Beyond the Gaussian Copula: Stochastic and Local Correlation

X. Burtschell[1] J. Gregory[2] & J-P. Laurent[3]

October 2005

Abstract

We consider "copula skew models" that account for the correlation smile in the pricing of synthetic CDO tranches. These can be viewed as stochastic or local correlation models and are extensions of the well-known one factor Gaussian copula model. We analyse these models through their conditional default probability distributions and marginal compound correlations. We also give some examples of using a particular stochastic correlation model to fit the market, illustrating the stability of the parameters over time.

JEL Classification: C 31, G 13
Key words: default risk, CDOs, correlation smile, factor copulas, stochastic correlation, local correlation.

Introduction

The factor copula approach has proved to be a powerful tool for pricing CDOs within a semi-analytical framework (see Gregory & Laurent [2003], Andersen, Sidenius & Basu [2003], Hull & White [2004], Laurent & Gregory [2005]). While the factor approach might also be used in the framework of intensity or structural models, it has been predominantly coupled with a copula approach. Though copula models fail to provide satisfactory dynamics of credit spreads and exhibit various kinds of unsatisfactory time instability, they usually allow independent specification of the dependence structure between the default times and the marginal credit curves, the latter being considered as market observables (single name CDS or index CDS).

Thus, despite the previous drawbacks, the one factor Gaussian copula (see Li [2000]) has become a market standard for the pricing of CDOs. The development of a liquid market of standardized CDO tranches

[1] BNP Paribas, xavier.burtschell@bnpparibas.com,

[2] Barclays Capital, jon.gregory@barclayscapital.com,

[3] ISFA Actuarial School, University Claude Bernard of Lyon & BNP Paribas, laurent.jeanpaul@free.fr, http://laurent.jeanpaul.free.fr

The second author was at BNP Paribas when this paper was written. The authors thank C. Donald, M. Leeming and R. Sharp for computational assistance and helpful discussions. They also thank A. Conze and M. Musiela for comments, M. Shchetkovski and F. Rhazal for pointing out an error in the first draft. The usual disclaimer applies.

has shown the Gaussian copula model to be inadequate and there has in the last year been a build up of literature on correlation skew modelling and a reliance by market practioners on base correlation approaches. The major developments and movements in the tranche market during 2005 and the already well-known deficiencies of the market standard modelling approach have led to even stronger calls than before for second generation models for CDO pricing based on realistic dynamics of market variables (i.e. credit spreads). This is clearly an important direction but, perhaps not surprisingly, researchers have tried to extend gradually the standard model to incorporate some skew features. The aim of this paper is clearly related to the latter approach, where we will give an overview of the general framework of "copula skew models", explain some connections between different ideas and describe results using one such model. This approach can be considered as complementary or an alternative to the base correlation approach which is currently the predominant (and perhaps sole) method used to industrially price and risk manage CDO and CDO squared tranches. We see this as a next step on a long path leading to next generation models, which will exist outside the copula framework.

Whereas in the one factor Gaussian copula correlation is a deterministic parameter, further extensions involve stochastic correlation (Burtschell et al. [2005], Schloegl [2005]) or stochastic risk exposure (see Andersen and Sidenius [2005]). The risk exposures may or may not be associated with a factor structure and may or may not be factor dependent. When the correlation is stochastic and independent of the factor, we will talk of a *stochastic correlation* model. When the correlation depends upon the factor, we will talk of a *local correlation* model (see Turc et al. [2005]). This terminology parallels the one used for equity derivatives.

As can be seen in Burstchell et al. [2005], Hull and White [2005], Sharke [2005], the distributions of conditional default probabilities are the main drivers of CDO prices within factor copula models. Such distributions also correspond to large portfolio loss distributions. We will thus further provide some emphasis on these quantities.

The paper is organized as follows. The first section recalls some characteristics of Gaussian and factor copula models and discusses how they may be inverted in order to derive a marginal compound correlation. Section 2 details the case of stochastic correlation models with emphasis a simple approach incorporating idiosyncratic and systemic risk. In section 3, we consider state dependent correlation models with a focus on random factor loading and local correlation models. Section 4 is dedicated to comparisons between the different approaches and a study of a particular stochastic correlation model.

1 Factor copulas and compound correlation

1.1 Gaussian and factor copulas

In the one factor Gaussian copula with flat correlation parameter ρ, we consider some latent variables:

$$V_i = \rho V + \sqrt{1-\rho^2}\bar{V}_i,\ i = 1, \dots, n, \tag{1.1}$$

where $V, \bar{V}_i, i = 1, \dots, n$ are independent Gaussian random variables. We denote by $F(t) = Q(\tau_i \le t)$ the marginal default probability for time horizon t[4] and define default times τ_i of names $i = 1, \dots, n$ by setting

[4]For notational simplicity, we consider here equal credit spreads and a unique correlation parameter ρ (flat correlation).

$\tau_i = F^{-1}(\Phi(V_i))$. The default times are conditionally independent upon V and the conditional default probabilities are given by:

$$p_t^{i|V} = Q(\tau_i \le t \mid V) = \Phi\left(\frac{\Phi^{-1}(F(t)) - \rho V}{\sqrt{1-\rho^2}}\right). \tag{1.2}$$

Tranche	Market	Gaussian
[0-3%]	24.0%	19.3%
[3-6%]	82.5	234.7
[6-9%]	26.5	82.0
[9-12%]	14.0	32.9
[12-22%]	8.75	6.99
[22-100%]	3.53	0.05

Table 1. Illustration of the inability of a flat correlation model to fit the market tranche quotes. The market quotes are as of the 30-Aug-05. The correlation in the Gaussian copula model was chosen via a least squares fit to the traded tranche premiums. Note that the [22-100%] tranche is not traded but can be implied from the rest of the capital structure and level of the index.

The previous one factor Gaussian copula model leads to a semi-explicit pricing of CDO tranches but does not match the market prices (see Table 1) and has therefore been extended in various directions. As in Hull and White [2004] or Kalemanova et al. [2005], one can consider other distributions for the factor V and the idiosyncratic risks $\bar{V}_i$, such as Student t or NIG. Other default times models to be discussed below or in Burstchell et al [2005], Mortensen [2005], Willeman [2005] remain in the factor copula framework. There exists a factor V such that default times τ_i, $i = 1, \dots, n$ are conditionally independent upon V. Loss distributions and CDO tranche premiums then depend upon the specification of the conditional default probabilities $Q(\tau_i \le t \mid V) = p_t^{i|V}$. For simplicity, we moreover assume that the distribution of $p_t^{i|V}$ is name invariant (homogeneity assumption) and denote by G the common distribution function[5]. In the case of the Gaussian copula, we can easily derive G as:

$$p \in [0,1] \to G(p) = Q\left(p_t^{i|V} \le p\right) = \Phi\left(\frac{\sqrt{1-\rho^2}\Phi^{-1}(p) - \Phi^{-1}(F(t))}{\mid \rho \mid}\right). \tag{1.3}$$

Assuming a constant recovery rate of δ, the pgf of the aggregate loss $L = (1-\delta)\sum_{i=1}^n 1_{\tau_i \le t}$ is given by:

$$\psi_L(u) = E[u^L] = \sum_{i=1}^n Q(L = i(1-\delta))u^{i(1-\delta)} = E\left[\left(p_t^{\cdot|V} u^{1-\delta} + 1 - p_t^{\cdot|V}\right)^n\right] = \int_0^1 \left(pu^{1-\delta} + 1 - p\right)^n dG(p).$$

Therefore, for $i = 0, \dots, n$,

$$Q(L = i(1-\delta)) = \int_0^1 \binom{n}{i} p^i (1-p)^{n-i} dG(p),$$

[5] for notational simplicity, we omit the time dependence in G.

and the distribution the distribution function of the aggregate loss, G_L is given through:

$$G_L(i(1-\delta)) = Q(L \leq i(1-\delta)) = \sum_{k=0}^{i} Q(L = k(1-\delta)).$$

We also recall another useful result related to large portfolio approximations, firstly stated by Vasicek [1987] in the one factor Gaussian copula framework. In the framework of factor copula models the aggregate loss L converges a.s. to $(1-\delta)p_t^{|V}$ as n tends to infinity. As a consequence, in the case of a zero-recovery rate, G_L converges weakly to G[6]. For large portfolios, we will then make the approximation, $G_L(p) = G\left(\frac{p}{1-\delta}\right)$ for $0 \leq p \leq 1$.

From the previous results, one can readily price CDO tranches. To simplify the exposition, we will further neglect discounting effects. The up-front premium of an $[\alpha, \beta]$ CDO tranche can be obtained from the distribution function of the aggregate loss G_L as: $1 - G_L(\beta) + \int_\alpha^\beta \left(\frac{p-\alpha}{\beta-\alpha}\right) dG_L(p)$. Integrating by parts, the premium can also be written as: $\int_\alpha^\beta \left(\frac{1-G_L(u)}{\beta-\alpha}\right) du$, which can be approximated for large n by $\int_\alpha^\beta \left(\frac{1-G(u/(1-\delta))}{\beta-\alpha}\right) du$.

The distribution function G also allows Monte Carlo simulation of default indicators $1_{\tau_i \leq t}$, $i = 1, \dots, n$. This can be useful for instance for the pricing of CDO squared. We firstly draw some random variable U uniformly distributed over $[0,1]$. Then $G^{-1}(U)$, where G^{-1} denotes the generalized inverse of G, admits distribution function G. We can then simulate $1_{\tau_i \leq t}$, $i = 1, \dots, n$ by drawing independently n Bernoulli random variables with parameter $G^{-1}(U)$.

