Universität Karlsruhe (TH)

Institute for Statistics and Mathematical Economic Theory

Chair of Statistics, Econometrics and Mathematical Finance
Prof. Dr. S.T. Rachev

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

Author: Jochen Papenbrock
Voltastr. 40
60486 Frankfurt am Main
Germany
E-Mail: J.Papenbrock@gmx.de

Frankfurt, November 1, 2006

Table of Contents

List of Tables

List of Figures

1 Acknowledgements

First of all, I would like to thank Prof. Dr. S.T. Rachev. His classes have always been an inspiring and efficient form of knowledge transfer. I was impressed by the international research environment and the friendly atmosphere provided by him and his chair members.
Second, I would like to thank my mentor and tutor Dr. Markus Höchstötter for his outstanding advisory qualities and his motivating support.
Since a lot of this work is based on the MATLAB code of Dr. Christian Menn I would like to thank him for providing the lines and also giving advice.
For the first part of the whole project I would like to thank Christoph Wolff and Matthias Küchler from Dresdner Bank AG, Frankfurt. They helped me set up the basic model upon which the later extensions at the Chair of Statistics, Econometrics and Mathematical Finance could be built. I appreciate their model advice and support in the programming language C++.
Finally, I would like to thank Prof. Dr. Marliese Uhrig-Homburg and Michael Kunisch both from the Institute for Finance, Banking, and Insurance. This is where I was first educated in financial engineering and especially in the field of correlation dependent credit derivatives.

In the private area I would like to thank my parents and my sister with her family. Without their help and everlasting support this would not have been possible at all.
I would like to mention Daniel Stock for his mental and technical support. It has always been a pleasure to discuss and study with him.
Finally I would like to mention Alexander Deierling for his all-embracing longtime support which could even be maintained across continents.

2 Introduction

The development of credit risk modeling in financial institutions has improved rapidly due to unusually high corporate default rates during 2001 and 2002 and upcoming quantitatively sophisticated bank capital regulations. Also, financial institutions realize their expertise in selection and management of credit portfolios including the efficient transfer of portfolio credit risk to capital markets. There has been a clear trend in the credit derivatives market, namely the growth of more complex and model-driven trading strategies and credit risk transfer activities.
An example is the emergence of synthetic Collaterized Debt Obligations (CDOs) which transfer the risk of a pool of single-name Credit Default Swaps (CDS). The initiator of such a deal may sell credit risk protection on several names by a corresponding number of CDS contracts and hedge that risk by purchasing protection via CDO loss tranches which leads to the term synthetic CDO. This transfer can either be a funded cash market security or an unfunded swap contract.

For issuers of synthetic CDOs, however, there is the difficulty of placing certain parts of the capital structure in the market. This may have led to further developments in the CDO market like 'single-tranche CDOs' (STCDOs). These instruments allow for a deal to be customized according to the CDO investors' needs with respect to the reference portfolio as well as the specific share of the overall loss distribution. Another development in the CDO market involves what is known as 'single tranche trading'. Dealers manage their 'short' position in the issued CDO tranche without actually acquiring the credit risk associated with the entire pool. This approach is known as 'delta-hedging' because of its similarity to hedging an option position.

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 like iTraxx or CDX. These CDS indices provide standard benchmarks against which other more customized pools of exposures can be assessed. Also, they serve as building blocks for derivatives 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 Gaussian copula within the structural approach by Merton (1974). 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 produces a single implied parameter to summarize the average correlation among various borrowers' default times. However, the fact that index tranches are quoted very frequently and with relatively narrow bid-ask spreads has uncovered several shortcomings of the existing pricing models for CDOs. Especially, the Gaussian copula model does not fit market prices very well as was reported from various sources. The model underperformance can be observed in the pronounced correlation smile when implied CDO tranche correlations are plotted as a function of tranche attachment points.
A promising modification might be the employment of different distribution assumptions taking 'statistical fingerprints' of empirical asset returns into account. An example is the model by Hull and White (2004) whose extension is known as 'Double Student-t copula' (Student-t distributed systematic and idiosyncratic factors). This approach is able to reproduce market quotes of standardized index tranches in a much better way than the factor Gaussian copula. Analyzing the interplay of the factor distributions reveals that the occurrence of extreme events in the distribution tails seems to be responsible for the more adequate market fit due to more realistic joint default behavior. But according to Duffie (2004) the latest focus on portfolio credit risk models that are designed to appropriately capture correlations in default timing is not sufficient. Many pricing and risk management applications call for modeling correlated default times as well as correlated changes in credit spreads. So effort should be directed to models that both reproduce market quotes of standardized index tranches like iTraxx or CDX in a much better way than the Gaussian copula model and simultaneously incorporate the underlying dynamic evolution of credit spreads. This is helpful in dynamically delta-hedge single tranches and price more exotic derivatives like options on tranches or forward starting CDOs, for example.

A first step in this direction is the model by Hull, Predescu, White (2005). Basically, it can be regarded as the dynamic version of the Merton approach allowing for intermediate defaults in the spirit of Black und Cox (1976). It consists of a Monte-Carlo simulation of properly discretized multivariate stochastic processes in the form of correlated geometric Brownian motions, realized by a similar decomposition as in the factor Gaussian copula approach. The discrete-time intermediate default barriers can be calibrated to the individual credit spread or CDS spread curves according to the algorithm of Hull and White (2001).
When the CDS spread of the iTraxx or CDX index is assumed to be representative for the single-name CDSs in the portfolio, this quote can be translated into a deterministic intensity in the exponential model to simplify further calculations. As a consequence,

a representative default time distribution can be generated to compute representative intermediate default barriers.
By this dynamic construction of intermediate defaults, the distance-to-default of the underlying portfolio names can be measured and transformed into the joint evolution of credit spreads during the life of CDO transaction. According to the authors, the resulting dynamic model is mathematically equivalent to the factor Gaussian copula model. This is due to the fact that the only source of dependence is linear correlation. However, it has to be remarked that both models - the factor Gaussian copula as well as the basic HPW model - are empirically dissatisfying as can be stated by the pronounced correlation smile.

In our approach we suggest a different distribution family known as α-stable or stable Paretian within the modeling framework of Hull, Predescu, White (2005). This class was first suggested for financial applications by Benoir Mandelbrot in the early sixties. It provides numerous advantageous statistical features, outstanding market fits and favorable modeling properties. This was impressively underlined in the book 'Stable Paretian Models in Finance' by Rachev and Mittnik (2000).
The special case of the characteristic stable distribution parameter $\alpha = 2$ simply resembles the Gaussian case, but for $\alpha < 2$, stable distributions exhibit increasing tail-heaviness the smaller α gets. From empirical studies, an α in the range from 1.6 to 1.9 seems to be adequate for financial returns but this results in infinite moments of order > 1 which basically means that the variance does not exist. As a consequence, the whole concept of covariance/correlation breaks down, and furthermore, infinite second order moments result in infinite exponential first order moments which basically means that there is no mean for the Merton-style firm value.
Fortunately, according to Menn and Rachev (2005b) there is a concept called 'smooth truncation' that fulfills requirements to apply α-stable distributions in the framework of Black/Scholes/Merton. Loosely spoken, the authors construct a composed distribution that preserves α-stable distributions in the center part and smoothly replaces the tails - in a mathematical sense - by truncated normal distributions. The result is a distribution family called 'Smoothly Truncated Stable Distributions' (STS). It has excellent properties for financial applications: finite moments of arbitrary order, support of the concept of covariance/correlation and placement of far more probability mass in the tails of the distribution than the Gaussian and even the Student-t distribution which was stated by tail probability studies and QQ plots.
These promising features prompted us to implement the STS family in order to replace the normal distribution assumption in the HPW model. For this reason, an algorithm for evaluations of α-stable distributions had to be employed since there is no closed form

expression of α-stable distributions except for a few cases like the Gaussian. There are several approaches to this problem ranging from integral representations for the density function to Fast Fourier Transforms (FFT) regarding the characteristic function. We chose to implement the latter one according to the approach by Menn and Rachev (2004b): 'Calibrated FFT-based Density Approximations for α-Stable Distributions', since it affords a good trade-off between speed and accuracy. Their Simpson rule based Fourier transformation provides relatively high accuracy in the center part for approximations of α-stable distributions but the tail areas have to be approximated by Bergström series expansions. The authors apply a calibration method for a grid of α-stable parameter combinations to find the optimal positions of the splice points between the Fourier transformation area and the series expansion area. In order to optimize accuracy, they benchmark their procedure with respect to a freely-available high-accuracy version of the integral representation approach by Nolan (1997).
However, for the HPW model architecture it is necessary to evaluate the cumulative distribution function so the generated density points are simply interpolated by cubic splines for integration purposes. Finally, a simulation of α-stable distributions is realized by the efficient algorithm of Chambers et al. which can be regarded as the generalization of the famous Box/Müller method for the Gaussian case.

At this point there remains the smooth truncation of the α-stable distribution in the tails. In order to replace the normal distribution in the HPW model by STS distributions, some normalization arrangements have to be made: the mean of the composed distribution has to be zero and the variance has to be one. This can be accomplished by a proper choice of the cut-off points defining the truncation position in order to have a standardized STS distribution with the following three parameters:

- the well-known α with a slightly different meaning in STS distributions,[1]
- a parameter σ that can be interpreted as a measure of how much of the composed distribution is α-stable regarding the center part of the composed distribution and how much is normal in the tails,
- and finally the asymmetry parameter β which is inherited from α-stable distributions.[2]

The sampling of STS distributions is conducted by a combination of the Chambers/Mallows/Stuck and the Box/Müller method.

[1] This is due to the fact that α still drives the tail-heaviness but the heavier the tails the more to the center the truncation has to be accomplished in order to standardize the STS distribution.

[2] The meaning of β is slightly different in STS distributions in comparison to plain α-stable distributions. Once again, this is due to the fact that the choice of β influences the position of truncation in the STS standardization procedure.

It has to be remarked that the HPW model and the new extension with STS distributions are able to deal with a completely heterogeneous portfolio regarding correlations, LGDs, and credit spread or CDS spread curves. However, due to simplifications we restrict ourselves to a completely homogeneous portfolio with a representative CDS spread curve, a representative deterministic recovery rate and an average correlation parameter. The free parameters of the STS distribution plus the average correlation parameter are calibrated to the five iTraxx tranche quotes simultaneously after the intermediate default barriers have been calibrated to the CDS spread curve. This is possible due to the strict separation of the distribution assumption and the barrier calibration being based on quantiles.

For calibration purposes we employ an intuitive version of a genetic/evolutionary algorithm as this optimization technique requires no gradient computations. It resembles a heuristic search procedure according to an evolutionary concept known from Darwin's survival-of-the-fittest theory. It comprises concepts such as natural selection, sexual selection and mutation.
There is an evenly spaced initial grid of free parameter choices and for each combination a single Monte-Carlo simulation of the extended HPW model is carried out to compute a fitness measure that summarizes the ability to match all tranche quotes simultaneously. Due to its sensitivity of parameter choices, special weight is put on the first tranche. In genetic algorithms, combinations with low fitness have to leave the population and better solutions in terms of fitness are coupled. Additionally, arbitrary solutions are mutated in that one parameter is changed a little.

After the calibration procedure, we could observe that our STS extension of the HPW model does not only produce lower errors in comparison to the basic HPW model but there also is a good fit to all tranche quotes simultaneously without any outliers. This leads to the assumption that the interplay of the extreme value STS factor distributions seems to be much more appropriate than the Gaussian distribution in the HPW model.
Simultaneously, our approach preserves the architectural advantages of the basic HPW model concerning the dynamic evolution of credit spreads. Due to the fact that the default process underlying the joint credit spread movement leads to the good fit of market data, it can be assumed that the credit spread dynamics are more realistic in terms of distance-to-default measures and rating migrations than in the Gaussian case.

The diploma thesis has the following structure: chapter 3 gives an overview of the

latest developments of the portfolio credit derivative business with emphasis put on the economics of CDOs. Chapter 4 refers to the factor Gaussian copula model which is the standard market approach. In chapter 5 we describe the Hull/Predescu/White model which is the basis of our extension. Chapter 6 gives an overview of the properties of α-stable distributions, outlines density approximations and simulation procedures, and finally closes with an explanation of the construction and standardization procedure of STS distributions. The content of chapter 7 is the valuation of synthetic CDOs based on the premium/protection leg concept borrowed from insurance business. Finally, chapter 8 presents the model outcomes after the calibration procedure in comparison to market quotes. Also, it delivers the best choice of free parameters and discusses further improvements in modeling and implementation.

3 The Portfolio Credit Derivative Business

Modern financial institutions that are involved in lending business face several sources of complex risks. An example is the emergence of a new generation of structured credit products that are exposed to the joint performance of multiple credit entities.
While only a few years ago the only possibility to manage the credit risk of a large bank was based on the acceptance/rejection of a new borrower, now credit risk can be managed directly by the use of (portfolio) credit derivatives and securitization with a variety of collateral assets.[3]
During the normal course of lending business, the arrival of a certain number of defaults is to be expected. But major risks arise if either the number of defaults exceeds expectations or when the number of defaults is due to expectation, but credit events tend to cluster, so there are several defaults occurring closely after each other.[4] In order to manage this risk, a number of new financial instruments have been introduced which are explicitly designed to trade and manage the risk of portfolio default dependencies. The most prominent representatives are ranked basket derivatives and credit portfolio derivatives on a percentile basis like CDOs.
Due to the introduction of these instruments, a transformation has taken place from credit management as a passive measurement and monitoring function into the active management of the credit exposure of a bank in order to utilize the new possibilities to optimize the risk/return profile of the credit book. So there is a need for quantitative models that incorporate more realistic and convenient methods for quantifying correlations of credit risks across borrowers. Current approaches try to incorporate both correlated default times and correlated fluctuations of credit spreads to effectively manage the risk of credit portfolios, as well as price and dynamically hedge credit portfolio derivatives of different degrees of structural complexities.[5]

3.1 Modeling Correlation of Default-Timing

The development of credit risk management is fueled by several technologies for miscellaneous credit risk transfer activities. Probably the most important credit derivative instrument is the credit default swap (CDS), in which one party (the 'protection seller') acquires the credit risk associated with a specific reference entity over a fixed time horizon in exchange for a fee from the counterparty (the 'protection buyer'). CDSs are used for hedging credit risk and serve as building blocks in creating more complex structured

[3] See Rogge and Schönbucher (2003), p.2.
[4] See Rogge and Schönbucher (2003), p.1.
[5] See for example Duffie (2004).

products.[6]

A second important credit derivative instrument is the synthetic CDO, in which the credit risk of a portfolio of single-name CDSs written on the portfolio names is transferred.[7] So in synthetic CDOs, the originator may gain exposure to the credit risk of a variety of names in the market by selling protection on numerous entities via a corresponding number of single-name CDSs. This realizes an exposure to a variety of names. The initiator hedges that risk by purchasing protection via CDO loss tranches which leads to the term synthetic CDO. This transfer can either be a funded cash market security or an unfunded swap contract. In contrast, cash CDO tranches are cash market securities collateralized by loans, bonds and other debt-related products. The technique of tranching in synthetic CDO transactions means that the losses associated with the portfolio of exposures are allocated separately to individual tranches. The allocation mechanism depends on priority rules established at the initiation of the CDO. These rules are simplified in comparison to complex subordination schemes to be found in cash CDOs. Synthetic CDOs rather exhibit similarities to percentile basket credit derivatives.

The riskiest tranche, which is the first to absorb any losses, is the 'equity', 'first-loss', or 'junior' tranche. At the other extreme, there are 'senior' and 'super-senior' tranches. These will only be hit by losses after subordinated tranches have absorbed their maximum loss. In between are the 'mezzanine' tranches. The ability to construct a CDO synthetically enables this technology to be applied to any set of exposures whose credit risk can be transferred via the CDS market.

The first generation of synthetic CDOs involved the issuance of tranches representing the full capital structure of the securitization.[8] This means that there is an equity tranche (e.g., absorbing the first 3% of losses of the portfolio notional amount), a mezzanine tranche (e.g., absorbing losses between 3% and 7%), and senior and super-senior tranches (e.g., absorbing losses between 7% and 100% of the portfolio notional amount).

For issuers of synthetic CDOs, however, there is the difficulty of placing certain parts of the capital structure, for example the high-risk equity tranche or a large super-senior tranche, in the market. This may have led to further developments in the CDO market like 'single-tranche CDOs' (STCDOs). These instruments allow for a deal to be customized according to the CDO investors' needs. Investors may select all aspects of the reference portfolio as well as the specific portion of the loss distribution to which they wish to be exposed to. If CDO issuers themselves have acquired the credit risk associated with the entire pool of exposures, this implies that they retain those por-

[6] See Committee on the Global Financial System (2005), p. 14.
[7] See Committee on the Global Financial System (2005), footnote 2, p. 13.
[8] See Committee on the Global Financial System (2005), p. 15.

tions of the capital structure that are not issued.
Another development in the CDO market involves what is known as 'single tranche trading'.[9] Dealers manage their 'short' position in the issued CDO tranche without actually acquiring the credit risk associated with the entire pool. This approach is known as 'delta-hedging' because of its similarity to hedging an option position.[10] A portfolio and a tranche are defined and the buyer and the seller of protection agree to exchange the cash flows that would have been applicable as if a synthetic CDO had been set up. So in these kind of deals, the underlying portfolio of CDSs is never created and merely represents a synthetic reference portfolio used to calculate corresponding cash flows.

From an economic perspective, tranching of credit portfolios is appealing, because it allows the credit risk associated with a pool of exposures to be divided up and allocated to parties based on their underlying risk preferences. CDO structures thus create custom exposures that investors desire and cannot achieve in any other way. These custom exposures fit into investors' various risk appetites and capital constraints. For example, some investors are more efficient holders of speculative-grade assets and some have a comparative advantage holding investment-grade assets. Tranches are sold to investors most suited to hold that characteristic risk.[11]

From a disclosure perspective, synthetic CDOs and other tranched credit risk products are challenging, because notional amounts are not a sufficient measure of risk. Several innovations and developments are useful to assess these risks in a better way. Among them is the market of single-name CDSs that has gained a significant amount of liquidity in recent years. In terms of outstanding notional, the market represents about 85% of the credit derivatives market, which has a total outstanding notional in excess of $ 4000 billion.[12] The CDS market is most liquid for CDS contracts with 5-year maturities. There is an increasing effort by dealers to build more continuous credit curves up to ten year maturities.[13]
There are two main reasons why CDS contracts are more liquid than corporate bonds.[14] The first is due to standardization because definitions like the one of credit events, for example, are clearly defined in the ISDA credit derivatives definition.[15] Second, CDS

[9] See Hull, Predescu, White (2005), p.3 and p.5.
[10] See Committee on the Global Financial System (2005), p. 16
[11] See Lucas and Sam (2001), p.11.
[12] See Elizalde (2005a), p. 1
[13] See Committee on the Global Financial System (2005), p. 15.
[14] See Amato und Gyntelberg (2005), p. 74.
[15] Credit events that trigger payment to the protection buyer include bankruptcy, failure to pay, repudiation and material restructuring of debt. Payoffs can either be settled by cash or in physical form.

contracts allow market participants to go long credit risk without an initial cash payment, as well as sell credit risk more efficiently than with corporate bonds. Also, the direct sale of loans and bonds may sometimes comprise client relationships or secrecy, or can be costly because of contractual restrictions on transferring the underlying names.[16]

By means of single-name CDSs, a reliable marking-to-market of individual credit risk becomes possible. This means, that the market credit risk is now measurable, and should therefore be managed. Beside the risk of outright defaults, credit risk representatives like credit spreads include fluctuations in response to market conditions. Credit spreads and CDS spreads contain the market's opinion on the default risk of the obligor under consideration and they provide a new objective, market-based early-warning instrument for changes in the default risk of an obligor.[17]

The observed growth and liquidity of the CDS market has fueled the creation of CDS indices. In this way, broad-based credit risk can be traded and hedged by means of an index CDS being representative for the underlying contracts. Also, indices provide a standard benchmark against which other more customized pools of exposures can be assessed. Finally, indices can be used as building blocks for constructing other products like CDS index tranches, as will be outlined.[18]

The liquidity of CDS index contracts is primarily enhanced by the liquid market of single-name CDS. Additionally, there exists a group of dealers committed to market making. Also, the index composition plays an important role in acceptance and associated liquidity of a CDS index.[19]

Two important reference indices of CDS portfolios have been created: DJ iTraxx for Europe and Asia and DJ CDX for North America and emerging markets. They are gaining importance due to clear geographical focus, relatively stable sector-rating compositions, standardized maturities for each index, and contract types in funded and unfunded form.

One of the most significant developments in financial markets in recent years resulting from the introduction of CDS indices has been the creation of liquid instruments that allow for the trading of credit risk correlations.[20] Prime among these instruments are standardized CDS index tranches. Broadly put, index tranches give investors - i.e. sellers of credit protection - the opportunity to take on exposures to specific segments of the loss distribution of the portfolio of CDSs comprising the index.

