Beyond the Gaussian Copula: Stochastic and Local Correlation

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

First version: October 2005
Revised: January 2007

Abstract

We consider stochastic correlation models that account for the correlation smile in the pricing of synthetic CDO tranches. These can be viewed as tractable extensions of the one factor Gaussian copula model. We analyse these models through their conditional default probability distributions. We also give some examples of using a three states stochastic correlation model to fit the market and discuss some risk management issues. We provide some analytical computations within the large homogeneous portfolio approximation. Eventually, we compare the stochastic correlation model with another popular state dependent correlation model, namely the random factor loading model.

JEL Classification: C 31, G 13
Key words: default risk, CDOs, correlation smile, factor copulas, stochastic correlation, local correlation, compound correlation, random factor loadings, large homogeneous portfolio approximations.

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 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).

[1]BNP Paribas, xavier.burtschell@bnpparibas.com,
[2]Barclays Capital, jon.gregory@barclayscapital.com,
[3]ISFA Actuarial School, Université Lyon I & 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 useful comments, M. Shchetkovski and F. Rhazal for pointing out an error in the first draft. Comments from the two referees and the journal editor are gratefully acknowledged. All remaining errors are ours.

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 has shown the Gaussian copula model to be inadequate and there has been a build up of literature on correlation skew modelling and a reliance by market practitioners 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 (Andersen and Sidenius [2005b], Burtschell et al. [2005], Schloegl [2005]) or state dependent risk exposure (Andersen and Sidenius [2005a], Laurent and Gregory [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 *state dependent correlation* model. Examples of such models are provided by Turc et al. [2005] or Andersen and Sidenius [2005a]. In this paper, we provide a thorough analysis of a three states stochastic correlation model and subsequently compare it with the random factor loading model. More precisely, the paper is organized as follows.

- The first section focuses on the main features of stochastic correlation models. We specialize this framework to a three states specification of the correlations. Since the previous model belongs to the factor copula class, we show that the probability generating function (pgf thereafter) of the aggregate loss can be easily computed, leading to semi-analytical CDO tranche premiums. Stochastic orders properties are used to show monotonicity results of CDO tranche premiums with respect to model parameters.
- Section 2 provides some market fits of the stochastic correlation model. We also study the stability of calibrated parameters through time, including the 2005 correlation crisis and discuss the credit deltas associated with the different tranches.
- Section 3 provides analytical computations of the distribution function of a large homogeneous portfolio, of zero-coupon CDO tranche premiums and of the *marginal compound correlation curve*, these being derived under the large homogeneous portfolio approximation.
- Section 4 aims at comparing stochastic and state dependent correlation models. We recall some useful properties of the random factor loading model, including large homogeneous portfolio approximations. We subsequently compare the large homogeneous portfolio distributions within the stochastic correlation and the random factor loading models. We also compare the two models through the marginal

compound and local correlation curves. This shows that while both models can be reasonably well fitted to market quotes, they would lead to significant differences for non standard tranches.

1 Factor copulas and stochastic correlation

We detail here the computation of the loss distributions associated with a three states stochastic correlation model. This model is an extension of the one factor Gaussian copula model and the same semi-analytics techniques also apply here. The model involves three parameters and, thanks to stochastic orders theory, we briefly discuss in this section how the parameters drive the CDO tranche premiums.

1.1 One factor Gaussian copula

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 standard Gaussian random variables. We denote by $F(t) = Q(\tau_i \leq t)$ the marginal default probability for time horizon t[4] and define default times τ_i of names $i = 1, \dots, n$ by setting $\tau_i = F^{-1}(\Phi(V_i))$ where Φ is the standard Gaussian distribution function. The previous one factor Gaussian copula model leads to a semi-explicit pricing of CDO tranches (see Gregory and Laurent [2003], Andersen et al [2003], Hull and White [2004]) but does not match the market prices (see Table 1).

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 iTraxx market tranche quotes. The market quotes are as of the 30-Aug-05 and apart from the equity tranche are expressed in basis points par annum (bp pa). As usual, the equity tranche is quoted as an up-front premium, in addition to the contractual 500 bp pa. The coupon payment frequency is quarterly. Maturity date is 20 June 2010 (iTraxx series 3). 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.

Generally, the one factor Gaussian copula model underprices the equity and senior tranches and overprices the mezzanine. The one factor Gaussian copula model has therefore been extended in various directions. A

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

first idea is to use more complex correlation structures as in Gregory and Laurent [2004]. Hull and White [2004], Kalemanova et al. [2005] consider other distributions for the factor V and the idiosyncratic risks $\bar{V}_i$, such as Student t or Normal Inverse Gaussian. Burtschell et al [2005] review some popular factor copula models.

1.2 Stochastic correlation models

In the following, we 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.

Let us firstly state the general structure of a stochastic correlation model. It involves some latent variables V_i:

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

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

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

$$\tau_i = F_i^{-1}\left(\Phi(V_i)\right), \tag{1.2}$$

where Φ is the Gaussian cumulative density function. Stochastic correlation models are related to mixtures of Gaussian copulas. We refer to Burtschell et al. [2005] for further discussion[5]. Let us remark that we remain in a copula and thus static framework. Such a static framework is well suited for European type options and thus for CDOs but does not extend readily to the pricing of payoffs that depend on the path of the aggregate loss such as forward starting CDOs. We emphasize that this feature is not specific to the stochastic correlation model but applies to all one factor copula approaches.

1.3 Three states stochastic correlation

For practical implementation, let us consider the following model which is an extension of Burtschell et al [2005] or Schloegl [2005]. For simplicity, we restrict to the case where the dependence structure is symmetric:

$$V_i = \left((1-B_s)(1-B_i)\rho + B_s\right)V + \left((1-B_s)\left((1-B_i)\sqrt{1-\rho^2}+B_i\right)\right)\bar{V}_i \quad i = 1,\dots,n, \tag{1.3}$$

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 random 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:

$$\tilde{\rho}_i = (1-B_s)(1-B_i)\rho + B_s, \tag{1.4}$$

[5] Although this model corresponds to a mixture of copulas, this is not the same idea as applied in base correlation or the "Mixture of Expected Losses" approach suggested by Liang and Li [2005]. As shown by Piterbarg [2003], such a model is under specified leading to inconsistencies in pricing. In the stochastic correlation approach we propose we still can define a unique distribution of default times.

which is indeed associated with a stochastic correlation model[6]. The marginal distribution of $\tilde{\rho}_i$ is discrete, taking value 0 with probability $q(1-q_s)$, ρ with probability $(1-q)(1-q_s)$ and 1 with probability q_s.

The state where $\tilde{\rho} = 0$ can be viewed 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 in 2005. 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.

The underpricing of the senior tranches when using a Gaussian copula is another issue. 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[7]. When $B_s = 1$, the correlation parameter is equal to 1. This corresponds to some sort of systemic risk. 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[8].

1.4 Stochastic orders

Let us recall that in our stochastic correlation model, $\tilde{\rho}_i = (1-B_s)(1-B_i)\rho + B_s$. It can be proved that increasing ρ, increasing q_s or decreasing q leads to a decrease of equity tranche premiums (and an increase of base correlation).

