Price Calibration and Hedging of Correlation Dependent Credit Derivatives using a Structural Model with α-Stable Distributions

Jochen Papenbrock
Chair of Econometrics, Statistics and Mathematical Finance, School of Economics and Business Engineering, University of Karlsruhe and Buttler Consulting, Munich
E-mail: jochen.papenbrock@buttler-consulting.de

Svetlozar T. Rachev
Chair-Professor, Chair of Econometrics, Statistics and Mathematical Finance, School of Economics and Business Engineering, University of Karlsruhe, Postfach 6980, 76128 Karlsruhe, Germany and Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106-3110, USA
E-mail: rachev@statistik.uni-karlsruhe.de*

Markus Höchstötter
Department of Econometrics and Statistics, University of Karlsruhe

Frank J. Fabozzi (corresponding author)
School of Management, Yale University, 135 Prospect Street, New Haven, CT 06511,
Phone: (203) 432-2421, E-mail: frank.fabozzi@yale.edu

*Svetlozar Rachev gratefully acknowledges research support by grants from Division of Mathematical, Life and Physical Sciences, College of Letters and Science, University of California, Santa Barbara, the Deutschen Forschungsgemeinschaft, and the Deutscher Akademischer Austausch Dienst.

Price Calibration and Hedging of Correlation Dependent Credit Derivatives using a Structural Model with α-Stable Distributions

Abstract

The emergence of CDS indices and corresponding credit risk transfer markets with high liquidity and narrow bid-ask spreads has created standard benchmarks for market credit risk and correlation against which portfolio credit risk models can be calibrated. Integrated risk management for correlation dependent credit derivatives, such as single-tranches of synthetic CDOs, requires an approach that adequately reflects the joint default behavior in the underlying credit portfolios. Another important feature for such applications is a flexible model architecture that incorporates the dynamic evolution of underlying credit spreads. In this paper, we present a model that can be calibrated to quotes of CDS index-tranches in a statistically sound way and simultaneously has a dynamic architecture to provide for the joint evolution of distance-to-default measures. This is accomplished by replacing the normal distribution by smoothly truncated α-stable (STS) distributions in the Black/Cox version of the Merton approach for portfolio credit risk. This is possible due to the favorable features of this distribution family, namely, consistent application in the Black/Scholes no-arbitrage framework and the preservation of linear correlation concepts. The calibration to spreads of CDS index tranches is accomplished by a genetic algorithm. Our distribution assumption reflects the observed leptokurtic and asymmetric properties of empirical asset returns since the STS distribution family is basically constructed from α-stable distributions. These exhibit desirable statistical properties such as domains of attraction and the application of the generalized central limit theorem. Moreover, STS distributions fulfill technical restrictions like finite (exponential) moments of arbitrary order. In comparison to the performance of the basic normal distribution model which lacks tail dependence effects, our empirical analysis suggests that our extension with a heavy-tailed and highly peaked distribution provides a better fit to tranche quotes for the iTraxx IG index. Since the underlying implicit modeling of the dynamic evolution of credit spreads leads to such results, this suggests that the proposed model is appropriate to price and hedge complex transactions that are based on correlation dependence. A further application might be integrated risk management activities in debt portfolios where concentration risk is dissolved by means of portfolio credit risk transfer instruments such as synthetic CDOs.

JEL classification: G12, G13

Keywords: Collateralized debt obligations, credit default swaps, index tranches, pricing, risk management, heavy tails, α-stable distribution

1 Introduction

Credit risk modeling has developed rapidly since the late 1990s. This has been fostered by the significant growth in the credit derivatives market of more complex and model–driven trading strategies and credit risk transfer activities. The market for synthetic collateralized debt obligations (CDOs), a vehicle that transfers the risk of a pool of single–name credit default swaps (CDS), is an example.

The development of such technologies has been fueled by the growth and liquidity of the CDS market and the creation of broad–based credit risk indices such as iTraxx or CDX. These CDS index products provide standard benchmarks against which other more customized pools of credit exposure can be assessed. Moreover, they serve as building blocks for other products such as CDS index tranches. These standardized tranches of a CDS index portfolio render possible a marking–to–market of credit risk correlations. By means of a standard model, their competitive quotes – in terms of cost of protection of a single tranche – are translated into so–called "implied correlations".

The current standard for price quotation of credit portfolio products – such as CDOs – is the one-factor Gaussian copula. It is a tool to aggregate information about the impact of default correlation on the performance of a rather static credit portfolio. Given a representative estimate of the term structure of credit spreads and a representative loss given default (LGD), the market-standard version of this copula is characterized by a single parameter to summarize all correlations among the various borrowers' default times. However, the fact that index tranches are quoted frequently and with relatively narrow bid–ask spreads has aided market participants in identifying several shortcomings of the existing pricing models for CDOs. In particular, the Gaussian copula model does not fit market prices very well.[1] The model underperformance can be observed due to the pronounced correlation smile when implied CDO tranche correlations are plotted as a function of tranche attachment points.

One possibility to resolve these shortcomings is to consider heavy–tailed distributions. In comparison to the normal distribution, heavy–tailed distributions incorporate the more frequent occurrence of extreme events in empirical asset returns. In the multivariate case, they exhibit measures of dependence that go beyond the concept of linear correlation. For example, certain tail dependence effects replicate an increase in credit default clustering during times of economic recession. The consideration of such effects may lead to improved risk management applications with respect to pricing and hedging accuracy.

An example of a heavy–tailed distribution is the double Student–t copula proposed by Hull and White (2004) where the interaction of heavy–tailed systematic and idiosyncratic factors lead to a default environment that is based on two effects: The basic linear dependency known as the only source of dependency in the Gaussian framework and tail dependence effects that create extreme systematic co–movements of firm values, combined with extreme idiosyncratic outcomes. This model exhibits a good overall fit to standardized index tranches, since prices are closer to the market quotes.[2]

The heavy-tailed copula model suggested by Hull and White, however, has two shortcomings.[3] The first is that the tail-fatness cannot be changed continuously. The second is that the maximum tail-fatness occurs when the Student–t distribution has 3 degrees of freedom.[4]

Another important aspect for integrated pricing and credit risk management applications is the employment of dynamic approaches that incorporate both an adequate modeling of default dependency as well as the joint evolution of credit spreads.[5] The latter aspect is preserved by the structural model of Hull, Predescu and White (2005) (HPW) that is a dynamic Merton–style approach in the flavor of Black and Cox (1976) incorporating intertemporal defaults.

On this modeling basis we enrich the desirable features of the dynamic structural model with a complex

[1]See for example Burtschell et al. (2005), p. 17.
[2]See Burtschell et al. (2005), p. 17.
[3]See Wang, Rachev and Fabozzi (2007), p. 12
[4]For these reasons, Wang, Rachev and Fabozzi (2007) introduce two new one-factor heavy-tailed copula models: (1) the one–factor double–t distribution with fractional degrees of freedom copula model and (2) the one–factor double mixture distribution of t and Gaussian distribution copula model. In each model, there is a parameter to continuously control the tail–fatness of the copula function. Moreover, the maximum tail-fatnesses of our two models are much larger than that for Hull and White's one-factor double–t copula model.
[5]See Duffie (2004).

dependence environment similar to the one in the double Student–t copula. We consistently replace the Gaussian distribution assumption in the underlying factor model with a smoothly truncated α–stable (STS) distribution.[6] This distribution family can be applied in the Black–Scholes no–arbitrage model due to finite moments of arbitrary order.

According to tail probability studies, the STS distribution family assigns much more probability mass to the tails than the Gaussian and even the Student–t distribution. Also, Wang, Rachev and Fabozzi (2007) have used continuous parameters for heavy–tailedness in copula functions to analyze time–dependent model performance in pricing standard tranches. They show that the Student–t copula model with non–fractional degrees of freedom is outperformed. Both reasons led to our choice to employ the STS distribution. It exhibits a continuous parameter to control leptokurtosis with less restriction than the Student–t distribution. Moreover there is a continuous parameter for asymmetry.

This combination of the HPW model with an engineered distribution leads to a considerable improvement in both modeling of default timing as well as joint credit spread movements in credit portfolios. In this way, established concepts such as the Black–Scholes no–arbitrage and linear correlation can be preserved while simultaneously providing improvements in price quality, hedging accuracy, and risk management effectiveness. Two applications include integrated pricing and management for portfolio credit risk as well as pricing and hedging of even complex structures like CDO^2 (i.e., CDOs of CDOs) or options on single-tranche CDOs.

