INTENSITY GAMMA: A NEW APPROACH TO PRICING PORTFOLIO CREDIT DERIVATIVES

MARK S. JOSHI AND ALAN M. STACEY

ABSTRACT. We develop a completely new model for correlation of credit defaults based on a financially intuitive concept of business time similar to that in the Variance Gamma model for stock price evolution. Solving a simple equation calibrates each name to its credit spread curve and we show that the overall model can be calibrated to the market base correlation curve of a tranching CDO index. Once this calibration is performed, obtaining consistent arbitrage-free prices for non-standard tranches, products based on different underlying names and even more exotic products such as CDO² is straightforward and rapid.

1. INTRODUCTION

During the last few years, there has been a great growth in the market of portfolio credit derivatives. Such derivatives pay off according to the defaults of a number of reference assets. A feature of such derivatives is therefore that the co-dependence of the assets' default times is a strong driver of their price. In this paper, we introduce a new mechanism for achieving default dependence which has attractive features including ease of calibration, time-homogeneity and the ability to match market prices.

The scope of our model is to price derivatives where the payoff is determined by the default times of a number of reference entities. We do not consider products which contain optionality nor whose payoff depends upon the credit spreads of the reference assets.

Suppose we have m reference assets $A_1, A_2, \dots, A_m$. We let τ_j denote the default time of asset A_j . (If A_j does not default before some fixed time horizon then, in effect, we take $\tau_j = \infty$). At the time of the default, there will be a loss rate, L_j , on asset A_j which will determine a recovery rate, $R_j = 1 - L_j$. A portfolio credit derivative, D , will then pay a sequence of cash-flows, (C_i) , the size and timing of each flow being a function of the recovery rates and default times. Whilst our work is extensible to the case of stochastic recovery rates, we restrict our attention here to that of deterministic recovery rates as is common practice in pricing exotic credit derivatives.

The most important examples of portfolio credit derivatives are CDO tranches where the payoff starts when the total loss on the portfolio reaches a certain level (or *attachment point*), continues from that point in proportion to the further losses that are incurred and

then ceases when the losses reach a certain higher level (the *detachment point*). Other important examples include n th-to-default products which pay off at the time of the n th default and in proportion to the loss incurred by that default (if the default occurs before the expiry of the product). These basic products are discussed very widely in the credit derivatives literature; see, for example, [13] for further details.

Returning to our general portfolio credit derivative, if we denote today's price of a riskless zero-coupon bond with maturity T by $P(T)$, and assume that interest rates are deterministic, then the price of the derivative will be

$$\mathbb{E} \left(\sum_i P(t_i(\tau_1, \dots, \tau_n)) C_i(\tau_1, \dots, \tau_n) \right)$$

where t_i is the timing of the i th cash flow and C_i its size. The expectation will be taken in some pricing (or risk-neutral) measure which specifies the joint distribution $\phi(\tau_1, \dots, \tau_n)$ of the default times. The purpose of a model is to specify what this measure is.

The pricing measure will be constructed to recover the prices of credit default swaps on single names. This means that the distribution of the default time for each individual name is determined. The purpose of a portfolio credit model is therefore purely to determine the codependence of defaults. For this reason, copula models have become popular. These models work by mapping each default time monotonically to some reference random variable with a particular distribution such as Gaussian [8] or student-t [1] in such a way as to match the individual default probabilities, determining the joint distribution of these reference random variables in some standard way (e.g. multivariate Gaussian) and mapping back again to obtain the joint distribution of the default times. See also [13] for a more detailed explanation. The Gaussian copula model has become sufficiently popular that the pairwise constant correlation required to match market prices is now known as implied correlation in the same way that implied volatility is quoted in other markets. The fact that different products with the same underlyings but different attachment and detachment points, or different maturities, trade with vastly different correlations shows that the market does not wholly believe the Gaussian copula model. This phenomenon is known as the “correlation smile” or “skew”.

In particular, there are now well-known indices of liquidly traded CDOs, the I-Traxx and Dow Jones CDX, which quote prices for standardized tranches of various baskets of reference names. It has recently become popular to work with “base correlation”, that is the implied correlations of equity tranches¹, and if one plots the implied correlation as a function of detachment point, one typically obtains an upwards sloping concave curve; see Section 6 for figures. For a discussion of base correlation including the advantages of using this approach, see [12].

