Gamma process dynamic modelling of credit

*The existing generation of credit derivatives models is unsatisfactory because they generally contain arbitrage, cannot describe the dynamics of the process, and are hard to extend beyond vanilla products. **Martin Baxter** has created a new tractable family of credit models, based on the gamma process, which allows arbitrage-free pricing of correlated credits in a dynamic model framework that can straightforwardly handle bespoke baskets and exotic products*

The standard Gaussian copula model, with its overlay of base correlation, is useful but not ideal. It is essentially a static look-up table that does not model the dynamics of the process. It is hard to extend to bespoke baskets or other products, and it readily admits arbitrage.

Various models have been proposed to address these problems. These include both structural models that drive the value of the firm (or proxy) and reduced-form models that drive default intensity rates or the loss distribution. Much academic research has favoured structural models, while practitioners tend to use reduced-form models.

This article focuses on practitioner requirements, but also presents a family of structural models. The model respects the practitioner's needs both for intuitive dynamics and for ease of technical implementation and calibration. It is a jump-based model that allows both global jumps and idiosyncratic jumps. This is important because some credit jump events are global (such as the credit crunch in summer 2007) and others are name-specific (such as Parmalat and Railtrack).

Model description

The economic ideas behind the new model are twofold. First, the tails of distributions are important because default events are tail events. Gaussian tails are too light to match the market, but jump processes introduce heavier tails. Second, jumps can happen for either market-wide or name-specific reasons.

Our new models are based on the gamma process as a common building block. Let $\Gamma(t;\,\gamma,\,\lambda)$ denote the pure-jump increasing Levy process with intensity measure:

$$\nu(x) = \frac{P(\Delta X_t \in dx)}{dx\, dt} = \gamma x^{-1} \exp(-\lambda x), \quad x > 0$$

The parameter γ controls the rate of jump arrivals and λ controls the inverse jump size. Jumps of size $[x, x + dx]$ occur as a Poisson process with intensity $\nu(x)dx$. Applebaum (2004) is a good reference for more details on Levy processes and their stochastic calculus and Winkel (2004) is a helpful brief introduction.

The gamma process has marginal distributions that follow the continuous gamma distribution. Its density at time t is:

$$f_t(x) = \Gamma(\gamma t)^{-1} \lambda^{\gamma t} x^{\gamma t - 1} \exp(-\lambda x), \quad x > 0$$

The key idea for creating dependency between credits is the following lemma.

■ **Lemma (multivariate Levy process).** For any Levy process $X(t)$, we can construct an n-dimensional multivariate Levy process with equal marginal distributions of $X(t)$ and correlation ϕ. Take a global factor X_g and idiosyncratic factors $\tilde{X}_i (i = 1, \ldots, n)$ to be independent identically distributed copies of $X(t)$, and define the ith process to be the sum:

$$X_i(t) = X_g(\phi t) + \tilde{X}_i((1-\phi)t)$$

Then each $X_i(t)$ has the same distribution as $X(t)$. The parameter ϕ is the proportion of the movement of an entity due to global events. If $X(t)$ has finite second moments, then ϕ is the correlation between $X_i(t)$ and $X_j(t)$.

The proof is trivial from the properties of a Levy process.

A first application gives our primary model, which we call the gamma model. We model that an entity's structural value under the risk-neutral measure is given by the exponential martingale:

$$S_t^i = S_0^i \exp\left(X_t^i + \mu t\right), \quad \text{with}$$

$$X_t^i = -\Gamma_g(t;\phi\gamma,\lambda) - \Gamma_i(t;(1-\phi)\gamma,\lambda) \quad \text{and} \quad \mu = \gamma \log\left(1+\lambda^{-1}\right)$$

where the two gamma processes Γ_g and Γ_i are independent global and idiosyncratic factors, with equal jump size but different intensities.

■ **Model definition.** We define a default threshold curve k_t^i, and the entity is deemed to have defaulted at the first time that the structural value hits or goes below the threshold. The default time is defined to be:

$$\tau_i = \inf\left\{t : S_t^i \le k_t^i\right\}$$

The default threshold curve k_t^i can be calibrated to match the survival probabilities of the entity from its credit default swap curve.