The paper is organized as followed. The structural model by Hull–Predescu–White is outlined in Section 2. Section 3 enhances the structural model by replacing the Gaussian distribution assumption with the STS distribution assumption. The applied calibration and valuation framework for index tranches is explained in Section 4. Section 5 reports the fitting properties of the model to standardized index tranches and Section 6 summarizes the contents of this paper.

2 The Hull–Predescu–White Model

2.1 Outline of the model

As in the case of modeling all derivative instruments, moving from the general principles of pricing to that of pricing a specific type of derivative one must consider the specific contractual feature. For the CDS index tranches, this means taking into account the fact that a tranche's outstanding notional amount declines stochastically over time. A fixed spread is paid to the protection seller on the decaying tranche notionals with the payments being made quarterly. So for each intermediate payment day between inception and termination, a loss distribution has to be assessed incorporating only discrete default times, in the simple case.

In practice, a one–factor Gaussian copula is often used to model intertemporal stochastic cash flows.[7] Since default events in the pool of names between premium payment days are relevant for pricing, credit default distributions for respective time horizons ranging from t_0 – the beginning of the deal – to the premium payment dates t_j are generated in the first step. In the second step, the stochastic changes between the payment days are used to price the tranches. The standard one–factor Gaussian copula model can therefore be regarded as static since there is no dynamic evolution of the underlying distance–to–default measures or, similarly, credit spreads.[8]

The structural model of Hull, Predescu and White (2005) however is much richer because the portfolio behavior is modeled chronologically until maturity. This is accomplished by a factor model based on the approach of Black and Cox (1976). Their extension of the static Merton model has a first passage time structure where a default event is triggered as soon as the value of the assets of a company drops below a continuous barrier level for the first time. This is realized by a general diffusion process of an obligor's default variable and an appropriately chosen barrier function that is made consistent with the underlying default time distribution. In the Hull–Predescu–White (HPW) extension, the default variables of the underlying obligors follow *correlated* diffusion processes and the barrier for

[6] Menn and Rachev (2005b).

[7] Similar to the concept of implied volatility in option pricing, the Gaussian copula has become the market standard to communicate prices of synthetic CDO tranches.

[8] See Hull, Predescu and White (2005), p. 11.

each obligor is calibrated in a way that it is made consistent with the respective marginal default time distributions. During the discrete simulation of the correlated processes, the common factor M adopts different values being constant in the specific time span of one process increment. In this way the default environment changes over time. As a by–product of this procedure, the joint evolution of correlated credit spreads is obtained.

In accordance with the standard market model, the HPW model is set up with a diffusion process for the value of the firm V_i of obligor i as follows:

$$dV_i(t) = \mu_i V_i(t)dt + \sigma_i V_i(t)dW_i(t)$$

with

$$V_i(T) = V_i(t)e^{\left(\mu_i - \sigma_i^2/2\right)(T-t)+\sigma_i X_i(t,T)}$$

and

$$X_i(t,T) \stackrel{d}{=} W_i(T) - W_i(t).^9$$

The expected return of the firm is μ_i, σ_i is the instantaneous standard deviation, and $W_i(t)$ is a Brownian motion under the real measure. Variable X_i can be imagined as some function of the value of the assets or the creditworthiness of company i. The resulting barrier equation is:

$$D_i(t) = \frac{\ln K_i - \ln V_i(t) - \left(\mu_i - \sigma_i^2/2\right)(T-t)}{\sigma_i},$$

with K_i as the notional repayment at maturity T in the Merton context. Hull and White (2001) present a discretized version of the model that can be solved numerically. This is necessary for extensions with distributions that do not exhibit closed–form expressions.

The model is set up in terms of the risk–neutral default probability density $q(t)$. This means that $q(t)\Delta t$ is the probability of default between t and $t + \Delta t$ as seen at time zero. In contrast, the hazard (default intensity) rate $\lambda(t)$ is defined as the probability of default between t and $t + \Delta t$ as seen at time t conditional on no earlier default. The two quantities provide the same information about the default probability environment and they are related by

$$q(t) = \lambda(t)e^{-\int\limits_t^T \lambda(\tau)d\tau},$$

when the exponential model for the default time distribution is employed. We later assume when we extend the model that default probabilities for entities of a homogeneous portfolio are generated by the same Poisson processes with constant "risk–neutral" default intensity λ so that:

$$Q(t) = 1 - e^{-\lambda t} \qquad \text{and} \qquad Q(t, t + \Delta t) = e^{-\lambda t} - e^{-\lambda(t+\Delta t)}.$$

With this assumption, we can derive the representative default intensity λ from the quoted CDS index spread as shown in Section 4.1. This allows us to compute "intermediate" default time distribution slices like $Q(t, t + \Delta t)$ to avoid the interpolation of risk–neutral default probabilities since the Black–Cox default barrier methodology is extremely sensitive to the applied interpolation method of the risk-neutral default probabilities retrieved from CDS or credit spread curves. In our calibration procedure, we use the simple exponential model to retrieve a default probability distribution function.

2.2 Construction of the Discrete Default Barriers

The default barrier algorithm is conveniently modeled in a synchronized way to the time grid of CDS premium payment dates $t_j,\ j = 1, \ldots, J$. In general, finer time grids will make the model arbitrarily close to an environment where defaults can happen at any time. In our computations, the default probability distribution is discretized so that defaults are modeled to happen at times t_j and further, they are associated with the midpoints $\frac{t_{j-1}+t_j}{2}$ in the pricing part in Section 4. Due to simplification, accrual effects between premium payment days will be neglected.

[9] With $\stackrel{d}{=}$ meaning equality in distribution.

The objective is to determine a default barrier for each company such that the default event is triggered when the firm's diffusion process first hits the barrier at this time. The barrier must be chosen so that the first passage time probability distribution is the same as the default probability densities $q(t)$. It is assumed that $X_i(0) = 0$ and that the risk–neutral process for $X_i(t)$ is a Wiener process with zero mean and unit variance per year. Additionally, the following definitions have to be made:

- the time grid is equidistant with $\delta = t_j - t_{j-1}; \quad j = 1, \ldots, J$,
- the risk–neutral first passage time probability for the interval $[t_{j-1}, t_j]$ is $Q_i(t_{j-1}, t_j); \quad j = 1, \ldots, J; \quad i = 1, \ldots, n$,
- the value of the default barrier for company i at time t_j is $D_i(t_j)$,
- $f_{ij}(x)\Delta x$ denotes the probability that $X_i(t_j)$ lies between x and $x + \Delta x$ and there has been no default prior to time t_j.

These definitions imply for the probability of first passage at time t_j that

$$Q(t_j) = 1 - \int_{D_{ij}}^{\infty} f_{ij}(x) dx.$$

Both D_{ij} and $f_{ij}(x)$ can be determined from $Q_i(t_{j-1}, t_j)$. The first barrier is found by the first increment $X_i(t_1)$ which is distributed zero mean and variance δ. As a result,

$$f_{i1}(x) = \varphi\left(\frac{x}{\delta}\right) \quad \text{and} \quad Q_i(t_0, t_1) = Q_i(t_1) = \Phi\left(\frac{D_{i1}}{\sqrt{\delta}}\right).$$

This implies that

$$D_{i1} = \sqrt{\delta}\Phi^{-1}(Q_i(t_1)).$$

The first barrier has been identified. If the distribution under consideration is not normal and there is no inverse evaluation method available, the barrier can be found by standard numerical procedures.

The probability that, in t_1, the process is in a survival position above the first barrier D_{i1} and that it will default in t_2 has to be equal to the probability of first hitting the barrier between t_1 and t_2. For determining the barrier D_{i2}, in our algorithm we find an approximation to the solution by nested intervals up to a certain tolerance level. The general equation for payment times $t_j,\ j = 2, \ldots, J$ is

$$Q(t_{j-1}, t_j) = \int_{D_{i,j-1}}^{\infty} f_{i,j-1}(u)\ \Phi\left(\frac{D_{ij} - u}{\sqrt{\delta}}\right) du. \tag{1}$$

The value for $f_{ij}(x)$ for all x above barrier D_{ij} is

$$f_{ij}(x) = \int_{D_{i,j-1}}^{\infty} f_{i,j-1}(u)\ \varphi\left(\frac{x - u}{\delta}\right) du. \tag{2}$$

where φ and Φ denote the standard normal probability density and distribution function, respectively.

