$

The factors γ_{it} depend on the current survival probabilities, which means that, in general, the breakeven correlation level is dependent on the spread trajectories. However, a suitable choice for the single spread dynamics will allow a perfect hedging with a constant breakeven correlation. Our first-to-default can be perfectly and continuously hedged if and only if:

- The real spread dynamics are given by:

$$dQ_{it} = \bar{\sigma}_i \xi_t \phi\left(\Phi^{-1}\left(Q_{it}\right)\right) dW_{it} + \mu_{it} dt \tag{5}$$

where $\bar{\sigma}_i$ denotes some positive constant and ξ a deterministic function such that $\lim_{t\to T}\int_0^t \xi^2_u du = +\infty$. Moreover, under the historical measure, W_{it} is a Brownian motion and μ_{it} is a general $\mathcal{F}$-measurable process.

- The pricing correlation satisfies:

$$\rho^{model}_{12} = \left(\frac{2\bar{\sigma}_1\bar{\sigma}_2}{\bar{\sigma}^2_1 + \bar{\sigma}^2_2}\right)\rho^S_{12} \tag{6}$$

It is worth noting that the equations (5) describing the spread dynamics can be integrated:

$$Q_{it} = \Phi\left(\exp\left(\frac{\bar{\sigma}^2_i \int_0^t \xi^2_u du}{2}\right)\Phi^{-1}\left(Q_{i0}\right) + \int_0^t \exp\left(\frac{\bar{\sigma}^2_i \int_u^t \xi^2_s ds}{2}\right)\left\{\bar{\sigma}_i \xi_u dW_{iu} + \frac{\mu_{iu}}{\phi\left(\Phi^{-1}\left(Q_{iu}\right)\right)} du\right\}\right)$$

Under the risk-neutral measure Q, the equation holds with μ_{it} zero. As expected, we check that $\lim_{t\to T} Q_{it} \in \{0, 1\}$.

Interpretation in terms of extended one-period Merton models

The simplest way to find the spread dynamics induced by the Gaussian copula model is to use a one-period structural model. Let us assume an asset value process $(A_{it})_{t\in[0,T]}$ for every name i. In the simplest Merton model[2], name i defaults if and only if its asset value at maturity is smaller than a threshold b_{iT}. Equivalently, the survival event of i at horizon T is defined by $(\tau_i > T) = (A_{iT} > b_{iT})$. If asset values follow Brownian motions, it can be checked easily that the associated (risk-neutral) dynamics of Q_{it} are:

$$dQ_{it} = \frac{1}{\sqrt{T-t}}\phi\left(\Phi^{-1}\left(Q_{it}\right)\right) dW_{it} \tag{7}$$

where $W_{it} = A_{it}$ are correlated Brownian motions. Clearly, the latter dynamics are a particular case of processes that follow (5), with $\bar{\sigma}_i = 1$ and $\xi_t = 1/\sqrt{T-t}$.

The more general dynamics (5) can be recovered with suitable asset value processes:

$$A_{it} := \int_0^t \eta_i(u) dW_{iu} \tag{8}$$

where η_i is a deterministic function such that $\int_0^T \eta_i^2 < +\infty$. The processes W_{it} are still correlated Brownian motions. We can set:

$$\eta^{(\alpha)}_i(t) := \bar{\sigma}_i\left(1 - \frac{t}{T}\right)^{\alpha} \exp\left(\frac{\bar{\sigma}^2_i T}{2(1+2\alpha)}\left[\left(1-\frac{t}{T}\right)^{1+2\alpha} - 1\right]\right) \tag{9}$$

for some constants $\alpha > -0.5$ and $\bar{\sigma}_i > 0$, which simplifies to:

$$\eta_i(t) := \bar{\sigma}_i\left(1 - \frac{t}{T}\right)^{\bar{\sigma}^2_i T/2 - 1/2} \tag{10}$$

for $\alpha = -0.5$. As previously, default events before T occur if and only if A_{iT} is below some boundary. Simple calculations show that the two latter specifications generate spread dynamics of the type (5), with various functions ξ. A structural model with the latter asset value processes provides richer survival probability dynamics in which the 'volatility-type' parameter $\bar{\sigma}_i$ can be interpreted as an 'acceleration of time'.