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

[7] 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.

[8] We can of course generalise this to a state of simply high correlation which part solves of the less economically intuitive features of comonotonicity. 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. Another 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. There are ways to get around this (see Elouerkhaoui [2003]), although this will destroy the natural simplicity of the approach. 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.

The proof is based on stochastic orders theory: let us firstly 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 exist 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[9]. 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[10]. 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 (see Burtschell et al [2005] for a detailed discussion).

1.5 Computation of loss distributions

We show here that the pgf of the accumulated losses can be easily computed up to a one dimensional numerical integration. Since CDO tranche premiums only involve loss distributions over different time horizons (see Gregory and Laurent [2003]), this paves the way for a semi-analytical computation of CDO tranche premiums.

In the above stochastic correlation model, default times are independent upon V, B_s. We are thus in a factor copula framework. Though this is seemingly a two factor framework, we show that the dimensionality of the model remains equal to one.

We denote the conditional default probabilities by $p_t^{i|V,B_s} = Q(\tau_i \leq t \mid V, B_s)$. Since $Q(\tau_i \leq t \mid V, B_s) = Q\left(V_i \leq \Phi^{-1}(F_i(t)) \mid V, B_s\right)$, we can write the conditional default probabilities as:

$$p_t^{i|V,B_s} = (1 - B_s) \times \left((1 - q)\Phi\left(\frac{\Phi^{-1}(F_i(t)) - \rho V}{\sqrt{1 - \rho^2}} \right) + qF_i(t) \right) + B_s 1_{V \leq \Phi^{-1}(F_i(t))}. \tag{1.5}$$

As can be seen in Burtschell et al. [2005], Hull and White [2005], the distributions of conditional default probabilities are the main drivers of CDO prices within factor copula models. Not surprisingly, this is also the case of the three states stochastic correlation model.

Let us now detail the computation of the pgf of the accumulated losses $L(t) = \sum_{i=1}^n M_i 1_{\tau_i \leq t}$, where $M_i = 1 - \delta_i$ is the risk exposure for name i and δ_i the corresponding recovery rate, which is assumed here

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

[10] 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.

to be deterministic[11]:

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

Let us remark that:

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

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

As a consequence, the term $E\left[u^{L(t)} \mid B_s = 0\right] = \int_{\mathbb{R}} E\left[u^{L(t)} \mid V = v, B_s = 0\right] \varphi(v) dv$, where φ is the Gaussian density, can be computed by some quadrature technique.

Let us now consider 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[12] of the conditional on $B_s = 1$ pgf $E\left[u^{L(t)} \mid B_s = 1\right]$ as:

$$1 - F_1(t) + (F_1(t) - F_2(t))\, u^{M_1} + \ldots + (F_{i-1}(t) - F_i(t))\, u^{M_1 + \ldots + M_{i-1}} + F_n(t) u^{M_1 + \ldots + M_n}. \tag{1.6}$$

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.

2 Market fits

We provide here some empirical results related to the ability of the three states stochastic correlation model to fit the market and to the risk management within such a model.

2.1 Calibration to market quotes

We now show some 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

[11] 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.

[12] 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 $\int_{\mathbb{R}} E\left[u^{L(t)} \mid V = v, B_s = 1\right] \varphi(v) dv$. 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.

spread of the index at the time of computation[13]. As for iTraxx tranches, the maturity is 20 June 2010 (iTraxx series 3). We also plot compound correlations (columns four and five). The compound correlation is the parameter to be put in the Gaussian copula model in order to price correctly the corresponding tranche[14]. The premium of the [22-100%] tranche is implied from the rest of the capital structure and level of the index[15]. 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[16].

Tranche	Market premium	Model premium	Market compound correlation	Model compound correlation
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 (1.5) 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 premium	Model premium	Market compound correlation	Model compound correlation
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%

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

[14] As for the marginal compound correlation, this parameter may not be unique. Unlike base correlation, it only depends on the price of a given tranche.

[15] Unless the default-free interest rates are equal to zero, we cannot compute the premium of the super senior tranche in a model-free framework.

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

Table 3. Fit of model characterised by equation (1.5) 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 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 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 premium	Model premium	Market compound correlation	Model compound correlation
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 (1.5) 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 premium	Model premium	Market compound correlation	Model compound correlation
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 (1.5) to CDX market data on 10-June-2005. $q_s = 0.189$, $q = 0.89$, $\rho^2 = 99\%$.

Tranche	Market premium	Model premium	Market compound correlation	Model compound correlation
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 (1.5) to iTraxx market data on 16-May-2005. $q_s = 0.246$, $q = 1$.

Tranche	Market premium	Model premium	Market compound correlation	Model compound correlation
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 (1.5) to CDX market data on 16-May-2005. $q_s = 0.205$, $q = 0.90$, $\rho^2 = 99\%$.

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.

2.2 Dynamics of implied parameters

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 also plotted the 5 year iTraxx index (series 3) over the same period. We clearly see the effect of the market dislocation of May returning, which is illustrated by a decrease in both idiosyncratic (q) and systemic risk (q_s). 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 2005 and yet the market tranche premiums (and consequently model parameters) moved significantly in this period. It seems unreasonable to suggest that any model would capture this behaviour with stable parameters.