Equations (1) and (2) can be solved numerically in the following way: For time grid point $j = 1, \ldots, J$ we consider K values for $X_i(t_j)$ between D_{ij} and a multiple of $\sqrt{t_j}$. In this way we bound the half–open intervals on the vertical line dynamically according to the deviation of the respective distribution. We define x_{ijk} as the kth value of $X_i(t_j)$ $(1 \le k \le K)$ and π_{ijk} as the probability that $X_i(t_j) = x_{ijk}$ with no earlier default. The discrete versions of equations (1) and (2) are

$$Q(t_{j-1}, t_j) = \sum_{k=1}^{K} \pi_{i,j-1,k} \Phi\left(\frac{D_{ij} - x_{i,j-1,k}}{\sqrt{\delta}}\right)$$

and

$$\pi_{ijl} = \sum_{k=1}^{K} \pi_{i,j-1,k}\ p_{ijkl},$$

where p_{ijkl} is the probability that X_i moves from $x_{i,j-1,k}$ at time t_{j-1} to x_{ijl} at time t_j. This can be accomplished with the following equation

$$p_{ijkl} = \Phi\left[\frac{0.5(x_{ijl} + x_{i,j,l+1}) - x_{i,j-1,k}}{\sqrt{\delta}}\right] - \Phi\left[\frac{0.5(x_{ijl} + x_{i,j,l-1}) - x_{i,j-1,k}}{\sqrt{\delta}}\right]$$

for $1 < l < K$. For $l = K$ we use the same equation with the first term on the right hand side equal to 1 to represent the unbounded integral

$$p_{ijkK} = 1.0 - \Phi\left[\frac{0.5(x_{ijl} + x_{i,j,l-1}) - x_{i,j-1,k}}{\sqrt{\delta}}\right].$$

When $l = 1$ we use the same equation with $0.5(x_{ijl} + x_{i,j,l-1})$ set equal to D_{ij} to define the first interval in the survival region

$$p_{ijk1} = \Phi\left[\frac{0.5(x_{ijl} + x_{i,j,l+1}) - x_{i,j-1,k}}{\sqrt{\delta}}\right] - \Phi\left[\frac{D_{ij} - x_{i,j-1,k}}{\sqrt{\delta}}\right].$$

In this way, for $1 < l < K$ there is assigned a certain probability mass of the process to be in the interval $[0.5(x_{ijl} + x_{i,j,l-1}), 0.5(x_{ijl} + x_{i,j,l+1})]$ at time t_j with the representative midpoint x_{ijl}. This is conditional on survival up to time t_{j-1} which is quantified by the probability $\pi_{i,j-1,k}$ for the representative midpoint $x_{i,j-1,k}$.

2.3 Simulation and Dynamic Credit Spreads

There exists an analytic expression of the probability of first hitting the barrier between times t and $t + \Delta t$.[10] When suppressing indices we have

$$Q(t, t + \Delta t) = \Phi\left(\frac{D(t + \Delta t) - X(t)}{\sqrt{\Delta t}}\right) + e^{2(X(t) - D(t))\frac{\mu - \sigma^2/2}{\sigma}}\Phi\left(\frac{D(t - \Delta t) - X(t)}{\sqrt{\Delta t}}\right) \tag{3}$$

which, in our case, will be given for $\Delta t = t_j - t_{j-1},\ j = 1, \ldots, J$.

The process for the mean zero and variance δ state variable X_i is

$$dX_i(t) = a_i dM(t) + \sqrt{1 - a_i^2}\, dZ_i(t)$$

when asset correlations are incorporated. In the Monte–Carlo implementation, we approximate this by

$$\Delta X_i = a_i \Delta M + \sqrt{1 - a_i^2}\, \Delta Z_i, \tag{4}$$

where ΔM and ΔZ_i are distributed *i.i.d.* $\mathcal{N}(0, \delta)$. The variables a_i, M, and Z_i in this model have a slightly different meaning than in the one–factor Gaussian copula approach due to the different model set–ups. Nevertheless, the correlation between the processes followed by the assets of companies i_1 and i_2, respectively, is $a_{i_1} a_{i_2}$.

While the Gaussian copula is a reasonable approximation to the HPW model, it is limited due to the fact that the only means of expressing dependence structures is given by the correlation coefficients. However, this is insufficient in most realistic cases when marginal distributions are used that are heavy–tailed. For example, the Student–t extension to the standard model outlined above exhibits tail dependence effects which cannot be modeled in the Gaussian case. In the HPW approach, the increments are constructed as the discrete convolution of two heavy–tailed variables. So the occurrence of extreme events enters the model at multiple stages. In case of the distribution being short of a closed form, (3) is substituted by probabilities of default before maturity derived from the discrete barrier algorithm. Hence, one is able to compute the joint evolution of the dynamic credit spread.

An extended barrier algorithm can easily be applied under the heterogeneous portfolio assumption: The calibration will be carried out with each of the marginal default time distributions and whenever the process X_i hits the specific barrier, a recovery rate R_i is assigned. The computational performance of the simulation is not affected but discrete barriers have to be calibrated for each underlying.

[10] See Hull, Predescu and White (2005), p. 7.

3 Extension with Smoothly Truncated Stable Distributions

In option pricing it is essential that the property of not only finite moments but also the existence of finite exponential moments is guaranteed. This is fulfilled by the Gaussian distribution. Since Fama (1965) and Mandelbrot (1963), however, it has been widely accepted that asset returns are asymmetric and leptokurtic with heavy tails. Their proposed alternative was the Pareto or α- stable distribution class. This class applies four parameters enabling the distribution to model excatly these features found in asset returns. The deficiency, however, of the α-stable distributions becomes apparent when computing moments, not even to mention exponential moments, since, under certain conditions to be presented, they do not exist.

Therefore, we employ a new class of probability distributions called smoothly truncated α–stable (STS) distributions. This distribution combines the modeling flexibility of stable distributions with the existence of arbitrary moments and thereby qualifies for applications in the Black–Scholes–Merton framework.

In this section, we will first briefly outline the characteristics of α–stable distributions and then present the method by Menn and Rachev (2004b) who offer a calibrated Fast Fourier Transform (FFT) based density approximation of α–stable distributions. The numerical generation of the cumulative α–stable distribution function is essential for the the smooth truncation and standardization procedures needed to set up the STS distribution. Finally, we present a method to simulate STS distributions based on the method to generate α–stable samples by Chambers et al. (1976).

3.1 The Stable Distribution Family

3.2 Stable random variables

Stable distributions are characterized by four parameters. The tail index, index of stability, or characteristic parameter α is responsible for the shape of the distribution in the tails as well as around the median. It determines the rate at which the tails of the distribution taper off. When $\alpha = 2$, a Gaussian distribution results. When $\alpha < 2$, the variance is infinite and the tails are asymptotically equivalent to a Pareto law (i.e., they exhibit a power–law behavior). Distributions with $1 < \alpha < 2$ parameters have unbounded variance but bounded mean. Those with α between 0 and 1 have both unbounded variance and mean. In general, moments of order δ exist up to $\delta < \alpha$.
Furthermore, skewness is accounted for by the parameter $\beta \in [-1, 1]$. Scale is modeled by the parameter $\sigma > 0$, and, finally, a measure of location is given by parameter $\mu \in \mathbb{R}$ which, for $\alpha > 1$, represents the mean.

Moreover, stable distributions possess the property of domains of attraction. If an empirical distribution is in the domain of attraction of a stable law, it has properties which are close to those of the specified stable law. The domain of attraction is completely determined by the tail behavior of the distribution and as a result the stable law is the ideal model if the true distribution has the appropriate tail behavior.

According to the stability property, appropriately centralized and normalized sums of iid α–stable random variables are again α–stable. This in turn means that α–stable distributions lie in their own domain of attraction which is a desirable property. Due to the *Generalized Central Limit Theorem* (GCLT), the stable class provides limit distributions for scaled sums of infinite variance random variables.