The economic idea underlying the model is that the tradable value process is a positive martingale that has a general upward trend but suffers from down jumps as (bad) news arrives. Many of these jumps are small, but some might be quite large.

■ **Implementation approximation.** The proper model as defined above is theoretically attractive, but awkward to implement. To simplify the implementation, we make the following

A. Results of calibrating CDX NA IG 8 on August 22, 2007

	Five-year CDX			Seven-year CDX			Ten-year CDX		
Tranche	Market	Gamma model	Crawley cat model	Market	Gamma model	Crawley cat model	Market	Gamma model	Crawley cat model
0–3%	49.3%	54.5%	49.6%	60.3%	59.7%	58.6%	65.0%	68.8%	66.7%
3–7%	206.0	208.0	202.8	337.0	379.1	347.8	670.0	673.0	664.4
7–10%	52.0	65.6	55.8	124.0	129.6	109.5	180.0	201.2	176.0
10–15%	30.0	35.4	29.4	57.0	71.3	62.3	87.0	106.4	90.7
15–30%	17.5	18.3	14.6	30.0	37.4	33.0	49.0	49.3	43.9
30–100%	8.0	4.4	9.2	14.0	9.3	13.4	16.0	11.5	16.4
Fit score		9.2	2.1		9.7	5.4		11.5	4.2
Gamma		5.0%	52.5%		5.0%	8.0%		5.0%	6.3%
Phi		26.4%	13.7%		39.2%	28.9%		37.8%	28.1%
Cat rate (bp)			8.8			9.3			10.6

approximation. The entity is deemed to have defaulted by time t if $S_t^i \le k_t^i$, or equivalently if $X_t^i \le \theta_t^i$, where k_t^i and/or θ_t^i are recalibrated using the approximation to match exactly the entity's market survival probabilities. This 'European' default condition is a tractable approximation to the exact barrier condition above. We will check the accuracy of this approximation below.

■ **Limiting cases.** The first of two interesting limiting cases is when gamma tends to infinity (keeping $\lambda = \sqrt{\gamma}/\sigma$, for some volatility σ), so that the structural value process converges to log-Brownian motion, $S_t^i = S_0^i \exp(\sigma W_t^i - \frac{1}{2}\sigma^2 t)$, by the central limit theorem. The Gaussian copula is thus a special case of the gamma model for high gamma. At the other extreme, as gamma tends down to zero (setting $\lambda = \exp(-h/\gamma)$, for some intensity h), then the value process becomes a single-jump martingale driven by an exponential default time:

$$S_t^i = S_0^i \exp(ht) I(\tau_i > t)$$

where $\tau_i \sim \exp(h)$ is a random variable. These two extremes correspond economically to cases where there is high spread volatility and diffusive spread behaviour before default (high gamma), such as Argentina, and where there is default with little or no warning in terms of spread movements (low gamma), such as Parmalat. The actual behaviour will be between these two extremes, with the gamma parameter controlling the relative likelihood of diffusive default versus jump default.

■ **Crawley catastrophe model.** We can create a variant of the model by adding a low-intensity high-impact global 'catastrophe' term to create disaster scenarios that might help reprice senior tranches. The model's form is that of a gamma model plus an (infinite) Poisson catastrophe term:

$$X_t^i = -\Gamma_g(t;\phi\gamma,\lambda) - \Gamma_i(t;(1-\phi)\gamma,\lambda) - Y.C(t)$$

where $C(t)$ is a very low-intensity global Poisson process whose intensity is called the catastrophe rate, and Y is an arbitrarily large number.

Other variants are also possible by adding upwards gamma jumps, Brownian motion terms or combinations of these. All the models share the common feature of gamma process down-jumps that contain both global and idiosyncratic jumps.

Pricing baskets

We are now ready to use our models to price tranches on baskets. Let us start with the basic gamma model. It is sufficient to model the log-value process $X(t)$, which comprises a global gamma process and an idiosyncratic gamma process:

$$X_t^i = -\Gamma_g(t;\phi\gamma,\lambda) - \Gamma_i(t;(1-\phi)\gamma,\lambda)$$

The default thresholds θ_t^i are calibrated to match the market survival probabilities of entity i at the swap times t (see 'Implementation approximation' above). Because λ is just a scale factor for X, and thus for θ, it can be ignored.

