Peter Jäckel*

The Discrete Gamma Pool model

First version: 19th September 2007
This version: 16th November 2008

Abstract

We propose a model for the dynamics of losses and spreads on portfolios for the purpose of pricing exotic variations of synthetic collateralised tranche obligations such as *Loss Triggered Leveraged Super-Senior* notes, *multi-callable CDOs*, and, by implication of the latter, *options on forward starting CDOs*. Also, we discuss how features such as *the counterparty's right to deleverage upon a loss trigger event* in a leveraged super senior can be understood as an embedded Bermudan swaption, and how this can be catered for in a numerical implementation.

1 Introduction

The market for correlation credit derivatives has seen a tremendous growth both in terms of volume, and in terms of product range. In recent years, various new products have become popular that exhibit not only sensitivity with respect to co-dependence of the respective obligor basket, as well as their respective credit spreads, but also with respect to *volatility of spreads*, such as *Loss Triggered Leveraged Super-Senior* (LTLSS) synthetic tranche obligations. Even *multi-callable collateralized debt obligations* have been considered for bespoke transactions. In this article, we describe an alternative model that caters for the specific features that are necessary for adequate pricing and risk management of products such as the LTLSS and CDO options.

The new model draws on ideas from Baxter's Gamma-Phi model [Bax06, Bax07], large homogeneous pool approximations [Vas87], as well as the Black-Karasinski model for credit default hazard rates [BK91]. As a whole, it combines a series of characteristics that are desirable for the modelling of stochastic spread sensitive credit correlation products. With respect to the stochasticity of spreads, it resembles other models such as discussed in [SPA05, CRT06, Lip06, LI07a, LI07b] to some extent, but with more emphasis on the simplifying benefits of top-down approaches as, e.g., also done in [GG05, EGG06, BPT07]. The closest known models are discussed in [AH07, LM07] which all are presented in a finite differencing setting and are based on a two-dimensional state space framework with one dimension representing a process driving some form of hazard rate, and the other dimension representing portfolio loss. In contrast to other top-down-esque frameworks, this model does, though, retain the concept of an explicit credit correlation parameter by the aid of using the large homogenous portfolio (LHP) approximation to create a discrete time step loss transition probability function. Since the LHP approximation of Baxter's Gamma model inherits the correlation parameter ϕ, the Discrete Gamma Pool model thus has a very natural and intuitive correlation parametrisation as part of its intrinsic structure. Because of this ambivalence of being on the one side based on a dynamic portfolio loss and loss intensity process, and on the other hand on an explicit correlation modelling approach we consider this model to be *half way* between *top-down* and *bottom-up*. The explicit representation of correlation makes it comparatively user friendly when, for instance, applied to optionality products on synthetic collateralised debt obligations and tranches. In addition, the model also allows for a total portfolio default scenario, which has been found to be crucial for fair representation of super-senior risk by other authors, too [Bax06, Bax07, BPT07, HW05].

*OTC Analytics

2 The building blocks

2.1 Baxter's Gamma-Phi model [Bax06, Bax07]

This model is, at its most fundamental level, based on the concept of a *distance-to-default* process for each of the underlying reference obligors. The distance-to-default process is chosen to be a Gamma process [PW84, DGS91, CT04], which is a strictly increasing and strictly discontinuous Lévy process whose law for shape parameter γ for time horizon t is given by

$$\mathsf{Prob}\left\{x(t) < x^{\star}\right\} \;=\; \Gamma(\gamma t, x^{\star}) \tag{1}$$

wherein $x(t)$ is the Gamma process and

$$\Gamma(\gamma t, x^{\star}) \;=\; \int_{0}^{(x^{\star})_+} g(\gamma t, x)\,\mathrm{d}x \tag{2}$$

with

$$g(\alpha, x) \;=\; x^{\alpha-1}\mathrm{e}^{-x}\big/\,\Gamma(\alpha) \tag{3}$$

and $\Gamma(\cdot)$ being Euler's Gamma function which satisfies

$$\Gamma(n+1) \;=\; n! \tag{4}$$

for non-negative integers n. Within this reduced form setup, default of the underlying obligor is indicated by the distance-to-default process first breaching a level $x^{\star}(t)$ which is defined as a function of time such that the probability of survival matches a given target function $p_{\text{survival}}(t)$. We show some sample curves for Gamma processes in figures 1 for $\gamma = 1$ and 2 for $\gamma = 100$, respectively. It is apparent in figure 2 that, as Baxter also remarks in [Bax06, Bax07], that the distance-to-default process becomes a standard Wiener process as $\gamma \to \infty$.

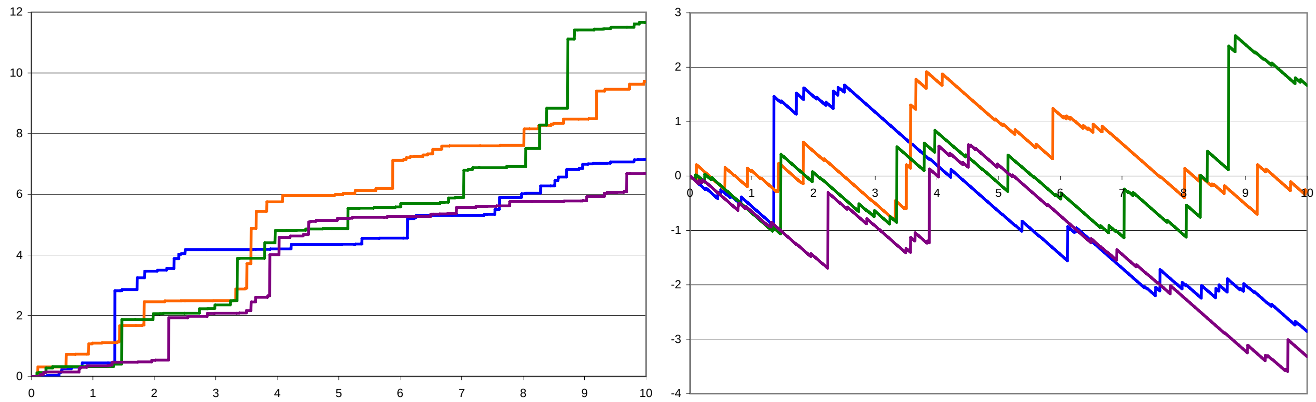

Figure 1: Four sample paths of a Gamma process for $\gamma = 1$. The right hand side figure shows the same data as the one on the left after subtraction of the drift γt.

In a single random variate approximation, the distance-to-default process can be represented by a deterministic function of a single uniform variate u that is representative of the Gamma process to some specific time horizon T such that the *marginal* law of the process is the same as in (1). This *constant quantile* approximation is given by

$$x_{\text{constant quantile}}(t) \;=\; \Gamma^{-1}(\gamma t, u) \tag{5}$$

wherein $x_{\text{constant quantile}}(t)$ stands for the single random variate approximation to $x(t)$. Within this setting, calibration of the default barrier to a given survival probability function $p_{\text{survival}}(t)$ is simply given by the analytical expression

$$x^{\star}_{\text{constant quantile}}(t) \;=\; \Gamma^{-1}(\gamma t, p_{\text{survival}}(t))\;. \tag{6}$$

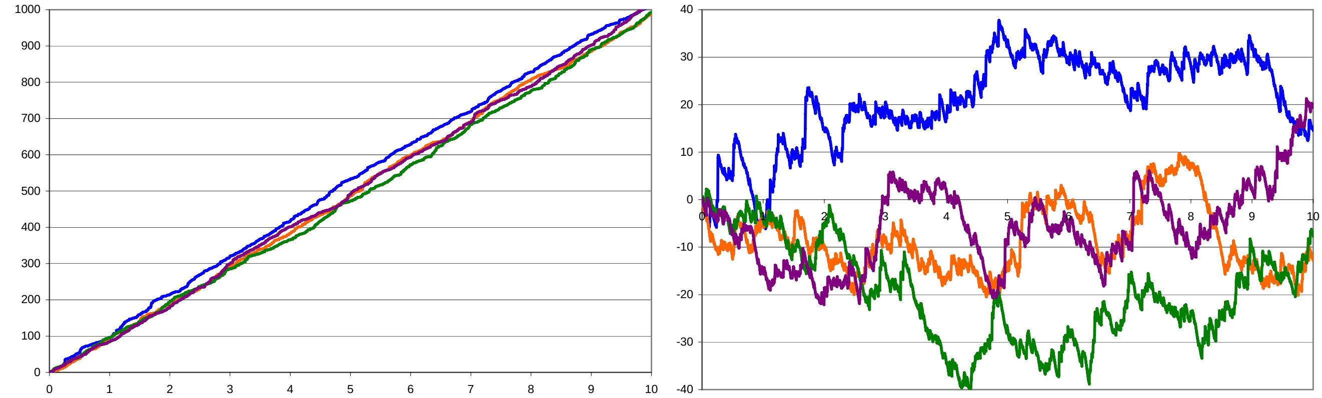

Figure 2: Four sample paths of a Gamma process for $\gamma = 100$. The right hand side figure shows the same data as the one on the left after subtraction of the drift γt.

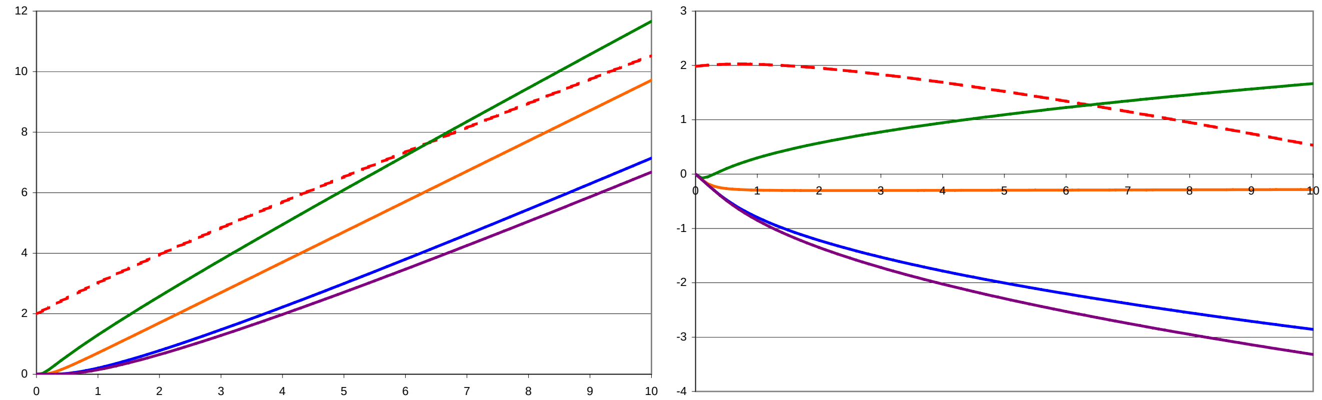

Figure 3: Constant quantile approximation (5) to the four Gamma process paths in figure 1 for $\gamma = 1$. The default barrier (dashed line) is defined by $x^\star(t) = \Gamma^{-1}(\gamma t, \mathrm{e}^{-\lambda t})$ with $\lambda = 5\%$, resulting in the survival probability being $\mathrm{e}^{-\lambda t}$. The right hand side figure shows the same data as the one on the left after subtraction of the drift γt.

The constant quantile approximations for the paths in figures 1 and 2 are shown in figure 3 and 4, respectively.

An alternative way to define a single random variate approximation to a Gamma process path, and this is what Baxter refers to as the *European approximation*, is to define the process along its entire path by its terminal value, i.e. to set

$$x_{\text{constant value}}(t) \;=\; x(T) \;=\; \Gamma^{-1}(\gamma T, u)\;. \tag{7}$$

for some terminal time horizon T, typically the maturity of a given contract. For this *constant value*

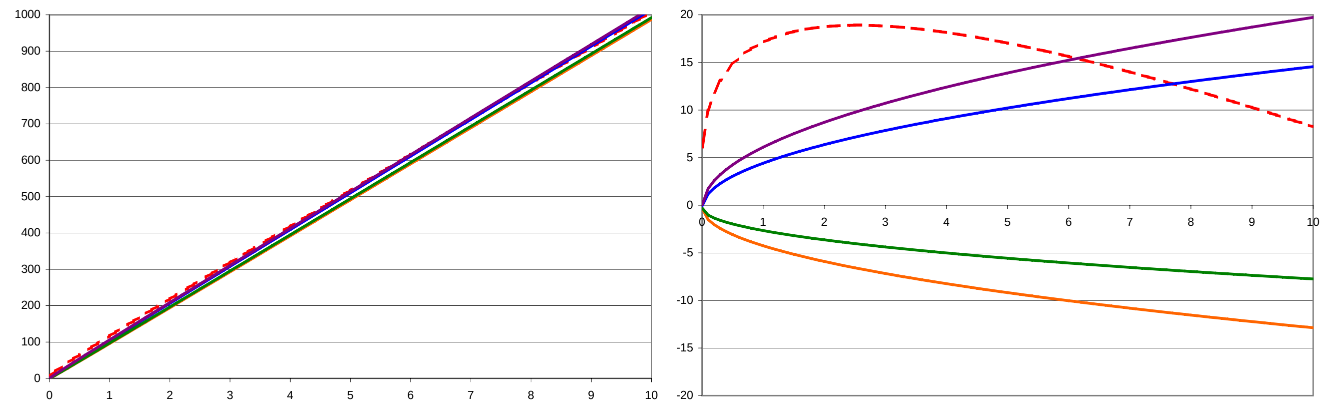

Figure 4: Constant quantile approximation (5) to the four Gamma process paths in figure 2 for $\gamma = 100$. The default barrier (dashed line) is defined by $x^\star(t) = \Gamma^{-1}(\gamma t, \mathrm{e}^{-\lambda t})$ with $\lambda = 5\%$, resulting in the survival probability being $\mathrm{e}^{-\lambda t}$. The right hand side figure shows the same data as the one on the left after subtraction of the drift γt.

approximation, the analytical calibration of the default barrier level is

$$x^\star_{\text{constant value}}(t) \;=\; \Gamma^{-1}(\gamma T, p_{\text{survival}}(t)) \;. \tag{8}$$

The constant value approximations for the paths in figures 1 and 2 are shown in figure 5.

