Portfolio Credit Risk Modeling during the Subprime Crisis

Erik Svensson

June 18, 2011

Abstract

Dependence, or correlation, modeling was at the heart of the last decade's booming market for complex, multi-underlying credit derivatives. Portfolio credit risk and pricing models were used for trading, structuring and rating of such derivatives and securities. The subprime credit crisis, which evolved into a financial crisis of global reach, unveiled that these models often failed to adequately account for systemic shocks to the economy. The crisis also caused drastic changes in risk perception among participants in credit derivatives markets and therefore provides an interesting setting for studies of pricing models. In this thesis we evaluate a Duffie-Gârleanu type model of portfolio credit risk by calibrating it to a range of European credit derivatives, both before and during the subprime crisis. We study the implied parameter set resulting from the calibration, and how it evolved during 2007 and 2008.

Stockholm School of Economics
Master's Thesis in Finance

Tutor: Assistant Professor Cristian Huse

Acknowledgements: I would like to thank my tutor Cristian Huse for his helpful guidance and support.

Contents

1 Introduction **3**

2 An Overview of Credit Derivative Securities **5**
2.1 Credit Default Swap, CDS 5
2.2 Credit Indices 5
2.3 Credit Index Tranches/Synthetic CDO 5

3 Portfolio Credit Loss Modeling **7**

4 Data **10**
4.1 Credit Derivatives Data 10
4.2 Interest Rate Data 10

5 The Model **11**
5.1 Risk-Neutral Default Intensities 11
5.2 Portfolio Loss Distribution 12
5.3 Making the Model Parsimonious 12
5.4 Economic Interpretation of the Parameters 13
5.5 Pricing Framework 14

6 Empirical Method **15**
6.1 Fitting Criterion 15
6.2 Fitting Procedure 16

7 Model Fit **17**

8 Results - Calibrated Parameters **23**

9 Conclusions **28**

10 Suggestions for Future Research **29**

References **30**

A Credit Derivatives Pricing **32**
A.1 CDS Pricing 32
A.2 Credit Index Pricing 33
A.3 Credit Tranche Pricing 33

1 Introduction

Volumes of credit derivative securities grew at extreme pace in the years preceeding the 2007-08 credit crisis. The outstanding amount of credit default swaps (CDS), the most common single entity credit derivative, grew from 920 *billion* USD to 62 *trillion* USD between 2001 and 2007—thus, the market became roughly 66 times bigger in just six years. In parallell, the instruments traded and issued became increasingly complex—achronyms such as CDO, CMO and RMBS[1] quickly became standard jargong in the financial markets. What these abbreviations have in common is that they all represent securities that involve pooling and packaging of credit risk (be it mortgage bonds, bank loans, card receivables or credit default swaps) from different loans or issuers, in contrast to a corporate bond or CDS which only involves the credit risk of one entity. The basic recipe was often the same: use as ingredients credits or loans that had seemingly little to do with each other, e.g. mortgages from different parts of a country or credit default swaps written on companies in different industry sectors, pool them together and issue securities or derivatives with different seniority (tranches) against the pool that served as collateral. Because of the assumed diversification in the pool, rating agencies awarded high credit ratings to senior tranches of pools of not so good loans. It is an instructive application of a fundamental aspect of modern portfolio theory—the risk-return trade-of can be enhanced with a diversified portfolio. This process allowed investment banks to manufacture securities with higher yield than regular bonds with a similar rating and perceived credit risk. Not surprisingly, these became very popular among end investors (e.g. insurance companies) that in the low interest rate environment were struggling to find attrative yields in their fixed income portfolios. With the financial markets' vocabulary, the *correlation* in the pool was low, meaning that defaults in the pool are usually not related to each other. Both the rating of credit derivative securities and the following trading, pricing and risk management done by investment banks was centered around correlation assumptions. In particular, the so called Gaussian copula model (Li, 2000) became the industry's standard weapon of choice. However, both scholars and industry experts have dedicated much effort to developing increasingly sophisticated pricing models for these securities.

As mentioned above, the fundamental idea that motivatives issuing securities backed by a collateral pool is to reduce the risk through diversification without decreasing the yield. In a very large portfolio, the idiosyncratic (or firm specific) risk tends to zero. Thus, one can argue that the concept is sound from a theoretical point of view. What remains is systemic risk—macroeconomic shocks that affect all loans simultaneously. The possibilities to reduce risk by diversification ultimately depend on how sensitive the portfolio is to systemic risks, or in other words, how correlated the loans are. In hindsight, the correlation, or dependence, assumptions behind the ratings, prices and yields of portfolio credit derivative securities were on many occasions too naïve, and in particular the possibility and consequences of a macroeconomic shock were underestimated. The focus of this thesis is correlation modelling and in particular modelling of systemic events. Market participants—investment banks, rating agencies and end investors—became increasingly aware of such systemic risks as the crisis unfolded, and this awareness was reflected in market prices of securities. In particular, we use market data from the credit crisis and study a Duffie-Gârleanu type model of a credit portfolio. The model, which was first described by Duffie and Gârleanu (2001), is a bottom-up model which implies that the default dynamics in the portfolio are constructed by combining the default dynamics for the individual loans, subject to a specification of the dependence between them. The dramatic developments in risk perception among market participants during 2007 and 2008 provides an interesting backdrop for evaluating pricing models for portfolio credit derivatives—the model's response to such changing

[1] Collateralized Debt Obligation, Collateralized Mortgage Obligation and Residential Mortgage Backed Security, respectively.

market conditions is interesting in at least two aspects. First, is the model, when subject to certain restrictions that limit the number of free parameters, flexible enough to describe a broad range of market scenarios for multiple reference securities simultaneously? Second, is the model's response to such changing market conditions intuitive and economically interpretable? The aim and contribution of this thesis is to answer these questions with respect to a Duffie-Gârleanu type model.

The outline of this thesis is as follows. Section 2 gives an overview of credit derivatives markets, with an emphasis on the securities that are relevant for this thesis. Section 3 gives a background on research and literature on portfolio credit loss modeling, where we point out certain aspects of the credit derivatives markets where simpler models often fail. The data set is described in Section 4 whereas Section 5 covers the Duffie-Gârleanu model setup and constraints that are imposed in order to limit the number of free parameters. In Section 6 we discuss the calibration of the model. Our results are described in Section 7 (model fit to market data) and Section 8 (estimated parameters), respectively. Finally, Section 9 concludes and summarizes the results and Section 10 contains a few suggestions for future research.

2 An Overview of Credit Derivative Securities

In this section we provide a brief overview of the credit derivatives that are relevant in this thesis. The reader familiar with such contracts may skip this section. In Appendix A we provide valuation formulas for these contracts.

2.1 Credit Default Swap, CDS

The credit default swap, or CDS, is the most common type of single-name credit derivative. It was originally brought to the financial markets as a tool to help corporate bond investors *hedge* their credit risk. However, investors can also enter into CDS contracts without actually owning underlying bonds—the CDS is thus also a means of *speculation* on the future perceived credit risk of the reference entity. More recently, CDSs have also been used as building blocks in more complex credit derivative securities, most notably the synthetic CDO (these are described below).

Typically, the *protection buyer* receives a payment from the *protection seller* in case of default of the *reference entity* before the maturity of the contract. This is referred to as the *default leg* of the CDS. In most cases, the protection buyer is not compensated for the recovery value on the reference bond. In turn, the seller of protection receives periodic coupon payments from the buyer of protection until a default event occurs. This cash flow stream is referred to as the *premium leg* of the CDS.

2.2 Credit Indices

Credit indices are designed to allow buying and selling of protection on a basket of references entities. Thus, these indices allow investors to express their views on the general level of credit worthiness in the basket. By going long and short indices with different underlying characteristics investors can also pursue relative value strategies. Intuitively, one can think of the index as a portfolio of CDS contracts. The protection buyer periodically (often quarterly or semi-annually) receives payments corresponding to the losses in the portfolio in that period. The protection seller receives fixed rate payments on the remaning (non-lost) notional amount.

The most popular credit indices are CDX.NA.IG and iTraxx Europe which are equally weighted indices of 125 investment-grade rated companies from North America and Europe, respectively. Their compositions are updated every six months when new *series* are released. Different maturities, from three to ten years, are also available.

2.3 Credit Index Tranches/Synthetic CDO

Collateralized Debt Obligations (CDO) are securities backed by a specified pool of collateral assets. A CDO *tranche* is constructed by 'tranching out' a specific piece of the credit risk of the underlying collateral portfolio. Tranches are assigned different credit ratings and carry different yield depending on their seniority with respect to the collateral pool, and investors can choose their preferred risk-reward relation. The notional amount backing a specific tranche begins to decrease once the portfolio loss reaches the lower attachment point of the tranche and is written down successively until the loss hits the upper attachment point. It resembles a call spread position on the portfolio loss. The least senior tranche is usually referred to as the *equity* tranche—investors in this tranche absorb the first losses in the pool and are therefore compensated by a yield that is high relative to the pool's average.

In a synthetic CDO, the underlying collateral are not physical assets (e.g. autoloans, credit card receivables or mortgages). Rather, it is typically backed by a portfolio of credit default swaps. One can intuitively think of these synthetic CDO tranches as a

portfolio of CDS contracts where losses are allocated to tranches according to their seniority. Protection sellers in more senior tranches will receive a smaller spread/coupon compared to the average spread in the portfolio. On the other hand, the probability that the seller will have to compensate the protection buyer is also smaller as their are other more junior tranches that absorb losses. The protection buyer receives periodic payments corresponding to the losses allocated to the tranche during the period. Thus, the protection legs begins to pay off when losses exceed the lower attachment point and provides a stream of cash flows until losses reach the upper attachment point. The protection seller receives periodic fixed payments on the *remaining* notional amount in the tranche. Thus, fixed payments occur until losses exceed the upper attachment point.

Since a credit index (such as CDX.NA.IG or iTraxx Europe) is basically a portfolio of CDS contracts, and synthetic CDO tranches can be backed by such a portfolio, it is also natural to consider tranched versions of the credit indices. These tranches allow market participants to express their views on the future correlation between the individual names in the index as the tranche's yield sensitivity with respect to correlation assumptions depends on its seniority. Therefore, given a specified model and market prices of different tranches it is also possible to deduce the market's implied view on the model's parameters (most importantly the 'correlation'). Relative value trades, going long and short different tranches, are a particular efficient way to construct correlation trades as it to some extent neutralizes the risk of changing general credit spreads.

3 Portfolio Credit Loss Modeling

The key to portfolio credit loss modeling is the mechanism that induces default dependence between individual names in a loan portfolio, a phenomenon that has been documented empirically by e.g. Das, Freed, Geng, and Kapadia (2002). In parallell with the recent growth of multi-name credit derivatives markets, default dependence modeling has received increasing attention from both academic researchers and industry practitioners, and many alternative models have been proposed to this end. The aim of this section is to provide the reader with a brief overview of the exising literature on portfolio credit loss models and some of the current research in this active field of research.