¹an *equity tranche* is one with an attachment point at 0

As in other markets, there is a desire to match all the liquidly traded instruments with a model that requires only a single set of parameters, in order to price more complex instruments consistently with their natural hedges. In other words, we wish to match the correlation smile. This has given rise to the need for more sophisticated models. If one works entirely with the Gaussian copula model, using different correlations to price different tranches, it is not easy even to interpolate and extrapolate prices for other tranches of the same basket with non-standard attachment points in an arbitrage-free way, let alone obvious how to price different kinds of portfolio credit derivatives such as CDO². Additionally, as the Gaussian copula model arises from a rather arbitrary procedure rather than through financial reasoning it can give rise to rather odd dynamics for credits spreads; it also has the undesirable feature of time-inhomogeneity.

In order to create a new model with more natural dynamics and with ability to match correlation smiles, we consider what causes default correlation. At times of stress in the economy, through contagion and generally poor trading conditions, more companies default. We can therefore think in terms of good and bad years. In a good year, the chance of default for each company is small, and in bad ones, it is high. This means that defaults cluster in bad years, and this gives rise to default correlation. We implement this concept via the concept of stochastic *business time* or *information arrival* which measures shocks to the market. In a bad year there are a lot of shocks, that is a lot of business time elapses, whereas in a good year very little business time passes.

We are partly inspired by the work of Madan and collaborators on the Variance Gamma model, [9, 10, 11], where a gamma process is used to model business time and the logarithm of a stock price is a Brownian motion when viewed as a function of business time. In the case of a single defaulting asset, Schoutens [3] has recently looked at using a firm value model based on the Variance Gamma process.

In our model the default event for each asset will have a certain intensity which represents the probability of default for each unit of business time (not calendar time) which passes. Conditional on the information arrival or business time process, (I_t) , defaults of different names are independent. There is a similarity here to a one-factor Gaussian copula model where conditional on a single Gaussian random variable, default times are independent. We give a more detailed description of our model in Section 3.

Thus the choice of the stochastic time process determines the assets' co-dependence. What properties should stochastic time have? Clearly, it must be increasing and preferably strictly so. For tractability, we will want a Markovian process with independent increments, and in order to achieve a time-homogeneous model we will also want the increments to be stationary. The most well-known and widely used process with these properties is the gamma process. It turns out that a single gamma process does not give enough flexibility to match market smiles but it is hardly any more work to use the sum of a (small) finite number of independent gamma processes, which will also have these properties. We include

in our sum the degenerate case of a gamma process with zero variance that corresponds to time moving at a constant rate. Note that our choice of using a linear of gamma processes is largely motivated by tractability and familiarity and one could equally well work with any of a wide class of *subordinators*; see Section 2.

We refer to the model just outlined, in which business time is the sum of a finite number of gamma processes, as the *Intensity Gamma* model. We show in this paper how to construct an Intensity Gamma model, how to calibrate it to recover single name default probabilities, how to use it for the effective pricing of portfolio credit derivatives including complicated products such as CDO² and we also show that it can match market smiles.

There are, of course, other approaches to obtaining default correlation. One can introduce auxiliary default processes which cause multiple defaults as in [5], thus causing default correlation. Another appealing approach is to introduce stochastic credit spreads, thus causing default correlation to arise from periods where all spreads are high simultaneously. The main problem with such models is obtaining enough correlation whilst retaining a tractable model.

2. THE GAMMA PROCESS AND SUBORDINATORS

An $\mathbb{R}_+$ -valued random variable, X , is said to have a *gamma distribution* with parameters $\gamma > 0$ and $\lambda > 0$ if it has density function f_X given by

$$f_X(x) = \frac{\lambda^\gamma}{\Gamma(\gamma)} x^{\gamma-1} e^{-\lambda x}, \quad (x > 0).$$

We write $X \sim \Gamma(\gamma, \lambda)$. The mean of X is then γ/λ and its variance is γ/λ^2 . When γ is an integer then X has the same distribution as the sum of γ independent exponentially distributed random variables each with parameter λ and it is often helpful to think of gamma random variables as simply generalizing this to non-integer γ .