[16] See Duffie und Gârleanu (2001), p. 4.
[17] See Schönbucher (2005), p.2.
[18] See Committee on the Global Financial System (2005), p. 15.
[19] See Amato und Gyntelberg (2005), p. 75.
[20] See Amato und Gyntelberg (2005), p.73.

Each tranche has a different sensitivity to credit risk correlations among entities in the index and a tranche's notional size is characterized by standardized attachment and detachment points, which mark certain percentiles of the portfolio loss. Such a tranche can be regarded as an option - or a combination of them - on the portfolio loss which resembles the underlying and the attachment/detachment points function as strikes.

The standardization of index tranches may prove to be a significant step further towards more complete markets. The emergence of index tranches therefore fills a gap in the ability of the markets to transfer certain types of portfolio credit risk across individuals and institutions. The liquid markets for index-CDSs and for STCDOs on an index portfolio render possible the marking-to-market of portfolio credit risk correlations. By means of a standard model, the competitive quotes of index tranches are translated into so-called 'implied correlations'. They are extracted according to techniques similar to the concept of implied volatility for plain vanilla options. In particular, regarding the idea that correlation within a copula model can be seen as the volatility in a standard Black/Scholes option framework, it is straight forward to calibrate smile and skew.

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 relatively 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. These individual default variables of the homogeneous portfolio are dependent on one systematic risk factor which results both in a complexity reduction due to the factor-model and also in conditionally independent defaults, whose properties simplify computations, as will be outlined at a later stage.

Implied tranche correlations can then be extracted from appropriate market data with this standard model to communicate market prices. Market participants interpolate between implied correlations when pricing customized deals that are not actively traded.[21] But the fact that index tranches are quoted very frequently and with relatively narrow bid-ask spreads has uncovered several shortcomings of the existing pricing models for CDOs.[22] The model underperformance can be observed by the pronounced correlation smile when implied CDO tranche correlations are plotted as a function of tranche attachment points.

[21] See Hull and White (2005).
[22] See Schönbucher (2005).

Many market participants prefer to use base correlations rather than tranche correlations. These quotes represent 'cumulative' tranches since the attachments points of all introduced tranches are set at 0%. The resulting approximate linearity suggests that interpolation between base correlations is accurate and base correlations somehow provide more accurate pricing than tranche correlations. However this is questionable since the valuation of a CDO tranche is highly sensitive to the exact position of points on the base correlation skew.
Hull and White (2005) therefore argue that the market's focus on implied correlations is misplaced. If there were models available to fit all tranche market prices in a better way, the correlation smile would nearly vanish and a consistent economic modeling could be obtained since all tranche spreads were matched simultaneously with the same parameter of linear correlation. Rogge and Schönbucher (2003) additionally remark that a credit risk model that is used for trading must be much more accurate than a model that is just used to assess the overall risk of a portfolio or an institution: Prices must be found for both the bid and the offer side of the market, and these prices cannot be set too conservatively, or there will be no trading. On the other hand, prices that are too aggressive or exhibit any systematic deficiencies will be mercilessly be exploited by market participants.[23]

Besides the Gaussian copula model, there exist numerous approaches characterized by different copula flavors and many parametric degrees of freedom to specify. These models were designed to capture the correlations in default timing in a more realistic way and/or simplify risk assessments by semi-analytic pricing expressions to avoid slowly converging simulation procedures.

The quality of these models is improved by one or more of the following features:

- Introduction of multi-factor models - using a group correlation structure according to region or industry sector with high correlation within a group and low correlation between groups leads to a more realistic dependence structure.

- Large portfolio approximations - assumptions of asymptotic models with homogeneous infinitely large portfolio characteristics simplify analytical derivations of loss distributions, for example.

- Relaxing the restriction of constant correlation and recovery rates, and considering heterogeneous portfolios - different obligors exhibit individual exposure towards systematic risk and in times of recession, correlation seems to increase and

[23] See Rogge and Schönbucher (2003), p. 2.

recovery rates tend to decrease. Authors like Altman et al. (2002), Andersen und Sidenius (2004), Hull and White (2004), and Hull, Predescu, White (2005) take these aspects into consideration. Laurent and Gregory (2003) use the fast Fourier transform method to calculate the conditional loss distributions of each of the companies constituting the reference portfolio. Andersen, Sidenius, Basu (2003) develop an intuitive recursion algorithm that is faster due to the computational burden associated with the evaluation of the characteristic function in the former approach. Hull and White (2004) offer an intuitive, iterative algorithm that is robust and flexible and makes use of probability bucketing for building up the portfolio loss distribution. In comparison to Andersen, Sidenius, Basu (2003) the buckets are not equidistant but can be chosen fine-grained around exhaustion points.

- Different copulas for more realistic default behavior in credit portfolios - in reality, joint defaults are not exclusively driven by linear correlation but also by the occurrence of extreme events which leads to fat tails in the credit loss distribution.[24] It turns out that the double Student-t copula model by Hull and White (2004) with the same heavy tailed distributions for systematic and idiosyncratic risk performs very well in market price fitting. This model will be reviewed in section 4.3.

  Independently, in their comparative analysis of CDO pricing models Burtschell et al. (2005) report that the double Student-t copula model has very good calibration features to the CDO market in comparison to other models like Gaussian, Student-t, stochastic correlation, Clayton and Marshall-Olkin copulas. Hull and White (2004) comment that in this Merton-style default process the different factors 'compete' against each other for extreme outcomes which finally leads to the good calibration features.

  It can be stated that the particular interplay between factors representing different risk sources - and simultaneously allowing for extreme events - results in remarkable market fits. It seems to be a plausible assumption that credit (portfolio) derivative models incorporating observed statistical 'fingerprints' of financial asset returns like heavy tails, skewness and leptokurtosis will in general lead to better market calibration qualities. This is also valid for the model presented in this thesis.

- There exists a new approach by Hull and White (2005) who employ the technique of implied copulas. Their copula model can be regarded as 'perfect' in that it hits

[24] An overview is given by Burtschell et al. (2005), where copulas like Student-t, double-t, Clayton, and Marshall Olkin are considered.

the tranche quotes exactly. The main idea is to use conditional hazard rates. The hazard-rate-path probability distribution is the only information they need about the underlying copula in order to value a CDO tranche or similar instruments. The authors also compute implied hazard rate paths for the Gaussian copula and the double Student-t copula with four degrees of freedom. As expected, the latter is more realistic in that uncertainty about the hazard rate increases with the passage of time. This is another explanation why two heavy tailed distributions for the factors fit the market data more accurately.

The model by Hull and White (2005) has numerous further advantages.[25] Among them is the valuation of CDO^2 (CDO on a CDO) and other transactions where the payoffs depend in a complex way on the number of defaults in one or more portfolios. This is due to the fact that market prices are fit exactly. However, it is not appropriate for some instruments like a one-year option on a five-year CDO. In this example, the transaction depends on the development of credit spreads in the first year. This brings us to the necessity to model the dynamic evolution of credit spreads which will be outlined in the following section.

3.2 Dynamic Evolution of Correlated Credit Spreads

Early industry models were essentially static as they only modeled the default risk over a fixed time horizon and were incapable of capturing the timing risk of defaults, which is an essential risk in all cash-flow based debt securitizations like CDOs. The key contribution to solve this problem was made by Li (2001) who extended the fixed-time Gauss copula model to an arbitrary time-horizon model so that the timing risk of defaults could also be incorporated. The modeling of the dependency between default events up to a fixed time was shifted to the dependency between default times over a certain horizon. This makes it straightforward to calibrate to a set of term structures of survival probabilities. These advantages made the Gaussian Copula model the standard model for pricing of CDOs and basket credit derivatives today.

However, limitations regarding the dynamics of credit spreads apply to the entire class of static factor-based models. However, Rogge and Schönbucher (2003) remark that to allow for dynamic hedging and risk management, a quantitative model must be able to reflect not only the default risk, but also the market's price dynamics accurately and thus capture the full range of credit risk codependencies.

Regarding nowadays' practice to use single-name CDS for the hedging of portfolio credit derivatives, realistic price dynamics for these instruments are needed and re-

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

quire calibration to ensure that the model prices are arbitrage-free with respect to the hedging instruments. They conclude that modern portfolio default risk models need not only capture default dependencies over a time-horizon in a realistic manner, but rather incorporate the dynamics both of the timing of defaults as well as the dynamics of credit spreads.

Schönbucher (2005) remark that one reason for the introduction of CDS indices and the corresponding index tranches was the creation of hedge instruments for the management of the risk of the more exotic portfolio credit derivatives like options or forward contracts of CDO tranches. Sidenius, Piterbarg, Andersen comment that an ideal CDO model incorporates the dynamics of credit spreads into CDO modeling while simultaneously maintaining exact calibration to CDO markets.

However, according to the Bank for International Settlements, such activity reflects a clear trend in the credit derivatives market, namely the growth of more complex and model-driven trading strategies and transaction structures.[26] The pricing and risk management of these more complex products and strategies require reliance on credit risk models and in particular on assumptions about the extent of default correlation between different reference entities. This is reflected in the emergence of what is referred to as 'correlation trading desks'. The correlation desks make markets in these complex products and strategies, while managing the overall risk exposure associated with the dealer's position. This trend encompasses the growth of

- standardized single-tranche CDOs and so-called bespoke single-tranche CDOs,
- other less common products such as ranked basket credit derivatives (first-to-default and n-th-to-default basket CDS) and more exotic portfolio credit derivatives,
- STCDOs with embedded options (options to cancel or extend STCDOs),
- outright options on STCDOs (option to enter a single-tranche swap to leverage the STCDO risk/return profile),
- CDOs using CDO tranches as collateral (CDO^2),
- forward starting STCDOs and many more.[27]

These more exotic credit portfolio derivatives call for modeling correlated default times as well as correlated changes in credit spreads. In most current models, the main focus is on default risk, with little attention paid to the evolution of credit spreads. As an

[26] See Committee on the Global Financial System (2005), p. 16.
[27] See Committee on the Global Financial System (2005), p. 16 and Schönbucher (2005), p. 2.

example, the importance of a proper modeling of the default behavior can be highlighted in the valuation of CDO^2 and other transactions where the payoffs depend on the number of defaults in one or more portfolios in a complex way. An example for the importance of dynamic evolutions of credit spreads is the valuation of a one-year option on a five-year CDO because this depends on credit spreads between years one and five conditional on what we observe happening during the first year.[28]

Another example for the need of dynamic models is buying the CDS index which is useful to gaining exposure to market wide credit spreads. There exist 'options on the CDS index' or so-called 'portfolio credit default swaptions' (i.e. options to buy or sell the index default swap), allowing investors to leverage this exposure and provide a tool for gaining exposure to market wide credit spread volatility.
The impact of the market value of a credit portfolio on a future option exercise date depends on which of the underlying firms, if any, will have defaulted by that date, and also on the credit spreads of the remaining firms on that date.
Summarizing these developments, it can be stated that one reason for the introduction of CDS indices and the corresponding index tranches was the creation of dynamic hedge instruments for the management of the risk of the more exotic portfolio credit derivatives.[29]

In terms of reproducing quotes of standardized traches, a 'spread dynamics model already' meeting the default time modeling requirements quite well is the multivariate hitting time model and its extensions by Hull, Predescu, White (2005) (HPW). It can be described as a dynamic structural model in Merton-style, with an extension in the flavor of Black und Cox (1976), admitting inter temporal defaults whenever the value of the firm assets falls below a certain barrier. The default of an obligor in the reference portfolio is represented by a geometric Brownian diffusion process that is constructed to hit the default barrier during the Monte-Carlo simulation procedure, according to statistical input parameters. The portfolio names' default processes are decomposed into a factor structure for systematic and idiosyncratic risk drivers which directly allows for modeling of credit portfolios. The barriers are calibrated to the marginal default time distributions and with individual recovery assumptions, the valuation of a heterogeneous portfolio can be conducted.
According to Hull, Predescu, White (2005) this model is different from the one-factor Gaussian copula model where the realization of a single systematic factor governs the default environment in all future time periods. In their dynamic approach, however,

[28] See Hull and White (2005), p. 10.
[29] See Schönbucher (2005), p. 2.

default environments change over time. With the simplifying assumption that once a company's assets are less than the barrier they stay in default status, the standard market model is a reasonable approximation to the Gaussian HPW model under the same default correlation structure. This relation is naturally breached for distributions other than Gaussian.
During one generation of a future scenario in the dynamic version, the default process of an individual obligor might exhibit some positive distance to the calibrated default barrier at some point in time during the life of the CDO. It is then possible to compute the remaining probability of default until CDO maturity from that position. This information can, in turn, be used to compute the respective credit spreads. Dynamic re-evaluations for all credit entities during one future scenario generation thus represent the correlated evolution of credit spreads in the portfolio.

Recapitulating, there is a demand for models with high-dimensional input vectors of statistically meaningful parameters that capture the portfolio default process in an economically sound way. Also, some successful structural models decompose central risk drivers into common and individual parts that provide extreme event occurences to incorporate empirical findings of financial asset returns. This approach results in a good fit to the market while simultaneously preserving calibration to the underlying marginal default probability distributions. And finally, the same models should reveal the underlying joint evolution of credit spreads for hedging purposes and for the valuation of more complex credit portfolio risk transfer activities.
Our extension to the HPW model fulfills both requirements. It is outlined in chapters 5 and **??**. The special focus of the rest of this chapter lies on the risks inherent in traditional CDOs, synthetic CDOs, and their derivative instruments.

3.3 The Mechanics, Economics, and Risks of CDOs

The institutional term Asset Backed Securities (ABS) subsumes the abstract class of financial claims towards a reference portfolio of assets.[30] ABS are structured fixed income securities that are backed by financial claims and they are classified by their reference assets. In practice, a special purpose trust is set up which takes title to the assets and the cash flows are passed through to the investors in the form of an asset-backed security. ABS are a modern form of refinancing and risk transfer. The types of assets that can be securitized range from residential mortgages to credit card receivables.
In this context, traditional CDO transactions consist of claims towards the cash flows of a credit portfolio. The term CDO refers to the transaction itself and/or the special

[30] See Burghof et al. (2000), p. 24.

purpose vehicle (SPV) that issues different obligations. CDO transactions deploy the technique of tranching which is based on the loss layer concept. Tranches are differentiated by allocation of portfolio loss and cash flow timing. The different tranches form the liability side of the SPV that is a separate legal entity. In a true-sale transaction the reference credit portfolio represents the asset side of SPV. The SPV might also gain exposure to portfolio credit risk in a synthetic way by a number of CDS contracts on different names.

The junior tranche incorporates the lowest rank and absorbs the risk of payment delays and the losses of the credit portfolio up to a certain level, which reduces risks of tranches of higher priority like the mezzanine and the senior tranches with the highest level of priority. The junior tranche investors receive the cash flow residuals of the transaction after the claims of higher priority tranches have been satisfied, so it is sometimes called equity tranche. In this representation, the tranches consist of both different qualities of abstract liabilities and a piece similar to an asset with equity characteristics. Since the expected total loss of the reference portfolio often lies within the loss layer of equity tranche it is referred to as First Loss Piece (FLP).

Rank and volume of a certain tranche resemble the individual characteristics. Volumes are chosen in a way that refinancing cost is minimized and the constraints of tranche investors are taken into account when customized. In this way, tranches of low rank protect higher tranches, so investors receive a higher coupon due to higher risk borne.

Financial institutions that are involved in an originating process of CDOs wish to transfer portfolio credit risk to other market participants. According to Krahnen (2005), CDOs are an instrument for investment banks and lending institutions to transfer parts of the risk of a credit portfolio to so-called 'remote investors'. This is due to the bridging effect between specialized originating financial institutions and relatively uninformed remote investors in a well-structured CDO transaction. Both parties, however, are interested in optimizing their risk/return profiles since investment in certain tranches might raise investment income and contribute to portfolio diversification.

The originating financial institution has several intentions to initiate a CDO transaction. Following the lines of Lucas and Sam (2001), these are speculative profit opportunities, gaining regulatory capital relief and enhancement of the balance-sheet structure under aspects of return and liquidity. Krahnen (2005) also mentions a reduction of exposure towards systematic shocks of the economy marking times of recessions.

This section describes the risks and economics of CDOs, especially synthetic CDOs. First of all, Gibson (2004) shows that mezzanine and equity tranches typically contain

a small fraction of the notional amount of the reference portfolio, but resemble highly leveraged products concerning the underlying credit risk.
The reference portfolio in a synthetic CDO is entirely made up of CDSs. This investment class developed as an outgrowth of cash CDOs that have a reference portfolio consisting of cash assets such as corporate bonds and loans. In cash CDOs, pricing and risk management is conducted by traditional securitization techniques, whereas synthetic CDOs rely on techniques to price credit derivatives like CDSs. Nevertheless, both risk transfer types are characterized by their tranching of credit risk.
The CDO issuer owns a portfolio of CDSs and is thus acts as a seller of protection. This risk is hedged by selling CDO tranches to investors, so the CDO issuer basically acts as an intermediary.
According to Duffie und Gârleanu (2001), in perfect capital markets, CDOs would serve no purpose as the cost of constructing and marketing a CDO would inhibit its creation. In practice, CDOs address some important market imperfections which can be depicted clearly in the case of cash CDOs. First, banks have regulatory capital requirements that they have to meet, so it may be valuable to securitize and sell some portion of their assets to hedge credit risk and/or reduce costly regulatory capital. Second, asset illiquidity may be overcome by securitization and thereby raise the total valuation of the portfolio. Illiquidity generally inhibits meeting a bank's diversification and risk/return targets.
A CDO transaction can further be characterized by the motivation of the originator. The *balance-sheet* CDO is designed to remove credit-related assets from the balance sheet, achieving capital relief and also increasing the valuation of the assets through an increase in liquidity. In a synthetic transaction, the bank does not actually transfer ownership of the loans and bonds to the SPV - a legally independent offshore entitiy absorbing the collateral portfolio completely and set up only for the purpose of a CDO transaction. The direct sale to the SPV may sometimes comprise client relationships, so the credit risk transfer is conducted synthetically instead.
The latest developments, however, are *'arbitrage'* deals that are driven by the needs of credit investors rather than commercial banks. This is due to the fact that the cost of a credit investor to assemble a portfolio of bonds and loans to create certain risk/return profiles is higher than creating a CDO. The CDO's spread income can compensate investors in the CDO tranches including transactions costs. In contrast, the cost of directly investing into a portfolio of bonds and loans is dominated by the high bid-ask spreads reflecting their illiquidity.
Arbitrage CDOs are inspired by equity tranche investors who often also originate the transaction. These investors hope to achieve a leverged return between the after-default

yield on assets and the financing cost due debt tranches.[31] This potential spread is the arbitrage of this transaction.[32] The originating institution purchases the CDO's assets from a variety of sources from the open market and starts the CDO transaction.
The distinction commonly drawn between balance-sheet and arbitrage CDOs ignores the fact that the asset seller in a balance-sheet CDO also enjoys potential 'arbitrage' profits when retaining the equity tranche. There is not much difference between the holder of the equity tranche in a CDO that buys the assets or an equity tranche investor in a CDO that buys assets the equity investor originated.[33]

Synthetic CDO tranches can either be *funded* or *unfunded* to offer tailored securities to meet investors' preferences.[34] An unfunded CDO is simply a multi-name CDS. The other version resembles a funded note, where the buyer of protection initially receives a pool of collateral securities from the protection seller and pays an upfront fee, in addition to paying a quarterly premium. In an unfunded contract, the protection buyer is exposed to counterparty risk, whereas in a funded contract, the protection buyer is exposed only to the risk of credit deterioration in the reference pool.
The pricing of synthetic CDOs described in chapter 7 is constructed like a tranche swap contract, but neglects counterparty risk due to simplification.[35] The premium, expressed as a fix tranche spread, is paid on the balance of the tranche principal remaining in the tranche after losses have been paid.

3.3.1 The Risk of CDO Tranches

A useful concept when assessing the risk of tranches is leverage which will be outlined by the following example. In a stylized hypothetical synthetic CDO transaction with three tranches, the CDO issuer hedges 100 CDS contracts by selling the three tranches at the par spread. Furthermore, the deal has the following features: equity tranche nominal 3%, mezzanine 7% and senior the remaining 90%.[36] The issuer will initially pay more than the underlying portfolio generates - the net position exhibits negative carry. But over time, occuring defaults erode the principal of the high-yielding equity tranche, which suffers from the first losses[37], and the negative carry diminishes as premium payments to the equity tranche holders are dependent on the stochastically

[31] See Lucas and Sam (2001), p. 6.
[32] See Gibson (2004), p. 3, footnote 3, remark that the terminology has taken hold despite the absence of a true arbitrage.
[33] See Lucas and Sam (2001), p. 6.
[34] See Amato und Gyntelberg (2005), p. 76.
[35] See Gibson (2004), p. 6, footnote 10.
[36] Further quite realistic characteristics are: maturity 5 years, quarterly payments, single-name CDS spread 60 basis points, portfolio notional $ 10 Mio., recovery 40%, default hazard rate 1% p.a., asset correlation 30%, constant continuously compounded interest rate 5% p.a.
[37] See Gibson (2004), p. 10.

decaying tranche nominal. In the end, the present discounted value of the carry is zero because all tranches are fairly priced at their par spreads. So there exists a risk due to timing of defaults, since early defaults reduce payments required on the high-spread equity tranche, whereas late defaults have the opposite effect.