Among industry practitioners the copula approach has become a very popular framework for pricing and hedging multi-name credit derivatives. In these models, a single name defaults when its return falls below a certain threshold that can be inferred from market CDS spreads so that the correct marginal default probabilities are reproduced. Given the marginal default probabilities of the individual names, a copula imposes a dependence structure between default times in the portfolio. The industry standard model is the Gaussian (or normal) copula[2] model (Li, 2000) which has become the Black-Scholes model of portfolio credit risk (O'Kane and Livesey, 2004). In this model, the default risk of each name in the portfolio can be decomposed into two sources; idiosyncratic risk and economy-wide, or systemic, risk. The loadings on the two factors determine how strongly connected to the general state of the economy a particular name is. In order to make the model parsimonious it is often assumed that all names have the same factor loadings. In this setting, one parameter, the correlation, completely specifies the dependence structure between default times in the portfolio. However, the Gaussian copula model does not provide a satisfactory fit to the market prices of CDS index tranches. In general, each tranche spread requires a different correlation, the so called *compound correlation* of the tranche. This phenomenon is similar to the implied volatility smiles observed in option markets. However, compound correlations suffer from e.g. uniqueness problems, see e.g. O'Kane and Livesey (2004), which motivates the theoretically more appealing concept of *base correlations*, defined as the implied correlations of equity tranches (McGinty, Beinstein, Ahluwalia, and Watts, 2004). Indeed, the shortcomings of the Gaussian copula model are often presented in terms of the base correlation skew, meaning that *equity* tranches with increasing detachment levels must be priced with different correlations. In addition, different base correlations are usually required to fit tranche spreads of different maturities.

Not surprisingly, a lot of research has been directed towards finding alternative models of portfolio credit risk, in particular models that are consistent with base correlation skews. Put more intuitively, the challenge has been to find models that can produce portfolio loss distributions consistent with those implied from market data. Andersen and Sidenius (2005) propose extensions to the Gaussian copula model with random recovery rates, and, more importantly, random factor loadings (RFL). In its simplest version, the RFL model assumes a two-point distribution for the correlation parameter in the Gaussian copula model. In this setting the correlation is dependent on the state of the economy, typically with increasing default dependence and systemic risk during "bad times". van der Voort (2007) takes a different view by introducing additional idiosyncratic risk in the Gaussian copula framework, motivated by the recent corporate scandals of e.g. Enron and Parmalat. Other authors have proposed completely different

[2]A copula is a concept from the field of probability theory. It defines the *coupling* structure of marginal probability density functions of random varibles into a multi-dimensional probability function. The dependence between the random variables is completely determined by the copula. Arguably, the normal or Gaussian is the most famous copula. Other examples are the t-, Gumbel and Clayton copulas. See McNeil, Frey, and Embrechts (2005) for an introduction to the subject.

families of copulae; Mashal and Naldi (2002) use the t-copula, Hull and White (2004) the double t-copula and Kalemanova, Schmid, and Werner (2007) the Normal Inverse Gaussian (NIG) copula.

The alternative copula models mentioned above do improve the fit to market quotes on CDS index tranches compared to the Gaussian copula model[3]. Nevertheless, from a theoretical point of view they all suffer from the same problem—they are *static* models of credit risk, meaning that default risks do not evolve dynamically, see e.g. Das, Freed, Geng, and Kapadia (2002) for empirical evidence. Besides being unrealistic and that copula parameters often lack economic interpretation, it is also argued that the static property of these models make them inappropriate for certain derivatives products, e.g. options on CDOs and forward starting CDOs, see e.g. Feldhütter (2008) or Mortensen (2006). Thus, there has been and is a need to move beyond the copula models. An alternative modeling framework, with dynamic default risk, are those based on stochastic (default) intensities. These models[4] were introduced for the purpose of single-name default modeling by Lando (1994), Jarrow and Turnbull (1995) and Duffie and Singleton (1999)[5]. They were later extended to a multi-name setting by Duffie and Gârleanu (2001) who provide an analytical expression for the portfolio loss probability under the assumption of homogenous default probabilities. In this framework, default dependence is induced by modeling individual default intensities as the sum of two intensity processes; one idiosyncratic and one common. The homogenous portfolio assumption was relaxed to allow heterogenous names by Mortensen (2006) who argues that homogenity among the individual names severly restricts the amount of possible default correlation in the portfolio. Mortensen (2006) also introduces a computationally efficient (Eckner, 2007b) semi-analytical valuation method, calibrates the model to CDO tranche spreads on a single day and compares the fit to various copula models.

Following this line of work, Eckner (2007a) performs daily recalibrations of a similar model and also studies risk premia in CDO tranche spreads. Feldhütter (2008) estimates the model using a longer time period, studies the relation between estimated single-name models and Mortensen's (2006) specification of the portfolio model and compares its ability to hedge CDO equity tranches to the Gaussian copula model. Other empirical studies related to CDO pricing are Longstaff and Rajan (2008) and Bhansali, Gingrich, and Longstaff (2008). They model the portfolio loss directly, without involving the single names[6], as a combination of independent Poisson counters representing idiosyncratic, sectorwide and economywide risk, respectively. Through calibration they are able to determine the contribution of each risk source to credit index spreads. Bhansali, Gingrich, and Longstaff (2008) cover data until the end of 2007 and their findings suggest that the subprime credit crisis primarily has been driven by the model's economywide systemic factor. Another paper relating to the subprime crisis is Scheicher's (2008) study of the relationship between tranche spreads and various macroeconomic indicators.

The aim of this thesis is to contribute to the empirical literature on multi-name credit derivatives modeling. In particular, we study estimation of a Duffie-Gârleanu

[3]Equivalently, the alternative copula models produce portfolio loss distributions more similar to those implied from market data than the Gaussian copula.

[4]These models are often referred to as *reduced form* models, reflecting that default events are not linked to any observable firm characteristics, e.g. the firm's debt-to-equity ratio. Such models are referred to as *structural* models, e.g. Merton (1974) and Leland and Toft (1996).

[5]See for instance Longstaff, Mithal, and Neis (2005) for an empirical application of this type of model.

[6]Models of the portfolio loss distribution that do not involve the constituent single names are sometimes referred to as *top-down* model. In contrast, models that involve connecting single name default risks are called *bottom-up* models. One important advantage of bottom-up models is that one can infer the market's view on default correlation (given a tractable model) and use this information to price bespoke derivatives.

type model using an extensive data set that covers the recent subprime crisis. Thus, the thesis extends the research of Eckner (2007b) and Feldhütter (2008) since we consider data from periods characterised by very different market environments. In addition, the data set used in this thesis covers European credit derivatives data, whereas the two mentioned papers used corresponding American data. It can also be seen as a complement to the recent working paper by Bhansali, Gingrich, and Longstaff (2008) since we consider a model involving the single names, rather than a top-down approach.

4 Data

4.1 Credit Derivatives Data

This thesis uses a dataset that consists of daily data on the 5-, 7-, and 10-year iTraxx Europe CDS indices, their standard tranches and constituent single name CDS contracts. The same set of companies applies to all index maturities in a given series. The standard tranches are the 0-3%, 3-6%, 6-9%, 9-12% and 12-22% tranches, respectively. New series are released on a semi-annual basis, in March and September. For each date in the dataset we consider the *on-the-run* series, i.e. the most recently issued version, except that Series 9 is used for a slightly extended period (until 20 October 2008).

Tranche	Lower attachment	Upper attachment
Equity	0%	3%
Junior Mezzanine	3%	6%
Mezzanine	6%	9%
Senior Mezzanine	9%	12%
Senior	12%	22%
Super Senior	22%	100%

Table 1: Tranche structure of the iTraxx Europe CDS index.

Prices for all tranches, except the equity tranche (0-3%), are quoted as annual spreads with quarterly payments. Equity tranche prices are instead quoted as up-front premiums. For example, an up-front premium of 20% means that the protection seller receives 20% of the tranche notional amount up-front *and* 500 basis points per year with quarterly payments.

4.2 Interest Rate Data

In addition, in order to discount cash flows we use EURIBOR and Euro swap rates to construct zero coupon bond yield curves. In particular, we use the 3-, 6-, and 9-months EURIBOR rates and the 1- to 10-year swap rates. The yield curve is obtained through a standard bootstrapping procedure. For intermediate dates we compute discount factors by cubic spline interpolation of zero coupon bond *yields*.

The dataset covers the period 2007-03-21 to 2008-10-20. All data has been collected from the Bloomberg system.

5 The Model

This section outlines a default intensity based model for the *joint* distribution of default times in a portfolio. The conditional probability of default within a small time interval Δt at time t is

$$\mathbb{Q}(\tau \leq t + \Delta t | \tau > t) \approx \lambda_t \Delta t,$$

where λ_t is the default intensity at time t and τ is the time of default. Unconditional default probabilities can thus be calculated through

$$\mathbb{Q}(\tau \leq t) = 1 - E^{\mathbb{Q}}\left(e^{-\int_0^t \lambda_s ds}\right). \tag{1}$$

A computationally tractable model of credit risk requires efficient calculations of (1). Choosing the default intensity process λ_t from the class of *affine processes* allows one to perform these calculations either through explicit analytical expressions or with reasonably efficient numerical methods. More specifically, the model used in this thesis is built upon a simple and computationally tractable affine process commonly referred to as *basic Affine Jump Diffusion (AJD)*. Mathematically, a stochastic process Z_t is a basic AJD under the probability measure $\mathbb{Q}$ if it satisfies the following stochastic differential equation (SDE):

$$dZ_t = \kappa(\theta - Z_t)dt + \sigma\sqrt{Z_t}dB_t^{\mathbb{Q}} + dJ_t^{\mathbb{Q}} \tag{2}$$

where $(B_t^{\mathbb{Q}})_{t\geq 0}$ is a standard Brownian motion and $(J_t^{\mathbb{Q}})_{t\geq 0}$ is a compound Poisson process with constant jump intensity l and exponentially distributed jumps with mean μ. The jump process is independent of the Brownian Motion. We require that $\kappa\theta \geq 0$ and $\mu \geq 0$. A basic AJD has the following important characteristics:

- It is *mean reverting*. The parameter θ is the long run equilibrium level and κ determines the speed of reversion towards the equilibrium level.
- It has a *diffusive* component that is governed by the Brownian motion $(B_t^{\mathbb{Q}})_{t\geq 0}$. The volatility of the diffusive part is determined by σ.
- It can experience *jumps*—these are governed by the compound Poisson process $(J_t^{\mathbb{Q}})_{t\geq 0}$ which is independent of the diffusive term. The jump process has two properties—average jump size and jump frequency.