1.2 marginal compound correlation

Since the market does not price with the Gaussian copula, it has become a standard practice to extract compound or base correlations from market prices[7]. Turc et al. [2005] have recently introduced the concept of *marginal compound correlation* as a limit case of the usual compound correlation of CDO tranches.

Let us consider a mezzanine tranche with attachment-detachment points of α and $\alpha+\eta$. Neglecting discounting effects, the upfront premium is given by $\frac{1}{\eta}\int_\alpha^{\alpha+\eta}(1-G_L(u))\,du$. For "small" η, the payoff corresponds to a digital option on the aggregate loss at the maturity of the tranche, with strike α, while the upfront premium is equal to $1-G_L(\alpha)$[8]. On the other hand, for a large homogeneous portfolio, the distribution function of a large homogeneous portfolio can be approximated by the distribution function of the conditional default probability G. The up-front premium of the above mezzanine tranche can then be approximated by $1 - G\left(\frac{\alpha}{1-\delta}\right)$.

The upfront premium of such a $[\alpha, \alpha]$ mezzanine tranche of a large homogeneous portfolio computed under the one factor Gaussian copula assumption with correlation parameter ρ and neglecting the discounting effects

[6] This means that $G_L(p)$ converges to $G(p)$ for all continuity points of G. In the examples below, G is continuous apart from 0 and 1. On the other hand, since $L \leq 1$ and $p_t^{|V} \leq 1$, $G_L(1) = G(1) = 1$.

[7] Thanks to the theory of stochastic orders, it can be proved that increasing ρ leads to a decrease in the upfront premium of an equity tranche (see Burtschell et al. [2005]). Thus, base correlation is unique whenever it exists, which may not be the case.

[8] Thanks to the right-continuity of G_L.

is given by: $1-\Phi\left(\frac{\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{\alpha}{1-\delta}\right)-\Phi^{-1}(F(t))}{|\rho|}\right)$ (see equation (1.3)). The *marginal compound correlation* for that mezzanine tranche is denoted by $\bar{\rho}(\alpha) \in [0,1]$[9] and is such that:

$$G\left(\frac{\alpha}{1-\delta}\right)=\Phi\left(\frac{\sqrt{1-\bar{\rho}(\alpha)^2}\Phi^{-1}\left(\frac{\alpha}{1-\delta}\right)-\Phi^{-1}(F(t))}{\bar{\rho}(\alpha)}\right). \tag{1.4}$$

$\bar{\rho}(\alpha)$ is the correlation parameter to be put in the Gaussian copula model in order to price correctly the previous $[\alpha,\alpha]$ mezzanine tranche. Given a marginal compound correlation $\alpha \in [0,1] \to \bar{\rho}(\alpha)$, we can easily derive G and price basket credit derivatives such as CDOs or CDO squared, either by semi-analytical techniques, by large portfolio approximations or by Monte Carlo.

We now consider the inverse problem of finding a marginal compound correlation from the distribution of conditional default probabilities G. To ease the discussion, we assume here that $\delta = 0$. As an example of the difficulties involved, let us assume that $F(t) = 0.4$ and G being associated with a discrete distribution with probability mass of 0.5 at 0 and 0.5 at 0.8[10]. For $0.5 < \alpha < 0.8$, $G(\alpha) = 0.5$, thus $G(\alpha) < 1 - F(t)$ and there is then no solution to equation (1.4)[11]. Unfortunately, marginal compound correlation may not exist for consistent pricing models.

Equation (1.4) can equivalently be written as:

$$\bar{\rho}(\alpha)\Phi^{-1}\left(G(\alpha)\right)+\Phi^{-1}\left(F(t)\right)=\sqrt{1-\bar{\rho}(\alpha)^2}\Phi^{-1}(\alpha)^{12}.$$

By squaring up the previous equality, we obtain the following second order equation, from which we must consider the positive roots[13]:

$$\left(\Phi^{-1}\left(G(\alpha)\right)^2+\Phi^{-1}(\alpha)^2\right)\bar{\rho}^2+2\Phi^{-1}\left(G(\alpha)\right)\Phi^{-1}\left(F(t)\right)\bar{\rho}+\Phi^{-1}\left(F(t)\right)^2-\Phi^{-1}(\alpha)^2=0. \tag{1.6}$$

Let us firstly discuss the existence of a solution to the previous equation. We denote $\Delta' = \Phi^{-1}(\alpha)^2 \times \left(\Phi^{-1}\left(G(\alpha)\right)^2+\Phi^{-1}(\alpha)^2-\Phi^{-1}\left(F(t)\right)^2\right)$. Equation (1.6) has real roots if and only if $\Delta' \geq 0$ but this may not be the case as illustrated by the following example. We assume that $p_t^{i|V}$ is discretely distributed with

[9] Since only $|\rho|$ is involved in the upfront premium of the mezzanine tranche, we can restrict to correlation parameters within 0 and 1. Let us also remark that $\alpha \to \Phi\left(\frac{\sqrt{1-\rho^2}\Phi^{-1}(\alpha)-\Phi^{-1}(F(t))}{\rho}\right)$ is not a distribution function for negative ρ's.

[10] It can be checked that $\int_0^1 p dG(p) = F(t)$ which means that the expected conditional default probability is the marginal default probability.

[11] The right hand term of equation (1.4) decreases from 1 to $1-F(t)$ as $\bar{\rho}(\alpha)$ increases from 0 to 1.

[12] We recall that for notational simplicity, $\delta = 0$.

[13] Let us emphasize that any solution of:

$$\bar{\rho}(\alpha)\Phi^{-1}\left(G(\alpha)\right)+\Phi^{-1}\left(F(t)\right)=-\sqrt{1-\bar{\rho}(\alpha)^2}\Phi^{-1}(\alpha), \tag{1.5}$$

is also a solution of equation (1.6). The roots of equation (1.6) are the roots of equation (1.4) plus the roots of equation (1.5). In the numerical examples below, we firstly write the roots of equation (1.6) and then check that they indeed solve equation (1.4).

probability 0.5 to be equal to 0 and 0.5 to be equal to 0.5. This leads to a marginal default probability of 0.25. G is constant on]0,0.5[and equal to 0.5. Then $\Delta' < 0$ for $\alpha \in]0.25, 0.5[$.

Let us now discuss whether equation (1.6) admits a unique solution. If $\Phi^{-1}\left(F(t)\right)^2 > \Phi^{-1}(\alpha)^2$, and provided that $\Delta' > 0$, there are two real roots of the same sign. We further meet some practical cases where we actually have two positive roots to equation (1.6) that appear to be also solutions of equation (1.4).

The good news is that for any given $\alpha \in [0,1]$, there are at most two marginal compound correlations $\bar{\rho}(\alpha)$[14]. However, there may be zero or two marginal compound correlations. The non uniqueness is not really a mathematical problem, since we can arbitrarily choose one of the two solutions, which leads by construction to the same value of $G(\alpha)$. As seen from above, only $G(\alpha)$ is involved in loss distributions and price computations. However, it obviously casts some doubt about the economic meaning of such a quantity and the ability to smoothly interpolate between tranches.

2 Stochastic correlation

The inability of the Gaussian Copula Model to fit the market tranches, as illustrated in Table 1 is well-known. Generally, we can say that it underprices the equity and senior tranches and overprices the mezzanine. We feel that the concept of local correlation to be described below is a useful and insightful one but is not a suitable way in which to build and parametrize a skew model. We now consider a general class of stochastic correlation models which we believe offer different ideas and like stochastic volatility models aim to explain rather than simply fit the market. We could consider stochastic correlation to be the simplest of the copula skew models but given the obvious shortcomings of the copula framework, we consider the simplest most intuitive approach to be preferable.

One nice feature of stochastic correlation models[15] is that the marginal distributions of the latent variables remain Gaussian. This eases calibration and implementation of the models. The general structure of a stochastic correlation model involves some latent variables V_i:

$$V_i = \tilde{\rho}_i V + \sqrt{1-\tilde{\rho}_i^2}\bar{V}_i,\ i = 1, \ldots, n, \tag{2.1}$$

where n is the number of names, $V, \bar{V}_i, i = 1, \ldots, n$ Gaussian random variables, all these being jointly independent, $\tilde{\rho}_i$ are some random variables taking values in [0,1], independent from the $V, \bar{V}_i, i = 1, \ldots, n$.

Thanks to the independence between $\tilde{\rho}_i$ and $V, \bar{V}_i$, for any $i = 1, \ldots, n$, given $\tilde{\rho}_i$, V_i follows a standard Gaussian distribution. Thus, the marginal distribution of V_i is also standard Gaussian. We can then still model the default times τ_i, $i = 1, \ldots, n$ as:

$$\tau_i = F_i^{-1}\left(\Phi(V_i)\right), \tag{2.2}$$

where Φ is the Gaussian cdf. Stochastic correlation models are related to mixtures of Gaussian copulas. We refer to Burtschell et al. [2005] for further discussion.

[14] Since these are solution of the second order equation (1.6).

[15] Compared with state dependent factor exposure or local correlation models.