Unfortunately, the application of stable laws in finance is at a disadvantage because of the lack of closed–form expressions for their probability density and cumulative distribution functions for most parameter values. Hence, numerical approaches have to overcome this deficiency.[11]

The α–stable distribution can be most naturally and conveniently described by its characteristic function $\phi(t)$ – the inverse Fourier transform of the probability density function. The most popular parameterization of the characteristic function of $X \sim S_\alpha(\sigma, \beta, \mu)$, i.e. an α–stable random variable

[11] See, for example, Zolotarev (1966).

with parameters α, σ, β, and μ, is given by Samorodnitsky and Taqqu (1994)

$$\log\phi(t) = \begin{cases} -\sigma^\alpha |t|^\alpha \left(1 - i\beta \text{sign}(t)\tan(\frac{\pi\alpha}{2})\right) + i\mu t \ , & \alpha \neq 1, \\ -\sigma|t| \left(1 + i\beta\frac{2}{\pi}\text{sign}(t)\ln|t|\right) + i\mu t \ , & \alpha = 1. \end{cases} \tag{5}$$

The representation in formula (5) is discontinuous at $\alpha = 1$ and $\beta \neq 0$. This can be overcome by shifting the variables by some amount depending on α and β. However, our representation fulfills our requirements since in most of the cases, α for financial data is larger than 1.5.[12] The parameter μ is equal to the mean.

The concept of smooth tail truncation allows for the preservation of the properties of α–stable distributions in the "center" of the engineered distribution, whereas an exponentially declining function replaces the power decaying tails of the stable law in order to guarantee the existence of arbitrary moments. Before we will explain the construction, the properties and implementational aspects of the STS distribution, an efficient algorithm for density approximations for stable non–Gaussian distributions will be outlined.

3.2.1 Density Approximation of Stable Distributions

The unambiguous relationship between the density function and the characteristic function is exploited by the FFT approach.[13] Concerning the computational speed, the FFT–based approach is faster for large samples, whereas the direct integration method favors small data sets as it can be computed at any arbitrarily chosen point. The FFT–based approach is not as universal as the direct integration method – it is efficient only for large α's and only as far as the probability density function calculations are concerned. When computing the cumulative distribution function, the former method must numerically integrate the density, whereas the latter takes the same amount of time in both cases.

We therefore decided to implement a simplified version of the calibrated FFT–based density approximation by Menn and Rachev (2004b) who employ an adaptive Simpson rule for the quadrature of the Fourier inversion integral. Since this approach lacks precision in the tails, they follow the suggestion of DuMouchel (1971) to use some additional asymptotic series expansion developed by Bergström (1952) in order to receive efficient tail approximations. The accuracy of the method is optimized with respect to values obtained by Nolan's STABLE.exe for a grid of parameter values of α and β. This is sufficient for stable distributions since they are scale and translation (i.e. shift) invariant. Density evaluations departing from the FFT grid nodes and the generation of the cumulative distribution function are performed by cubic spline interpolations. In comparison to Nolan's program, the approach results in a significant reduction of the computation time while simultaneously preserving satisfactory accuracy.[14]

3.2.2 Simulation of Stable Random Variables

The complexity of the problem of simulating sequences of α–stable random variables comes from the fact that there are no analytic expressions for the inverse $F^{-1}(x)$. A more elegant and efficient solution for standardized skewed α–stable distributions was proposed by Chambers et al. (1976). The method reduces to the well–known Box–Müller method for Gaussian distributions in the case of $\alpha = 2$ (and $\beta = 0$), and is based on a certain integral formula derived by Zolotarev (1966).

We can easily simulate a stable random variable for all admissible values of the parameters α, β, σ, and μ, with random variable X being standard α–stable distributed using the following property: if $X \sim S_\alpha(1, \beta, 0)$ then, for $\alpha \neq 1$, $Y = \sigma X + \mu$ for $\alpha \neq 1$ is $S_\alpha(\sigma, \beta, \mu)$–distributed.

3.3 Smoothly Truncated Stable Distributions

Guaranteeing a finite mean for the asset price, the class of STS distributions share with stable distributions some realistic features such as leptokurtosis and skewness which has been observed in asset return behavior. Despite the fact that STS distributions possess light tails in the mathematical sense,

[12]See Rachev et al. (2005), figure 1.
[13]See, for example, Höchstötter, Rachev and Fabozzi (2005).
[14]This is comprehensively quantified in Menn and Rachev (2004b).

they provide a flexible tool to model extreme events since a reasonable amount of probability is assigned to extreme events. Technically, however, tail dependence of STS distributions is zero due to exponential tails.

STS distributions are obtained by smoothly replacing the upper and lower tail of an arbitrary α–stable cumulative distribution function by two appropriately chosen normal tails. The result is a continuously differentiable probability distribution function with support on the whole real line. By this construction, the density of an STS distribution consists of three parts: Left of some lower truncation level a and right of some upper truncation level b, it is described by two outer normal densities and in the center the density equals the one of a stable distribution. If the stable distribution in the center is symmetric around zero, the means of the two normal distributions only differ in sign while the variance is equal. However, this does not apply to a skewed stable center distribution.

Due to the finite moment generating function which results from truncation, STS distributions lie in the domain of attraction of the Gaussian law. Owing to the amount of probability of extreme events, the speed of convergence to the normal distribution is extremely slow. It can be stated that the family of STS distributions provides impressive modeling flexibility and turns out to be a viable alternative to many popular heavy–tailed distributions.

STS distributions form a six parameter distribution family $S_{\alpha}^{[a,b]}(\sigma, \beta, \mu)$, where a and b are the truncation points of the α–stable distribution. The parameters (μ_i, σ_i) of the two normal distributions, respectively, are uniquely defined by construction.

In the HPW framework reviewed above, the only imposed conditions on the factors are a continuous probability distribution function with support on the whole real line having zero mean, unit variance, and a finite moment generating function. A properly standardized STS distribution uniquely defined by the vector of stable parameters $\theta = (\alpha, \beta, \sigma, \mu)$ fulfills these requirements. Truncation levels for all four stable parameters can always be efficiently calculated by moment matching conditions such that the resulting distribution is standardized.

Analogous to stable random variables, there is an interpretation for STS distributions between parameter α and the probability for extreme events. The latter increases monotonically with decreasing α, decreasing β, and decreasing σ. Keeping the other stable parameters constant, the left truncation level a decreases and the right truncation level b increases monotonically with increasing α. This follows mathematical intuition since for small values of α, the stable center distribution is extremely heavy–tailed and has to be cut off near the mode to arrive at a unit variance. Since σ represents the scale parameter of the stable distribution part, the variation of the center distribution increases with increasing σ: The truncation has to be accomplished in a certain range around the mode to guarantee a variance of one.[15]

Regarding the implementation of the density estimation, the modules from the calibrated FFT–density approximation for the center with Bergstöm series expansion for the tails can be utilized to perform the necessary interpolation and integration procedures on the basis of cubic splines. For the random sample generation, the algorithm by Chambers–Mallows–Stuck can be used in combination with an algorithm for Gaussian samples for the tail distributions.

3.3.1 STS Distributions in the HPW Model

We extend the HPW model by standardized STS distributed factors so that

$$\Delta X_i = a_i \Delta M + \sqrt{1 - a_i^2}\Delta Z_i \tag{6}$$

and $\Delta M/\sqrt{\delta}$ and $\Delta Z_i/\sqrt{\delta}$ have independent standardized STS distributions with the same parameters α, β, and σ. This is in accordance with the factor extension of HPW, so that the correlation between the assets is $\rho = a_{i_1} a_{i_2}$ for each different pair of assets.

Hull and White (2004) have shown that the double Student–t copula approach with same tail index

[15] See Menn and Rachev (2005b), p. 11.

for both factors results in a good market fit.[16] For this reason we use the same parameters for the distributions of ΔM and ΔZ_i.