5.1 Risk-Neutral Default Intensities

Under the risk-neutral measure $\mathbb{Q}$ the default intensity $\lambda_{i,t}$ of obligor i is described by

$$\lambda_{i,t} = X_{i,t} + a_i Y_t, \tag{3}$$

where $X_{i,t}$ represents the idiosyncratic and Y_t the systemic contribution to the risk-neutral default intensity. Hence, in this model default dependence between the firms in the portfolio is induced through the common (systemic) factor Y. The factor loading $a_i > 0$ is the sensitivity of obligor i to the general state of the economy. Duffie and Gârleanu (2001) assumed that $a_i = 1,\ i = 1, \ldots, m$, which means that all obligors are equally sensitive to the common factor. This assumption was relaxed by Mortensen (2006), with the motivation that it restricts the amount of possible default dependence in the portfolio. As in Duffie and Gârleanu's (2001) original paper, and the more recent work by Mortensen (2006), Eckner (2007a; 2007b) and Feldhütter (2008), $X_1, \ldots, X_m$ and Y are independent under $\mathbb{Q}$ and described by basic AJDs, according to

$$dX_{i,t} = \kappa_i(\theta_i - X_{i,t})dt + \sigma_i\sqrt{X_{i,t}}dB_t^{(i)} + dJ_t^{(i)} \tag{4}$$

$$dY_t = \kappa_Y(\theta_Y - Y_t)dt + \sigma_Y\sqrt{Y_t}dB_t^{(Y)} + dJ_t^{(Y)}. \tag{5}$$

In this model, the jump processes $J^{(Y)}$ and $J^{(i)}$ have intensities l_Y and l_i and mean jump sizes μ_Y and μ_i, respectively. Thus, the default intensity λ_i of obligor i can experience jumps both due to firm-specific and systemic events. The dependence on the systemic process is governed by the loading a_i.

5.2 Portfolio Loss Distribution

In order to compute prices for multi-name credit derivatives, e.g. CDOs, one has to compute probabilities from the joint portfolio distribution. Moreover, the computations have to be carried out efficiently. Fortunately, the model outlined above allows computing the required probabilities semi-analytically, due to a method first described by Mortensen (2006). The basic idea is to compute the portfolio loss *conditional* on the common factor. Due to the properties of AJDs it is possible to perform this computation analytically. Then, one can integrate out the common factor using its density function. The latter step requires numerical integration. For implementation details we refer to (Eckner, 2007a).

5.3 Making the Model Parsimonious

First, a restriction is imposed on the systematic factor loadings a_i. Specifically, we require that

$$\frac{1}{m}\sum_{i=1}^{m} a_i = 1. \tag{6}$$

See (Eckner, 2007a) for a motivation of this constraint. Further, let the systematic factor loading a_i of reference entity i be given by the ratio of the i-th entity's 5-year CDS spread to the average 5-year CDS spread across the portfolio,

$$a_i = \frac{c^{\text{cds}}_{i,t,5}}{\text{Avg}(c^{\text{cds}}_{i,t,5})}. \tag{7}$$

First, we note that this defintion of a_i ensures that the condition (6) is fulfilled. Second, it is key to note that a_i is actually a time-varying quantity. Since the model has a large number of free parameters in the general form outlined above it is necessary to impose additional restrictions in order to estimate its parameters and make the estimates economically interpretable. We follow Eckner (2007a) and impose the following restrictions

$$\kappa_i = \kappa_Y \equiv \kappa \tag{8}$$

$$\sigma_i = \sqrt{a_i}\sigma_Y \equiv \sqrt{a_i}\sigma \tag{9}$$

$$\mu_i = a_i\mu_Y \equiv a_i\mu \tag{10}$$

$$\omega_1 = \frac{l_Y}{l_i + l_Y} \tag{11}$$

$$\omega_2 = \frac{a_i\theta_Y}{a_i\theta_Y + \theta_i} \tag{12}$$

Now, consider once again the default intensity $\lambda_{i,t}$ of obligor i (3). Recall that it is a sum of two basic AJDs; one representing idiosyncratic risk, the other representing systemic risk. In general, the sum of two basic AJDs is not itself a basic AJD. However, the parameter restrictions that we impose above actually ensure that this is the case, hence resulting in a model that is tractable from a computational point of view even though it combines systemic and idiosyncratic risk. After imposing the conditions above, we end up with a model that can be described by the parameter vector $\Theta^{\mathbb{Q}} = \{\kappa, \theta_Y, \omega_1, \omega_2, \mu, L, \sigma, Y_0\}$ in addition to the instantaneous idiosyncratic default intensities $X_{i,0}$.

5.4 Economic Interpretation of the Parameters

In this section we provide a brief description of the model parameters, and their economic interpretations.

κ - Mean Reversion Speed

κ is the common mean reversion speed of both the idiosyncratic and systemic processes. When κ is high, default intensities revert to their long-run level fast, whereas it takes more time when κ is low.

θ_Y - Long-term Default Intensity

θ_Y is the long-term default intensity of the systemic process—the mean reversion property will cause the default intensity to revert to this level.

ω_1 - Jump Event Dependence

ω_1 is an important model parameter. As decribed above, the main challenge of portfolio credit risk modelling is to model default dependence realistically. That is indeed one shortcoming of many popular models, such as the Gaussian copula framework, and a reason to consider more sophisticated models, e.g. Duffie-Gârleanu. ω_1 is a measure of dependence between jumps in the model. Specifically, ω_1 is the fraction of jump intensity due to the systematic default intensity process. A high value means that the individual names are more likely to *simultaneously* experience a jump in default probability, reflecting macro-economic shocks. At the extreme, $\omega_1 = 1$ means that *any* jump will affect all names in the portfolio.

ω_2 - Long-term Default Intensity Dependence

The parameter restrictions ensure that the total default intensity (3) of obligor i is another basic AJD. The long-run default intensity of that basic AJD is $\theta_i + a_i\theta_Y$. Thus, it follows from the definition (12) of ω_2 that it is the fraction of total long-run default intensity that comes from the systemic process. Hence, a higher ω_2 means more correlation between default intensities as the systemic factor becomes increasingly important.

μ - Mean Jump Size

The parameter μ is the mean default intensity jump size of the systemic process. Intuitively, large jump size (μ) combined with strong jump default dependence (ω_1 close to 1) should be a key factor driving the more senior tranche spreads higher.

L - Total Jump Intensity

The total jump intensity L is the sum of jump default intensities from the systemic and individual default intensity processes. It it thus a measure of the total probability that the default intensity of a particular name will experience a jump. Note that the jump intensity of the systemic process, l_Y, is easily recovered from $l_Y = \omega_1 L$.

σ - Default Intensity Volatility

σ is a measure of the (diffusive) volatility of the default intensity process. Since it affects the default intensity through a Brownian process, σ should not usually cause large sudden fluctuations. These shocks are instead captured with the jump processes.

Y_0 - Instantaneous Systemic Default Intensity

Y_0 is the instantaneous default intensity of the systemic process, and hence a measure of the short term probability of a systemic "default".

5.5 Pricing Framework

The default intensity based model outlined above allows us to compute default probabilities that are necessary for the valuation of credit derivatives. However, we must also specify a framework for recovery rates in case of defaults - the important variable is often the *expected loss* which intuitively is contingent on (i) expected default rates and (ii) expected recovery rates. We adopt a framework that is considered as fairly standard and widely used, the interest reader is referred to e.g. Mortensen (2006) or Eckner (2007a). The validity of these assumptions have been subject to both academic and industry research, the general conclusion being that these assumptions are fairly innocuous (Eckner2007b).

Assumption 1: Under the risk neutral measure $\mathbb{Q}$,

1. Default intensities and interest rates are independent of each other.
2. Recovery rates are independent of default intensities.
3. Recovery rates are independent of each other.
4. The expected recovery rate is 40 %.

6 Empirical Method

Specifically, the objective of this thesis is to study

(i) whether a Duffie-Gârleanu type model, as outlined in the previous section, can describe European corporate credit derivative markets during normal (i.e. before the credit crisis) and distressed market conditions, respectively

(ii) how a changing market environment is transfered to the model parameters, and

(iii) if the model's response to new market conditions is economically interpretable.

For a given set of dates in the data set we determine the model parameter set that, according to some measure, minimizes the differences between market prices and model prices. Since we are interested in how the model responds to changing market conditions we will re-estimate the model for each date of interest. We can then track how the optimal parameter set and the quality of the model fit evolves through time, and draw conclusions.

6.1 Fitting Criterion

Perhaps the most natural measure of fit, and certainly something that most market participants can relate to, is to compare the sum of squared differences between market and model implied spreads for individual CDSs, the index and tranches simultaneously. This measure should give a sense of the overall ability of the calibrated model to replicate market prices of traded instruments. Here, model implied spreads refer to spreads that can be calculated given a certain set of model parameters. One question that arises is whether to measure differences between market and model prices in an absolute or relative sense. The effect of using relative differences is that the more senior tranches, which by definition have lower spreads, are given the same weight/importance as tranches with higher spreads, since each difference is scaled by dividing with the size of the market spread. A small error on a very senior tranche can thus be magnified compared to when an absolute difference is used. Empirically, through experimenting with the two methods, we found that using relative differences works better. It is also a neutral choice, since the calibration is not biased in favor of any instrument—all observations at a given time are treated as equally important.

Let $c^{\text{tr}}_{j,t,M}(S)$ for $S \in \{\text{I}, \text{M}\}$ denote the time t spread of the j-th tranche with maturity M (5,7 or 10 years) as implied by the model (I) or the market (M). Further, let $c^{\text{cds}}_{i,t,M}(S)$ and $c^{\text{idx}}_{t,M}(S)$ denote the spread at time t of the i-th CDS and index with maturity M, respectively. Pricing errors are then defined as per below, where X denotes the considered set of maturities:

$$\text{RMSE}^{\text{tr}}_t(\Theta^{\mathbb{Q}}) = \sqrt{\frac{1}{J}\frac{1}{|X|}\sum_{j=1}^{J}\sum_{M\in X}\left(\frac{c^{\text{tr}}_{j,t,M}(I) - c^{\text{tr}}_{j,t,M}(M)}{c^{\text{tr}}_{j,t,M}(M)}\right)^2} \tag{13}$$

$$\text{RMSE}^{\text{cds}}_t(\Theta^{\mathbb{Q}}) = \sqrt{\frac{1}{N}\frac{1}{|X|}\sum_{i=1}^{N}\sum_{M\in X}\left(\frac{c^{\text{cds}}_{i,t,M}(I) - c^{\text{cds}}_{i,t,M}(M)}{c^{\text{cds}}_{i,t,M}(M)}\right)^2} \tag{14}$$

$$\text{RMSE}^{\text{idx}}_t(\Theta^{\mathbb{Q}}) = \sqrt{\frac{1}{|X|}\sum_{M\in X}\left(\frac{c^{\text{idx}}_{t,M}(I) - c^{\text{idx}}_{t,M}(M)}{c^{\text{idx}}_{t,M}(M)}\right)^2} \tag{15}$$

The final error metric, or fitting criterion, is the sum of pricing errors for tranches, CDS contracts and indices:

$$C_t(\Theta^{\mathbb{Q}}) = \text{RMSE}^{\text{tr}}_t(\Theta^{\mathbb{Q}}) + \text{RMSE}^{\text{cds}}_t(\Theta^{\mathbb{Q}}) + \text{RMSE}^{\text{idx}}_t(\Theta^{\mathbb{Q}}). \tag{16}$$

6.2 Fitting Procedure

The goal of the model estimation is to fit the basic AJD model outlined in Section 5.3 to market data for CDS contracts, credit indices and credit tranches. This is achieved by minimizing (16). The procedure runs as follows:

1. Given a parameter vector $\Theta^{\mathbb{Q}}$ individual CDS spreads are calibrated by varying the initial intensities $\lambda_{i,0}$, $1 \leq i \leq N$, subject to $\lambda_{i,0} \geq a_i Y_0$. The fitting criterion/error metric is (14).
2. The parameter vector $\Theta^{\mathbb{Q}}$ is varied in order to minimize (16). In each update of $\Theta^{\mathbb{Q}}$ and Y_0, Step 1 is repeated.

