Estimating Credit Contagion in a Standard Factor Model

Daniel Rösch* — Leibniz University of Hannover

Birker Winterfeldt† — University of Regensburg

First Version: September 8, 2006
This Version: July 21, 2008

Abstract

State-of-the-art credit risk portfolio models as well as the New Basel Capital Accord consider only symmetric dependencies between borrowers in a portfolio, such as correlations. Recently, asymmetric dependencies have been introduced by Davis/Lo (2001) among others. However, statistical estimation techniques and empirical evidence on contagion are still rather scarce. The present paper provides a simple credit risk portfolio model extension to credit contagion and shows how its parameters can be easily estimated and tested. We apply our methodology to a dataset provided by Moody's Default Risk Service and find significant contagion effects. By sensitivity analyses we show how contagion can seriously affect credit losses.

Key Words: Credit Risk Models, Credit Contagion

JEL Classification: G20, G28, C51

*Professor Dr. Daniel Rösch (corresponding author), Institute of Banking and Finance, Faculty of Economics and Management, Leibniz University of Hannover, Königsworther Platz 1, 30167 Hannover, Germany
Phone: +49-511-762-4668, Fax : +49-511-762-4670
Email: daniel.roesch@finance.uni-hannover.de, Internet: http://www.danielroesch.de

†Dr. Birker Winterfeldt, Department of Statistics, Faculty of Business and Economics, University of Regensburg, 93040 Regensburg, Germany

We would like to thank participants of the workshop Risikomanagement 2006, University of Innsbruck and the Credit Risk Workshop 2007, University of Ulm for helpful comments on an earlier version of the paper.

1 Introduction

Among the most important positions on the asset side of a financial institution's balance sheet are credit risky securities, and a major task for risk managers and analytics is the appropriate modeling and forecasting of the inherent credit risk. Typically, banks and other institutions apply credit risk models for this purpose, either purchased from a vendor such as CreditMetrics or CreditRisk+ or internally developed, see e.g. Finger (1998), Credit Suisse First Boston (1998), or Bluhm/Overbeck/Wagner (2003) for overviews. Credit risk models use borrower default probabilities, losses given default, exposure sizes, correlations among borrowers, and other parameters as input variables, and derive forecasts for loss or market value distributions using analytical approaches or simulation techniques.

A common feature of most state-of-the-art credit risk models is the treatment of the dependencies between borrowers in a symmetric way, e.g. via correlations, dependencies on common risk factors or more general dependency measures in a copula framework. As many studies show, the strength of dependency crucially affects the shape of the derived distributions and key risk figures such as Value-at-Risk or Expected Shortfall, see Frey/McNeil/Nyfeler (2001) and Hamerle/Rösch (2005). This symmetric view can also be found in the New Basel Capital Accord where a key driver for banks' regulatory capital is correlation, see Basel Committee on Banking Supervision (2004).

More recently, researchers and practitioners have argued that, besides symmetric dependencies due to common risk factors, asymmetric dependencies might exist as well. These effects are often called infection or contagion effects because the dependency works in one direction only. For example, the default of a large automobile company might cause financial distress for its suppliers, while conversely the automobile company might not be affected by a default of one of its suppliers.

One of the first models that explicitly model credit contagion was developed by Davis/Lo (2001, DL hereafter). They consider a portfolio where the default of any company may infect any other company in the portfolio. Under this kind of contagion, the portfolio loss distribution can easily be derived. An extension of the model is provided by Egloff/Leippold/Vanini (2004), who use neural-network-like connections between borrowers which allow for a variety of inter-firm infections. However, this model cannot be as easily applied in practice as the simple DL model because detailed information regarding the microstructural dependencies is needed. Another model by Neu/Kühn (2004, NK hereafter) incorporates contagion effects into a CreditMetrics-like credit risk model, thereby linking contagion with state-of-the-art models.

The DL and NK models build the starting point for our analysis. We first discuss some limitations of the former model when it is applied to real-world data. Then, we develop a simple contagion extension of factor models that are used in most credit risk models. The extension is similar to NK, but is not as portfolio constrained as their approach.

Our main contribution consists of deriving a framework for empirical estimation and calibration of contagion effects. We apply our methodology to rating data from Moody's Default Risk Service and conduct sensitivity analyses. Our main findings are that contagion effects are significant and can seriously affect loss distributions.

The rest of the paper is organized as follows: Section 2 provides a short discussion of the DL model. In section 3 we describe our contagion model extension of usual credit factor models and derive the estimation framework. Section 4 provides a description of the data used for our analysis and presents the empirical results. Section 5 concludes.

2 Reviewing the Davis/Lo-Model

Davis/Lo (2001) were among the first to model contagion effects in a bond portfolio. They assume that any bond may default either directly or may be infected by any defaulting bond in the portfolio. p denotes the probability of a direct default, n the number of bonds in the portfolio, and q the probability with which a defaulting bond infects another bond. Then the expected default rate $E\left[DR\right]$ in a portfolio equals:

$$E\left[DR\right] = 1 - (1-p)\left(1-pq\right)^{n-1}. \tag{1}$$

As can be taken from equation (1) the expected default rate depends not only on the parameters p and q but also on the portfolio size n. Therefore, these parameters cannot be interpreted on a stand alone basis. The probability of any firm to be infected, and thus the expected default rate of the portfolio, increases with the number of firms in the portfolio.