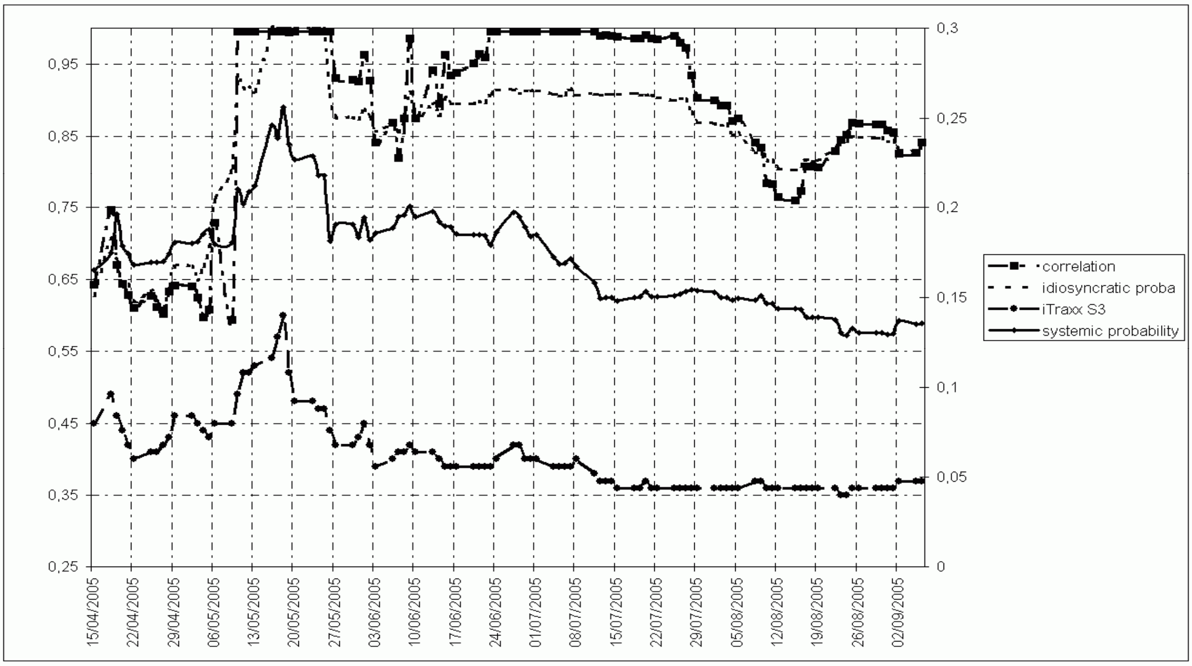

Figure 1. Time series of 5Y iTraxx series 3 calibrated systemic (q_s), idiosyncratic (q) and correlation (ρ^2) parameters. Systemic probability (q_s) on right axis.

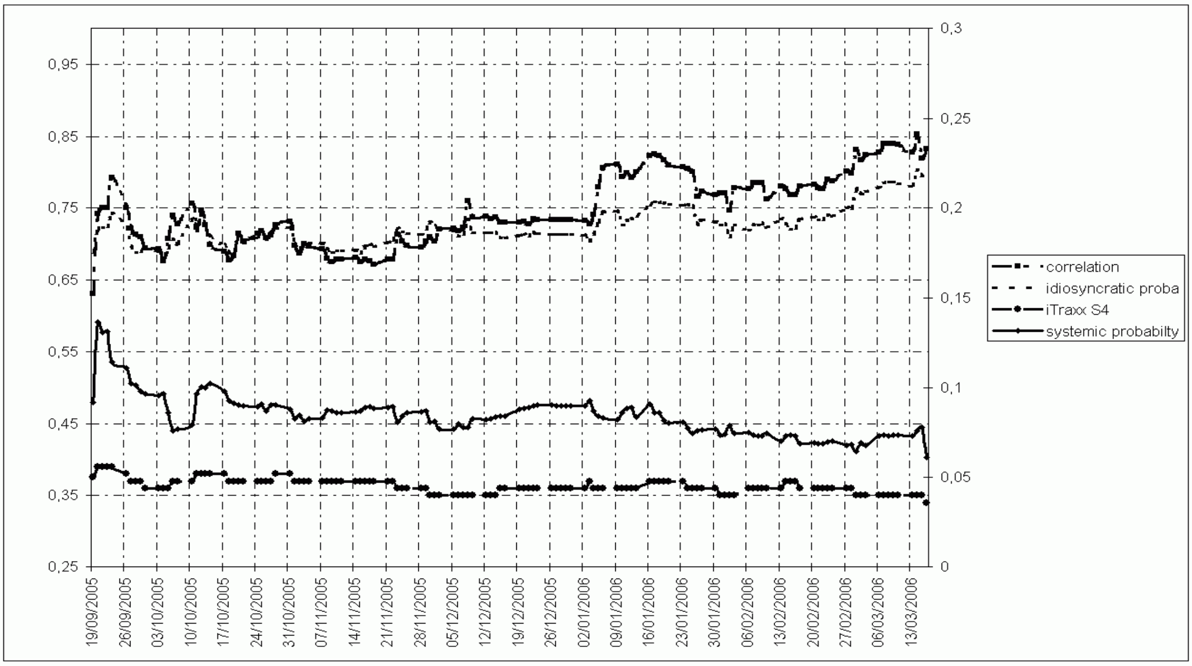

Figure 2. Time series of 5Y iTraxx series 4 calibrated systemic (q_s), idiosyncratic (q) and correlation (ρ^2) parameters. Systemic probability (q_s) on right axis.

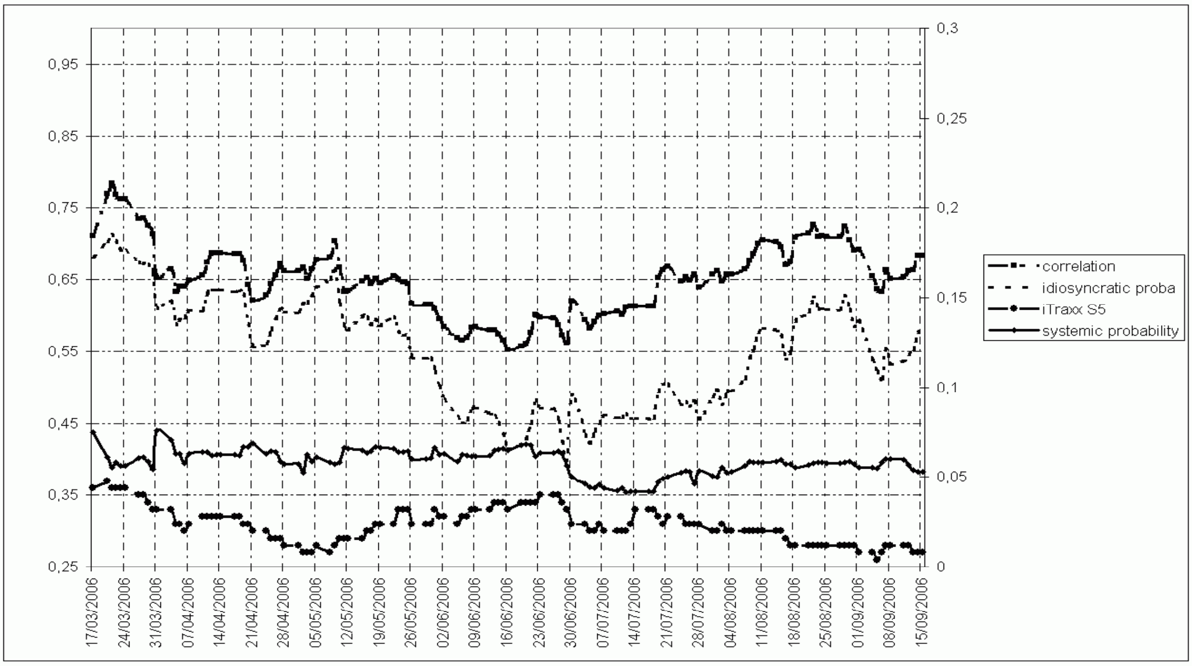

Figure 3. Time series of 5Y iTraxx series 5 calibrated systemic (q_s), idiosyncratic (q) and correlation (ρ^2) parameters. Systemic probability (q_s) on right axis.

Figures 2 and 3 represent the same calibrated parameters, idiosyncratic risk, q, systemic risk, q_s and correlation, ρ^2 from the 19th of September 2005 to the 16th of March 2006 (Figure 2) and from the 17th of March 2006 to the 15th of September 2006 (Figure 3). This is associated respectively with the 5 year iTraxx index series 4 and the 5 year iTraxx index series 5 that are also plotted on the graphs. It can be seen that the systemic probability q_s kept going down until recently. This shows the impact of leverage super senior issuance. Idiosyncratic probability (q) and correlation parameter (ρ^2) moved in intermediate ranges with a peak in March 2006.

