Fast CDO Tranche Pricing using Free Loss Unit Approximations

Douglas Muirden

Credit Quantitative Research,

Royal Bank of Scotland,

135 Bishopsgate, London EC2M 3UR.

douglas.muirden@rbs.com

Original Version : May 2010. This Version : January 2011

Abstract

Volatility in Credit Markets has highlighted that effective risk management of a CDO tranche trading book requires analysis across a wide range of potential scenarios. The Basel III Comprehensive Risk Measure (CRM) requires repricing a trading book over many thousands of simulations, while counterparty risk calculations such as Potential Future Exposure (PFE) are generally analyzed by revaluation across a wide range of future scenarios. Fast CDO tranche pricing has therefore become increasingly important. In many cases accuracy can be sacrificed for speed, but previously documented approximations are too inaccurate or are not applicable to the random recovery valuation models currently in use by most banks. This paper presents an approximation which retains high accuracy in extreme cases, and can be used efficiently with random recovery models and inhomogeneous portfolios.

1 Introduction

Volatility in Credit Markets has highlighted that effective risk management of a CDO tranche trading book requires analysis across a wide range of potential scenarios. Traditional risk calculations for CDO tranches have concentrated on local sensitivities to market data, many of which could be calculated efficiently using semi-analytic methods. The non-linear nature of CDO tranche prices mean that local sensitivities have limited use when predicting the effect of large market moves. Effective risk management in volatile markets therefore also requires revaluation of the book under different market data scenarios with significant shifts to spreads and rates.

Counterparty risk has become an important area of focus, and calculations such as Potential Future Exposure (PFE) are generally best analyzed by revaluation of trades under a wide range of future scenarios. Furthermore, the Basel III capital adequacy requirements specify a Comprehensive Risk Measure (CRM) calculation based on the 99.9% tails of the distribution of potential future losses. Accurate calculation of this measure is essential, otherwise banks will be required to hold highly punitive levels of regulatory capital. To achieve the required degree of precision involves repricing a trading book across many thousands of simulations so the ability to recalculate CDO tranche prices quickly has therefore become increasingly important.

In many cases accuracy can be sacrificed for speed, but previously documented approximations become less reliable in extreme cases, and these cases often correspond to the most important scenarios for the analysis.This paper presents an approximation which retains high accuracy in extreme cases, and can be used efficiently with random recovery models and inhomogeneous portfolios. It can be used on its own without switching to other approximations hence avoiding discontinuities, and is accurate enough on real-world portfolios that it can be used without setting up ad-hoc fallback rules to switch to slower recursion pricing on problematic tranches. Inaccuracies on small portfolios are common to all these approximations however and these remain.

2 Normal and Poisson Approximations

One popular speed-up for factor-based models has been to apply approximations to the evaluation of tranche prices conditional on the market factor. For models using random recoveries, ElKaroui et al. (2008) propose a Normal approximation with a Stein correction. When dealing with the restricted case of deterministic and homogeneous recoveries the authors suggest combining this with a Stein-corrected Poisson approximation, but since the market now requires random recovery models for consistent pricing this must be discarded. This leaves the Normal approximation which can be very inaccurate for base tranches, tight spreads, and/or low correlation. The Stein correction term in practice is very small and not significant.

Conditional on the market factor the loss distribution can be written as $L = \sum L_i X_i$ where $X_i \sim Bernoulli(p_i)$ with p_i the conditional probability of default on name i and L_i the loss on default. The X_i are independent and both p_i and L_i are functions of the market factor. A normalized distribution is assumed such that $\sum L_i = 1$ and $L \in [0,1]$.

When the expected portfolio loss $\mathrm{E}[L]$ is not too low or high a Normal approximation is reasonable but otherwise the actual loss distribution accumulates at 0 and 1 respectively so a bounded distribution would be more appropriate in these cases. A Poisson distribution is a natural proxy for the true conditional loss distribution when expected losses are low. It can also be used when expected losses are near to 1 by expressing the distribution in terms of $1-L$ rather than L. Extreme low or high expected losses will always arise since we are integrating across the market factor which can take on any value.

For illustrative purposes the following figures show the behaviour of these approaches in approximating a standard homogeneous binomial. In practice the true conditional loss distribution will be a form of inhomogeneous binomial. Note that for qualitative comparisons only the discrete distributions have been normalized by grid size. The tranche prices themselves give the true quantitative comparison. For these illustrations the Stein corrections, which are in any case quite small, are not included.

Figure 1 compares a 100-name homogeneous loss distribution with its approximating Normal for different levels of expected portfolio loss. The approximation evidently deteriorates when the expected portfolio loss approaches a distribution boundary.