Despite its intuitiveness and simplicity, the model has some limitations when applied to real-world portfolios. First of all, it is assumed that all bonds within a portfolio may be infected by other bonds in the same portfolio only. In reality, contagion effects cross portfolio borders. Bonds outside the portfolio are not considered in the model. When two identical portfolios consisting of 50 bonds each (e.g. $p = 5\%$ and $q = 5\%$, expected default rate according to model: 15.97%) are combined into a single portfolio ($n = 100$), the model results in a jump of the expected default rate to 25.85% for constant parameters p and q. This result indicates that the model holds only in a portfolio of firms with no connections to firms outside the portfolio.

Another shortcoming becomes apparent when estimating the unknown parameters p and q from an empirical time series of defaults. The parameters can easily be estimated from historical data e.g. by maximum-likelihood estimation. However, if the portfolio size varies over time the resulting parameter estimates cannot be interpreted. Another constraint is the mathematical operability. The distribution function requires the computation of a sum of binomial coefficients which becomes cumbersome for large portfolios.

In the following, we propose an alternative credit contagion model which can be estimated from historical data, is mathematically simple, and can be applied to portfolios of every size.

3 The Models

3.1 Standard Credit Risk Factor Model Specification

Our model is an extension to one of the most popular credit factor representations as it is used in CreditMetrics and also in the Basel II Capital Accord. We assume a default mode frame-work with a discrete-time horizon. Consider a continuous variable $R_{i,t}$ of borrower i in time period t ($i \in \mathbb{I}_t$, $t = 1, \ldots, T$) which may be interpreted as some creditworthiness index, e.g. the return on the firm's assets. $\mathbb{I}_t$ is the set of firms in time period t. Then the credit default event is modeled as the event that the creditworthiness of the firm crosses some threshold c_i, i.e.

$$R_{i,t} < c_i \quad \Leftrightarrow \quad D_{i,t} = 1, \tag{2}$$

where:

$$D_{i,t} = \begin{cases} 1 & \text{borrower } i \text{ defaults in period } t \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

is the default indicator ($i \in \mathbb{I}_t$, $t = 1, \ldots, T$). In the CreditMetrics (and the Basel II) framework the creditworthiness indexes are assumed to follow Gaussian copulas, that is:

$$R_{i,t} = \sqrt{\rho} \cdot F_t + \sqrt{1-\rho} \cdot U_{i,t} \tag{4}$$

where $F_t \sim N(0,1)$ and $U_{i,t} \sim N(0,1)$ are both normalized i.i.d. random variables and independent from each other ($i \in \mathbb{I}_t$, $t = 1, \ldots, T$). F_t is a systematic risk factor which drives all credit qualities jointly while $U_{i,t}$ are idiosyncratic, borrower-specific risk factors. $\sqrt{\rho}$ is the factor loading of the systematic factor with ρ representing the asset correlation.

Given a realization of the systematic risk factor, the conditional probability of default is:

$$\pi_i(f_t) = P(R_{i,t} < c_i | F_t = f_t) = \Phi\left(\frac{c_i - \sqrt{\rho} \cdot f_t}{\sqrt{1-\rho}}\right) \tag{5}$$

with expectation (the "probability of default"):

$$\pi_i = \int_{-\infty}^{\infty} \Phi\left(\frac{c_i - \sqrt{\rho} \cdot f_t}{\sqrt{1-\rho}}\right) d\Phi(f_t) = P(R_{i,t} < c_i) = \Phi(c_i) \tag{6}$$

where $\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution. In the following, we assume for ease of exposition that the borrowers are homogenous w.r.t. the parameters c and ρ within a risk segment.

3.2 Credit Contagion

We now assume that, besides a set of firms which follow the above standard specification, a distinct set of firms exists. These firms additionally depend upon the first group via credit contagion. That is, default events in the first group infect firms in the second group, causing an increase of their default probabilities but not vice versa. All firms can be assigned to either the infecting firms ("I-firms") or infected firms ("C-firms"). This strict ex ante segmentation, which is also used by Jarrow/Yu (2001), eliminates looping defaults. Let $\mathbb{I}_t$ be the set of infecting firms at time t, and N_t^I be their cardinal number, i.e. $N_t^I = |\mathbb{I}_t|$. The process for the infecting firms is then given as above by ($i \in \mathbb{I}_t$, $t = 1, 2, \ldots, T$):

$$R_{i,t}^I = \sqrt{\rho} \cdot F_t + \sqrt{1-\rho} \cdot U_{i,t}. \tag{7}$$

Analogously, let $\mathbb{C}_t$ be the set of infected firms at time t, and N_t^C be their number. Then we extend the standard factor specification for the infection of the second group of firms to ($j \in \mathbb{C}_t$, $t = 1, 2, \ldots, T$):

$$\begin{aligned} R_{j,t}^C &= \sqrt{\rho} \cdot F_t + \sqrt{1-\rho} \cdot U_{j,t} - \beta \cdot \frac{\sum_{i \in \mathbb{I}_t} D_{i,t}^I(f_t)}{N_t^I} \\ &= \sqrt{\rho} \cdot F_t + \sqrt{1-\rho} \cdot U_{j,t} - \beta \cdot \frac{D_t^I(f_t)}{N_t^I}. \end{aligned} \tag{8}$$