We also computed the correlations between the calibrated parameters and the relevant iTraxx index. This is reported in Table 8. We can notice that there is a high degree of positive dependence between the correlation parameter ρ^2 and the idiosyncratic probability parameter q. The index level and the systemic probability parameter q_s were also highly correlated (iTraxx series 3 and 4) though we could see a decrease over the more recent period (iTraxx series 5). Other pairwise correlations may look large for some sub-periods, but without any stability through time. The previous statistical analysis may be helpful for risk management purpose and capital reserve policy. It also gives some insights about the specification of a fully dynamic credit risk model.

	S3	S4	S5
$\rho(F(t), \rho^2)$	2%	-38%	-13%
$\rho(F(t), q)$	10%	-41%	-10%
$\rho(F(t), q_s)$	87%	72%	23%
$\rho(q, \rho^2)$	94%	91%	86%
$\rho(q, q_s)$	30%	-43%	34%
$\rho(\rho^2, q_s)$	23%	-43%	-1%

Table 8. Dynamics of calibrated parameters. iTraxx series.

2.3 Credit deltas

Figure 4 plots the credit deltas for the different tranches in the iTraxx market on 31-August-2005. The credit delta of a given tranche with respect to one name is computed the following way: we shift the 5 year credit spread of that name, compute the change in present value for the corresponding CDS, say ΔPV_i and the change in present value for the CDO tranche, say ΔPV_{CDO}. The nominal is equal to $n = 125$. The credit delta for the CDO tranche is then computed as $\frac{\Delta PV_{CDO}}{\Delta PV_i}$. This provides the amount of CDS to be held in order to hedge the present value of the CDO tranche against small changes in credit spreads, provided that the parameters ρ^2, q_s, q remain unchanged.

Not surprisingly the deltas are larger for the equity tranche since other tranches are call spreads. Also credit deltas of the [3-6%] tranche are larger than those of the [6-9%] and so on which is not surprising either since the latter tranche is more out of the money than the former. We can notice that the credit deltas of the equity tranche increase with the level of name spreads: once again, this can be easily understood since higher spread names are likely to default first and contribute more to the present value of the equity tranche. Conversely, the credit deltas of the less junior tranches tend to decrease with the level of the credit spreads. We can notice rather large bumps in the shape of the credit deltas. This can be explained as a consequence of the comonotonic state in a heterogeneous credit spread framework. It can be easily checked that the credit deltas are a weighted average of the deltas within two simpler stochastic correlation models, that is $(1-q_s)\delta_{\tilde{\rho}_i=(1-B_i)\rho} + q_s\delta_{\tilde{\rho}_i=1}$; $\delta_{\tilde{\rho}_i=(1-B_i)\rho}$ corresponds to the vector of credit deltas when the stochastic correlations are equal to $\tilde{\rho}_i = (1-B_i)\rho$, $i = 1, \ldots, n$ (which is a two states stochastic correlation model) and $\delta_{\tilde{\rho}_i=1}$ corresponds to the vector of credit deltas with comonotonic default dates. In a comonotonic model, defaults dates are perfectly ordered, the name with the highest spread defaulting first, the second default being associated with the second highest spread and so on. As a consequence, a first to default swap is associated with a delta equal to 100% for the name with the highest spread and a zero delta for the other names. A second to default swap is associated with a delta equal to 100% with respect to the second highest credit spread name and a zero delta for the other names. Considering that a CDO tranche with a nominal equal to the number of names is similar to the sum of k-th to default swaps, in a comonotonic model, the corresponding credit deltas are equal to 100% within a given range of credit spreads and equal to zero elsewhere. Having in mind that the systemic probability was equal to 13% on 31-August-2005, this leads to a discontinuity in the credit deltas of the same magnitude as can be seen in Figure 4. Though the irregular shape of the credit deltas can be well-explained, this is clearly an unpleasant feature of the model as far as risk management is concerned. Let us remark that this irregularity would vanish if all credit spreads were equal (homogeneous portfolio). With this in mind, we can slightly adapt the stochastic correlation model for it to behave homogeneously in the comonotonic state[17]. This small adjustment to the stochastic correlation model is associated with smooth deltas and does not lead to extra computational difficulty.

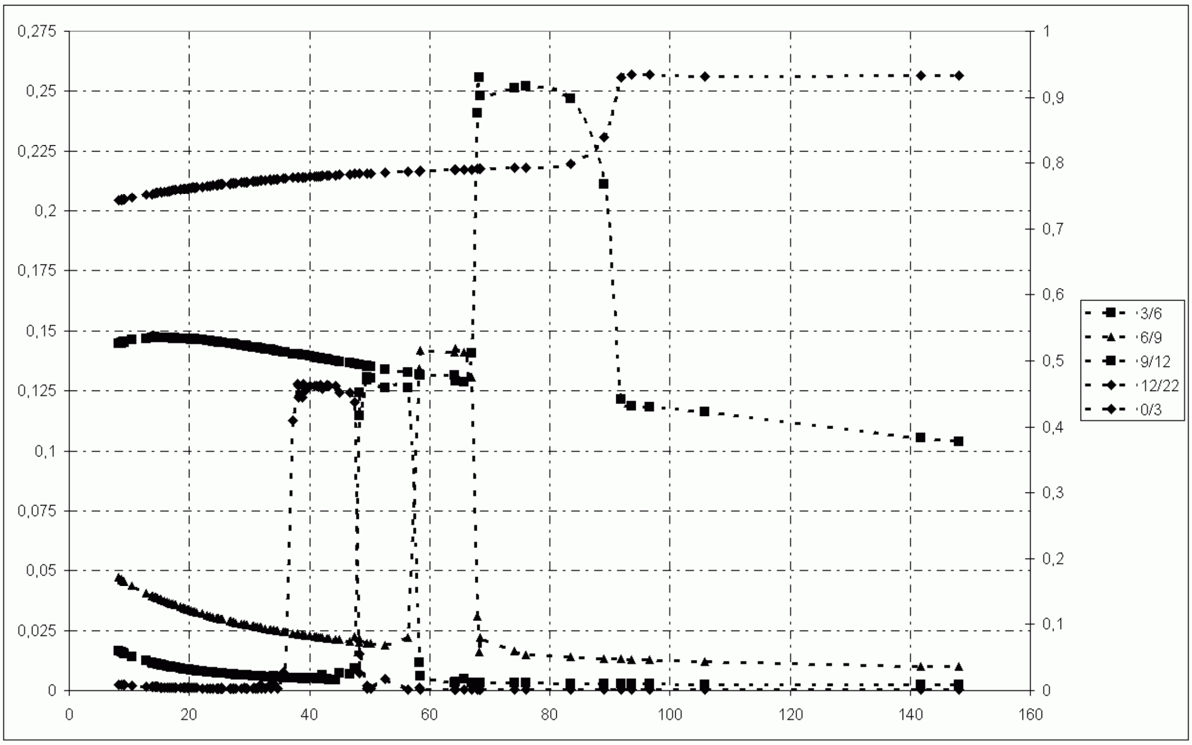

[17]Default dates are then defined as $\tau_i = (1-B_s)H_i^{-1}(\Phi(V_i)) + B_s F^{-1}(\Phi(V))$ where $H_i(t) = \frac{F_i(t)-q_s F(t)}{1-q_s}$, $i = 1, \ldots, n$, $F(t) = \frac{1}{n}\sum_{i=1}^{n} F_i(t)$ and $V_i = (1-B_i)\rho V + \sqrt{1-(1-B_i)\rho^2}\bar{V}_i$, $i = 1, \ldots, n$. It can be checked that this leads to a proper model provided that q_s is not too large.