There are two ways to influence timing risk: First, the deal spread could be set at a lower level than the par spread. Second, equity tranche spread might be deferred to a reserve account for later payment dates through structural means. In cash and some synthetic CDOs cash flows are subordinated through some prespecified payment waterfall similar to traditional securitization. In this way, negative carry might be avoided entirely, since cash flows are diverted to first pay structural fees and high-rank tranche spreads. If cash is available after satisfaction of prioritized payments, equity spreads and asset management fees are paid.

This example just gives an idea of the possible variety of contractual designs to perform this kind of securitization. The structure of a CDO transaction can be highly customized and there are dramatical effects on the returns that CDO investors earn. There are three important elements of a CDO's structure[38]:

- the above mentioned payment waterfall determines the priority of cash flows assigned to tranches;
- 'trigger' provisions that divert cash flows from more junior tranches to senior tranches to enhance credit quality. These are par value tests for notional and interest proceeds;
- covenants restricting the choice of credits by bounds to average credit rating or average spread or reinvestment warrants.

There exist models for risk assessment and valuation of these structures more common to cash CDOs.[39] The model described in this paper would not be adequate for such transactions. It is more appropriate for transactions with stylized structures like synthetic CDOs (and more or less exotic derivatives on this product), that are structured like swaps and can be accurately valued by using models exclusively considering loss percentiles. For these transactions there exist several risk measures[40]:

- the sensitivity of a tranches's value to a change in the credit quality of the underlying portfolio;
- a tranche's expected loss being a function of the the portfolio loss until the CDO's maturity;

[38] See Gibson (2004), p. 17, and Lucas and Sam (2001), p. 18.

[39] See the model of Duffie und Gârleanu (2001) that is able to capture a lot of the structure described above.

[40] See Committee on the Global Financial System (2005), p. 47.

- the unexpected loss of a tranche representing a deviation from the expected loss.

According to Gibson (2004), p. 12., all three risk measures characterize the three tranches in the same way. The equity tranche, with only 3% of the notional exposure, is by far the riskiest. The senior tranche, with 90% of the notional exposure, is the least risky. The equity and mezzanine tranches together account for only 10% of the CDO's notional amount but 70 to 90% of the risk, depending on the risk measure used. The developed risk measures can be used to compute the leverage of CDO tranche. Leverage is computed by dividing each tranche's risk, expressed as a percent of notional amount, by the risk of the entire portfolio. When computed this way, a cash investment in bonds or loans has leverage equal to one by construction. The leverage numbers suggest that the hypothetical equity tranche has 13 to 17 times the risk of a cash investment in bonds or loans. The mezzanine tranche has 5 to 7 times the risk of a cash investment. The senior tranche has much less risk than a cash investment.
Credit spreads on the CDO's reference portfolio and a tranche's location in the CDO's capital structure affect the risk and leverage measures presented above. If a tranche moves higher in the CDO's capital structure it has more credit enhancement. The author finds that the riskiness of the reference portfolio has little effect on leverage, but credit enhancement has a strong effect.

3.3.2 Information Asymmetry and CDOs

There are several reasons why equity tranches are often retained by originating banks.[41] First, the CDO originator has the desire to create such risk/return and diversification profiles resulting from an equity tranche investment. Due to construction, the CDO originator has to be invested in this tranche to realize 'arbitrage' profits. Second, the retention can be regarded as an instrument of signaling to reduce informational asymmetry.
According to Duffie und Gârleanu (2001), the effect of adverse selection can be discovered in the transfer of bank loans or junk bonds. There is an informational asymmetry between the potentially better-informed seller of such assets and the potentially uninformed buyer, so there is a price reduction sometimes called a lemon's premium. The technique of tranching may dissolve this informational asymmetry and thus reduce the lemon's premium: A large senior tranche is relatively immune to the effects of adverse selection since a large number of entities have to default to seriously influence the tranche's performance. But this is only the case if the business cycle is in a recession and systematic risk is very high. That scenario is independent of the credit selection

[41] See Gibson (2004), p. 17.

abilities of the originator, so the issuer can retain significant fractions of smaller subordinated tranches that are subject to adverse selection.

Moral hazard in the context of CDOs is revealed in the CDO manager's incentives to select high-quality assets, to engage in costly enforcements of covenants and other restrictions concerning the behavior of the obligor.[42] If the originator has a 100-percent equity interest in the asset cash flows, there are incentives to more effectively select and monitor assets. The issuer can retain a portion of subordinate tranches or in arbitrage CDOs management fees may be subordinate to the issued tranches. Having own money at stake in a first-loss position demonstrates a degree of commitment to profitability and builds and maintains a reputation for being a 'reliable CDO issuer'. In this light, investors' willingness to pay more for tranches is increased and the total valuation of the tranche structure is increased in comparison to an unprioritized pass-through structure.

In the context of standardized CDO tranches these information asymmetry induced risks are dissolved. The portfolio of underlyings is not picked by an originator but there is public information about the composition of the index portfolio. This reduces analyses to a number of ingredients involved in pricing individual tranches. These are default probability, default severity (or recovery) and default correlation. There exist sophisticated single-name credit models to estimate the impact of default probability and recovery assumptions. Correlation, however, deals with the distribution of joint defaults so there is an additional complexity involved when it comes to risk assessment of portfolios. This feature will be outlined in the following section.

3.3.3 The Impact of Correlation

In the context of stylized synthetic CDOs, the diametrically opposed motivations of high and low rank tranche investors can be illustrated by a simple option's model of tranches.[43] These transactions are strongly related to the structure of a percentile basket derivative, consisting of a transaction with equity, mezzanine and senior tranche in this example. The rank of a certain tranche is simply defined by percentile levels at which the investors bear losses stemming from the reference credit portfolio. The nominal of each tranche is defined by a loss layer which is represented by an attachment point L_L and a detachment point L_U resembling percentiles of the total initial value of the portfolio which is equivalent to the maximal total loss possible. The equity tranche can be viewed as a zero bond with nominal $L_U - L_L$ and an attachment point of 0% as it is the first loss position. At the beginning of the transaction the investor as a seller of protection pays the nominal to the SPV, which is the buyer of protection. At maturity

[42] See Duffie und Gârleanu (2001), p. 5.
[43] See Esposito (2002), p. 2ff.

the investor receives the residual of the nominal reduced by losses that he has to bear. If the realized percentile loss L of the portfolio is larger than the detachment point L_U of the equity tranche the investors's nominal is completely used up for protection and he receives no repayment. If $L < L_U$ he receives the nominal reduced by $L_U - L$. Since the SPV is protected from the first losses up to L_U, the equity tranche investors receive a premium for taking over that risk. The residual value of the tranche T_{Equity} can be expressed in the framework of options:

$$\begin{aligned} T_{Equity} &= L_U - \min(L_U, L) \\ &= L_U + \max(-L_U, -L) \\ &= \max(0, L_U - L) \\ &= \mathrm{Put}(L, L_U). \end{aligned}$$

So the residual value of the tranche resembles a put option on the portfolio loss with strike L_U. In a similar way, the mezzanine tranche can be represented by detachment/attachment points. The corresponding investor is obligated to compensate for occurring portfolio losses larger than the mezzanine attachment point which is equivalent to the detachment point of the equity tranche, but only up to the level of the mezzanine detachment point which in turn is equivalent to the attachment point of the senior tranche. This results in the following residual value of the mezzanine tranche:

$$\begin{aligned} T_{Mezzanine} &= \max(\min(L_U - L_L, L_U - L), 0) \\ &= \max(L_U - L_L + \min(0, L_L - L), 0) \\ &= L_U - L_L + \max(\min(0, L_L - L), -(L_U - L_L)) \\ &= L_U - L_L + \min(0, L_L - L) + \max(0, L_L - L_U - \min(0, L_L - L)) \\ &= L_U - L_L - \max(0, L - L_L) + \max(0, L_L - L_U + \max(0, L - L_L)) \\ &= L_U - L_L - \max(0, L - L_L) + \max(0, \max(L_L - L_U, L - L_U)). \end{aligned}$$

Since $L_L - L_U$ is negative this results in

$$\begin{aligned} T_{Mezzanine} &= L_U - L_L - \max(0, L - L_L) + \max(0, L - L_U)) \\ &= L_U - L_L - \mathrm{Call}(L, L_L) + \mathrm{Call}(L, L_U). \end{aligned}$$

The repayment of the mezzanine tranche can be interpreted as a zero bond with nominal $L_U - L_L$ and a combination of two European calls with the same underlying L. Finally, since the detachment point of the senior tranche corresponds to 100% of the initial portfolio value the residual value is

$$\begin{aligned} T_{Senior} &= 1 - L_{lower} - \max(0, L - L_{lower}) \\ &= 1 - L_{lower} - \mathrm{Call}(L, L_{lower}). \end{aligned}$$

In general, long options exhibit a positive vega, i.e. both put long and call long benefit from an increase in loss volatility. In this example, the equity tranche investors benefit from increased loss volatility since they are long a put option, in contrast to senior investors who are short an option. We see ambivalent coherences for the mezzanine tranche since calls are short and long in that combination.

Beside the different reactions of the equity and senior tranches towards changing loss volatility another major parameter underlines this diametrically opposed positions: the amount of diversification in the reference portfolio which can be measured by default correlation. This is an important measure in the credit portfolio context which can be seen from an example in Schönbucher (2000). We consider two obligors i and j and a fix time horizon T. An indicator variable shows if obligor i has defaulted until time T. In case of default the variable jumps from 0 to 1. That is represented by the notation $\mathbf{1}_i$. The expected value of the indicator function is just the default probability of obligor i:

$$\mathbb{E}(\mathbf{1}_i) = p_i.$$

The variance of the default indication is computed by

$$Var(\mathbf{1}_i) = \mathbb{E}\left(\mathbf{1}_i - \mathbb{E}(\mathbf{1}_i)\right)^2 = p_i \cdot (1 - p_i).$$

The probability that i and j default before T is $\mathbb{E}(\mathbf{1}_i \cdot \mathbf{1}_j) = p_{i,j}$. The notation of conditional default probabilities are $p_{i|j}$ and accordingly $p_{j|i}$. The linear coefficient between the events of default is $Corr(\mathbf{1}_i, \mathbf{1}_j) = \varrho_{i,j}$. The connection is given by Bayes' Theorem:

$$p_{i|j} = \frac{p_{i,j}}{p_j}, \qquad p_{j|i} = \frac{p_{i,j}}{p_i}.$$

The linear coefficient of correlation is computed by

$$\varrho_{i,j} = \frac{p_{i,j} - p_i p_j}{\sqrt{p_i(1-p_i)p_j(1-p_j)}}.$$

Since the denominator of this expression is positive it is important that the joint probability of default is greater than the product of the individual default probabilities to actually result in a positive default correlation. The joint default probability is given by

$$p_{i,j} = p_i p_j + \varrho_{i,j}\sqrt{p_i(1-p_i)p_j(1-p_j)}.$$

The default of i is conditional to the one of j:

$$p_{i|j} = p_i + \varrho_{i,j}\sqrt{\frac{p_i}{p_j}p_i(1-p_i)(1-p_j)}.$$

Assuming realistic numbers for the default correlation of 10% and for the default probabilities of 1% it follows from the above that

$$p_{i,j} = 0,01 * 0,01 + 0,1 * 0,01 * 0,99 = 0,00109 \approx p^2 + \varrho p \approx \varrho p$$

and

$$p_{i|j} = 0,01 + 0,1 * 0,99 = 0,109 \approx \varrho.$$

It can be deduced that both the joint and the conditional default probability is dominated by the default correlation coefficient. According to figure 1 the influence of default correlation on the portfolio loss distribution is enormous. When default cor-

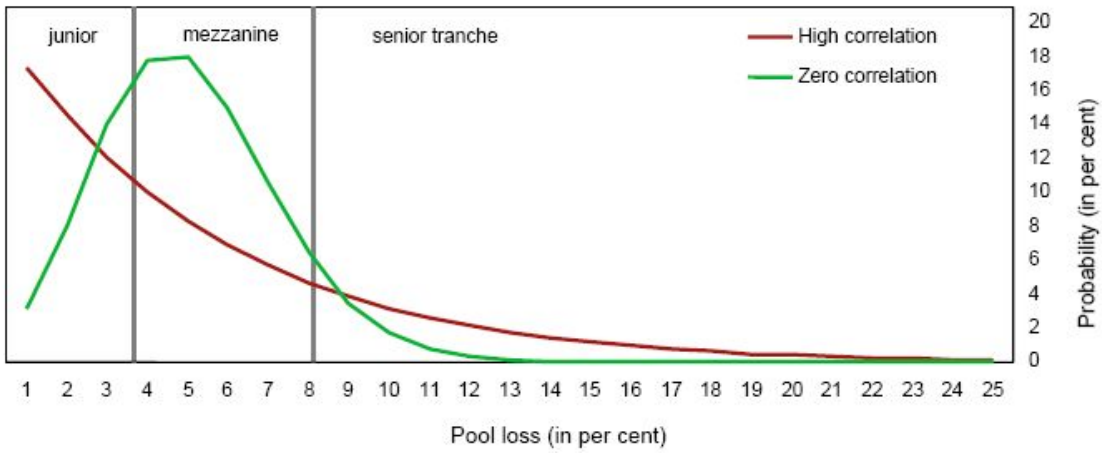

Figure 1: Impact of default correlation on the portfolio loss distribution for high and zero correlation. Source: Committee on the Global Financial System (2005a), p. 18.

relation is low, investors are able to analyse companies on a case-by-case basis. In an untranched portfolio, the correlation between names defaulting has an effect on the value of the portfolio only when correlation is high. It is only when portfolios are tranched that the relative value of default correlation becomes meaningful.

According to Gibson (2004), the effect of correlation on CDO tranches is intuitive. The more the defaults within a portfolio become correlated, the more the portfolio behaves like a single credit. So the probability of the equity tranche being wiped out becomes more similar to the probability of the most senior tranches being wiped out.

A higher correlation of defaults implies a greater likelihood that losses will wipe out the equity and mezzanine tranches and inflict losses on the senior tranche. Thus, the value of the senior tranche falls as correlation rises. Conversely, higher general correlation also makes the extreme case of very few defaults more likely. Thus, the value of the equity tranche rises as correlation rises. Equity tranche investors gain more in a scenario with very few defaults than they lose from a scenario with many defaults (they are only exposed to the first few defaults). Mezzanine tranches are subject to both effects, which can broadly cancel each other out and make mezzanine tranches less sensitive to correlation.

Since CDO tranches are sensitive to correlation which in turn typically depends on the state of the business cycle, the correlation risk can also be characterized by business

cycle risk.[44] In models underlying the Basel II proposal as well as many other credit portfolio models, a credit's default risk depends on a common, macroeconomic risk factor and an idiosyncratic risk factor. From a perspective of equity markets, this concept has been commonplace for many years, and part of the driving forces behind the latest developments in correlation trading and CDOs seem to be an increased willingness and ability to quantify and measure the effect of diversification in credit portfolios.[45]
The hypothetical portfolio is analyzed again to examine the impact of business cycle risk on the tranches.[46] For this purpose, the expected loss of a tranche is computed conditional on a certain value of the common factor representing three states of the economy: boom, growth according to trend, and recession.[47]
The equity tranche expects to bear defaults of about half its notional, even in a trend growth macroeconomic scenario, and looses its entire notional in a recession. The mezzanine tranche suffers no losses in a boom and almost none in a trend growth scenario. Since it bears most of the portfolio's expected loss in a recession it can be regarded as a leveraged bet on the business cycle. The senior tranche suffers very little loss even in a recession. At around the 96th percentile of the common factor, the senior tranche principal is significantly eroded by additional losses.
Nevertheless, all parties involved in a CDO transaction might benefit from the resulting investment and diversification opportunities to enhance their risk/return profiles. Leverage has to be taken care of and identical ratings do not guarantee identical risk characteristics: They may have equal loss probability but unequal severity due to correlation effects.[48] Effects regarding correlation are particularly strong for some of the most recent innovations in credit markets, namely single-tranche CDOs and first-to-default basket swaps, which are discussed now.

3.3.4 Single-Tranche CDOs and Correlation Trading

In order to conduct a CDO transaction, underwriters are obligated to sell the full tranche structure of the portfolio. Since the tranche's profiles differ heavily, there might be difficulties in transferring the full structure to the capital market. To circumvent these limitations, dealers have recently begun to offer single-tranche CDOs.[49] These products are customized in a way that the investor can select aspects of the reference portfolio as well as the specific portion of the portfolio loss distribution they wish to be exposed to. Dealers may purchase the entire portfolio and hold the unsold tranches

[44] See Committee on the Global Financial System (2005), p. 49.
[45] See Committee on the Global Financial System (2005), p. 38.
[46] See also Gibson (2004), p. 23.
[47] This is realized by setting the common factor driving defaults at its 10th, 50th, and 90th percentiles, respectively.
[48] See Committee on the Global Financial System (2005), p.26.
[49] See Gibson (2004),p. 19.

but this may contradict their investment strategy. Instead, dealers may hedge their exposures on just the single tranche that is sold. When selling the tranche the dealer has bought protection to that extent. This position has to be hedged by selling credit protection on the entire reference portfolio in an amount that offsets the exposure to spread movements and defaults. As described before, the delta of a tranche resembles the sensitivity towards small moves in credit spreads. Instead of selling 100% protection on the entire portfolio a smaller fraction of protection is sold. Consider the hypothetical portfolio with nominal $1 billion: In a hedge concerning the entire portfolio, there would be sold $10 million of protection on each name. Since the equity tranche with a 3% detachment point exhibits the highest delta, $4.83 of protection have to be sold on each name to eliminate the delta risk of this tranche. For the mezzanine tranche covering the next 7% of the portfolio, credit protection of $4.71 are necessary and the remaining 90% of the losses in the senior tranche an be hedged by $2.17 per name in the reference portfolio. In a heterogeneous portfolio the deltas and the amount of protection has to be computed for each underlying portfolio name. Once again, the difference in leverage of tranches can be seen from these numbers. Computing the size of each tranche's hedge position relative to the tranche notional would give similar results as to the leverage measures described before.

It has to be remarked here, that these amounts change over time as the levels of credit spreads change. In single-tranche CDO it is very obvious that these products require dynamic hedging and consistent models to support such deals. There are additional risks in such a deal that have to be considered: There is no need to compute such hedge positions of each underlying CDS in a traditional CDO, so there is additional model risk involved. When dynamically adjusting hedges, there is liquidity risk, so trading costs may rise for a perfect hedge. Also, there might be large movements of credit spreads referred to as convexity or gamma risk similar to the option's market which are not protected by delta hedges. If spreads even move considerable in a very small time period, dealers are additionally exposed to jump-to-default risk. Delta hedging by definition requires an approximation of continuous rebalancing of very small changes, and therefore can not hedge properly against discontinuous moves like the jump-to-default. Once again, the correlation risk has to be kept in mind: If hedges were set up with wrong correlation assumptions which influences the value of the single tranche, there is a large impact on the price and hedge parameters.

The impact of large credit spread fluctuations can be measured on a mark-to-market profit&loss basis.[50] When the gamma risk of a sold single-tranche CDO, hedged with single-name CDS, is considered, there are seveal aspects that have to be noted. There are completely different reactions of each tranches regarding the sensitivity to a shock

[50] See Committee on the Global Financial System (2005), p. 53.

of a single credit spread and a shock to all reference entities. Also, as time passes by towards maturity, hedging creates more and more difficulties due to illiquidity in the CDS market for maturities shorter than 5 years.

Summarizing, a model is needed that is calibrated both to the underlying CDS spreads as well as to quotes of standardized index tranches. The latter requirement ensures a reflection of the broad-based credit risk inherent in the market portfolio. If we assume that the market index provides a high level of diversification, customized tranches on different reference portfolios might exhibit a different level of diversification due to to concentration: The more discrepancy is observed between both portfolios, the worse the correlation assumption. Additionally, in the model there has to be a dynamic evolution of correlated credit spreads to be able to hedge dynamically.

Correlation trading is the most sophisticated segment of the credit derivatives market. Reportedly, large dealers and hedge funds are most active in this market segment.[51]

3.3.5 Facts about CDS indices

A CDS index contract is an insurance covering default risk on the pool of names in the index.[52] The buyer of protection on the index is obligated to pay the same premium on all the names in the index (called the fixed rate). In the case of a credit event due to bankruptcy or failure to pay, the respective entity is removed from the index and the contract continues with a reduced notional until maturity, but the fixed rate remains unchanged. Each entity has a notional of $1 Mio. and if one entity defaults, the CDS index notional is reduced by 1/125 of the original notional and by one defaulted entity. This fixed rate is slightly lower than the average of the underlying CDS spreads.[53] This is due to risky companies with high spreads that are not expected to be paid as long as for the high quality companies due to higher default risk. Therefore high default companies should carry more weight.