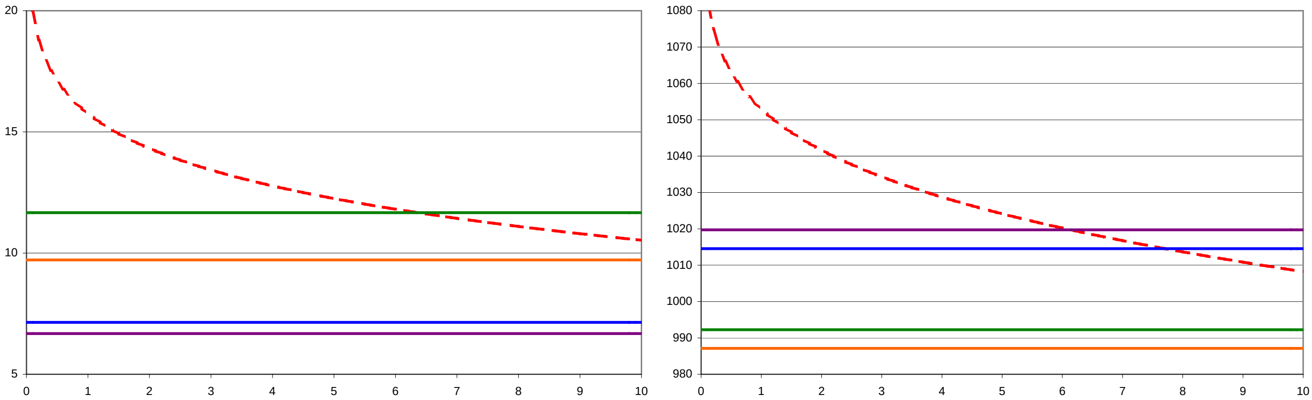

Figure 5: Constant value approximation (7) to the Gamma process paths in figures 1 for $\gamma = 1$ (left) and 2 for $\gamma = 100$ (right). The default barrier (dashed line) is defined by $x^\star(t) = \Gamma^{-1}(\gamma T, \mathrm{e}^{-\lambda t})$ with $\lambda = 5\%$ and $T = 10$, resulting in the survival probability being $\mathrm{e}^{-\lambda t}$.

Codependence between single name distance-to-default processes is introduced into this model by decomposing the contribution of uncertainty to the process into a component that is common to all names, and an idiosyncratic component. It is assumed that a proportion of ϕ with $\phi \in [0, 1]$ of the uncertainty originates from the common component, and the remaining proportion of size $(1 - \phi)$ is given by an independent idiosyncratic component. If we identify the concept of *uncertainty* with information arrival, which is nothing other than the flow of time, then, by virtue of the fact that the Gamma process is a Lévy process with identically distributed independent increments, we can postulate that the total single name distance-to-default process $x_i(t)$ is given by a linear combination of the common process x_{common} and an idiosyncratic part $\tilde{x}_i$:

$$x_i(t) \;=\; x_{\text{common}}(\phi t) + \tilde{x}_i((1-\phi)t) \tag{9}$$

where we assume that $x_{\text{common}}(t)$ and $\tilde{x}_i(t)$ are Gamma processes for the same shape parameter γ. It is important to note that equation (9) does *not* mean that one process is looked up at time ϕt, and the other one at time $(1-\phi)t$. Instead, it is to be read that the common process provides a certain percentage of the information flow driving the distance-to-default process, and the other one the remainder. An alternative to understand this is to assume that the name-specific distance-to-default process $x_i(t)$ receives a continuous stream of small random upwards shocks. At any point in time, the small positive shock originates from a common Gamma process with probability ϕ, and from an idiosyncratic Gamma process that is specific to this name with probability $(1 - \phi)$. The random choice variable at any point in time as to whether the individual's shock originates from the common process or from the idiosyncratic process is common to all underlying reference obligations. In distribution, a set of individual processes constructed in this fashion behaves exactly as constructed by (9).

Within the constant value approximation (7) and (8), we have

$$x_{\text{common}} \;\sim\; \Gamma(\gamma\phi T, \cdot) \qquad \text{and} \qquad \tilde{x}_i \;\sim\; \Gamma(\gamma(1-\phi)T, \cdot) \;. \tag{10}$$

Note that in the constant value approximation the final maturity T only serves, effectively, to linearly rescale the Gamma parameter, whence we shall drop it from hereon, i.e.

$$\gamma \cdot T \quad \longrightarrow \quad \gamma \;.$$

The probability of survival of a single name conditional on the common factor is thus given by

$$\mathsf{Prob}\left\{x_i < x_i^\star(t) \middle| \, x_{\text{common}}\right\} \;=\; \mathsf{Prob}\left\{x_{\text{common}} + \tilde{x}_i < x_i^\star(t) \middle| \, x_{\text{common}}\right\} \tag{11}$$

$$= \Gamma\left(\gamma(1-\phi), x^\star(t) - x_{\text{common}}\right) \quad (12)$$

$$= \Gamma\Big(\gamma(1-\phi), \Gamma^{-1}(\gamma, p_i(t)) - \Gamma^{-1}(\gamma\phi, u)\Big) \quad (13)$$

with u being the uniform common factor, i.e.

$$x_{\text{common}} = \Gamma^{-1}(\gamma\phi, u) , \quad (14)$$

and $p_i(t)$ being the unconditional survival probability of the respective reference obligor. It is worth noting that this model for codependent survival probabilities and thus default times converges to the setting of the standard single correlation number Gaussian copula model [Vas87, Li00] in the limit of $\gamma \to \infty$. In this case, the parameter ϕ in Baxter's notation equals the square root of the flat portfolio correlation figure.

As a final note in this section we would like to point out that, in our notation, the common factor's meaning, both as a uniform or as a Gamma variate as related by (14), represents how *bad* the global economy is. This means that high values of x_{common} reduce conditional survival probabilities, which may differ in other author's notation. This point is worth mentioning particularly in the context of a model based on an asymmetric distance-to-default process such as the Gamma process whereas it is typically glanced over for the Gaussian copula model since there, the alignment of the common factor becomes irrelevant when one eventually arrives at the calculation of conditional survival probabilities. For the Gamma-Phi model this is not so: the choice of alignment of the common factor makes a distinct difference for the overall behaviour of the model, and indeed, choosing it not the way we explain above[1] creates an entirely different model whose resulting loss distributions would be irreconcilable with market observations.

2.2 Large homogeneous pool approximations

For a homogeneous portfolio so large that individual obligor's contribution to the overall portfolio value can be seen as near-infinitesimal, the proportion of the portfolio that defaults over any finite time horizon is equal to one minus the single name survival probability itself [Vas87]. If we assume that the loss level in a portfolio of unit total value at the beginning of a time step is denoted as ℓ, and that all defaults recover at the same recovery rate R, then, conditional on a uniform common factor, the incremental loss over a time step of size Δt is

$$\Delta\ell = (1-R)(1-\tfrac{\ell}{1-R})(1-p(\Delta t)|\,u) . \quad (15)$$

The terms on the right hand side of (15) represent, in order: proportion of the portfolio notional that can be lost $(1-R)$, proportion of the portfolio notional that has not defaulted to date $(1-\frac{\ell}{1-R})$, and the conditional default proportion over the time step, the latter being one minus the conditional single name survival probability $p(\Delta t)|\,u$ (conditioned on the uniform common factor value u). Solving (15) for the conditional survival probability yields

$$p(\Delta t)|\,u = 1 - \tfrac{\Delta\ell}{1-R-\ell} . \quad (16)$$

For the Gamma-Phi model, this conditional survival probability is given by

$$p(\Delta t)|\,u = \Gamma\Big(\gamma(1-\phi), \Gamma^{-1}(\gamma, p(\Delta t)) - \Gamma^{-1}(\gamma\phi, u)\Big) \quad (17)$$

where $p(\Delta t)$ is the unconditional single name survival probability over the time step Δt. Substituting this back into (16) gives us a one-to-one relationship between the incremental loss and the common factor

$$\Gamma\Big(\gamma(1-\phi), \Gamma^{-1}(\gamma, p(\Delta t)) - \Gamma^{-1}(\gamma\phi, u)\Big) = 1 - \tfrac{\Delta\ell}{1-R-\ell} . \quad (18)$$

Given the uniform law of u, this can be resolved for the distribution of the loss increment:

$$\mathsf{Prob}\left\{\ell(t+\Delta t) < \ell(t)+\Delta\ell\right\} = \Gamma\Big(\gamma\phi, \Gamma^{-1}(\gamma, p(\Delta t)) - \Gamma^{-1}(\gamma(1-\phi), 1 - \tfrac{(\Delta\ell)_+}{1-R-\ell(t)})\Big) . \quad (19)$$

[1] The way we choose the orientation of the uniform common factor is consistent with Baxter's choice.

2.3 The Black-Karasinski model for credit default hazard rates

The Black-Karasinski model was originally designed for interest rate derivatives [BK91]. One of its drawbacks is that its continuous dynamics imply that the expectation of rollover returns diverges since short rates evolve lognormally [HW93, SS97], which in turn implies that interest rate futures are infinite. This is because when a short rate r follows a geometric Brownian motion, the expectation

$$\mathsf{E}\left[\mathrm{e}^{\int_t^T r(t')\,\mathrm{d}t'}\right] \tag{20}$$

is undefined, and it is this term that links futures for interest rates with forward interest rates under any model [CIR81]. However, if we replace $r(t)$ with $\lambda(t)$ and interpret the short rate as an instantaneous hazard rate, no problems arise since expressions such as (20) do not come up in the context of credit derivatives modelling based on the generic concept of a Cox process, i.e. a reduced form default model given by a Poisson process with stochastic intensity.

For our purposes, we define the Black-Karasinksi model for the credit default hazard rate $\lambda(t)$ as given by

$$\lambda(t) \;=\; \hat{\lambda}(t) \cdot \mathrm{e}^{-\frac{1}{2}\mathsf{V}_0[y(t)]+y(t)} \tag{21}$$

wherein $y(t)$ is an Ornstein-Uhlenbeck process $y(t)$ whose dynamics are consistent with the stochastic differential equation

$$\mathrm{d}y \;=\; -\kappa \cdot y \cdot \mathrm{d}t + \sigma \cdot \mathrm{d}W \;, \tag{22}$$

and $\hat{\lambda}(t)$ is a deterministic hazard rate amplitude function. For constant σ and κ, the variance function $\mathsf{V}_0\left[y(t)\right]$ is given by

$$\mathsf{V}_0\left[y(t)\right] \;=\; \sigma^2 \cdot \left(1 - \mathrm{e}^{-2\kappa t}\right) / \left(2\kappa\right) \,. \tag{23}$$

3 The Discrete Gamma Pool model

The Discrete Gamma Pool model is a model for the dynamics of *losses* and *spreads* on a portfolio of defaultable reference assets. The evolution of losses and spreads is defined for discrete steps in time (typically two weeks). Given a homogeneous portfolio hazard rate amplitude $\hat{\lambda}$, we assume that a homogeneous portfolio single name survival probability over any given time step from t to $t + \Delta t$ is given by

$$p(\Delta t) \;=\; \mathrm{e}^{-\lambda(t)\Delta t} \,. \tag{24}$$

This stepwise single name survival probability is the generator of the single time step loss distribution as given by equation (19), i.e.

$$\Psi(\ell(t+\Delta t)) \;=\; \Gamma\left(\gamma\phi\,,\,\Gamma^{-1}(\gamma, \mathrm{e}^{-\lambda(t)\Delta t}) - \Gamma^{-1}(\gamma(1-\phi), 1 - \tfrac{(\ell(t+\Delta t)-\ell(t))_+}{1-R-\ell(t)})\right) \tag{25}$$

with $\ell(t)$ being the loss at time t, and $\ell(t + \Delta t)$ being the loss incurred in the portfolio at the later time $t + \Delta t$. We give examples for $\Psi(\cdot)$ in figures 6 and 7. The analysis shows that the probability function $\Psi(\cdot)$ remains zero for a positive offset $\Delta\ell_{\text{min}}$ given by

$$\Delta\ell_{\text{min}} \;=\; \left(1 - \Gamma(\gamma(1-\phi), \Gamma^{-1}(\gamma, \mathrm{e}^{-\lambda\Delta t}))\right) \cdot (1 - R - \ell) \,. \tag{26}$$

This means that, for finite γ and for $\ell < 1 - R$ (which merely means that the portfolio still contains undefaulted reference obligations), there is a minimum loss the portfolio is guaranteed to incur over any finite time step Δt. This is essentially a consequence of the large portfolio limit approximation. In the limit of $\gamma \to \infty$, the Gaussian copula model is recovered which does not imply that there must be a minimum loss over any finite time step. In practice, however, the minimum incremental loss is so small for any realistic parameters that it is of no practical concern.