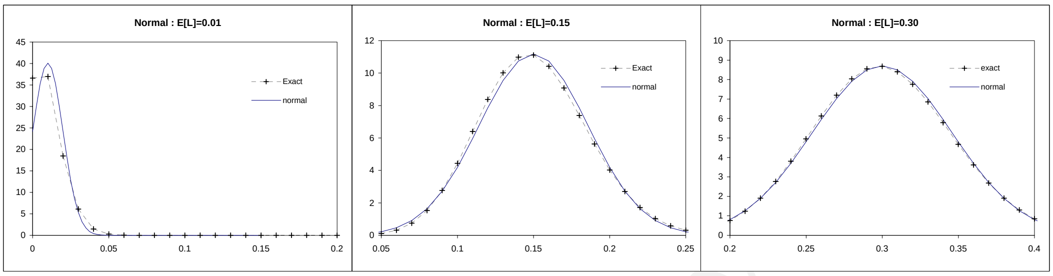

Figure 1: Normal approximation to the Binomial.

Figure 2 compares the same 100-name loss distribution with the standard Poisson approximation. As portfolio expected loss increases the accuracy deteriorates.

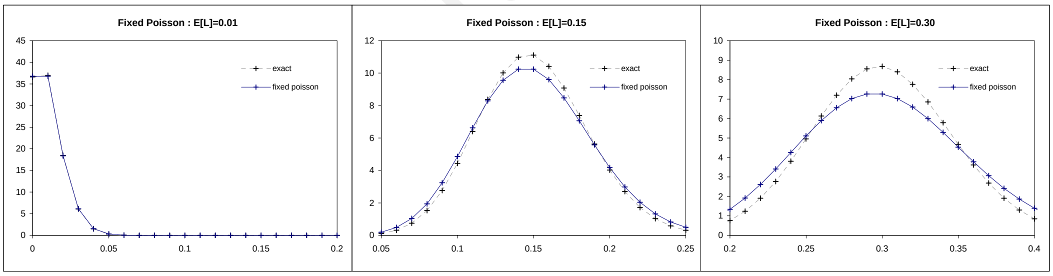

Figure 2: Standard Poisson approximation to the Binomial.

The ranges of accuracy of the Poisson and Normal are complementary so a threshold for expected loss can be specified at which the approximation changes from Poisson to Normal. For the example here this would typically be set around 0.10 to 0.15 although the two approximations are not necessarily very close in this changeover region. More seriously, if recoveries are inhomogeneous the distribution will be sparse over a grid with a small loss unit and the standard Poisson

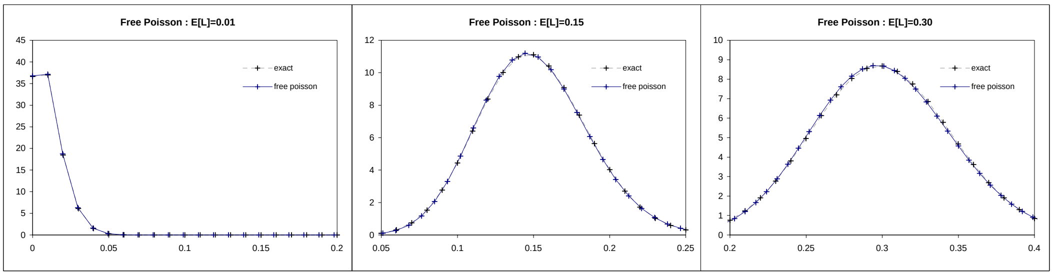

Figure 3: Free Poisson approximation to the Binomial.

approach becomes problematic, so for random recoveries or non-homogeneous L_i only the Normal approximation is used.

For a Poisson approximation we use $L \simeq \delta N$ where δ is a loss unit scaling the Poisson counter N. Implicit in standard treatments dealing with Poisson approximation is that the approximation loss unit δ is equal to that of the actual distribution, $\mathrm{GCD}[L_i]$. Allowing δ to be a free parameter however we get a more flexible distribution which can be used across the whole range of the distribution not just when $\mathrm{E}[L]$ is low, as seen in figure 3. Of course we are approximating a distribution with one discretization with another distribution having a different discretization, but there is nothing inherently wrong with this, it is simply a question of the accuracy of the end result. In fact at low expected losses, the free Poisson is very similar to the standard fixed Poisson and in this homogeneous example the grid size is very close to that of the true distribution. As expected loss increases the grid size decreases and the free Poisson smoothly changes over to be very close to Normal. Since we are integrating across the market factor, across different approximation loss units, this will also tend to smooth out any discretization bias.

