Dynamic modelling of single-name credits and CDO tranches

Martin Baxter[1]

Nomura Fixed Income Quant Group

20 March 2006

1. Introduction

Credit default events can happen in a variety of ways. Sometimes a credit slides towards default over a period of time, while other credits default instantly and without warning. Sometimes a credit's problems are unique to itself, but at other times many credits are influenced by common trends or events.

No matter how the default event happens, it is important to remember that such credit events are relatively rare. The most likely outcome for (almost) any entity is survival. Practically, this means that the study of credit is the study of the tails of random distributions.

These two separate observations make it necessary to use models that have both varied dynamics and the correct tail behaviour. This paper will present such a model. The model fits both CDS and CDO prices, is strongly intuitive, and is tractable to implement. Also it provides a credible dynamic of CDS spread evolution, opening up the possibility of pricing more exotic tranche-linked trades.

There are two basic ideas behind the new model. Firstly, the tails of the normal distribution are too light to match the reality of market tails. In the equity market this is a recognised truth and low-strike equity premia are much higher than log-normal models would indicate. It can also be seen in the problems that simple Brownian structural models have when they underestimate short-term default intensities. In common with other recent research, the solution for this problem is to add jumps to the processes. This makes the tails of the distribution heavier and more realistic.

The second idea, although itself simple, is based on the desire to get the correct complex dynamics. Existing continuous models have no jump terms, and existing jump models have no continuous terms. The new model has both. It has both a continuous Brownian-motion term and a discontinuous Variance-Gamma jump term. And each of these terms is further divided into a global (systemic) factor and an idiosyncratic (name-specific) factor. So in total there are four terms: global continuous, idiosyncratic continuous, global jump, and idiosyncratic jump.

We can give an intuition for each of these factors along with relevant examples from the recent past.

	Global factor	**Idiosyncratic factor**
Continuous Brownian component	General trends, often driven by the credit cycle or equity index markets	Entity-specific trend, such as Argentina, Delphi
Discontinuous Variance-Gamma jump component	Sudden news at the global level, such as the 11-Sep-2001 attack, General Motors downgrade.	Firm-specific news, such as Parmalat, Railtrack

[1] Nomura International plc, 1 St Martin’s-le-Grand, London EC1A 4NP.
Email: work@martinbaxter.co.uk

The model is structural, in that it assumes that credits are driven by a hidden process that is a proxy for the value of the firm. This process is modelled by a Brownian-Variance-Gamma (BVG) process as described above. There will now also be two "correlations" – one for the continuous movements, and another for the jumps.

The model enjoys the property of being a genuine arbitrage-free dynamic model of spread evolution. It is consistent with CDS curves, has a small number of relatively stable parameters, and reprices CDO tranches.

2. Single-name credits

The model for an individual credit's evolution is similar to independently developed work by others such as Schoutens (2006) and Moosbrucker (2006) in that it is a structural model with jumps. Unlike other published models it also includes a continuous term.

The model takes the form of a new process, called *Brownian-Variance-Gamma* (BVG), which is the sum of a Brownian term and a Variance-Gamma process. The form we use is

$$X(t;\rho,\gamma,\lambda) = \sqrt{\rho}\,W_t - \Gamma_1(t;\gamma,\lambda) + \Gamma_2(t;\gamma,\lambda)$$

where:

- o $W(t)$ is a Brownian motion
- o parameter ρ is the square of the volatility. The choice of this notation will be clearer in the CDO section, when it will be used as a correlation
- o each $\Gamma_i(t;\gamma,\lambda)$ is an independent gamma process with jump parameters γ (jump intensity) and λ (inverse jump size)

The process $X(t)$, which is a proxy for the value of the firm, is a Lévy process. This means it is a Markov process, with stationary independent increments. It is more general than the Lévy processes usually used in finance as it contains both continuous and discontinuous terms. Its advantages include a small number of parameters and computational feasibility. For simplicity, we have restricted the up and down jumps to be the same size. This can be relaxed without technical difficulty, but the increase in the number of parameters does not improve the results significantly.

For details of Lévy processes, see the very good book by Applebaum (2004) and a useful brief introduction by Winkel (2004), as well as the paper by Madan et al. (1998) which introduced the Variance-Gamma process for modelling equity skew.