A key property of gamma random variables is that if X and Y are independent with distributions $\Gamma(\gamma_0, \lambda)$ and $\Gamma(\gamma_1, \lambda)$ then $X + Y$ has distribution $\Gamma(\gamma_0 + \gamma_1, \lambda)$. Inspired by this observation it is possible to construct a cadlag² stochastic process $(X_t)_{t \geq 0}$ with

- (1) $X_0 = 0$;
- (2) independent increments, i.e., for any $0 \leq t_0 < t_1 < \dots < t_n$ the random variables $X_{t_1} - X_{t_0}, \dots, X_{t_n} - X_{t_{n-1}}$ are independent; and
- (3) there exist γ, λ such that $X_{t+s} - X_s \sim \Gamma(\gamma t, \lambda)$ for any $s, t \geq 0$.

Such a process is called a *gamma process* with parameters γ and λ .

Although it is not obvious from the above description, a gamma process is a pure jump process. To state this precisely, let us say that a jump occurs at time t if $X_t - \lim_{s \rightarrow t^-} X_s > 0$,

²*Continué à droite, limité à gauche*, a standard technical condition. See, for example, [15]

in which case the magnitude of the jump is simply this difference. Then an increasing cadlag process is said to be a *pure jump process* if $X_t - X_0$ is equal to the sum of the magnitudes of all the jumps occurring in the interval $(0, t]$. A gamma process path has infinitely many jumps in a finite interval, but for any $\epsilon > 0$ there are only finitely many jumps of size more than ϵ . A full understanding of this property and its relation to the definition of a gamma process given above requires a certain amount of sophistication in which the characteristic functions of random variables are an important tool; see [4].

More generally, a cadlag process satisfying (1) and (2) and with (3) replaced by the more general condition that the distribution of $X_{t+s} - X_s$ is independent of s is called a *Lévy process*. An increasing Lévy process is known as a *subordinator*. See, again, [4] for a great deal of further information about these topics.

Subordinators are natural models for business time or information arrival, as described in the Introduction. We will model information arrival by a process I_t that is the sum of n gamma processes, the i th of which has parameters (γ_i, λ_i) , plus a constant drift a ³; we call such a process a *multigamma process*. In fact we will find that taking $n = 2$ gamma processes gives us enough flexibility to match market data precisely, although we will also illustrate what happens when just one gamma process is used.

3. THE INTENSITY GAMMA DEFAULT MODEL

The defaults are driven by the information arrival process I_t . Each name i defaults at a rate $c_i(t)$ *per unit of information arrival*. Therefore the probability that a name survives to time T , conditional on (I_t) , is

$$\exp \left(- \int_0^T c_i(t) dI_t \right).$$

Note that the same information process drives all the defaults, although one could develop a model with different information processes only affecting names in certain sectors or regions. The default rate $c_i(t)$ is a function of calendar time t , not “business time” I_t ; there are arguments in favour of both, but the former makes calibration of each name much easier.

We in fact take $c_i(t)$ to be piecewise constant. So if it takes the value c between times t_1 and t_2 then the probability of survival up to t_2 (given that it has survived up to time t_1) and conditional on the path (I_t) is just

$$\exp \left(- c(I_{t_2} - I_{t_1}) \right).$$

³which can be regarded as the limit of a gamma process as $\lambda \rightarrow \infty$ and $\gamma/\lambda \rightarrow a$

4. CALIBRATION TO VANILLA CREDIT DEFAULT SWAPS

Market quotes are readily available for the credit default swap spreads of individual names for a range of maturities such as 1, 3, 5 and 10 years. Making assumptions of deterministic interest rates and recovery rates, and also assuming that default intensities are piecewise constant, one can infer survival curves for each name, i.e., the probability that the name has defaulted by any time t . (Note, one really only needs interest rates to be independent of defaults.)

Now suppose that we have somehow chosen the parameters of the information process, i.e., the parameters $(\gamma_j, \lambda_j)_{j=1}^n$ for each gamma process and the drift a ; we will indicate in Section 6 how these parameters are determined. The default functions c_i for each name i are chosen so as to match the market implied survival probabilities. We encapsulate the survival curves as a sequence of probabilities corresponding to particular fixed times between any two of which the default intensity is constant.