The main indices have been consolidated into a single family for North America and Emerging Markets (DJ CDX) and for Europe and Asia (DJ iTraxx).[54] The indices comprise the most liquid single-name CDS and a group of global dealers is committed to market making. Besides the geographical segmentation, there exist sub-indices differentiated by investment grade/high yield, by sector (financials, industrials, consumer, etc.), and maturity (5/10 years). The indices trade OTC in funded or unfunded form (swap contracts).

A once created index remains static over its lifetime. The payment or coupon dates are the standard CDS dates: 20th of March, June, September, and December. On the

[51] See Gibson (2004), p. 23.
[52] See Amato und Gyntelberg (2005), p. 74.
[53] See Hull, Predescu, White (2005), p. 6.
[54] The indices are managed by Dow Jones.

roll dates (the March and September dates each year), new index compositions will become on-the-run to ensure that the reference portfolio represents the aggregate market and that the on-the-run contracts have 5 or 10 year maturity. Intuitively, the fixed rate is also adapted besides reference entities and maturities. Fresh index contracts have no influence on older deals. Their specificities like maturity, fixed rate and pool composition remain unchanged as fixed in the past, except for defaults, of course.

The most important indices are the DJ CDX North America investment grade index and the iTraxx index for Europe. Each index consists of 125 most important CDSs. Index tranches are also standardized regarding the composition of the pool already described 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.[55] The premiums on the standardized mezzanine and senior tranches are running spread with no upfront payment. By contrast, there exists an upfront payment with the equity tranche as a percentage of tranche notional, in addition to paying a running spread premium of 500 basis points. The payment of a relatively large upfront-payment changes the exposure to timing risk compared to the non-equity tranches. The pricing of standardized tranche swap contracts and more details about index tranches will be provided in chapter 7.

Due to the attention of many dealers towards index tranches, investors describe liquidity in the indices as excellent, even when the single-name CDS market is less liquid.[56] Trading in index tranches provides a useful tool for taking on or hedging macro risk exposures. Since index tranches are similar to customized single-tranches, both instruments in combination allow for hedging correlation risk. This further improves transparency in customized tranches and makes this investment class more attractive. Additionally, dealing in sub-indices for different sectors or risk exposures and bets on the outperformance or underperformance of the major index can be taken, just as in equity markets.

After this introduction into the portfolio credit derivative business we will now proceed with the description of the standard market model that is used to communicate prices: the Gaussian copula.

[55] See Amato und Gyntelberg (2005), p. 77, footnote 10.

[56] See Committee on the Global Financial System (2005), p. 56.

4 The Factor Gaussian Copula Model

To classify the structural model of Hull, Predescu, White (2005) in the context of other models, we first have to describe the static One-factor Gaussian copula which resembles the current market standard model. The model class provides a structural link between the default probability and the corresponding fundamental financial variables of the firm. Information regarding the firm's credit structure, precisely assets and liabilities, is linked to the company's default behavior. The resulting fundamental default variables are considered in the context of a portfolio by making such variables dependent on one or more, generally unobserved, common factors. This leads to a complexity reduction and the concept of conditional independence can be applied which yields substantial simplification: If it is assumed that defaults of different titles in the credit portfolio are independent conditional on a common market factor, it is simpler to compute the aggregate portfolio loss distributions. The factor (or conditional independence) approach allows to use semi-analytic computation techniques avoiding time consuming Monte Carlo simulations. Examples are the approaches described by Gregory and Laurent (2004) who use fast Fourier transformation techniques, as well as Hull and White (2004) and Andersen, Sidenius, Basu (2003) who apply recurrent and iterative numerical procedures to build up the loss distribution for the pool of reference instruments. These copula models are employed to generate loss distributions at different points in time since cash-flow based valuation as in CDOs demands for such architectures.[57] A second approach within the structural framework was introduced by Black und Cox (1976), allowing defaults at any time until maturity. Within this text, this model will be referred to as 'dynamic Factor Gaussian copula' since we will show how the Hull/Predescu/White model generates joint defaults in a common way and simultaneously allows for a dynamic evolution of correlated credit spreads.

The factor Gaussian copula model is based on the approaches for option pricing by Black und Scholes (1973) and the extension for credit risk by Merton (1974) to value corporate liabilities. The focus of modeling is the evolution of the firm's structural variables, such as asset and debt values, to determine the time of default. In Merton's model a firm defaults if, at the time of servicing the debt, its assets are below its outstanding debt. The dynamics of the assets V_i of company i follow a continuous-time

[57] According to Elizalde (2005b), p. 3., reduced form models, in contrast, rely solely on the market prices of the firms' defaultable instruments to extract both their default probabilities as well as their credit risk dependencies. This leads to easy calibration properties but the link between credit risk and information regarding the firms' financial situation reflected by balance sheets is missing and the extension to the portfolio context is not straightforward. The time of default in intensity/reduced form models is determined by the first jump of an exogenously given jump process. The parameters governing the default hazard rate are inferred from market data.

diffusion under the physical or real probability measure $\mathbb{P}$. The firm-value process is driven by geometric Brownian motion

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

with μ_i as the total expected return, σ_i as the (relative) instantaneous volatility, and $W_i(t)$ as a standard Brownian motion under the real measure $\mathbb{P}$. The capital structure of company i consists of equity and a zero bond with nominal K_i. The value of the firm V_i is the sum of both. Under these assumptions the equity resembles a European call option with maturity T and a strike of K_i. The company defaults if the nominal repayment K_i at maturity T cannot be served by the the assets $V_i(T)$ at time T. Consequently, a default of this zero bond can only happen at T and the default probability is

$$P(V_i(T) < K_i).$$

Due to Itô's lemma, the asset value at time T can be expressed as a function of the current asset value:

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

The variable $x_i(t,T)$ has a standard normal distribution and is given by

$$x_i(t,T) = \frac{W_i(T) - W_i(t)}{\sqrt{T-t}}.$$ [58]

It can be interpreted as a default variable of the company over a time horizon t to T because at maturity it is considered, if it is below the default barrier D_i. The default of the company in T follows the equivalence

$$V_i(T) < K_i \Leftrightarrow x_i(t,T) < D_i(t,T)$$

with

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

Therefore, the default probability for the time span $T-t$ is a standard normal distributed variable $\Phi(\cdot)$ and the following holds

$$p_i = P(x_i \leq D_i) = \Phi(D_i).$$

This is a latent variable model, since the initial firm value $V_i(t)$, the strike K_i, μ_i, and σ_i are not directly observable. For the purpose of pricing, this model can be used in a risk-neutral world with the measure $\mathbb{Q}$. In this case, a non-dividend paying company

[58] The Brownian motion follows a normal distribution with zero mean and standard deviation $\sqrt{T-t}$. That's why the model is called Gaussian.

has a firm-value process with drift r_i, which is the risk-free rate.

The next step is to specify the default barrier $D_i(t,T)$. Under the assumptions that company i has a rating of BB[59] and the one-year default probability of a BB company is 1.06% the default barrier then is $D_{BB}(t = 0, T = 1)$ and with standard normal distributed default variable, this results in the default probability

$$P(x_i \leq D_{BB}) = \Phi(D_{BB}) \tag{1}$$

and the barrier

$$D_{BB}(0,1) = \Phi^{-1}(1.06\%).$$

Since the only random variable affecting the default/non-default status of a firm is x_i, the correlation structure between the firms' default probabilities have to be introduced through the random variables $x_1, \ldots, x_n$ with n entities in the underlying portfolio. It is assumed that the asset return correlation coefficient of each pair of random variables x_i and x_j, $i \neq j$, is $\rho_{i,j}$. Then the joint default probability of company i with rating BB and company j with rating AA is

$$p_{i,j} = P(x_i \leq D_{BB}, x_j \leq D_{AA}) = \int_{-\infty}^{D_{BB}} \int_{-\infty}^{D_{AA}} \varphi_{\rho_{i,j}}(x_i, x_j) dx_i dx_j,$$

with $\varphi_{\rho_{i,j}}(\cdot)$ as the density of the bivariate standard normal distribution. This concept can be extended to a multivariate standard normal distribution of the order of entities in the credit portfolio. This results in the probability distribution of joint defaults in the portfolio.

The disadvantage of this approach is that the default barrier D_{BB} is related to the given default probabilities of rating agencies. But these figures are only available for one year and for this reason default barriers concerning other horizons can not be generated. But for valuation and pricing, risk-neutral default probabilities are required. Rating agency given default probabilities however exhibit the real measure. An approach to extend this framework is presented in section 4.2 about copulas.

Another disadvantage of this method is that with n entities in the portfolio we have to estimate $n(n-1)/2$ pairwise asset return correlations. For the portfolio context of default securities, specializing to correlation matrices with a factor structure can yield significant improvements in computational speed. In practice, the variance/covariance matrix is estimated by factor analysis or regression of equity returns or credit spreads of the respective companies.

[59] This example is from Gupton, Finger and Bhatia (1997), p. 88.

4.1 Factor Models for Credit Portfolios

In the framework of factor copula models there is a special factor reduction needed that identifies abstract factors providing no direct link to macroeconomic indicators like the S&P 500 index or the unemployment rate. The model requires a low number of independent driving systematic factors that fully explain the correlations.[60] This leads to conditional loss distributions and as a next step the distributions are integrated over the systematic independent factor(s) to arrive at unconditional loss distributions. Again the default variables x_i are assumed to be standard normal. In this example the reduction is realized with one systematic factor M and the idiosyncratic factors Z_i that exhibit the standard normal distribution as well. This factor model goes back to the work of Vasicek (1977) and has the form

$$x_i = a_i M + b_i Z_i \tag{2}$$

with a_i and b_i as factor loadings for systematic and idiosyncratic risk.[61] All the Z_i are independent from each other and also exhibit independence to the systematic factor M. The amount of systematic risk that can not be explained by the factor loadings a_i is pooled in the factor loadings $b_i = \sqrt{1-a_i^2}$. Unity is according to the standard deviation of the x_i. Equation (2) changes to

$$x_i = a_i M + \sqrt{1-a_i^2} Z_i. \tag{3}$$

With $\mathcal{N}(0,1)$ and the independence of systematic and idiosyncratic risks this results in the distribution of the returns:

$$a_i\mathcal{N}(0,1) + b_i\mathcal{N}(0,1) \sim \left(\sqrt{a_i^2+b_i^2}\right)\mathcal{N}(0,1) \sim \left(\sqrt{a_i^2+(1-a_i^2)}\right)\mathcal{N}(0,1) \sim \mathcal{N}(0,1).$$

The correlation between the x_i (for $i \neq j$) is computed by:

$$\mathrm{Corr}(X_i, X_j) = \frac{\mathrm{Cov}\left(a_i M + \sqrt{1-a_i^2} Z_i, a_j M + \sqrt{1-a_j^2} Z_j\right)}{\sqrt{\mathrm{Var}\left(a_i M + \sqrt{1-a_i^2} Z_i\right)} \cdot \sqrt{\mathrm{Var}\left(a_j M + \sqrt{1-a_j^2} Z_j\right)}} \tag{4}$$

The nominator of equation (4) is changed to:

$$\begin{aligned} &\mathrm{Cov}\left(a_i M, a_j M\right) + \mathrm{Cov}\left(a_i M, \sqrt{1-a_j^2} Z_j\right) + \mathrm{Cov}\left(a_j M, \sqrt{1-a_i^2} Z_i\right) \\ &+ \mathrm{Cov}\left(\sqrt{1-a_i^2} Z_i, \sqrt{1-a_j^2} Z_j\right). \end{aligned} \tag{5}$$

[60] A corresponding algorithm that is based on principal component analysis is presented in Andersen, Sidenius, Basu (2003).

[61] The framework is well known in the Basel II context.

The last three summands are zero as Z_i and Z_j are independent from each other and also independent of M. Since the standard deviations of M and Z_i are 1 by normalization, the product in the denominator of equation (4) changes to:

$$\sqrt{\mathrm{Var}(a_i M + \sqrt{1-a_i^2} Z_i)} = \sqrt{a_i^2 + 1 - a_i^2} = \sqrt{1} = |1|. \tag{6}$$

We get

$$\mathrm{Corr}(x_i, x_j) = \frac{\mathrm{Cov}(a_i M, a_j M)}{1} = a_i a_j \mathrm{Var}(M) = a_i a_j.$$

These equations can easily be extended to m independent systematic factors $M_1, M_2, \ldots, M_m$. Equation (3) can be noted as

$$x_i = a_{i1} M_1 + a_{i2} M_2 + \ldots + a_{im} M_m + \sqrt{1 - a_{i1}^2 - a_{i2}^2 - \ldots - a_{im}^2} Z_i. \tag{7}$$

The correlation between the x_i now is $a_{i1}a_{j1} + a_{i2}a_{j2} + \ldots + a_{im}a_{jm}$. This can be shown with an example of 2 factors M_1 and M_2:

$$\begin{aligned} x_i &= a_{i1} M_1 + a_{i2} M_2 + \sqrt{1 - a_{i1}^2 - a_{i2}^2} Z_i \\ x_j &= a_{j1} M_1 + a_{j2} M_2 + \sqrt{1 - a_{j1}^2 - a_{j2}^2} Z_j. \end{aligned}$$

Similar to equation (5) there only remain those covariances, whose value is not 0. This is due to independence of the Z_is. Also the systematic factors are independent of each other by the construction of the special factor model. The denominator of equation (4) is populated with m systematic factors and by construction it is equal to 1. This results in

$$\mathrm{Corr}(x_i, x_j) = a_{i1}a_{j1}\mathrm{Var}(M_1) + a_{i2}a_{j2}\mathrm{Var}(M_2) = a_{i1}a_{j1} + a_{i2}a_{j2}. \tag{8}$$

Because of the square root it has to be valid that $a_{i1}^2 + a_{i2}^2 + \ldots + a_{im}^2 \leq 1$.
It should be noted that the asset correlation coefficient is not equal to the default correlation coefficient. Due to the construction that default is only triggered as soon as the asset value return falls below a prespecified barrier, the default correlation coefficient is much smaller than the asset correlation.

4.2 The One-Factor Gaussian Copula

In the sense of an abstract definition of the problem to value basket/portfolio credit derivatives, a series of payments that have a functional dependency $f(\cdot)$ to the vector of times to default $\boldsymbol{\tau}$ has to be considered.[62] The derivation of the joint risk-neutral default probability

$$Q(\boldsymbol{\tau} \leq \mathbf{T}) = Q(\tau_1 < t_1, \ldots, \tau_n < t_n) \tag{9}$$

[62] See Andersen, Sidenius, Basu (2003), p. 67.

is explained in this section. This probability can be used in the framework of the Asset Pricing Theory. At time 0 the value C of the basket/portfolio credit derivative is the expected value of the function $f(\boldsymbol{\tau})$

$$C(0) = \mathbb{E}(f(\boldsymbol{\tau})). \tag{10}$$

It is explicitly assumed that $f(\cdot)$ incorporates the discount of future cash flows. The first derivative of equation (9) produces the density φ of the default time vector $\boldsymbol{\tau}$. This results in today's value of the credit basket derivative

$$C(0) = \int_{\mathbb{R}^n \in [0,T]} f(\boldsymbol{t})\varphi(\boldsymbol{t})d\boldsymbol{t}, \tag{11}$$

This is an n-dimensional integral of all default times until maturity of the credit basket derivative.

There are methods to extract the marginal risk-neutral default probabilities $(Q_i(t))_{t\geq 0}$ from credit spread or CDS spread data. This is realized by the concept of the copula. The advantage is that the joint multivariate default time distribution is decomposed into the marginal default time distribution and another part that represents the dependence structure. This means that the marginal distributions can be chosen independently and according to frameworks for individual obligors. The concept of copulas will be introduced now with the help of Sklar's Theorem [63]: Let G be an n-dimensional distribution function with marginals $F_1, \ldots, F_n$. Then there exists an n-Copula C so that for all $x \in \mathbb{R}^n$:

$$G(x_1, \ldots, x_n) = C(F_1(x_1), \ldots, F_n(x_n)). \tag{12}$$

If $F_1, \ldots, F_n$ are continuous functions then the copula C is clearly defined. This theorem can be used vice-versa:

If C is an n-Copula and $F_1, \ldots, F_n$ are distribution functions then the function G from above is defined as an n-dimensional distribution function with marginals $F_1, \ldots, F_n$.

If F is a continuous monotonically increasing function, the inverse function can be computed. Equation (12) can thus be transformed to:

$$C(u_1, \ldots, u_n) = G(F_1^{-1}(u_1), \ldots, F_n^{-1}(u_n)).$$

This shows that a copula function is a mapping of an n-dimensional standard uniform distribution to the interval [0,1].

In the case of independence of the $x_1 \ldots, x_n$ the simplest form is a product copula. Another famous representative is the Gauss copula:

$$C(u_1, \ldots, u_n) = \Phi_{\Sigma,n}(\Phi^{-1}(u_1), \ldots, \Phi^{-1}(u_n)),$$

[63] See Embrechts, Lindskog, McNeil (2003), S. 4.

with Φ as the standard normal distribution and $\Phi_{\Sigma,n}$ as multivariate standard normal distribution with covariance matrix Σ. The t-Copula has the following shape:

$$C(u_1, \ldots, u_n) = t_{\Sigma,n,\nu}(t_\nu^{-1}(u_1), \ldots, t_\nu^{-1}(u_n)).$$

The product copula enters numerous models such as Hull and White (2004) or Laurent and Gregory (2003) because of the inherent conditional independence that makes approaches like analytic algorithms possible.
The combination of the recently presented concepts realize the generation of complete loss distributions of a credit portfolio. These are the prerequisites to price certain tranches. Let H be the cumulative distribution function of the idiosyncratic factor in equation (3). Solving for Z_i we receive

$$P(x_i < D_i(T)|M = m) = H\left[\frac{D_i(T) - a_i m}{\sqrt{1 - a_i^2}}\right]. \tag{13}$$

$D_i(T)$ represents the individual default barrier of obligor i at time T in the sense of Merton. Conditional to factor M the firm's default indicator variables x_i are independent from each other. Choosing the standard normal distribution Φ for the Z_i and M we get

$$P(x_i < D_i(T)|M = m) = \Phi_{Z_i}\left[\frac{D_i(T) - a_i m}{\sqrt{1 - a_i^2}}\right]. \tag{14}$$

The step from the Merton-approach to the survival time copula was undertaken by Li (2001) for two obligors and by Laurent and Gregory (2003) for n obligors. In Merton (1974) the exogenous default barrier $D_i(T)$ for one obligor is given for a fix one-year time horizon and is accessed through default probabilities under the real measure. But for the fair spread computation of tranches default barriers in quarterly intervals that are based on the risk-neutral default probabilities $Q_i(\tau_i < t_i)$ have to be known.
There has to be a percentile transformation so that the distribution F of the returns x_i is mapped to the distribution of default times.[64] Especially the default barrier $x_i = D_i(t_i)$ is transformed to $\tau_i = t_i$ by means of the mapping $\tau_i = Q_i^{-1}[F_i(D_i(t_i))]$.
Now the copula approach can be applied for the joint distribution of default times.[65] We define the times to default as $\tau_i = Q_i^{-1}(u_i)$ with u_i realization of a standard uniform random variable. The inverse Q_i^{-1} can be computed as the distribution function $Q_i(\tau_i < t_i)$ is strictly monotonously increasing. From Skla's Theorem about copulas, any arbitrary joint continuous distribution function can be split into a copula and the marginal distributions. That's why

$$C(u_1, \ldots, u_n) = Q(\tau_1 < t_1, \ldots, \tau_n < t_n)$$

[64] See Hull and White (2004), p. 10.
[65] See Schmidt and Ward (2002), p. 111.

defines a joint distribution on $[0,1]$. The joint distribution of $(\tau_1, \ldots, \tau_n)$ can be written as

$$Q(\tau_1 < t_1, \ldots, \tau_n < t_n) = C(Q_1(t_1), \ldots, Q_n(t_n)).$$

The mapping $D_i(t_i) = F_i^{-1}(Q_i(t_i))$ admits the desired transformation:

$$\begin{aligned} &Q(X_1 < D_1(t_1), \ldots, X_n < D_n(t_n)) = Q(\tau_1 < t_1, \ldots, \tau_n < t_n) \\ &= C(Q_1(t_1), \ldots, Q_n(t_n)) = C(u_1, \ldots, u_n) \\ &= G\left(F_1^{-1}(u_1), \ldots, F_n^{-1}(u_n)\right) = G\left(F_1^{-1}(Q_1(t_1)), \ldots, F_n^{-1}(Q_n(t_n))\right). \end{aligned}$$

The conditional default probability $P(x_i < D_i(t_i)|M)$ from (13) changes to

$$P(\tau_i < T|M = m) = H\left[\frac{F_i^{-1}(Q_i(T)) - a_i m}{\sqrt{1 - a_i^2}}\right].$$

