

APPLYING IMPORTANCE SAMPLING TO PRICING SINGLE TRANCHES OF CDOs IN A ONE-FACTOR LI MODEL

MARK S. JOSHI

ABSTRACT. It is shown that importance sampling can be effectively applied to the pricing of a single tranche of a CDO. In particular, by shifting the mean of the common factor, it is demonstrated that the price can be estimated to an accuracy of approximately one percent with about ten thousand paths in a large range of cases. This is achieved at minimal extra computational complexity.

1. INTRODUCTION

Single-tranche synthetic CDOs with one hundred or more underlying names have recently become very popular. They present, however, significant challenges in the pricing and the computation of sensitivities. Whilst Monte Carlo is an easy method to implement, it tends to result in slow convergence, with the problem being exacerbated by short maturity times. Here we demonstrate how to apply importance sampling to pricing in the Li model for single factor models to achieve rapid convergence with stability across maturities.

We recall that in the the Li model, individual names are assumed to default according to a Poisson process with a deterministic intensity, known as the *hazard rate*. The key difficulty in basket credit derivative pricing is to decide how to correlate these default processes. In the Li model, this is done by using a Gaussian copula for the default times; the model is therefore also known as the Gaussian copula model. Roughly, this means that the default times are converted into Gaussians, correlated as Gaussians and then mapped back to default times.

In a single factor model, the correlation matrix of the Gaussians is typically taken to be constant off the diagonal; although more complicated forms are possible. This allows a particularly simple method of simulating the Gaussians in a Monte Carlo simulation. If we have N names and the correlation is ρ , we draw a single Gaussian, Z , to be the common factor and N independent Gaussians Z_j , $j = 1, \dots, N$. We

then set

$$(1.1) \quad W_j = \sqrt{\rho}Z + \sqrt{1 - \rho}Z_j,$$

to obtain the correlated variates. Note that we have order $\mathcal{O}(N)$ operations. Note also that conditional on a value of Z the variates W_j are independent. This observation is key in quasi-analytic approaches to pricing tranching credit derivatives, [4].

Let name j have hazard rate λ_j . The cumulative default probability of asset j is then

$$(1.2) \quad E_{\lambda_j}(\tau) = 1 - e^{-\int_0^\tau \lambda_j(s) ds}.$$

A simple algorithm to construct correlated defaults with the Gaussian copula is therefore to proceed as above to construct the variates W_j and then set

$$(1.3) \quad \tau_j = E_{\lambda_j}^{-1}(N(W_j)),$$

where N denotes the cumulative normal function.

Once one has the default times, one obtains the pay-off and discounts the cashflows appropriately. The price is then found by averaging across many draws.

The problem lies in the fact that many paths will result in zero pay-off except for fixed spread payments. This is particularly the case for tranches with seniority and for short times to expiry. For example, suppose hazard rates are one percent, there is one year to go and ten defaults are needed from a hundred name basket to obtain a non-trivial pay-off. If defaults are independent the probability of the requisite number of defaults occurring is approximately $7.29E - 8$. Thus only a tiny fraction of paths will be useful. Once correlation has been put in, the number of non-trivial paths increases but it will still be very small.

The problem is therefore a natural candidate for importance sampling, where the underlying probability distribution is reweighted to sample the areas where the greatest change occurs. Here, in particular, we want to ensure that enough defaults occur for most paths to give rise to a non-trivial pay-off. For a general discussion of importance sampling in the context of derivatives pricing, see [3].

We have previously applied importance sampling to the pricing of nth to default swaps for small baskets [6]. However, the approach to importance sampling presented there does not obviously generalize to small factor models for large baskets, and so here we present a new technique.

The key to our technique is to shift the mean of the common factor in a judicious fashion. We specify the details in Section 2. Once this has been done, convergence becomes stable and robust with respect to maturity. We present numerical results in Section 3.