Therefore, the Gaussian copula pricing formula (2) can be recovered from the specification of any one-period Merton models as given by equation (8). Particularly, the pricing/breakeven correlations are those between asset values only when these asset values follow Brownian motions ($\eta_i = 1$ for every name i). More generally, for instance under (9) and (10), 'volatility' effects appear. Then, breakeven correlations are functionals of these terms and of the correlations between the underlying asset values (see (6)).

Note that a digital worst-of put written on several stocks is formally identical to a first-to-default, if we identify asset values and stock prices (see Fermanian & Vigneron, 2009). Then, as a sub-product, we have exhibited a consistent framework to price and hedge these equity derivatives perfectly.

Extension to first-*p*-to-default derivatives and *m*-factor Gaussian copula models

The results for first-to-default baskets on two names can be extended significantly. We now consider a basket of n names and a first-p-to-default payout that is defined by:

$$\psi^{(p)}_T\left(\bar{\delta}\right) = (1-R)\min\left(\sum_{j=1}^n \delta_j, p\right)$$

and where $p \in \{1, \dots, n\}$. It should be stressed that first-p-to-default derivatives are the basic building blocks for all other European-style payouts, which depend on default events only. They can also be seen as the direct equivalent of the base tranches [0, x%] of a synthetic CDO.

The previous first-p-to-default will be priced under an m-factor Gaussian copula model. It means that there exists a 'systemic' standard Gaussian random vector $\mathbf{X}$ in $\mathbb{R}^m$, such that the default events are independent conditional on $\mathbf{X} = \mathbf{x}$. Then, the conditional default probability of j is:

$$p_{j|\mathbf{x}} = Q\left(\tau_j \le T \,|\, \mathbf{X} = \mathbf{x}, \mathcal{F}_t\right) = \Phi\left(\frac{\bar{\Phi}^{-1}\left(Q_{jt}\right) - \beta'_j\mathbf{x}}{\sqrt{1 - \left|\beta_j\right|^2}}\right) \tag{11}$$

where β_j is a vector in $\mathbb{R}^m$, the pricing 'beta' factor of j. As previously, we set:

$$d_j := d_j(T, \mathbf{x}) := \frac{\bar{\Phi}^{-1}\left(Q_{jt}\right) - \beta'_j\mathbf{x}}{\sqrt{1 - \left|\beta_j\right|^2}}$$

which implies $p_{j|\mathbf{x}} = \Phi(d_j)$ and $q_{j|\mathbf{x}} = 1 - p_{j|\mathbf{x}} = \bar{\Phi}(d_j)$. Thus, we are able to evaluate the joint law of default events before T, and to value every first-p-to-default. Its current present value is denoted by $\psi^{(p)}_t$.

As for the first-to-default, the theta is zero. Moreover, between

[2] *In contrast to more complex structural models, like first-hitting times*

t and $t + dt$, the change in value of a delta-hedged first-p-to-default is given by:

$$d\pi_t^{(p)} = d\psi_t^{(p)} + \sum_{j=1}^{n} \Delta_{jt}(1-R)dQ_{jt}$$
$$= \frac{1}{2}\sum_{i,j}\frac{\partial^2 V_t^{(p)}}{\partial Q_{it}\partial Q_{jt}} d\langle Q_{it}, Q_{jt}\rangle$$

By a tricky integration by parts, Fermanian & Vigneron (2009) obtained:

$$d\pi_t^{(p)} = \frac{(1-R)}{2}\sum_{i<j}\left(\gamma_{it}^2 + \gamma_{jt}^2\right)A_{ij}\left[\beta_i'\beta_j - \frac{2\gamma_{it}\gamma_{jt}}{\gamma_{it}^2+\gamma_{jt}^2}\rho_{ij}^S\right]dt \quad (12)$$

where we have set:

$$A_{ij} := \sum_{\substack{\delta_k;\, k\le n,\, k\ne i,j \\ \Sigma_k \delta_k = p-1}} \int \prod_{k\ne i,j} p_{k|\mathbf{x}}^{\delta_k} q_{k|\mathbf{x}}^{1-\delta_k} \frac{\phi(d_i)\phi(d_j)\phi(\mathbf{x})}{\sqrt{1-|\beta_i|^2}\sqrt{1-|\beta_j|^2}} d\mathbf{x}$$

As before in the two-name derivation, ρ_{ij}^S denotes the correlation between the Brownian motions that drive the spread moves of i and j in the real world.