If both systematic and idiosyncratic factors have standard normal distribution, we get

$$P(\tau_i < T|M = m) = \Phi_{Z_i}\left[\frac{\Phi_{x_i}^{-1}(Q_i(T)) - a_i m}{\sqrt{1 - a_i^2}}\right].$$

From this approach we can compute the probability of $\binom{n}{k}$, $k \leq n$, defaults in the portfolio. If we deal with a homogeneous portfolio, the distribution of defaults is given by the well-known binomial distribution. Submitting subscript indices due to homogeneity, we arrive at the conditional probability for k out of n defaults until time T:

$$P(k|M = m) = \binom{n}{k} \cdot P(\tau < T|M = m)^k \cdot [1 - P(\tau < T|M = m)]^{n-k}.$$

By the law of iterated expectations, the probability of exactly k defaults is the average of the conditional probabilities of k defaults. The average is built over the possible realizations of M and weighted with the probability density function $\phi(m)$. If both factors have normal distribution, the probability of k defaults is

$$P(k) = \int_{-\infty}^{\infty} \binom{n}{k} \left(\Phi_Z\left(\frac{\Phi_X^{-1}(\cdot) - a \cdot m}{\sqrt{1 - a^2}}\right)\right)^k \left(1 - \Phi_Z\left(\frac{\Phi_X^{-1}(\cdot) - a \cdot m}{\sqrt{1 - a^2}}\right)\right)^{n-k} \varphi(m) dm.$$

With this formula, CDO tranches of a very simplified portfolio can be priced. It can easily be extended to several independent common factors. The independence of the common factors due to construction leads to simplifications: The multivariate distribution can be expressed as a product of the marginals. Numerical procedures have to be applied to realize the integration over the common factor. For heterogeneous portfolios other procedures have to be set up like the semi-analytic approaches by Hull and White (2004), Andersen, Sidenius, Basu (2003) or Laurent and Gregory (2003). The only free parameter in the factor Gaussian copula for homogeneous portfolios is

$a = \sqrt{\rho}$. The impact of a change of this parameter can most suitably be studied when considering ranked basket credit derivatives like n-th-to-default CDSs. For test reasons, Hull/White consider a homogeneous credit portfolio of 10 entities, with default intensity of 0.01 and deterministic recovery rates of 40% for all firms. They compute n-th-to-default CDS tranche spreads in basis points for $n = 1, \ldots, 10$ on this underlying portfolio. The cost of protection for different n can represent the correlation risk inherent in tranches as n-th-to-default contracts exhibit similarities to tranches of different seniority.[66] As correlation between the default variables is increased from $a = \sqrt{0.0}$ via $\sqrt{0.3}$ to $\sqrt{0.6}$, default correlation rises with the effect of lowering cost of protection for 'subordinated' n-th-to-default CDS and lifting cost of protection for high rank n-th-to-default CDS. It can be seen from this table:

n	$a = \sqrt{0.0}$	$a = \sqrt{0.3}$	$a = \sqrt{0.6}$
1	603	440	293
2	98	139	137
3	12	53	79
4	1	21	49
5	0	8	31
6	0	3	19
7	0	1	12
8	0	0	7
9	0	0	3
10	0	0	1

Table 1: Spreads in basis points for n-th-to-default CDSs for different correlations and constant default intensity $\lambda = 0.01$. Source: Hull and White (2004), exhibit 3, p.13.

The impact of correlation on tranches of different seniorities can be observed in table 2.

iTraxx IG	0-3%	3-6%	6-9%	9-12%	12-22%
a=$\sqrt{0.0}$	44.3%	69	0	0	0
a=$\sqrt{0.15}$	31.5%	258	64	18	2
a=$\sqrt{0.25}$	24.5%	294	110	49	11
a=$\sqrt{0.3}$	21.2%	300	127	64	18
Market Quotes	27.6%	168	70	43	20

Table 2: Market quotes for Spreads if iTraxx index tranches. The quotes are in basis points and the equity quote is a percentage notional functioning as upfront fee. Source: Hull and White (2004), exhibit 12, p. 20.

An additional observation can be made here if market quotes and the predictions of the factor Gauss copula with different copula correlations are considered. The model

[66] See Hull and White (2004), p. 17.

exaggerates cost of protection for intermediate tranches while cost of protection for equity and very senior tranches are underestimated. This is distinctively expressed as the well-known correlation smile. Therefore, to match market data a model is required that increases cost of protection for equity and very senior tranches and reduces cost for intermediate tranches.[67] As empirical findings show, distributions producing more outcomes of extreme events are more appropriate in financial modeling than the Gaussian distribution which is regarded as light-tailed in the mathematical sense.[68] In the following we will present such an extension with Student-t distributions exhibiting heavy tails.

4.3 The Double Student-t Copula

The presented multivariate firm-value model can be extended by certain thick-tailed distributions. It can be conjectured from empirical findings that the dependence structure between asset returns is much more complex than represented by the Gaussian copula - this approach only represents the degree of linear dependence. For this reason we introduce the concept of tail-dependence which is a local dependence measure that reflects nonlinear functional dependencies and especially the amount of extreme outcomes. The correlation coefficient lacks the property of invariance under increasing changes of variables.[69] This means that a strictly increasing transformation of the marginal variables influences the dependence measure. Since the copula of a multivariate distribution uniquely describes its dependence structure, copula-based measures of dependence are favored. The copula as a unique and intrinsic dependence measure is scale-invariant and so is the tail dependence.[70] Demonstratively, the bivariate tail dependence measures the amount of dependence in the upper and lower quadrant tail of the distribution. Technically, the probability of the first variable to be large is measured, conditional on the second variable being large, too.

In the context of a factor model, the tail dependence between the default variable x_i and the systematic and idiosyncratic factors is of interest to understand the dependence structure. Malevergne and Sornette (2004) analyse these tail dependence effects in factor models. They only consider the tail dependence between the default variable and the systematic factor with fixed factor loading, but their dependency analysis can be extended to the idiosyncratic factor as well. The authors show that there is absence of tail dependence, if the systematic factor distribution is rapidly varying, which is the case for the Gaussian, exponential and any distribution decaying faster than any power-law. For regularly varying factors they provide a factor tail-dependence formula.

[67] See Hull and White (2004), p. 19.
[68] See Rachev and Mittnik (2000).
[69] See 'Financial Risk and Heavy Tails' by Bradley/Taqqu in Rachev, S., ed. (2003).
[70] See Embrechts, Lindskog, McNeil (2003).

It applies for any power-law tail decay behavior like the Student-t or α-stable distributions which will we outlined in chapter 6.1.
The mechanism of action in the standard model in combination with heavy-tailed distributions will be outlined now. As an example we chose the one-factor Student-t copula introduced for CDO pricing by Hull and White (2004). The Student-t distribution - just like the normal distribution - belongs to the class of elliptical distributions which means that the multivariate distribution is uniquely defined by the variance-covariance matrix.[71] The distribution has a degree-of-freedom parameter ν and as $\nu \to \infty$ the normal distribution is recaptured. For $\nu < \infty$ the distribution exhibits regularly varying functions for the tails which allows to categorize distributions as heavy-tailed. The smaller the degree-of-freedom parameter, the heavier the tails. For $\nu > 2$ it exhibits a finite variance of $\nu/(\nu - 2)$.
Again, a homogeneous credit portfolio of 10 entities, with default intensity of 0.01 and deterministic recovery rates of 40% for all firms is considered. The authors test Student-t distributions of different degrees of freedom for the systematic and idiosyncratic factors. The label 'double' Student-t copula is deduced from this construction as both factors are Student-t distributed. Variations are conducted assuming either 5 or ∞ for the degree of freedom for the systematic/idiosyncratic factor distributions considering all combinations. The results can be seen in table 4.3.

n	∞/∞	$5/\infty$	$\infty/5$	$5/5$
1	440	419	474	455
2	139	127	127	116
3	53	51	44	44
4	21	24	18	22
5	8	13	7	13
6	3	8	3	8
7	1	5	1	5
8	0	3	0	4
9	0	2	0	2
10	0	1	0	1

Table 3: Impact of different combinations of Student-t distributions for the systematic and idiosyncratic factors on n-th-to-default CDS spreads in bps on an underlying homogeneous credit portfolio of 10 entities, with default intensities of 0.01 and deterministic recovery rates of 40% for all firms.

The authors deliver their own interpretation of these results: Using heavier tails for the systematic factor and keeping the idiosyncratic factor normally distributed ($5/\infty$) lowers the cost of protection in an n-th-to-default CDS if n is small, and increases it if n is large in comparison to the base case ∞/∞. Using heavier tails for the idiosyncratic

[71] See Bluhm et al. (2003), p. 106f.

risk and keeping the systematic factor normal ($\infty/5$) raises the cost of protection in an n-th-to-default CDS if n is small, and lowers it if n is large.

This can be demonstratively shown in the following way: First of all, the properly normalized factor model for Student-t distributed systematic and idiosyncratic factors is

$$x_i = a_i\sqrt{\frac{\nu-2}{\nu}}M + \sqrt{1-a_i^2}\,\sqrt{\frac{\widetilde{\nu}-2}{\widetilde{\nu}}}Z_i \qquad \forall i \leq n, \tag{15}$$

where M is the systematic Student-t distributed factor with ν degrees of freedom, and Z_i with $\widetilde{\nu}$ is the idiosyncratic Student-t distributed factor, respectively. It has to be remarked here, that the linear correlation between x_i and x_j is the same as in the baseline model due to the normalization of the Student-t distributions. Due to construction, the linear correlation between the default variables is still consistent with the data of the variance-covariance-matrix underlying the factor structure and thus the calibration features for the default barriers are maintained despite the selection of a different factor distribution, concerning ν and $\widetilde{\nu}$.

Now, for one demonstrative simulation procedure, two consecutive drawings of random systematic and idiosyncratic factors are undertaken and weighted with the respective factor loadings to receive an outcome for the normalized Student-t distributed variable x_i. When M has heavy tails and the Z_i are normal ($5/\infty$), an extreme value for a particular x_i is more likely to arise from an extreme value of M than an extreme value of Z_i. It is therefore more likely to be associated with extreme values for other x_i. On the one hand, this should increase default correlation, since one extreme negative outcome will draw a large number of credits below the default barrier. The normally distributed Z_i will be able to draw the default variable above the barrier in very rare cases, since the normal distribution almost never produces extreme outcomes. On the other hand, an extreme positive outcome of M has the effect, that almost no credit defaults, which would result in a reduction of default correlation. As can be seen from the case $5/\infty$ in table 4.3, the increasing default correlation dominates, since n-th-to-default CDS spreads are higher for $n = 4, \ldots, 10$ and lower for small $n = 1, \ldots, 3$.

Assuming now the case, that M is normally distributed and the Z_i have heavy tails, an extreme value for a particular x_i is more likely to arise from an extreme value of the Z_i than an extreme value of M. It is therefore less likely to be associated with extreme values for the other x_i. Assuming now a drawing from the normal systematic factor does not result in an extreme value (as almost always the case), defaults have to come from negative extreme outcomes of the idiosyncratic factors, which might appear more often, since this factor has heavy tails. However, the overall effect is a decrease of default correlation, since most of the default variables x_i will be in a normal region not indicating default, and only a few extreme outcomes will cause single credits to default with most of the other credits still being alive. This can be stated by the cost

of protection in column $\infty/5$ in table 4.3. Cost of protection rises sharply for the first-to-default CDS, declines for $n = 2, \ldots, 5$, and is the same for the rest, always compared to the base case ∞, ∞.

Now a final step is to choose both factors as heavy tailed. A remarkably observation can be made here: Compared to the base case, cost of protection for the first-to-default CDS rises, prices for n-th-to-default CDS with $n = 2$ and 3 are lower, and prices for $n = 4, \ldots, 10$ rise again. This phenomenon has direct connection to the correlation smile known from the standard one-factor Gauss copula, if we interpret the first-to-default CDS as an equity tranche, the second- and third-to-default as mezzanine, and the n-th-to-default CDS with $n = 4, \ldots, 10$ as senior. According to Hull and White (2004), to match market data, it requires a model that increases the fair spread for the equity and senior tranches and reduces it for intermediate tranches.[72] The only model in their paper that accomplishes this fit is the introduction of heavy tails for both factors. This special and favorable effect on the tranches is even tractable by different degrees of freedoms for the factors. The authors suggest 4 degrees of freedom for both factors after their calibration tests. They report good calibration results which is underlined by the analysis of (Burtschell et al. (2005)), where the so-called double t-copula model is attested a good market fit as well.

The authors show that the same degrees of freedom for both factors deliver the best empirical fit to quotes of synthetic standard tranches. This fact also results in improved algorithmic implementation features, as Student-t distributions with the same degree of freedom exhibit sum stability. So for barrier computations the quantiles of the x_i can easily be found.

The overall effect of two heavy-tailed factors can also be recognized from the resulting default correlation in table 4.3. We simulate the default correlation between two entities for a variation of the degrees of freedom, provided that they are identical for both factors in all cases. Factor loadings of $\sqrt{0.3}$ and default intensities of 0.01 are assumed.

degrees of freedom	default correlation
∞	0.0475
7	0.0605
6	0.0795
5	0.0832
4	0.1143
3	0.1576

Table 4: Variation of the degrees of freedom - being identical for both systematic and idiosyncratic factors - in 500.000 simulations for each case with factor loadings of $\sqrt{0.3}$ and default intensities of 0.01.

[72] See Hull and White (2004), p.19.

For lower degrees of freedom the default correlation is raised, with the linear correlation between the default variables being the same. This is due to tail dependence effects not captured by parametric measures of dependence.[73] In the double Student-t copula model however, the increased default correlation has a different origin, since we keep constant the factor loadings and thus the linear correlation between default variables. In this approach, a completely different default behavior is generated consisting of many scenarios where a large number of obligors survive and quite a few single defaults happen which raises cost of first-loss protection. And there are scenarios of default clustering with quite a number of entities in default represented by higher cost of high-rank protection. Scenarios with just a few defaults are rare since cost of intermediate loss levels is lower compared to the base case (∞/∞).
This generation of a specific default behavior leads to a better fit to market data as can be seen from the following table 5: The interplay between two heavy tailed distri-

iTraxx IG	0-3%	3-6%	6-9%	9-12%	12-22%
market data	27,6%	168	70	43	20
double Student-t copula	25,5%	171	69	42	23
one-factor Gauss copula	21,2%	300	127	64	18

Table 5: Spreads predictions of iTraxx tranches by means of double Student-t copula and one-factor Gauss copula with constant correlation of 0.3 in bps with 5-year maturity.

butions leads to the desired behavior of lifting cost of protection for equity and senior tranches and lowering cost for intermediate tranches. If exact calibration to all standard index tranches was possible, there would be no implied correlation smile and no common level of correlation implied from the quotes of standardized tranches.[74] The strict separation of factor loadings and the other degrees of freedom leaves many opportunities to calibration while simultaneously preserving consistency to the underlying CDS data.
The only model in the analysis that accomplishes this fit is the introduction of heavy tails for both factors. This special favorable effect on the tranches is tractable by different degrees of freedom assigned to the factor distributions. The authors suggest 4 degrees of freedom for both factors after their calibration tests.
They report good calibration results which is underlined by the analysis of Burtschell et al. (2005), where the double Student-t copula model is attested a good market fit in a relation to a variety of other models like simple Student-t, Clayton and Marshal-Olkin.

[73] See Malevergne and Sornette (2004) for tail dependence effects in factor models.
[74] See Duffie (2004), for example.

There is a high degree of tractability in calibration: Heavy tails increase cost of protection for equity and senior and decrease it for intermediate tranches. On the other hand, when increasing the factor loading, cost of protection rises for low tranches and rises for higher more senior-like tranches. In this way the 'curve' of the 5 quotes of standard tranches can be fit by different mechanisms, influenced by the choice of parameters which are factor loading and degree of tail-heaviness.

These insights will be picked up in chapter **??** for smoothly truncated stable distributions. We will now present the structural model of Hull, Predescu, White (2005). It can be seen as a dynamic generalization of the standard model.

5 The Hull/Predescu/White Model

Investors of synthetic CDO tranches can be seen as protection sellers as they suffer from credit portfolio losses if these occur in the percentile range of the specific tranche. The compensation payments to the buyer of protection are represented by the initial tranche nominals being scaled down by respective portfolio losses so that the outstanding tranche principals are stochastically decaying.

On the other hand, protection sellers receive a periodic premium for risk bearing which is expressed as an annual spread payed on the outstanding tranche nominal. The actual payments are made quarterly as fractions of the fixed annual spread. Since losses may reduce tranche nominals, actual premium payments are stochastically decaying in time. At the beginning of the transaction both parties - protection buyer and protection seller - ideally decide to trade at a specific spread that is fixed for the life of the CDO. Under expectations about the future loss behavior of the reference portfolio, they agree on the so-called fair spread that is the break-even of the payments to be made by both parties: loss compensation and premiums. A pricing approach of such deals is presented in chapter 7 in detail.

In the valuation of synthetic CDOs, it is convenient to assume discrete default times. The pricing thus requires at least 20 joint loss distributions at the coupon dates, since coupons are paid quarterly at the payment days $t_1 = \frac{1}{4}$ year to $t_{20} = 5$ years. This means that each tranche has a different exposure to the timing of defaults, so the default development in time has to be taken into consideration when pricing CDOs.

In the standard market model described in the last chapter, the realization of the common factor M at time zero governs the default outcomes of all future time periods so that the default environment is the same for the whole life of the CDO.[75] This means that loss distributions are always generated as seen from time t_0: Imagine we simulated the loss distribution for the time horizon $[t_0, t_1]$, then we knew the outstanding tranche nominals at time t_1. Additionally, the resulting payments by protection seller and protection buyer could be determined. Now, for the next coupon payment in t_2 the change of the outstanding tranche nominals since the last coupon payment date t_1 has to be specified in the simulation. In the one-factor Gauss copula the loss distribution for the time horizon $[t_0, t_2]$ is simulated with the same M. The difference between the two generated loss distributions thus governs the outstanding tranches' nominals between $[t_1, t_2]$ which in turn affects the cash flows of seller and buyer. The standard one-factor Gaussian copula model can therefore be regarded as static since there is no dynamic evolution of underlying credit spreads.

The structural model of Hull, Predescu, White (2005) however is much richer because

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

a bad default environment in one time span can be followed by a good default environment in the next time span. This is accomplished by an approach that was introduced by Black und Cox (1976) that resembles an extension of Merton's model. Their 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 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 co-product of this procedure, the joint evolution of correlated credit spreads is obtained, so dynamic hedging of certain credit portfolio positions and consistent pricing of derivatives of CDOs is achieved.
In accordance to the standard model, the HPW approach is set up with a diffusion process for the value of the firm:

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

Now X_i has a different meaning:

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

The variable $X_i(t,T)$ has the distribution

$$X_i(t,T) = W_i(T) - W_i(t),$$

so it is not normalized by $\sqrt{T-t}$. This variable can be imagined as some function of the value of the assets or the creditworthiness of company i. The latter can alternatively be imagined as the usual discrete credit ratings that are replaced by continuous measures. So X_i is some function of the default measure of company i.
The resulting barrier equation is the following:

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

The default barrier is a function of time to make the model consistent with exogenously specified initial default probabilities. According to Hull and White (2001), it can be argued that one reason for company i's default barrier being a function of time is that its capital structure is more complicated than in the simple Merton setting. This

introduces some non-stationarity into the default process and is a price that must be paid to make the model consistent with the risk-neutral default probabilities.
The continuous default barrier methodology was intensively studied by Avellaneda and Zhu (2001). They show that the calibration of the default barrier leads to a free boundary problem for the associated partial differential equation. Also, they propose a new interpretation of the default barrier model in terms of a risk-neutral distance-to-default process for the firm, and show that the calibration procedure to a set of default probabilities is equivalent to specifying an appropriate ´excess drift' for the distance to default process. As soon as the model is fully calibrated, this excess drift can be interpreted as a market price of risk for the value of the firm that is associated with the firm's perceived creditworthiness, consistently with Merton (1974).[76]
Hull and White (2001) present a discretized version of the model that can be solved numerically. This is favorable for extensions with different distributions that do not exhibit closed-form expressions, as we will see in the case of smoothly truncated stable distributions in chapter **??**.
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 measures provide the same information about the default probability environment and they are related by

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

when the exponential model for the default time distribution is employed. We later assume for our extensions 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:

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

Since the Black/Cox default barrier methodology requires interpolation of the risk-neutral default probabilities retrieved from CDS or credit spread data, it has to be taken care of, that different interpolation methods for the default probability density function do not produce significant changes in the barriers generated by the model. In our calibration procedure we use the simple exponential model to back out a default probability distribution function. Due to this assumption, we can derive the representative default intensity λ from the quoted CDS index spread, as shown in section 7.1. This brings us into the position to compute 'intermediate' default time probabilities

[76] See Avellaneda and Zhu (2001), p. 2-5.

like $Q(t, t + \Delta t)$ to avoid the interpolation of risk-neutral default probabilities.