3 Free Poisson Approximation

A normalized distribution is assumed such that $\sum L_i = 1$ and $L \in [0,1]$. We have $L \simeq \delta N$ where δ is a free loss unit parameter and $N \sim \mathit{Poisson}(\lambda)$. Let $\mu = \mathrm{E}[L]$ and $\sigma^2 = \mathrm{Var}[L]$. Since $\mathrm{E}[N] = \mathrm{Var}[N] = \lambda$, matching moments gives $\delta = \sigma^2/\mu$ and $\lambda = \mu^2/\sigma^2$.

The Poisson probability mass function $p(k) = \mathrm{Prob}[N = k] = e^{-\lambda}\lambda^k/k!$, the distribution function is $F(k) = \mathrm{Prob}[N \leq k] = \sum_{i=0}^{k} p(i)$, and we have the partial expectation $\sum_{i=0}^{k} i p(i) = \lambda(F(k) - p(k))$.

The expected loss of a base tranche detaching at K in this approximation is then given by

$$\mathrm{ETL}(K) = K + (\mu - K)F(k) - \mu p(k) \tag{1}$$

where $k = \lfloor K/\delta \rfloor$, the integer part of K/δ.

An important implementation detail is for the distribution function $F(k)$. The parameters for this can get large so rather than direct calculation by summing $p(k)$ the standard identity $F(k) = 1 - G(\lambda, k+1)$ is used where G is the incomplete gamma ratio function defined by $G(x,\alpha) = \gamma(x,\alpha)/\gamma(\infty,\alpha)$ where $\gamma(x,\alpha) = \int_0^x e^{-t} t^{\alpha-1} dt$. Calculation of the incomplete gamma ratio function is treated in DiDonato and Morris (1987).

To capture the boundary at 1, when $\mu > 0.5$ the tranche is valued as an option on $1 - L$ rather than L.

For Poisson approximations the Stein correction depends on the difference between the mean and variance for the true loss distribution. Since the free Poisson can also be viewed as scaling the true loss distribution in such a way that mean and variance are equal, it consequently has a zero Stein correction by construction.

4 Free Binomial Approximation

The same approach can be developed starting from a homogeneous Binomial approximating distribution. We use $L \simeq \delta B$ where $B \sim \mathit{Binomial}(n, p)$. Since the binomial is defined on the integers $0, \ldots, n$ it is natural to use $\delta = 1/n$ leaving us two free parameters, n and p. $\mathrm{E}[B] = np$ and $\mathrm{Var}[B] = np(1-p)$, so matching moments gives $p = \mu$ and $\delta = 1/n = \mu(1-\mu)/\sigma^2$, where $\mu = \mathrm{E}[L]$ and $\sigma^2 = \mathrm{Var}[L]$.

With the Binomial probability function $f(k) = \mathrm{Prob}[B = k] = \binom{n}{k} p^k (1-p)^{n-k}$, the distribution function is $F(k) = \mathrm{Prob}[B \leq k] = \sum_{i=0}^{k} f(i)$, and we have the partial expectation $\sum_{i=0}^{k} i f(i) = np(F(k) - \frac{n-k}{n} f(k))$

The expected loss of a base tranche detaching at K in this approximation is then given by

$$\mathrm{ETL}(K) = K + (\mu - K)F(k) - \mu \frac{n-k}{n} f(k) \tag{2}$$

where $k = \lfloor nK \rfloor$.

Again the distribution function $F(k)$ can be evaluated either by direct calculation by summing $p(k)$ or if necessary by $F(k) = 1 - I_p(k+1, n-k)$ where $I_x(a,b)$ is the incomplete beta ratio function defined by $I_x(a,b) = B_x(a,b)/B_1(a,b)$ where $B_x(a,b) = \int_0^x t^{a-1}(1-t)^{b-1} dt$. Calculation of the incomplete beta ratio function is treated in Press et al. (1992).

References

DiDonato, A. R. and Morris, A. H. (1987). Incomplete gamma function ratios and their inverse. *ACM TOMS*, 13: 318–319. http://www.netlib.org/toms/654.

ElKaroui, N., Jiao, Y., and Kurtz, D. (2008). Gauss and Poisson Approximation: Application to CDOs Tranche Pricing. *Journal of Computational Finance*.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C*. Cambridge University Press. www.library.cornell.edu/nr/cbookcpdf.html.