2.1 correlation regimes and idiosyncratic risk

We assume here that $V, \bar{V}_i, \tilde{\rho}_i, i = 1, \dots, n$ are jointly independent. We denote by $\tilde{F}$ the distribution function of $\tilde{\rho}_i$. For simplicity, we assume thereafter that the $\tilde{\rho}_i$ are identically distributed. The latent variables V_i are conditionally independent upon V. We thus deal with a one factor copula. The conditional default probabilities can be written as:

$$p_t^{i|V} = Q(\tau_i \leq t \mid V) = \int_0^1 \Phi\left(\frac{\Phi^{-1}(F_i(t)) - \rho V}{\sqrt{1-\rho^2}}\right) d\tilde{F}(\rho). \tag{2.3}$$

Let us consider a simple example within the previous framework, associated with a binary distribution for the risk factor exposures, $\tilde{\rho}_i = (1 - B_i)\rho + B_i\beta$, which leads to:

$$V_i = ((1 - B_i)\rho + B_i\beta) V + \sqrt{1 - ((1 - B_i)\rho + B_i\beta)^2}\bar{V}_i, \ i = 1, \dots, n, \tag{2.4}$$

where B_i are Bernoulli random variables, $V, \bar{V}_i, i = 1, \dots, n$ Gaussian random variables, all these being jointly independent. We denote by $q = Q(B_i = 1)$, ρ and β lie between 0 and 1. Depending on B_i being equal to 0 or 1, we have a correlation parameter equal to ρ or to β. There are two possible regimes for a given name, one in which it[16] is correlated to the factor V with different correlation parameters.

From iterated expectations theorem and the independence between V and B_i, we have:

$$p_t^{i|V} = Q\left(V_i \leq \Phi^{-1}(F_i(t)) \mid V\right) = (1-q)Q\left(V_i \leq \Phi^{-1}(F_i(t)) \mid V, B_i = 0\right) + qQ\left(V_i \leq \Phi^{-1}(F_i(t)) \mid V, B_i = 1\right),$$

which leads to:

$$p_t^{i|V} = Q(\tau_i \leq t \mid V) = (1 - q)\Phi\left(\frac{\Phi^{-1}(F_i(t)) - \rho V}{\sqrt{1-\rho^2}}\right) + q\Phi\left(\frac{\Phi^{-1}(F_i(t)) - \beta V}{\sqrt{1-\beta^2}}\right), \tag{2.5}$$

which is simply an average of the conditional (on V) default probabilities corresponding to two one factor Gaussian copula models. From this expression of the conditional default probabilities, one can go through the computation of the pgf of the number of defaults or of the loss distribution and then to the computation of basket default or CDO tranche premiums.

As for the computation of the distribution function of $p_t^{i|V}$, let us introduce the real function J such that $J(v) = (1-q)\Phi\left(\frac{\Phi^{-1}(F_i(t))+\rho v}{\sqrt{1-\rho^2}}\right) + q\Phi\left(\frac{\Phi^{-1}(F_i(t))+\beta v}{\sqrt{1-\beta^2}}\right)$ for all v in $\mathbb{R}$[17]. J is increasing from 0 to 1. We denote by G the distribution function of $p_t^{i|V}$. Let $z \in \mathbb{R}$. Then $G(J(-z)) = Q\left(p_t^{i|V} \leq J(-z)\right) = Q(J(-V) \leq J(-z)) = \Phi(-z)$. The graph $\{(J(-z), \Phi(-z)), z \in \mathbb{R}\}$ can be used to plot the distribution function of $p_t^{i|V}$.

[16] "it" refers to the latent variable V_i. Let us remark that we equivalently write:

$$V_i = (1 - B_i) \times \left(\rho V + \sqrt{1-\rho^2}\bar{V}_i\right) + B_i\left(\beta V + \sqrt{1-\beta^2}\bar{V}_i\right), \ i = 1, \dots, n,$$

which involves one factor Gaussian models with different correlations. It can be checked that the joint distribution of the V_i is not Gaussian and thus, while we deal with a factor copula, we do not deal anymore with a Gaussian copula.

[17] for notational simplicity, we omit the dependence upon time horizon t and name i.

Perhaps the most obvious use of a model of this type is to incorporate a state of independence by setting (say) $\beta = 0$. We can think of this as the incorporation of idiosyncratic risk as names defaulting in this state do so with no other impact on the other names. An obvious example of the importance of this is evidenced by the spread widening and downgrades of Ford and GMAC early this year. A knock-on effect of this was higher equity and lower mezzanine tranches premiums in the tranche market coming from a perception of increased idiosyncratic risk in the underlying portfolio. The above specification allows us to apply idiosyncratic risk on a name by name basis. For example, names with wide spreads might be associated with large proportions of idiosyncratic risk.

2.2 dependent correlation and systemic risk

In addition to idiosyncratic risk, the inability to fit the market comes from the overpricing of the mezzanine or equivalently the underpricing of the senior tranches. We can see this clearly when looking at the market premiums of increasingly senior tranches which do not decrease rapidly, suggesting the presence of some lower bound. It can also be observed from the premiums of the senior tranches on iTraxx and CDX which cover losses of [22-100%] and [30-100%] respectively. These tranches are not traded but their premium can be extracted from the rest of the capital structure and the index level. Their premiums are not negligible, even though these tranches can withstand tens of credit events[18].

With the latter point in mind, the obvious component still missing in the model is some sort of systemic risk. One way in which this could be incorporated is via a systemic event such as described in Tavares et al. [2004] or Trinh et al. [2005]. In this framework, with a certain probability, all names will go to default simultaneously. This has the attraction that presumably this systemic probability is very closely related to the super senior tranche premium. An intuitive problem with the approach is that the systemic risk is unrelated to portfolio size and always related to default of the entire pool, irrespective of the size[19]. There are also technical issues, most obviously that there is a cap on the systemic risk spread which is the tightest spread in the underlying portfolio. Even if this is not a direct problem for pricing it complicates the characterisation of the greeks, for example the credit delta may have an unusual behaviour. Furthermore, the fact that the systemic risk will remain unchanged as spreads move is questionable.

An alternative approach that we will follow involves the incorporation of comonotonic risk which corresponds to a state of 100% correlation. In the model, this corresponds to perfect correlation between default times. We could question this on economic grounds since perfectly correlated names may default years apart and we become certain on the ordering of defaults (more risky names always default before less risky names)[20]. However, this approach is more convenient than the above representation for names with different spreads and raises no technical problems. Furthermore the effect of comonotonicity is exactly to increase both equity and senior tranche premiums which is what we obviously must do (see Table 1).

[18] Assuming a recovery of 30%, the iTraxx super senior tranche can withstand 39 credit events whilst 53 credit events will not yet cause a loss on CDX. Of course, lower recoveries will decrease these numbers.

[19] There are ways to get around this (see Elouerkhaoui [2003]), although this will destroy the natural simplicity of the approach.

[20] We can of course generalise this to a state of simply high correlation which part solves of the less economically intuitive features of comonotonicity.

For that purpose, we consider the following model. For simplicity, we restrict to the case where the dependence structure is symmetric:

$$V_i = ((1-B_s)(1-B_i)\rho + B_s)\,V + \Big((1-B_s)\Big((1-B_i)\sqrt{1-\rho^2} + B_i\Big)\Big)\,\bar{V}_i \quad i = 1,\dots,n, \tag{2.6}$$

where $\bar{V}_1,\dots,\bar{V}_n, V, B_s, B_1,\dots,B_n$ are independent, $\bar{V}_1,\dots,\bar{V}_n, V$ are standard Gaussian variables, B_s, $B_1,\dots,B_n$ are Bernoulli random variables and $0 \le \rho \le 1$. We denote by $q = Q(B_i = 1)$ and $q_s = Q(B_s = 1)$. Thus:

$$\tilde{\rho}_i = (1-B_s)(1-B_i)\rho + B_s, \tag{2.7}$$

which is indeed associated with a stochastic correlation model[21]. The marginal distribution of $\tilde{\rho}_i$ is discrete, taking value 0 with probability $q(1-q_s)$, ρ_i with probability $(1-q)(1-q_s)$ and 1 with probability q_s. The conditional upon V and B_s distribution of $\tilde{\rho}_i$ is discrete taking value B_s with probability q and $(1-B_s)\rho+B_s$ with probability $1-q$. Whenever $B_s = 1$, the correlation is equal to 1. While in the previous subsection the correlations were stochastic but independent, in the latter model the correlation parameters are themselves (positively) correlated. As mentionned before, rather than a comonotonic state, we could have an additional (high) correlation parameter. This is clearly a way to extend the model but we have found that in fitting to the market this value will always be very close to 1.

Default times are independent upon V, B_s. We denote the conditional default probabilities by $p_t^{i|V,B_s} = Q(\tau_i \le t \mid V, B_s)$. Since $Q(\tau_i \le t \mid V, B_s) = Q\big(V_i \le \Phi^{-1}(F_i(t)) \mid V, B_s\big)$, we can write the conditional default probabilities as:

$$p_t^{i|V,B_s} = (1-B_s) \times \left((1-q)\Phi\left(\frac{\Phi^{-1}(F_i(t)) - \rho V}{\sqrt{1-\rho^2}}\right) + qF_i(t)\right) + B_s 1_{V \le \Phi^{-1}(F_i(t))}. \tag{2.8}$$

Let us now consider the computation of the pgf of the accumulated losses $L(t) = \sum_{i=1}^n M_i 1_{\tau_i \le t}$, where $M_i = 1-\delta_i$ is the risk exposure for name i and δ_i the corresponding recovery rate, which is assumed here to be deterministic[22]:

$$\psi_{L(t)}(u) = E\left[u^{L(t)}\right] = q_s E\left[u^{L(t)} \mid B_s = 1\right] + (1-q_s) E\left[u^{L(t)} \mid B_s = 0\right].$$

The term $E\left[u^{L(t)} \mid B_s = 0\right]$ corresponds to the pgf within the two states stochastic correlation model with one idiosyncratic state (risk exposure equal to zero) and can be computed by some quadrature technique as $\int E\left[u^{L(t)} \mid V = v, B_s = 0\right]\varphi(v)dv$, where φ is the Gaussian density and:

$$E\left[u^{L(t)} \mid V, B_s = 0\right] = \prod_{i=1}^n \left(q_t^{i|V,B_s=0} + p_t^{i|V,B_s=0} u^{M_i}\right),$$

$p_t^{i|V,B_s=0} = (1-q)\Phi\left(\frac{\Phi^{-1}(F_i(t))-\rho V}{\sqrt{1-\rho^2}}\right) + qF_i(t)$ and $q_t^{i|V,B_s=0} = 1 - p_t^{i|V,B_s=0}$.