Figure 4. CDO tranche deltas with respect to the level of credit spreads computed on 31-August-2005. Nominal is equal to 125. 5 year credit spreads on the x axis are expressed in bp pa. Credit deltas of the equity tranche on right axis.

3 Large homogeneous portfolio approximations

Starting from Vasicek [1987], it has been shown that loss distributions could be greatly simplified in the case of large homogeneous portfolios and a factor framework. Under the homogeneity assumption, i.e. $F_1(t) = \ldots = F_n(t) = F(t)$, $\delta_1 = \ldots = \delta_n = \delta$[18] and provided that the number of names n is large, the aggregate loss $L(t) = (1-\delta) \times \frac{1}{n}\sum_{i=1}^{n} 1_{\tau_i \le t}$ has the same distribution as $(1-\delta)p_t^{\cdot|V,B_s}$[19]. As a consequence, up to a scaling factor, the distribution of the losses within a large homogeneous portfolio and the distribution of the conditional default probability are the same.

3.1 Large homogeneous portfolio distributions

Within the stochastic correlation model, the distribution function of conditional default probabilities G is given by[20]: $p \in [0,1] \to G(p) = Q(p_t^{\cdot|V,B_s} \le p)$ and:

$$G(p) = (1-q_s)\left(1_{qF(t)<p<1-q+qF(t)}\Phi\left(\frac{1}{\rho}\left(\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{p-qF(t)}{1-q}\right)-K_t\right)\right)+1_{p\ge 1-q+qF(t)}\right)+q_s\Phi(-K_t), \quad (3.7)$$

for $0 \le p < 1$, where $q_s = Q(B_s = 1)$ and $q = Q(B_i = 1)$, $K_t = \Phi^{-1}(F(t))$. 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: $\psi_L(u) = \int_0^1 \left(pu^{1-\delta}+1-p\right)^n dG(p)$.

We have $Q\left(p_t^{\cdot|V,B_s} = 0\right) = q_s\Phi(-K_t) = q_s(1-F(t))$ and $Q\left(p_t^{\cdot|V,B_s} = 1\right) = q_s\Phi(K_t) = q_sF(t)$. 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 probabilities[21]. Thus, increasing q_s should result in increasing the premiums of equity and senior tranches and decreasing the premiums of mezzanine tranches, leading to a more pronounced correlation smile.

The distribution function is constant between 0 and $qF(t)$ and between $1-q+qF(t)$ and 1[22]. 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

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

[19] For notational simplicity, we will further remove the name dependence in the superscripts.

[20] The proof is detailed in appendix A.

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

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

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 at $qF(t)$ 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).

3.2 Zero-coupon CDO tranche premiums

We can provide a closed form expression of the premium of a zero-coupon CDO in the large homogeneous portfolio approximation. Indeed, for an option on the aggregate loss, with maturity t, exercise price k we can express the up-front premium as :

$$q_s(1-k)F(t) + (1-q_s)(1-q)\left(\Phi_{2,\rho}(\Phi^{-1}(F(t)), V^*) - k^*\Phi(V^*)\right) \tag{3.8}$$

for $qF(t) < k < 1 - q + qF(t)$ where:

$$k^* = \frac{k - qF(t)}{1-q}, \quad V^* = \frac{1}{\rho}\left(\Phi^{-1}(F(t)) - \sqrt{1-\rho^2}\,\Phi^{-1}(k^*)\right), \tag{3.9}$$

and $\Phi_{2,\rho}$ is the bivariate Gaussian distribution function with correlation parameter ρ. The proof is postponed to Appendix B.

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

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

with $K_t = \Phi^{-1}(F(t))$. This result can be used for the computation of risk measures which, as the Expected Shortfall, involve weighted averages of quantiles of the loss distribution. Let us notice that combining the previous results, we can derive the Expected Shortfall of the aggregate loss in closed-form when using the large homogeneous portfolio approximation.

3.3 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[23]. 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. In this subsection, we firstly review the main features of marginal compound correlation. We then show that the marginal compound correlation can be computed analytically within the stochastic correlation model.

Let us consider a mezzanine tranche with attachment-detachment points of α and $\alpha + \eta$. Neglecting discounting effects, the upfront premium is given by $\frac{1}{\eta}\int_\alpha^{\alpha+\eta}(1 - G_L(u))\,du$ where G_L denotes the aggregate

[23] 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.

loss distribution function. 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 $1 - G_L(\alpha)$[24]. On the other hand, for a large homogeneous portfolio, the distribution function of the aggregate losses can be approximated by the distribution function of the conditional default probability G (see previous subsection). 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 is given by: $1 - \Phi\left(\frac{\sqrt{1-\rho^2}\Phi^{-1}\left(\frac{\alpha}{1-\delta}\right) - \Phi^{-1}(F(t))}{|\rho|}\right)$[25]. The *marginal compound correlation* for that mezzanine tranche is denoted by $\bar{\rho}(\alpha) \in [0, 1]$[26] and is such that:

$$G\left(\frac{\alpha}{1-\delta}\right) = \Phi\left(\frac{\sqrt{1-\bar{\rho}(\alpha)^2}\Phi^{-1}\left(\frac{\alpha}{1-\delta}\right) - \Phi^{-1}(F(t))}{\bar{\rho}(\alpha)}\right). \quad (3.10)$$

$\bar{\rho}(\alpha)$ is the correlation parameter to be put in the Gaussian copula model in order to price correctly the previous $[\alpha, \alpha]$ mezzanine tranche.

In Appendix C, we show that for any given $\alpha \in [0, 1]$, $\bar{\rho}(\alpha)$ can be analytically computed from equation (3.10). We need to solve for a second order equation. Thus, there are at most two marginal compound correlations $\bar{\rho}(\alpha)$[27]. 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.

4 Stochastic and state dependent correlation

In order to give a broader perspective on the three states stochastic correlation model, we compare it with another popular extension of the one factor Gaussian copula, namely the random factor loading model (Andersen & Sidenius [2005a]). For the paper to be self-contained, we firstly review the main properties of the model. We direct the reader to Andersen & Sidenius [2005a] for further details. We then compare large homogeneous portfolio distributions and marginal compound correlation curves. Eventually, we compute the local correlation curves associated with the two models.