$D_t^I(f_t) = \sum_{i \in \mathbb{I}_t} D_{i,t}^I(f_t)$ is the number of defaulting infecting firms at time t, β denotes an unknown coefficient that measures the impact of contagion on the default probability. The effect of contagion on firm j is then β times the default rate of the infecting firms. Note that if β equals zero there is no contagion at all and the model reduces to the standard factor model.

While the probabilities of default for the infecting firms are still given by the "autonomous" probabilities (5) and (6), the probabilities of default of the infected firms now depend additionally on the default rate of the contaminating firms. Conditional on the risk factor *and* the number of defaulting infectors $D_t^I(f_t) = d_t^I(f_t)$, one obtains the conditional probability:

$$\pi^C\left(f_t, d_t^I(f_t)\right) = \Phi\left(\frac{c - \sqrt{\rho} \cdot f_t + \beta \cdot \frac{d_t^I(f_t)}{N_t^I}}{\sqrt{1-\rho}}\right). \tag{9}$$

Therefore, the default probabilities of the infected firms increase with the contagion coefficient β.

3.3 Model Estimation

After outlining the model framework, we will calibrate the models from observed data. We suggest a maximum-likelihood approach. For standard factor models, i.e. without contagion, this approach has been used e.g. by Gordy/Heitfield (2000) for the Gaussian specification of CreditMetrics, or Frey/McNeil (2003) and Hamerle/Rösch (2007) for other Bernoulli mixture models, as used by CreditRisk+ or CreditPortfolioView.

The likelihood is based on the respective probabilities of observing particular numbers of defaulted infectors and defaulted infected firms. Conditional on the common systematic factor f_t, the probability of observing $d_t^I(f_t) = \sum_{i \in \mathbb{I}_t} d_{i,t}^I(f_t)$ defaulting infecting firms ($d_t^I(f_t) = 0, 1, \ldots, N_t^I$) is given by:

$$P(d_t^I(f_t)|f_t) = \binom{N_t^I}{d_t^I(f_t)} \cdot \pi^I(f_t)^{d_t^I(f_t)} \cdot \left(1 - \pi^I(f_t)\right)^{(N_t^I - d_t^I(f_t))}, \tag{10}$$

where:

$$\pi^I(f_t) = \Phi\left(\frac{c - \sqrt{\rho} \cdot f_t}{\sqrt{1-\rho}}\right) \tag{11}$$

is the homogenous conditional default probability of an infecting borrower.

Conditional on the common risk factor *and* the default frequency of the infectors, we obtain the probability of observing $d_t^C = \sum_{j \in \mathbb{C}_t} d_{j,t}^C(f_t)$ defaulting contaminated firms ($d_t^C(f_t) = 0, 1, \ldots, N_t^C$) as:

$$P(d_t^C(f_t)|f_t, d_t^I(f_t)) = \binom{N_t^C}{d_t^C(f_t)} \cdot \pi^C(f_t, d_t^I(f_t))^{d_t^C(f_t)} \cdot \left(1 - \pi^C(f_t, d_t^I(f_t))\right)^{(N_t^C - d_t^C(f_t))}, \quad (12)$$

where $\pi^C(f_t, d_t^I(f_t))$ is the homogenous conditional default probability of the infected firms from equation (9).

Due to the rule of conditional probability where the joint probability of two events A and B is given by $P(A \cap B) = P(A|B) \cdot P(B)$, the joint probability of observing $d_t^I(f_t)$ defaulting infectors and $d_t^C(f_t)$ defaulting infected firms is:

$$P(d_t^I(f_t), d_t^C(f_t)|f_t) = P(d_t^I(f_t)|f_t) \cdot P(d_t^C(f_t)|f_t, d_t^I(f_t)). \quad (13)$$

Finally, if we observe these default patterns for a whole time series of independent years, the log-likelihood function is:

$$l(c, \rho, \beta) = \sum_{t=1}^{T} \ln \left\{ \int_{-\infty}^{\infty} P(d_t^I(f_t), d_t^C(f_t)|f_t)\, \varphi(f_t)\, df_t \right\}. \quad (14)$$

For a given time series of default data this function is optimized with respect to the parameters c, ρ, and β.

3.4 Model Extensions

The contagion model as presented in section 3.2 takes into account just one segment consisting of infecting and infected firms. All infecting firms are assumed to influence the default processes of all infected firms. This universal dependency structure does not seem adequate in reality, where the bankruptcy of a firm in one business sector is not likely to directly affect the default probabilities of firms in other business sectors. Therefore, it seems adequate to assume contagion channels within defined sectors only. However, the default processes of all firms are driven by the same systematic risk factor F_t. Moreover, the assumption of homogeneous default probabilities of all firms within a business sector is an unnecessary restriction that we will drop.