5.1 Construction of the Discrete Default Barriers

After this preprocessing we will outline the discrete time version of the Black/Cox default barrier methodology by Hull and White (2001). By increasing the time grid points within the CDO horizon, this model can be made arbitrarily close to a model where defaults can happen at any time. Avellaneda and Zhu (2001) find, that the barriers in the continuous model are slightly below those in the discrete model which is due to the fact that defaults can only happen at discrete points in time.
Conveniently, the time grid is synchronized to the coupon payment days t_j, $j = 1, \ldots, J$. This means that the default probability distribution is discretized so that defaults are modeled to happen at times t_j, but they are associated with the midpoints $\frac{t_{j-1}+t_j}{2}$ in the pricing part.
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 mus 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 $D_i(t_j)$,
- $f_{ij}\Delta X_i(t_j)$ 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 that

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

Both D_{ij} and $f_{ij}(x)$ can be determined inductively from the $Q_i(t_{j-1}, t_j)$. The first barrier is found by the first increment of $X_i(t_1)$ that has 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 the process in t_1 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. All values and distributions are known except for the barrier D_{i2}. In our algorithm we find an approximation to the solution by a bisection method, where we declare an initial search region that is iteratively split into halves to identify the new search region until a certain tolerance level is met. The general equation for times t_2 to t_J is

$$Q(t_{j-1}, t_j) = \int_{D_{i,j-1}}^{\infty} f_{i,j-1}(u)\ \Phi\left(\frac{D_{ij} - u}{\sqrt{\delta}}\right) du. \tag{16}$$

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

$$f_{ij} = \int_{D_{i,j-1}}^{\infty} f_{i,j-1}(u)\ \varphi\left(\frac{x - u}{\delta}\right) du. \tag{17}$$

Equations (16) and (17) 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 $X_i(t_j)$ $(1 \leq k \leq K)$ and π_{ijk} as the probability that $X_i(t_j) = x_{ijk}$ with no earlier default. The discrete versions of equations (16) and (17) 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 $x_{i,j-1,k}$.

5.2 Simulation and Dynamic Credit Spreads

There exists an analytic expression of first hitting the barrier between times t and $t + \Delta t$.[77] When suppressing indices we have:

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

By the discrete barrier procedure, the continuous barrier in the Black/Cox model is replaced so that it is only defined at times $t_j,\ j = 1, \ldots, J$. The default probabilities in each interval (t_{j-1}, t_j) are the same as those for the continuous barrier.

Under a recovery assumption the equation can be used to determine a credit spread or a CDS spread. At each t_j the remaining default probability of each credit until maturity can be calculated and translated into a credit spread. In this way, there is a dynamic evolution of the underlying credit spread at discrete points in time during the life of the CDO.

The process for the 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, on each simulation trial there are discrete paths sampled for M and Z_i for each time increment $\delta_j = t_j - t_{j-1}$. We assume that time progresses in equidistant intervals δ so that the subindex j can be omitted:

$$\Delta X_i = a_i \Delta M + \sqrt{1 - a_i^2} \Delta Z_i, \tag{18}$$

where ΔM and ΔZ_i are independent normal distributions with zero mean and a variance of δ.

It has to be remarked here, that the variables a_i, M and Z_i have a different meaning than in the one-factor Gaussian copula approach. Nevertheless, the correlation between the processes followed by the assets of companies i_1 and i_2 is $a_{i_1} a_{i_2}$.

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

The authors comment, that there are some reasons to suppose that the correlation environment generated by the factor Gaussian copula is a reasonable approximation to the HPW extension.[78] With the simplifying assumption that once the companies' assets drop below the barrier they remain less than the barrier, the Gaussian copula and the HPW model give the same joint probabilities of default for the same factor loadings. Even for relatively large credit spreads the prices given by the two models are fairly close. So the standard model seems a reasonable approximation for the HPW model in case of the same asset correlation.

It has to be mentioned that this approximation is only valid for the Gaussian case as the only source of default dependence is the linear correlation coefficient. If we employ distributions for the process that exhibit heavy tails or produce considerable amounts of extreme events, there is clearly a difference between the two approaches. For example, the Student-t extension to the standard model outlined above exhibits tail dependence effects. For different time horizons during the CDO's life the default variable is always characterized by the sum of two Student-t factor distributions. In the HPW model extension there is a discrete process with each increment being the sum of two heavy-tailed distributions. So the occurrence of extreme events enters the model at multiple stages.

The equation (**??**) for the computation of the dynamic credit spreads cannot be used for distributions lacking closed-form expressions. Nevertheless, if the process X_i has reached a certain position in one future scenario generation at some time $t_j < t_J$, the remaining probability to hit one of the discrete barriers until maturity t_J can be computed by the discrete barrier algorithm. This in turn allows to compute the credit spreads at each t_j so the model is capable of displaying the joint evolution of dynamic credit spreads.

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

In the base case HPW model, the correlation parameter and the recovery rate are independent of the systematic factor M. The authors apply two extensions that introduce stochastic correlation and stochastic recovery rates.[79] This is realized by a coupling of the default environment M with those parameters. There is empirical evidence to suggest that asset correlations are stochastic and increase when default rates are high. On the other hand, empirical findings show that recovery rates are negatively depen-

[78] See Hull, Predescu, White (2005), p. 11.
[79] See Hull, Predescu, White (2005), p. 14,15.

dent on default rates.[80]

According to expectations, the base case HPW model as the generalization of the standard model does not provide a good fit to market prices of CDO tranches. The authors calibrate their model to the market quote of the equity tranche as it reacts most sensitive to correlation: The mezzanine tranche is greatly overpriced and the super-senior tranches are highly underpriced.

The advanced model of stochastic correlation leads to a significantly better fit at the 1% level.[81] The authors suggest that further research quantifying the relationship between asset correlations and default rates could lead to better pricing models. An extension of the basic strutural model to incorporate stochastic recovery rates produces a marginally better fit to market prices and the authors comment that it can be regarded as statistically - but not economically - significant.

In our extension we abstain from extensions like stochastic correlations and recoveries but concentrate on a distribution class producing more extreme events than in the Gaussian case. The following chapter will engage in this approach. Nevertheless, our extension could be further improved by the introduction of stochastic correlation and stochastic recovery.

[80] For these findings see for example Andersen und Sidenius (2004) and Altman et al. (2002).

[81] See Hull, Predescu, White (2005), p. 17.

6 An Extension with Smoothly Truncated Stable Distributions

The factor structure used throughout credit portfolio models preserves the separability between factor loading and other model parameters like the tail index representing tail-thickness. This affords a degree of robustness to possible errors in distributional assumptions. For example, the default barriers of an obligor in the HPW model can be calibrated consistently to the (risk-neutral) default probabilities of different time horizons due to normalization arrangements of the heavy-tailed distribution, even if the systematic and idiosyncratic risk factor distribution assumptions are incorrect.

There exist certain prerequisites for arbitrary distributions to be used in the outlined Black/Scholes/Merton framework for asset pricing. The distributions employed have to possess finite moments of arbitrary order which then also applies to exponential moments. The latter property is important for log-returns of financial assets in order to guarantee a finite mean for the firm-value and - as a consequence - finite prices for options on the firm's assets which resembles the character of credit portfolio loss tranches.

Since the seminal work of Fama (1965) and Mandelbrot (1963) it is widely accepted among researchers, that asset return distributions are asymmetric, highly peaked around the mean and heavy tailed.[82] Various authors like e.g. Rachev and Mittnik (2000) have suggested to replace the normal by the stable distribution. They have turned out to meet statistical characteristics of financial data with high accuracy to empirical observations.

The class of stable distributions forms an ideal alternative to the Gaussian law due to their combination of the stability property with modeling flexibility. The class disposes of parameters directly affecting the fundamental shape of the distribution, like skewness and probability mass in the tails. The main drawback, however, of stable non-Gaussian distributions is their infinite variance.[83]

Therefore, we employ a new class of probability distribution called smoothly truncated alpha-stable distributions (STS) family. It combines the modeling flexibility of stable distributions with the existence of arbitrary moments and therefore qualifies for applications in the Black/Scholes/Merton framework.

In this chapter, we will first outline the characteristics of α-stable distributions and we present the method of 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 additionally provided

[82] See for example Rachev and Mittnik (2000).
[83] See Menn and Rachev (2004a), p. 3.

by their approach. This serves as a building block since α-stable cdf evaluations facilitate the smooth truncation and standardization of STS distributions. Finally we present a method to simulate STS distributions based on the method to generate α-stable samples by Chambers et al..

Further information about STS distributions and their application in option pricing can be found in Menn and Rachev (2004a), Menn and Rachev (2005a) and Menn and Rachev (2005b).

6.1 The Stable Distribution Family

Since the pioneering work of Mandelbrot and Fama, the normal assumption is rejected and a more general family of distributions is proposed - the class of stable laws.[84] The Gaussian distributions resembles a limiting case of the stable distribution family. Stable non-Gaussian laws are also called stable Paretian or Lévy stable. This is due to the Pareto-like power decay of the tail of the distribution and due to Paul Lévy who carried out the fundamental research of characterizing the family of non-Gaussian stable laws.

It is argued that financial asset returns are the cumulative outcome of a vast number of pieces of information and individual decisions arriving almost continuously in time.[85] As such, since the pioneering work of Louis Bachelier in 1900, they have been modeled by the Gaussian distribution. The strongest statistical argument for it is based on the Central Limit Theorem (CLT), which states that the sum of a large number of independent, identically distributed variables from a finite-variance distribution will tend to be normally distributed. However, underlined by empirical research, financial asset returns usually have heavier tails and thus the *generalized* CLT, which states that stable laws are the only possible limit distributions for properly normalized and centered sums of independent, identically distributed random variables, has to be applied. In other word, if the sum of a large number of iid random variables has a non-degenerate limiting distribution after normalizing location and scale, the limiting distribution must be a member of the stable class.

Stable distributions are attractive because they have the desirable 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.

Another attractive feature is the stability property. There exists an important shape

[84] See Racheva-Iotova and Stoyanov (2004).
[85] See Rachev and Mittnik (2000).

parameter which governs the properties of the distribution. It is called the index of stability and is denoted by α. 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 its own domain of attraction. The stability property is a desirable property in portfolio theory as a portfolio of assets with α-stable returns has again α-stable returns. Additionally it reduces the computational burden when summing α-stable distributions.

The α-stable distribution requires four parameters for complete description: an index of stability $\alpha \in (0, 2]$ also called the tail index, tail exponent or characteristic exponent, a skewness parameter $\beta \in [-1, 1]$, a scale parameter $\sigma > 0$ and a location parameter $\mu \in \mathbb{R}$. The tail exponent α 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.

Figure 2 illustrates the effect of changes in α on the shape of an α-stable probability density function in the center of the distribution and figure 3 illustrates the tail

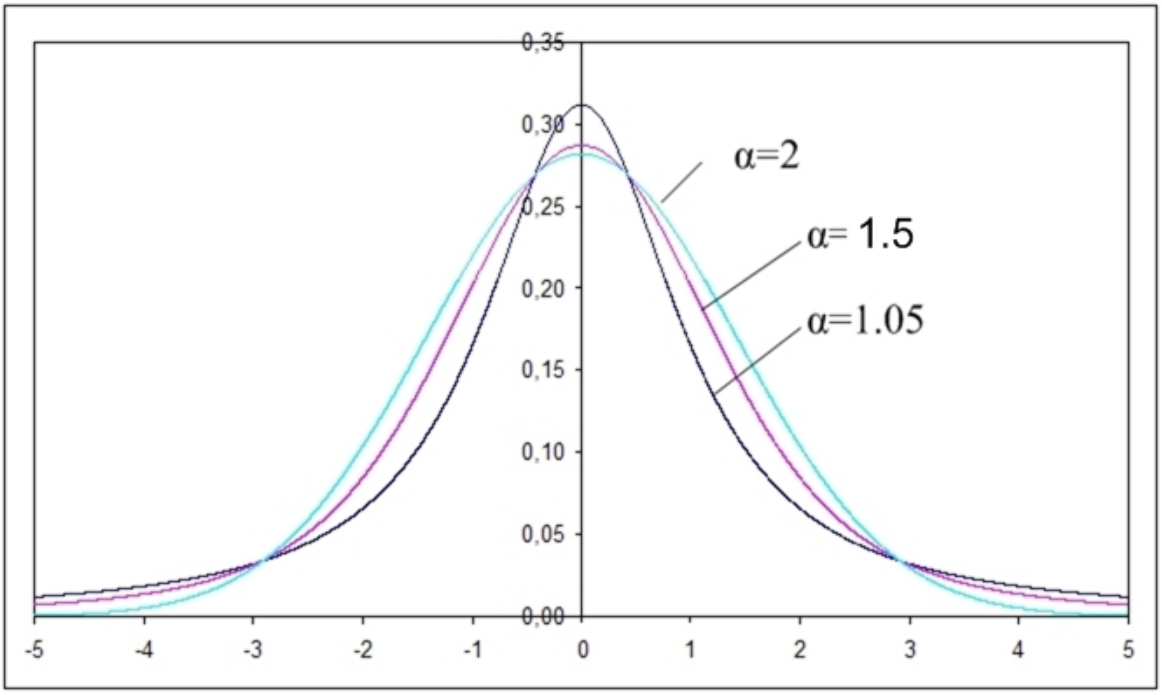

Figure 2: Shape of α-stable distributions in the center part. Variation of α.

behavior. A change in α does not simply affect the 'spread' of the distribution, as would a change in variance.[86] Rather, it moves probability mass between the shoulder and the tails of the distribution. Distributions with α parameters that lie between one and two have unbounded variance but bounded mean. Those with parameters that

[86]See Gordy and Heitfield (2001), p.12.

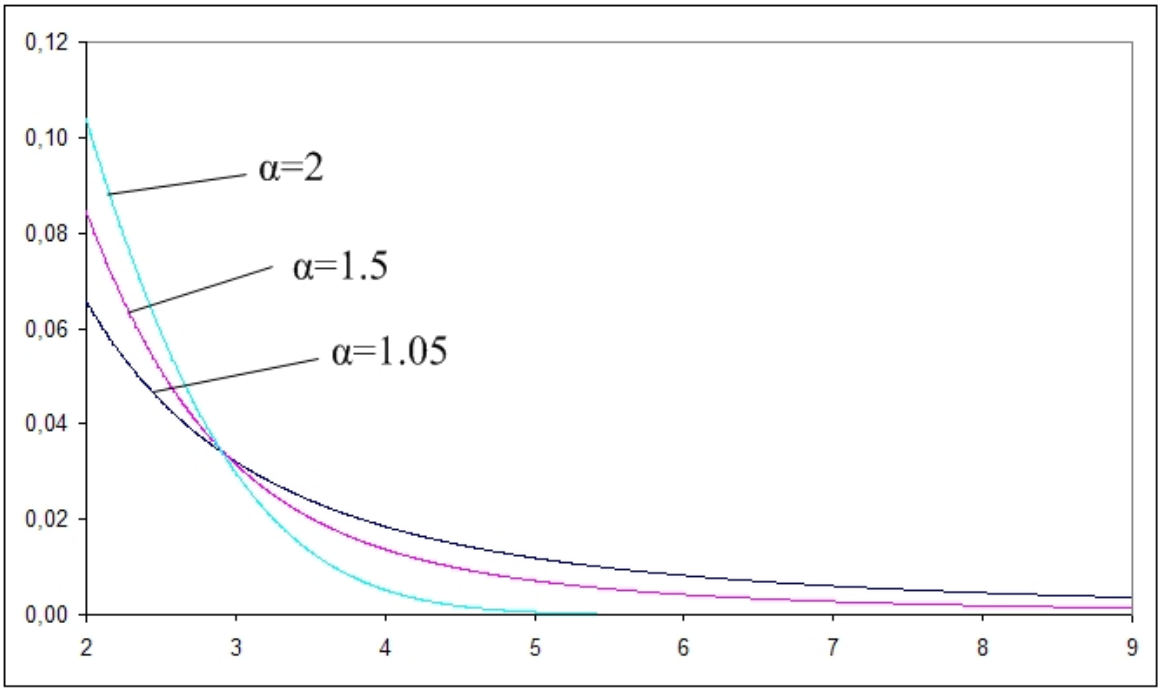

Figure 3: Shape of α-stable distributions in the tails. Variation of α.

lie between zero and one have both unbounded variance and mean. The influence of the skewness parameter β is visualized in figure 4. Unfortunatly, the application of stable laws in finance is handicapped by the lack of closed-form expressions of their propability density and cumulative distribution functions - with the exception of three special cases. These exceptions include the well known Gaussian ($\alpha = 2$) law, whose density function is given by:

$$f_G(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-\frac{(x-\mu)^2}{2\sigma^2}\right\},$$

and the lesser known Cauchy ($\alpha = 1$, $\beta = 0$) and Lévy ($\alpha = 0.5$, $\beta = 1$) laws. There exist numerous density approximations by different authors and the oldest approach was developed by Zolotarev (1966). It is based on an integral representation of the density function and it was extended by Nolan (1997) 1997 who offers a free download of the Fortran program.[87] This program called STABLE.exe exhibits high relative accuracy for default tolerance settings: About 10^{-13} for values used in typical financial applications like approximating asset return distributions. Naturally, the accuracy of the program can be increased at the cost of computational time.

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. Much confusion has been caused by these different representations. The variety of formulas is caused by a combination of historical evolution and the numerous

[87] See http://academic2.american.edu/ jpnolan/stable/stable.html

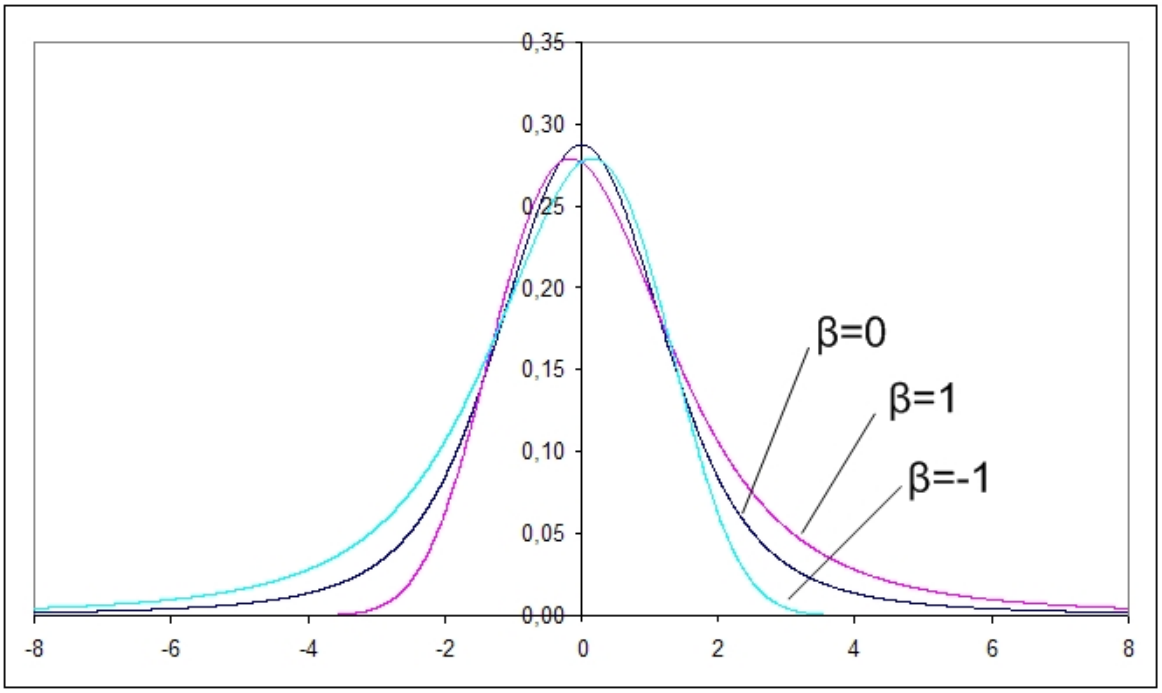

Figure 4: Shape of α-stable distributions in the center part. Variation of β.

problems that have been analyzed using specialized forms of the stable distributions. The most popular parameterization of the characteristic function of $X \sim S_\alpha(\sigma, \beta, \mu)$, i.e. an α-stable random variable with parameters α, σ, β and μ, is given by Samorodnitsky and Taqqu (1994):

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

Note, that the traditional Gaussian scale parameter $\sigma_{Gaussian}$ is not the same as σ in the above representation. The relation is the following: $\sigma_{\text{Gaussian}} = \sqrt{2}\sigma$. The representation above is discontinuous at $\alpha = 1$ and $\beta \neq 0$, but this drawback is not an inherent property of the class of α-stable distributions and appears because of the special form of the characteristic function. There exist other representations for statistical and numerical work, but we employ this representation since the α's range for financial data usually is $1.6 < \alpha < 1.9$.