Similarily to (Eckner, 2007b), Step 2 in the fitting procedure is implemented by fitting each parameter separately and iteration over the set of parameters.

7 Model Fit

In this section we briefly discuss the model's fitting capabilities. Relevant figures and tables can be found on the following pages, 18 to 22. A summary of the model fit is provided in Table 2 and in Figure 6 on page 22 - interestingly, the model fit has an improving trend during the sample period. Figures 1 to 3 illustrate the market and model prices for the five-, seven- and ten-year tranches, respectively. The x-axis tick marks represent the bi-annual dates when the data rolls to a new series. Beginning with the five-year tranches (Figure 1, page 18) we note that in general the model seems to agree reasonably well with market prices across all five tranches—the increasing market spreads are also reflected in model prices. The model overprices the equity tranche compared to the market during the first part of the sample, corresponding to Series 7. Interestingly, the fit then improves. The model also slightly overprices the mezzanine (6-9%) and senior mezzanine (9-12%) tranches. This bias is more pronounced for Series 9 (March - October 2008). Moving on to the seven-year tranches (Figure 2, page 19) the general picture is similar. The fit to the Series 7 seven-year equity tranche is better than the corresponding five-year fit. For the ten-year tranches (Figure 3, page 20) the picture is different. The model underprices the junior mezzanine (3-6%) and mezzanine (6-9%) tranches compared to market data. This is not observed in Figures 1 or 2. However, the fit to the equity tranche seems better.

The model and market prices for the indices are shown in Figure 4 on page 21. Until the beginning of 2008 there is a clear upward-sloping term structure of index spreads, a feature that the model is not able to reproduce to full extent. The fit then improves as the term structure flattened and spreads increased sharpy during the first quarter of 2008. Then, during the later part of the sample, the model overprices the index spreads compared to the market for all three maturities. What probably causes this effect is shown in Figure 5 which illustrates the *basis*, i.e. the difference between the five-year index spread and the average five-year CDS spread across the index members. The basis became significantly negative during the summer of 2008, implying that the average CDS spread is higher than the index spread. In terms of supply and demand this translates to more demand for and/or less supply of protection-writing on single names compared to the broad index. Probably both effects were present. What does this imply for the calibration of the model? In this model the index is represented by a portfolio of CDS contracts. Therefore, the probability of losses on the index notional is explicitly linked to the losses on the underlying CDS contracts. Hence, per definition, the model index spreads should track the average model CDS spread—it does not take into account the supply and demand differences.

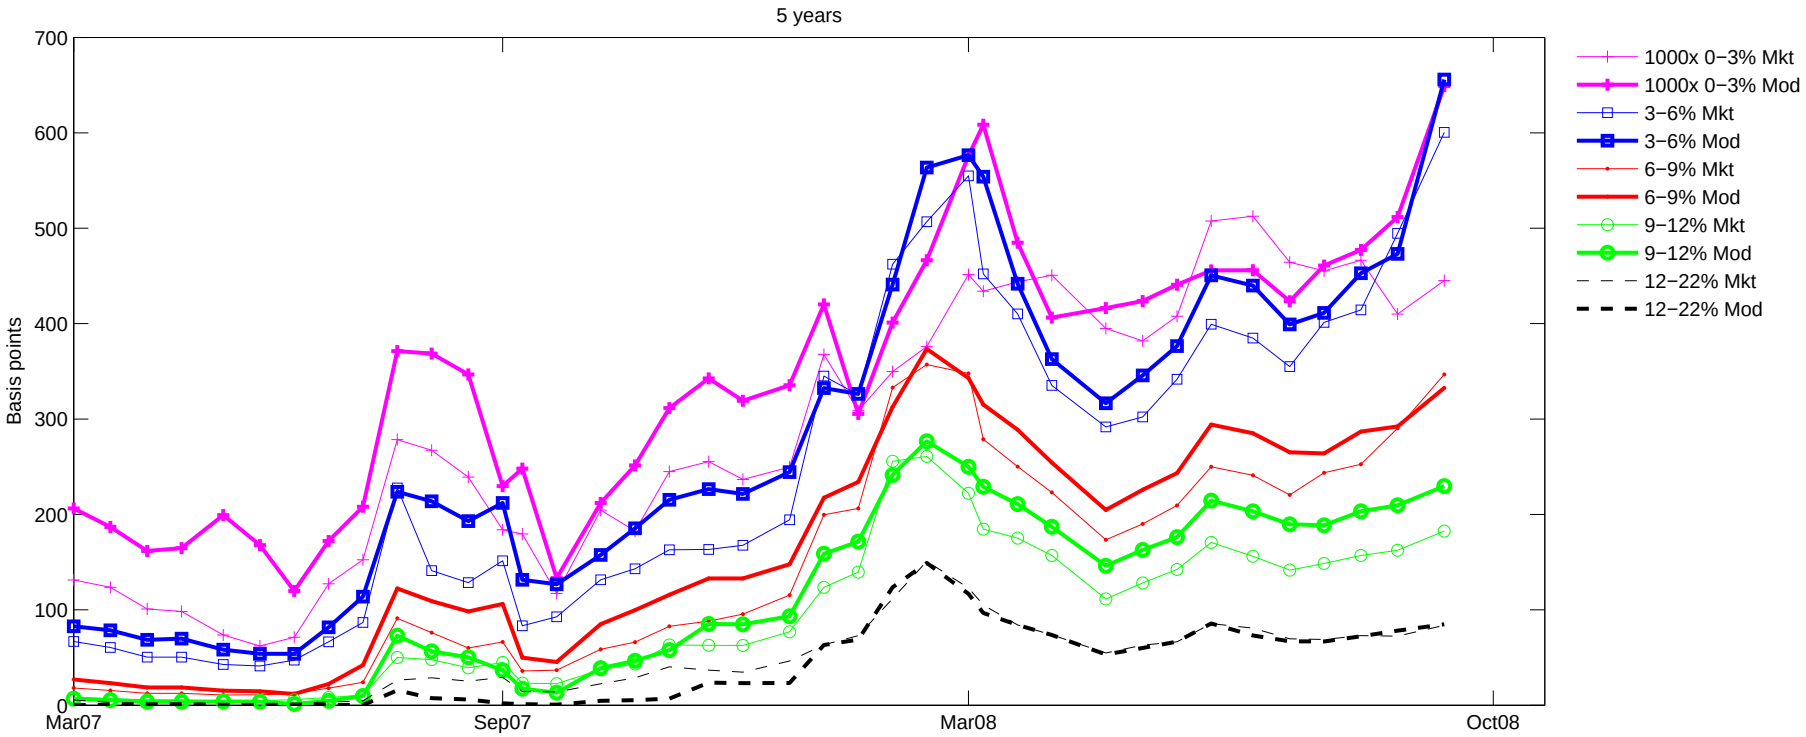


Figure 1: Model (**bold**) and market spreads for the 5-year maturity tranches. The equity tranche upfront payment has been multiplied with 1000 in order to make its size comparable to the spreads of the other tranches. There are approximately ten business days between each sample date.

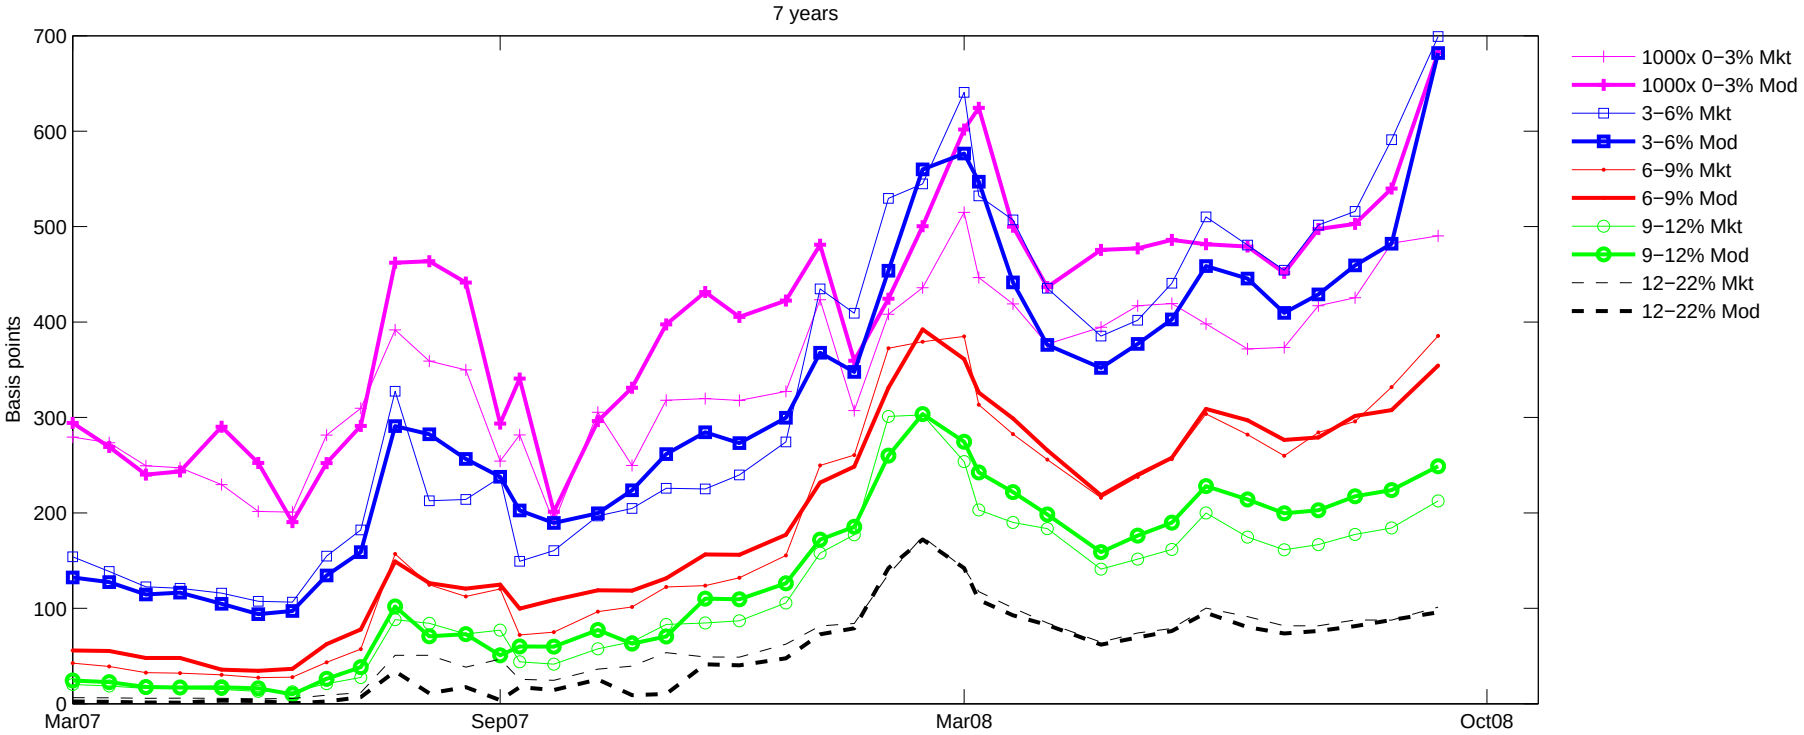


Figure 2: Model (**bold**) and market spreads for the 7-year maturity tranches. The equity tranche upfront payment has been multiplied with 1000 in order to make its size comparable to the spreads of the other tranches. There are approximately ten business days between each sample date.