For example, we may think of the standard credit factor model segmenting by rating grade, and credit contagion within a business sector. We may therefore separate firms within a homogeneous rating grade into different sectors and divide each sector into those firms which infect other firms and those which are affected by contagion. Figure 1 illustrates the contagion channels within a business sector.

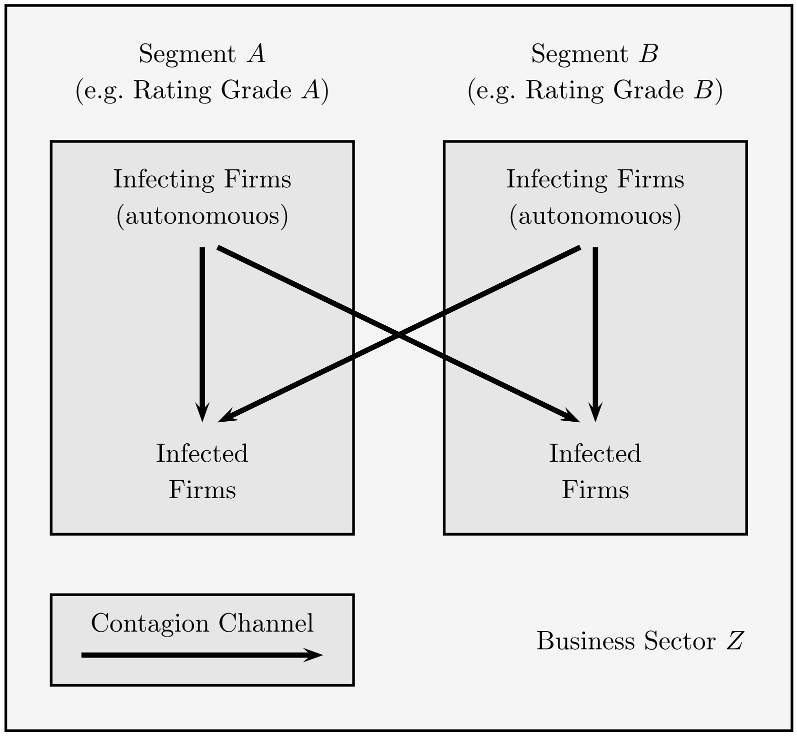

Figure 1: Contagion channels

3.5 Simulation Study

Before we present the empirical results of our estimation model we will demonstrate the robustness of the estimation procedure in a simulation study. Consider a synthetic portfolio consisting of 8,000 firms which can be mapped unambiguously to one of three sectors (X, Y, Z) as well as to one of two rating grades (A, B). 1,000 firms belong to the sector X, 2,000 to the sector Y, and 5,000 to the sector Z. In a first step, it is assumed that within each sector $\theta = 20\%$ of the firms may infect the remaining 80%, no matter whether or not the firms belong to the same rating grade. Within each sector/contagion-group combination 50% of the firms belong to segment A (rating grade A) and 50% to segment B (rating grade B). Each segment is assumed to be homogeneous in terms of default probabilities and correlations. Table 1 illustrates the number of firms in the respective segments.

		rating grade A	rating grade B	sum
sector X	infecting firms	100	100	200
	infected firms	400	400	800
sector Y	infecting firms	200	200	400
	infected firms	800	800	1,600
sector Z	infecting firms	500	500	1,000
	infected firms	2,000	2,000	4,000
sum		4,000	4,000	8,000

Table 1: Number of firms in respective segments (simulation study)

We now conduct a simulation study consisting of the following steps:

1. Generate a time series of simulated defaults in the portfolio according to the stochastic model described in section 3.4. The length of the simulated time series equals 20 years which is also the length of the time series used for the empirical analysis.
2. Estimate the parameters c^A, c^B, ρ^A, ρ^B, and β from the simulated default data using a maximum-likelihood function.
3. Repeat steps 1 to 2 10,000 times.

The assumed probabilities of default, the asset correlations as well as the contagion factor β can be taken from table 2. The parameter values are chosen arbitrarily, different values lead to qualitatively similar results.

	rating grade A	rating grade B
default threshold c	-1.64485	-1.28155
corresponding autonomous default probability π	0.05	0.10
asset correlation ρ	0.20	0.10
contagion factor β	2.00	
proportion of infecting firms θ	0.20	

Table 2: True Parameters in synthetic portfolio (simulation study)

Table 3 summarizes characteristics of the parameter estimates. While the (autonomous) probabilities of default are estimated accurately, the asset correlations show a small downward bias. These results are in line with the findings of Gordy/Heitfield (2000).

	mean	std.dev.
autonomous default probability $\hat{\pi}^A = \Phi(\hat{c}^A)$	0.0500302	0.0115710
autonomous default probability $\hat{\pi}^B = \Phi(\hat{c}^B)$	0.0999514	0.0130323
asset correlation $\hat{\rho}^A$	0.1894615	0.0488790
asset correlation $\hat{\rho}^B$	0.0950753	0.0281452
contagion factor $\hat{\beta}$	2.0124706	0.1165849

Table 3: Parameter estimates (simulation study 1)

The estimation of the contagion factor β is of particular interest. As can be taken from table 3, the average estimated contagion factor ($\bar{\hat{\beta}} = 2.01$) virtually equals the true value of 2.00. Figure 2 shows a histogram of the parameter estimates $\hat{\beta}$. No abnormalities are visible.