The Gaussian copula approach to pricing portfolio credit derivatives was introduced by Li, [7]. The main motivation for the model is more that it provides a simple method of introducing default correlation rather than that it provides a financially plausible mechanism. The pricing of CDOs without Monte Carlo simulation in one-factor models using quasi-analytic methods was studied by Laurent and Gregory, [4]. Their methods rely on the portfolio being homogeneous in terms of recovery rates and on the pay-off being insensitive to the order that assets default in. Similar techniques were later studied by Hull and White, [5], under similar assumptions. In [1], various techniques including Monte Carlo without importance sampling are discussed.

The techniques we present here have the advantage over quasi-analytic methods that they do not require the recovery rate of each asset to be the same (or have the same distribution), and that they can be applied to products for which the pay-off depends upon which asset has defaulted as well as how many have defaulted. One example of such a product is a CDO-squared that is a CDO in which the underlying assets are also CDOs.

2. IMPORTANCE SAMPLING

Our approach is to shift the common factor, Z , in such a way that most paths become significant and are concentrated in the region where most variation occurs. We then have to compensate with a likelihood ratio factor to account for our measure change. Our approach is to pick the shift in such a way that the average number of defaults is in the middle of the range where the number of defaults matter.

For example, suppose we have one hundred names with zero recovery rate. The product we wish to price is a single mezzanine tranche which covers losses between 10 and 20. We then want the number of defaults per path to average 15.

We take a shift, μ , which we indicate how to compute lower down, and our algorithm changes as follows.

- Draw a variate X from a standard normal Gaussian.
- Set $Z = \mu + X$.

- Proceed as before.
- Multiply the discounted final pay-off for the path by $e^{0.5\mu^2 - \mu Z}$.

The final multiplication is by the standard likelihood ratio for shifting means, see eg [3]. Note that the additional computational effort per path is tiny. One simply has to add a constant, and to compute a single exponential which is then multiplied by. When working with a hundred name basket, it will be hard even to detect the extra CPU time.

The crucial issue is, of course, how to choose μ . Suppose the pay-off of our credit derivative varies for between k and l defaults and there are N names. We then target

$$m = \frac{k + l}{2}$$

defaults. Our target default probability, conditional on the draw of Z , is therefore

$$p = m/N.$$

We choose μ in a way that this holds true when $X = 0$ which is the median case. Given a hazard rate λ_j and a maturity time T , we can compute the natural default probability, $p_j = E_{\lambda_j}(T)$, as mentioned above. Asset j will default in the algorithm if and only if

$$\sqrt{\rho}Z + \sqrt{1 - \rho}Z_j \leq N^{-1}(p_j).$$

If we take $x = 0$, we obtain that a default occurs if and only if

$$Z_j \leq \frac{N^{-1}(p_j) - \sqrt{\rho}\mu}{\sqrt{1 - \rho}}.$$

We therefore simply solve for the unique μ_j such that the righthand side of this equation is equal to $N^{-1}(p)$ to make name j default with probability p .

However, we can only choose one μ for all j rather than choosing the desired value for each one. We therefore set

$$(2.1) \quad \mu = \frac{1}{N} \sum_{j=1}^N \mu_j.$$

The procedure we have used to arrive at μ is a little ad hoc. However, the precise value of μ is not particularly important: by picking a μ in the right area we get a varying number of defaults with an average in the range that causes the pay-off to vary; changing the mean a little up or down does not destroy this property. In any case, we can expect

quadratic behaviour near a minimum so small perturbations will not have much effect.

Note that our entire approach relies on the fact that there is a non-trivial correlation between assets. In the case of zero correlation, our approach will not work as there is no common factor.