[21] It can be easily checked that $(1-B_s)\Big((1-B_i)\sqrt{1-\rho^2} + B_i\Big) = \sqrt{1-\tilde{\rho}_i^2}$.

[22] Thus, the distribution of $L(t)$ is discrete. One could easily cope with stochastic recovery rates and compute the characteristic function of the aggregate losses in that framework.

Let us now compute the term $E\left[u^{L(t)} \mid B_s = 1\right]$. For notational simplicity, we assume that the names are ordered, with name 1 associated with the highest default probability and name n associated with the lowest default probability, i.e. $F_1(t) \geq \ldots \geq F_n(t)$. Conditionally on $B_s = 1$, $\tilde{\rho}_i = 1$, $V_i = V$, thus $\tau_i = F_i^{-1}(\Phi(V))$, $i = 1, \ldots, n$. In the comonotonic case, the accumulated losses can only take $n+1$ values: $0, M_1, M_1 + M_2, \ldots, M_1 + \ldots + M_n$. Conditionally on $B_s = 1$, the loss equal to $M_1 + \ldots + M_{i-1}$ occurs with probability $F_{i-1}(t) - F_i(t)$. The probability of no losses occurring is $1 - F_1(t)$ while the probability of a loss equal to $M_1 + \ldots + M_n$ is $F_n(t)$. This provides an analytical expression[23] of the conditional on $B_s = 1$ pgf $E\left[u^{L(t)} \mid B_s = 1\right]$ as:

$$1 - F_1(t) + (F_1(t) - F_2(t))\, u^{M_1} + \ldots + (F_{i-1}(t) - F_i(t))\, u^{M_1 + \ldots + M_{i-1}} + F_n(t) u^{M_1 + \ldots + M_n}. \tag{2.9}$$

This completes the computation of the pgf of the loss distribution $\psi_{L(t)}$. Going back to the loss distribution is then rather straightforward, especially when the risk exposures are equal $M_1 = \ldots = M_n = M$ since $\psi_{L(t)}(u)$ is a polynomial of u^M whose coefficients can be easily computed by recursion. When the recovery rates are not constant, the previous approach can be adapted through some bucketing technique.

2.3 distribution of conditional default probabilities

We now look after the distribution of the conditional default probabilities $p_t^{\cdot|V,B_s}$. For notational simplicity, we will further remove the name dependence in the superscripts. Under the homogeneity assumption, i.e. $F_1(t) = \ldots = F_n(t) = F(t)$, $\delta_1 = \ldots = \delta_n = \delta$[24] and provided that the number of names n is large, the aggregate loss $L(t) = (1 - \delta) \times \frac{1}{n} \sum_{i=1}^{n} 1_{\tau_i \leq t}$ has the same distribution as $(1 - \delta) p_t^{\cdot|V,B_s}$. We can thus compute CDO or CDO squared tranche premiums and marginal compound correlations out of the distribution function of $p_t^{\cdot|V,B_s}$[25].

The distribution function of conditional default probabilities G is given by: $p \in [0,1] \to G(p) = Q(p_t^{\cdot|V,B_s} \leq p)$ and:

$$G(p) = (1 - q_s)\left(1_{qF(t) < p < 1 - q + qF(t)} \Phi\left(\frac{1}{\rho}\left(\sqrt{1 - \rho^2}\,\Phi^{-1}\left(\frac{p - qF(t)}{1 - q}\right) - K_t\right)\right) + 1_{p \geq 1 - q + qF(t)}\right) + q_s \Phi(-K_t), \tag{2.10}$$

for $0 \leq p < 1$, where $q_s = Q(B_s = 1)$ and $q = Q(B_i = 1)$, $K_t = \Phi^{-1}(F(t))$ (see Appendix).

Let us remark that when $q_s > 0$, 0 and 1 have positive probability: $Q\left(p_t^{\cdot|V,B_s} = 0\right) = q_s \Phi(-K_t) = q_s(1 - F(t))$ and $Q\left(p_t^{\cdot|V,B_s} = 1\right) = q_s \Phi(K_t) = q_s F(t)$. Since $q_s = Q\left(p_t^{\cdot|V,B_s} \in \{0,1\}\right)$, increasing q_s leads to increase the weights associated with extreme probabilities and decrease the probabilities of intermediate

[23] We might have proceeded as above, starting from $p_t^{i|V,B_s=1} = 1_{V \leq \Phi^{-1}(F_i(t))}$, $i = 1, \ldots, n$ and then compute $E\left[u^{L(t)} \mid B_s = 1\right]$ as $\int E\left[u^{L(t)} \mid V = v, B_s = 1\right] \varphi(v) dv$. However, when the default probabilities $F_i(t)$ differ, $E\left[u^{L(t)} \mid V = v, B_s = 1\right]$ is a non-smooth function of v and many points are required in the numerical integration. The above approach is then more accurate and faster.

[24] As above, the recovery rates are assumed to be deterministic.

[25] Let us emphasize that in the homogeneous case, we can compute the pgf of the aggregate loss only through the distribution function of conditional default probabilities G. Indeed, we have: $\psi_L(u) = \int_0^1 \left(pu^{1-\delta} + 1 - p\right)^n dG(p)$.

probabilities[26]. Increasing q_s should result in increasing the premiums of equity and senior tranches and decreasing the premiums of mezzanine tranches.

The distribution function is constant between 0 and $qF(t)$ and between $1-q+qF(t)$ and 1[27]. The inner interval $]qF(t), 1-q+qF(t)[$ is associated with a smooth distribution function. This interval always contains the marginal default probability $F(t)$. As q decreases from 1 to 0, the inner interval increases from $]F(t), F(t)[$ to $]0,1[$. Decreasing the idiosyncratic probability q extends the support of the continuous part of the distribution. This should decrease premiums of intermediate mezzanine tranches and increase premiums of junior mezzanine and senior tranches. When $q < 1$ and $0 < \rho^2 < 1$, there are no probability masses at $qF(t)$ and $1-q+qF(t)$. If $\rho^2 < 50\%$, G has a right derivative $qF(t)$ and a left derivative at $1-q+qF(t)$ equal to zero. In this case, the distribution function G is differentiable at $qF(t)$ and at $1-q+qF(t)$. If $\rho^2 > 50\%$, then the right derivative at $qF(t)$ and the left derivative at $1-q+qF(t)$ are infinite. Even when $\rho^2 < 50\%$, there might be a sharp increase of G after $qF(t)$. When $q = 1$, there is a probability mass of $1-q_s$ at $F(t)$ (associated with the independence case).

We can also derive the (lower) quantile function G^{-1}. For $0 \leq u \leq q_s(1-F(t))$, then $G^{-1}(u) = 0$. For $1 - q_sF(t) \leq u \leq 1$, $G^{-1}(u) = 1$. For $q_s(1-F(t)) < u < 1-q_sF(t)$,

$$G^{-1}(u) = (1-q)\Phi\left(\frac{1}{\sqrt{1-\rho^2}}\left(\rho\Phi^{-1}\left(\frac{u-q_s\Phi(-K_t)}{1-q_s}\right)+K_t\right)\right)+qF(t).$$

2.4 stochastic orders and risk management

Increasing correlation in a stochastic correlation model leads to a decrease in equity tranche premiums. More precisely, let us consider two stochastic correlation models. We denote by $\tilde{\rho}_1, \dots, \tilde{\rho}_n$ and by $\tilde{\beta}_1, \dots, \tilde{\beta}_n$ the correlation parameters. We assume at this stage the $\tilde{\rho}_i$ $i = 1, \dots, n$ are iid as are the $\tilde{\beta}_i$, $i = 1, \dots, n$. We denote by $\tilde{F}_\rho$ the distribution function of the $\tilde{\rho}_i$'s and by $\tilde{F}_\beta$ the distribution function of the $\tilde{\beta}_i$'s. Let us assume that $\tilde{F}_\beta(u) \leq \tilde{F}_\rho(u)$ for all $u \in [0,1]$. This means that $\tilde{\rho}_i \leq \tilde{\beta}_i$ with respect to first order stochastic dominance. As a consequence, there exists some non-negative random variables $\nu_1, \dots, \nu_n$ independent from $V, \bar{V}_1, \dots, \bar{V}_n$ such that $\tilde{\beta}_i = \tilde{\rho}_i + \nu_i$, $i = 1, \dots, n$ where the previous equalities hold in distribution[28]. Conditionally upon $\tilde{\rho}_1, \dots, \tilde{\rho}_n, \nu_1, \dots, \nu_n$ the latent variables in the two models are conditionally Gaussian, with correlation parameters equal to $\tilde{\rho}_1, \dots, \tilde{\rho}_n$ and $\tilde{\rho}_1 + \nu_1, \dots, \tilde{\rho}_n + \nu_n$. We thus have more dependence with respect to the supermodular order in the second model[29].