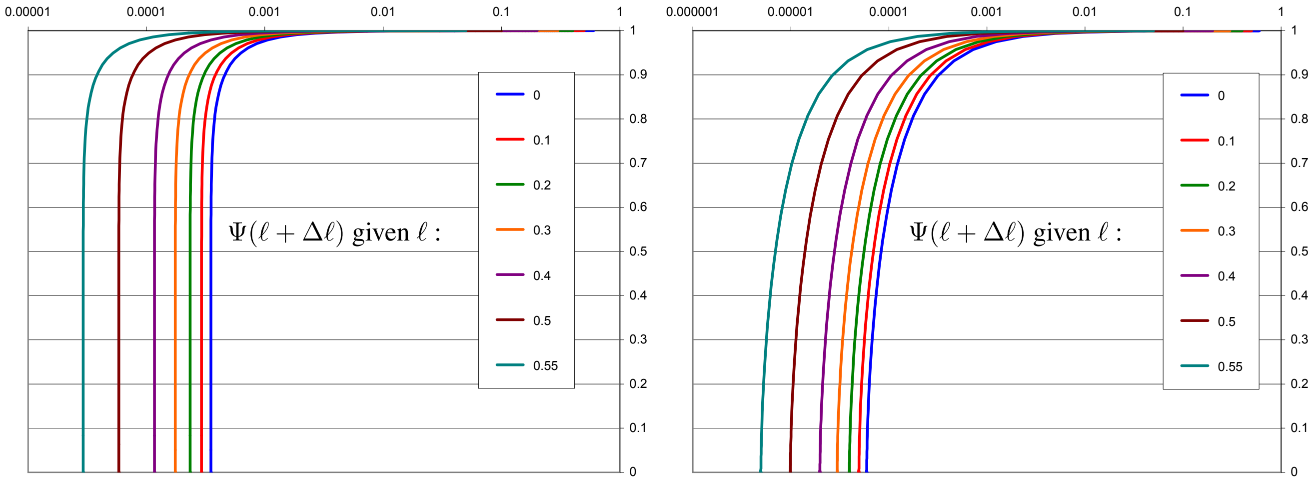

Figure 6: The cumulative loss transition probability (25) as a function of the incremental loss $\Delta\ell$ for $\gamma = 1$, $\lambda = 2\%$, $\Delta t = 1/26$, $R = 40\%$, and $\phi = 10\%$ (left figure) and $\phi = 60\%$ (right figure) and various initial loss levels ℓ as indicated in the legend.

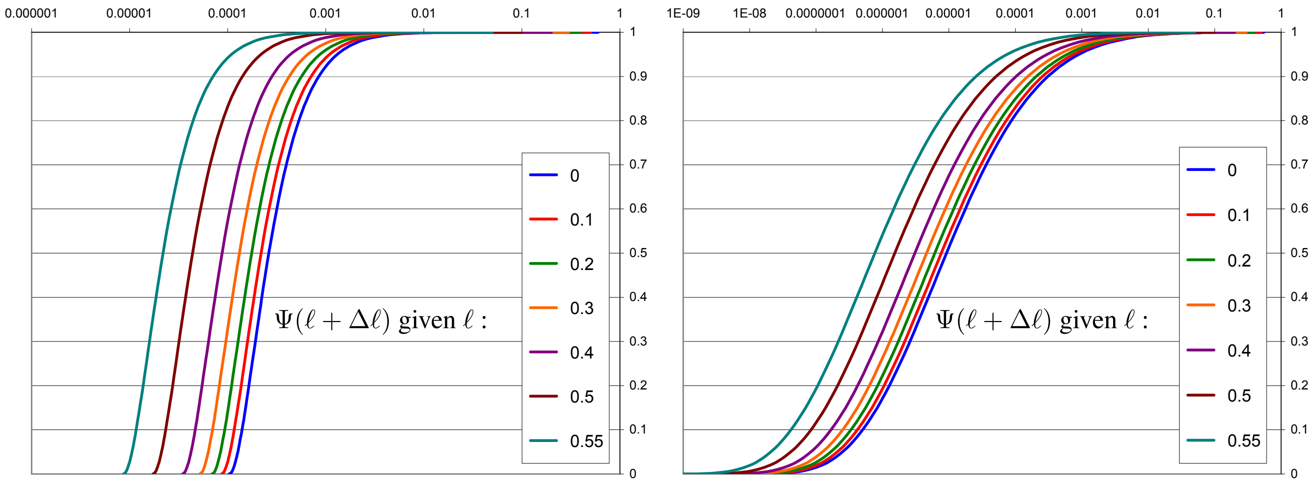

Figure 7: The cumulative loss transition probability (25) as a function of the incremental loss $\Delta\ell$ for $\gamma = 20$, $\lambda = 2\%$, $\Delta t = 1/26$, $R = 40\%$, and $\phi = 10\%$ (left figure) and $\phi = 60\%$ (right figure) and various initial loss levels ℓ as indicated in the legend.

Another interesting feature of the incremental loss law (25) is that there is a finite probability associated with the event that all names that are left in the portfolio default over the time step Δt given by

$$1 - \Psi(1-R) \; = \; 1 - \Gamma\left(\gamma\phi\,,\,\Gamma^{-1}(\gamma, \mathrm{e}^{-\lambda\Delta t})\right) \; . \tag{27}$$

Interestingly, this probability of total portfolio default does *not* depend on the current loss level, nor on the recovery rate. The fact that this model provides a total portfolio default risk with small associated risk-neutral probability is a *desirable feature*[2]. This becomes particularly important when super-senior tranches are to be considered given that market observable prices for super-senior risk tend to imply that there is indeed value associated with the writing down of the full underlying basket. We show a graph of the portfolio catastrophy probability in figure 8. We recognize in the figure the behaviour for the two limiting cases of $\gamma \to 0$, when the total portfolio default probability becomes a linear function of ϕ, and $\gamma \to \infty$ when, for all $\phi < 1$, the probability of total portfolio default vanishes. We can also see that the probability of total default is independent from γ in the case of $\phi = 1$ as one would expect it to be. When all names are one, the probability of total default is $1 - \mathrm{e}^{-\lambda\Delta t}$ as follows from (27).

The incremental loss density $\psi(\ell + \Delta\ell)$ given a current loss level ℓ can be derived by differentiation

[2] This phenomenon has also been observed in other models when calibrated to market observable tranche quotes [BPT07, HW05].

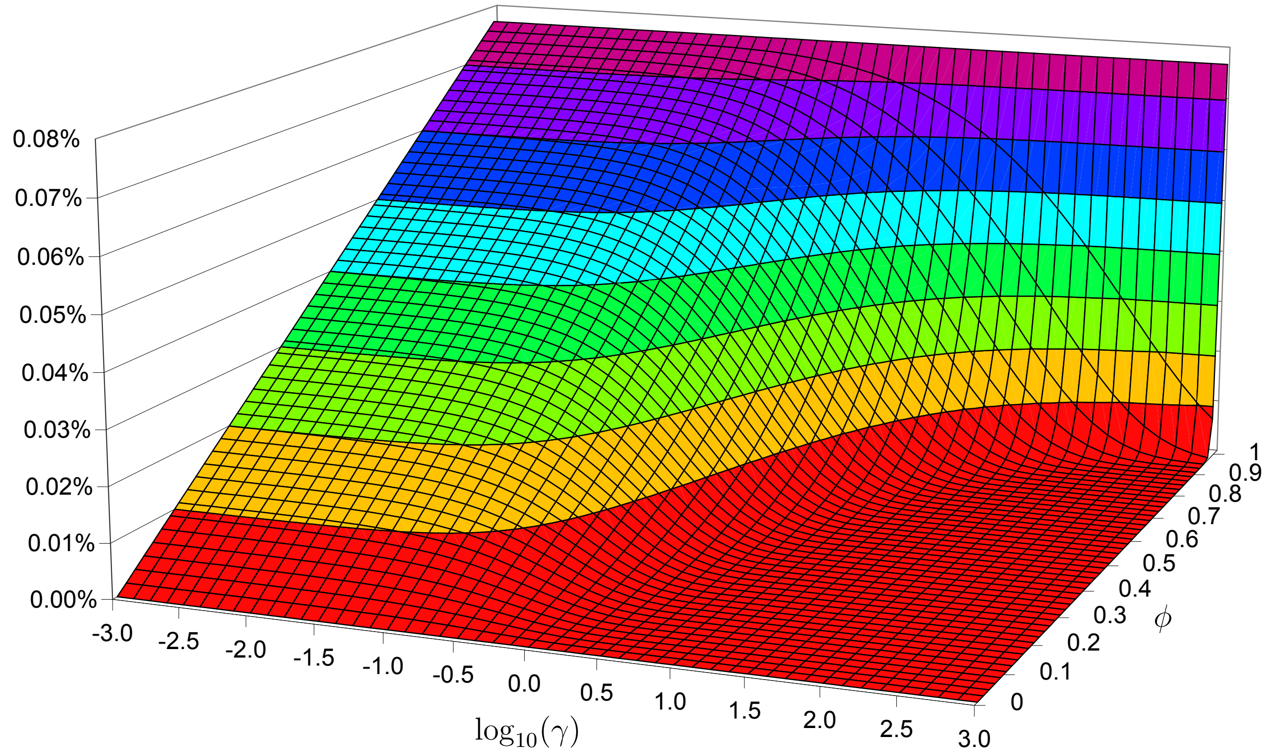

Figure 8: The total portfolio default probability (27) as a function of the shape parameter γ and the codependence parameter ϕ for $\lambda = 2\%$ and $\Delta t = 1/26$.

of (25) and is

$$\psi(\ell + \Delta\ell) \;=\; \mathbf{1}_{\{\Delta\ell > \Delta\ell_{\min}\}} \cdot \frac{g\big(\gamma\phi\,,\,\Gamma^{-1}(\gamma, \mathrm{e}^{-\lambda\Delta t}) - \Gamma^{-1}(\gamma(1-\phi), 1 - \frac{\Delta\ell}{1-R-\ell})\big)}{(1-R-\ell)\cdot g\big(\gamma(1-\phi)\,,\,\Gamma^{-1}(\gamma(1-\phi), 1 - \frac{\Delta\ell}{1-R-\ell})\big)} \tag{28}$$

with $g(\cdot,\cdot)$ as defined in (3). The densities for the cumulative probabilities shown in figures 6 and 7 are displayed in figures 9 and 10. What is not visible in these figures is that the incremental loss density has

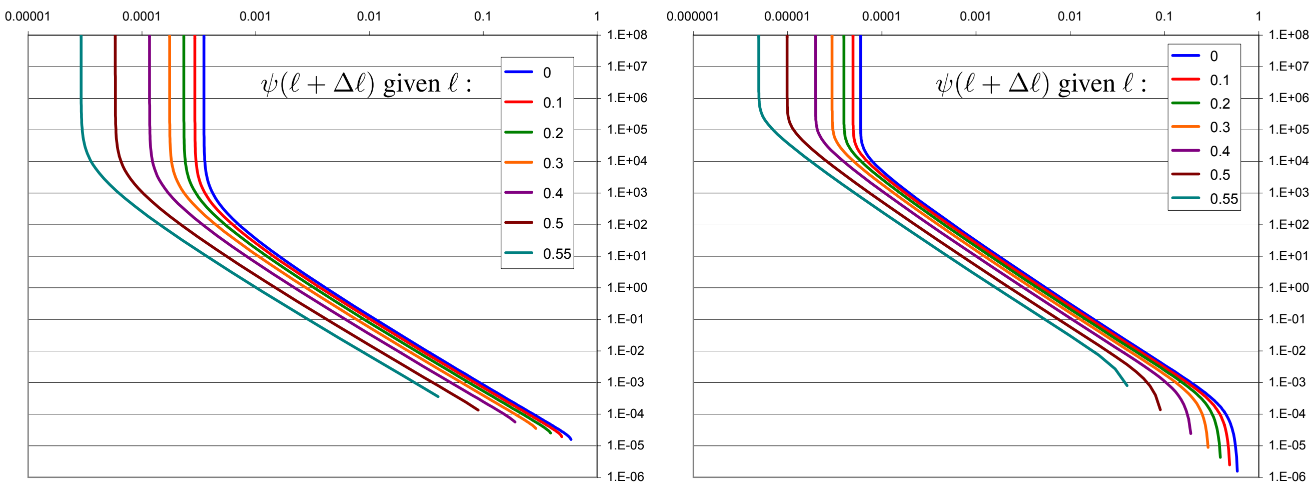

Figure 9: Incremental loss densities as defined in (28) as a function of the incremental loss $\Delta\ell$ for $\gamma = 1$, $\lambda = 2\%$, $\Delta t = 1/26$, $R = 40\%$, and $\phi = 10\%$ (left figure) and $\phi = 60\%$ (right figure) and various initial loss levels ℓ as indicated in the legend.

an additional Dirac function component at $\Delta\ell_{\max} = 1 - R - \ell$ with its weight given by equation (27) representing the total portfolio default probability.

Within the model's discrete dynamics' framework, loss distributions are the result of discrete iteration of the incremental loss transition probability matrix. This is one of the main differences between this model and the original Gamma-Phi model: there is no single random variate (constant quantile or

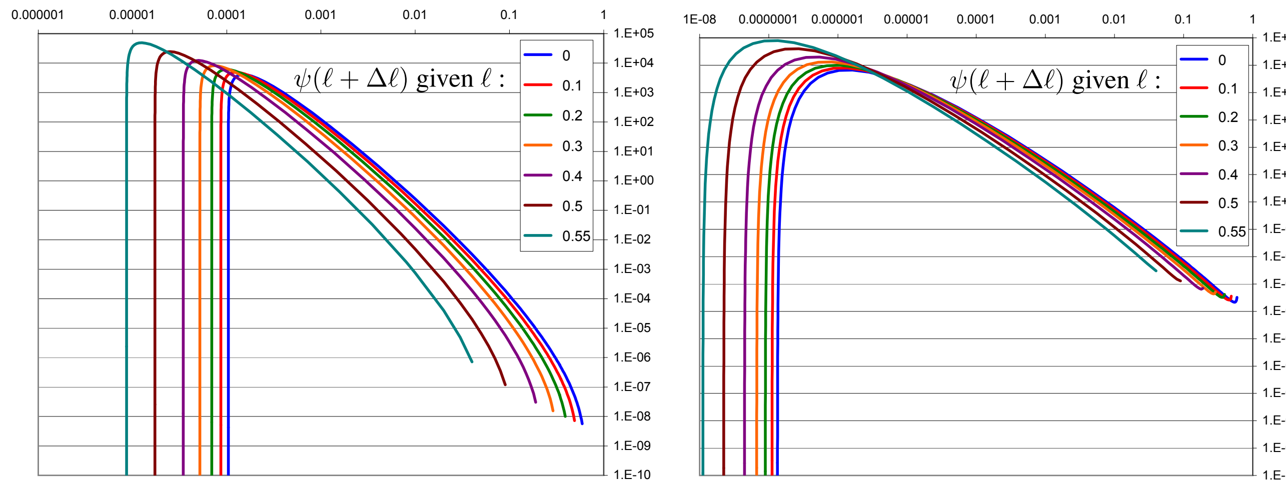

Figure 10: Incremental loss densities as defined in (28) as a function of the incremental loss $\Delta\ell$ for $\gamma = 20$, $\lambda = 2\%$, $\Delta t = 1/26$, $R = 40\%$, and $\phi = 10\%$ (left figure) and $\phi = 60\%$ (right figure) and various initial loss levels ℓ as indicated in the legend.

constant value) approximation intrinsic in the discrete dynamics. We merely use the single random variate approximation of Baxter's Gamma-Phi model to generate the Δt loss transition probabilities. This does, of course, imply that, essentially, the selected time step Δt becomes an intrinsic model parameter. We found, though, that, for the pricing of actual derivatives when calibrating to relevant securities, the sensitivity to changes in Δt is very moderate indeed and thus of no particular relevance. We show the discretised loss transition probabilities generated by the Discrete Gamma Pool model in figure 11 when the single name default hazard rate λ is constant at 2%.