Analogous results are obtained using different parameter combinations for the default probabilities, asset correlations, and the contagion factor. The PDs and the contagion parameter are always estimated accurately, the asset correlations show a small downward bias. All in all, the parameter estimation seems to work quite well.

In the next step, the parameter estimates are analyzed when no contagion effect is assumed in the generation of the time series of simulated defaults ($\beta = 0$). Except for β the parameter values of table 2 still hold. The resulting parameter estimates are presented in table 4.

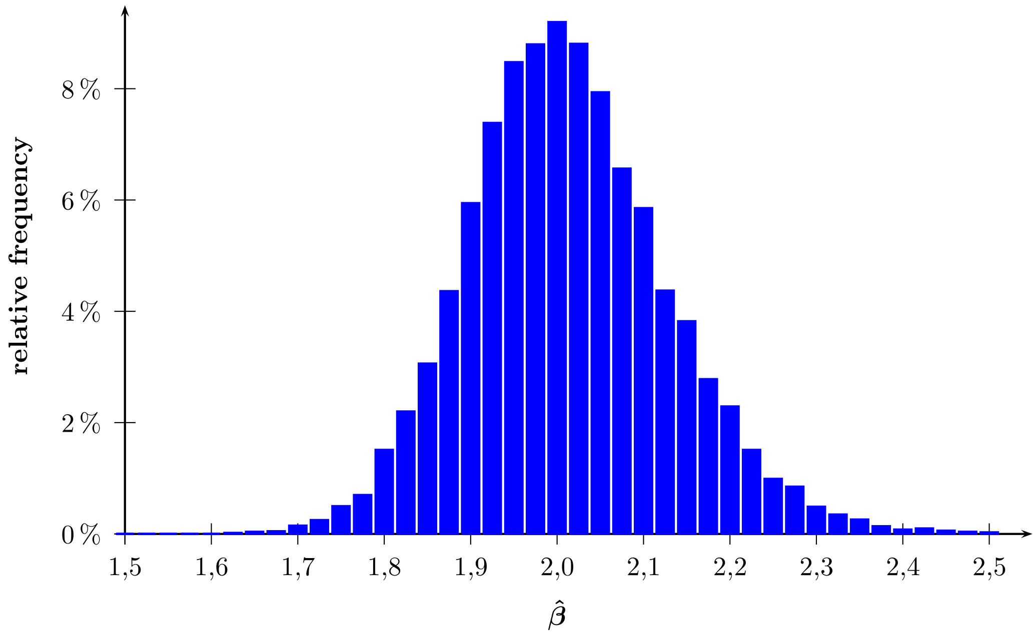

Figure 2: Histogram of parameter estimates $\hat{\beta}$

	mean	std.dev.
autonomous default probability $\hat{\pi}^A = \Phi(\hat{c}^A)$	0.0499327	0.0116084
autonomous default probability $\hat{\pi}^B = \Phi(\hat{c}^B)$	0.0998806	0.0128940
asset correlation $\hat{\rho}^A$	0.1891532	0.0500241
asset correlation $\hat{\rho}^B$	0.0948750	0.0281223
contagion factor $\hat{\beta}$	0.0032047	0.1085583

Table 4: Parameter estimates (simulation study 2)

The contagion factor continues to be estimated accurately ($\bar{\hat{\beta}} = 0.003$). This indicates that the likelihood function is suitable for both portfolios with and without contagion effects. If the analyzed time series of defaults does not contain contagion effects the estimated contagion factor is near zero. This implies that statistically significant contagion factors estimated from empirical data indeed indicate contagion effects.

In a last step of our simulation study, we analyze the parameter estimates if contagion effects are considered when generating the time series of defaults but only the parameters of the standard factor model of equation (5), i.e. without the contagion parameter, are estimated. Again, the parameters of table 2 are used. Table 5 describes the resulting parameter estimates.

It is apparent that all estimated parameters exceed their true values. This is firstly due to the fact that both ρ and β are parameters describing the dependency structure of the model. High values for the asset correlation as well as for the contagion factor increase the probability of joint default events and thus the tail weight is overestimated. Since the contagion channels are independent from the rating categories, this holds both for ρ^A and ρ^B. Secondly, by credit contagion the default probabilities of the infected firms exceed their autonomous default

	mean	std.dev.
autonomous default probability $\hat{\pi}^A = \Phi(\hat{c}^A)$	0.0716119	0.0192589
autonomous default probability $\hat{\pi}^B = \Phi(\hat{c}^B)$	0.1298055	0.0212883
asset correlation $\hat{\rho}^A$	0.2453088	0.0661608
asset correlation $\hat{\rho}^B$	0.1452702	0.0497609

Table 5: Parameter estimates (simulation study 3)

probabilities. If the contagion parameter is neglected in the estimation process the default probabilities have to be increased to ensure the correct expected loss. Again, due to the independence of the contagion channels from the rating categories both π^A and π^B are affected.