Convolutions of the STS distributions have to be computed numerically. The idea for the implementation is similar to the construction of the default barriers as explained in Section 2.2. There is a grid of intervals and a certain amount of probability is assigned to the midpoints. This applies for the left summand of the right hand side of equation (6). Conditional on those probabilities we build up the cumulative distribution for certain grid points with the distribution of the right summand of equation (6). The open interval distribution parts of the sum of the two factors are adapted in the same way as for the barrier computations to represent infinite support of the distributions. These operations could be extended to several independent systematic factors in the usual way, but we conveniently restrict ourselves to a one–factor model.[17]

4 The Valuation of Synthetic CDOs

The purpose of this section is to outline our valuation of synthetic CDOs. To create these structures, the owner of a portfolio of single–name CDS distributes the credit risk by creating loss tranches which, in return, are sold to investors.[18] A standardized index portfolio of CDSs is used as a reference portfolio with synthetic CDO tranches. The protection seller offers compensation for losses induced by credit events in this portfolio of reference entities. On the other hand, the owner of the portfolio as the protection buyer pays a periodic premium to the protection seller. The premium is expressed as an annual spread on the tranche's outstanding notional. Premiums are usually paid quarterly.

The pricing of the tranche spreads is accomplished by matching the discounted expectations of the payments of the protection seller and the protection buyer. This spread can be computed using an actuarial approach based on a fixed premium leg and a floating protection leg for different tranches, respectively.[19]

A further development in the market involves what is known as "single tranche CDOs". These deals are based on an arbitrary portfolio and some tranche where the buyer and seller of protection agree to exchange the cash flows that would have been applicable as if a synthetic CDO had been set up. The most important standard portfolios used for this purpose are the CDX IG, a portfolio of 125 investment–grade companies in North America, and the iTraxx IG, a portfolio of 125 European investment–grade companies.[20]

The CDO structure is similar to a derivative on a credit portfolio based on percentiles with the following attributes. The buyer of a tranche l with lower attachment K_l^L and higher detachment point K_l^U will bear all losses in the portfolio value in excess of $K_{l,L}$ and up to K_l^U percent of the initial value of the portfolio N_{total} such that the constructed CDO's loss exposure is limited to $K_l^U - K_l^L$ percent of the initial portfolio value. Table 1 summarizes the different attachment/detachment percentage levels for the two standard indices iTraxx IG and CDX IG.

PLACE TABLE 1 ABOUT HERE

Taking the risk–neutral default time probability distribution of the underlying names as given, we generate future scenarios for the loss behavior of the portfolio. Under the assumption that the only

[16] There exist similar models in practice and it is often assumed that both M and Z_i have distributions with the same tail index.

[17] Fortunately, performance can be strongly improved to restrict the grid to a smaller abscissa range. This is possible since the truncation produces negligible small values for the normal distributions in the tails due to their non–heavy–tailed character. All procedures mentioned so far – including the numerical convolution – consume 12 seconds for one specific parameter tuple (α, σ) in the symmetric case in C++ on a 1.5GHz processor and 512 MB of RAM.

[18] See Hull, Predescu and White (2005), p. 5.

[19] There is an exception concerning the up–front fee of the equity tranche which results in a different default time risk profile.

[20] See Amato und Gyntelberg (2005).

source of risk comes from the portfolio, the expected cash flows of the participating credit risk transfer parties can thus simply be discounted at the risk–free rate. Due to this assumption, the formulas for pricing synthetic CDOs do not differentiate between funded or unfunded transactions and the valuation can be set up similar to plain-vanilla CDSs.

The portfolio under focus is set up under the following conventional homogeneity assumptions to simplify computations:

- Independence of the firm's credit risk and the default–free interest rates under the risk–neutral measure.
- The correlation coefficient $\rho_{i,j}$ for one year between each pair of random variables X_i, X_j is the the same for any two firms $i \neq j$ and will be indicated as ρ. In the employed factor model this corresponds to $a = \sqrt{\rho}$.
- The default intensity λ generating the marginal default distributions is the same for all obligors.
- The loss given default – or correspondingly the recovery rate – is deterministic and the same for all companies.
- The initial notional of each credit in the portfolio is the same.

4.1 Intensity Calibration by CDS Market Quotes

Before we consider the pricing of synthetic CDOs, we present a simple method to extract a representative marginal default intensity λ from market quotes. We consider a CDS contract initiated at time 0 with maturity T. Let the premium payment dates be denoted as $0 = t_0 < t_1 < \ldots < t_J = T$. The CDS has notional N while s_{CDS} denotes the annual CDS spread. In order to determine the fair spread, the discounted premium and protection legs have to be computed by setting them equal under risk–neutral expectations.

In the case of default before maturity, the protection seller has to make compensatory payments amounting to $(1-R)N$, where R is the recovery–of–face–value rate at default time τ. Today's expected value of this payment is

$$\mathbb{E}PV_{prot}(0) = \mathbb{E}\left[B(0,\tau)\mathbf{1}_{\{\tau \leq T\}}(1-R)N\right], \tag{7}$$

where

$$B(0,\tau) = e^{-\int_0^\tau r_s ds} \quad \text{and} \quad E\left[\mathbf{1}_{\{\tau \leq T\}}\right] = Q(0,\tau) = 1 - e^{-\int_0^\tau \lambda_s ds}.$$

$PV_{prot}(0)$ represents the expected present value of the compensatory payments and $B(0,\tau)$ is the risk–neutral discount factor for time τ. In order to discretize this equation for the simple extraction procedure of λ, we have to make a transformation for payoffs at default first, since τ is unknown. Equation (7) thus becomes:

$$\begin{aligned}\mathbb{E}PV_{prot}(0) &= \mathbb{E}\left[B(0,\tau)\ \mathbf{1}_{\{\tau \leq T\}}\ (1-R)N\right] \\ &= \int_0^T B(0,t)\ (1-R)N\ dQ(0,t).\end{aligned}$$

These integrals represent the fact that payments are made when losses occur in continuous time. For the implementation, however, we assume that potential defaults can only happen at the premium payment days. So no intermediate defaults are admitted by the model. We then get as an approximation

$$\int_0^T B(0,t)\ (1-R)N\ dQ(0,t) \approx \sum_{j=1}^{J} B(0,t_j)\ (1-R)N\ \left[Q(0,t_j) - Q(0,t_{j-1})\right].$$

The valuation of the premium leg is slightly more complicated when accrued premiums are considered. At each CDS premium payment date the protection buyer has to make a payment if no default has occurred until that date. If a default event occurs, the protection buyer has to pay the fraction of the premium that has accrued since the last premium payment date at that specific default time τ. For simplification, accrued premiums are not considered and δ will be the accrual factor representing the constant 3–month period between premium dates. The following equation expresses the expectation of the present value of premium payments made:

$$\mathbb{E}PV_{prem}(0) = \mathbb{E}\left[\sum_{j=1}^{J} B(0,t_j)\mathbf{1}_{\{\tau>t_j\}} s_{CDS} N\delta\right] = s_{CDS} N\delta \sum_{j=1}^{J} B(0,t_j)\ (1 - Q(0,t_j))\ .$$

As a final result, the fair spread with deterministic recovery rate and constant deterministic intensity can be computed in the following way, as default is restricted to happen only at premium payment dates:

$$\begin{aligned}
s_{CDS} &= \frac{(1-R)N\sum_{j=1}^{J} B(0,t_j)\quad [Q(0,t_j) - Q(0,t_{j-1})]}{N\delta\sum_{j=1}^{J} B(0,t_j)\ (1-Q(0,t_j))} \\
&= (1-R)\ \frac{\sum_{j=1}^{J} B(0,t_j)\left(e^{-\int_0^{t_j}\lambda_s ds} - e^{-\int_0^{t_{j-1}}\lambda_s ds}\right)}{\sum_{j=1}^{J} B(0,t_j)\ e^{-\int_0^{t_j}\lambda_s ds}\ \delta} \\
&\overset{\lambda_s=\lambda}{=} (1-R)\ \frac{\sum_{j=1}^{J} B(0,t_j)\ e^{-\lambda t_j}\left(e^{-\lambda(t_{j-1}-t_j)} - 1\right)}{\sum_{j=1}^{J} B(0,t_j)\ (e^{-\lambda t_j})\ \delta} \\
&= (1-R)\ \frac{\sum_{j=1}^{J} B(0,t_j)\ e^{-\lambda t_j}\left(e^{-\lambda\delta} - 1\right)}{\sum_{j=1}^{J} B(0,t_j)\ (e^{-\lambda t_j})\ \delta} \\
&= (1-R)\ \frac{\left(e^{\lambda\delta} - 1\right)}{\delta}.
\end{aligned}$$

This expression can be inverted to derive the deterministic default intensity as a function of the CDS index spread:

$$\lambda = \frac{1}{\delta}\ \ln\left(\frac{s_{index}\delta}{1-R} + 1\right).$$

The resulting λ is utilized to compute the representative marginal default distributions in the exponential model for all companies in the reference portfolio.