Thanks to the previous analysis and since $\tilde{\rho}_i = (1-B_s)(1-B_i)\rho + B_s$, we conclude that increasing ρ, increasing q_s or decreasing q will lead to an increase of dependence with default times with respect to the supermodular and thus to a decrease of equity tranche premiums (and an increase of base correlation).

[26]When $q_s = 1$, $Q\left(p_t^{|V,B_s} = 0\right) = 1 - F(t)$ and $Q\left(p_t^{|V,B_s} = 1\right) = F(t)$ which corresponds to the comonotonic case.

[27]The probability of these two intervals is equal to zero.

[28]We simply set $\nu_i = \tilde{F}_\beta^{-1}\left(\tilde{F}_\rho\left(\tilde{\rho}_i\right)\right) - \tilde{\rho}_i$, $i = 1, \dots, n$.

[29]Increasing any non diagonal covariance term in a Gaussian vector with zero mean leads to an increase of dependence with respect to the supermodular order. We conclude by using the invariance of the supermodular order under mixing.

3 State dependent correlation

We are now going to consider some other extensions of the Gaussian copula, belonging to the class of one factor mean variance Gaussian mixtures. In this case, the latent variables are modelled by:

$$V_i = m_i(V) + \sigma_i(V)\bar{V}_i, \quad i = 1, \ldots, n \tag{3.1}$$

where $V, \bar{V}_i$ are independent Gaussian random variables. In this approach, default times are conditionally independent upon V which makes the semi-analytical approach still available (see Laurent & Gregory [2005]). The latent variables V_i are usually no longer Gaussian which makes calibration onto marginal credit curves a bit trickier. For simplicity, we consider from now-on that the credit spreads are equal, i.e. $F_1 = \ldots = F_n = F$. We will consider in greater detail two subcases, the random factor loading model introduced by Andersen & Sidenius [2005] where $\sigma_i(V)$ is constant and the local correlation model of Turc et al [2005], where $\sigma_i(V) = \sqrt{1 - m_i(V)^2}$.

3.1 local correlation

Turc et al. [2005] introduce a local correlation model, where the correlation is state-dependent. This approach is appealing since unlike the base correlation, it is associated with a proper model, and the interpretation and calibration of the local correlation are rather intuitive. However, as in the case of the marginal compound correlation, local correlation cannot always be calibrated onto arbitrage free CDO tranche premiums.

The latent variables are given by:

$$V_i = -\rho(V)V + \sqrt{1 - \rho^2(V)}\bar{V}_i, \quad i = 1, \ldots, n \tag{3.2}$$

where ρ is some function of V taking values in $[0, 1]$. Unless $\rho(V)$ is constant, V_i is not Gaussian. By conditioning upon V, we get the marginal distribution of V_i:

$$H(x) = Q(V_i \le x) = \int_{\mathbb{R}} \Phi\left(\frac{x + \rho(v)v}{\sqrt{1 - \rho^2(v)}}\right)\varphi(v)dv \tag{3.3}$$

for $x \in \mathbb{R}$, where Φ is the Gaussian cdf, φ the Gaussian density function. The default times are obtained from $\tau_i = F^{-1}(H(V_i))$ and the conditional default probabilities can be written as:

$$p_t^{i|V} = Q(\tau_i \le t \mid V) = \Phi\left(\frac{H^{-1}(F(t)) + \rho(V)V}{\sqrt{1 - \rho^2(V)}}\right). \tag{3.4}$$

3.2 random factor loadings

The random factor loading model is closely related to the local correlation approach. We briefly recall the modelling and direct the reader to Andersen & Sidenius [2005] for further details. The simplest form is associated with the following parametric modelling of latent variables:

$$V_i = m + (l1_{V<e} + h1_{V\ge e})V + \nu\bar{V}_i, \quad i = 1, \ldots, n \tag{3.5}$$

where $V, \bar{V}_1, \ldots, \bar{V}_n$ are independent standard Gaussian variables, l, h, e some input parameters, $l, h > 0$. This can be seen a random factor loading model, since the risk exposure $l1_{V<e} + h1_{V\geq e}$ is state dependent. m and ν are such that $E[V_i] = 0$ and $E\left[V_i^2\right] = 1$. This leads to $m = (l-h)\varphi(e)$ and:

$$\nu = \left(1 + m^2 - l^2\left(\Phi(e) - e\varphi(e)\right) - h^2\left(e\varphi(e) + 1 - \Phi(e)\right)\right)^{1/2},$$

where Φ is the Gaussian cdf, φ the Gaussian density function. The marginal distribution of V_i, that we denote by H_{RFL}, involves a bivariate Gaussian cdf:

$$H_{RFL}(x) = Q(V_i \leq x) = \Phi_2\left(\frac{x-m}{\sqrt{\nu^2+l^2}}, e, \frac{l}{\sqrt{\nu^2+l^2}}\right) + \Phi\left(\frac{x-m}{\sqrt{\nu^2+h^2}}\right) - \Phi_2\left(\frac{x-m}{\sqrt{\nu^2+h^2}}, e, \frac{h}{\sqrt{\nu^2+h^2}}\right),$$

where $\Phi_2(.,.,\rho)$ is the bivariate Gaussian cdf with correlation parameter ρ. The default times are then defined by $\tau_i = F^{-1}\left(H_{RFL}(V_i)\right)$, $i = 1, \ldots, n$, where F denotes the marginal distribution of the default times. The conditional default probabilities can be written as:

$$p_t^{i|V} = Q(\tau_i \leq t \mid V) = \Phi\left(\frac{1}{\nu}\left(H_{RFL}^{-1}(F(t)) - m - \left(l1_{V\leq e} + h1_{V>e}\right)V\right)\right). \tag{3.6}$$

The distribution function of the conditional default probabilities is given through:

$$G(p) = Q(p_t^{i|V} \leq p) = \Phi\left(\min\left(\frac{z(p)}{h}, -e\right)\right) + \left(\Phi\left(\frac{z(p)}{l}\right) - \Phi(-e)\right)1_{z(p)>-el}, \tag{3.7}$$

for $0 < p < 1$, where $z(p) = \nu\Phi^{-1}(p) - H_{RFL}^{-1}(F(t)) + m$. Let us assume that $e < 0$ and $h < l$. We can then write G^{-1} as:

$$G^{-1}(u) = \Phi\left(\frac{1}{\nu} \times \left(\left(1_{u\leq\Phi(-e)}h + 1_{u>\Phi(-e)}l\right)\Phi^{-1}(u) + H_{RFL}^{-1}(F(t)) - m\right)\right). \tag{3.8}$$

4 Comparison of stochastic and local correlation

4.1 distribution of conditional default probabilities

To illustrate the previous models, we firstly plot the distribution functions associated with the stochastic correlation and the random factor loading model (figure 1). As for the stochastic correlation model, the parameters are respectively $q_s = 0.14$, $q = 0.81$, $\rho^2 = 58\%$. The default probability $F(t) = 2.96\%$ is such that $\Phi^{-1}(F(t)) = -1.886$. As for the random factor loading model, we took $l = 85\%$, $h = 5\%$ and $e = -2$ (see Andersen and Sidenius [2005]) and $H^{-1}(F(t)) = -1.886$. The graph also shows the distribution function associated with the independence case.

Using the previous data usually provide reasonable fit to quoted CDO tranches. Let us remark that the stochastic correlation and random factor loading models are associated with overall rather similar patterns. However, there is a big discrepancy for the small probability region. The weights put on small probabilities, and thus on first losses are strikingly different. Though the equity tranches may be priced accordingly in the two models, it is likely that a first to default swap would be priced quite differently.

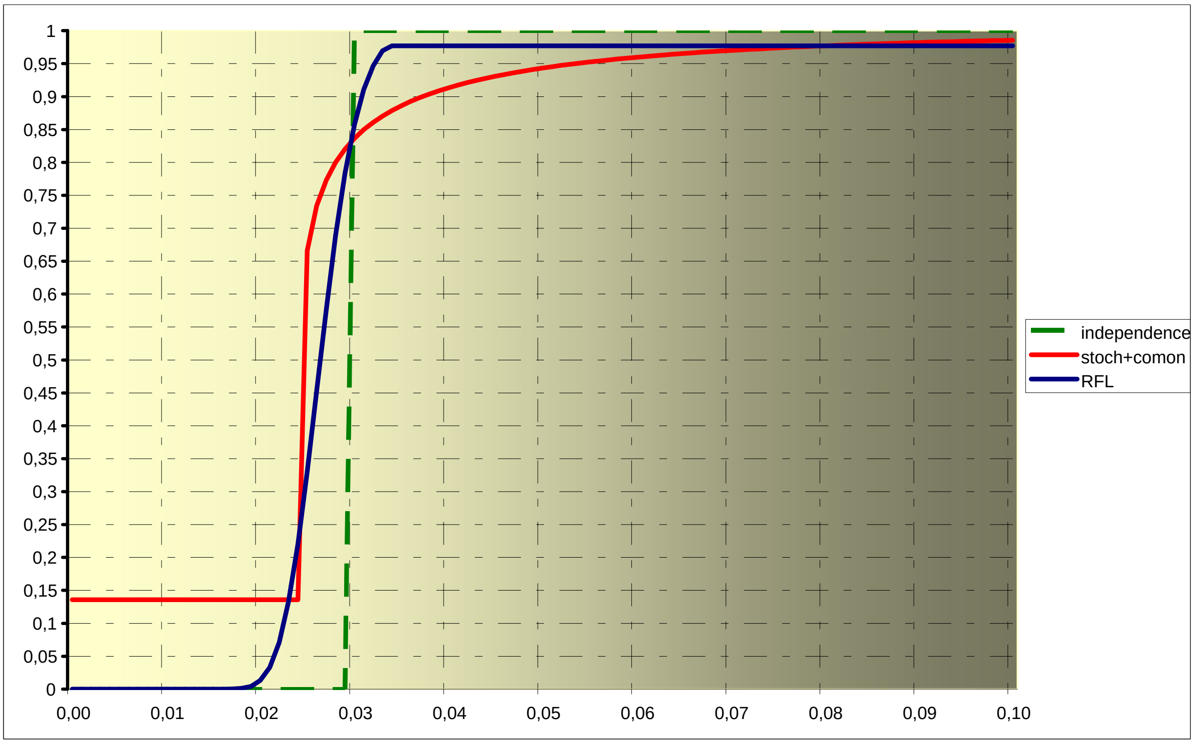

Figure 1. Distribution functions of conditional default probabilities.