Adapting a remark of Madan, our BVG process can also be represented (proof: via generating functions) as a time change of Brownian motion:

$$X(t) = W\left(\rho t + \Gamma(t;\gamma,\tfrac{1}{2}\lambda^2)\right).$$

This means that its marginal distribution can be expressed in the useful form

$$X(t) = Z\sqrt{\rho t + 2\lambda^{-2}\Gamma(\gamma t)}$$

where Z is a normal random variable and $\Gamma(\gamma t)$ is a gamma random variable with unit scale parameter. We can then approximate the default time at which the process goes below a threshold level θ as

$$\{\tau < t\} \cong \{W_*(\rho t + \Gamma(t;\gamma,\tfrac{1}{2}\lambda^2)) < \theta\},$$

where $W_*(t) := \inf\{W(s) : 0 \le s \le t\}$. Another less-accurate approximation is useful for CDOs and more exotic options:

$$\{\tau < t\} \cong \{W(\rho t + \Gamma(t; \gamma, \tfrac{1}{2}\lambda^2)) < \theta\},$$

where θ needs to be re-calibrated to match survival probabilities. In either case, we can calculate default probabilities by conditioning on the gamma variable and using standard Brownian formulas.

Parameter fitting

We can assume, due to scaling, that the factor ρ is one. The threshold θ can easily be calibrated to fit the curve's, say, 5y point. There only remains the jump parameters γ and λ to fit. This can be done over a set of credits, which may correspond to a credit basket correlation trade. Lambda controls the slope of the CDS curve: a small lambda implies a flatter curve, and a large lambda causes a steep curve. Gamma controls the shape of the front end of the curve: a small gamma implies a low front-end, and a higher gamma causes a higher front-end.

Alternatively the best global fit can be found for a basket of curves. The basket used may be either liquid or bespoke. Some fitted values and residual spread errors out to 10y are (market data of 9 March 2006):

Basket	Gamma	Lambda	75% percentile worst case error	90% percentile worst case error
CDX 125 S5	46.4%	15.4%	8.4 bp	13.7 bp
iTraxx 125 S4	46.0%	15.0%	7.0 bp	12.8 bp

For 75% of credits in each basket, the spread errors are less than 10bp, and fewer than 10% of credits have errors more than 15bp. We do not need a perfect fit, and this level of accuracy is fully adequate. Indeed it is actually quite a good fit for a two-parameter model compared against 125 credit curves. The market and model CDS spread curves for a typical credit, Washington Mutual (which has the median error in the CDX basket), are shown in the chart:

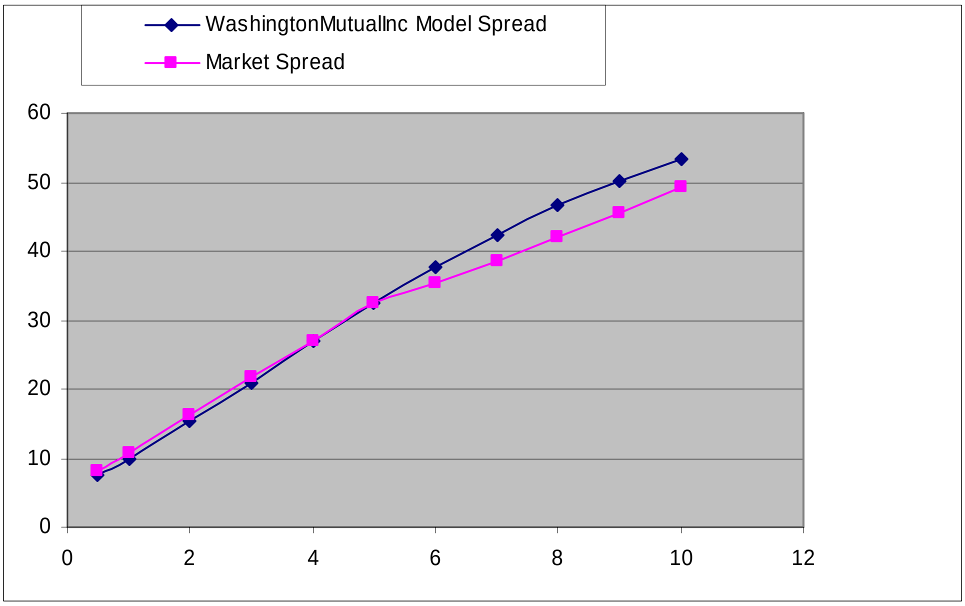

As it happens, there is a large 2-dimensional set of points in gamma-lambda space which all produce good quality fits. This is encouraging, but means that identification of the parameters above is tentative.