■ **Calibration.** The calibration process for a given basket and a particular maturity is as follows. We start by fixing λ at any arbitrary rate (for example, 100%). For any gamma-phi pair we can calibrate the thresholds θ_t^i for every name i and time t. The model is now parameterised and can be run to give the spreads of tranches within the basket's capital structure. A goodness-of-fit score can be calculated by comparing the model's tranche prices with the market's. We choose our objective function to be the root weighted-mean squared error (weighting senior tranches more heavily because their spreads are smaller):

$$V(\gamma,\phi) = \left(\sum_{\substack{tranche\\ j=1}}^{m} w_j \left(S_j^{market} - S_j^{model}\right)^2 / \sum_{k=1}^{m} w_k \right)^{1/2}$$

with weight $w(j) = 1/spread(j)$. We then vary (γ, ϕ) to find the minimum value of V. The implementation is described below. The results of calibrating to the quoted tranches on the CDX North America Investment Grade basket, on August 22, 2007 are shown in table A.

The fitting is relatively good but not exact, with the catastrophe model fitting better than the simple gamma model, especially for super-senior. The other main errors are for some equity tranches.

For practical reasons, it is necessary to match the market in liquid tranches precisely. This can be accomplished with some non-model price adjustment factors. This could be a 'base phi' approach as per base correlation, albeit with a much flatter curve, or some other scheme. The market prices for the CDX NA IG 8 tranches are collated from quotes received by Nomura from several other dealers. Each price is halfway between the best bid and best offer over the dealers polled.

■ **Risks.** Typically, gamma risk is negative for equity and senior tranches, and positive for mezzanine tranches (for bought protection positions). Phi risk, like base correlation risk, is negative for junior tranches and positive for mezzanine and senior tranches. In terms of the base correlation curve, phi controls the average curve level (increasing phi shifts the curve up in parallel) and gamma controls the slope of the curve (increasing gamma makes the curve flatter). Intuitively, increasing gamma moves the model

1 Gamma model CDX calibration

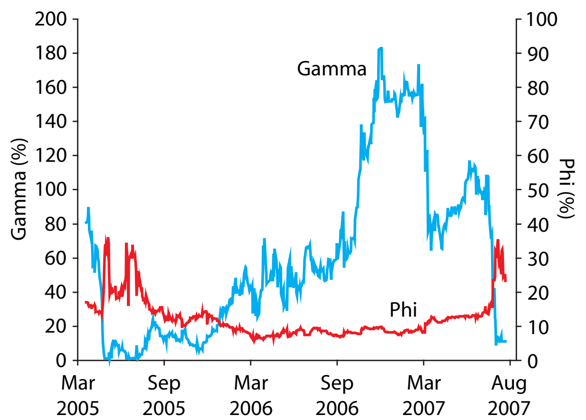

closer to the Gaussian copula, as above ('Limiting cases'), which decreases equity and senior spreads. Increasing gamma does not change entities' survival probabilities, because the default thresholds will recalibrate to keep them constant.

The gamma and phi risk 'vectors' are linearly independent, which is useful for both calibration and risk management. The fact that the gamma model has only two parameters, and that they both have an intuitive interpretation, is a notable feature of the model.

Parameter history. Parameter stability is also an important practical concern. If the fitted parameters vary widely from day to day, the gamma model will be less useful for hedging and will be less credible as a description of the process dynamics. Figure 1 shows the fitted gamma and phi parameters through time for the on-the-run CDX five-year basket from April 1, 2005 to August 22, 2007.

We see that the parameters are mostly stable from day to day, but there are trends and movements. One striking feature is the auto crisis of May 2005. The five-year gamma level was trading quite high at around 80% and crashed down to 10%. This caused equity (and senior) spreads to widen and mezzanine spreads to tighten, corresponding to a steepening of the base correlation curve. The phi levels did not change particularly during this period, though they became more volatile shortly afterwards. In terms of the gamma model, this event could be described as a gamma crash.

During 2006, phi was stable but gamma later rallied to relatively high levels. This corresponded to a flattening of the base correlation curve and an increase in mezzanine spreads relative to equity and senior tranches, which some called the 'reverse correlation crisis' of November 2006. Much of the gamma gains rapidly unwound during credit volatility in March 2007.