We now generate loss distributions from the true parameters (see table 2) as well as from the parameter estimates including and excluding the contagion factor (see tables 3 and 5), respectively. All exposures were set to one in order to focus on the contagion effect only. Table 6 contains the average default rates as well as the popular risk measures Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) for the different models. The risk measures are defined as follows:

Definition 1 (Value-at-Risk) $\mathrm{VaR}_\alpha(X) = \inf\{x \in \Re | P(X \leq x) \geq \alpha\}$.

Definition 2 (Conditional Value-at-Risk) $\mathrm{CVaR}_\alpha(X) = \mathrm{E}\{X | X \geq \mathrm{VaR}_\alpha(X)\}$.

model	mean	std.dev.	$\mathrm{VaR}_{0.999}$	$\mathrm{CVaR}_{0.999}$
true model	811.27	713.39	5,074.00	5,562.30
estimated model (including $\hat{\beta}$)	809.42	691.69	4,926.00	5,415.66
estimated model (excluding $\hat{\beta}$)	805.68	640.08	4,212.00	4,635.97

Table 6: Descriptive statistics of loss distributions

As can be taken from table 6 the true model and the estimated model including $\hat{\beta}$ result in approximately the same loss distributions. However, if the contagion parameter is neglected like in the standard factor model, the increased parameter estimates for the default probabilities nearly ensure the correct mean loss but the weight in the tail as measured by the VaR or the CVaR is considerably underestimated. The increased asset correlations do not adequately assess the tail risk. This is also illustrated by the histograms of the estimated loss distributions in figure 3.

4 Empirical Analysis

4.1 Empirical Data

For the empirical analysis, we use yearly bond default data for the years 1983 to 2003 published by Moody's Default Risk Service. In 1983, Moody's changed its rating methods resulting in a break in the time series. Therefore, no data before 1983 is used. The dataset contains

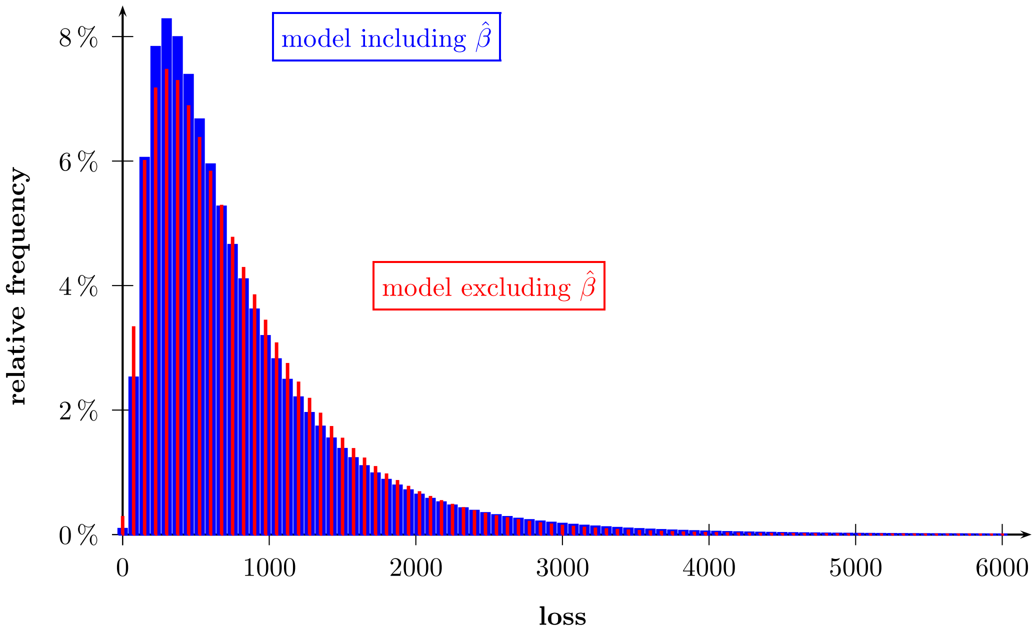

Figure 3: Loss distributions of estimated models

information regarding the respective face amounts of the exposures, the firms' countries of legal incorporation, the rating grades assigned by Moody's, and information on whether or not the bonds defaulted. Only bonds issued by US firms are used for our analysis. Altogether, we obtain a database with 30,643 firm-years including 820 defaults over the abovementioned time horizon.

All firms are mapped into two sub-portfolios named "investment grade" (rating grades Aaa to Baa) and "speculative" (rating grades Ba to C), respectively. The average default rate equals 2.68 percent, 0.16 percent, and 5.47 percent in the total portfolio, the investment grade portfolio, and the speculative portfolio, respectively. The time-dependency of the default rates can be seen in figure 4.