The calibration of each individual name is straightforward. If $c_i \equiv c$ between t_1 and t_2 then the conditional probability of survival to time t_2 given that the process has survived to time t_1 is

$$(4.1) \quad \int e^{-cx} f_I(x) dx$$

where f_I is the density function of $I_{t_2} - I_{t_1}$. The Laplace transform of a gamma distribution is simple and well-known so (4.1) is equal to

$$(4.2) \quad e^{-a\tau} \prod_{j=1}^n \frac{1}{(1 + c/\lambda_j)^{\gamma_j \tau}}.$$

where $\tau = t_2 - t_1$. One then sets (4.2) equal to the market-inferred probability and solves for c ; this is easily done numerically by taking logarithms and applying the Newton-Raphson method.

Note that we have given a method of calibrating to vanilla credit default swaps, given a business time process. When we are fitting correlation skews it will also be necessary to fit the business time process. This means that an inner calibration to credit default swaps will be carried out repeatedly: for each choice of business time process parameters that we try, it will be necessary to calibrate the single name curves. The fact that our single name calibration is very rapid is therefore essential.

5. PRICING PORTFOLIO CREDIT DERIVATIVES

When we refer to portfolio credit derivatives we mean products whose cashflows are determined by the default times and recovery rates of names in a particular portfolio. We do not therefore include products with optionality. However, products with complicated

payoffs such as CDO² are certainly within our scope. As indicated in the Introduction, we model recovery rates as deterministic.

Given the choice of parameters for the information arrival process, I_t , we price portfolio products by rapidly calibrating the c_i s to the CDS market as described in Section 4 and then running a Monte Carlo pricing engine. A path (I_t) out to the maturity of the product is randomly generated then an independent exponential random variable with mean 1 is generated for each name which determines the default time for that name: if V_i is the exponential random variable corresponding to name i , then name i will default at first time T_i for which

$$\int_0^{T_i} c_i(t) dI_t \geq V_i.$$

A naïve way to generate paths I_t would be to break the time interval of interest up into a number of small subintervals and somehow generate appropriate gamma random variables corresponding to the increments of each gamma process for each of these intervals. One would then treat the gamma process as having constant slope across each subinterval. This method is, however, rather slow and not particularly accurate.

Fortunately, there are more efficient and accurate ways to generate gamma paths. We use a method based on Cont and Tankov [4, Example 6.16]. We do not give details here but describe in general terms how the method operates. The magnitude and timing of largest jumps that occur in a gamma process over a fixed time interval can, it turns out, be rapidly and accurately simulated. These largest jumps are generated until a point is reached where the remaining jumps can be replaced by a (very small) addition to the drift term without any significant loss of precision. The resulting (approximation to a) random path is then stored as a finite number of jumps plus a positive drift.

Speeding up the pricing. The Monte Carlo pricing just described is reasonably rapid even without the use of variance reduction techniques. A typical tranche of a 125-name CDO with a 5-year maturity can be priced to three significant figures within about five seconds on a desktop PC with our C++ code. The intelligent use of quasi-random numbers reduces this to less than a second.

In the case of a certain very important class of products, a further speed increase can be achieved through the use of a control variate (see [6] for an explanation of this term). We say that a product is *permutation-invariant* if the cashflows do not change when the default times are permuted amongst the various names⁴. For such products it is possible to obtain a good analytic approximation to the price by assuming that c_i is constant for each name over the lifetime of the product (and chosen so as to match the market-implied survival probability at the maturity of the product). The survival probability at maturity time T , conditional on the path (I_t) , is then just $e^{-c_i I_T}$. Given the value of I_T one can

⁴Strictly speaking this is a property not only of the product but also the assumptions made about the deterministic recovery rates.

then rapidly and analytically obtain, for each k , the probability that there are precisely k defaults before maturity. The unconditional probability that there are precisely k defaults can then be obtained by integrating with respect to the distribution of I_T . One then makes the further simplifying assumption that all the defaults occur precisely in the middle of the lifetime of the product to determine which cashflows occur when there are k defaults and thereby price the product. This is exactly the same approach used to enable the analytic pricing of permutation-invariant products in the one-factor Gaussian copula model, see [7], although an additional integral is evaluated in that case to give an exact price by adjusting for the fact that the defaults do not actually all occur at the same time. In the case of the Intensity Gamma model this approach just gives an approximate price, but this can be effectively used as a control variate. It is possible that an improved control variate could be obtained with various more sophisticated choices of default timing taking into account the value of k and the average default intensity; however, we found that the unsophisticated method of using the middle was effective. For a typical 5-year tranching CDO this seems to result in a roughly 8-fold speed-up in pricing.