4.2 The Valuation of Index Tranches

A CDS index contract is insurance that covers default or other credit events as specified in the contract for a pool of reference entities in the index.[21] The buyer of protection on the index is obligated to pay the same premium on all the reference entities in the index (called the fixed rate) for as long as they have not been removed due to an event.

Once created, the components of the index are unchanged over the contract's tenor. The payment or premium payment dates are the standard CDS dates: 20^{th} of March, June, September, and December. Each index consists of the 125 most important CDSs. Index tranches are standardized regarding the composition of the pool and the tranche notionals. Quotations of standardized tranches reflect a high degree of liquidity and market forces are pushing towards two extremes: standardized index tranches with great liquidity used in active trading and bespoke tranches which are designed for buy–and–hold purposes that can be evaluated relative to an index.[22] The premiums on the standardized mezzanine and senior tranches are the spread with no upfront payment. By contrast, there exists an upfront payment for the equity tranche as a percentage of tranche notional, in addition to paying a running spread premium of 500 basis points.

We will now describe the standard market model that is used to compute prices. Let t denote the time passed since the CDO transaction was started, T the maturity of the CDO, N_{total} the initial portfolio value, and $Z_{total}(t)$ the percentage loss in the portfolio value at time t. The total loss at t then is $Z_{total}(t)N_{total}$. The loss suffered by the holder of tranche l from time 0 to t is a percentage $Z_l(t)$ of the portfolio notional value N_{total}

$$Z_l(t) = \min\Big[\max\big(Z_{total}(t) - K_l^L, 0\big), K_l^U - K_l^L\Big].$$

[21] See Amato und Gyntelberg (2005), p. 74.
[22] See Amato und Gyntelberg (2005), p. 77, footnote 10.

We consider a transaction initiated at time 0 with maturity T. Again, let the premium payment dates be denoted as $0 = t_0 < t_1 < \ldots < t_J = T$. The premium payment dates are on a quarterly basis, so $\delta = 0.25$ years.

In a Monte–Carlo simulation, for each generation of a future scenario, the respective losses of each tranche at all specified premium payment dates are stored. After all simulation procedures have been carried out, these values are averaged to obtain the expected percentage tranches losses

$$\mathbb{E}Z_l(t_j), \quad \text{for} \quad j = 0, \ldots, J \quad \text{and} \quad \forall\ l.$$

The expected present value of the protection leg is described by the following formula:

$$\mathbb{E}PV_l^{prot}(0) = \sum_{j=1}^{J} B\left(0, \frac{t_j + t_{j-1}}{2}\right)\Big(\mathbb{E}Z_l(t_j) - \mathbb{E}Z_l(t_{j-1})\Big)N_{total}.$$

The holder of tranche l receives a periodic premium payment with frequency δ years, amounting to $s_l\delta$ times the tranche's outstanding notional $N_l^{out}(t)$. However, the initial tranche notional is stochastically decaying in time induced by tranche losses. At time t_j the outstanding tranche notional is

$$N_l^{out}(t_j) = \Big(K_l^U - K_l^L - Z_l(t_j)\Big)N_{total}.$$

At premium payment dates t_j $(j = 1, \ldots, J)$ the expected average outstanding tranche notionals since the last premium payment dates have to be considered. The outstanding between payment dates t_{j-1} and t_j is simply the average of $N_l^{out}(t_{j-1})$ and $N_l^{out}(t_j)$. It will be denoted as $N_l^{out}(t_{j-1}, t_j)$ and it has to be taken into account that defaults are assumed to occur only at the midpoints between arbitrary premium payment dates. As a result, the expected average outstanding tranche notional between two premium payment dates is assembled in the following way:

$$\mathbb{E}N_l^{out}(t_{j-1}, t_j) = \left[K_l^U - K_l^L - \mathbb{E}Z_l(t_j) + \frac{\mathbb{E}Z_l(t_j) - \mathbb{E}Z_l(t_{j-1})}{2}\right] N_{total}.$$

This equation directly allows for the computation of the expected present value of the premium payments:

$$\mathbb{E}PV_l^{prem}(0) = \sum_{j=1}^{J} B(0, t_j)\ \mathbb{E}[N_l^{out}(t_{j-1}, t_j)]\ s_l\delta. \tag{8}$$

Finally, the equation for the constant over time fair spread s_l of tranche l is:

$$s_l = \frac{\sum_{j=1}^{J} B\left(0, \frac{t_j + t_{j-1}}{2}\right)\Big(\mathbb{E}Z_l(t_j) - \mathbb{E}Z_l(t_{j-1})\Big)N_{total}}{\sum_{j=1}^{J}\ B(0, t_j)\ \mathbb{E}[N_l^{out}(t_{j-1}, t_j)]\ \delta}.$$

There is a different quotation for the equity tranche. The protection seller receives the quoted upfront fee, expressed as a percentage f of the tranche principal, so that the investor purchases the equity tranche at the discount $f(K_{equity}^U - K_{equity}^L)N_{total}$. Additionally, a spread s_{Equity} of 500 basis points per year is paid on the outstanding tranche principal. Note that the overall consequence of this agreement is a different exposure of the equity tranche to default timing. Just as before, this discount is derived by setting equal the expected present values of the premium and the protection legs. Only the premium leg in equation (8) has to be changed to:

$$\mathbb{E}PV_{equity}^{prem}(0) = f(K_{equity}^U - K_{equity}^L)N_{total} \ + \sum_{j=1}^{J} B(0, t_j)\ \mathbb{E}[N_{equity}^{out}(t_{j-1}, t_j)]\ s_{equity}\delta.$$

5 Calibration and Results

For the calibration to the iTraxx IG index we consider the tranche quotes on April 11, 2005. The settlement date of the third series of this index is September 20, 2005 and matures on September 20, 2010. The index CDS spread on April 11, 2005 was 38.81 bps.

There are 125 equally weighted reference entities in the index. Concerning the marginal default distributions and recovery rates, we construct a homogeneous portfolio with the usual assumptions. We use the constant default intensity model to derive the marginal default distributions and assume a constant recovery rate of $R = 40\%$.[23] The applicable risk–free rate for tranches of the Europe–based iTraxx IG is the Euro zero curve.

Conveniently, we calibrate the equity tranche because its pricing is most sensitive to the model parameters. The input parameters are the factor loading a, and the tupel (α, σ) of the standardized STS factor distributions.

In the literature it is often proposed that the calculus–based method of Powell relying on multidimensional direction sets be employed.[24] We consider a version of the intuitive genetic algorithm (GA), instead, and provide additional information about the other tranche quotes in the objective or fitness function to obtain an overall fit with the main focus on the equity tranche quote. Table 2 shows the calibration results for the Gaussian and STS versions of the HPW model.

PLACE TABLE 2 ABOUT HERE

The market quote for the equity tranche is matched exactly by the two competing models. For the other tranche quotes, there is a large gap between market quotes and those quotes produced by the Gaussian version of the HPW model. The version with symmetric STS distribution, however, provides a good fit. Note that the senior tranche with 12% attachment and 22% detachment level is priced much more realistically than the Gaussian version is capable of. There is the same environment of linear correlation provided in our extension but simultaneously, there are additional effects that influence the joint default loss behavior. For example, due to extreme negative outcomes of the systematic factor, there are a large number of joint defaults which are observable more often than in the Gaussian case.

In the empirical fit it can also be observed that our model provides a close match to the rest of the tranches, including a perfect match to the price sensitive equity tranche. This remarkable overall-fit can be interpreted in the following way: The dynamic interplay of the heavy–tailed systematic and idiosyncratic factors results in scenarios that are characterized by groups of firms defaulting jointly in short time horizons. The frequency of occurrence and the number of defaulting firms in these scenarios seem to be adequate to match the cost of protection of all tranches simultaneously.