There is only space here to briefly mention the dynamics of curve evolution under this model. The dynamics are quite intuitive, with credits most likely to carry down their curve, but with a small chance that spreads will blow out.

We do not use this implied curve either to tell us where the market ought to be or to use it instead of market levels. (For actual calculations we will allow θ to depend on maturity to allow exact matching.) The curve is important because it tells us that the model is credible, and that it can be used as a basis for more complicated trades, such as baskets.

3. Basket pricing

We are now ready to use our model to price tranches on baskets. Prior to describing the new model, we can re-express the one-factor Gaussian copula model as a continuous-time model involving Brownian motions. The entity's "value" *X*(i) is the sum of a global factor and an idiosyncratic factor

$$X_i(t) = \sqrt{\rho}\, W^g(t) + \sqrt{1-\rho}\, W^i(t)\ ,$$

where *W*(*g*) and *W*(i) are independent Brownian motions, and ρ is the correlation between the global factor and the entity. Notoriously, there is no single value of this parameter which reprices all the tranches correctly.

Our new model is based on the BVG processes $X(\rho, \gamma, \lambda)$ we have described above. The model is again that *X*(i) is a BVG equal to the sum of a global BVG factor and an idiosyncratic BVG factor:

$$X_t^i(1,\gamma,\lambda) = X_t^g(\rho,\phi\gamma,\lambda) + \tilde{X}_t^i(1-\rho,(1-\phi)\gamma,\lambda)$$

The jump parameters γ and λ are as before. They can be calculated either by calibrating the CDS curves in the basket, or by calibration to tranche prices. The new parameters ρ and ϕ are generalised "correlations". The parameter ρ is the "Brownian correlation" which is the fraction of Brownian noise attributed to global factors. The other parameter ϕ is the "jump correlation" which is the fraction of incoming jumps due to global factors.

Other previous approaches to CDO modelling with jumps include Joshi and Stacey (2005), Luciano and Schoutens (2005), and Moosbrucker (2006). These have various ways of introducing dependency, but only use purely discontinuous processes.

The new approach is different from these earlier models in two important respects. Firstly the process contains both Brownian motion *and* a variance gamma. This allows jumps to be placed within the context of an otherwise continuously evolving market. Secondly, the new model allows both the continuous and discontinuous components to be divided between global and idiosyncratic effects. Not every jump is global, and not every continuous evolution is idiosyncratic.

This new model addresses many of the problems of the Gaussian copula "Base Correlation" method. It is an actual model with dynamics. It is arbitrage-free. Bespoke tranches can be confidently priced. Bespoke baskets can also be priced, subject to reasonable estimation of the four parameters, of which the two jump parameters can be referenced to the CDS curves in the basket. Parameters are stable and credible. Derived products can, in theory, be priced though much work remains on this point.

3.1 Fitting CDX

We work with the CDX 125 S5 basket. We calibrate the four parameters to the 5y tranche prices. For the market of 9 March 2006, the jump parameters are γ of 8.2% and λ of 15%. The correlations are ρ as 40.5% and ϕ as 9.7%.

Comparing prices for a range of tranches and maturities we have:

	5y CDX 125 S5		**7y CDX 125 S5**		**10y CDX 125 S5**	
Tranche	Market	Model	Market	Model	Market	Model
0% - 3%	33.0%	33.0%	51.3%	48.4%	59.0%	55.4%
3% - 7%	84.0	84.2	243.3	288.2	609.8	642.1
7% - 10%	17.4	17.2	36.6	35.2	101.7	177.0
10% - 15%	9.5	10.7	18.9	17.5	51.1	45.4
15% - 30%	4.0	5.0	6.2	8.1	13.4	13.8
30%-100%	1.9	0.7	2.8	0.9	5.0	1.4

The model fits the 5y market quite well, especially when we remember that it is using only four parameters. For longer maturities, to which we did not calibrate, there are larger differences especially for the mezzanine tranches and (consequently) the super senior.

