Remarks on Pricing Correlation Products

Harald Skarke[1]
Bank Austria Creditanstalt
Dept 8641, P.O. Box 35
1011 Vienna, Austria

July 17, 2005

Abstract

The standard approach for the risk neutral pricing of correlation products such as CDOs is to apply a one factor Gaussian copula model with certain variations. This model can be expressed as an average over scenarios where, conditional on a parameter, default time distributions for different obligors are independent. Using this description I show that the Gaussian copula approach leads to certain predictions concerning relative changes of default probabilities that are not only counterintuitive but also in disagreement with market data. These problems can be avoided within the framework of conditional independence by changing the forms of the conditional default time probability distributions away from what is dictated by the Gaussian copula assumption. This leads to a class of models that do not require modifications such as using the base correlation technique; this class contains the model recently suggested by Hull and White.

[1]harald.skarke@ba-ca.com

Remarks on pricing correlation products

Harald Skarke

1 General considerations and the present market standard

For the pricing of a structured product involving several credits, such as a tranche of a synthetic CDO or CDO–squared, or an n-th to default basket, it is not sufficient to have only the spread curves for the individual obligors. In addition, one has to account for the fact that the patterns of defaults and spread movements of the individual credits are not independent. With tranched indices such as iTraxx and CDX, one now has information on the market perception of these dependencies. It is an important problem to implement this data in a consistent way into a risk neutral pricing scheme.

In the following I shall make a number of simplifying assumptions. I ignore recovery rates – this is justified, for example, if we assume a fixed recovery for every CDS. I also assume that a single tranched index is used.

For the pricing of a correlation product it is sufficient to have a joint probability distribution $\mathbf{F}(t_1, \ldots, t_N)$ for the default times $t_1, \ldots, t_N$ of the N credit default swaps referenced by the deal. The pricing then becomes a two step process,

$$\left.\begin{array}{l}\text{Individual data}\\ \text{Correlation data}\end{array}\right\} \longrightarrow \text{ Joint distribution } \longrightarrow \text{ Price} \qquad (1)$$

The second step (distribution → price) is the standard calculation of expected values on loss and premium legs and will not be considered here. Note that the current standard practice for quoting CDO correlation data, namely the base correlation technique, is not of the type (1) as it works without having a generally valid joint distribution. A method for constructing such a distribution can lead to a consistent pricing method only if it obeys the following conditions.

- Consistency with individual data: the marginal distributions of $\mathbf{F}$ should be the distributions $F_1(t_1), \ldots, F_N(t_N)$ implied by the spread curves (and recovery assumptions) corresponding to obligors $1, \ldots, N$.
- Consistency with correlation data: if the method is applied to the index portfolio, it should lead to the correct index tranche prices.

The first of these conditions is equivalent to the statement that

$$\tilde{\mathbf{F}}(u_1, \ldots, u_N) := \mathbf{F}(F_1^{-1}(u_1), \ldots, F_N^{-1}(u_N)) \qquad (2)$$

is a copula function; conversely, a copula function $\tilde{\mathbf{F}}$ implies a distribution $\mathbf{F}$ via $\mathbf{F}(t_1, \ldots, t_N) = \tilde{\mathbf{F}}(F_1(t_1), \ldots, F_N(t_N))$. A standard approach for constructing copula functions is to take some well known multivariate distribution together with its marginal distributions and then to apply a transformation like the one above. The most commonly used copula function of this type is the Gaussian copula

$$\tilde{\mathbf{F}}_{\text{Gaussian}}(u_1, \ldots, u_N) := \mathbf{\Phi}_{\mathbf{0},\Sigma}(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_N)), \tag{3}$$

where Φ denotes the normalized Gaussian distribution and $\mathbf{\Phi}_{\mathbf{0},\Sigma}$ is the multivariate Gaussian with mean zero and correlation matrix Σ. For the pricing of CDO tranches, this approach is further narrowed down to using a correlation matrix where all mutual correlation numbers (i.e., all off-diagonal entries of Σ) have the same value ρ. In this case it turns out that the joint distribution $\mathbf{F}$ can be described in terms of a one parameter family of scenarios such that in each specific scenario the distributions for the individual oligors are independent:

$$d\mathbf{F} = \int dF_1(t_1|V) \ldots dF_N(t_N|V)\; dG(V) \tag{4}$$

(the integral referring to $dG(V)$ only). The conditional probabilities

$$F_i(t_i|V) = \Phi\left(\frac{\Phi^{-1}(F_i(t_i)) - \sqrt{\rho}V}{\sqrt{1-\rho}}\right) \tag{5}$$

depend on a random variable V that is normally distributed,

$$G(V) = \Phi(V). \tag{6}$$

If we work with a fixed time horizon, e.g. the termination date T of the deal that we want to price, we can interpret V as a variable describing general economic conditions during the interval $[0, T]$ in such a way that a large V means good economic conditions, hence low default probabilities; i.e., the $F_i(t_i|V)$ are monotonically decreasing in V.