4.2 marginal compound correlation

We show the marginal compound correlation associated with the three parameters stochastic correlation and random factor loading models, with the same set of parameters as above (figure 2). The recovery rate is equal to 40%. We see some global similarity between the two models. Looking in greater detail, we see some difference in the region of first losses, since the marginal compound correlation is much larger for the stochastic correlation model. This is consistent with the analysis of distribution functions and the probability mass at 0 in the case of the stochastic correlation model.

The marginal compound correlation exhibit a smile feature, with a rather strong irregularity where the correlations are much higher. After $p = qF(t)$, there is a sharp increase of the distribution function associated with the stochastic correlation model, leading to that sharp increase in the marginal compound correlation after $p = (1-\delta)qF(t)$. That shape is not inconsistent with what was observed in the market with compound correlations on 10 years iTraxx tranches. We notice that we have a zero marginal compound correlation for $p = (1-\delta)F(t)$. This feature is model independent. For attachment - detachment points above that threshold, there are two marginal compound correlations. We chose the smallest one consistently with market practice, which explains the discontinuity in the marginal compound correlations. Of course, both marginal compound correlations lead to the same value of G, which is the meaningful input.

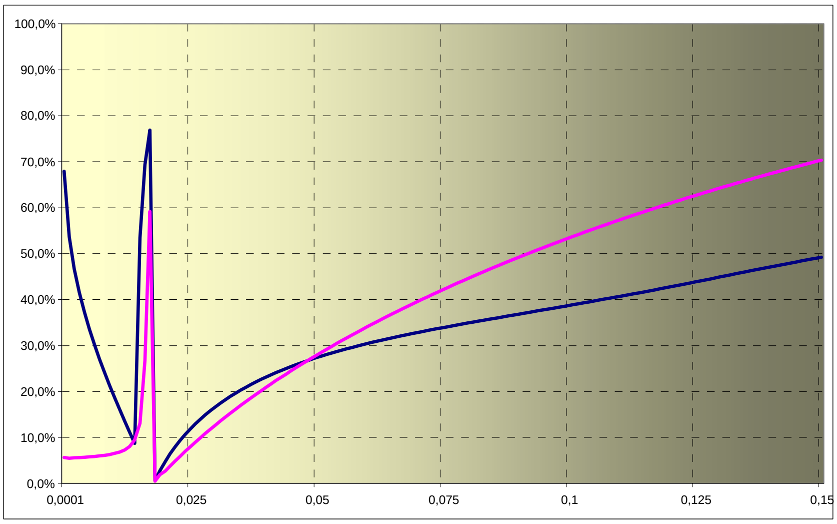

Figure 2. Marginal compound correlation associated with stochastic correlation and random factor loading models

4.3 local correlation

We chose to calibrate the local correlation function ρ onto the distribution of $p_t^{i|V}$, G, which embeds all relevant information about loss distributions at time t. Rather than using some arbitrary and possibly flawed interpolation of market prices or correlations[30], we have firstly fitted some model (either stochastic correlation or random factor loading), which produces arbitrage free prices and a reasonable fit to market quotes and in a second step, we consider the associated local correlation. The calibration has been completed according to Turc et al. [2005]. For self-completeness, we recall and detail the approach. We denote by:

$$\varepsilon(V) = \frac{H^{-1}(F(t)) + \rho(V)V}{\sqrt{1-\rho^2(V)}}$$

and assume that ε is increasing with V. From equation (3.4), we have $p_t^{i|V} = \Phi(\varepsilon(V))$. Since $p_t^{i|V} = \Phi(\varepsilon(V))$, for p in $[0,1]$, $G(p) = Q\left(p_t^{i|V} \leq p\right) = \Phi\left(\varepsilon^{-1}\left(\Phi^{-1}(p)\right)\right)$, which leads to $\varepsilon(v) = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right)$ for $v \in \mathbb{R}$. We thus need to look for a function ρ such that:

$$\frac{H^{-1}(F(t)) + \rho(v)v}{\sqrt{1-\rho^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right), \ \forall v \in \mathbb{R}, \tag{4.9}$$

with: $H(x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho(v)v}{\sqrt{1-\rho^2(v)}}\right)\varphi(v)dv$. The right-hand term is given while the left-hand term involves ρ either directly or indirectly through H.

We proceed with a fixed point algorithm, constructing a series of functions ρ_n converging to a solution of previous functional equation. We set $\rho_0(v) = \rho$, where ρ is some arbitrary correlation parameter (ρ_0 is then a constant function). We then set H_0 such that $H_0(x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho v}{\sqrt{1-\rho^2}}\right)\varphi(v)dv$, which leads obviously to

[30] It is well-known that careless smoothing of base correlations leads to negative tranche prices.

$H_0 = \Phi$. We then solve the following equation for all v's:

$$\frac{\Phi^{-1}(F(t)) + \rho_1(v)v}{\sqrt{1-\rho_1^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right). \quad (4.10)$$

which provides ρ_1[31]. H_1 is such that $H_1(x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho_1(v)v}{\sqrt{1-\rho_1^2(v)}}\right)\varphi(v)dv$. Step n simply involves solving for $\rho_n(v)$ in:

$$\frac{H_{n-1}^{-1}(F(t)) + \rho_n(v)v}{\sqrt{1-\rho_n^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right),$$

and then set H_n such that $H_n(x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho_n(v)v}{\sqrt{1-\rho_n^2(v)}}\right)\varphi(v)dv$. Whenever $H_n = H_{n-1}$[32], the algorithm has converged and we set $\rho = \rho_n$.

Equation (1.4) can be equivalently written as $\frac{\Phi^{-1}(F(t))+\bar{\rho}(p)\Phi^{-1}(G(p))}{\sqrt{1-\bar{\rho}(p)^2}} = \Phi^{-1}(p)$. By stating $p = G^{-1}(\Phi(v))$, we obtain $\frac{\Phi^{-1}(F(t))+\bar{\rho}(p)v}{\sqrt{1-\bar{\rho}(p)^2}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right)$. Comparing with equation (4.10), we get $\bar{\rho}(p) = \rho_1(v)$, which means that local correlation is directly related to the marginal compound correlation of some mezzanine tranche with an appropriate strike. The previous equation also shows that local correlation shares the same existence and uniqueness problems as marginal compound correlation.

Let us remark that ε is time dependent due to the terms $F(t)$ and G; thus the calibration of ρ involves a special maturity choice. There is clearly no guarantee that calibrating on two different time horizons would lead to the same local correlation.

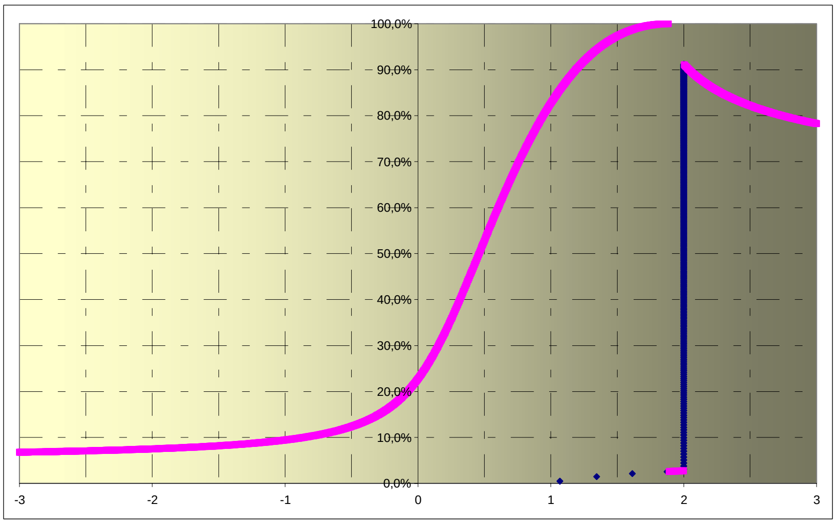

Figure 3. Local correlation associated with the random factor loading model

In the random factor loading model, we notice that for low levels of the factor V, $\rho(V)$ is close to 5%, while for high levels of V, $\rho(V)$ is close to 85% consistently with the inputs, as is the case for the discontinuity at $V = 2$. When V is in between 1 and 2, there are two local correlations, leading to the same CDO prices.

[31] The previous equation is quite similar to the one providing the marginal compound correlation. By squaring up, we obtain a second order algebraic equation, which we solve as usual.

[32] We only need to check that $H_n^{-1}(F(t)) = H_{n-1}^{-1}(F(t))$.