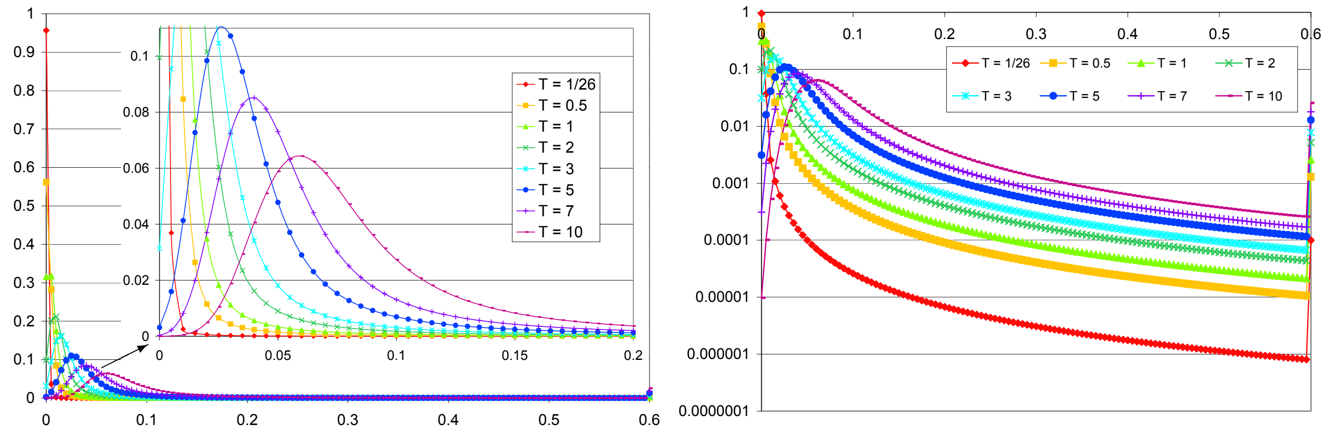

Figure 11: The discretised loss transition probabilities as generated by the Discrete Gamma Pool model for $\sigma = 0$ and $\gamma = 5/2$, $\lambda = 2\%$, $\Delta t = 1/26$, $R = 40\%$, and $\phi = 60\%$ on a linear scale (left, with inset enlargements) and on a logarithmic scale (right). The loss dimension was discretised in steps of 0.5%. Note that the actual data of relevance are only the marked points since the transition probabilities are discrete within the model's lattice implementation. The lines are merely in aid of guiding the eye.

The evolution of the hazard rate $\lambda(t)$ is given by a discrete approximation to equations (21) and (22). The temporal discretisation of these equations is to be seen as a standard Euler approximation to the stochastic differential equations, or, in terms of finite differencing implementations of the associated partial differential equations, as a standard first order explicit scheme. By virtue of the stochasticity of the underlying homogeneous portfolio single name hazard rate, the par spread on a credit default swap on the whole portfolio, or on any tranches, is also stochastic.

In summary, the Discrete Gamma Pool model is put together from a time discretised version of the large pool limit of Baxter's Gamma-Phi model and a time discretised version of a Black-Karasinski model for the hazard rate determining the single name survival probability over any given time step. In this sense, we view the mathematical problem of pricing contingent claims under the dynamics of the model as a discrete version of a partial integro differential equation of the form

$$\partial_t V - \kappa \cdot y \cdot \partial_y V + \tfrac{\sigma^2}{2} \cdot \partial_{yy} V + \mathsf{E}\left[V(t, y, \ell + \Delta\ell) + \mathcal{P}(t, \ell, \ell + \Delta\ell) \,\middle|\, \ell \right] - V \;=\; r \cdot V \tag{29}$$

with $V = V(t, y, \ell)$ being the value of the contingent claim at hand, r being the continuously compounded interest rate, $\mathcal{P}(t, \ell, \ell + \Delta\ell)$ standing for cashflows triggered by losses of size $\Delta\ell$ conditional on the current loss level in the overall portfolio being ℓ, and the expectation $\mathsf{E}[\cdot]$ being over the (unknown) Lévy measure associated with the dynamics of the loss ℓ. In our actual discrete time setting, the transition probabilities generated by that measure are given by equation (25). Their dependence on the spread process driver y arises from the fact that these transition probabilities depend on $\mathrm{e}^{-\lambda\Delta t}$, and that λ depends on y through (21).

Since we are keeping interest rates deterministic, it is straightforward to switch to the discounted value function

$$\tilde{V}(t, y, \ell) \;:=\; \mathrm{e}^{-\int_0^t r(t')\mathrm{d}t'} \cdot V(t, y, \ell) \tag{30}$$

which enables us to write the pricing equation without the right-hand side source term:

$$\partial_t \tilde{V} - \kappa \cdot y \cdot \partial_y \tilde{V} + \tfrac{\sigma^2}{2} \cdot \partial_{yy} \tilde{V} + \mathsf{E}\left[\tilde{V}(t, y, \ell + \Delta\ell) + \tilde{\mathcal{P}}(t, \ell, \ell + \Delta\ell) \,\middle|\, \ell \right] - \tilde{V} \;=\; 0\,. \tag{31}$$

Clearly, any cashflows that are entering the pricing equation are to be discounted by division through the money market account in complete analogy to the contingent claim value $\tilde{V}$:

$$\tilde{\mathcal{P}}(t, \ell, \ell + \Delta\ell) \;=\; \mathrm{e}^{-\int_0^t r(t')\mathrm{d}t'} \cdot \mathcal{P}(t, \ell, \ell + \Delta\ell)\,. \tag{32}$$

For further details as to the numerical analysis involved in the implementation of a finite differencing (lattice) algorithm for the Discrete Gamma Pool model, we refer the reader to [Jäc07].

4 Loss Triggered Leveraged Super-Senior notes

A conventional synthetic collateralised debt obligation is a trade in which a counterparty deposits a certain amount of collateral, and with that collateral, underwrites protection on a tranche with attachment point a and detachment point d of the total notional of a pool of reference obligations. A default of any of the reference obligations may give rise to losses in the tranche, and thus consumption of the counterparty's collateral, if, after adjustment for recovery on the defaulted reference obligations, the total loss in the portfolio exceeds the attachment point of the underwritten tranche. The proportion of collateral the counterparty forfeits upon an incremental loss in the portfolio of size $\Delta\ell$ is given by

$$\frac{(\min(d, \ell + \Delta\ell) - \max(\ell, a))_+}{d - a} \tag{33}$$

with ℓ representing the loss level in the portfolio prior to the default event giving rise to the incremental loss $\Delta\ell$. All cashflows that are made in this fashion in response to defaults in the portfolio constitute what is commonly referred to as the *protection leg* or the *benefit leg*. In return for the underwriting of the risk on the protection leg, the counterparty receives fees on the *premium leg*. These are, in the absence of default events, simply regular coupons at a previously agreed fixed rate. An exception to this rule is often made for *equity tranches*, i.e. tranches whose attachment point is zero, where, typically, part of the fee is paid upfront in a lump sum payment. For standard contracts of *amortizing type*, the payments are scaled down as losses erode the notional in the portfolio according to

$$\frac{(\min(d, 1 - \tilde{R}) - \max(\ell, a))_+}{d - a} \tag{34}$$

where $\tilde{R}$ represents the proportion of the original portfolio notional that has been recovered following the default events that led to the current loss level being ℓ. When all defaults recover at the same rate R, we have

$$\tilde{R} \;=\; \ell \cdot \frac{R}{1-R} \,. \tag{35}$$

The term $\min(d, 1-\tilde{R})$ in (34) can be seen as an *adjusted detachment point* whose difference from the contractual detachment point typically only matters for super-senior tranches, whence it is often referred to as a *super-senior correction*. We show in figure 12 the effect of the super-senior correction on the effective premium leg notional as a function of incurred losses in the portfolio.

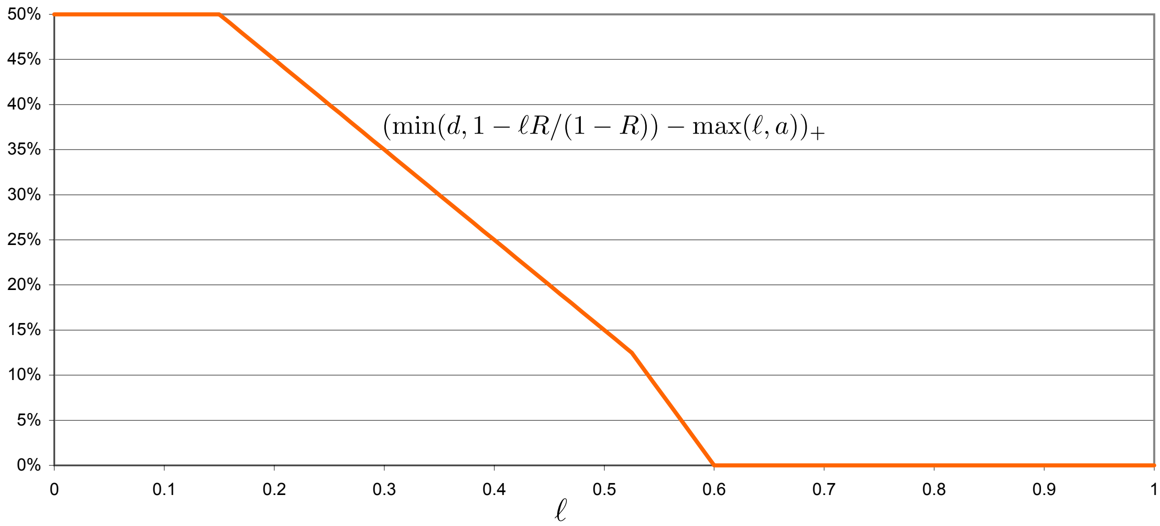

Figure 12: The effective premium leg notional as a function of incurred loss assuming a deterministic recovery rate $R = 40\%$ for a super-senior tranche attaching at $a = 15\%$ and detaching at $d = 65\%$. All figures are given as a proportion of the portfolio notional. The effect of the super-senior correction can be seen for $\ell > (1-d)(1-R)/R = 52.5\%$.

In addition to regular coupon payments that are adjusted for the erosion of notional as a function of incurred losses, standard contracts typically pay accrued interest on default. These payments are made at the next regular coupon date, i.e. are deferred until the next canonical payment time. Within the valuation framework (31), the accrued interest fits in as a type of loss induced payment $\tilde{\mathcal{P}}(t, \ell, \ell + \Delta\ell)$.

Based on the above explanations of payments on the protection and the premium leg, we can state in formal terms what the loss induced payment expression $\tilde{\mathcal{P}}(t, \ell, \ell + \Delta\ell)$ in equation (31) for those legs mean, respectively:

- **Protection Leg:**

$$\tilde{\mathcal{P}}_{\text{protection}}(t, \ell, \ell + \Delta\ell) \;=\; (\min(d, \ell + \Delta\ell) - \max(\ell, a))_+ \cdot \mathrm{e}^{-\int_0^t r(t')\mathrm{d}t'} \tag{36}$$

  The above expression is the amount that is deducted from the counterparty's collateral, which is equivalent to a payment the protection seller must make to compensate for the loss in the tranche.

- **Premium Leg:**

$$\tilde{\mathcal{P}}_{\text{premium}}(t, \ell, \ell + \Delta\ell) \;=\; s \cdot \tau(t) \cdot (\nu(a, d, R, \ell) - \nu(a, d, R, \ell + \Delta\ell)) \cdot \mathrm{e}^{-\int_0^{\hat{T}(t)} r(t')\mathrm{d}t'} \tag{37}$$

  The term $\hat{T}(t)$ is defined as the next canonical coupon payment time on or after t and the nominal fixed spread paid on the premium leg is given by s. The term $\tau(t)$ stands for the accrual factor from the last regular coupon payment time to the time of default. The premium leg effective payment notional function $\nu(a, d, R, \ell)$ is defined as

$$\nu(a, d, R, \ell) \;=\; (\min(d, 1 - \ell R/(1-R)) - \max(\ell, a))_+ \tag{38}$$

  and an example for this function was shown in figure 12.

Note that in both cases the payments are expressed in terms of proportion of notional associated with the entire portfolio, i.e. we assume that, for a standard contract, the collateral posted against the (a, d) tranche is $(d - a)$ times the portfolio notional. In addition to the loss induced payments $\tilde{\mathcal{P}}$, unlike the protection leg, the premium leg also incurs regular coupons whose discounted amount depends on the current loss level and is given by

$$s \cdot \tau \cdot \nu(a, d, R, \ell) \tag{39}$$

times the portfolio notional, assuming that the accrual factor for the full period is τ.