The strength of the method lies in the fact that due to the conditional independence of the $F_i(t_i|V)$ as exhibited in formula (4) efficient numerical methods are available. On the other hand this approach, as it is, simply does not work. With any choice of the correlation ρ at least some of the index tranches are priced very wrongly. Currently, there seem to be two major strategies for overcoming this problem. One of them uses the concept of 'base correlation' for the pricing of CDO tranches [3]. In this approach only tranches with an attachment point of zero are priced according to the scheme outlined above, with different correlation levels for different detachment points. A tranche with attachment point a and detachment point d is then priced as the difference between two tranches with attachment point 0 and detachment points d and a, respectively. While this seems to give reasonable results for a non–standard tranche of a portfolio that is very similar to the index portfolio, it is not quite clear how the base correlation data should be mapped between portfolios with different characteristics. Another disadvantage of base correlation is the fact that it uses explicitly the structure of the CDO in such a way that it cannot work for other structures such as, e.g., CDO-squareds. A different approach [1] introduces a number of extra parameters to make the correlation V–dependent; in this way it is possible to reproduce what is known as the correlation smile or skew.

2 Further problems with the Gaussian copula

Let us now explore the following questions. What can we say about eq. (4) in general? And given eq. (4), what are the implications of choosing the particular forms of eqs. (5) and (6)? Applying the first of the two consistency conditions given above to eq. (4) leads to

$$\int F_i(t_i|V)dG(V) = F_i(t_i). \tag{7}$$

Note that the shape of the function $G(V)$ has no meaning per se since we can always transform to a different function $\tilde{G}(\tilde{V})$ via using $\tilde{V} = \tilde{G}^{-1}(G(V))$, provided we change $F_i(t_i|V)$ to

$$\tilde{F}_i(t_i|\tilde{V}) = F_i(t_i|V(\tilde{V})) \quad \text{with} \quad V(\tilde{V}) = G^{-1}(\tilde{G}(\tilde{V})). \tag{8}$$

Consider the default probability $F_i(T|V)$ of obligor i during the lifetime $[0, T]$ of the deal, conditional on V. It exhibits a relative change proportional to

$$\beta_i := -\frac{1}{F_i(T|V)}\frac{\partial F_i(T|V)}{\partial V} \tag{9}$$

under a small change of V. I have included the minus sign because of the negative dependency of $F_i(T|V)$ on V (the state of the economy) whereas in general "betas" are rather thought of as the tendency of a certain spread to move with the average spread, which should be negatively correlated with the state of the economy. Given the above argument about the scaling of V, a single β_i is not a very meaningful number per se, but the quotient β_i/β_j for two different credits is independent of such a rescaling and gives a clear indication of which of the default probabilities is expected to change more under a change of V. For the case (5) of the Gaussian copula, one finds

$$\beta_i = \frac{\sqrt{\rho}}{\sqrt{1-\rho}}\phi\left(\frac{\Phi^{-1}(F_i(T)) - \sqrt{\rho}V}{\sqrt{1-\rho}}\right) \Big/ \Phi\left(\frac{\Phi^{-1}(F_i(T)) - \sqrt{\rho}V}{\sqrt{1-\rho}}\right), \tag{10}$$

so the behaviour of β_i is dominated by the function $\phi(x)/\Phi(x)$. This function is actually decreasing, implying that the default probability of a riskier obligor should exhibit a smaller relative change under a change of V than the default probability of a less risky one. For example, the relative change for an obligor with 0.5% default risk should be more than twice that of an obligor with 20% default risk. This is in contradiction to the general belief that high spread companies are more likely to default because of economic conditions whereas low spread companies are expected to default, if they do so at all, because of idiosyncratic risk factors such as fraud.