An example is the extreme negative outcome of the systematic factor: At first, such a scenario almost never occurs in the Gaussian model and second – to further develop this exemplified scenario – there might be a reasonable amount of idiosyncratic heavy-tailed factors with extreme *positive* outcomes, which is almost never displayed by normally distributed idiosyncratic factors. This in turn means that the size of the default cluster due to the systematic impact may be reduced at the same time by some surviving companies due to their extreme positive idiosyncratic factor outcomes. This is just an example of the complex default environment created by our model. It can be stated that there is an adequate implicit micro structure of default scenarios provided by our model as the close match of the model quotes to empirical data shows. This can also be seen in Figure 1. For graphical illustration, it attributes the spreads of tranches 2 to 5 from Table 2 to the tranches' detachment points and then interpolates.

PLACE FIGURE 1 ABOUT HERE

As the spread lines reveal, the market quotes and the spreads given by the STS–HPW model show much more resemblance than the Gaussian HPW alternative. In comparison to the market, the Gaussian model exaggerates the cost of protection for tranches 2 to 4 and underestimates the cost of

[23]This is due to standardized deal conventions.

[24]See for example Press, Teukolsky, Vetterling and Flannery (1992), p. 412.

protection for tranche 5 – the senior tranche. In contrast, the senior tranche is nearly matched in the STS–HPW model and there is also a good overall fit to other tranches provided.

Due to the calibration flexibility of our genetic algorithm, the error weights of a few selected tranches can be increased in the optimization function. To a certain degree, this renders possible the closest matchings of some quotes while other tranche quotes may then exhibit more distant positions to what the market believes is correct in terms of cost of protection.

The Gaussian HPW model was calibrated to the equity tranche with a factor loading $a = \sqrt{0.2278}$. The best parameters in the STS extension of the HPW model are $\alpha = 1.8834$, $\sigma = 0.5873$, and $a = \sqrt{0.2565}$. This is equivalent to truncation levels of around ± 11.0938 for the standardized STS distributions. Note that the resulting STS distribution with the remarkable fit to empirical data is dominated by a stable distribution in the center part with α being well below 2, i.e. a leptokurtic, non-Gaussian. It has to be remarked that the solution space for the genetic algorithm included the Gaussian case which was not chosen in the evolutionary process generated by the genetic algorithm. This suggests that extreme outcomes of systematic and idiosyncratic risk factors are necessary to lead to such empirical results.

Considering the status of indices such as the iTraxx or CDX and their standard tranches as benchmarks for credit risk and correlation, our model provides high economic value: As we use the Black–Cox–type structural model we establish a dynamic relation between the default process and the financial variables of the underlying companies. Additionally, we incorporate empirical findings of asset return distributions such as asymmetry and leptokurtosis. In this way, we extend the approach based solely on linear correlation by the incorporation of extreme events in the systematic risk driver that represent default clustering in times of recession, and more generally, there are extreme positive and negative outcomes of systematic and idiosyncratic risk drivers that create a special, desirable default environment. Empirically, this is essential to take into consideration as institutions and markets are complex feedback–driven systems.

Our empirical results and those of many other researchers show that the Gaussian hypothesis has to be rejected in most cases. The world is not fully informed or acts rationally as suggested by the idea of the *Homo Oeconomicus*. Only in some phases of the market we can assume “normal” behavior or Gaussianity. For these reasons it is important to use heavy–tailed distributions in financial modeling. The STS distribution is a good solution as it reproduces empirical findings of financial asset returns concerning extreme values. Moreover, in contrast to the α-stable class, this distribution class is compatible with the convenient pricing approaches due to the finite–moment generating function. Since our optimization algorithm leaves room for α–parameterizations close to or equal to 2, our model is flexible to cover normal market phases as well as extreme market scenarios.

In future risk management applications, more detailed information about portfolio credit risk could be provided besides basic information sets such as default time distributions or general asset return correlation. A tail index of a heavy-tailed distribution applied in an enhanced version of the standard Gaussian copula model could be such an additional information set.

A further improvement in matching all index tranche quotes simultaneously, even for different maturities, could be accomplished by the use of another empirically meaningful parameter of the STS distribution: β. Incorporating asymmetric STS distributions will lead to different dependence effects in the upper and lower distribution tails. The full potential of STS distributions in terms of reflecting empirical asset return properties such as heavy tails or asymmetry, can easily be exploited in the model presented in this paper. The asymmetric left and right truncation points of standardized STS distributions can be determined by adequate optimization procedures. The additional parameter β can be found together with the other optimization parameters in the evolutionary environment of the genetic algorithm.[25] Since the introduction of heavy tails led to such improvements, the consideration of both asymmetry and heavy tails might already lead to a perfect match of all tranche quotes

[25] An additional parameter will slow down the optimization speed of the genetic algorithm but variation of the evolutionary search might speed the procedure up. Also, further effort could be undertaken with respect to the implementation of the model. Operations to build up the evaluation function of standardized STS distributions and the barrier calibration for a homogeneous portfolio require about 35 seconds. The Monte–Carlo simulation with 10,000 scenarios consumes about 105 seconds. This part leaves room for various performance improvements.

simultaneously. This is critical for efficient pricing and hedging.

There are other reasonable ways to improve the empirical quality of the model:

- Correlation and recovery rates may be modeled as stochastic quantities. The first reflects empirical research showing that default correlations are positively dependent on default rates. The second is empirically intuitive since recovery rates are negatively dependent on default rates.[26]
- A relaxation of the homogeneous portfolio assumption constitutes an additional computational burden. However, it might also lead to a better quote fit since important market information is not simply "averaged out".
- The constant intensity calibration of the marginal default time distributions to the CDS index quote and the interpolation method of the respective zero curve for discounting are approximations. More sophisticated models could be employed, instead.

6 Conclusion

In this paper, we introduced an enhanced version of the HPW model for true integrated pricing and hedging of correlation dependent credit derivatives. Our model generalizes a Black–Cox–type structural model for credit portfolios with respect to the Gaussian distribution assumption. Instead of the normal distribution, we employed the smoothly truncated α–stable (STS) distribution of the factor diffusion processes.

This served as an alternative to the commonly used Gaussian or Student–t distributions. A comparison of the symmetric STS and the Gaussian version of the HPW model showed the remarkable advantage of the STS distribution in pricing standard tranches of CDS indices such as iTraxx. That is, the STS adapted favorably to characteristics observed in financial markets such as asymmetry and heavy–tailedness in a tractable way, as respective parameters are continuous and create an empirically relevant distribution shape. The Gaussian distribution was found to be incapable of handling this. Even the Student–t distribution was found useless since it fails to account for asymmetry and exhibits non-fractional parameters for heavy–tailedness.

Our approach had two advantages over the standard static Gaussian copula model for pricing standard tranches of CDS indices such as the iTraxx. It offered the benefit of describing underlying credit spreads and default behavior dynamically which was based on the Black-Cox approach. Additionally, the STS distribution in this model framework leads to a nearly perfect match with tranche quotes of the iTraxx index. The remarkable empirical fit could be attributed to the properties of the underlying α–stable distribution reflecting the observed leptokurtic and asymmetric behavior of asset returns.

Our model was economically and statistically viable as it combined the ability of the structural model to connect financial and default variables on the one hand with the empirical features of the α–stable distribution on the other. This formed a transparent, comprehensive, and tractable model with appealing theoretical and empirical features. The dynamic interplay of risk factors due to tail-dependence effects lead to default scenarios that were much richer in structure than the framework of linear correlation. The latter environment was also provided by our model but, additionally, there was a special dependence environment among reference entities that captures effects like default clustering in times of recession. This effect could be clearly attributed to the reasonable amount of probability mass in the tails of the STS distribution.

Correlated dynamic fluctuations of credit spreads in the HPW model rendered possible diverse pricing and risk management applications of credit portfolios and certain corresponding derivatives. With the STS extension of the HPW model, we presented a coherent framework for adequate pricing and hedging of many correlation dependent credit derivatives such as synthetic CDOs, options on CDOs or CDO^2. Complex risk management activities based on integrated modeling of credit portfolios and their corresponding risk mitigation instruments can be accomplished. This is due to the full calibration of our convenient Merton–style approach to representative markets. For example, since our precise and dynamic model is superior to the ordinary Gaussian copula it is now possible to efficiently dissolve concentration risk in bond portfolios by the use of correlation dependent credit derivatives

[26] Hull, Predescu and White (2005) successfully extend their model in these directions.

when portfolio and risk transfer instruments are both modeled in our calibrated framework.