All firms are assigned to one of eleven Moody's broad industries. The size of the industries in the database differs considerably, with the largest (industrial) containing 20,448 firm-years and the smallest (sovereign) containing just 134 firm-years. For the following analyses, the industries banking, finance, insurance, real estate finance, securities, and thrifts are combined in the sector banking/insurance. The industries other non-bank, public utility, sovereign, and transportation constitute the sector other. The most numerous industry, "industrial", builds the third sector. The number of firms in the respective sectors can be taken from table 7.

sectors	absolute frequency	relative frequency
banking/insurance	6,308	20.59%
industrial	20,448	66.73%
other	3,887	12.68%

Table 7: Number of firms in respective sectors (empirical data)

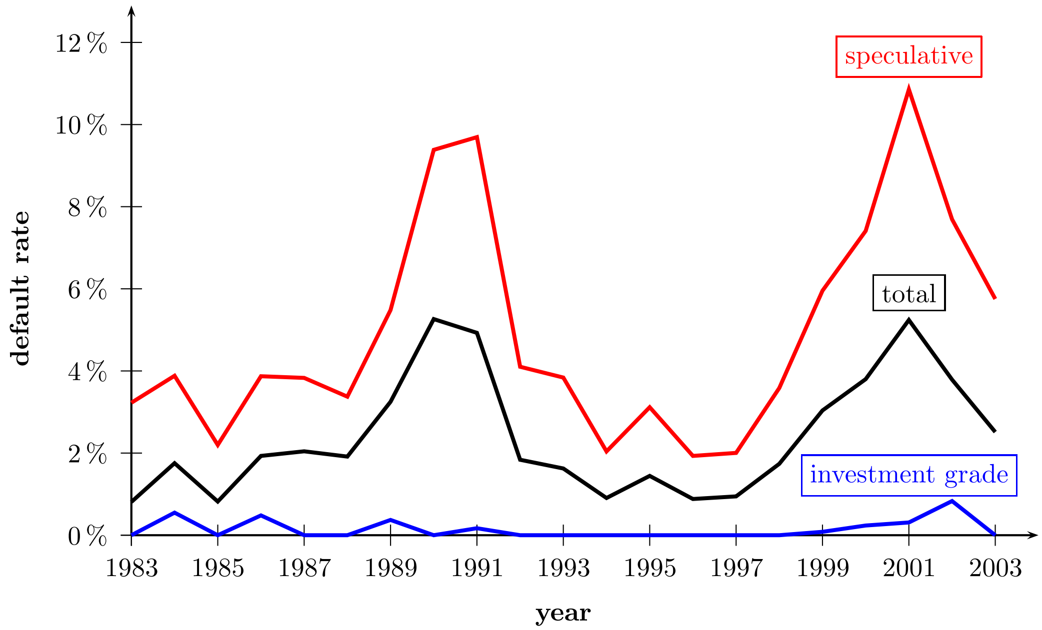

Figure 4: Yearly default rates (empirical data)

4.2 Empirical Estimation Results

We now turn to the analysis of the empirical data. First of all, some assumptions regarding the contagion channels have to be made. As already pointed out, we assume that credit contagion occurs within business sectors. The three sectors banking/insurance, industrial and other presented in section 4.1 serve this purpose. In a first step, we assume that within each sector the 20% biggest firms (measured by exposure size) may infect the remaining 80%, no matter whether or not the firms belong to the same rating grade. The intuition behind this assumption is that the default of a large bond may set signals for the entire sector. The ratio 20/80 is chosen rather arbitrarily and follows Pareto's principle stating that in many areas 20 percent of something are responsible for 80 percent of the results. Although the dataset contains the exact rating grades of all firms, the firms have to be mapped into one of two rating categories (investment grade and speculative) to ensure a sufficiently high number of defaults per rating category for a robust parameter estimation. Table 8 illustrates the number of firms in the respective segments.

Apparently, the number of infecting firms in the segment investment grade is overproportionally high. This is plausible and can be traced back to the correlation of the rating categories and the firm sizes that determine whether a firm belongs to the infecting or the infecting firms in our model. In average, bigger firms show lower probabilities of default than smaller firms.

The resulting parameter estimates for the two segments investment grade and speculative can be taken from table 9. All parameter estimates are significant at the 5% level with the exception of the asset correlation in the investment grade portfolio. This is due to the low number of observed defaults in this subportfolio resulting in a high standard error. The most interesting parameter is the contagion factor β, which equals $\hat{\beta} = 3.5532$ and is highly

		investment grade	speculative	total
banking/	infecting firms	1,128	123	1,251
finance	infected firms	4,063	994	5,057
industrial	infecting firms	2,402	1,679	4,081
	infected firms	5,469	10,898	16,367
other	infecting firms	630	140	770
	infected firms	2,411	706	3,117
total		16,103	14,540	30,643

Table 8: Number of firms in respective segments (empirical data)

	estimate	std.err.	p-value
autonomous default probability (investment grade) $\hat{\pi}^{\text{inv}}$	0.001171	0.000478	0.0235
autonomous default probability (speculative) $\hat{\pi}^{\text{spec}}$	0.04422	0.005140	< 0.0001
asset correlation (investment grade) $\hat{\rho}^{\text{inv}}$	0.1342	0.08593	0.1341
asset correlation (speculative) $\hat{\rho}^{\text{spec}}$	0.03972	0.01486	0.0090
contagion factor $\hat{\beta}$	3.5532	1.2989	0.0127

Table 9: Parameter estimates (empirical data, $\theta = 20\%$)

significant. Obviously, contagion effects indeed play a role in empirical credit portfolios and therefore should not be ignored when assessing credit risk. Mapping firms to infecting and infected firms by exposure size is a simplifying assumption. In reality, banks will have better means, allowing for alternative allocations of the firms. Moreover, setting the proportion of infecting firms θ to 20% needs justification. Here, it was chosen arbitrarily to permit the parameter estimation. However, if the infection parameter proves to be statistically significant in this very simple set-up this result should also hold for more sophisticated mappings.