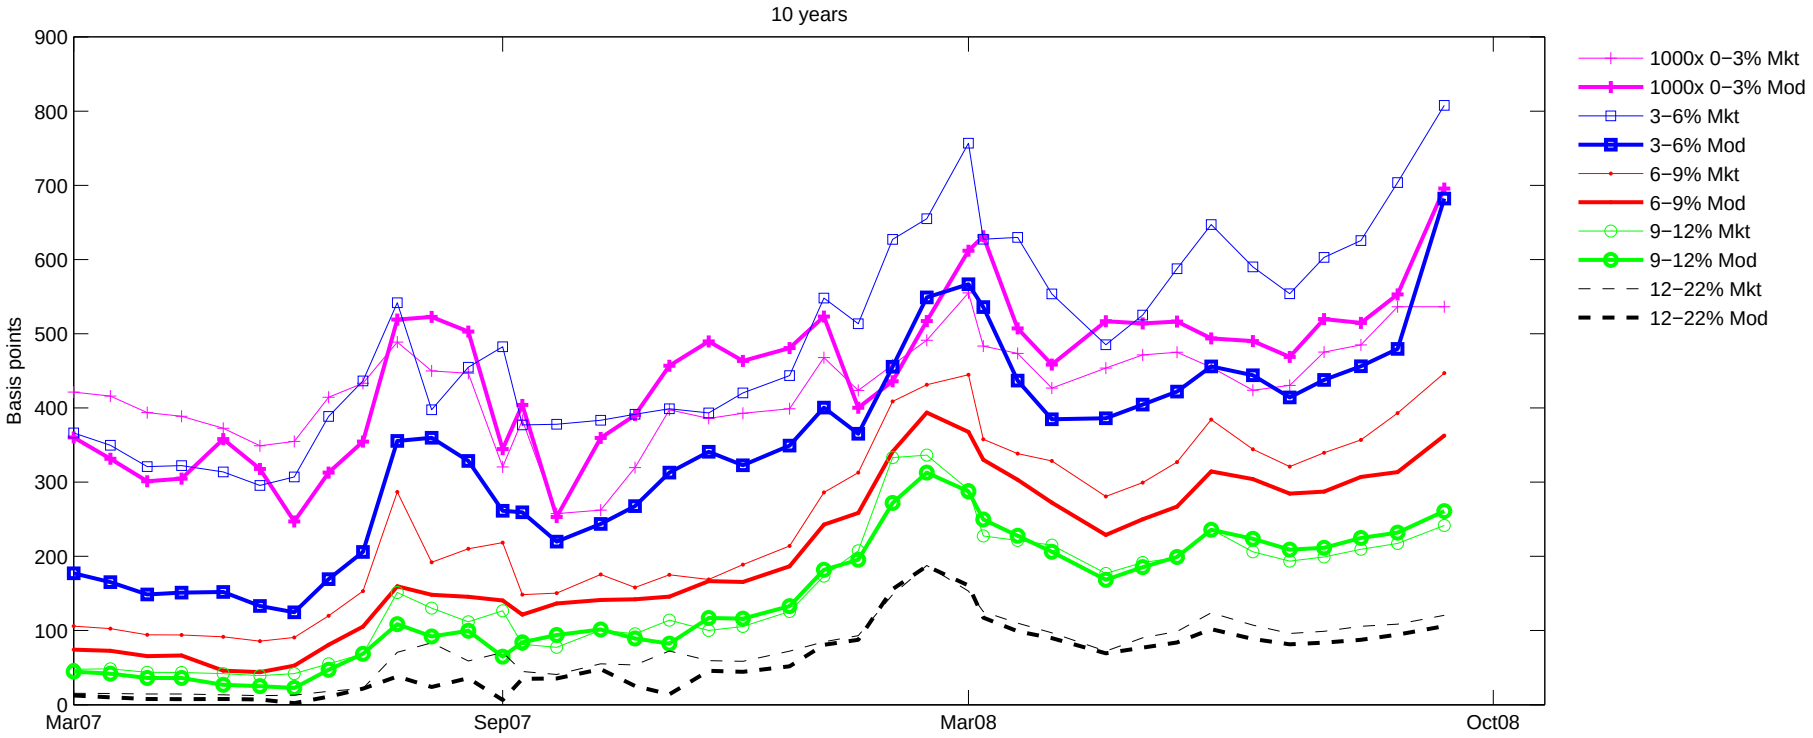


Figure 3: Model (**bold**) and market spreads for the 10-year maturity tranches. The equity tranche upfront payment has been multiplied with 1000 in order to make its size comparable to the spreads of the other tranches. There are approximately ten business days between each sample date.

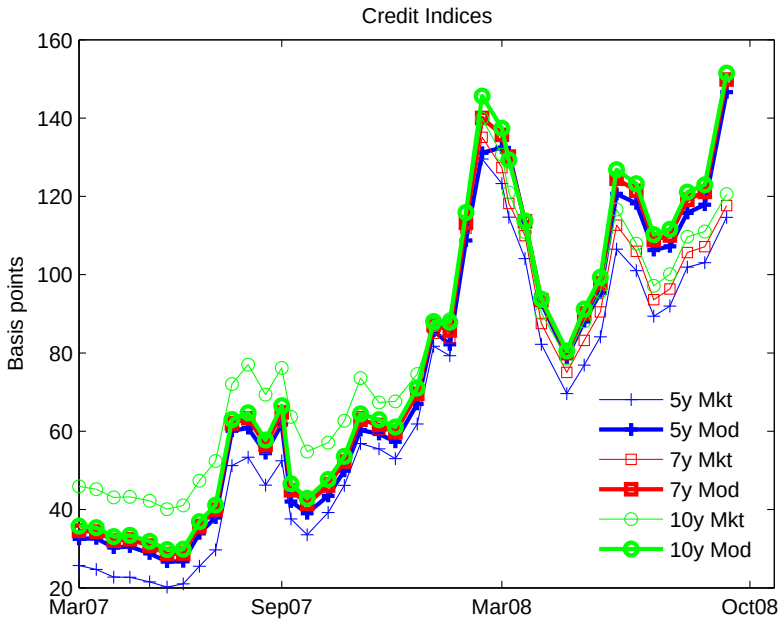


Figure 4: Model (**bold**) and market spreads for the credit indices with maturities 5, 7 and 10 years.

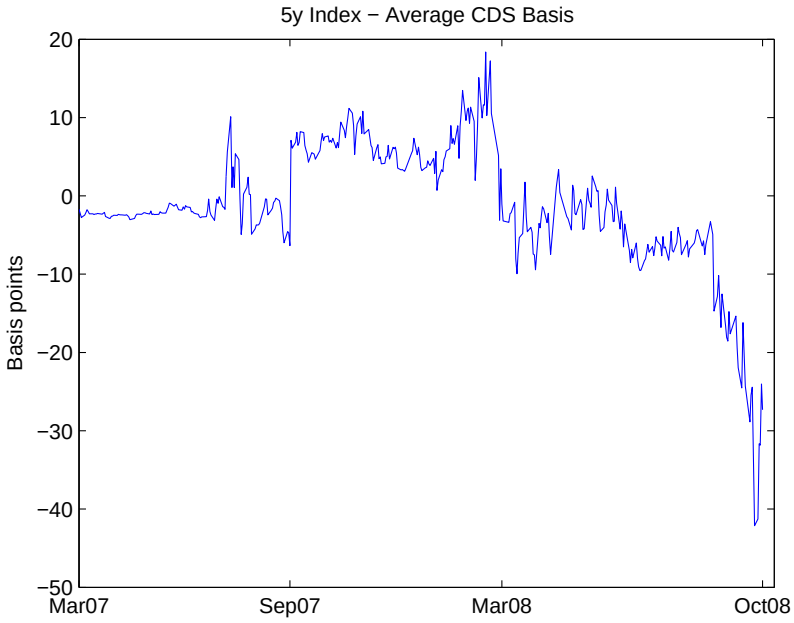


Figure 5: Difference (basis) between the five year index spread and the average CDS spread of the index members, measured in basis points.

Series 7	Tr. 0-3%	Tr. 3-6%	Tr. 6-9%	Tr. 9-12%	Tr. 12-22%	Index	Avg. CDS
5y	0.65	0.31	0.44	0.29	0.86	0.26	0.09
7y	0.13	0.11	0.26	0.19	0.64	0.05	0.14
10y	0.16	0.46	0.35	0.23	0.46	0.26	0.34
Series 8	Tr. 0-3%	Tr. 3-6%	Tr. 6-9%	Tr. 9-12%	Tr. 12-22%	Index	Avg. CDS
5y	0.23	0.22	0.26	0.19	0.44	0.07	0.10
7y	0.18	0.14	0.17	0.18	0.26	0.04	0.15
10y	0.14	0.26	0.13	0.09	0.21	0.14	0.24
Series 9	Tr. 0-3%	Tr. 3-6%	Tr. 6-9%	Tr. 9-12%	Tr. 12-22%	Index	Avg. CDS
5y	0.19	0.13	0.13	0.26	0.03	0.16	0.00
7y	0.24	0.12	0.03	0.17	0.06	0.10	0.05
10y	0.14	0.26	0.14	0.05	0.12	0.07	0.12

Table 2: Summary of model fit, in terms of relative pricing errors, $\frac{\text{abs}(c(I)-c(M))}{c(M)}$, where $c(I)$ and $c(M)$ are the implied model and market spreads, respectively. The 0-3% tranche pricing error refer to the upfront payment.