The resulting joint evolution of credit spreads was based on the interaction of heavy-tailed distributions in the factor model which was precisely calibrated to the market. This also resulted in a high model quality of credit spread dynamics.

Therefore, since heavy–tailedness and asymmetry play an important role in realistic markets, we strongly prefer distributions such as the STS in financial modeling.

References

Amato, J.D.; Gyntelberg, J. (2005): *CDS index tranches and the pricing of credit risk correlations*; BIS Quarterly Review, March 2005.

Bergström, H. (1952): *On some Expansion of Stable Distributions*; Arkiv for Matematik 2, p. 375-378.

Black, F.; Cox, C. (1976): *Valuing Corporate Securities: Some Effects of Bond Indenture Provisions*; Journal of Finance, Vol. 31(2), p. 351-367.

Burtschell, X.; Gregory, J.; Laurent, J.-P. (2005): *A Comparative Analysis of CDO pricing models*; Working Paper, BNP Paribas.

Chambers, J. M.; Mallows, C. L.; and Stuck, B. W. (1976): *A Method for Simulating Stable Random Variables*; Journal of the American Statistical Association, Vol. 71, p. 340-344.

Committee on the Global Financial System (2005):*Credit Risk Transfer*; Bank of International Settlements, Basel.

Duffie, D. (2004): *Time to Adapt Copula Methods for Modelling Credit Risk Correlation*; RISK, April, p. 77.

DuMouchel, W.H. (1971): *Stable Distributions in Statistical Inference*; Dissertation, Department of Statistics, Yale University.

Gordy, M.; Heitfield, E. (2001): *Of Moody's and Merton: a structural model of bond rating transitions*; download: http://www.bis.org/bcbs/events/oslo/gordyheit.pdf.

Höchstötter, M.; Rachev, S.T., Fabozzi, F.J. (2005): *Distributional Analysis of the Stocks Comprising the DAX 30*; download: http://www.statistik.uni-karlsruhe.de/292.php

Hull, J.; Predescu, M.; White A. (2005): *The Valuation of Correlation-Dependent Credit*; Probability and Mathematical Statistics, Vol. 25, p. 363-383.

Hull, J.; Predescu, M.; White A. (2005): *The Valuation of Correlation-Dependent Credit Derivatives Using a Structural Model*; Working paper, University of Toronto.

Hull, J.; White, A. (2001): *Valuing Credit Default Swaps II: Modeling Default Correlations*; Journal of Derivatives, Vol. 8(3), p. 12-22.

Hull, J.; White, A. (2004): *Valuation of a CDO and n-th-to-default CDS without Monte Carlo Simulation*; Journal of Derivatives, Vol. 12(2), p. 8-23.

Menn, C.; Rachev, S. T. (2004a): *A New Class of Probability Distributions and its Application to Finance*; download: http://www.statistik.uni-karlsruhe.de/292.php

Menn, C.; Rachev, S. T. (2004b): *Calibrated FFT-Based Density Approximations for α-Stable Distributions*; forthcoming: Computational Statistics and Data Analysis, download: http://www.statistik.uni-karlsruhe.de/292.php

Menn, C.; Rachev, S. T. (2005a): *A GARCH Option Pricing Model with α-Stable Innovations*; European Journal of Operations Research, Vol. 163(1), p. 201-209.

Menn, C.; Rachev, S. T. (2005b): *Smoothly Truncated Stable Distributions, GARCH-Models, and Option Pricing*; technical report, download: http://www.statistik.uni-karlsruhe.de/292.php

Nolan, J.P. (1997): *Numerical calculation of stable densities and distribution functions*; Communications in Statistics - Stochastic Models 13, p. 759-774.

Press, W. H.; Teukolsky, S. A.; Vetterling, W. T. and Flannery, B. P. (1992): *Numerical Recipes in C: The Art of Scientific Computing*; 2nd edition, Cambridge University Press. North Holland.

Rachev, S.T.; Mittnik, S. (2000): *Stable Paretian Models in Finance*; Series in Financial Economics and Quantitative Analysis. John Wiley & Sons, Inc., New York, London.

Rachev, S., Stoyanov, S., Biglova, S., Fabozzi, F. (2005): *An Empirical examination of daily stock return distributions for U.S. stocks*, in: *Data Analysis and Decision Support, Springer Series in Studies in Classification, Data Analysis, and Knowledge Organization*, Daniel Baier, Reinhold Decker, and Lars Schmidt-Thieme (eds.), Berlin: Springer-Verlag, p. 269-281, 2005.

Racheva-Iotova, B.; Stoyanov, S. (2004): *Univariate stable laws in the field of finance - approximation of density and distribution functions*; Journal of Concrete and Applicable Mathematics 2.

Samorodnitsky, G.; Taqqu, M.S. (1994): *Stable Non-Gaussian Random Processes*; Stochastic Modeling. Chapman & Hall, London, New York.

Wang, D.; Rachev, S.T., Fabozzi, F.J. (2007) *Pricing of Credit Default Index Swap Tranches with One-Factor Heavy-Tailed Copula Models*; download: http://www.statistik.uni-karlsruhe.de/292.php

Zolotarev, V.M. (1966): *On representation of stable laws by integrals*; Selected Translations in Mathematical Statistics and Probability 6, 84-88.

Index tranche no.	Attachment in %	Detachment in %
iTraxx IG		
1	0	3
2	3	6
3	6	9
4	9	12
5	12	22
CDX IG		
1	0	3
2	3	7
3	7	10
4	10	15
5	15	30

Table 1: iTraxx IG Index Tranches and CDX IG Index Tranches

iTraxx IG 5–year	0-3%	3-6%	6-9%	9-12%	12-22%	Sum of Errors
Market Quotes	24.70%	160	49	22.5	13.75	
HPW model: Gaussian	24.70%	246.79	80.75	29.24	5.55	
		(0.54)	(0.65)	(0.30)	(0.60)	2.09
HPW model: STS	24.70%	158.10	55.93	29.59	15.82	
		(0.01)	(0.14)	(0.32)	(0.15)	0.62

Table 2: Spread predictions of iTraxx tranches in the Gaussian and STS versions of the HPW model. The numbers in parentheses represent the relative errors referencing to the market quotes.

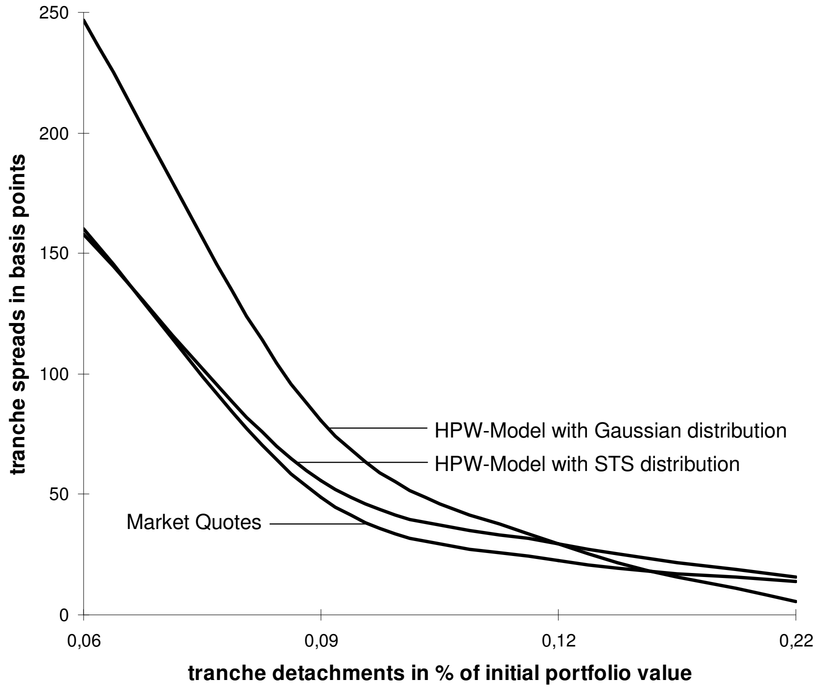

Figure 1: Graphical illustration of Table 2, visualizing tranche spreads as a function of detachments for the two models and the market.