In summer 2007, as credit spreads widened, gamma crashed again and phi increased sharply. This corresponded particularly to the large increases in spreads of senior tranches as the base correlation curve became steeper at the senior end.

B. Tranche delta ratios

Tranche	Base corr absolute	Base corr relative	Gamma model
0–3%	23.8	27.5	27.9
3–7%	5.2	5.2	5.2
7–10%	1.2	1.1	1.1
10–15%	0.5	0.4	0.4
15–30%	0.2	0.3	0.1

Risk management. Under the gamma model, there are two sorts of risk – credit spread risk and correlation skew risk. Credit deltas can be easily calculated both for each individual name and for the basket index. This is possible because the model is a 'bottom-up' model, which describes each name's behaviour. The delta ratios under the gamma model for five-year CDX can be compared against base correlation (using either absolute or relative attachment-mapping). The values, as at December 2006, are quite similar but not identical (see table B).

The skew parameter risks are described above, and they provide a new way of managing risk on collateralised debt obligation (CDO) portfolios. Gamma and phi risk can be added up over all tranches in a portfolio, even including bespoke tranches, bespoke baskets and exotic products. The total gamma and phi risks can then be monitored and managed by trading in liquid tranches to achieve the desired exposure. This is an improvement on base correlation risk management, which is essentially inventory-based and which finds it difficult to net skew risk across either the capital structure or bespoke baskets/products.

During the May 2005 crisis, some investors had on the 'positive carry' trade, which sold equity protection and bought mezzanine protection. This trade was about flat in credit spread risk and long correlation. Although it was unknown at the time, this was also a long gamma position. As gamma crashed, so the position lost money. Having had gamma-phi risk management could have given a possible warning signal.

On a practical note, the gamma model has been used for official pricing and front-office risk management at Nomura since August 2006. This period includes the correlation crises of November 2006 and March 2007, as well as the credit crunch of summer 2007.

Implementation and applications

Outline implementation. Although not a copula model, the gamma model can be implemented in a similar way to the Gaussian copula. The main difference is that the normal distribution function is replaced with the gamma distribution function, which has fast and accurate approximations. See, for example, sections 6.2 (incomplete gamma) and 4.5 (quadrature integration) of Press *et al* (1988). The inverse distribution function can be found with interval bisection followed by a few Newton-Raphson iterations to polish it.

The calibration for the gamma model is relatively straightforward. There are only two parameters (gamma and phi), and they both have non-trivial and different effects on the tranche spreads. We use an optimiser similar to the Levenberg-Marquardt method, described in section 15.5 of Press *et al* (1988). About half a dozen iterations are enough to get a good calibration.

European approximation compared with barrier method. We can test the European implementation approximation described above against the true barrier price. We implemented the barrier method using numerical tree methods, which is feasible but runs significantly more slowly than the European formula.

We can repeat the calibration for the CDX basket using the barrier method. The model tranche prices are almost identical to the prices from the European approximation. All the differences are less than a 0.1-basis-point spread or 0.1% upfront, except for 10 years, where the maximum difference was a 0.6bp spread. The fitted values of the parameters that achieve this are different, as we might expect, but not markedly. The gammas were up to 3.5% different and the phis up to 2.4% different. So using the better formula results in a small change in co-ordinates, but no significant difference to the model's prices.

Copulas and arbitrage. Although the gamma model implementation is similar to a copula implementation, the model is not a copula. A copula joins together some individual marginal distributions to make a multivariate marginal distribution, but it does not provide inter-temporal dynamics or a natural probability filtration. For instance, a copula cannot in general give the conditional probability that a credit defaults at a later time given the basket's state at an earlier time. Paradoxically, the Gaussian copula is an exception to this restriction, but when base correlation is introduced the limitation returns.

In contrast, any structural model introduces a natural filtration generated from its value process, and any conditional path-dependent probability is theoretically available. Credit spread dynamics are also created. For instance, the gamma model will show a spread widening as the value process approaches the default barrier and a spread tightening as it drifts away.

Arbitrage is another unpleasant consequence of the base correlation approach. As there is no single correlation number, there can be no equivalent martingale measure that reprices all tranches. The absence of such a measure guarantees the existence of arbitrage. See, for example, section 6.6 of Baxter & Rennie (1996). Practically, this arbitrage can appear as non-monotonic or negative tranche par-spreads. Our new Levy models do not allow arbitrage because there is a martingale measure, which removes this important problem.