6.1.1 Basic Properties of Stable Distributions

Property 1. *(Scale) Let $X \sim S_\alpha(c, \beta, \mu)$ and $a \in \mathbb{R}$, $a \neq 0$. Then*

$$aX \sim S_\alpha(|a|c, \text{sign}(a)\beta, a\mu), \qquad \alpha \neq 1.$$

Property 2. *(Shift) Let $X \sim S_\alpha(\sigma, \beta, \mu)$ and $a \in \mathbb{R}$. Then $X + a \sim S_\alpha(\sigma, \beta, \mu + a)$.*

These two properties lead to the scale and shift invariance principles of the stable law. It means that for computational implementation it is sufficient to implement the standardized stable density, as for a random variable $X \sim S_\alpha(\sigma, \beta, \mu)$ it is valid that $(X - \mu)/\sigma \sim S_\alpha(1, \beta, 0)$.

Property 3. *(Sum of α-stable distributions) Let X_1 and X_2 be independent random variables such that $X_1 \sim S_\alpha(\sigma_1, \beta_1, \mu_1)$ and $X_2 \sim S_\alpha(\sigma_2, \beta_2, \mu_2)$. Then $X_1 + X_2 \sim S_\alpha(\sigma, \beta, \mu)$, with*

$$\sigma = (\sigma_1^\alpha + \sigma_2^\alpha)^{1/\alpha}, \qquad \beta = \frac{\beta_1 \sigma_1^\alpha + \beta_2 \sigma_2^\alpha}{\sigma_1^\alpha + \sigma_2^\alpha}, \qquad \mu = \mu_1 + \mu_2.$$

Property 4. *(moments of an α-stable distribution) Let $X \sim S_\alpha(\sigma, \beta, \mu)$, $0 < \alpha < 2$ and $p \in \mathbb{R}_+$. Then*

$$\begin{aligned} \mathbb{E}|X|^p &< \infty \quad \text{for any} \quad 0 < p < \alpha \\ \mathbb{E}|X|^p &= \infty \quad \text{for any} \quad \alpha \leq p. \end{aligned}$$

As a consequence, if $\alpha \leq 1$, then the corresponding α-stable distribution does not have a finite first moment. For $1 < \alpha < 2$ however, the location parameter μ equals the mathematical expectation of $X \sim S_\alpha(\sigma, \beta, \mu)$ due to the existing first moment. Additionally, these properties show that a stable non-Gaussian distribution has infinite variance and this leads to the concept of smooth tail truncation.[88] It allows for the preservation of the properties of the stable law in the 'center' of the distribution, whereas an exponentially declining function replaces the power decaying tails of the stable law in order to guarantee the existence arbitrary moments. Before we will explain the construction, the properties and implementational related aspects of the STS distribution, an efficient algorithm for density approximations for stable non-Gaussian distributions will be outlined.

6.1.2 Density Approximation of Stable Distributions

The density function/characteristic function relationship can be exploited by the fast Fourier transform (FFT). For data points falling between the equally spaced FFT grid nodes, an interpolation technique has to be used. Concerning the computational speed, the FFT based approach is faster for large samples, whereas the direct integration method favors small data sets since it can be computed at any arbitrarily chosen point. The FFT based approach is not as universal as the direct integration method

[88] See Menn and Rachev (2005b), p. 3.

- 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.

Numerical approximation or direct numerical integration lead to a drastic increase in computational time and loss of accuracy.

We therefore decided to implement a simplified version of the calibrated FFT-based density approximation by Menn and Rachev (2004b). The authors employ an adaptive Simpson rule for the quadrature of the Fourier inversion integral. Since this approach lacks of 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 get 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 it is scale and translation/shift invariant. As the computational effort of the FFT-routine is virtually independent of the number of evaluations on the equidistant grid, desired levels of accuracy should determine the computational cost. In comparison to Nolan's program, the approach results in a significant reduction of the computation time while simultaneously preserving satisfactory accuracy which is comprehensively quantified in their study. Our simplified version will be outlined and resulting accuracy will be discussed now.

The density approximation of the Fourier inversion integral is convergent and can therefore be restricted to the compact interval

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \phi(t) dt \approx \frac{1}{2\pi} \int_{-a}^{a} e^{-itx} \phi(t) dt. \tag{19}$$

This integral can be evaluated using various quadrature formulas. There is a classical algorithm in Rachev and Mittnik (2000), e.g., that uses the left point algorithm (LPA) for integral approximations. First, the authors determine the accuracy of the midpoint algorithm (MPA) in comparison to the highly accurate counterparts calculated by Nolan's program. They find that the MPA is about twice as good as the LPA. As a next step, an evaluation of the density with MPA and LPA for only $N/2$ grid points is undertaken. The results are weighted setting $f(x) = \frac{1}{3} f_{LPA}(x) + \frac{2}{3} f_{MPA}(x)$. This approach can be seen as the application of the Simpson rule (SPA) to the Fourier inversion integral in equation (19). The generation of a reduced number of density estimates does not result in a serious limitation in the context of the FFT-approach. The authors examine the accuracy and find that the relative error versus the Nolan densities decreases about two magnitudes in comparison to the LPA/MPA cases. Further accuracy is obtained by searching an exact weight between MPA and LPA with

respect to Nolan's values for a grid of parameter values of α and β. The approximation is performed by a polynomial regression of order 2. This approach is called adaptive SPA and the relative error is about four magnitudes smaller than the MPA/LPA error. We skip this last point and implement the straight SPA in C++ using the FFT approach from 'Numerical Recipes in C'.[89] The discrete Fourier transform appears to be an $O(N^2)$ process in comparison to the FFT-algorithm performing in $O(Nlog_2N)$. This is due to the splitting of a discrete Fourier transform of length N to the sum of two discrete Fourier transforms of length $N/2$. When the original number of points is a power of two, this step can be repeated recursively to finally arrive transforms of length 1. The successive subdivisions of the data into even and odd are tests of least significant bits and the idea of bit reversal can be exploited. This makes the bookkeeping of the recursive application extraordinary simple: Adjacent pairs of one-point transforms are combined to a two-point transform and so on, until the first and second halves of the whole data set are combined into the final transformation. Performance is further improved by replacing complex numbers as their respective real and imaginary real-valued parts which is a favorable storage arrangement for computations. Evaluations departing from the grid nodes is taken care of by the interpolation function. For large absolute values, the linear methods seem appropriate as the curvature of f almost vanishes, but it will play a significant role around the mode of the distribution due to the bell-shape. For this reason, the authors propose a cubic spline interpolation. Nevertheless, the accuracy still becomes unacceptable for large absolute values on the abscissa. To overcome this deficiency the already mentioned concept of tail approximation by an asymptotic Bergström series expansion is applied. An adapted form of this representation is chosen for the alpha-stable parameterization used in the FFT-approach:

Lemma. *(Bergström Series Expansion) For $1 < \alpha \leq 2$, the density $f_{\alpha,\beta}(x)$ of a standardized stable distribution has the following asymptotic expansion for $x \to \infty$:*

$$f(x) = \frac{1}{\pi} \sum_{k=1}^{\infty} ca_k (cx)^{-k\alpha-1},$$

with

$$\begin{aligned} \beta^* &= \frac{2}{\pi(\alpha-2)} \arctan(\beta \cdot \tan(\pi\alpha/2)) \\ c &= \cos(\pi\beta^*(\alpha-2)/2)^{1/\alpha} \\ a_k &= (-1)^{k-1} \frac{\Gamma(1+k\alpha)}{k!} \sin\left(\frac{k\pi}{2}(\alpha+\beta^*(\alpha-2))\right). \end{aligned}$$

The Bergström expansion for all x-values of the FFT-grid are calculated and the tail area is defined by splicing points whose position is dependent on the parameter grid

[89] See Press, Teukolsky, Vetterling and Flannery (1992), p. 504.

of α and β for a standardized stable distribution. The transition between center and tail of the distribution is set equal to that position of grid points where the relative error of the FFT is larger than the relative error of series expansion with respect to Nolan's values. The resulting error is continuous concerning the x-values and the parameters as the relative errors of the FFT-approximation and the Bergstöm expansion vary smoothly provided that the expansion order n and the grid order N are kept constant. The resulting grid evaluations are continuously expanded by means of the cubic spline interpolation. In order to force the obtained probability density to integrate to unity, the whole density approximation is divided by the value of its integral.
We restrict ourselves to symmetric distributions so the respective splice points are defined for a grid of α-parameter values ranging from $\alpha = 1.5$ to 1.9 with 0.1 increments. Recapitulating, we implement the straight SPA-FFT approach with Bergström series expansion for symmetric α-stable distributions with $N = 2^{15}/2$, $n = 15$, and an abscissa-range of $[-256, 256]$ with interval spacing $h = 0.03125$. Finally, computations are forced to integrate to unity. This complete operation for one specific α (β is zero due to the lack of asymmetry implementations) takes about 1 second in C++ on a 1.5GHz processor and 512 MB of RAM.

6.1.3 Simulation of Stable Random Variables

The complexity of the problem of simulating sequences of α-stable random variables stems from the fact that there are no analytic expressions for the inverse $F^{-1}(x)$ nor the cumulative distribution function $F(x)$. A more elegant and efficient solution for standardized skewed α-stable distributions was proposed by (Chambers, J.M., Mallows, C.L. and Stuck, B.W. (1976). A Method for Simulating Stable Random Variables, Journal of the American Statistical Association 71: 340-344.). 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).
For $\alpha \in (0, 2]$ and $\beta \in [-1, 1]$, the method for computer generation of a random variable $X \sim S_\alpha(1, \beta, 0)$ is the following:[90]

- generate a random variable V uniformly distributed on $(-\frac{\pi}{2}, \frac{\pi}{2})$ and an independent exponential random variable W with mean 1;
- for $\alpha \neq 1$ compute

$$X = S_{\alpha,\beta} \cdot \frac{\sin\left[\alpha(V + B_{\alpha,\beta})\right]}{\left[\cos(V)\right]^{1/\alpha}} \cdot \left[\frac{\cos\left[V - \alpha(V + B_{\alpha,\beta})\right]}{W}\right]^{(1-\alpha)/\alpha},$$

[90] See Weron (1996), p. 167.

where

$$\begin{aligned} B_{\alpha,\beta} &= \frac{\arctan(\beta \tan\frac{\pi\alpha}{2})}{\alpha} \\ S_{\alpha,\beta} &= \left[1+\beta^2\tan^2\left(\frac{\pi\alpha}{2}\right)\right]^{1/(2\alpha)} \end{aligned}$$

We can easily simulate a stable random variable for all admissible values of the parameters α, β, σ and μ using the following property: if $X \sim S_\alpha(1,\beta,0)$ then $Y = \sigma X + \mu$ for $\alpha \neq 1$ is $S_\alpha(\sigma,\beta,\mu)$-distributed. Although many other approaches have been proposed in the literature, this method is regarded as the fastest and the most accurate.

6.2 Smoothly Truncated Stable Distributions

The class of STS distributions goes back to the work of Menn and Rachev (2004a). It is compatible with the Black/Scholes/Merton framework due to finite moments of arbitrary order. As a consequence, log-returns of a financial asset described by an STS distribution guarantee a finite mean for the asset price. Simultaneously, it is a powerful tool in describing the distribution of financial variables which are typically leptokurtic and possibly exhibit skewness. Despite the fact that STS distributions possess light tails in the mathematical sense, they provide a flexible tool to model extremal events since a reasonable amount of probability is assigned to extreme events. The convenient measure of tail dependence for regularly varying distributions is not applicable due to the exponential decay of the normally distributed tails of STS distributions.

The authors quantify the occurrence of extreme events by means of QQ-plots and tail probabilities in comparison to other likewise standardized well-know distributions. It turns out that the standardized STS distribution assigns more probability to the tails than the Student-t and far more than the Gaussian.

STS distributions are obtained by smoothly replacing the upper and lower tail of an arbitrary 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, the means of the two normal distributions only differentiate by sign and the variance is equal. However, this does not apply for 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. Due to the amount of probability of extreme events, the speed of convergence to the normal distribution

is extremely slow. This is quantified by the authors by means of tail probabilities in comparison to the normal distribution.[91] Since it could be argued that the normal distribution is not a challenging benchmark for comparing tail probabilities, the authors provide an illustration by comparing the tail probabilities of standardized STS distributions with other popular heavy-tailed distributions like standardized generalized extreme value distributions, the already mentioned standardized Student-t distribution and the standardized skewed Student-t distribution. The quantile-quantile plot in the examined range being representative for financial data assigns the highest tail probabilities to the STS distribution. It can be stated that the family of STS distributions provides an impressive modeling flexibility and turns out to be a viable alternative against many popular heavy-tailed distributions.

In the following we will define the density of an STS distribution according to Menn and Rachev (2005b).

Density of STS distributions. *Let g_θ denote the density of an α-stable distribution with parameter vector $\theta = (\alpha, \beta, \sigma, \mu)$ and h_i, $i = 1, 2$ denote the densities of two normal distributions with mean μ_i and standard deviation σ_i. Furthermore, let $a, b \in \mathbb{R}$ be two real numbers with $a \leq m \leq b$, where m denotes the mode of g_θ. The density of a STS distribution is defined by:*

$$f(x) = \begin{cases} h_1(x) \text{ for } x < a \\ g_\theta \text{ for } a \leq x \leq b \\ h_2(x) \text{ for } x > b. \end{cases} \tag{20}$$

In order to guarantee a well-defined continuous probability density, the following conditions are imposed:

$$h_1(a) = g_\theta(a) \text{ and } h_2(b) = g_\theta(b) \tag{21}$$

and

$$p_1 := \int_{-\infty}^{a} h_1(x)dx = \int_{-\infty}^{a} g_\theta(x)dx \text{ and } \int_{b}^{\infty} h_2(x)dx = \int_{b}^{\infty} g_\theta(x)dx =: p_2. \tag{22}$$

The density in equation (20) is continuous and bell-shaped with a smooth distribution function and therefore the chosen name is justified. STS distributions form a six parameter distribution family $S_\alpha^{[a,b]}(\sigma, \beta, \mu)$. The parameters (μ_i, σ_i) of the two involved normal distributions are uniquely defined by the two equations (21) and (22). They can be obtained from the following two equations:[92]

$$\begin{aligned} \sigma_1 &= \frac{\varphi(\Phi^{-1}(p_1))}{g_\theta(a)} \quad \text{and} \quad \mu_1 = a - \sigma_1 \Phi^{-1}(p_1) \\ \sigma_2 &= \frac{\varphi(\Phi^{-1}(p_2))}{g_\theta(b)} \quad \text{and} \quad \mu_2 = b + \sigma_2 \Phi^{-1}(p_2). \end{aligned}$$

[91] See Menn and Rachev (2005b), p. 54.

[92] The derivation of these formulas is provided in Menn and Rachev (2005b), Appendix A, p. 37.

It can be realized that σ_1 equals the ratio of two special density values. Also, due to the relationship $p_1 = G_\theta(a)$, the quantity a equals the p_1-quantile of the stable cumulative distribution G_θ. Analogous arguments work for σ_2 with the relation $p_2 = 1 - G_\theta(b)$.
It can be shown that a STS distribution is scale and translation invariant since this property is quasi inherited from α-stable and normal distributions[93]: For $c, d \in \mathbb{R}$ and $X \sim S_\alpha^{[a,b]}(\sigma, \beta, \mu)$ it is valid that the random variable $Y = cX + d$ as an affine transformation of the random variable X is again STS distributed, i.e. $Y \sim S_{\tilde{\alpha}}^{[\tilde{a},\tilde{b}]}(\tilde{\sigma}, \tilde{\beta}, \tilde{\mu})$. In the case of $\alpha \neq 1$ the impact of the affine transformation on the distribution parameters is summarized in the following equations:

$$\tilde{a} = ca, \tilde{b} = cb + d, \tilde{\alpha} = \alpha, \tilde{\sigma} = |c|\sigma, \tilde{\beta} = \text{sign}(c)\beta, \tilde{\mu} = c\mu + d.$$

For $\alpha \neq 1$, the parameters a, b and μ can be identified as location parameters, whereas σ serves as a scale parameter.
The moments of this composed distribution class are represented by the following formulas:

$$\mathbb{E}X = \int_{-\infty}^{a} x h_1(x)dx + \int_{a}^{b} x g_\theta(x)dx + \int_{b}^{\infty} x h_2(x)dx \tag{23}$$

$$\mathbb{E}X^2 = \int_{-\infty}^{a} x^2 h_1(x)dx + \int_{a}^{b} x^2 g_\theta(x)dx + \int_{b}^{\infty} x^2 h_2(x)dx, \tag{24}$$

denoting X as the density of a STS distribution as previously and h_1 (h_2) as the density for $x < a$ ($x > b$).

In the HPW framework presented above, the only imposed conditions on the factors are a continuous probability distribution function with support on the whole line having zero mean, unit variance and a finite moment generating function. A properly standardized STS distribution fulfills theses requirements. According to the originators of the class of STS distributions, extensive numerical experiments have shown that the subclass of standardized STS distributions is uniquely defined by the vector of stable parameters $\theta = (\alpha, \beta, \sigma, \mu)$ due to moment matching conditions.[94] Assuming that the claimed relation holds true, the truncation levels given by the four stable parameters can always be effectively calculated by means of equations (23) and (24) such that the resulting distribution is standardized.[95]
Figure 5 shows the influence of the four distribution parameters on the truncation levels a and b for the standardization process. The authors deliver an interpretation

[93] See Menn and Rachev (2005b), p. 9, 10 and 38.
[94] See Menn and Rachev (2004a), p. 7.
[95] See Menn and Rachev (2005b), p. 38-40 for a detailed elaboration of the procedure.

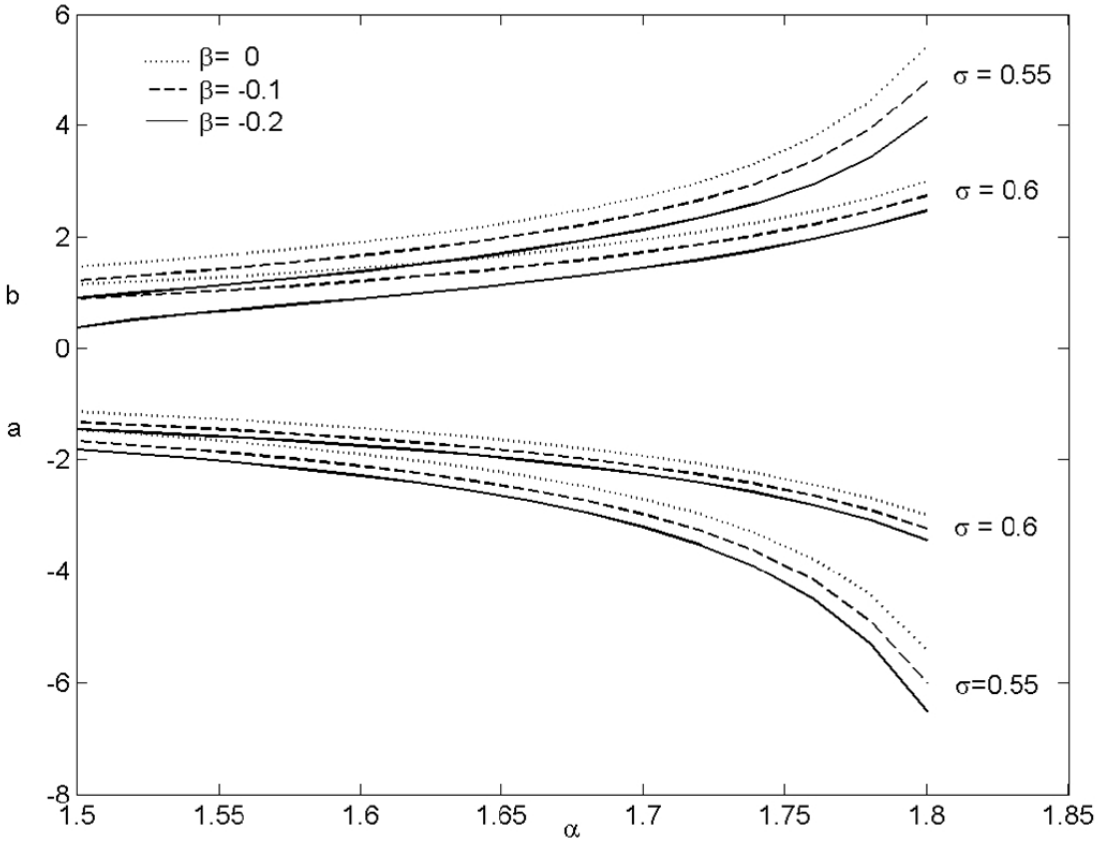

Figure 5: Truncation levels a and b for fixed stable parameter sets for standardized STS distributions. Source: Menn and Rachev (2005b), p. 53.

for the relations between the relevant stable parameters and the probability for extreme events: The latter increases monotonically with increasing α, decreasing β and decreasing σ.[96] 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 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 lower range around the mode to guarantee a variance of one.[97] As there exists no closed-form expressions for the density g_θ, the mean and variance given by the two moment generating functions of an arbitrary STS distribution can only be evaluated by the help of numerical integration. The authors claim that due to modern computer power, the operations to be carried out do not imply any limitation to their approach. This can be stated by our C++ implementation to facilitate arbitrary pdf or cdf evaluations of a standardized STS distribution on the real interval $[-256, 256]$: Computations for this procedure take about 2 seconds. It has to be remarked here that our implementation applies for symmetric standardized STS distribution in the intervals $\alpha \in [1.50, 1.87]$ and $\sigma \in [0.54, 0.90]$ for the pre-specified abscissa range from $[-256, 256]$. The cut-off levels thus simply differentiate by sign due to symmetry. We assume that in case of asymmetric standardized STS distributions, a standard algorithm for the non-linear optimization problem to find the two cut-off levels is even faster than our slow-but-simple two-step search procedure in the symmetric case. Our implementation could easily be enhanced to asymmetric STS distributions as the standardized equations are formulated with respect to both arbitrary α *and* arbitrary β in the familiar ranges.