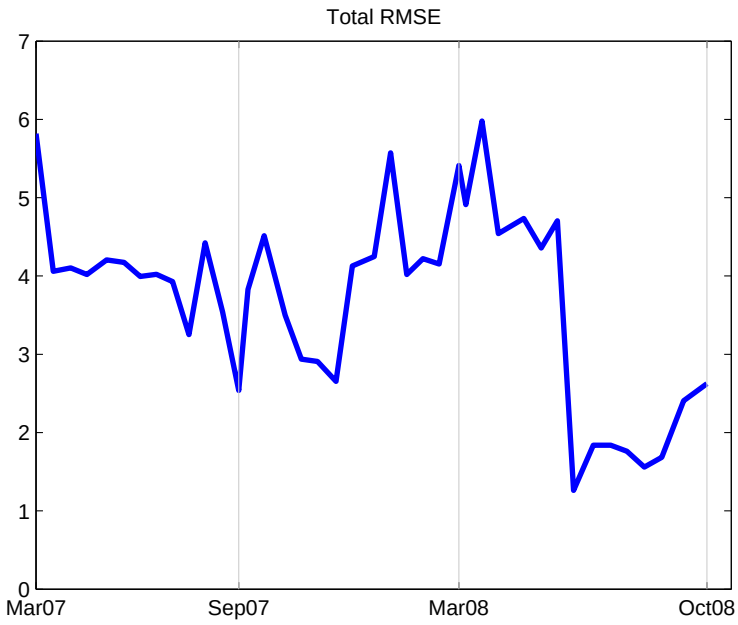


Figure 6: The total RMSE (16) for each model calibration date.

8 Results - Calibrated Parameters

This section describes our main findings, i.e. the optimal parameter sets determined for each model calibration date. We show time series describing how parameters have changed during the sample period, and discuss them in connection with important events that took place as the credit crisis unfolded. We find that some parameters did not respond in any particular way to the dramatically changing market conditions, whereas other parameters have clearly been related to the market's development during the crisis. We invite the reader to review Section 5.4, where the role of each parameter is discussed in the order of appearance. On a more qualitative note we observed that the behaviour of the calibrated parameters are similar when the fitting criterion is modified slightly, e.g. if absolute instead of relative pricing errors are used, or if the CDS constracts are not used as calibration instruments. This suggests that the results are somewhat robust.

κ - Mean Reversion Speed

κ, the mean reversion speed toward the long-run default intensity level, decreased during the credit crisis, see Figure 7 below. One can interpret this as increasing uncertainty about the evolution of future default rates as higher mean reversion implies weaker drift of default intensities to their long-run level. This seems intuitively correct. In particular, low values occur in March (bail out of Bear Sterns) and September (Lehman Brothers' default) 2008.

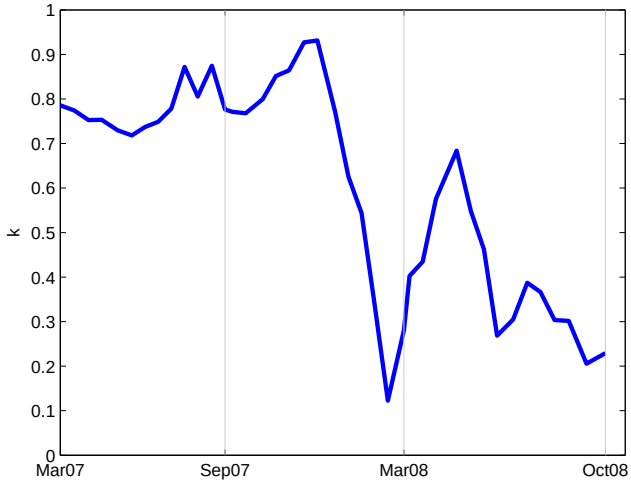


Figure 7: κ measures the strength of reversion to the long run default intensity.

θ_Y - Long-term Default Intensity

θ_Y is the long-term default intensity of the systemic process. Consider again the dynamics of the systemic process (5) and the following equivalent specification:

$$\begin{aligned} dY_t &= \kappa(\theta_Y - Y_t)dt + \sigma\sqrt{Y_t}dB_t^{(Y)} + dJ_t^{(Y)} \\ &= \kappa\theta_Y dt - \kappa Y_t dt + \sigma\sqrt{Y_t}dB_t^{(Y)} + dJ_t^{(Y)} \end{aligned}$$

.

This suggests that what matters as far as the dynamics are concerned are actually the two quantities $\kappa\theta_Y$ and κ. Therefore, one can suspect that θ_Y is a parameter that should be somewhat difficult to infer from the calibration process. Figure 8 illustrates that the behaviour is quite volatile, and that it is difficult to draw any meaningful conclusions in relation to the changing economic environment during the crisis. Therefore, should the model be used in practice, it might be a good idea to lock θ_Y *a priori* or introduce a penalty term in the fitting criterion in order to avoid large sudden swings in the calibrated parameter set.

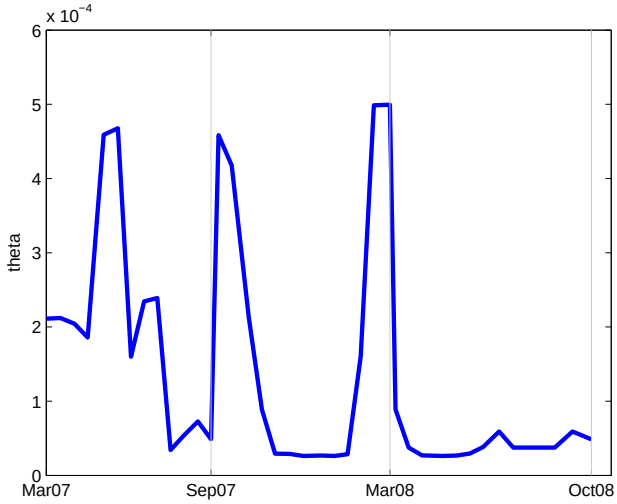


Figure 8: θ_Y is the long-term default intensity of the systemic process.

ω_1 - Jump Event Dependence

As described earlier, in the chosen model setting, the parameter ω_1 is a measure of the amount of default dependence between the constituents in the portfolio. Specifically, it measures the amount of default dependence due to *jumps* in the default intensity process. The higher the value, the higher the degree of default correlation, since a larger fraction of the total jump event intensity comes from the systemic jump process $dJ_t^{(Y)}$. Figure 9 below shows a time series of ω_1. Clearly, it has been steadily increasing during the sample period, corresponding to greater importance of broad macro-economic risk factors. Before the credit crisis about 50 % of total jump risk was systemic, but this fraction increased to 90 %. We can also identify peaks at key events—the "beginning" of the crisis (August 2007), the bail out of Bear Sterns (March 2008) and Lehman Brothers' default (September 2008).

One can also relate the higher degree of default dependence suggested by the model to the increasing spreads on the senior CDO tranches. On a relative basis, spreads on senior tranches increased more than on the junior ones, reflecting a shift from idiosyncractic risk towards adverse development in the macro-economic backdrop. Should a jump event occur in the systemic default intensity process, and ω_1 is close to unity, then many portfolio constituents are likely experiencing this jump, and this should clearly affect the spreads on the senior tranches.

ω_2 - Long-term Default Intensity Dependence

As described earlier, ω_2 is a measure of the amount of long run default intensity that is due to the common process. A high value means that in the long run a larger fraction

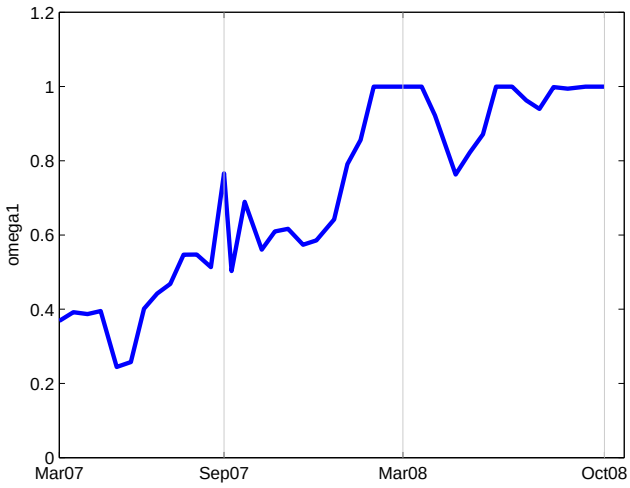


Figure 9: ω_1 is a measure of the default dependence between the portfolio constituents and has been steadily increasing.

of the default intensity is generated from the common risk factor. Therefore, it is also a measure of default dependence. As such, one would expect an increasing time series, similar to what we observed for ω_1 in Figure 9 above. Figure 10 reveals that this is not quite the case, but also that the series is more volatile compared to ω_1. To conclude, ω_2 is a parameter that seems more difficult to infer through the calibration procedure that was used. One possible reason is that it is not as important as other parameters for pricing the securities used in the calibration. Should this model be used in practice, one solution could be to set ω_2 *a priori*, or introduce a penalty term in the fitting criterion that controls the deviations of ω_2 from the latest calibrated value. This conclusion is similar to what we observed above regarding the long-term default intensity θ_Y.

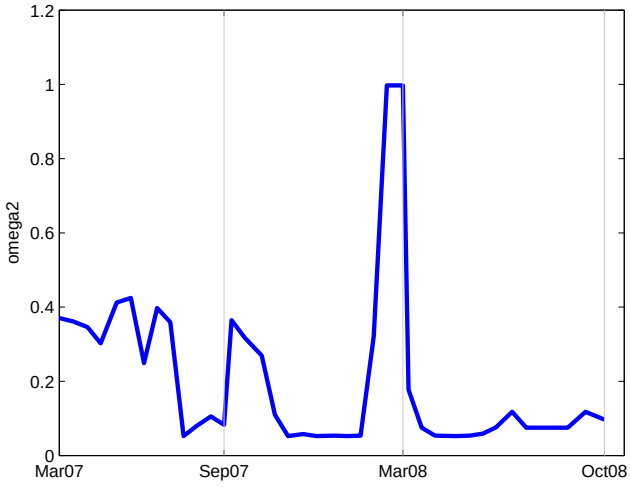


Figure 10: ω_2 is a measure of the default dependence between the portfolio constituents but the behaviour is less clear compared to ω_1.

μ - Mean Jump Size

There are two factors that influence the nature of jump events in the default intensity process—the mean jump size μ and the total jump intensity L. The important quantity is actually μL which one could think of as an expected jump size. In other words, it does not matter whether there are many small jumps (L large - μ small) or a few large jumps (L small - μ large). Intuitively, the uncertainty of future default intensities should have increased during the crisis. Therefore, it is not unreasonable to assume that market participants began to "price in" a greater risk of future credit shocks during the crisis. And, as argued above, μL is a measure of this risk, so we should expect that at least one of the parameters μ and L shows an increasing behaviour. Figure 11 below illustrates that this is indeed the case for the mean jump size μ.

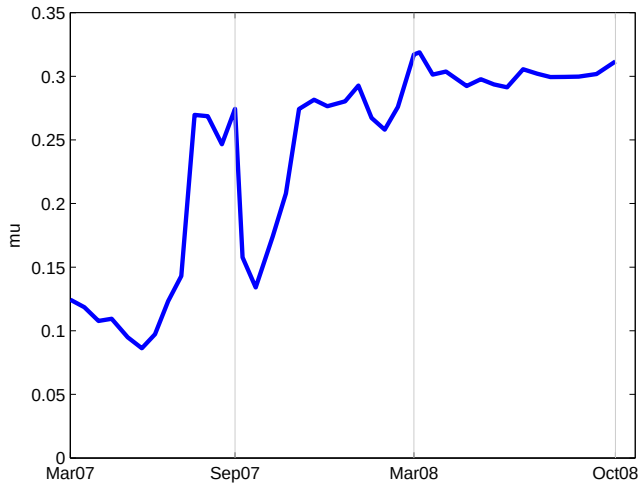


Figure 11: μ is the mean default intensity jump size.