3. NUMERICAL RESULTS

In this section, we present a range of numerical results. We study tranches of a homogeneous basket for simplicity, although in this case, there are other quasi-analytic approaches which are preferable to Monte Carlo [4]. We take a basket of 100 names. We take each name to have zero recovery rate. In the graphs that follow we plot the standard deviation of our Monte Carlo simulation of the price of a tranche between two levels for varying expiries and correlations. As it is the error as a fraction of price that is important, we normalize the standard deviation by dividing by the price. Note we take the price of the protection leg only, in order to avoid cancellation effects. In Figure 3, we present the results of one case without importance sampling to contrast the other results. We only present one case as without importance sampling a very large number of paths is required to obtain results, and thus to run a simulation for many maturities is very time consuming. Note that whilst 5-year trades are most liquidly traded, the difficulties with pricing decrease with maturity, and we therefore focus on shorter term deals. In any case, a 5-year deal will decrease in length throughout its life so it is essential that a numerical method can cope with such short term deals. An important feature of our results is that the normalized standard deviation is stable even for short term deals. Across the range of figures 1 – 6, we see that the normalized standard deviation is typically around 2. Even in stress cases, the standard deviation rises to only 16, in these cases the value is miniscule in any case. With a normalized standard deviation of 1, we will require ten thousand paths to obtain the price with a standard error of one percent of the price. This can be done in approximately one second. With a standard deviation of 2 this means we need about four seconds to get the price within one percent. Note that in [1], the typical number of paths used was 4×10^5 . In figure 3, see that for a mezzanine tranche without importance sampling; the normalized standard deviation is substantially higher everywhere and blows up much more rapidly for short maturities.

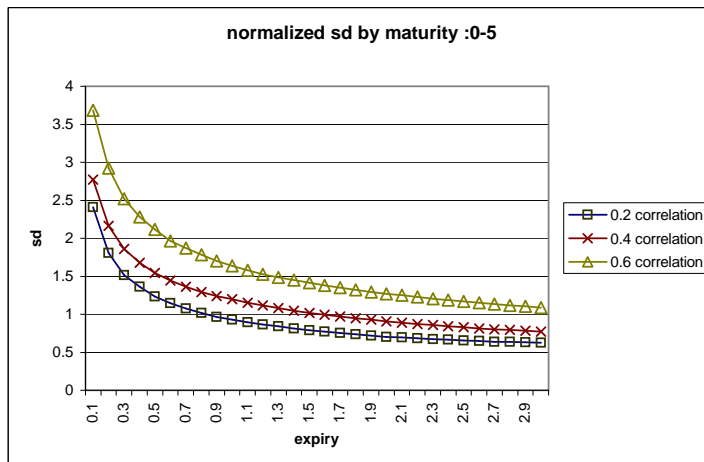


FIGURE 1. The normalized standard deviation of an equity tranche ranging from 0 to 5.

4. ELLIPTIC COPULAS

Whilst the Gaussian copula has recently become a popular model for quoting prices, it is not so clear that it is generally believed for pricing, and similar to the Black-Scholes model it has become a method of easily communicating prices, with implied correlation becoming the analogue of implied volatility. For example, in [1] the student-T copula is used. See also [8], [9]. The student-T copula is an *elliptic* copula.

A common feature of elliptic copulas is that they are conditionally Gaussian. A random variable is drawn from some distribution which specifies variance, and then the correlated Gaussians with this variance are constructed. The correlated Gaussians are then transformed to uniforms via the appropriate cumulative distribution function.

The correlated Gaussians can be constructed in the same manner as for the Gaussian copula. We can therefore apply the techniques of this paper by shifting the mean of the common factor. This could be done either by fixing a constant shift that bring the default times into

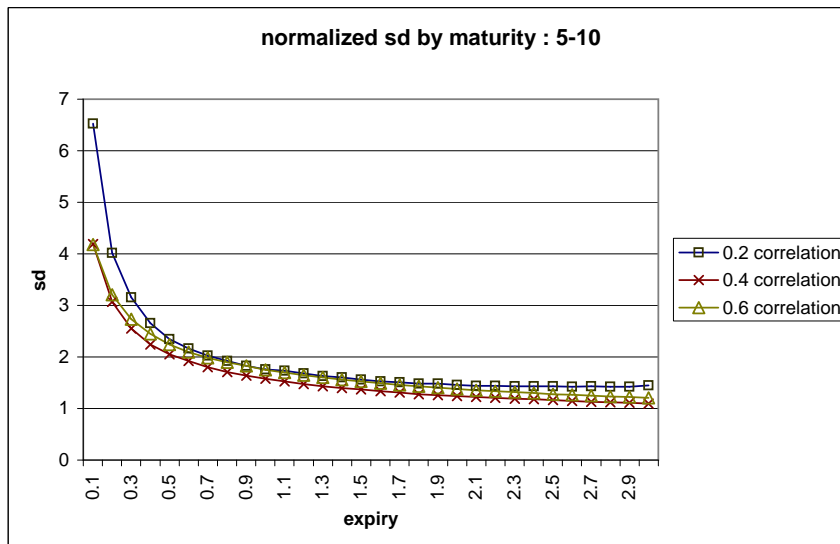


FIGURE 2. The normalized standard deviation of a mezzanine tranche ranging from 5 to 10.