A *leveraged super-senior tranche* is a contract very similar to the standard tranche we described above. The difference is that, instead of posting collateral for the full tranche notional given by $(d - a)$ times the portfolio notional, the counterparty only provides a proportion of $1/m$ of the nominal tranche notional, for some leveraging factor $m > 1$. This results in an asymmetric situation between the premium and the prottection leg: whilst the counterparty receives s on the full super-senior tranche, they are only exposed to losses up to $(d - a)/m$. We show this in figure 13. Given the sub-par protection provided

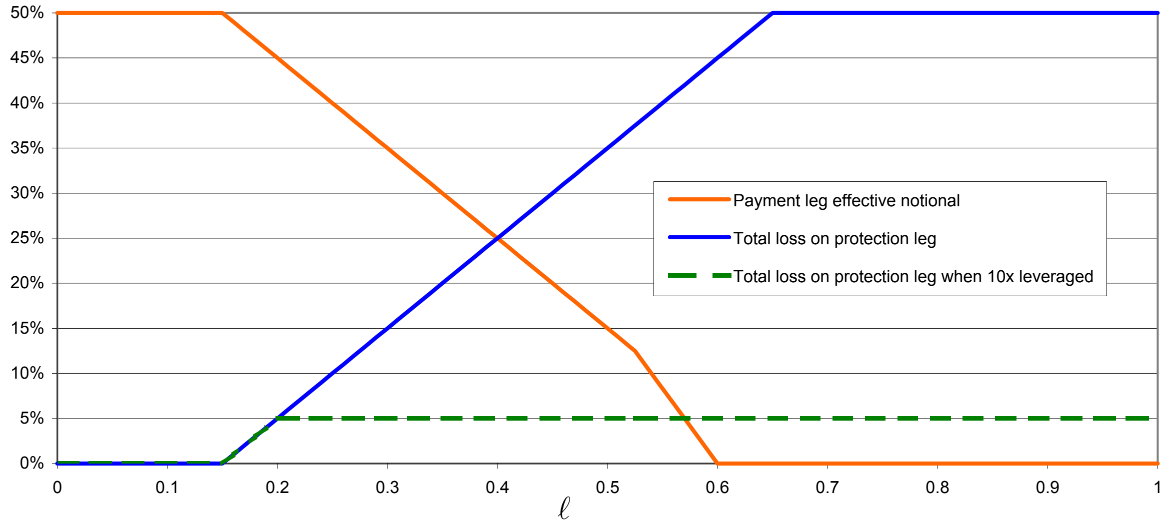

Figure 13: The long dashed line shows the total loss a $10\times$ ($m = 10$) leveraged super-senior tranche attaching at $a = 15\%$ and detaching at $d = 65\%$ suffers as a function of incurred loss in the portfolio. The solid blue line is the loss the same tranche incurs when fully collateralised ($m = 1$).

by a counterparty that provides only $m\times$ down-leveraged collateral, the spread s on the premium leg is naturally smaller than it would be on a fully collateralised deal. In addition, the buyer of protection typically demands further risk-mitigating features to be added to the deal. These features most commonly entail either or both of:-

- **Spread triggered unwind clauses.** A clause of this type usually states that, as and when the underlying standard super-senior tranche's market observable par spread rises to or beyond a contractually agreed level, the deal is unwound. This is done based on the assumption that the protection buyer is holding a back-to-back fully collateralised deal with a third party which is unwound at the observed market par spread. Any losses incurred in the unwinding of the underlying standard super-senior tranche are deducted from the counterparty's collateral, whose remainder, if positive, is then returned to the counterparty. Should, against expectations, the unwinding result in some positive net proceeds, this is typically *not* passed on to the counterparty. If we denote the standard super-senior tranche's premium leg net present value for unit (100%) spread as A (for *annuity*) and the net present value of the protection leg as B (for *benefit*), the financial position of the counterparty at the time of unwinding is

$$\min(0, (s \cdot A - B)) \quad . \tag{40}$$

  A requirement for a spread-triggered unwind clause is that there is some form of *market-observable* par spread for the underlying super-senior tranche. In practice, this tends to limit the choice of underlying portfolio considerably. However, since the net present value of the underlying super-senior tranche is very strongly linked to its par spread, this type of risk-mitigation clause tends to leave very little risk for the protection buyer whence it is not uncommon that a standard tranche pricing model is deemed sufficient for the hedging and risk management of Spread Triggered Leveraged Super-Senior notes.

- **Loss triggered unwind clauses.** In this case, an unwind clause exactly as in the spread triggered version is added to the contract, with the difference that not the par spread on the underlying tranche, but *losses incurred in the underlying portfolio* give rise to early termination of the deal, as and when they breach a contractually agreed level. Since losses are much more readily observable than spreads in an objective fashion for bespoke portfolios, this variation opens up the full universe of reference obligations for standard tranches to become part of a portfolio for a Loss Triggered Leveraged Super-Senior contract. The downside of this type of deal is that the unwind cost is largely unknown in advance since spreads on tranches tend not to move in any predictable way in line with losses, or vice versa. This is the reason why a comparatively sophisticated model is needed in order to trade, hedge, and risk manage this type of contract.

A further layer of complexity for LTLSS contracts is that, if agreed at inception, at the time of a loss trigger event, the counterparty may have the *option to deleverage* instead of agreeing to an unwinding of the trade. This means that, when the trigger level is breached, the counterparty may post the full collateral against the tranche. This may be of interest to the counterparty particularly, but not only, in situations when it is expected that the unwinding of the underlying standard super-senior tranche generates positive proceeds which are, as part of the LTLSS contract, not passed on to the counterparty. If the net position of the underlying tranche, i.e. $s \cdot A - B$, is sufficiently positive, the counterparty can attempt to lock in the positive value by posting the full collateral, and entering into a back-to-back trade on the same underlying super-senior tranche with a third party.

To make matters even more complicated, a loss triggered leveraged super-senior trade with embedded deleveraging option may have not just one, but several specified loss trigger levels, with *partial deleveraging options*. As each level is breached, the counterparty has the option to deleverage to the next level by posting more collateral up to a previously agreed next higher amount, or to unwind. And finally, there is the possibility for the layered trigger levels to vary with time in order to cater for the fact that the relative probabilities of touching one and the same loss trigger level, and the relative probabilities with potentially incurred subsequent unwind costs depend not only on the trigger level, but also on the residual time to maturity in the underlying super-senior tranche. We show an example of such a trigger level term structure with three tiers of leverage, namely 12, then 6, then 3, in figure 14. In this example, the counterparty holds, effectively, three subsequent options to deleverage upon loss trigger events. At any time, though, if the deal is unwound, the counterparty's maximum exposure to losses is always limited to the amount of collateral that is posted at that point in time. This maximum downside limit to the amount of collateral posted also applies to the extremely unlikely case that in a sudden event all or part of the underlying super-senior tranche notional is wiped out. The trigger features embedded in the contract are, essentially, designed to protect the protection buyer in an LTLSS deal from these highly unlikely gap events.

4.1 Analysis of the LTLSS with the Discrete Gamma Model

The dynamics of the Discrete Gamma Pool model, by their nature of allowing for losses as well as spread moves to occur separately, make it suitable for the valuation and hedging of Loss Triggered Leveraged Super-Senior notes. Without the option to deleverage, both Monte Carlo simulations and finite differencing (lattice) implementations can be used for the valuation. When the contract does allow for deleveraging, however, the product bears similarity to a complex variation of a Bermudan swaption: upon the event of a loss trigger, and only then, the counterparty holds the right to decide between two

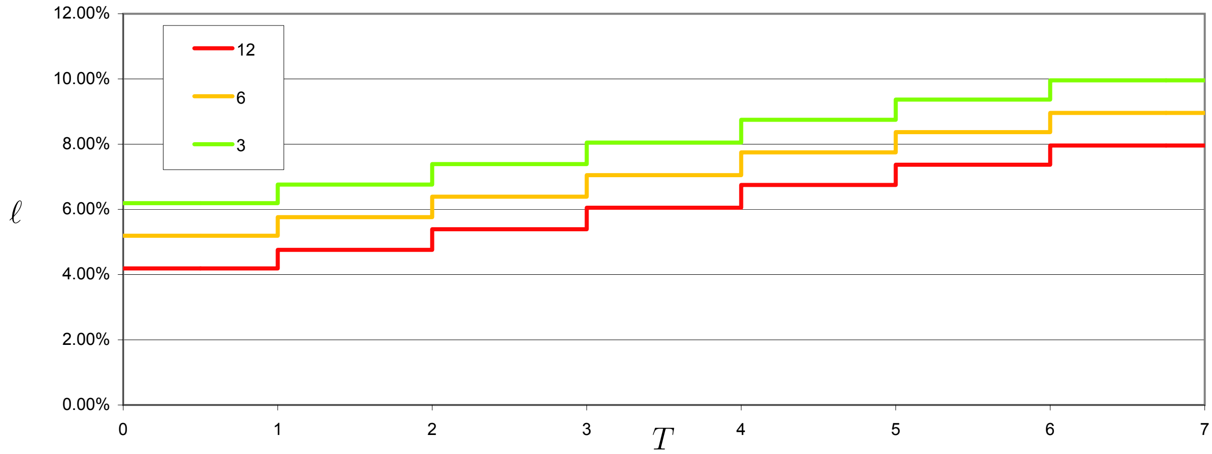

Figure 14: Sample loss trigger levels for a three-tiered Loss Triggered Leveraged Super-Senior tranche. Below the lowest line denoted as "12", the counterparty must keep 1/12 of the tranche notional's worth of collateral posted. When this line is breached by losses in the portfolio, the counterparty has to either agree to double their collateral, and thus deleverage to a leverage level of 6, or the deal is unwound. When the line denoted as "6" is breached, the counterparty can deleverage down to a tranche notional to collateral ratio of 3, else the deal is terminated at the potential expense of the counterparty. Finally, as the line denoted as "3" is exceeded by portfolio losses, the counterparty must post the full tranche notional's worth of collateral, or the contract terminates, again with unwind costs at the expense of the posted collateral.

possible outcomes there and then. This alone makes it akin to some form of conditional option. However, since LTLSS contracts typically have more than one leveraging layer, a similar event may arise again when one of the next higher trigger levels is exceeded. The valuation of such a contract fits well into the pricing framework of multi-excercisable options such as *flexi-caps* that is fairly well known in the area of interest derivatives pricing. A flexi-cap is a contract that permits the holder to exercise k out of l subsequent caplet contracts with $k < l$. The pricing of such deals is readily achieved by the aid of a lattice based numerical implementation, and suitably selected *transition rules*. At the heart of the valuation is the concept that the pricing must necessarily entail the valuation of not one, but *several* contracts in parallel. For a flexi-cap, the number of parallel calculations needed is k: at the lowest layer, we price the option that has only one exercise right left. At each point in time that is observed along the discretised time line, a choice between the backward induction value, and the proceeds from exercising there and then, is made. At the next layer, it is assumed that two rights to exercise are left, and at each point in the backwards induction one compares the value arrived at from the backwards evolution of the partial (integro) differential equation of the contract value with the sum of the proceeds from exercising the currently fixing caplet *plus* the value of the next lower layer. Subsequent layers follow in complete analogy, until we arrive at layer k, which represents the value of the k-out-of-l flexi-cap. This methodology of effectively extending the state space in an additional direction can also be applied to the lattice based valuation of Asian (American or European), Hindsight, Lookback, and many other types of exotic options [Wil98].

In the context of loss triggered leveraged super-senior notes, we start with the definition of a sequence of auxiliary pricing legs. Denote $V_0(t, y, \ell)$ as the value of the underlying super-senior tranche from the counterparty's point of view, i.e.

$$V_0(t, y, \ell) \;=\; s \cdot A(t, y, \ell) - B(t, y, \ell) \tag{41}$$

with s being the contractually agreed spread, $A(t, y, \ell)$ being the risky annuity which is defined as the value of the premium leg assuming unit (100%) spread, and $B(t, y, \ell)$ being the net present value of the benefit leg. Further, assume that the leverage levels, from lowest to highest, are given by the numbers m_l with $m_0 = 1$ and $m_1 < m_2 < m_3 < \ldots$, etc. For instance, in our example in figure 14 we would have $m_1 = 3$, $m_2 = 6$ and $m_3 = 12$. Also, denote the loss trigger levels bordering the respective domains as θ_l, with $\theta_1 > \theta_2 > \theta_3 > \ldots$ and so on. The auxiliary valuation leg $V_l(t, y, \ell)$ for $l > 0$ is defined as the leveraged super-senior deal that has all of the inner leverage layers up until and including $\#l$, but not beyond. The value $V_2(t, y, \ell)$, for instance, is a deal that starts with initial collateral $(d - a)/m_2$

assuming that there have been no losses in the portfolio yet, or if so (e.g. if this was a replacement trade), losses do not exceed the defined loss level θ_2. The transition rules that connect the different layers of the product and enable valuation of the outermost one, the actual loss triggered leveraged super-senior trade we really wish to value, are given by the condition

$$V_l(t,y,\ell) \;=\; \max(V_{l-1}(t,y,\ell), \min(0, \max(V_0(t,y,\ell), -c_l(\ell)))) \quad \forall \quad \ell \geq \theta_l(t) \tag{42}$$

for any layer $l > 0$, and for any t, y, and ℓ. The auxiliary *remaining collateral* function $c_l(\ell)$ in (42) is defined as

$$c_l(\ell) \;=\; \min((d-a)/m_l, \max(a + (d-a)/m_l - \ell, 0)) \tag{43}$$

In practical terms, this means that any lattice based numerical implementation must, upon backwards induction from a later time horizon $t + \Delta t$ to time t (during which loss induced payments $\mathcal{P}$ must be included), override the induced expectations for valuation leg $V_l(t,y,\ell)$ on all (y,ℓ)-discretisation nodes for which $\ell \geq \theta_l(t)$.