Regarding the software implementation, the modules from the calibrated FFT-density approximation with Bergstöm series expansion can be recycled to perform necessary interpolation and integration procedures on the basis of cubic splines. Also, the Chambers/Mallows/Stuck simulation algorithm can be utilized as will be outlined now. First of all, for a standardized STS random variable, a region has to be defined from which of the three responsible distributions it has to be sampled. This is accomplished by means of a drawing from a standard uniform variable Z. If $Z < p_1$, random numbers are generated from $N(\mu_1, \sigma_1)$ as long as one sample is in the range smaller than a. The same applies for the case $Z > p_2$ with the respective parameters of the right hand side truncated normal distribution. The same rejection method is carried out by the help of the Chambers/Mallows/Stuck method for the stable parameter set $(\alpha, \beta, \sigma, \mu)$

[96] See Menn and Rachev (2004a), p. 8.
[97] See Menn and Rachev (2005b), p. 11.

if $Z \in [p_1, p_2]$.
All single modules of our approach have been explained so we can proceed with the actual model extension.

6.2.1 STS Distributions in the HPW Model

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

$$\Delta X_i = a_i \Delta M + \sqrt{1 - a_i^2} \Delta Z_i \tag{25}$$

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. As the results of Hull and White (2004) have shown, the double Student-t copula approach with same tail index for both factors results in a good market fit.[98] For this reason we use the same parameters for the distributions of ΔM and ΔZ_i.
Since the normalized sum of the iid STS random variables is attracted by the Gaussian law and STS distributions are not infinitely divisible, the distribution of ΔX_i is not in the class of STS distributions any more. What we can conclude is that it is continuous, defined on the whole real line, has zero mean and variance of δ, has exponential decay far out in the tails, and still exhibits a reasonable amount of tail probability.
Due to its construction, the convolution of the two STS distributions has to be computed numerically. The idea for the implementation is similar to the construction of the default barriers as in section 5.1. 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 (25). Conditional on those probabilities we build up the cumulative distribution for certain grid points with the distribution of the right summand of equation (25). 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 restrict ourselves to a one-factor model.
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.

[98] In literature there exist similar models and it is often assumed that both M and Z_i have distributions with the same tail index.

7 The Valuation of Synthetic CDOs

The purpose of this chapter is to outline the valuation of synthetic CDOs. To create these structures, the arranger of the deal sells a portfolio of single-name CDS in the market and then distributes the credit risk to investors by creating loss tranches.[99] In this context, a standardized index portfolio of CDSs is used as a reference portfolio and the synthetic CDO structure is set up. The protection seller offers compensation for losses induced by credit events in this portfolio of reference entities. On the other side, the protection buyer pays a periodic premium to the seller. The premium is expressed as an annual spread on the outstanding tranche nominal. Premiums are usually paid as quarterly coupons. The purpose of the pricing process is to find those tranche spreads that overreach none of the parties involved: expectations of the payments of protection seller and protection buyer have to be equal, so in this swap contract there exists no up-front fee. This so-called fair spread can be computed with an actuarial approach based on a fixed/premium and a floating/protection leg for different tranches.[100]
A further development in the market involves what is known as 'single tranche trading'. A portfolio and a tranche are defined and 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 are the CDX IG, a portfolio of 125 investment grade companies in the United States, and iTraxx, a portfolio of 125 European investment grade companies.[101] The employed CDO structure is similar to a percentile portfolio credit derivative 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}. CDO tranching allows holders of each tranche to limit their loss exposure to $K_l^U - K_l^L$ percent of the initial portfolio value. The following tables will summarize the different attachment/detachment percentage levels for the two standard indices CDX IG and iTraxx IG.

iTraxx IG tranche no.	attachment in %	detachment in %
1	0	3
2	3	6
3	6	9
4	9	12
5	12	22

Table 6: iTraxx IG Index Tranches

[99] See Hull, Predescu, White (2005), p. 5.
[100] There is an exception concerning the up-front fee of the equity tranche. Nevertheless, this resembles a fair transaction as the tranche spread is reduced accordingly.
[101] See Amato und Gyntelberg (2005).

CDX IG tranche no.	attachment in %	detachment in %
1	0	3
2	3	7
3	7	10
4	10	15
5	15	30

Table 7: CDX IG Index Tranches

In order to employ the technique of risk-neutral valuation in the famous Black-Scholes framework, there have to be made some assumptions about the market. We take a probability measure $\mathbb{Q}$ as given, so the processes of non-dividend-paying assets can be discounted at the risk-free short rate r_t under this equivalent martingale measure. Absence of arbitrage is a necessary requirement for the existence of at least one equivalent martingale measure and market completeness guarantees its uniqueness.[102] Taken the risk-neutral default time probability distribution of the underlying names as a given, we generate future scenarios for the loss behavior of the portfolio. Under the assumption, that the only source of risk stems from the portfolio, the resulting cash flows of the participating credit risk transfer parties can thus simply be discounted at the risk-free rate. This implies that there exists no counterparty risk that could affect the obligations to compensate or pay premiums. 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 conventional homogeneity premises to simplify computations. In the following, we will provide a list of the assumptions:

- 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 will be the same for all obligors.
- The loss given default - or correspondingly the recovery rate - will be deterministic and the same for all companies.
- The initial nominal of each credit in the portfolio will be the same.

[102] See for example Elizalde (2005b), p. 15.

7.1 Intensity Calibration by CDS Market Quotes

Before we consider the pricing of synthetic CDOs, we present a simple method to extract a marginal default intensity λ from market quotes that is applied representatively for all reference entities. It is convenient to us the quoted CDS index spread with the same maturity as the tranches under consideration.[103] From this quote we will extract λ by a similar approach as the actuarial default leg/premium leg method for synthetic CDO pricing.

It has to be remarked here, that we could as well utilize the CDS quotes of all underlyings to arrive at the marginal default distributions in the exponential model. In this case, there was an individual barrier calibration for every underlying due to heterogeneity. Due to simplification in terms of a more homogeneous portfolio we utilize the representative CDS index spread instead. The presented method will be valid for any CDS contract under the established assumptions.

Consider a CDS contract initiated at time 0 with maturity T. Let the coupon payment dates be denoted as $0 = t_0 < t_1 < \ldots < t_J = T$. The insured credit has nominal N and s_{CDS} denotes the annual CDS spread. In order to determine the spread, the premium leg - as a function it - and the protection leg have to be computed. The discounted risk-neutral expectations of the two legs have to be equal in order to receive the fair spread.

In the case of a 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

$$\mathbb{E}PV_{prot}(0) = \mathbb{E}\left[B(0,\tau)\mathbf{1}_{\{\tau\leq T\}}(1-R)N\right], \tag{26}$$

where

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

$PV_{prot}(0)$ 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

[103] Hull, Predescu, White (2005) explain in footnote 1, p. 6, why the index CDS spread is actually slightly lower than the average of the credit default swap spreads.

at default first, since τ is unknown. Equation (26) thus changes to:

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

These integrals represent the fact that payments are made when losses occur in continuous time. For the ease of implementation, we assume that potential defaults can only happen at the coupon 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 coupon date, the protection buyer has to make a premium payment in case no default has occurred until that date. In case of a default event, at that specific default time τ the protection buyer has to pay the fraction of the premium that has accrued since the last coupon date. For simplification, accrued premiums are not considered and δ will be the accrual factor resembling the constant 3-month period between premium dates. The following equation expresses the expectation of the present value of premium payments made:

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

As a final result, the fair spread with deterministic recovery rate and deterministic intensity can be computed in the following way, since default are restricted to happen

only at coupon dates and a constant accrual factor:

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

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.

7.2 The Valuation of Index Tranches

With these inputs we can now proceed to value index tranches of synthetic CDOs. Let t denote the time passed since the CDO was originated, 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 origination at time 0 to t is a percentage $Z_l(t)$ of the portfolio notional value N_{total}:

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

The correctness of this formula can be shown by a distinction of cases:

$$Z_l(t) = \begin{cases} 0 & \text{if } Z_{total}(t) < K_l^L, \\ Z_{total}(t) - K_l^L & \text{if } Z_{total}(t) \geq K_l^L, \\ K_l^U - K_l^L & \text{if } Z_{total}(t) \geq K_l^U. \end{cases}$$

In the following we will outline the actuarial premium leg/ protection leg approach for synthetic CDO tranches. We consider a transaction initiated at time 0 with maturity T. Again, let the coupon payment dates be denoted as $0 = t_0 < t_1 < \ldots < t_J = T$. The predetermined frequency of the coupon payment dates is usually on a quarterly basis, so $\delta = 0.25$ in years. Without loss of generalization, we restrict ourselves to

discrete default times in synchronization with the described version of the HPW model: Defaults are considered at the coupon payment dates whereas they are assumed to actually occur at the midpoints between coupon payment days.
In the Monte-Carlo simulation, for each generation of a future scenario, the respective losses of any tranche at all specified coupon dates are stored. After all simulation procedures have been carried out, these values are averaged to arrive at the expected percentage tranches losses

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

In a portfolio credit default swap contract, tranche investors adopt the position of protection sellers. In our calculations we assume that the only source of risk stems from the credit portfolio and coupon payments and loss compensations are always paid at payment dates without any counterparty risk. 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 discount is chosen in correspondence to our assumptions of default timing.
The tranche investors have to be compensated for bearing the risk of such compensatory payments. The holder of tranche l receives a periodic coupon payment with frequency δ years, amounting to $s_l\delta$ times the outstanding tranche nominal $N_l^{out}(t)$. The spread does not vary during the life of the CDO and it is usually quoted in basis points per annum. However, the initial tranche nominal is stochastically decaying in time induced by tranche losses. At time t_j the outstanding tranche nominal is

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

At coupon payment dates t_j $(j = 1, \ldots, J)$ the expected average outstanding tranche nominals since the last coupon dates have to be considered. The outstanding between coupon 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 coupon dates. As a result, the expected average outstanding tranche nominal between two coupon 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 coupon payments:

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

Finally, this is the equation for the fair spread s_l of tranche l:

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

There is a 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 $f(K^U_{equity} - K^L_{equity})N_{total}$. Additionally, a spread s_{Equity} of 500 basis points per year is paid on the outstanding tranche principal. It has to be remarked 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 equating the expected present values of premium and protection leg. Only the premium leg has to be changed from above like

$$\mathbb{E}PV^{prem}_{equity}(0) = f(K^U_{equity} - K^L_{equity})N_{total} \; + \; \sum_{j=1}^{J} \; B(0,t_j)\; \mathbb{E}[N^{out}_{equity}(t_{j-1},t_j)]\; s_{equity}\delta.$$

8 Calibration and Results

For the calibration to the iTraxx IG index we consider the tranche quotes on 11-April-2005. The settlement date of the third series of this index is 20-September-2005 and matures on 20-September-2010. The index CDS spread on that day is 38.81 bps.
There are 125 equally weighted obligors 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 40%. The applicable risk-free rate for tranches of the Europe-based iTraxx IG is the € zero curve.
Conveniently, we calibrate the equity tranche because the pricing is most sensitive to the model parameters. This resembles a non-linear optimization problem without information about the gradient. The input parameters are the factor loading a, and the tupel (α, σ) of the standardized STS factor distributions.
In literature it is often proposed to employ the calculus-based method of Powell relying on multidimensional direction sets.[104] We consider a version of the intuitive genetic/evolutionary algorithm (GA).[105] These are intelligent search methods based on a biological evolution concept like Darwin's survival-of-the-fittest theory as they comprise concepts such as natural selection, sexual selection and mutation. Also, they are applicable for numerous problems and heuristically process information generated at previous stages of the search process.
The GA has the characteristics of maintaining a population of solutions and results in simultaneous determination of multiple minima rather than a single local minimum. We provide additional information about the other tranche quotes in the objective or fitness function to receive an overall fit with the main focus on the equity tranche quote.
This is realized by a weighted sum of the relative errors of the model-generated cost of protection of the tranches with respect to market quotes. The highest weight is assigned to the equity tranche relative error. The results can be seen in table 8. The Gaussian HPW model was calibrated to the equity tranche with a factor loading $a = \sqrt{0.22779}$. The best parameters in the STS extension of the HPW model are $\alpha = 1.88335$, $\sigma = 0.58727$, and $a = \sqrt{0.25647}$. This is equivalent to truncation levels of around -11.09375 and 11.09375 for the standardized STS distributions. The extension of the HPW model provides a remarkable fit to the quotes in comparison to the Gaussian case.
It has to be emphasized that the incorporation of skewed STS distributions will presumably lead to even better results due to the additional statistically meaningful parameter

[104] See for example Press, Teukolsky, Vetterling and Flannery (1992), p. 412.
[105] See for example J. H. Holland, Genetic Algorithms, Scientif. Am. 44 (1975).

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

β. Also, the extension to stochastic correlation and stochastic recovery rates - as discussed in Hull, Predescu, White (2005) - may lead to even better results. A relaxation of the homogeneous portfolio assumption indeed constitutes additional computational burden but it might also lead to a better quote matching since important market information is not simply 'averaged out'. Finally, the constant intensity calibration of the marginal default time distributions to the index CDS quote can only be seen as an approximation and more sophisticated models should be employed instead.
It is up to future research to arrive at a full calibration to all tranche quotes in the STS and other extensions of the HPW model. The result would be an economically and statistically viable model. This can be motivated by four aspects:

- The structural approach establishes a relation between the default process and the financial variables of the underlying companies. Incorporation of as much economic structure as possible should always be the focus of financial modeling.

- The STS distribution reproduces empirical findings of financial asset returns concerning extreme values characteristics and asymmetry. Also, the distribution class is compatible with the convenient pricing approaches according to the finite moment generating function.

- If full calibration was possible both to the marginal default distributions as well as to tranche quotes, there would be a consistent risk assessment of market credit portfolios and certain derivatives like CDO^2. This is due to the explanatory power of the inherent default behavior, represented by the implied parameters.

- Correlated dynamic fluctuations of credit spreads in the HPW model render possible diverse pricing and risk management applications of credit portfolios and certain corresponding derivatives.

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. On the one hand, the Monte-Carlo simulation may be accelerated by reducing the variance of simulation estimates.[106] On the other hand, there are many starting points for improving the implementation in C++. Our approach represents the credit portfolio as 125 instantiations of a single-name credit class in an object-oriented way. This entails a large number of method-calls strongly reducing performance. Also, the rejection sampling method for STS distributions could be improved except for the Chambers/Mallows/Stuck algorithm for the stable center part, as it is already regarded as the fastest and the most accurate method. Improvements of the pricing part can be neglected as computations hardly consume any time.

[106] As Hull, Predescu, White (2005) point out in footnote 7 on page 9, their use of antithetic paths does not reduce standard errors appreciably, but it reduces time spent for sampling.

References

Altman, E.I.; Brady, B.; Resti, A.; Sironi, A. (2002): *The Link Between Default and Recovery Rates: Implications for Credit Risk Models and Procyclicality*; Working paper, Stern School of Business, New York University, April 2002.

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

Andersen, L.; Sidenius, J. (2004): *Extensions to the Gaussian Copula: Random Recovery and Random Factor Loadings*; Working paper, Bank of America, Juni 2004.

Andersen, L.; Sidenius, J.; Basu,S. (2003): *All Your Hedges in One Basket*; Risk, November 2003.

Avellaneda, M; Zhu, J. (2001): *Modeling the Distance-to-Default Process of a Firm*; RISK Dec. 2001.

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

Black, F.; Scholes, M. (1973): *The Pricing of Options and Corporate Liabilities*; Journal of Political Economics, Mai 1976, p. 637-659.

Bluhm, C. (2003): *CDO Modeling: Techniques, Examples and Applications*; Arbeitspapier, HypoVereinsbank, Dezember 2003.

Bluhm, C. et al. (2003): *Credit Risk Modeling*; Chapman and Hall/CRC.

Burghof, H.-P. et al. (2000): *Kreditderivate*; Schäffer-Poeschel, Stuttgart.

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.

Duffie, D.; Gârleanu, N. (2001): *Risk and Valuation of Collateralized Debt Obligations*; Working Paper, Graduate School of Business, Stanford University.

Duhamel, P.; Vetterli, M. (1990): *Fast Fourier Transforms: A Tutorial Review and a State of the Art*; Signal Processing 19, 259-299.

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

Duffie, D.; Singelton, K. (1999b): *Simulating Correlated Defaults*; Working Paper, Graduate School of Business, Stanford University.

Elizalde, A. (2005): *Credit Default Swap Valuation: An Application to Spanish Firms*; CEMFI and UPNA.

Elizalde, A. (2005b): *Credit Risk Models IV:anderstanding and pricing CDOs*; CEMFI and Universidad Publica de Navarro, download: www.cemfi.es/∼elizalde/.

Embrechts,P.; Lindskog, F.; McNeil, A. (2003): *Modelling dependence with copulas and applications to risk management*; Handbook of heavy tailed distributions in finance, edited by Rachev ST, published by Elsevier/North-Holland, Amsterdam.

Esposito, M. (2002): *Basic Insights in Pricing Credit Derivatives*; Proceedings of the 2nd International Conference 'Managing Credit and Market Risk', Verona.

Gibson, M.S. (2004): *Understanding the Risk of Synthetic CDOs*; Finance and Economics Discussion Series 2004-36, Board of Governors of the Federal Reserve System (U.S.).

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.

Gregory, J.; Laurent, J.-R. (2004): *In the Core of Correlation*; University of Lyon & BNP Paribas.

Gupton, G.M.; Finger, C.C.; Bhatia, M. (1997): *CreditMetrics - Technical Document*; New York, J.P. Morgan & Co. Incorporated.

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.

Hull, J.; White, A. (2005): *The Perfect Copula*;

Krahnen, J.P. (2005): *Der Handel von Kreditrisiken: Eine neue Dimension des Kapitalmarktes*; Perspektiven der Wirtschaftspolitik, Vol. 6(4), p. 499-519.

Laurent, J.-R.; Gregory, J. (2003): *Basket Default Swaps, CDOs and Factor Copulas*; Working paper, ISFA Actuarial School, University of Lyon.

Li, D.X. (2001): *On Default Correlation: A Copula Approach*; The Journal of Fixed Income, Vol. 9, p. 43-54.

Lucas, D.; Sam, T. (2001): *CDO Handbook*; Global Structured Finance Research Report, JP Morgan Securities Inc.

Malevergne, Y.; Sornette, D. (2004): *Tail Dependence of Factor Models*; The Journal of Risk 6(3), p. 71-116.

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

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/technical_reports.

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/technical_reports.

Merton, R.C. (1974): *On the Pricing of Corporate Debt: The Risk Structure of Interest Rates*; Journal of Finance, Vol. 2, 449-471.

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

Rachev, S.T. (2003): *Handbook of Heavy-tailed Distributions in Finance*; North Holland. 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.

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.

Rogge, E.; Schönbucher, P.J. (2003): *Modelling Dynamic Portfolio Credit Risk*; working paper, February 2003 (published, e.g., on http://www.defaultrisk.com)

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

Schönbucher, P.J. (2000): *Factor Models for Portfolio Credit Risk*; Working Paper, Department of Statistics, University of Bonn.

Schönbucher, P.J. (2005): *Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives*; published on http://www.defaultrisk.com.

Schmidt, W.; Ward, I. (2002): *Pricing default baskets*; Risk, Januar 2002.

Sidenius, J.; Piterbarg, V.; Andersen, L. (2005): *A new framework for dynamic credit portfolio loss modelling*; published on http://www.defaultrisk.com.

Vasicek, O. (1977): *An equilibrium characterization of the term structure*; Journal of Financial Economics, Vol. 5, p. 177-188.

Weron, R. (1996): *On the Chambers-Mallows-Stuck method for simulating skewed stable random variables*; Statistics and Probability Letters 28, 165-171. *See also:* **R. Weron**, *Correction to: On the Chambers-Mallows-Stuck method for simulating skewed stable random variables*; Research Report HSC/96/1, Wroclaw University of Technology, 1996, http://www.im.pwr.wroc.pl/ hugo/Publications.html.

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