a reasonable range for the median draw of the variance, or by doing a mean shift conditional on the variance. The second will be more effective but also more time consuming. We leave it to other authors to explore the optimal strategy.

5. GREEKS

As well as computing the price, it is also important to know the option price sensitivities both for risk management and for hedging purposes. Of course, any technique that results in faster pricing will help with computing the Greeks, particularly if they are done by bumping model parameters. One would, however, prefer something more sophisticated.

One approach is to use the likelihood ratio method of Broadie and Glasserman, [2]. This was applied effectively in [6] to the pricing of nth to default swaps for small baskets. With this method, a sensitivity is computed by multiplying the pay-off by a term dependent on the Greek

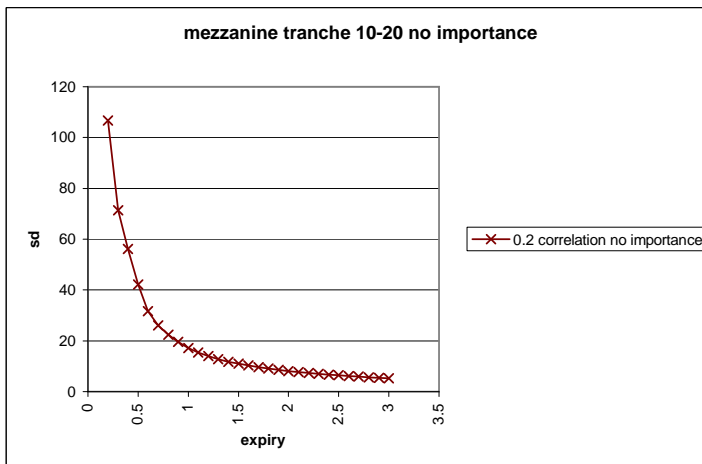


FIGURE 3. The normalized standard deviation of a mezzanine tranche ranging from 10 to 20 without importance sampling.

computed. The great advantages of likelihood ratio are that only one simulation is required to compute all the Greeks and that the method is pay-off independent, so that once all the method has been coded for one pay-off it immediately works for any pay-off.

Unfortunately, with the Gaussian copula the likelihood ratio terms tend to be rather rapidly varying which results in slow convergence. In the case we examine here, there are two obvious ways to proceed. We can either compute the joint density of the defaults unconditionally, or we pick the common factor, Z , and then work with the conditional density. The computations for the unconditional case are given in [6], and we therefore concentrate on the conditional case here. In any case, we found that the two techniques gave similar rates of convergence.

In this case, we can write the expectation to be evaluated in the form

$$(5.1) \quad \int \phi(z) \left(\int \psi_z(\tau_1, \dots, \tau_n) F(\tau_1, \dots, \tau_n) d\tau_1 \dots d\tau_n \right) dz,$$




FIGURE 4. The normalized standard deviation of a mezzanine tranche ranging from 10 to 20.

where $\phi(z)$ is the standard normal density, F is the discounted pay-off and ψ_z is density of the default times conditional on z .

As conditional on a given value z , the default times are independent, the density, ψ_z , factors into a product of terms $\psi_z^{(i)}(\tau_i)$ which are the densities of the individual defaults.