The transition rule (42) can be readily understood if we unroll it from the inside out, and compare with the contract definitions. The innermost term $\max(V_0(t,y,\ell), -c_l(\ell))$ is nothing other than a flooring of the unwind costs, given by the net present value of the unleveraged super-senior contract $V_0(t,y,\ell)$, at minus the current collateral value. The $\min(0,\cdot)$ expression around this part represents the fact that the counterparty will never receive any potential positive proceeds from the unwinding of the underlying unleveraged super-senior contract. The final choice $\max(\cdot,\cdot)$ between this value, and the value of the next layer contract $V_{l-1}(t,y,\ell)$ is simply the choice between unwinding and posting more collateral: upon handing over more securities as collateral, the deal changes nature into one with lower leverage and higher exposure of the counterparty to losses exactly as represented by the next lower valuation leg $V_{l-1}(t,y,\ell)$.

5 Multi-callable CDOs and Bermudan CDO options

A multi-callable CDO tranche is a contract in which one counterparty has the right to terminate all future payments on a set of exercise dates. When the counterparty holding the prerogative of early termination also holds the loss protection, and pays the premium, this is referred to as a multi-callable *payer's* CDO tranche. Likewise, if the holder of the exercise rights is receiving the coupons, we call it a *receiver's* CDO tranche. The net position of a holder of a callable payer's CDO tranche can also be viewed as having entered into a payer's CDO tranche plus the holding of a Bermudan option to enter into a receiver's CDO tranche, i.e. a Bermudan option on a receiver's CDO. When there is only a single exercise opportunity, by virtue of put-call parity, the net position is equal to just an option to enter into a payer's CDO tranche on the exercise date. The underlying of this option can be called a *forward starting CDO*. A forward starting CDO is to be seen as the contract to take over all further payments in a running CDO previously held by a different agent at some point in the future, assuming that at the forward start date all currently due payments (past coupons as well as compensation for past losses in the tranche) have been settled. A forward starting CDO may therefore, if defaults occur prior to the forward start date, by the time it becomes active, have effective attachment, detachment, and live portfolio notional figures that are not what was contractually defined at the very beginning when the forward starting deal was entered into[3]. Note that the aforementioned put-call parity only holds for callable CDOs that have a single exercise date. For multi-callable CDO tranches, a callable payer's tranche is *not* equal to a Bermudan option on a forward starting payer's CDO. It is, though, equal to a payer's CDO tranche plus a Bermudan option on a receiver's CDO tranche.

Valuation of a Bermudan option on a CDO tranche on a finite differencing lattice can be done in complete analogy to the discussion of the loss triggered leveraged super-senior note. Assuming again

[3]This definition of a forward starting CDO differs from an alternative, more complex synthetic version, where any defaults prior to the forward start date are contractually defined to be treated as if they recover at 100%.

that $A(t, y, \ell)$ is given as the the risky annuity leg which is defined as the value of the premium leg assuming unit (100%) spread, and $B(t, y, \ell)$ is the net present value of the benefit leg on any given filtration node on the lattice, the Bermudan option leg V_{option}, following the rolling back in the finite differencing step, undergoes the transition rule

$$V_{\text{option}}(t, y, \ell) \leftarrow \max\left(V_{\text{option}}(t, y, \ell) \,,\, \omega \cdot (B(t, y, \ell) - s \cdot A(t, y, \ell)) \right) \tag{44}$$

on any of the exercise dates, with $\omega = +1$ for options on payer's tranches, and $\omega = -1$ options on receiver's tranches, both struck at s, respectively. Naturally, at the initialisation of the option leg at the end of the time line discretisation used in the backwards induction, $V_{\text{option}}(t, y, \ell)$ is set to zero for all (t, y, ℓ).

6 Calibration

The Discrete Gamma Pool model provides a number of parameters that can be used for the calibration of the model to the most relevant hedge quantities, all of which, with the exception of R, can be made a function of time in order to bootstrap to multiple maturities. The hazard rate amplitude $\hat{\lambda}$ can be used to calibrate the model to match the value of protection on the entire portfolio. When implemented as a finite differencing lattice, the model is very stable whence, based on the initial guess of

$$\hat{\lambda} \approx \frac{s_{0\%-100\%}}{1-R} \tag{45}$$

and the analytical knowledge that the value of portfolio protection is guaranteed to be zero when $\hat{\lambda} = 0$, three to four iterations of any suitable root finding algorithm typically result in a highly accurate (discrepancy below 0.0001% of portfolio protection value) calibration to the portfolio protection leg.

Given any choice of γ, and assuming that we always calibrate to the portfolio protection leg value as discussed above, the model can also be calibrated to the protection leg of one specifically given tranche by iterating over ϕ. We found that typically only four to five iterations over ϕ are needed for highly accurate calibration, starting with the two initial guesses given by the square root of the Gaussian copula model's correlation number plus/minus 5% (suitably capped or floored if necessary). During this procedure, the number of iterations in the sub-nested root-finding for $\hat{\lambda}$ can be reduced to only two as long as we take as the initial guess always the last value from the previous outer iteration, and we still arrive at high calibration accuracy to both the portfolio protection leg and the super-senior tranche protection leg value.

The above procedure does, of course, not simultaneously also calibrate to the entire portfolio's par spread, nor the super-senior tranche's. The former can usually be calibrated, if desired, by the small change in the definition of the stochastic hazard rate from (21) to

$$\lambda(t) = \hat{\lambda}(t) \cdot e^{-\frac{1}{2}\mathsf{V}_0[y(t)]+\mu t+y(t)} \tag{46}$$

and fitting μ until the portfolio par spread is matched in addition to the portfolio protection leg as well as the tranche's protection leg. This change to the definition of λ does, in fact, allow the model to reflect a feature typically incurred by a Gaussian copula model: market-observable CDS quotes are usually upward sloping with maturity, resulting in a very steep increase for deterministic instantaneous forward default rates. Since a term structure of hazard rates affects the premium and the protection leg in a slightly different way, a change in the slope changes the par spread when one keeps the protection leg value constant. Having said that, it may or may not be desirable to calibrate the portfolio's par spread in addition to the portfolio protection leg, depending on whether protection is bought or sold on the underlying super-senior tranche. Exactly the same consideration applies to the calibration of the specific tranche's protection leg. Alternatively, ϕ could be calibrated such that the super-senior tranche's par spread is matched to given data.

As for the parameter γ, it can be used for calibration of the overall loss distribution. This last degree of freedom can be used to match whichever additional quantity is seen as most suitable for the specific trade at hand. Examples include the protection value or spread of an equity tranche detaching at the first loss trigger level, mezzanine tranches spanning the loss trigger levels, etc. In order to permit calibration to multiple maturities, all of λ, ϕ, and γ can also be functions of time.

6.1 Calibration to iTraxx Series 8

As an example for the use of piecewise constant term structures for $\lambda(t)$, $\gamma(t)$, and $\phi(t)$, we have calibrated the discrete gamma pool model to a set of sample quotes for iTraxx Europe Series 8 tranches shown in figure 15. Using the instantaneous volatility $\sigma = 70\%$, the mean reversion strength $\kappa = 10\%$,

BBIT3 P308 n Curncy **BBIT**
Screen Printed
15:41 **iTRAXX Tranches** PAGE 1 / 2

Series 8 5Y	Bid	Offer	Ref	Time	BC	Delta
0-3% + 500 bp	1) 24.25	25.25	65	9:33	16) 42.09	31) 10.50
3-6%	2) 245.00	255.00	65	9:33	17) 54.17	32) 4.50
6-9%	3) 143.00	151.00	65	9:33	18) 61.07	33) 2.75
9-12%	4) 94.00	99.00	65	9:33	19) 66.47	34) 1.75
12-22%	5) 54.00	59.00	65	9:33	20) 79.40	35) 1.10
Series 8 7Y	**Bid**	**Offer**	**Ref**	**Time**	**BC**	**Delta**
0-3% + 500 bp	6) 32.00	33.00	72	9:33	21) 40.65	36) 8.00
3-6%	7) 338.00	348.00	72	9:33	22) 51.33	37) 4.75
6-9%	8) 189.00	196.00	72	9:33	23) 58.16	38) 3.00
9-12%	9) 123.50	129.50	72	9:33	24) 63.62	39) 2.10
12-22%	10) 74.00	78.50	72	9:33	25) 76.59	40) 1.30
Series 8 10Y	**Bid**	**Offer**	**Ref**	**Time**	**BC**	**Delta**
0-3% + 500 bp	11) 37.50	38.50	79	9:33	26) 41.20	41) 6.00
3-6%	12) 494.50	509.50	79	9:33	27) 46.53	42) 5.50
6-9%	13) 239.00	248.00	79	9:33	28) 53.57	43) 3.50
9-12%	14) 143.00	150.00	79	9:33	29) 60.12	44) 2.30
12-22%	15) 82.50	88.50	79	9:33	30) 75.62	45) 1.40

Figure 15: Tranche quotes for iTraxx Europe Series 8 from Bloomberg page BBIT3 on January 11$^{\text{th}}$ 2008.

the model recovery rate $R = 39.92\%$, the time step $\Delta t = 1/26$, and the loss discretisation size $\Delta\ell = (1 - R)/125 = 0.48064\%$, we calibrated the model to the mid quotes[4] for the 0%–3% and 0%–22% tranches, as well as to the portfolio spread which is 66.11bp for 5Y, 73.80bp for 7Y, and 81.11bp for 10Y. The resulting term structures for $\hat{\lambda}$, ϕ and γ are shown in figure 16. Note that the value for $\hat{\lambda}$

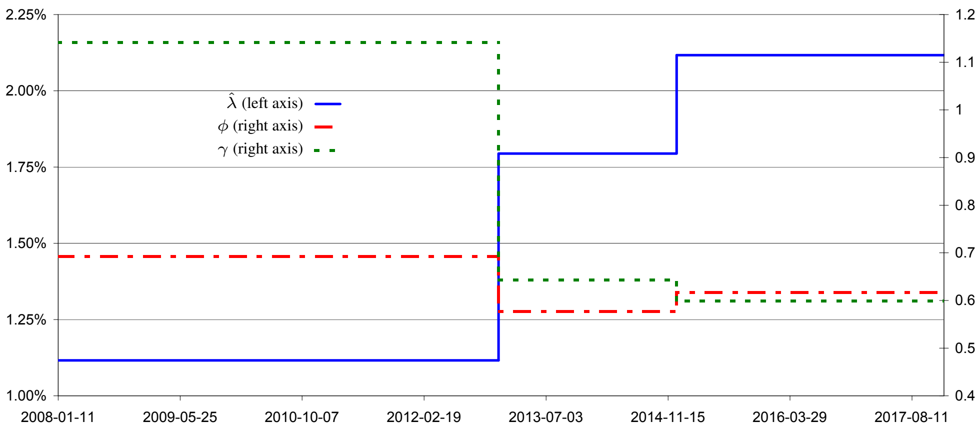

Figure 16: Parameter term structures for calibration to the quotes shown in figure 15.

to five years (1.12%) is very close to the simple approximation of spread divided by one minus recovery rate (which is 1.10%). Using these parameters, we then also computed the model's implied values for the protection and the premium legs of the remaining quoted tranches, i.e. 3%–6%, 6%–9%, 9%–12%,

[4]The mid quotes are taken to be the average of bid and offer quotes.

and 12%–22%, and subsequently backed out the Gaussian copula implied base correlation numbers for the respective base tranches. We show the result in figures 17 to 19. The first thing we note is that the

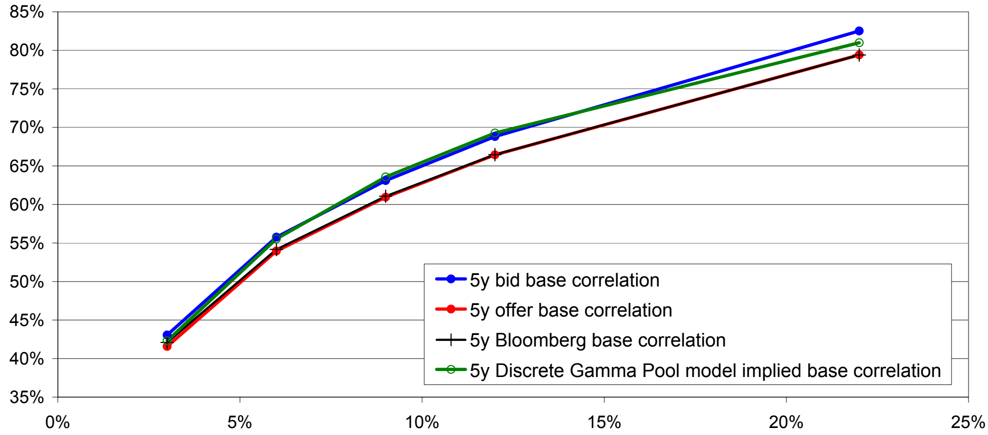

Figure 17: The discrete gamma pool model's calibration results for iTraxx Europe Series 8 expiry 2012-12-20 (nominal maturity 5Y) from Bloomberg page BBIT3 on January 11th 2008.

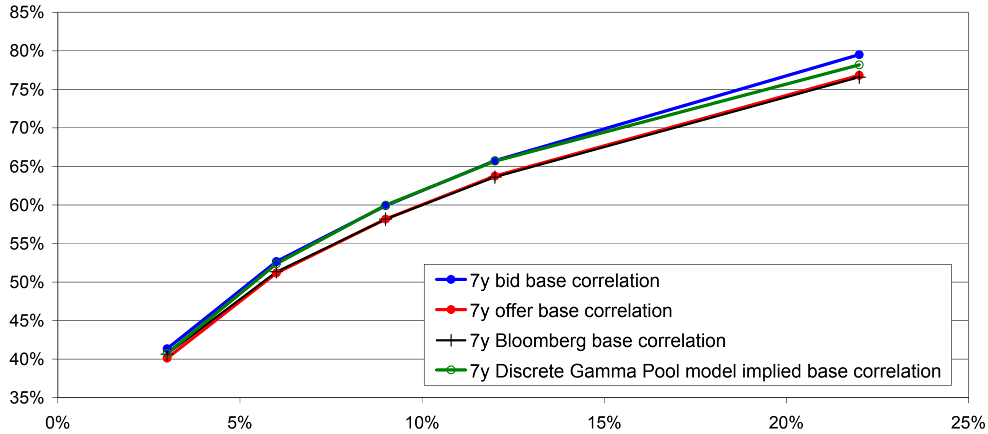

Figure 18: The discrete gamma pool model's calibration results for iTraxx Europe Series 8 expiry 2014-12-20 (nominal maturity 7Y) from Bloomberg page BBIT3 on January 11th 2008.