6. SMILE MATCHING AND TIME HOMOGENEITY

Credit indices, based on a basket of liquid names covering a range of maturities, regions and levels of credit quality, have evolved rapidly in recent years. Prices for standardized CDO tranches for the most important of these baskets are readily available.

If we wish to price a portfolio credit derivative based on some bespoke basket of names, using the methods described so far, it remains to choose the parameters $(\gamma_i, \lambda_i)_{i=1}^n$ and a for the information arrival process (I_t) . We do this by calibrating to market quoted CDO prices for an index with, as far as possible, similar properties to the bespoke basket in terms of diversity, region and credit quality. This is similar in spirit to the common practice, when using the Gaussian copula model, of somehow inferring base correlations for a bespoke basket from the correlations of a corresponding index. We do not address here in detail the problem, shared with other modelling techniques, of how to calibrate when there is no suitable index; see, however, the comments made in Section 7 about extensions of the model.

The parameters of the *multigamma* process, I_t , are chosen so as to match the market-quoted CDO tranche prices of a suitable index. Typically there are five quoted prices corresponding to the 3%, 6%, 9%, 12% and 22% attachment points, although for some indices different points are used. We work with the prices for five equity tranches and, as is currently standard in the market, refer to prices in terms of the corresponding base correlation as described in the Introduction.

Multiplying the multigamma process drift a by some factor m , and dividing each λ_i by the same m is equivalent to multiplying all the gamma paths by m . Such a constant

multiple washes out in the calibration of the individual names (all the c_i s get divided by exactly m) so there is effectively one redundant parameter in the specification of the multigamma process. We therefore set $a = 1$. This gives a number of free parameters equal to twice the number of gamma processes, $2n$.

Taking $n = 1$ enables one to match two prices exactly⁵. In practice one can actually obtain a curve which matches much of the base correlation skew tolerably well. See Figures 1 and 2. Note that in these cases it is only the very senior parts of the curve which are poorly matched by a single gamma process (plus constant drift) and this is particularly marked in the case of the CDX curve because the highest quoted detachment point is 30%. This actually fits in well with the general view that the super senior tranches (i.e. the complements of the large equity tranches) have a much higher market price than would be justified from any considerations of expected payoff. Furthermore, the actual prices of these pieces are much less correlation sensitive than for the lower tranches.

When we use two gamma processes, however, we can see on the same figures that we have enough flexibility to get a very good fit to the whole base correlation curve.

FIGURE 1. The base correlation smile for the 5-year US CDX together with the base correlations implied by the gamma process (plus a drift) which best matches the prices and also those implied by the best sum of two gamma processes.

For one gamma process the parameters can be adjusted by hand until a good match is obtained. For two processes the effect of varying the four free parameters is really too

⁵While it is reasonably clear that that one obtains a $2n$ -dimensional manifold of possible price vectors, it is by no means clear that any given set of $2n$ prices can be matched. In practice it seems that we can match $2n$ market prices, but we do not address this issue from a theoretical standpoint.

FIGURE 2. The base correlation smile for the 5-year ITraxx Europe together with the best fit from a gamma process and from two gamma processes.

complicated for this and so an optimizer is needed. We made use of the Downhill Simplex Method; see [14].

Although Figures 1 and 2 give just two examples of calibration, both prior to the credit market crisis of May 2005, we obtained similar results post-May and also for other indices. The one case where we did not obtain a near-perfect fit is the high-yield index after May 2005, where the base correlation curve was not monotonic and showed a pronounced smile shape. In that case, the best fit that we could obtain with two gamma processes had significantly less curvature than the market smile, although it was very comfortably within the bid-offer spread. Increasing the number of gamma processes did not seem to give a significant improvement.

Having derived the parameters of the multigamma process I_t so as to match the quoted index tranche prices, one can then use these parameters to derive prices for non-standard tranches and, indeed, for quite general portfolio credit derivatives on the same basket using the Monte Carlo methodology described in Section 5. It also seems reasonable to use the same multigamma process when pricing products with a different underlying basket which has similar properties to the index to which one has calibrated. Since these prices are all derived from a common pricing measure they are both consistent with the index prices and arbitrage-free.