The behaviour (10) of the Gaussian copula is also in quantitative disagreement with market data, as the following argument shows. I have used data for the iTraxx series 2 from January 2004 to mid March 2005 in the following way. For each of the 125 credit default swaps, I have calculated the regression line for the relative movement of the spread $s_{i,t}$ (i indexing the credit and t denoting time in days) over the relative movement of the average spread of the iTraxx portfolio, using time steps of one week; in formulas, the regression was for

$$y_{i,t} = \frac{s_{i,t} - s_{i,t-7}}{(s_{i,t} + s_{i,t-7})/2} \quad \text{over} \quad x_t = \frac{a_t - a_{t-7}}{(a_t + a_{t-7})/2} \quad \text{with} \quad a_t = \frac{1}{125}\sum_{i=1}^{125} s_{i,t}. \tag{11}$$

Let us call the slopes of these lines empirical betas. It turns out that of the 125 empirical betas, exactly one is negative and one is larger than two, the remaining 123 being in the range (0,2). If we interpret a movement of the average spread of the index roughly as a change in the market participants' perception of the distribution of V, then the application of the Gaussian copula would indicate that high spread names should have, on average, lower betas than low spread names. Exactly the opposite is true. The single obligor with negative beta (of -0.03) has a spread of 11.7 bp, the single obligor with a beta of more than two (2.09, to be precise) has a spread of 105.4 bp. The twenty smallest spread names all have betas less than one, and the four names with spreads larger than 80 bp all have betas greater than one. A linear regression, this time for the empirical betas over the logarithms of time averaged spreads, has a positive slope of 0.36.

These problems are not addressed by the two aforementioned variations on the Gaussian copula approach. While this should not matter too much if the portfolio underlying the deal that is to be priced is very similar to the index portfolio, it severely affects the ability of these approaches to cover wider ranges of, for example, average spreads.

3 Suggestions for improvements

So, how can we improve the situation? Clearly, the values (9) for the betas should become different from those implied by (10). This can only be achieved if we depart from the specific form (5) of the conditional probabilities. The simplest possible modification would be to demand that all betas be equal. This does not work for formula (9) as it stands because of the limitation $F_i(T|V) \in [0,1]$. It does work, however, if we replace the probabilities $F_i(t|V)$ by intensities (hazard rates) $h_i(t|V)$ — for small default probabilities, relative changes in the hazard rates will lead to approximately proportional relative changes in the default probabilites

$$F_i(t) = 1 - e^{-\int_0^t h(\tilde{t})d\tilde{t}} \approx \int_0^t h(\tilde{t})d\tilde{t}, \tag{12}$$

so we do get the desired behaviour. Using our freedom in reparametrizing V, we can set the β_i to 1 which results in

$$h_i(t|V) = h_i^0(t)e^{-V}. \tag{13}$$

Then one has to choose the $h_i^0(t)$ and $G(V)$ in such a way that the consistency condition (7), which becomes

$$1 - \int e^{-\int_0^t h_i^0(\tilde{t})e^{-V}d\tilde{t}} dG(V) = F_i(t), \tag{14}$$

is obeyed, and that the model reproduces the prices of the index tranches. This corresponds more or less to a model that was, unfortunately for the present author, suggested recently by Hull and White [2]. To see this, transform to $\tilde{V} = e^{-V}$, which is what Hull and White denote by λ. Somewhat confusingly, in their paper λ rather than h_i is time dependent, which would not allow to choose the time dependenies for the different obligors' default intensities independently; this would actually lead to a violation of the consistency condition (7).

A simple variation of this is to allow different constant betas for different credits, resulting in

$$h_i(t|V) = h_i^0(t)e^{-\beta_i V}. \tag{15}$$

In that case one could either use historically observed betas (but with caution, given the short time span of observations), or one could make a simple adjustment to take into account the fact that larger spreads tend to lead to larger betas. In general there is of course an infinity of different choices for $h_i(t|V)$ or $F_i(t|V)$ and $G(V)$ that obey the consistency conditions, with many of them performing better than the Gaussian copula; the choices (13) and (15) are just very simple examples.

The calibration of such a model will certainly pose a number of technical problems, but none of them should be too difficult to overcome. Here it may be very useful to exploit the possibility of rescaling V. In practice one will probably choose some very simple functions (e.g., piecewise linear) for the $F_i(t|V)$ or $h_i(t|V)$, and a discrete distribution for V. One such possibility for the latter would be to define V as the number of defaults in the index portfolio, implying that $F_i(t|V)$ is the default time distribution for the i'th credit conditional on the index exhibiting V defaults in $[0, T]$. For models of the types (13) and (15) some iterative procedure, starting with h_i^0 that are just the hazard rates implied from the input data, should work well.

References

[1] Andersen, Leif and Jakob Sidenius, "Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings", The Journal of Credit Risk, Vol 1 Nr. 1 p. 29.

[2] Hull, John and Alan White, "The perfect copula", on `www.defaultrisk.com`.

[3] McGinty, Lee, Eric Beinstein, Rishad Ahluwalia and Martin Watts, "Introducing Base Correlations", Credit Derivatives Strategy, JPMorgan.