A Appendix : Pricing Differences

Prices for the following figures were calculated under a Random Recovery model using actual market data for standard detachments on the CDX series 9 portfolio, using the Stein-corrected Normal and the Free Poisson approximation as well as the standard recursion calculation for reference. Prices were calculated for each tranche over a grid of different flat correlations and spread scaling factors. The scaling factors were applied to all spreads in the market data, so for example for a scaling factor of 0.5 all spreads were halved before running price calculations. The values shown are relative asset leg pricing errors, $(P_{approx} - P_{ref})/P_{ref}$ where P_{approx} is the approximation pricing and P_{ref} is the standard recursion pricing of expected discounted tranche losses.

At low spread levels and low correlations the Normal-based approximation significantly overvalues the 0-100 portfolio expected loss so this affects all base tranches, see figures 4 and 5. It is exacerbated at short maturities see figure 6 (note the scales change between figures). Base tranche pricing is improved dramatically with the Free Poisson approximation. Inaccuracies remain as can be seen by looking at mezzanine tranches, figures 7 and 8. The Free Poisson is still a significant improvement over the Stein Normal however.

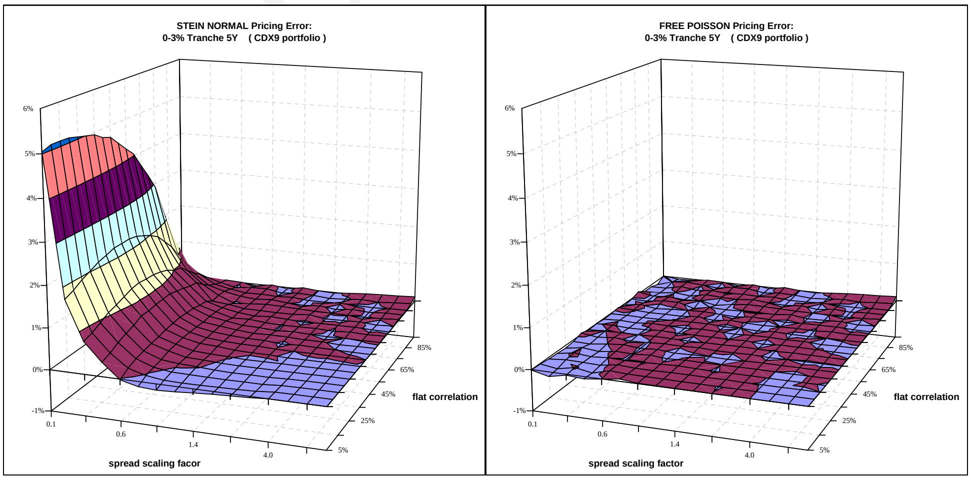

Figure 4: Relative pricing errors : 0-3% Tranche 5Y (CDX9 Portfolio).

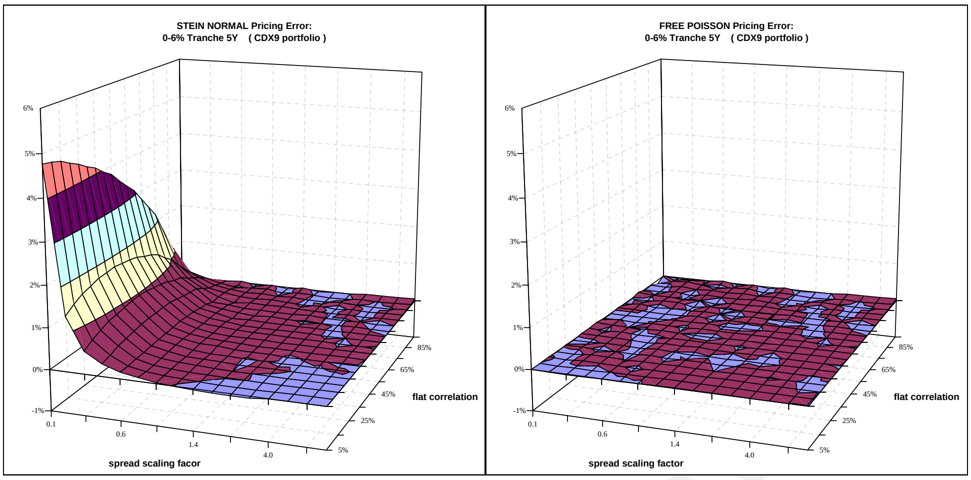

Figure 5: Relative pricing errors : 0-6% Tranche 5Y (CDX9 Portfolio).