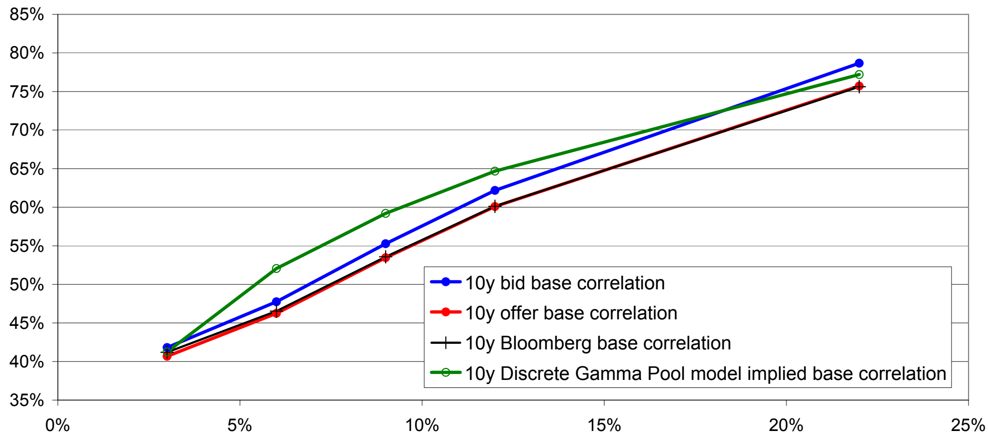

Figure 19: The discrete gamma pool model's calibration results for iTraxx Europe Series 8 expiry 2017-12-20 (nominal maturity 10Y) from Bloomberg page BBIT3 on January 11th 2008.

base correlation figures we took from Bloomberg in the quote table 15 are not near the centre between our bid and offer base correlation curves. We attribute this to the fact that we used our own set of CDS (mid) quotes for the underlying names which we omit here for the sake of brevity, and that we applied no basis correction. This, however, is of no consequence to the discussion of the quality of calibration of the model. The second thing we note is that, with calibration only fixing the 0%–3% and the 0%–22% base correlation, the model seems to calibrate well the intermediate quoted detachment levels, too, for 5Y and 7Y. For the nominal 10Y maturity, the calibration is somewhat suboptimal for the intermediate detachment points. However, specifically for the pricing and risk management of credit correlation products, considering that we could have chosen *any two points* to calibrate to, we consider the quality of fit still well satisfactory.

7 Numerical examples

In this final section, we give details of a specific example for a fully fledged loss triggered leveraged super-senior tranche. Since this product has much more intricate internal transition rules than a vanilla multi-callable CDO, or its counterpart, a Bermudan CDO option, we omit any examples for CDO options for the sake of brevity since all that is learned from an LTLSS case study readily encompasses all we need to know to value CDO options.

The scenario we selected is for a five year maturity ($T = 5$) loss triggered leveraged super-senior note and comprises the following constant parameters:-

$$\begin{array}{cccccc} \hat\lambda = 2\% & \sigma = 50\% & \kappa = 10\% & R = 40\% & \gamma = 5/2 & \phi = 60\% \\ s = 0.20\% & a = 15\% & d = 55\% & \Delta t = 1/26 & \tau = 1/2 & r = 4.5\% \\ \theta_1 = 6\% & \theta_2 = 5\% & \theta_3 = 4\% & m_1 = 2 & m_2 = 5 & m_3 = 10 \end{array} \quad (47)$$

The value of protection on the portfolio is 5.09% of the portfolio's notional in this case. The par spread on the entire portfolio is 120.5 basis points which shows how well the initial guess for $\hat\lambda$ works in this case when calibration is required:

$$\lambda(0) \; = \; \hat\lambda \; = \; 200\text{bp} \; \approx \; 200.83\text{bp} = \frac{120.5\text{bp}}{1 - 40\%} \; = \; \frac{s_{0\%\text{–}100\%}}{1 - R} \quad (48)$$

All calculations were carried out with a loss discretisation of $\Delta\ell = 0.25\%$.

In figure 20 (A) we show the value of the underlying super-senior tranche for this setup from the point of view of the protection seller as a function of the loss in the portfolio and the initial hazard rate value $\lambda(0)$ assuming none of the other parameters (not even $\hat\lambda$) changes. The value of protection on this

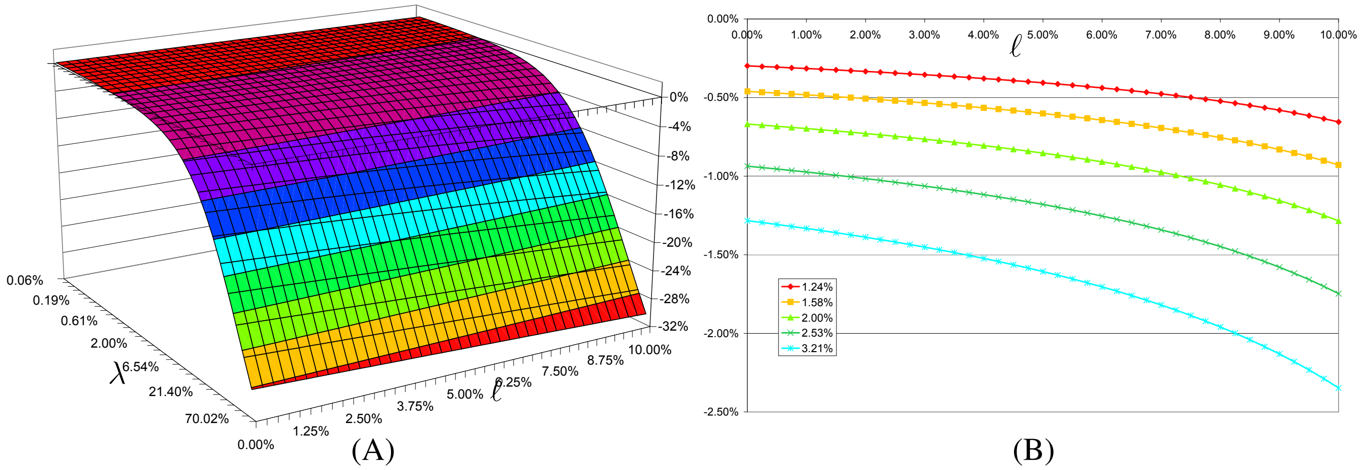

Figure 20: (A) the value of the super-senior tranche as a function of portfolio loss ℓ and current hazard rate λ. (B) the value of the super-senior tranche as a function of portfolio loss ℓ for five different values of the current hazard rate λ as indicated in the legend.

tranche is 101.8 basis points of the portfolio notional, and the par spread is 58.24 basis points. Since it is difficult to see any dependence of the super-senior tranche value on losses, we show this for five

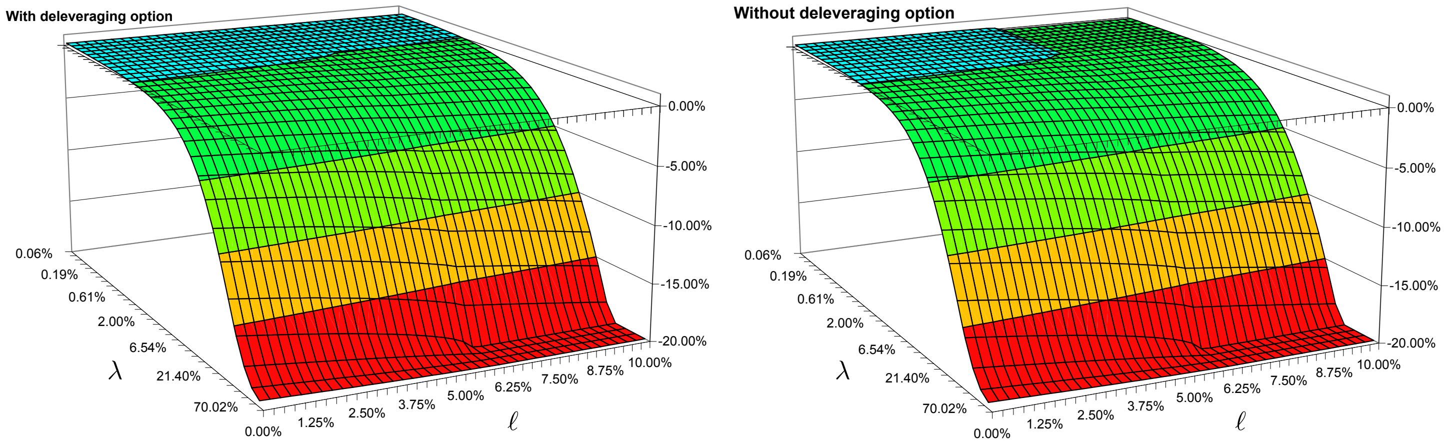

(A) (B)

Figure 21: The value of the auxiliary leveraged super-senior contract leg V_1 as a function of portfolio loss ℓ and current hazard rate λ with (A) and without (B) the option to deleverage.

different levels of λ in figure 20 (B). Next, we show in figure 21 the value of the auxiliary leveraged super-senior contract valuation leg V_1 as defined on page 14 in section 4.1, i.e. the value of a leveraged super-senior contract with a single trigger level at 6% and leverage $m_1 = 2$. It can clearly be seen that the value is floored at -20% which is what can be expected from a $2\times$ leveraging on a tranche of size $d - a = 40\%$. Also, we can observe upon careful inspection that the value of the contract with deleveraging option is higher than that without for very low spreads (i.e. λ) and losses in excess of the trigger level. This becomes very obvious in figure 22 (A) where the difference is shown by itself. In

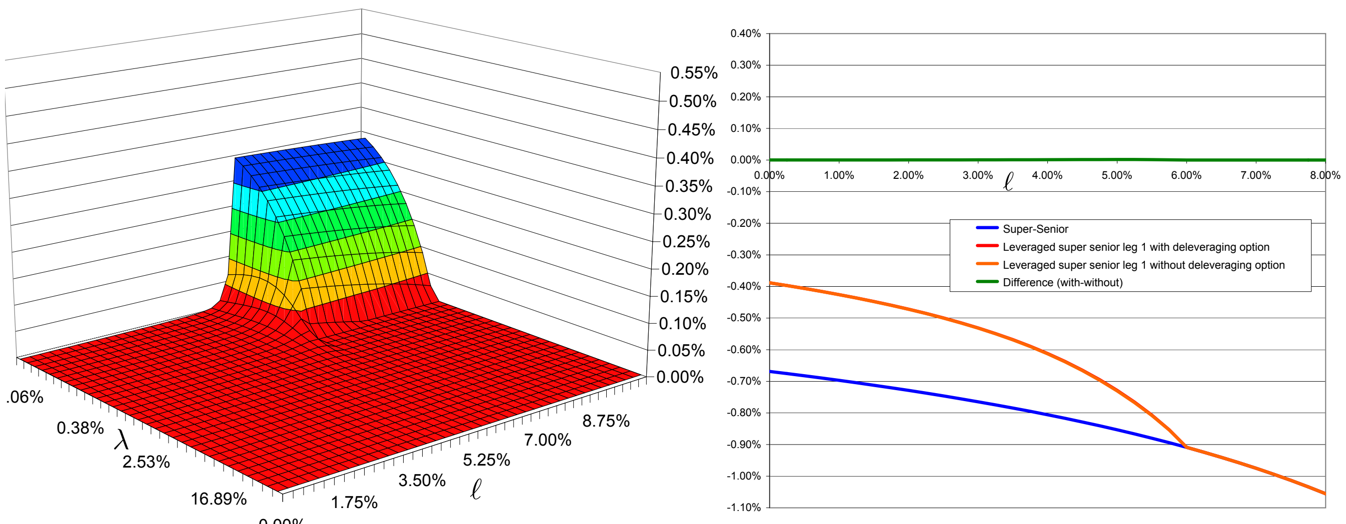

(A) (B)

Figure 22: (A) the difference in value between the auxiliary leveraged super-senior contract leg V_1 with and without the deleveraging option as shown in figure 21. (B) the value of the auxiliary leveraged super-senior contract leg V_1 in comparison to the underlying super-senior tranche itself. Note that the two lines for the leveraged value with and without deleveraging option are in complete overlap in this figure.

figures 23 and 24, we show the same as in figures 21 and 22, but for the auxiliary leveraged super-senior contract leg V_2. Finally, we show in figures 25 and 26 the same analysis for the actual loss-triggered leveraged super-senior contract V_3 we intended to value. It is interesting to note that the difference in value of the contract with and without deleveraging option as shown in figures 22 (A), figures 24 (A), and figures 26 (A) becomes more and more exotic in its profile the further down the nested hierarchy of deleveraging options one looks. In contrast, however, the value profile of the leveraged super-senior contract with deleveraging option itself, remains benign in the sense that it is smooth as a function of spreads and, particularly, as a function of losses. This can be seen by comparing figures 25 (A) and 25 (B) where one can see a drop of value near the first trigger level in (B) but not in (A). This means that, whilst giving the counterparty the option to deleverage increases the value to the counterparty, it smooths out the value profile and thus improves the stability of any hedge.