L - Total Jump Intensity

L is the total jump intensity—it represents the probability of an individual name in the portfolio experiencing a jump in default intensity. The calibrated time series of L is shown in Figure 12. It has been rather flat during the crisis—this suggests that the probability of a single name experiencing a jump in it's default intensity has not changed much. This is not a totally convincing result, but as argued above, the important quantity is the product μL rather than μ or L individually. Since μ increased, the jumps are on average larger. We also highlight the changes in systemic jump intensity, l_Y, which is not an actual parameter of the model but is easily recovered through the relation $l_Y = \omega_1 L$. Given what we know about ω_1 and L, it should come as no surprise that it has been increasing. This is illustrated in Figure 13. Thus, the probability of systemic shocks has increased, which seems reasonable from an economic point of view.

σ - Default Intensity Volatility

The credit crisis obviously caused great uncertainty regarding the state of the general economy and the credit quality of many corporates. Accordingly, it is not unreasonable to assume that the volatility of default intensities, as measured by the parameter σ, increased during the crisis. Figure 14 illustrates the calibrated values for σ. The changes are not that pronounced and quite volatile, compared to the re-pricing observed in jump

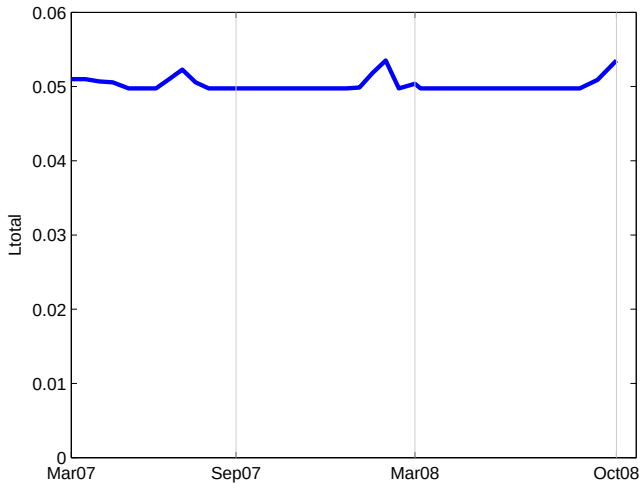


Figure 12: L is the total jump intensity - i.e. a measure of the probability that a single name will experience a default intensity shock.

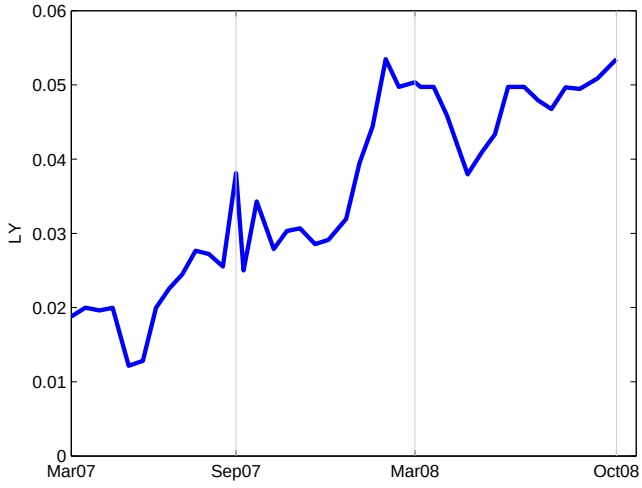


Figure 13: l_Y is the systemic jump intensity.

risks (Figures 9 and 11). This suggests that markets became more concerned about system-wide adverse shocks to the economy rather than generally higher volatility of credit spreads.

Y_0 - Current Systemic Default Intensity

In the model that we consider, Y_0 represents the current, or instantaneous, default intensity of the systematic factor. It is therefore directly related to the short term default probability. Therefore, it is reasonable to expect that Y_0 should have increased during the crisis. Figure 15 below illustrates this; there is an increasing trend with peaks in August 2007 and March 2008.

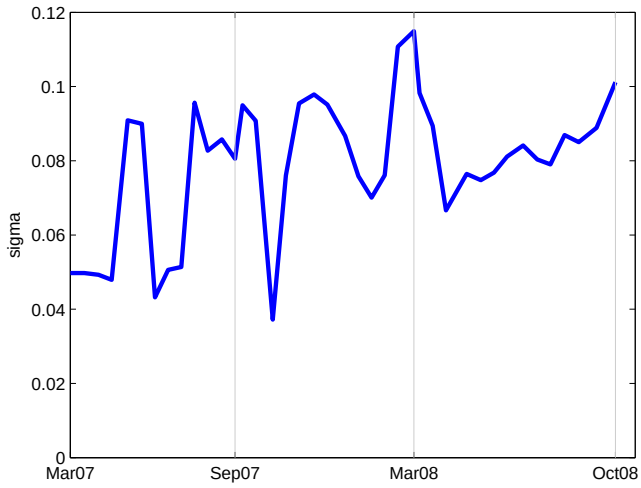


Figure 14: σ is the (diffusive) default intensity volatility.

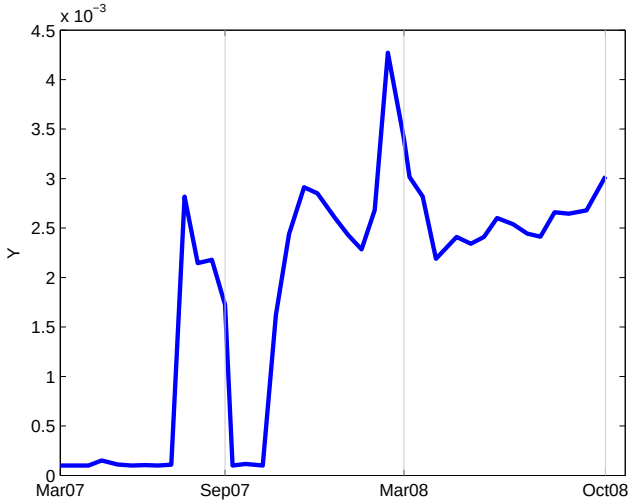


Figure 15: Y_0 is the current (instantaneous) default intensity of the systemic process.

9 Conclusions

Sections 7 and 8 highlighted our main findings; the model's fit to market data and its implied parameters from the calibration. In this section we summarize these results and provide some suggestions on how the model framework could be used in practice. For clarity, we recall from Section 6 that the objectives of the thesis are to assess

(i) whether a Duffie-Gârleanu type model (subject to certain conditions that make it parsimonius) can describe European corporate credit derivative markets during normal (or non-distressed, i.e. before the credit crunch) and distressed market conditions, respectively

(ii) how a changing market environment is transfered to the calibrated model parameters, and

(iii) if the model's response to new market conditions is economically interpretable.

Our results from Section 7 suggest that the Duffie-Gârleanu type model of portfolio credit risk, with seven free parameters (in addition to the idiosyncratic current intensities $X_{i,0}$), is indeed able to capture market prices of credit derivative instruments during both calmer and more volatile market regimes. Moreover, at a given point in time, the model is able to fit market prices both across instruments and maturities. This is a clear advantage of the model compared to e.g. the Gaussian copula which requires different implied parameters (correlation) depending on tranche specification and maturity. On the other hand, it is not possible to obtain a perfect fit to market data since the number of free parameters are limited (and fewer than the number of reference market quotes used at each point in time). What seems clear from Section 8 is also that all seven model parameters have not been equally "reactive" to the changing market conditions during the crisis, given the market quotes that were available for the calibration. In other words—some parameters are more important than others. In particular, our findings suggest that the parameters that are related to jumps in default intensities, i.e. credit risk shocks, are more important. These parameter show clear trends and do not fluctuate wildly during the sample period. A practical solution could therefore be to set some of the other parameters, e.g. the long-term default intensity θ_Y, *a priori*. This should ensure greater day-to-day stability in the optimal parameter set, which is a key feature if the model is used for trading and hedging purposes. However, if this principle is adopted, it is also key to run a parallell unconstrained model calibration to detect possible long-term trends in the fixed parameters. As has been discussed in Section 8, the implied parameter set's time series behaviour has in general been economically reasonable—it suggests increasing volatility of default intensities, higher dependence in the portfolio and increasing jump risks. Overall, we therefore believe that the model has performed quite well with respect to the three goals described above. In particular, we find the model's combination of (a) ability to fit market data, and (b) model parameters that are economically interpretable *and* respond logically to changing market conditions, very appealing. In our view, this makes the model particularly interesting for risk management purposes. A portfolio of credit derivatives could be stress tested against different changes in the implied parameter set, resulting in sensitivities of the portfolio's value with respect to the different parameters. Since the parameters are quite intuitive and economically interpretable it should be possible to obtain a good qualitative understanding of the portfolio risk, in addition to quantitative risk measures.

10 Suggestions for Future Research

Our results suggest that the model is flexible enough to describe a wide range of market scenarios. In addition, the model parameters are economically intuitive and several of them behaved as one would expect during the credit crisis. Clearly, the model has many desirable features and advantages. This thesis has not explored the technical aspects of the model calibration in great detail and much research could be done on that topic. Rather, we chose a fitting criterion that was as simple as possible as a starting point for evaluating the model. As mentioned above, from a market participant's (e.g. a trading desk or a risk manager), not only is it important that the model can be calibrated to market data, but it is also key that the optimal parameter set doesn't change too much during a short period of time for no obvious reason. First, one interesting topic for future research would be to look at adding a penalty term reflecting the deviation from the latest optimal parameter set to the fitting criterion. A question that naturally arises is how much weight that should be assigned to the penalty term compared to the fit to market prices. Second, the fitting criterion itself is an interesting research topic—one could experiment with different definitions of the pricing error. Again, the goal would be to find a specification that combines (i) good fit to market data across instruments and maturities and (ii) a calibrated parameter set that isn't too volatile.

References

Andersen, L., and J. Sidenius (2005): "Extensions to the Gaussian copula: random recovery and random factor loadings," *Journal of Credit Risk*, 1(1), 29–70.

Bhansali, V., R. Gingrich, and F. A. Longstaff (2008): "Systemic credit risk: What is the market telling us?," Working Paper, UCLA Anderson School.

Das, S., L. Freed, G. Geng, and N. Kapadia (2002): "Correlated default risk," Working Paper, Santa Clara University.

Duffie, D., and N. Gârleanu (2001): "Risk and Valuation of Collaterized Debt Obligations," *Financial Analysts Journal*, 57(1), 41–59.

Duffie, D., and K. Singleton (1999): "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies*, 12(4), 687–720.

Eckner, A. (2007a): "Computational techniques for basic affine models of portfolio credit risk," Working Paper, Stanford.

——— (2007b): "Risk Premia in Structured Credit Derivatives," Working Paper, Stanford.

Feldhütter, P. (2008): "An Empirical Investigation of an Intensity-Based Model for Pricing CDO Tranches," Working Paper, Copenhagen Business School.