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

[25] In the case of the Gaussian copula, we can easily derive G as:

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

.

[26] 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.

[27] Since these are solution of the second order equation (5.1).

4.1 Random factor loadings

We can think of the following extension of the Gaussian copula, belonging to the class of one factor mean variance Gaussian mixtures, where the latent variables are modelled by:

$$V_i = m_i(V) + \sigma_i(V)\bar{V}_i, \quad i = 1, \dots, n. \tag{4.1}$$

$V, \bar{V}_i$ are independent standard 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]).

For simplicity, we consider from now-on that the credit spreads are equal, i.e. $F_1 = \ldots = F_n = F$. In the random factor loading model introduced by Andersen & Sidenius [2005a], $\sigma_i(V)$ is constant, while in the local correlation model of Turc et al [2005], $\sigma_i(V) = \sqrt{1 - m_i(V)^2}$.

The simplest form of random factor loadings is associated with the following parametric modelling of latent variables:

$$V_i = m + \left(l 1_{V<e} + h 1_{V \geq e}\right) V + \nu \bar{V}_i, \quad i = 1, \dots, n \tag{4.2}$$

where $V, \bar{V}_1, \dots, \bar{V}_n$ are independent standard Gaussian random variables, l, h, e some input parameters, $l, h > 0$. This can be seen a random factor loading model, since the risk exposure $l 1_{V<e} + h 1_{V \geq e}$ is state dependent. m and ν are such that $E\left[V_i\right] = 0$ 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 cumulative density function, φ the Gaussian density function. The marginal distribution of V_i, that we denote by H_{RFL}, involves a bivariate Gaussian cumulative density function:

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

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

$$p_t^{i|V} = Q(\tau_i \leq t \mid V) = \Phi\left(\frac{1}{\nu}\left(H_{RFL}^{-1}(F(t)) - m - \left(l 1_{V \leq e} + h 1_{V>e}\right) V\right)\right). \tag{4.3}$$

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

$$G(p) = Q(p_t^{i|V} \leq p) = \Phi\left(\min\left(\frac{z(p)}{h}, -e\right)\right) + \left(\Phi\left(\frac{z(p)}{l}\right) - \Phi(-e)\right) 1_{z(p) > -el}, \tag{4.4}$$

for $0 < p < 1$, where $z(p) = \nu\Phi^{-1}(p) - H_{RFL}^{-1}(F(t)) + m$[28].

[28] Let us assume that $e < 0$ and $h < l$. We can then write G^{-1} as:

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

As for the stochastic correlation model, this result is useful for the computation of risk measures.

4.2 Comparison of large homogeneous portfolio distributions

For further comparisons, we plot the distribution functions of the conditional default probabilities associated with the stochastic correlation and the random factor loading models (Figure 5). We recall that these also correspond to the distribution functions of large homogeneous portfolios. 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 [2005a]) and $H^{-1}_{RFL}(F(t)) = -1.886$. The graph also shows the distribution function associated with the independence case.

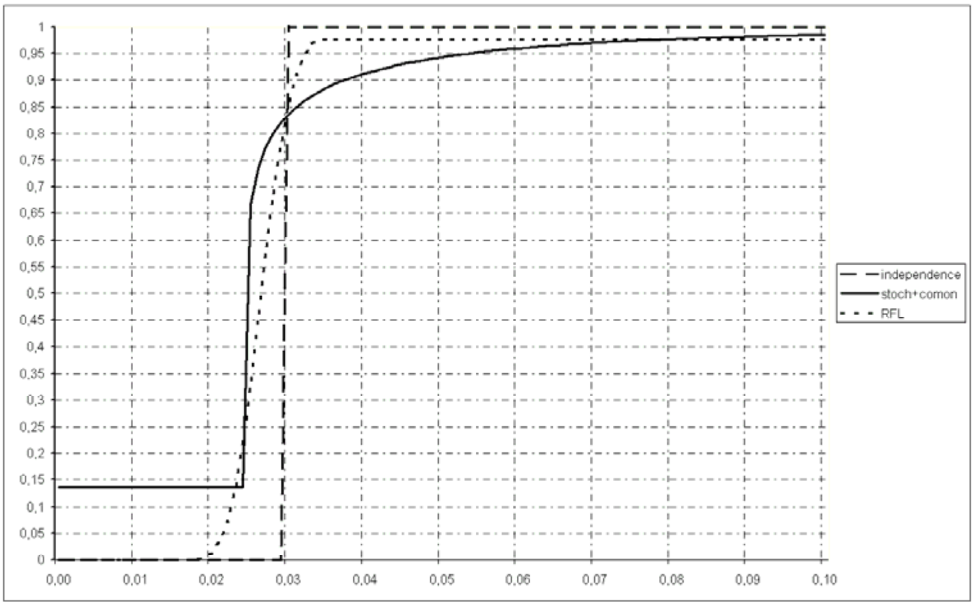

Figure 5. Distribution functions of conditional default probabilities.

Using the previous data usually provides 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 or a [0,1%] equity tranche would be priced quite differently.

4.3 Marginal compound correlation curves

We now show the marginal compound correlation curves associated with the three parameters stochastic correlation and random factor loading models, with the same set of parameters as above (Figure 6). The recovery rate is equal to 40%.

The marginal compound correlations exhibit a smile feature, with a rather strong peak in the middle. Let us recall that 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 $p = (1-\delta)F(t)$. Both models exhibit this feature. 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.

Let us remark some difference in the region of first losses where 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.

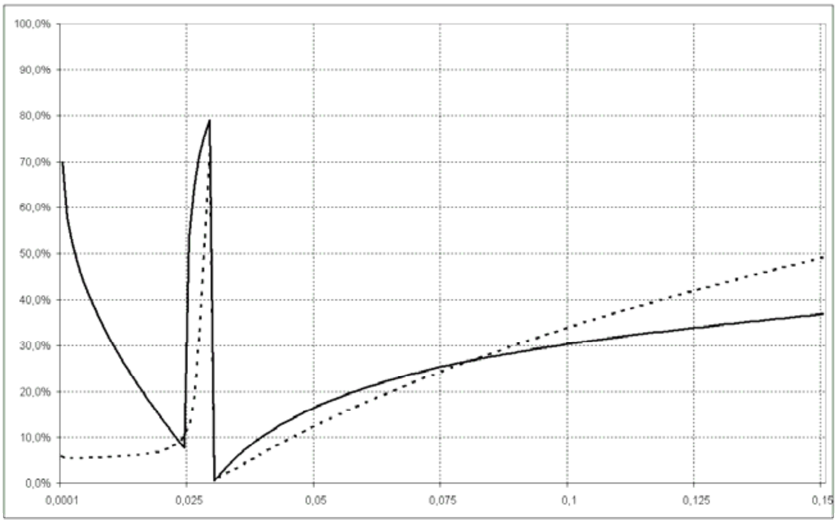

Figure 6. Marginal compound correlation (vertical axis) associated with stochastic correlation and random factor loading models with respect to the level of attachment - detachment point (horizontal axis)