Again, we should note the parallel of the formula (12) with (1). The terms A_{ij} above have the same status as cross-gamma factors. In brackets, we observe differences between the pricing correlations $\beta_i'\beta_j$ and the realised spread correlations ρ_{ij}^S. But correlation levels do not provide the full picture. Credit spread volatilities matter too, through the additional factors γ_{it}.

From these hedging errors, we deduce the 'flat correlation' levels (identical for all the names) that neutralise the portfolio mark-to-market variations between two successive dates. They are still called instantaneous breakeven correlations. They depend on the current spread levels, the instantaneous volatilities and the correlations between credit spread moves. With the previous notation, the breakeven correlation between t and $t + dt$ is:

$$\rho_{BE,t}^2 = \sum_{i<j}\frac{2A_{ij}\gamma_{it}\gamma_{jt}}{\sum_{i<j}A_{ij}\left(\gamma_{it}^2+\gamma_{jt}^2\right)}\rho_{ij}^S =: \sum_{i<j} w_{ijt}\rho_{ij}^S$$

where all the 'weights' w_{ijt} belong to [0, 1] but do not add up to one, except when all the γ_{it} are equal. Actually, since the previous weights w_{ijt} depend on $\rho_{BE,t}$ itself, we need to solve an implicit equation to find the breakeven correlation.

The breakeven correlation levels depend on the 'strike' p through the coefficients A_{ij}. The latter coefficients depend on all the current survival probabilities Q_{kt}, $k = 1, \ldots, n$ too. Thus, it is unlikely one would find a single constant correlation level that would allow a pricing by replication for all the first-p-to-default. To do so, we need to allow for correlation heterogeneity and assume the spread dynamics exhibited in the previous section. Indeed, it is proved in Fermanian & Vigneron (2009) that any structured credit derivative[3] can be perfectly and continuously hedged if and only if:

- The real spread dynamics are given by (5).
- We can find pricing factors β_i such that, for every couple (i, j), $i \ne j$:

$$\beta_i'\beta_j = \left(\frac{2\bar{\sigma}_i\bar{\sigma}_j}{\bar{\sigma}_i^2+\bar{\sigma}_j^2}\right)\rho_{ij}^S \quad (13)$$

The latter condition seems to be strong. Even if it can be checked that $\Sigma := [2\bar{\sigma}_i\bar{\sigma}_j\rho_{ij}^S/(\bar{\sigma}_i^2 + \bar{\sigma}_j^2)]$ is a correlation matrix (that is, positive definite with ones on the main diagonal), this matrix is not necessarily induced by a m-factor correlation structure. Actually, we have seen that, if Σ is given by an m-factor correlation structure, it is necessary to price the structured product with an m-factor Gaussian copula model.

We can apply the latter result to our first-p-to-default derivative. If the spread dynamics in the real world are given by (5), then the previous quantities γ_{it} are now the product of a constant $\bar{\sigma}_i$ and a deterministic function of time ξ_t. Thus, the instantaneous breakeven correlation of a first-p-to-default is now:

$$\rho_{BE,t}^2 = \sum_{i<j}\frac{2A_{ij}\bar{\sigma}_i\bar{\sigma}_j}{\sum_{i<j}A_{ij}\left(\bar{\sigma}_i^2+\bar{\sigma}_i^2\right)}\rho_{ij}^S \quad (14)$$

Note that this 'flat' breakeven correlation is a true weighted average of the breakeven beta factors β_i:

$$\rho_{BE,t}^2 = \sum_{i<j}\frac{A_{ij}\left[\bar{\sigma}_i^2+\bar{\sigma}_j^2\right]}{\sum_{i<j}A_{ij}\left(\bar{\sigma}_i^2+\bar{\sigma}_j^2\right)}\beta_i'\beta_j$$