In contrast to the preceding remarks, if one works just within the Gaussian copula model it is much less clear how one should obtain arbitrage-free prices even for non-standard CDO tranches of the index. For example, if one just linearly extrapolates the base correlation

curve one can end up with a price for, say, the protection leg of the 0-35% equity tranche which is less than that for the 0-34% tranche! The pricing of more complicated products within the Gaussian copula framework, such as the currently popular CDO², is even trickier and it seems that at the moment a number of rather involved techniques are under development in various houses for this purpose.

The codependence in the intensity gamma model is time-homogeneous, since the amount of information arrival in the first year will be independent of and identically distributed to the amount which arrives in, say, the sixth year. The time inhomogeneities only arise from the time-to-maturity dependence of credit spreads, i.e., from time-homogeneity in market inputs. Time-homogeneity means, in particular, that if one moves forward in time within the model, any changes in the shape of correlation smiles as a function of time to maturity will arise only from credit spread time inhomogeneities or defaults which occur. This is not the case for the Gaussian copula model because the joint distribution of exponential random variables with a Gaussian copula does not have the memorylessness property.

The independence of the business time process across disjoint times has the consequence that for a given set of Intensity Gamma parameters, there will be significantly less correlation observed in a 7-year product than in a 5-year product. The weak law of large numbers⁶ shows that as $T \rightarrow \infty$ then I_T becomes more concentrated around its mean; if it were entirely concentrated at its mean then the default-in- $(0, T]$ events for different names would be completely independent. In Figure 3 we show the implied base correlation curves for a particular bespoke basket with different maturities, using the parameters resulting from calibration to the 5-year US CDX curve.

The decrease in correlation with increasing maturity is, we believe, a desirable and realistic property. We remark, however, that this change in correlation level is more marked than actually seems to be observed in the market. Although we would like to suggest that the market prices are, perhaps, poorly aligned in this respect, it would be unwise to recommend the use of Intensity Gamma as it stands to price across substantially different maturities, i.e., to calibrate to an index for one maturity and then to use those parameters to price a product with a very different maturity. With just a little work, however, it should be possible to extend Intensity Gamma to make it more suitable for this purpose. We indicate how to do this in the next section to which we now turn.

7. OBSERVATIONS AND FUTURE DIRECTIONS

The model we have introduced is an improvement on the Gaussian copula in that it naturally introduces correlation skew allowing calibration to the market whilst retaining tractability. It perhaps, raises as many questions as it answers, however, and in this section

⁶one can also see this directly from the distribution

FIGURE 3. The base correlation smile for a bespoke basket with different maturities, using the same Intensity Gamma parameters

we discuss what features of the market it does not capture and what research topics are thrown up by the model.

The process for business time is Markovian and the amount of time that passes during an interval (s, t) is independent of what happened during $(0, s)$. A consequence of this is that information arrival during $(0, s)$ (other than the name's default) will not affect the probability that the name will default during the period (s, t) . The model therefore does not feature dynamic credit spreads, and would not be suitable for products involving optionality.

Our interest in the model has been sparked by the need to obtain matches to the correlation smile and thereby infer prices. Another viewpoint is that one wants a model that accurately reproduces market behaviour of the underlying names; it is therefore an interesting question whether the model is a good proxy for the real-world behaviour of credit markets, although we do not attempt to address this here.

An important part of any model is to compute Greeks. We expect many of the techniques used for computing Greeks by Monte Carlo, [2], to go over to this case, and our control variate will be effective there too. Calculating sensitivities to changes in vanilla spread is fairly straightforward, but robust measurement of the effect of changes in index correlation is much more challenging as it requires recalibration of the model⁷: it is therefore an important future direction to develop effective methodologies for computing sensitivities.

⁷This problem is far from unique to the Intensity Gamma model.