3.2 Fitting iTraxx

Repeating the process for iTraxx 125 S4, we calibrate jump parameters γ of 15.5% and λ of 15%, and fit the correlations ρ as 10.8% and ϕ as 9.4%. The table of prices is:

	5y iTraxx 125 S4		**7y iTraxx 125 S4**		**10y iTraxx 125 S4**	
Tranche	Market	Model	Market	Model	Market	Model
0% - 3%	26.5%	26.5%	47.7%	48.2%	57.7%	58.9%
3% - 6%	66.4	66.4	220.6	269.9	614.4	695.6
6% - 9%	22.5	22.3	49.8	44.5	130.9	177.4
9% - 12%	11.8	12.4	27.4	24.6	53.9	49.0
12% - 22%	4.7	5.4	9.6	11.1	20.0	20.9
22%-100%	1.4	0.6	3.7	0.9	6.4	1.4

Again the 5y fit is good, but the longer-dated tranches are different in the mezzanine and super senior.

3.3 Remarks on fitting

The fittings shown above are typical for this four-parameter model with no term structure. The fitting is not perfect in two particular respects. Firstly, the CDO-fitted gamma and lambda parameters need not equal the CDS-fitted values, and the gamma parameter particularly does not. The CDS fitting above was tentative, but the gamma difference is an inconsistency.

Secondly, the longer-term dates are not as good as 5y. More accurate fittings for 7y and 10y can be achieved by having a term structure for the parameters, or by allowing the up-jump size to be smaller than the down-jump size. But the mezzanine/super senior feature still remains. We will study this in more detail.

4. Mezzanine tranche pricing

This section will address the mis-match between the market and model prices for longer-dated mezzanine tranches. The evidence suggests that the market price is not purely rational.

We will study the CDX 3%-7% tranche. At 7y the market price for this is about 244bp and the model price is 288bp. At 10y, the market price is 610bp and the model price is 642bp. We aim to show that the market prices are not consistent with a credible set of expected tranche notionals.

A useful quantity to study is the *forward single-period CDS spread*, or *swaplet spread*, defined as

$$Swaplet(i) = \frac{Df_i (N_{i-1} - N_i)}{Df_i \delta_i (N_{i-1} + N_i)/2}.$$

where *Df*(i) is the discount factor on the *i*th payment date and δ(i) is the accrual fraction for the relevant period. Here we also use a standard approximation for including accrued interest in the denominator. This is the fair price to pay for forward-starting single-period credit protection.

We can calculate the 10y strip of swaplet spreads for the tranche both under our new model and using market prices. Calculations under the new model are straightforward because it is arbitrage-free. Using market prices is harder, as there is no standard way to strip the curve, so we have tried to come up with a composite curve that matches the market's tranche prices and has relatively regular swaplet spreads.

The graph shows the swaplet spreads for both the model and the market.