1 Breakeven 'beta factors' with varying intensities and spread correlations as a function of the strike p (detachment point)

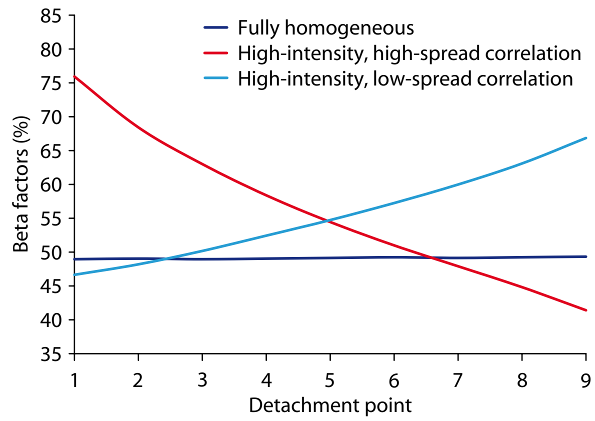

Implied correlation skew shapes

To illustrate the results above, let us assume there exists a basket of $n = 10$ underlying names, and that their spread dynamics follow:

$$dQ_{it} = \frac{\bar{\sigma}_i}{\sqrt{T-t}}\phi\left(\Phi^{-1}(Q_{it})\right)dW_{it} \quad (15)$$

where the Brownian motions are one-factor correlated: $E[dW_{it}dW_{jt}] = \beta_i^S\beta_j^S dt$, when $i \ne j$. In this experiment, we will consider $T = 5$ years maturity first-p-to-default, where p goes from one to nine. We have generated a particular trajectory of the survival probabilities Q_{it}, $i = 1, \ldots, 10$, $t \in [0, T_1]$, $T_1 < T$.[4] We have calculated the empirical corresponding breakeven correlations. By construction, they are the flat correlation levels that cancel the profit and loss variations of a delta-hedged first-p-to-default between $t = 0$ and $t = T_1$. We can play with three sets of parameters:

[3] *European, without running premiums and whose payout depends on the occurrence of default events only*
[4] *T_1 is here equal to 200 days*

2 Breakeven 'beta factors' for a heterogeneous basket along one (two) dimension(s)

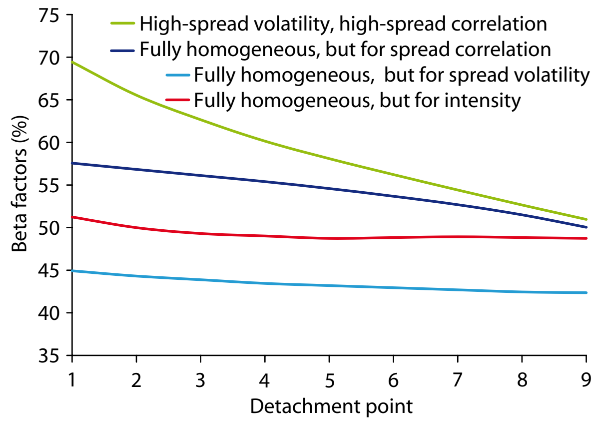

3 Breakeven 'beta factors' for a heterogeneous basket along all three dimensions

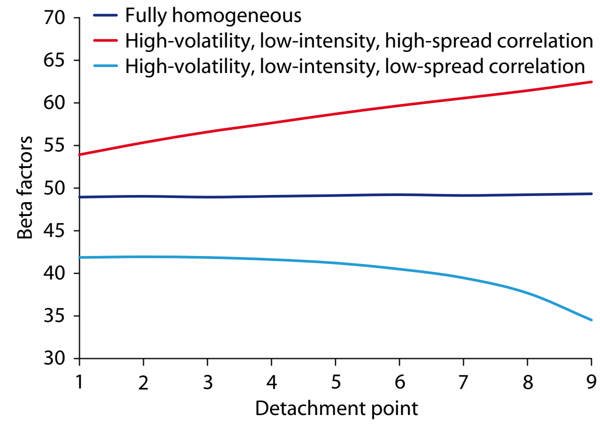

- The spot survival probabilities $Q_{i0}(T)$, or equivalently the default intensities levels λ_{i0} at $t = 0$.
- The spread correlation factors β_i^S.
- The volatility-like terms $\bar{\sigma}_i$.