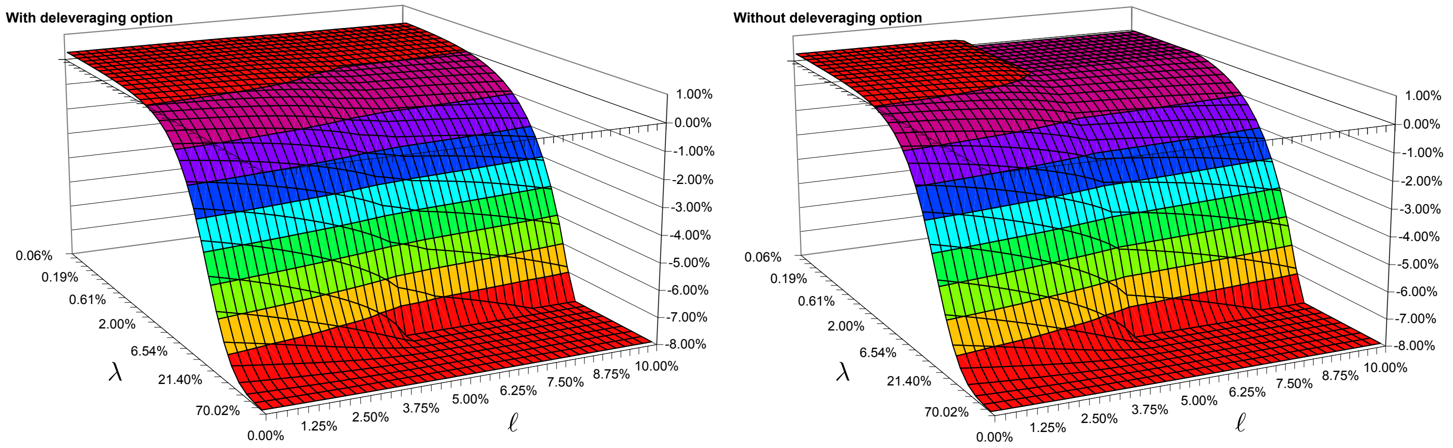

(A) (B)

Figure 23: The value of the auxiliary leveraged super-senior contract leg V_2 as a function of portfolio loss ℓ and current hazard rate λ with (A) and without (B) the option to deleverage.

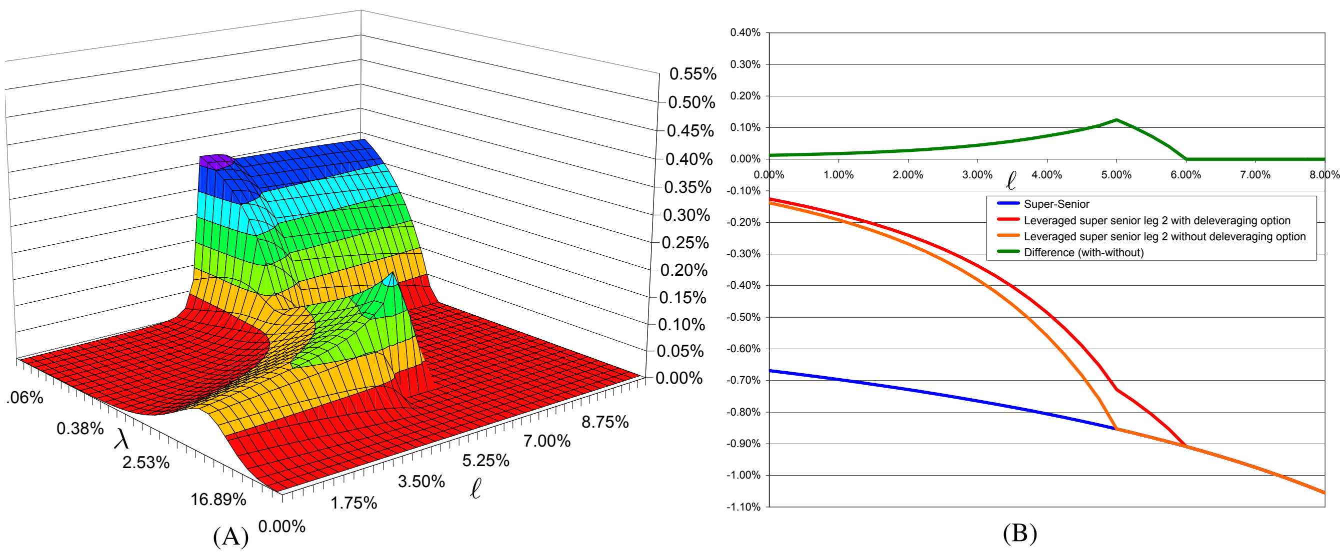

(A) (B)

Figure 24: (A) the difference in value between the auxiliary leveraged super-senior contract leg V_2 with and without the deleveraging option as shown in figure 23. (B) the value of the auxiliary leveraged super-senior contract leg V_2 in comparison to the underlying super-senior tranche itself.

References

[AH07] M. Arnsdorf and I. Halperin. BSLP: Markovian Bivariate Spread-Loss Model for Portfolio Credit Derivatives. Quantitative Research, JP Morgan, 2007. www.igorhalperin.com/

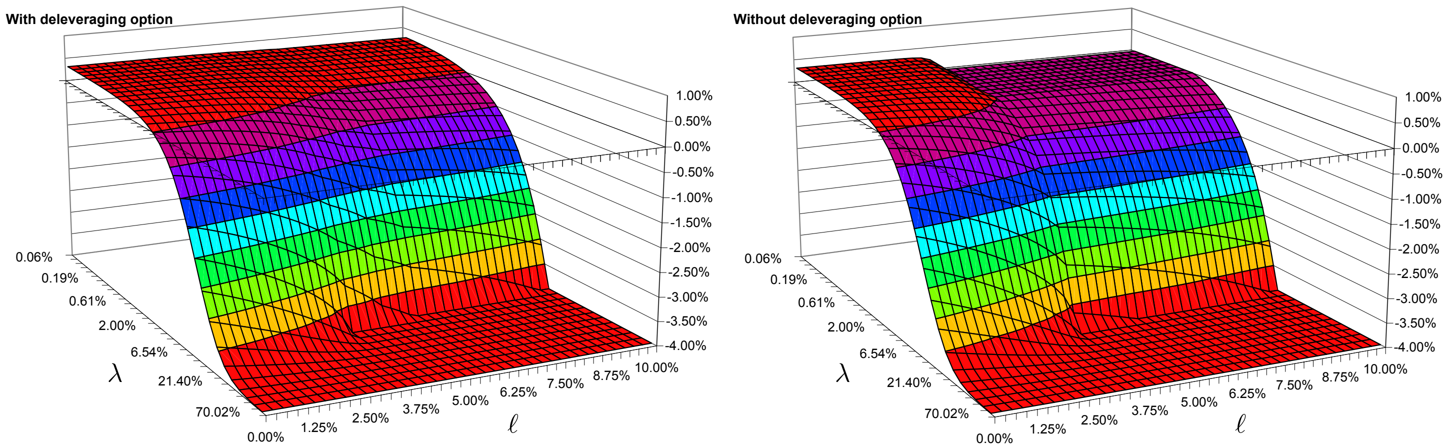

(A) (B)

Figure 25: The value of the leveraged super-senior contract leg V_3 as a function of portfolio loss ℓ and current hazard rate λ with (A) and without (B) the option to deleverage.

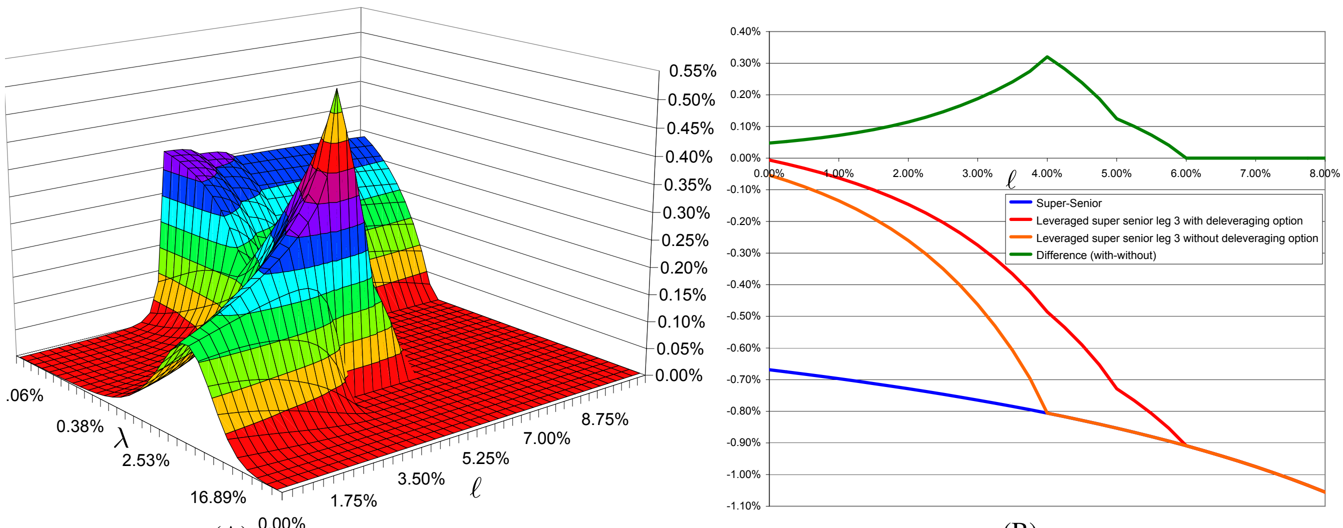

(A) (B)

Figure 26: (A) the difference in value between the leveraged super-senior contract V_3 with and without the deleveraging option as shown in figure 25. (B) the value of the leveraged super-senior contract V_3 in comparison to the underlying super-senior tranche itself.

CreditPortfolios/BSLP_pub.pdf.

[Bax06] M. Baxter. Lévy simple structural models. Technical report, Nomura, December 2006. www.nomura.com/resources/europe/pdfs/cdomodelling_dec06.pdf.

[Bax07] M. Baxter. Gamma process dynamic modelling of credit. *Risk*, 20(10):98–101, October 2007.

[BK91] F. Black and P. Karasinski. Bond and option pricing when short rates are lognormal. *Financial Analysts Journal*, pages 52–59, July/August 1991.

[BPT07] D. Brigo, A. Pallavicini, and R. Torresetti). Calibration of CDO Tranches with the dynamical Generalized-Poisson loss model, 2007. ssrn.com/abstract=900549.

[CIR81] J. C. Cox, J. E. Ingersoll, and S. A. Ross. The relation between forward prices and futures prices. *Journal of Financial Economics*, 9(4):321–346, December 1981.

[CRT06] A. Chapovsky, A. Rennie, and P. Tavares. Stochastic Intensity Modelling for Structured Credit Exotics. Technical report, Merrill Lynch, October 2006. pedro_t.tripod.com/finance/articles/StochIntensityForBaskets.pdf.

[CT04] R. Cont and P. Tankov. *Financial modelling with jump processes*. Chapman & Hall / CRC, 2004.

[DGS91] F. Dufresne, H. Gerber, and E. Shiu. Risk theory with the gamma process. *ASTIN Bulletin*, 21(2):177–192, 1991. www.casact.org/library/astin/vol21no2/177.pdf.

[EGG06] E. Errais, K. Giesecke, and L. Goldberg. Pricing Credit From The Top Down With Affine Point Processes. Working paper, Stanford University, 2006.

[GG05] K. Giesecke and L. Goldberg. A Top Down Approach To Multi-Name Credit. Working paper, Stanford University, 2005.

[HW93] M. Hogan and K. Weintraub. The Lognormal Interest Rate Model and Eurodollar Futures. Discussion paper, Citibank, New York, 1993.

[HW05] J. Hull and A. White. Valuing Credit Derivatives Using an Implied Copula Approach. Working paper, University of Toronto, 2005. www.rotman.utoronto.ca/~hull/DownloadablePublications/ImpliedCopulaPaper.pdf.

[Jäc07] P. Jäckel. The Implementation of the Discrete Gamma Pool Model. www.jaeckel.org/DiscreteGammaPoolModelImplmentation.pdf, September 2007.

[Li00] D. X. Li. On Default Correlation: A Copula Approach. *Journal of Fixed Income*, 9:43–54, March 2000.

[LI07a] A. Lipton and S. Inglis. Factor Models for Credit Correlation. Technical report, Merrill Lynch, 2007. kalx.net/fmcc.pdf.

[LI07b] A. Lipton and S. Inglis. Factor Models for Credit Correlation. *Risk*, 20(12):110–115, December 2007.

[Lip06] A. Lipton. Dynamic Credit Correlation Models. Technical report, Merrill Lynch, 2006. www-cfr.jbs.cam.ac.uk/archive/PRESENTATIONS/seminars/2007/Lipton.pdf.

[LM07] A.V. Lopatin and T. Misirpashaevy. Two-Dimensional Markovian Model for Dynamics of Aggregate Credit Loss, 2007. finmath.stanford.edu/seminars/documents/MarkovianProjectionMisirpashaev06Apr2007.pdf.

[PW84] H. Panjer and G. Willmot. Models for the distribution of aggregate claims in risk theory. *Transactions of the Society of Actuaries*, 36:399–477, 1984.

[SPA05] J. Sidenius, V. Piterbarg, and L. Andersen. A New Framework For Dynamic Credit Portfolio Loss Modelling. Technical report, November 2005. www.jakobsidenius.com/papers/spa_vp_1.pdf.

[SS97] K. Sandmann and D. Sondermann. A Note on the Stability of Lognormal Interest Rate Models and the Pricing of Eurodollar Futures. *Mathematical Finance*, 7(2):119, April 1997.

[Vas87] O. A. Vasicek. Probability of loss on loan portfolio. Working paper, KMV, 1987. www.moodyskmv.com/research/whitepaper/Probability_of_Loss_on_Loan_Portfolio.pdf.

[Wil98] P. Wilmott. *Derivatives*. John Wiley and Sons, 1998.