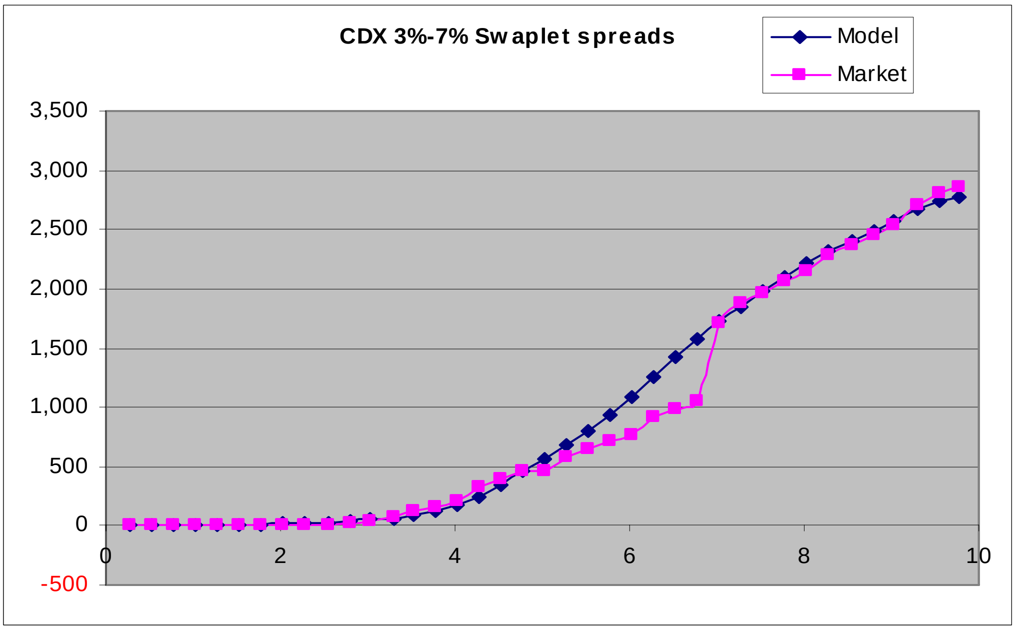

The market curve has problems around 7y. It is not steep enough between 5y and 7y, so has to jump after 7y to get back to the correct trend in order to price 10y correctly. This is caused by the 7y mezzanine price being low in comparison with the rest of the market.

Although this is not a strict arbitrage this behaviour is unintuitive and strongly suggestive that market prices are being driven by factors other than expectation pricing of losses under a consistent measure.

The cause of the market's problems could be demand and supply pressure, which can be a driver of the CDO market. In particular, there is much demand for investors for mezzanine tranches, which forces their spreads tighter.

There are also "technical" factors. The market is not particularly liquid or large and is dominated by a relatively small number of major institutions. Those institutions may well have some price-setting ability.

Also important may be the weakness of the current generation of models, which restricts risk management. The Gaussian copula does not allow cross-maturity hedging. So traders might be unhappy to hedge, say, 7y exposure with 5y baskets. The new model gives more hedging power and does allow such trades. The current lack of hedging opportunities makes it easier to build up unrelieved demand at various points on the maturity-attachment grid.

The mismatch between market and model for the super-senior tranche would also be mostly resolved if the market mezzanine prices adjusted to be closer to the model mezzanine prices.

5. Parameters and their risks

The basic model has just four parameters, apart from the CDS curves themselves. It is useful to gain an understanding of each of these parameters in turn, as well as knowledge of its effect on the various tranches.

5.1 Phi – jump correlation

The parameter ϕ is the correlation between an entity's jumps and global jumps. Typical calibrated values lie around the 10%-20% range. This parameter is the most important driver of senior tranche spread. Indeed the principal way that the senior tranche can default is that if there is a big correlated move in many credits. Big moves can only come from the jump component, and to make them correlated requires ϕ.

Increasing ϕ has the typical correlation impact of tightening equity tranches and widening senior tranches.

5.2 Rho – continuous correlation

The parameter ρ is the correlation between an entity's continuous movement and global continuous movement. Depending on the market, it can calibrate in a range of 0% - 50%, though the risk to this parameter is not high. This parameter behaves similarly to the flat correlation in the Gaussian copula model.

Increasing ρ makes the equity tranche tighter, and mezzanine tranches wider. It has no significant effect on senior tranches, which are relatively immune to continuous movements.

5.3 Gamma – jump intensity

The parameter γ controls the intensity of jumps (both global and idiosyncratic). The parameter is clearly important to pricing senior tranches, but two different effects are at work. If γ is small, then the jump term is dwarfed by the continuous term and the model reduces to the Gaussian copula (with correlation ρ). But conversely if γ is large then the central limit theorem applies to the jump term, which now approximates a Brownian motion. So the model again reduces to a Gaussian copula (but with correlation ϕ).

The effect of increasing gamma can be complex. It will tighten senior spreads, and sometimes tighten equity or junior mezz. Other spreads will widen.

5.4 Lambda – inverse jump size

The parameter λ is the (inverse) size of the jumps. For individual CDS curves, lambda controls the slope of the CDS curve. For tranches, increasing lambda makes the jump component less important, and so moves the model back to being a Gaussian copula.

Increasing lambda generally tightens equity and widens mezzanine. The senior effect is smaller and more variable. If rho is small (less than 5%), the lambda risk can change.

5.5 Risk summary

A table of the risks is useful both to summarise these comments, and to show that the four parameters together give some control of the various tranches.