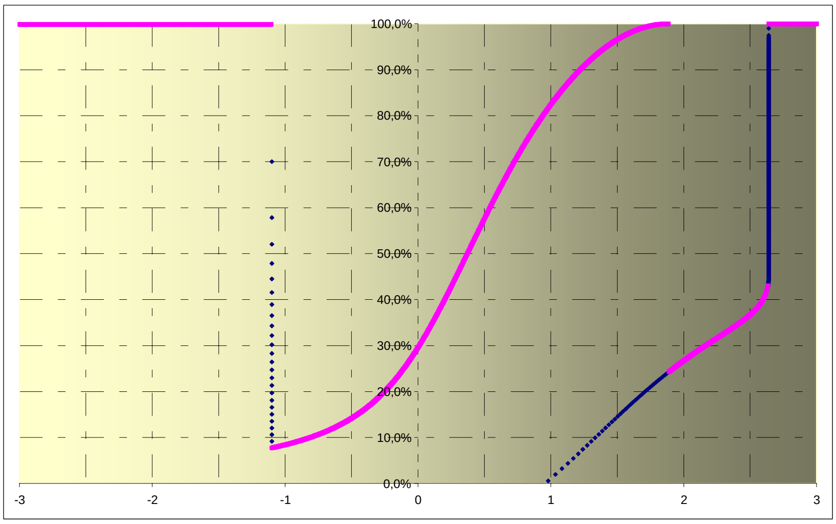

Figure 4. Local correlation associated with the stochastic correlation model

In the three parameters stochastic correlation model, for small and large values of V, the local correlation is equal to 1, corresponding to the existence of a comonotonic state. As in the previous case, we have two local correlations for values of V approximately between 1 and 2.

Figure 6 below plots the functions H_0 and H_1 for the stochastic correlation model. It can be seen that $H_0^{-1}(F(t)) = H_1^{-1}(F(t))$. This shows that the algorithm has converged only after one step and that actually, $\rho_1 = \rho$ and $H_1 = H$. We also checked that the function ε was increasing.

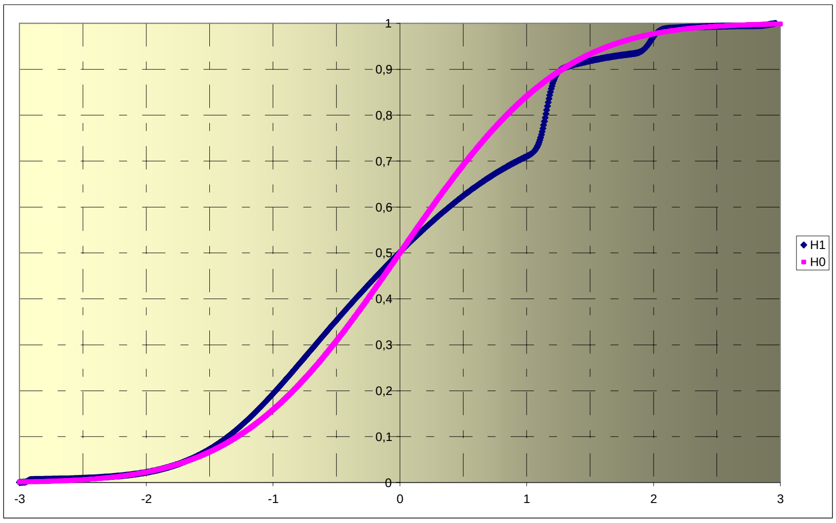

Figure 5. Check of convergence of the fixed point algorithm (stochastic correlation model)

In the random factor loading model as well as in the stochastic correlation model, we see that the local correlation is rather irregular and not unique, casting some doubt about the intuitiveness of the approach.

4.4 market fits

We now show some recent fits of the three parameter stochastic correlation model to five years iTraxx and CDX tranche premiums. We took market spreads on the date in question (column two) and assumed a flat recovery of 40%. The premiums were computed using the semi-analytical technique described above, taking into account the differences in individual spreads and discounting effects. For simplicity, we provide the average spread of the index at the time of computation[33]. We also plot compound correlations (columns four and five). We can see that the fit is rather good with close agreement on all tranches. Indeed it is possible to add name specific parameters and fit perfectly although our own view is that this is of limited value[34].

Tranche	Market	Model	Market ρ	Model ρ
Index	36			
[0-3%]	24%	25%	16%	14%
[3-6%]	83	84	4%	4%
[6-9%]	27	27	12%	12%
[9-12%]	14	14	17%	17%
[12-22%]	9	9	28%	28%
[22-100%]	4	2	63%	56%

Table 2. Fit of model characterised by equation (2.8) to iTraxx market data on 31-August-2005. All values in bp pa unless otherwise stated. As usual the equity tranche is quoted as an up-front premium, in addition to the contractual 500 bp pa. $q_s = 0.13$, $q = 0.84$, $\rho^2 = 73.5\%$.

Tranche	Market	Model	Market ρ	Model ρ
Index	50			
[0-3%]	40%	38%	10%	13%
[3-7%]	126	139	2%	2%
[7-10%]	36	39	12%	13%
[10-15%]	20	17	20%	19%
[15-30%]	10	10	34%	34%
[30-100%]	2	3	59%	65%

Table 3. Fit of model characterised by equation (2.8) to CDX market data on 31-August-2005. All values in bp pa unless otherwise stated. As usual the equity tranche is quoted as an up-front premium, in addition to the contractual 500 bp pa. $q_s = 0.15$, $q = 0.84$, $\rho^2 = 85.3\%$.

We now show a fit for the 10-Jun-2005 closer to the crisis earlier in the year. While the market was still in a rather dislocated state at this point, we still have a good fit of the model to the market. Let us remark

[33] Detailed data about individual credit spreads and default free rates is available upon request to the authors

[34] Unless of course this coincides with our view regarding the idiosyncratic risk of each name.

that the fitting parameters are related to the premiums of the junior super senior and super senior tranches (i.e. [12-22%] and [22-100%]). As discussed by Gregory et al. [2005], these two tranches should be closely related. However, their values have not historically always behaved in this way, perhaps not surprisingly as the [22-100%] is only valued implicitly. We can improve the fit by adjusting the recovery values (since the relative values of these tranches are very sensitive to recovery). However, we believe that this is rather needless and is unrelated to any risk management of recovery risk.

Tranche	Market	Model	Market ρ	Model ρ
Index	41			
[0-3%]	29%	29%	15%	17%
[3-6%]	109	113	2%	2%
[6-9%]	34	35	10%	10%
[9-12%]	23	21	18%	17%
[12-22%]	15	15	29%	30%
[22-100%]	4	4	62%	61%

Table 4. Fit of model characterised by equation (2.8) to iTraxx market data on 10-June-2005. $q_s = 0.195$, $q = 0.87$, $\rho^2 = 76.5\%$.

In the case of CDX, the best fit on 10-June-2005 is obtained with $\rho^2 = 99\%$. $\tilde{\rho}_i$ can then only take values 0 and 1. We can check that when $\rho^2 = 1$, the distribution of conditional default probabilities is discrete with probability masses at 0, $qF(t)$, $1 - q + qF(t)$ and 1.

Tranche	Market	Model	Market ρ	Model ρ
Index	57			
[0-3%]	45%	44%	11%	13%
[3-7%]	160	191	0%	1%
[7-10%]	41	42	9%	10%
[10-15%]	22	20	17%	16%
[15-30%]	14	14	34%	35%
[30-100%]	5	4	72%	67%

Table 5. Fit of model characterised by equation (2.8) to CDX market data on 10-June-2005. $q_s = 0.189$, $q = 0.89$, $\rho^2 = 99\%$.

As for the 16-May-2005, the best fit for iTraxx data is obtained with $q = 1$ (thus, the value of ρ^2 is irrelevant). In this case, $\tilde{\rho}_i$ can then only take values 0 and 1 and the distribution of conditional default probabilities is discrete with masses at 0, $F(t)$ and 1. It can be seen that the fit is rather poor for the equity and junior mezzanine tranches, while the most senior tranches are well priced. We also notice that during this troubled period, there was some difficulty in finding compound correlations for the [3-6%] iTraxx and [3-7%] CDX tranches. In this case, we reported a 0% compound correlation. This is consistent with the previous discussion about the existence of the marginal compound correlation.

Tranche	Market	Model	Market ρ	Model ρ
Index	54			
[0-3%]	47%	41%	9%	16%
[3-6%]	150	211	0%	0%
[6-9%]	48	41	7%	6%
[9-12%]	33	35	14%	15%
[12-22%]	24	25	28%	29%
[22-100%]	5	6	56%	62%

Table 6. Fit of model characterised by equation (2.8) to iTraxx market data on 16-May-2005. $q_s = 0.246$, $q = 1$.

Tranche	Market	Model	Market ρ	Model ρ
Index	76			
[0-3%]	62%	56%	6%	13%
[3-7%]	256	390	0%	0%
[7-10%]	56	59	6%	6%
[10-15%]	31	29	14%	14%
[15-30%]	19	20	32%	33%
[30-100%]	2	5	51%	66%

Table 7. Fit of model characterised by equation (2.8) to CDX market data on 16-May-2005. $q_s = 0.205$, $q = 0.90$, $\rho^2 = 99\%$.