Hull, J., and A. White (2004): "Valuation of a CDO and nth to Default CDS without Monte Carlo simulation," *Journal of Derivatives*, 12(2), 8–23.

Jarrow, R., and S. Turnbull (1995): "Pricing Options on Financial Securities Subject to Default Risk," *Journal of Finance*, 50(1), 53–85.

Kalemanova, A., B. Schmid, and R. Werner (2007): "The Normal Inverse Gaussian Distribution for Synthetic CDO Pricing," *Journal of Derivatives*.

Lando, D. (1994): "Three Essays on Contingent Claims Pricing," Ph.D. thesis, Cornell University.

Leland, H., and K. Toft (1996): "Optimal capital structure, endogenous bankcruptcy, and the term structure of credit spreads," *Journal of Finance*, 51, 987–1019.

Li, D. X. (2000): "On Default Correlation: A Copula Function Approach," The RiskMetrics Group, Working Paper 99-07.

Longstaff, F., S. Mithal, and E. Neis (2005): "Corporate yield spreads: default risk or liquidity? New evidence from the credit-default swap market," *Journal of Finance*, 60(5), 2213–2253.

Longstaff, F. A., and A. Rajan (2008): "An Empirical Analysis of the Pricing of Collaterized Debt Obligations," *Journal of Finance*, 63, 529–563.

Mashal, R., and M. Naldi (2002): "Pricing Multiname Credit Derivatives: Heavy Tailed Hybrid Approach," Working Paper, Columbia Business School.

McGinty, L., E. Beinstein, R. Ahluwalia, and M. Watts (2004): "Credit correlation: a guide," JP Morgan, Credit Derivatives Strategy.

McNeil, A. J., R. Frey, and P. Embrechts (2005): *Quantitative Risk Management*. Princeton University Press.

Merton, R. (1974): "On the pricing of corporate debt: the risk structure of interest rates," *Journal of Finance*, 29, 449–470.

Mortensen, A. (2006): "Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models," *Journal of Derivatives*.

O'Kane, D., and M. Livesey (2004): "Base Correlation Explained," Lehman Brothers, QCR Quarterly 2004-Q3/4.

Scheicher, M. (2008): "How has CDO market pricing changed during the turmoil? Evidence from CDS index tranches," Working Paper, ECB.

van der Voort, M. (2007): "Factor Copulas: External Defaults," *Journal of Derivates*.

A Credit Derivatives Pricing

For completeness, in this section we provide valuation formulas for CDS contracts, credit indices and tranches. These formulas are not unique to the basic AJD framework but hold for any model of portfolio credit risk as long as the widely adopted conditions 1 and 2 of Assumption 1 are valid.

A.1 CDS Pricing

Let t denote the current time and consider a credit default swap on a reference name i with maturity T, notional amount N_i and CDS spread at time t equal to $c^{\text{cds}}_{i,t,T}$. Further, let $\{t_j : 1 \leq j \leq n\}$ denote the set of payment dates when coupons payments are due. First, consider the fixed leg payments from the buyer of protection. The market value is

$$V_i^{\text{cds,Fixed}}(t) = E_t^{\mathbb{Q}}\left(\sum_{j:t_j>t} e^{-\int_t^{t_j} r_s ds} 1_{\tau_i > t_j} N_i c^{\text{cds}}_{i,t,T} \frac{t_j - t_{j-1}}{360}\right) \tag{17}$$
$$- N_i c^{\text{cds}}_{i,t,T} \frac{t - \max(t_j : t_j \leq t)}{360}$$

where the second term represents the accrued coupon between time t and the prior payment and τ_i is the default time of the reference entity i under the risk-neutral measure $\mathbb{Q}$, conditional on information available at time t. The market value of the protection leg, i.e. the payments made by the protection seller, equals

$$V_i^{\text{cds,Prot}}(t) = E_t^{\mathbb{Q}}\left(e^{-\int_t^{\tau_i} r_s ds} 1_{\tau_i \leq t_n} W(\tau_i)\right) \tag{18}$$

where

$$W(\tau_i) = N_i(1 - R_i) - N_i c^{\text{cds}}_{i,t,T} \frac{\tau_i - \max(t_j : t_j < \tau_i)}{360} \tag{19}$$

is the payment due from the protection seller in case of default at time τ_i. Note that this payment is in general stochastic since the recovery rate R_i could be random. The second term represents the accrued coupon payment between the time of default and the prior payment date.

Subject to the conditions in Assumption 1, the value of the fixed leg (17) can be simplified to

$$V_i^{\text{cds,Fixed}}(t) \approx N_i c^{\text{cds}}_{i,t,T} \sum_{j:t_j>t} B_t(t_j)\mathbb{Q}_t(\tau_i > t_j)\frac{t_j - t_{j-1}}{360} - N_i c^{\text{cds}}_{i,t,T}\frac{t - \max(t_j : t_j \leq t)}{360} \tag{20}$$

where $\mathbb{Q}_t(\tau_i > s)$ is the probability of survival at least until s conditional on all information available at time t and $B_t(s)$ is the time t discount factor with maturity s derived from a risk-free interest rate curve. Under the assumption that a default always occurs right between two fixed leg coupon payments, the value of the protection leg can be approximated by

$$V_i^{\text{cds,Prot}}(t) \approx N_i \sum_{j:t_j>t} B_t\left(\frac{\max(t_{j-1}, t) + t_j}{2}\right)\mathbb{Q}_t(\tau_i \in (\max(t_{j-1}, t), t_j]) \tag{21}$$
$$\times \left[(1 - E^{\mathbb{Q}}(R_i)) - \frac{c^{\text{cds}}_{i,t,T}}{2}\frac{(t_j - \max(t_{j-1}, t))}{360}\right].$$

Here, it is key to note that we have utilized that under Assumption 1 the recovery rate R_i is independent of other stochastic variables under the risk-neutral measure $\mathbb{Q}$.

Therefore, it is possible to simply replace the recovery rate with its expected value under $\mathbb{Q}$, typically assumed to be 40 %. The fair CDS spread is easily obtained by combining equations (20) and (21) and solving for the spread. Note that the solution is known in closed form since the survival probabilities, e.g. $\mathbb{Q}_t(\tau_i > s)$, are known in closed form in the proposed basic AJD framework.

A.2 Credit Index Pricing

Consider a credit index with maturity T and m reference names. The index is spread at time t is denoted by $c^{\text{idx}}_{t,T}$. The index notional N^{idx}_s at time $s \geq t$ is given by

$$N^{\text{idx}}_s = \sum_{i=1}^{m} N_i 1_{\tau_i > s}, \tag{22}$$

i.e. the sum of notionals N_i that have not defaulted at time s. Note that N^{idx}_s is a random variable as we do not know how much notional that will have survived at a given future time s. Under Assumption 1, the market value of the fixed leg (payments from protection buyer) is similar to the corresponding formula for a CDS:

$$\begin{aligned} V^{\text{idx, Fixed}}(t) \approx c^{\text{idx}}_{t,T} \sum_{j:t_j>t} B_t(t_j) \frac{t_j - t_{j-1}}{360} E^{\mathbb{Q}_t}(N^{\text{idx}}_{t_j}) \\ - N^{\text{idx}}_{t_{j-1}} c^{\text{idx}}_{t,T} \frac{t - \max(t_j : t_j \leq t)}{360} \end{aligned} \tag{23}$$

where the second term is the accrued coupon from the latest payment date. Assuming that defaults occur half-way between subsequent coupon dates, the market value of the protection leg is given by

$$\begin{aligned} V^{\text{idx,Prot}}(t) \approx \sum_{j:t_j>t} B_t\left(\frac{\max(t_{j-1}, t) + t_j}{2}\right) E^{\mathbb{Q}_t}(N^{\text{idx}}_{t_j} - N^{\text{idx}}_{\max(t,t_j-1)} \\ \times \left[(1 - E^{\mathbb{Q}}_t(R) - \frac{c^{\text{idx}}_{t,T}}{2} \frac{(t_j - \max(t_{j-1}, t))}{360}.\right] \end{aligned} \tag{24}$$

To obtain the fair index spread, we set (23) equal to (24) and solve for $c^{\text{idx}}_{t,T}$. The expected index notional under the risk-neutral measure is given by

$$E^{\mathbb{Q}}_t = (N^{\text{idx}}_s) = \sum_{i=1}^{m} N_i \mathbb{Q}_t(\tau_i > s). \tag{25}$$

The survival probabilities, $\mathbb{Q}_t(\tau_i > s)$, are known in closed-form in the basic AJD framework. Therefore, also the expected index notional (25) and the fair index spread have closed-form solutions.

A.3 Credit Tranche Pricing

Consider a credit tranche (the j-th tranche) with maturity T, spread $c^{\text{tr}}_{j,t,T}$ and lower and upper attachment points A_j and B_j, respectively. The definition of the remaining notional is slightly different than for a credit index, and is given by

$$N^{\text{tr}}_s = \sum_{i=1}^{m} N_i - L^{\text{tr}}_s \tag{26}$$

where

$$L^{\text{tr}}_s = \sum_{i=1}^{m} N_i (1 - R_i) 1_{\tau_i \leq s} \tag{27}$$

is the total portfolio loss until time s. The notional amount of tranche j at time s is denoted by $N^{\text{tr}}_{j,s}$ and given by

$$N^{\text{tr}}_{j,s} = N(B_j - A_j) - [\max(L^{\text{tr}}_s - A_j N, 0) - \max(L^{\text{tr}}_s - A_j N, 0)]. \tag{28}$$

Under Assumption 1, the value of the fixed leg is given by

$$\begin{aligned} V^{\text{tr,Fixed}}_j(t) \approx c^{\text{tr}}_{j,t,T} \sum_{k:t_k>t} B_t(t_k) \frac{t_k - t_{k-1}}{360} E^{\mathbb{Q}}_t(N^{\text{tr}}_{j,t_k}) \\ - N^{\text{tr}}_{j,t} c^{\text{tr}_{j,t,T}} \frac{t - \max(t_k : t_k \le t)}{360} \end{aligned} \tag{29}$$

where the second term represents the accrued coupon payment. Assuming that defaults occur half-way between subsequent coupon dates, the market value of the protection leg is given by

$$\begin{aligned} V^{\text{tr,Prot}}_j(t) \approx \sum_{k:t_k>t} B_t\left(\frac{\max(t_{k-1}, t) + t_k}{2}\right) E^{\mathbb{Q}}_t(N^{\text{tr}}_{j,t_k} - N^{\text{tr}}_{j,\max(t,t_k-1)}) \\ \times \left[1 - \frac{c^{\text{tr}}_{j,t,T}}{2} \frac{(t_j - \max(t_{j-1}, t))}{360}.\right] \end{aligned} \tag{30}$$

where the second term in the brackets represents the excepted accrued coupon at the time of default. The fair tranche spread is obtained by setting (29) equal to (30) and solving for $c^{\text{tr}}_{j,t,T}$. The expected tranche notional $E^{\mathbb{Q}}_t(N^{\text{tr}}_{j,s})$ is not known in closed form, but can be calculated with computationally efficient methods in the proposed Duffie-Gârleanu framework. Therefore, also the fair tranche spreads can be computed efficiently.