There are many ways in which the model could be extended. One realistic adjustment with a clear financial interpretation would be to introduce a stochastic delay between the arrival of information which causes a default and the actual occurrence of the default. A major shock, such as the events of September 2001, may well cause a number of companies to default some months, or even years, afterwards. In this adjusted model, one observes shocks immediately, but their effect is spread out over a period of time. Such an adjustment to the model would introduce some dependence across adjacent but disjoint time intervals and would thereby lessen the decrease in correlation with maturity described at the end of Section 6. Such a model could also avoid having simultaneous defaults and would have stochastic (although perfectly correlated) credit spreads. We cannot give more details here, but believe this worthy of investigation.

Another possible extension would be to have separate business time processes which only affect particular industrial sectors or geographic regions to achieve a more complicated correlation structure. One might wish to do this to obtain properties of the model which one thought were important to the pricing of a particular product that were not otherwise captured; it is hard to see that it is necessary for the pricing of ordinary CDOs if one can capture the market smile with the simpler model. Another motivation for use of multiple business time processes occurs when it would be natural to calibrate different parts of a portfolio to different indices. One could have one process calibrated to each relevant index and correlate the processes in some suitable way; we have had to do this in the case of bespoke portfolios containing a mixture of high-yield and investment grade names from different regions, with good results. And just as with the Gaussian copula model where the original model has been greatly refined to achieve greater speed there is much room to carry out similar improvements to Intensity Gamma.

The model gives an alternative method of pricing both bespoke tranches and more complicated products. It is an important question, which we leave to future work, how different these prices are from those obtained with the range of methodologies (many still under development) based on the Gaussian copula model and base correlation.

8. CONCLUSION

We have proposed a new model for pricing portfolio credit derivatives. It has a natural financial interpretation and can match the correlation smile without introducing unnatural inhomogeneities whilst retaining sufficient tractability to be usable for pricing. Perhaps the greatest practical strength of the model in its current form is that, once calibrated to index tranche prices, it can then be used to rapidly price any portfolio credit derivative without optionality on the same or a broadly similar basket of names including ones with more exotic payoff structures such as CDO².

REFERENCES

- [1] L. Anderson, J. Sidenius, and S. Basu, All your hedges in one basket, *Risk* **16**(11) 67–72 (2003).
- [2] M. Broadie, P. Glasserman, Estimating security derivative prices by simulation, *Management science*, **42**, 269–285, (1996).
- [3] J. Cariboni and W. Schoutens, Pricing Credit Default Swaps under Lévy Models, preprint (2004).
- [4] R. Cont and P. Tankov, *Financial Modelling With Jump Processes*, Chapman and Hall, Boca Raton (2004).
- [5] D. Duffie and K. Singleton, Modelling Term Structure of Defaultable Bonds, *Review of Financial Studies*, **12** 687–720 (1999).
- [6] P. Glasserman, *Monte Carlo Methods in Financial Engineering*, Springer, New York (2004).
- [7] J. Gregory, J. Laurent, Analytic approaches to the pricing and risk management of basket credit derivatives and CDOs, preprint (2004).
- [8] D. X. Li On Default Correlation: A Copula Function Approach, RiskMetrics Working Paper
- [9] D. Madan, E. Seneta, The Variance Gamma model for share market returns, *J. Business* **63**, 511–524 (1990).
- [10] D. Madan, F. Milne, Option pricing with V.G. martingale components, *Mathematical Finance* **1**(4), 39–55 (1991)
- [11] D. Madan, P. Carr, E.C. Chang, The Variance Gamma process and option pricing, *European Finance Review* **2** (1), (1998)
- [12] L. McGinty, E. Beinstein, R. Ahluwalia and M. Watts, *Credit Correlation: A Guide*, JP Morgan Credit Derivatives Strategy (2004).
- [13] P.J. Schönbucher, *Credit derivatives pricing models*, Wiley, Chichester (2003).
- [14] W.H. Press, S.A. Teutolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C*, 2nd edition, Cambridge University Press, 1992
- [15] L.C.G. Rogers and D. Williams, *Diffusions, Markov Processes and Martingales*, 2nd Edition, Cambridge University Press, Two Vols, (2000).

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Joshi, Mark S.;Stacey, Alan M.

Title:

Intensity gamma: a new approach to pricing portfolio credit derivatives

Date:

2006-03

Citation:

Joshi, Mark S. and Stacey, Alan M. (2006) Intensity gamma: a new approach to pricing portfolio credit derivatives.

Persistent Link:

<http://hdl.handle.net/11343/34305>