The dependence of the hazard rate, λ_i , will enter only into (5.1) via the dependence of $\psi_z^{(i)}$ on it. We therefore obtain the following expression for the sensitivity with respect to λ_i :

$$(5.2) \quad \int \phi(z) \left(\int \psi_z(\tau_1, \dots, \tau_n) \frac{\partial \log \psi_z^{(i)}}{\partial \lambda_i} F(\tau_1, \dots, \tau_n) d\tau_1 \dots d\tau_n \right) dz,$$

It remains to compute

$$\frac{\partial \log \psi_z^{(i)}}{\partial \lambda_i}.$$

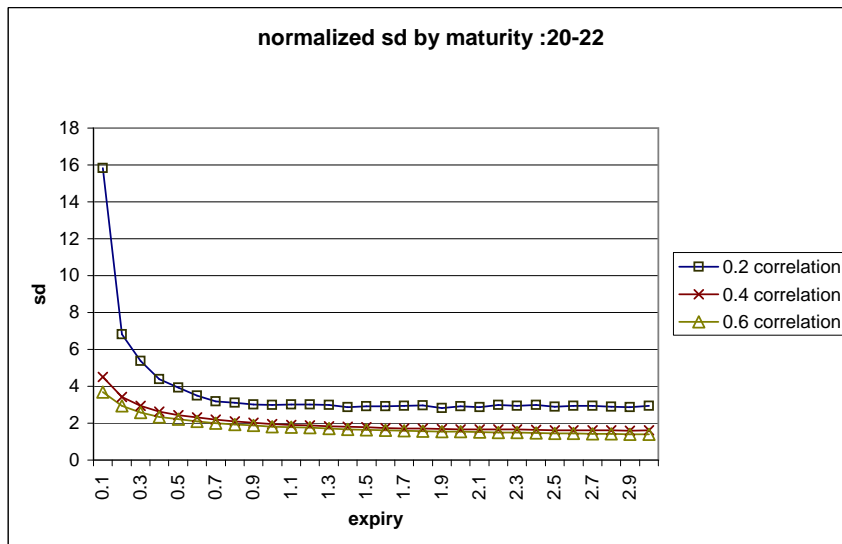


FIGURE 5. The normalized standard deviation of a mezzanine tranche ranging from 20 to 22.

However, this is straightforward; the distribution function, $\Psi_z^{(i)}(T)$, is easily seen to be

$$N\left(\frac{N^{-1}(E_{\lambda_i}(T)) - \rho^{\frac{1}{2}}z}{\sqrt{1-\rho}}\right).$$

This is then easily differentiated with respect to T to obtain the density function, and then to get the likelihood ratio one simply takes log and differentiates again, this time with respect to λ_i . We omit the expressions for brevity.

Whilst it is easy to carry out these differentiations, it is important to realize that we have to take two derivatives of the inverse cumulative normal function which results in a rapidly varying term. Thus while the likelihood ratio method can be easily applied to this case, the variance of the simulation is high. The ratio of the standard deviation to Greek is typically around 50 even with importance sampling, and we therefore need roughly $25^2 = 625$ times as many paths to achieve the same level of accuracy as for the price. However, as one computes all the

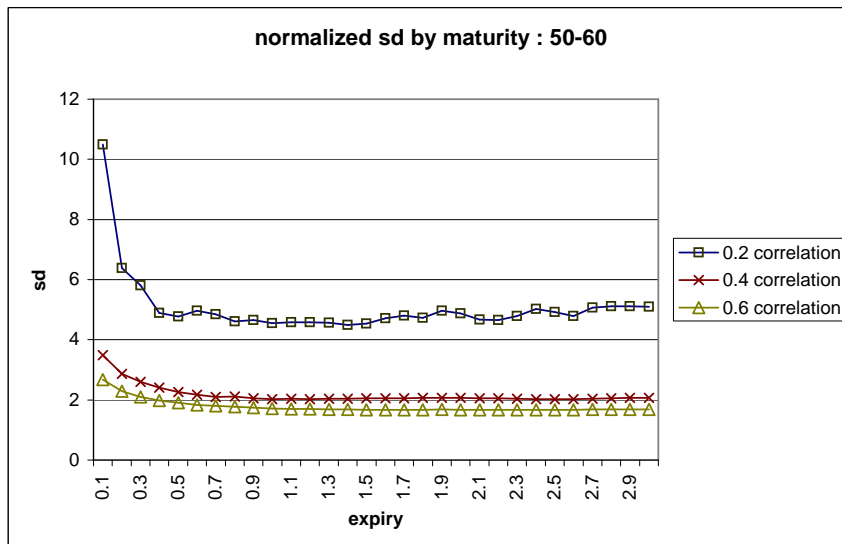


FIGURE 6. The normalized standard deviation of a senior tranche ranging from 50 to 60.

likelihood ratios simultaneously this is not necessarily prohibitive. Yet it is certainly a lot worse than one would desire, and further work needs to be done in this area.