Parameter	Equity	Mezzanine	Senior	Restrictions
Gamma – intensity	Variable	Variable	Tighter	Gamma > 3%
Lambda – inv. jump size	Tighter	Wider	Small	Rho > 5%
Rho – continuous correlation	Tighter	Wider	None	
Phi – jump correlation	Tighter	Wider	Wider	

6. Implementation

Implementation does not require Monte Carlo, though Monte Carlo simulation is certainly possible. Since the individual entities are conditionally independent given the global factor, typical existing implementations of the Gaussian copula (GC) can be converted to use the new model. The parts that need changing are:

- calculation of the threshold level. Under GC this uses the inverse normal distribution function, which becomes the inverse BVG distribution function.
- conditional survival probabilities. Under GC this uses the normal distribution function, which is now replaced with the BVG distribution function.
- integration against the global factor. With GC this is a numerical integration against the normal distribution, which now uses the BVG distribution.

In every case, all we need is the BVG distribution function. Using the representation of a BVG process marginal distribution as

$$X(t) = Z\sqrt{\rho t + 2\lambda^{-2}\Gamma(\gamma t,1)}\ ,$$

we can express the distribution function of X(t) as

$$P(X(t) \le x) = \int_0^\infty \Phi\left(\frac{x}{\sqrt{\rho t + 2\lambda^{-2} y}}\right) f(y;\gamma t)\, dy\ .$$

Here f(y;γ) is the density of the Gamma(γ) distribution. This integration can be done relatively quickly and easily. For those who integrate by Gaussian quadrature, it is convenient that the gamma distribution is a well-studied case. See, for example Numerical Recipes (1988), chapter 4.5 on the Gauss-Laguerre polynomials.

Also, the Gamma distribution itself can be quickly calculated. Numerical Recipes is again a useful reference with a good treatment in chapter 6.2.

7. Summary and Conclusions

Skew is caused by skew. That is, CDS and CDO pricing is controlled by the tails of the model distributions. But the normal distribution has tails which are too light. Of the various heavy-tailed distributions that might be suitable, we have chosen to use a sum of a continuous Gaussian and a discontinuous gamma. This distribution fits individual CDS curves well, increasing our confidence in the model. There are only two jump parameters so the model is not over-specified.

The model can also be used to price CDO tranches, along with a measure of "correlation". But correlation is not one number. There are two separate correlations - the first between the normal components and the second between the jump components. This model fits the market well, with some exceptions that can be explained.

The model is dynamic, arbitrage-free, and can be extended to other products. It also creates a framework that leaves the door open to further developments by practitioners and academics. Other models in the framework can be created by changing the precise form and parameterisation of the global and idiosyncratic factors to increase sophistication or pricing accuracy.

An important practical point is that the model does not require a Monte Carlo implementation (though it is possible to do so), and CDO products can be priced using an implementation similar to the present standards.

References

David Applebaum (2004) "*Levy Processes and Stochastic Calculus*", Cambridge University Press.

Mark Joshi and Alan Stacey (May 2005) "Intensity Gamma: a new approach to pricing portfolio credit derivatives", preprint. (http://www.quarchome.org/ig.pdf)

Elisa Luciano and Wim Schoutens (Dec 2005) "A Multivariate Jump-Driven Financial Asset Model", preprint. (http://perswww.kuleuven.ac.be/~u0009713/multivg.pdf)

Dilip Madan, Peter Carr and Eric Chang (1998) "The Variance Gamma Process and Option Pricing", *European Finance Review*, Vol2, No 1 (Sep), pp 79-105. (http://www.rhsmith.umd.edu/faculty/dmadan/vgoptj98.pdf)

Thomas Moosbrucker (Jan 2006) "Pricing CDOs with Correlated Variance Gamma Distributions", *Centre for Financial Research*, Univ. of Cologne, colloquium paper (http://www.cfr-cologne.de/downloads/kolloquium/Moosbrucker%20-%20041102.pdf)

William Press, Saul Teukolsky, William Vetterling, Brian Flannery (1988) "*Numerical Recipes in C*", Cambridge University Press.

Wim Schoutens (Jan 2006), "Jumps in Credit Risk Modelling", presentation at King's College London 31 Jan 2006.

Matthias Winkel (2004), "Introduction to Levy processes", graduate lecture at the Department of Statistics, Univ. of Oxford 22 Jan 2004. (http://www.stats.ox.ac.uk/~winkel/lp1.pdf)

Disclaimer

Credit Derivatives