An obvious question to be considered in a model is not only the ability to fit the market but also the stability of parameters over time. We show in Figure 1 the time series of the calibrated idiosyncratic risk, q, systemic risk, q_s and correlation, ρ^2 from the 15th of April 2005 to the 6th of September 2005. We clearly see the effect of the market dislocation of May returning, which is illustrated by a decrease in both idiosyncratic and systemic risk. Clearly over this rather volatile period, the parameters are not so stable. However, by looking at the more recent history we can be a little reassured that this is not so unreasonable. For example, the iTraxx index level was rather stable in the range 35 to 38 bps from mid-July to end of August and yet the market tranche premiums (and consequently model parameters) move significantly in this period. It seems unreasonable to suggest that any model would capture this behaviour with stable parameters.

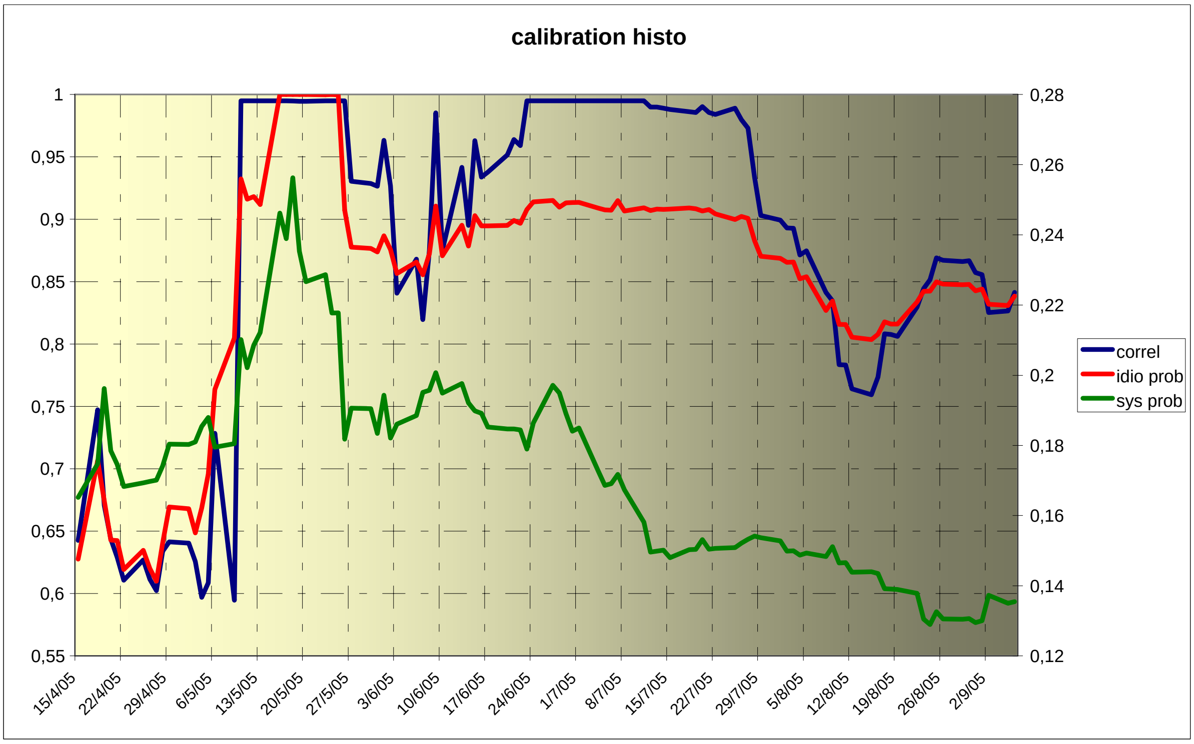

Figure 6. Time series of 5Y iTraxx calibrated systemic (q_s), idiosyncratic (q) and correlation (ρ^2) parameters. Systemic probability (q_s) on right axis.

Conclusion

In this article, we have shown how to adapt the one factor Gaussian copula and analysed some copula skew approaches to matching the market prices of synthetic CDO tranches.

We show that a simple class of stochastic correlation models can provide a reasonable good fit to the market at a single maturity. These are quite easy to implement, since the marginal distributions of latent variables remain Gaussian, and the parameters have intuitive interpretation. Stochastic orders theory is helpful for risk management purposes. We provided some dynamics of the calibrated parameters, which is a first step in the risk management process. The cornerstone of the pricing remains the factor copula approach, which is associated with semi-analytical pricing techniques and easy to use large portfolio approximations. This we believe is an alternative (or complementary) to the market standard base correlation approach.

It can be seen that distribution functions of conditional default probabilities are unambiguous modelling inputs, while local correlation suffers from existence and non uniqueness issues. Thus, a local correlation approach offers some intuition but is ultimately not an obvious way to parametrise of model.

Stochastic and local correlation approaches suffer from global consistency, related to calibration to different maturities, credit spread dynamics and hedging issues. Such issues may well need to be addressed outside a copula framework. Before this however, the copula skew approach is a valuable next step in understanding the pricing and risk management issues of bespoke credit portfolios. Stochastic and local correlation might thus either appear as transient and transitional since they do not sever the umbilical cord with the Gaussian copula and lead towards a mature modelling framework. But they do, we believe, give a useful way in which to analyse the problems of the correlation skew, move away from base correlation approaches and take the next step forward in correlation modelling.

References

Andersen, L., J. Sidenius & S. Basu, 2003, *All Your Hedges in One Basket*, *RISK*, November, 67-72.

Andersen, L. & J. Sidenius, 2005, *Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings*, Journal of Credit Risk, 1(1).

Burtschell, X., J. Gregory & J-P. Laurent, 2005, *A Comparative Analysis of CDO Pricing Models*, working paper, ISFA Actuarial School, University of Lyon & BNP-Paribas.

Elouerkhaoui, Y., 2003, *Credit Derivatives: Basket Asymptotics*, working paper, université Paris Dauphine.

Gregory, J., S. Delacote & C. Donald, *Relative Value of Super-Senior Tranches*, Structured Credit Relative Value Strategy BNP Paribas, 2005.

Gregory, J. & J-P. Laurent, 2003, *I Will Survive*, *RISK*, June, 103-107.

Hull, J. & A. White, 2004, *Valuation of a CDO and an n^{th} to Default CDS Without Monte Carlo Simulation*, Journal of Derivatives, 2, 8-23.

Hull, J. & A. White, 2005, *The Perfect Copula*, working paper, University of Toronto.

Kalemanova, A., B. Schmid & R. Werner, 2005, *The Normal inverse Gaussian distribution for synthetic CDO*, working paper.

Laurent, J-P. & J. Gregory, 2005, *Basket Default Swaps, CDOs and Factor Copulas*, Journal of Risk, 7(4), 103-122.

Li, D.X., 2000, *On Default Correlation: a Copula Approach*, Journal of Fixed Income, **9**, March, 43-54.

Mortensen, A., 2005, *Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models*, working paper, Copenhagen Business School.

Schloegl, L., 2005, *Modelling Correlation Skew via Mixing Copulae and Uncertain Loss at Default*, Presentation at the Credit Workshop, Isaac Newton Institute.

Sharke, H., 2005, *Remarks on Pricing Correlation Products*, Bank Austria Creditanstalt working paper.

Tavares, P., T-U. Nguyen, A. Chapovsky & I. Vaysburd, 2004, *Composite Basket Model*, working paper, Merrill Lynch.

Trinh, M., R. Thompson & M. Devarajan, 2005, *Relative Value in CDO Tranches: A View through ASTERION*, Quantitative Credit Research Quarterly, Lehman Brothers, vol. 2005-Q1.

Turc, J., P. Very & D. Benhamou, 2005, *Pricing CDOs with a Smile*, SG Credit Research.

Vasicek, O., 1987, *The Loan Loss Distribution*, working paper, KMV.

Willemann, S., 2005, *Fitting the CDO Correlation Skew: A Tractable Structural Jump-Diffusion Model*, working paper, Aarhus School of Business.

Appendix: distribution of conditional default probabilities

$$Q(p_t^{|V,B_s} \leq p) = (1-q_s)E\left[Q\left((1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p \mid V\right)\right]+q_sE\left[Q\left(1_{V\leq K_t}\leq p\mid V\right)\right],$$

for $0 \leq p < 1$. $1_{V\leq K_t} \leq p \Leftrightarrow V > K_t$. As a consequence, $E\left[Q\left(1_{V\leq K_t}\leq p\mid V\right)\right] = \Phi(-K_t)$. Let us assume that $q < 1$, $\rho > 0$. Then, $(1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)$ varies between $q\Phi(K_t)$ and $1-q+q\Phi(K_t)$. We recall that $\Phi(K_t) = F(t)$.

- if $qF(t) < p < 1-q+qF(t)$,

$$(1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p \Leftrightarrow -V \leq \frac{1}{\rho}\left(\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{p-q\Phi(K_t)}{1-q}\right)-K_t\right).$$

  Then $Q\left((1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p\mid V\right) = 1$ if $-V \leq \frac{1}{\rho}\left(\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{p-q\Phi(K_t)}{1-q}\right)-K_t\right)$ and 0 otherwise.

- if $p \leq qF(t)$, then $Q\left((1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p\mid V\right) = 0$.

- if $p \geq 1-q+qF(t)$, then $Q\left((1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p\mid V\right) = 1$.

As a consequence, we can write $E\left[Q\left((1-q)\Phi\left(\frac{K_t-\rho V}{\sqrt{1-\rho^2}}\right)+q\Phi(K_t)\leq p\mid V\right)\right]$ as:

$$1_{qF(t)<p<1-q+qF(t)}\Phi\left(\frac{1}{\rho}\left(\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{p-q\Phi(K_t)}{1-q}\right)-K_t\right)\right)+1_{p\geq 1-q+qF(t)}.$$

This leads to the stated expression.