Another possibility which we leave for future work is to use the path-wise method, [2], extended to discontinuous pay-offs, as was done in [6] for nth to default baskets.

6. CONCLUSION

We have introduced a new importance sampling method for the pricing of credit derivatives on large basket of names. We have shown that it is highly effective for obtaining the price of a tranced CDO. We have also examined the computations of Greeks and seen that the likelihood ratio method can be applied in this case.

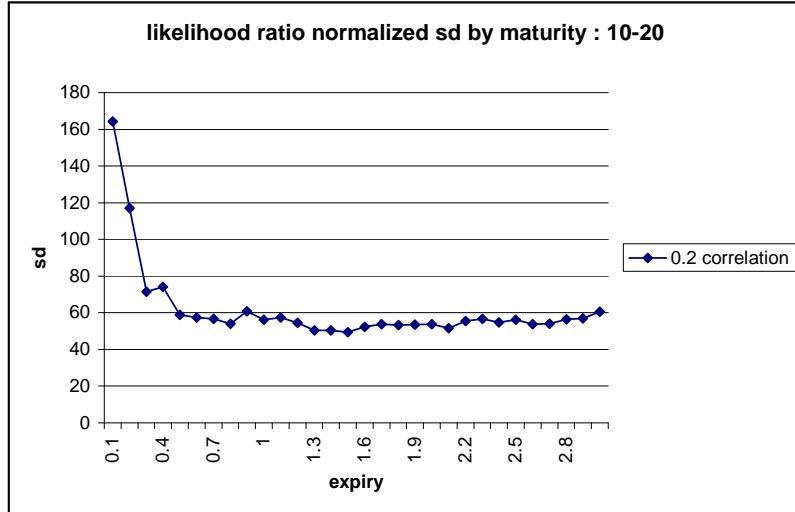


FIGURE 7. The normalized standard deviation of the sensitivity of a mezzanine tranche with respect to a hazard rate ranging from 10 to 20.

7. ACKNOWLEDGEMENTS

I am grateful to Dherminder Kainth for helpful conversations.

REFERENCES

- [1] L. Andersen, J. Sidenius, S. Basu, All your hedges in one basket, Risk Magazine November 2003, 67–70
- [2] M. Broadie, P. Glasserman, Estimating security derivative prices by simulation, *Management science*, 42, 269-285, (1996).
- [3] P. Glasserman, Monte Carlo methods in financial engineering, Springer Verlag 2003.
- [4] J. Gregory, A. Laurent, I will survive, Risk Magazine June 2003, 103–107
- [5] J. Hull, A. White, Valuation of a CDO and an nth to default CDS without Monte Carlo simulation, preprint October 2003
- [6] M. Joshi, D.S. Kainth, Rapid computation of prices and Greeks for nth to default swaps in the Li model, *Quantitative Finance*, volume 4, issue 3, (June 04), pages 266 - 275

- [7] D. X. Li, On Default Correlation: A Copula Function Approach, RiskMetrics Working Paper
- [8] R. Mashal and M. Naldi, Pricing Multiname Credit Derivatives: Heavy Tailed Hybrid Approach, Preprint, <http://www.columbia.edu/~rm586>, 2002
- [9] R. Mashal, A. Zeevi, Beyond correlation: extreme co-movements between financial assets, Working paper, Columbia Business School

QUARC, GROUP RISK MANAGEMENT, ROYAL BANK OF SCOTLAND, 280
BISHOPSGATE, LONDON EC2M 4RB