We start with the fully homogeneous basket case, in which the spot default probabilities, the spread correlation factors and the spread volatilities are the same for all the names. Here, we have chosen $\lambda_{i0} = 5\%$, $\beta_i^S = 50\%$ and $\bar{\sigma}_i = 50\%$ for all i. If the trajectories, $(Q_{it}(T))_{t\in[0,T_1]}$, $i = 1, \ldots, n$, we have simulated are not 'too volatile', then the previous (trajectory-dependent) factors A_{ij} should not be very different. In this case, the theoretical breakeven correlations is close to 50% for all strikes (that is, for all first-p-to-default, $p = 1, \ldots, 9$) (see equation (14)). Indeed, in figure 1, we do not observe any skew in this fully homogeneous case.

We move away from this core case to introduce heterogeneity among the 10 underlying names. We choose five values for each parameter, each to be used twice in the basket, and consider them in increasing or decreasing order. Initial default intensities are set at 2%, 3%, 5%, 8% and 10%; spread volatilities at 22%, 33%, 50%, 75% and 113%; and spread 'beta' factors at 22%, 33%, 50%, 75% and 99%. By combining the increasing or decreasing versions of the parameter sets, we build different heterogeneous portfolios.

Generally, heterogeneous spread correlations and spread dispersion produce breakeven correlation skews. Figure 1 shows the behaviour of these skews in the presence of opposite extremes in the intensities and spread correlations. Upward-sloping skews srise when higher correlations are paired with tight spread names and lower correlations with wide names, as commonly observed in the market. Moreover, downward-sloping skews arise if the reverse is true.

Figure 2 shows the effect of the introduction of heterogeneity in a single name, all other things being equal. This generates slightly downward-sloping skews. The skews can be made steeper by moving two types of parameter together, for instance by combining the larger spread volatilities with the larger spread correlations.

In figure 3, we play with the three dimensions simultaneously. For instance, by combining the increasing volatilities and spread correlations with decreasing default intensities, we recover the standard upward-sloping skew. By experimenting with different combinations of spread volatilities, correlations and intensities, it is possible to generate various shapes of breakeven correlation curves, illustrating the flexibility of the model.

Conclusion

We have shown that, under the 'no sudden default' assumption, the Gaussian copula model can be seen as a pricing model by replication. This can be achieved if and only if the underlying spread dynamics belong to a particular family and if the pricing correlations are defined as in equations (6) and (13). We have linked explicitly these parameters with the correlations and volatilities of spread moves. These results are available in the general case of first-p-to-default derivatives and m-factor Gaussian copula models. This approach can be nicely deduced from a Merton-style approach. We believe the advantages and shortcomings of the current standard model in the CDO world are understood more in depth now. There are numerous avenues for further research and extensions, such as term structures of default probabilities, recovery heterogeneity and/or randomness, and alternative factor models. ■

Jean-David Fermanian is professor of finance and statistics at École Nationale de la Statistique et de l'Administration (Ensae) in Paris. Olivier Vigneron is global co-head of credit hybrids trading and risk management at JP Morgan in London. They would like to thank colleagues and former colleagues at BNP Paribas and JP Morgan for numerous and valuable comments. All opinions and any errors are those of the authors, and do not reflect those of JP Morgan or Ensae. Email: Jean-David.Fermanian@ensae.fr, olivier.x.vigneron@jpmorgan.com

References

El Karoui N, M Jeanblanc-Picque and S Shreve, 1998
On the robustness of the Black and Scholes formula
Mathematical Finance 8, pages 93–126

Fermanian J-D and O Vigneron, 2009
On break-even correlation: the way to price structured credit derivatives by replication
Available at *www.defaultrisk.com*

Li D, 2000
On default correlation: a copula function approach
Journal of Fixed Income 9(4), pages 43–54