4.4 Local correlation curves associated with the two models

We show here that the concept of local correlation introduced by Turc et al. [2005] can be used for the purpose of comparing stochastic correlation and random factor loading models. In Turc et al. [2005] framework, the latent variables are given by:

$$V_i = -\rho(V)V + \sqrt{1-\rho^2(V)}\bar{V}_i, \quad i = 1, \dots, n \tag{4.5}$$

where $V, \bar{V}_1, \dots, \bar{V}_n$ are independent standard Gaussian random variables and $\rho(.)$ is some function of V taking values in $[0,1]$, known as the local correlation. By conditioning upon V, we get the marginal distribution of V_i:

$$H(x) = Q(V_i \le x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho(v)v}{\sqrt{1-\rho^2(v)}}\right)\varphi(v)dv \tag{4.6}$$

for $x \in \mathbb{R}$, where Φ is the Gaussian cumulative density function, φ the Gaussian density function. The conditional default probabilities can be written as:

$$p_t^{i|V} = Q(\tau_i \leq t \mid V) = \Phi\left(\frac{H^{-1}(F(t)) + \rho(V)V}{\sqrt{1-\rho^2(V)}}\right). \quad (4.7)$$

The local correlation can be used in a way which parallels the local volatility modelling in the equity derivatives market: this consists in a non parametric calibration of $\rho(.)$ on CDO tranche premiums computed under the three states stochastic correlation model (respectively the random factor loading model). Let us denote by G the distribution function of the conditional default probabilities in the three states stochastic correlation model (respectively the random factor loading model). To achieve the non parametric calibration, we must look for $\rho(.)$ such that the distribution function of $p_t^{i|V}$ in equation (4.7) is equal to G. Since only conditional default probability distributions are involved in the computation of CDO tranches, this procedure guarantees a perfect calibration to CDO tranche premiums computed under the three states stochastic correlation model (respectively the random factor loading model).

The previous procedure leads to the following equation involving $\rho(.)$:

$$\frac{H^{-1}(F(t)) + \rho(v)v}{\sqrt{1-\rho^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right), \ \forall v \in \mathbb{R}, \quad (4.8)$$

with: $H(x) = \int_{\mathbb{R}} \Phi\left(\frac{x+\rho(v)v}{\sqrt{1-\rho^2(v)}}\right)\varphi(v)dv$[29]. It can be shown that the local correlation can be viewed as a rescaled marginal compound correlation. Technical details about the previous inverse problem can be found in the paper by Turc et al [2005] and are provided with further discussion in Appendix D.

Figure 7 plots the local correlation curve associated with the 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. We have two local correlations for values of V approximately between 1 and 2, leading to the same CDO prices.

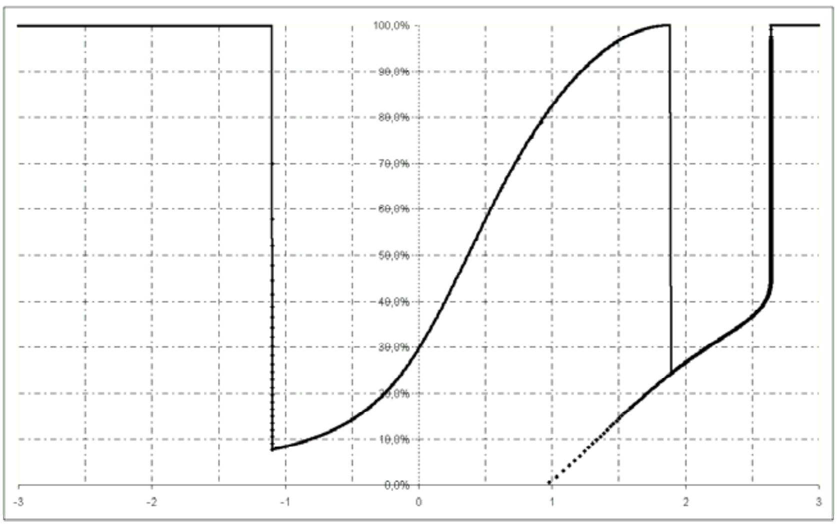

[29] Let us remark that the right-hand term is given while the left-hand term involves ρ either directly or indirectly through H.

Figure 7. Local correlation (vertical axis) associated with the stochastic correlation model with respect to the level of the latent factor V (horizontal axis)

Figure 8 plots the local correlation curve associated with 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 $V = 2$. As in the previous case, when V is in between 1 and 2, there are two local correlations.

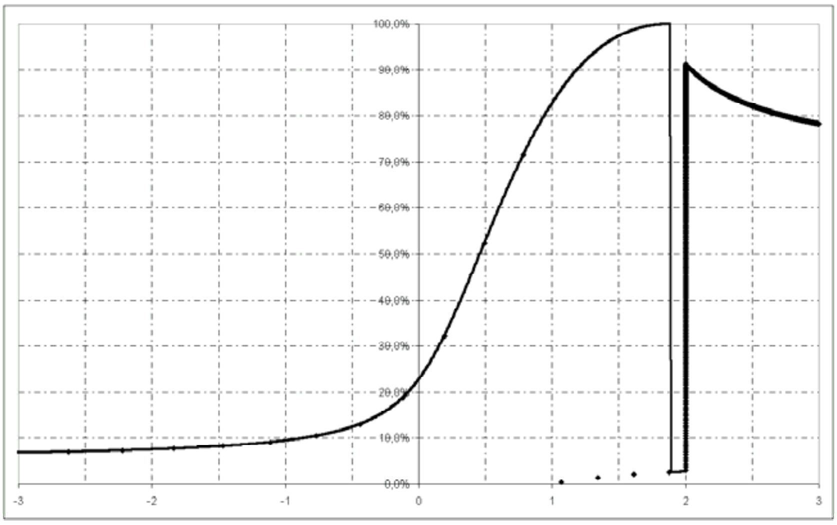

Figure 8. Local correlation (vertical axis) associated with the random factor loading model with respect to the level of the latent factor V (horizontal axis)

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 parametrize a 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, 2005a, *Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings*, Journal of Credit Risk, 1(1).

Andersen, L. & J. Sidenius, 2005b, *CDO Pricing with Factor Models: Survey and Comments*, Journal of Credit Risk, 1(3).

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.

Gregory, J. & J-P. Laurent, 2004, *In the Core of Correlation*, *RISK*, October, 87-91.

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.

Li, D.X., M. Liang, 2005, *CDO Squared Pricing Using Gaussian Mixture Model with Transformation of Loss Distribution*, working paper available on SSRN.

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

Piterbarg, V., 2003, *Mixture of Models: A Simple Recipe for a... Hangover?*, working paper available on SSRN.

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

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, *Probability of Loss on Loan Portfolio*, working paper, Moody's KMV.

Appendix A: distribution of conditional default probabilities within the stochastic correlation model

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

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.