■ **Extensions to bespoke baskets and exotic products.** The gamma model will price bespoke baskets, given suitable levels of gamma and phi for that basket. While they can never be known with absolute certainty, they can be reasonably estimated by reference to the liquid CDO market. Bespoke basket parameters can be estimated using suitable simple interpolation schemes, such as linear interpolation based on basket spread between liquid investment-grade and high-yield basket parameters. By contrast, base correlation interpolation has noted difficulties in its choice of attachment-point mapping scheme, which is both arbitrary and price-sensitive.

More exotic products, such as CDO-squared and long-short CDOs, can also be priced under the model. By conditioning on the global factor, the product can be priced using the conditional independence of the names. This approach to exotic products will automatically include correlation skew in the pricing, in a dynamic arbitrage-free way.

For some products, we need to calibrate simultaneously to more than one maturity. We do this by first calculating a term structure of gamma-phi parameters (γ_t, ϕ_t) for each marginal loss distribution at time *t*. We can then strip these parameters to build up the global and idiosyncratic gamma processes that match these marginals. Term-structure modelling is theoretically possible.

Pricing path-dependent products, such as forward-starting CDOs, requires both term-structure modelling and a dynamic model. A model such as the gamma model would be able to price these products if its spread dynamics were a reasonable description of reality. The possibility is interesting, but more work needs to be done.

References

Applebaum D, 2004
Levy processes and stochastic calculus
Cambridge University Press

Baxter M and A Rennie, 1996
Financial calculus: an introduction to derivative pricing
Cambridge University Press

Press W, S Teukolsky, W Vetterling and B Flannery, 1988
Numerical recipes in C
Cambridge University Press

Winkel M, 2004
Introduction to Levy processes
Graduate lecture at the Department of Statistics, University of Oxford, January 22, available at *www.stats.ox.ac.uk/~winkel/lp1.pdf*

Summary and conclusions

We have introduced a family of Levy process structural models, along with a simple correlation structure. The family is centred around the gamma model. The models fit the market CDO prices better than the Gaussian copula. They also describe the dynamics of the correlated processes in a consistent arbitrage-free way. The models can handle bespoke baskets and are capable of being extended to more complex products.

Among Levy models, the simple gamma model has a number of advantages. It fits relatively well, and is practical to implement and calibrate. It has only two parameters, which are stable, intuitive and allow parameter-based portfolio risk management.

It is, of course, possible to develop more sophisticated models within this framework to achieve better accuracy or more precise dynamics. An example is the catastrophe model, which is better at fitting super-senior prices. ■

Martin Baxter is a quantitative analyst at Nomura International in London.

He would like to thank Edmund Crawley and Artur Steiner for their help and ideas.

Email: work@martinbaxter.co.uk

Guidelines for the submission of technical articles

Risk welcomes the submission of technical articles on topics relevant to our readership. Core areas include market and credit risk measurement and management, the pricing and hedging of derivatives and/or structured securities, and the theoretical modelling and empirical observation of markets and portfolios. This list is not an exhaustive one.

The most important publication criteria are originality, exclusivity and relevance – we attempt to strike a balance between these. Given that *Risk* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the technical editor and his team, submissions are logged, and checked against the criteria above. Articles that fail to meet the criteria are rejected at this stage.

Articles are then sent to one or more anonymous referees for peer review. Our referees are drawn from the research groups, risk management departments and trading desks of major financial institutions, in addition to academia. Many have already published articles in *Risk*. Depending on the feedback from referees, the technical editor makes a decision to reject or accept the submitted article. His decision is final.

We also welcome the submission of brief communications. These are also peer-reviewed contributions to *Risk* but the process is less formal than for full-length technical articles. Typically, brief communications address an extension or implementation issue arising from a full-length article that, while satisfying our originality, exclusivity and relevance requirements, does not deserve full-length treatment.

Submissions should be sent to the technical team at *technical@incisivemedia.com*. The preferred format is MS Word, although Adobe PDFs are acceptable. The maximum recommended length for articles is 3,500 words, and for brief communications 1,000 words, with some allowance for charts and/or formulas. We expect all articles and communications to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations. Authors should allow four to eight weeks for the refereeing process.