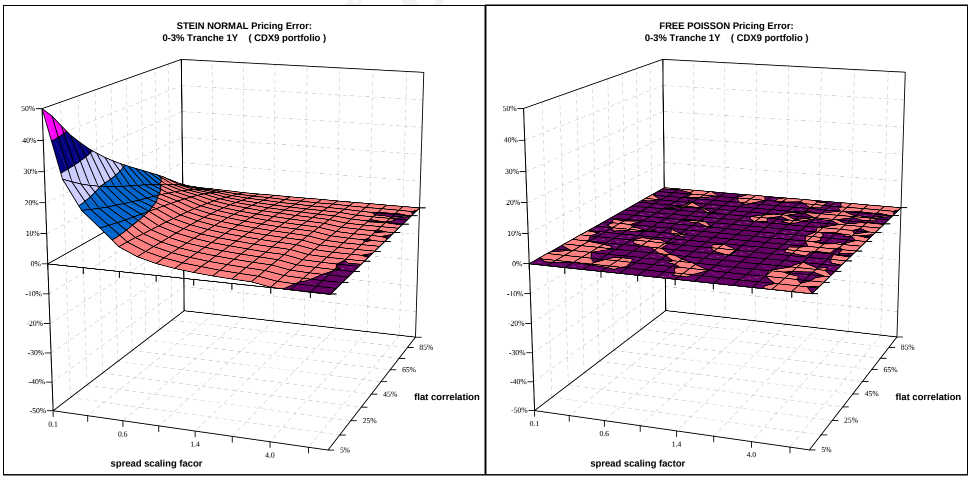

Figure 6: Relative pricing errors : 0-3% Tranche 1Y (CDX9 Portfolio).

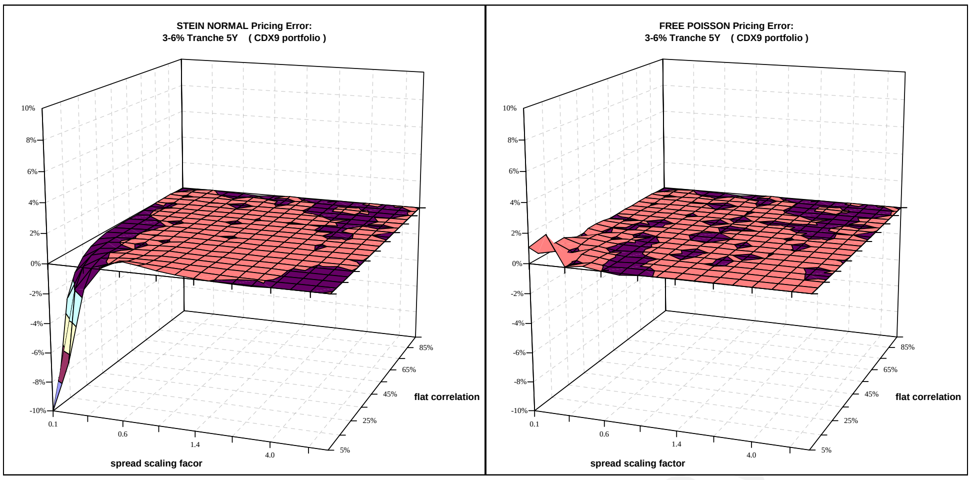

Figure 7: Relative pricing errors : 3-6% Tranche 5Y (CDX9 Portfolio).

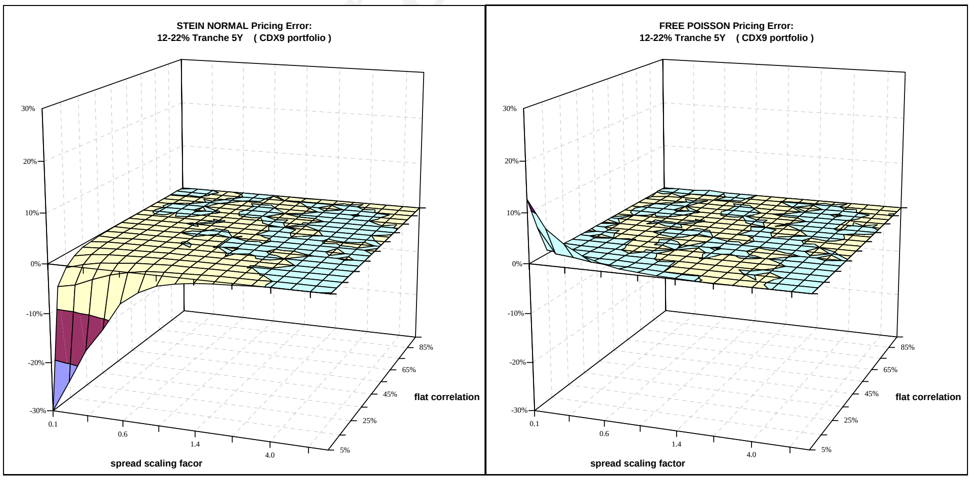

Figure 8: Relative pricing errors : 12-22% Tranche 5Y (CDX9 Portfolio).