Appendix B: computation of zero-coupon CDO premiums

We detail here the computation of zero-coupon CDO premiums using the large homogeneous portfolio approximation. We need to compute $\int (p-k)^+ dG(p) = E\left[\left(p_t^{\cdot|V,B_s}-k\right)^+\right] = q_s E\left[\left(p_t^{\cdot|V,1}-k\right)^+\right] + (1-q_s)E\left[\left(p_t^{\cdot|V,0}-k\right)^+\right]$.

Since $p_t^{\cdot|V,1} = 1_{V\leq\Phi^{-1}(F(t))}$, $E\left[\left(p_t^{\cdot|V,1}-k\right)^+\right] = (1-k)F(t)$. Let us now compute $E\left[\left(p_t^{\cdot|V,0}-k\right)^+\right]$ where $p_t^{\cdot|V,0} = (1-q)\Phi\left(\frac{\Phi^{-1}(F(t))-\rho V}{\sqrt{1-\rho^2}}\right)+qF(t)$. This can be written as $(1-q)E\left[\left(\Phi\left(\frac{\Phi^{-1}(F(t))-\rho V}{\sqrt{1-\rho^2}}\right)-k^*\right)^+\right]$

with $k^* = \frac{k-qF(t)}{1-q}$. We will further focus on the case where $0 < k^* < 1$ or equivalently where $qF(t) < k < 1-q+qF(t)$, the other cases leading to straightforward computations. $\Phi\left(\frac{\Phi^{-1}(F(t))-\rho V}{\sqrt{1-\rho^2}}\right) > k^*$ if and only if $V < V^* = \frac{1}{\rho}\left(\Phi^{-1}(F(t)) - \sqrt{1-\rho^2}\Phi^{-1}(k^*)\right)$. Let us remark that $\Phi\left(\frac{\Phi^{-1}(F(t))-\rho V}{\sqrt{1-\rho^2}}\right)$ can be written as $E\left[1_{W_i \le \Phi^{-1}(F(t))} \mid V\right]$ where $W_i = \rho V + \sqrt{1-\rho^2}\bar{V}_i$. As a consequence,

$$E\left[\left(\Phi\left(\frac{\Phi^{-1}\left(F(t)\right) - \rho V}{\sqrt{1-\rho^2}}\right) - k^*\right)^+\right] = \Phi_{2,\rho}(\Phi^{-1}(F(t)), V^*) - k^*\Phi(V^*)$$

Appendix C: computation of the marginal compound correlation

Let us 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[30]. For $0.5 < \alpha < 0.8$, $G(\alpha) = 0.5$, thus $G(\alpha) < 1 - F(t)$ and there is then no solution to equation (3.10)[31]. Unfortunately, marginal compound correlation may not exist for consistent pricing models.

Equation (3.10) can equivalently be written as:

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

We recall that for notational simplicity, $\delta = 0$. By squaring up the previous equality, we obtain the following second order equation, from which we must consider the positive roots:

$$\left(\Phi^{-1}\left(G(\alpha)\right)^2 + \Phi^{-1}(\alpha)^2\right)\bar{\rho}^2 + 2\Phi^{-1}\left(G(\alpha)\right)\Phi^{-1}\left(F(t)\right)\bar{\rho} + \Phi^{-1}\left(F(t)\right)^2 - \Phi^{-1}(\alpha)^2 = 0. \quad (5.1)$$

Let us emphasize that any solution of:

$$\bar{\rho}(\alpha)\Phi^{-1}\left(G(\alpha)\right) + \Phi^{-1}\left(F(t)\right) = -\sqrt{1-\bar{\rho}(\alpha)^2}\Phi^{-1}(\alpha), \quad (5.2)$$

is also a solution of equation (5.1). The roots of equation (5.1) are the roots of equation (3.10) plus the roots of equation (5.2). In the numerical examples below, we firstly write the roots of equation (5.1) and then check that they indeed solve equation (3.10).

Let us discuss the existence of a solution to equation (5.1). We denote by:

$$\Delta' = \Phi^{-1}(\alpha)^2 \times \left(\Phi^{-1}\left(G(\alpha)\right)^2 + \Phi^{-1}(\alpha)^2 - \Phi^{-1}\left(F(t)\right)^2\right).$$

Equation (5.1) has real roots if and only if $\Delta' \ge 0$ but this may not be the case as illustrated by the following example. We assume that $p_t^{i|V}$ is discretely distributed with probability 0.5 to be equal to 0 and 0.5 to be

[30] 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.

[31] The right hand term of equation (3.10) decreases from 1 to $1 - F(t)$ as $\bar{\rho}(\alpha)$ increases from 0 to 1.

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 (5.1) 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 previously showed some practical cases where they were actually two positive roots to equation (5.1) that appear to be also solutions of equation (3.10).

Appendix D: derivation of the local correlation function

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 (4.7), we have $p_t^{i|V} = \Phi\left(\varepsilon(V)\right)$. Thus, 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:

$$\frac{H^{-1}(F(t)) + \rho(v)v}{\sqrt{1-\rho^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right), \ \forall v \in \mathbb{R}, \tag{6.1}$$

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

$$\frac{\Phi^{-1}(F(t)) + \rho_1(v)v}{\sqrt{1-\rho_1^2(v)}} = \Phi^{-1}\left(G^{-1}\left(\Phi(v)\right)\right). \tag{6.2}$$

which provides ρ_1[32]. 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 $H_n = H_{n-1}$[33], the algorithm has converged and we set $\rho = \rho_n$.

Equation (3.10) 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 (6.2), we get $\bar{\rho}(p) = \rho_1(v)$, which

[32] 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.

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

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.

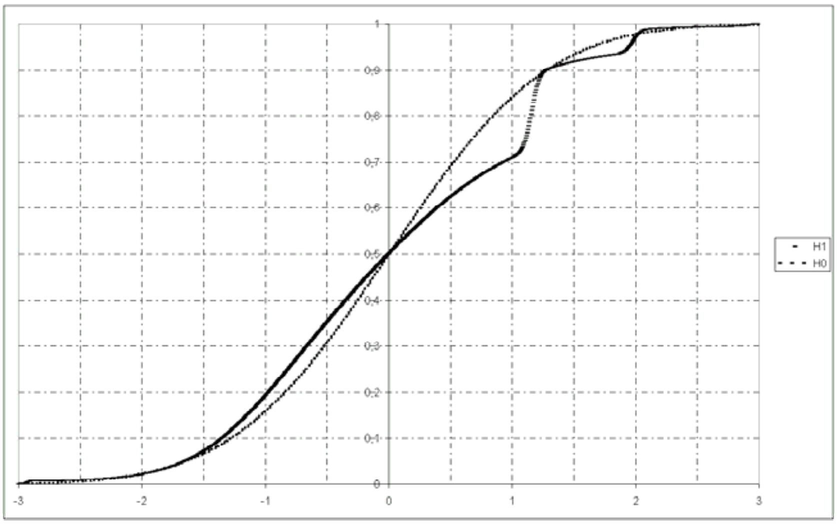

Figure 9. Check of convergence of the fixed point algorithm (stochastic correlation model)

Figure 9 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.