The tables 10 to 12 show the parameter estimates for varying values of the assumed proportion of infecting firms in the empirical dataset ($\theta = 10\%$, $\theta = 15\%$, $\theta = 25\%$). They illustrate that the statistical significance of the infection parameter was not incidental in table 9. All parameters (again except for the asset correlation in the investment grade segment) continue to be significant at the 5% level.

Concluding, table 13 shows the parameter estimates for the empirical data in the standard factor model in which no contagion parameter is considered. Analogously to the simulation study of section 3.5 the estimates of all remaining parameters increase.

Figure 5 illustrates the loss distributions in 2003, the last year of the empirical dataset, for the models including and excluding the contagion parameter, respectively. Again, all exposures were set to one in order to focus on the contagion effect only. The proportion of infecting firms was again assumed to be $\theta = 20\%$. In total, 2,189 firms are included in the dataset in 2003.

It is apparent that the model including the contagion parameter results in a fatter tail of the loss distribution. The portfolio risk is arguably underestimated if contagion effects are neglected. This information can also be taken from table 14 containing descriptive statistics of the two loss distributions.

	estimate	std.err.	p-value
autonomous default probability (investment grade) $\hat{\pi}^{\text{inv}}$	0.001220	0.000425	0.0095
autonomous default probability (speculative) $\hat{\pi}^{\text{spec}}$	0.04575	0.005378	< 0.0001
asset correlation (investment grade) $\hat{\rho}^{\text{inv}}$	0.09831	0.07055	0.1788
asset correlation (speculative) $\hat{\rho}^{\text{spec}}$	0.04208	0.01542	0.0123
contagion factor $\hat{\beta}$	3.3825	1.6215	0.0500

Table 10: Parameter estimates (empirical data, $\theta = 10\%$)

	estimate	std.err.	p-value
autonomous default probability (investment grade) $\hat{\pi}^{\text{inv}}$	0.001180	0.000424	0.0116
autonomous default probability (speculative) $\hat{\pi}^{\text{spec}}$	0.04454	0.005310	< 0.0001
asset correlation (investment grade) $\hat{\rho}^{\text{inv}}$	0.1041	0.07079	0.1572
asset correlation (speculative) $\hat{\rho}^{\text{spec}}$	0.04207	0.01529	0.0123
contagion factor $\hat{\beta}$	3.7552	1.4940	0.0206

Table 11: Parameter estimates (empirical data, $\theta = 15\%$)

	estimate	std.err.	p-value
autonomous default probability (investment grade) $\hat{\pi}^{\text{inv}}$	0.001174	0.000483	0.0246
autonomous default probability (speculative) $\hat{\pi}^{\text{spec}}$	0.04468	0.005025	< 0.0001
asset correlation (investment grade) $\hat{\rho}^{\text{inv}}$	0.1371	0.08535	0.1240
asset correlation (speculative) $\hat{\rho}^{\text{spec}}$	0.03829	0.01458	0.0162
contagion factor $\hat{\beta}$	3.2986	1.1912	0.0118

Table 12: Parameter estimates (empirical data, $\theta = 25\%$)

4.3 Sensitivity Analysis

To conclude, we conduct sensitivity analyses of our portfolio with respect to the contagion factor β. First of all, we generate the loss distribution of the empirical portfolio for the year 2003 using the maximum-likelihood estimates shown in table 9. All exposures were set to one in order to focus on the effect of the varying contagion parameter on the loss distributions only. Ceteris paribus, we then generate loss distributions using different parameter values for the contagion factor. The values are not chosen arbitrarily but are set to the upper limits of the confidence interval of the point estimate $\hat{\beta}$ at a confidence level of 95% and 99.9%, respectively. These limits equal $\hat{\beta}_{0.95} = 6.2626$ and $\hat{\beta}_{0.999} = 8.5532$, respectively. Figure 6

	estimate	std.err.	p-value
autonomous default probability (investment grade) $\hat{\pi}^{\text{inv}}$	0.001367	0.000535	0.0189
autonomous default probability (speculative) $\hat{\pi}^{\text{spec}}$	0.05086	0.005838	< 0.0001
asset correlation (investment grade) $\hat{\rho}^{\text{inv}}$	0.1323	0.07532	0.0944
asset correlation (speculative) $\hat{\rho}^{\text{spec}}$	0.05291	0.01784	0.0076

Table 13: Parameter estimates, not considering the contagion effect (empirical data)

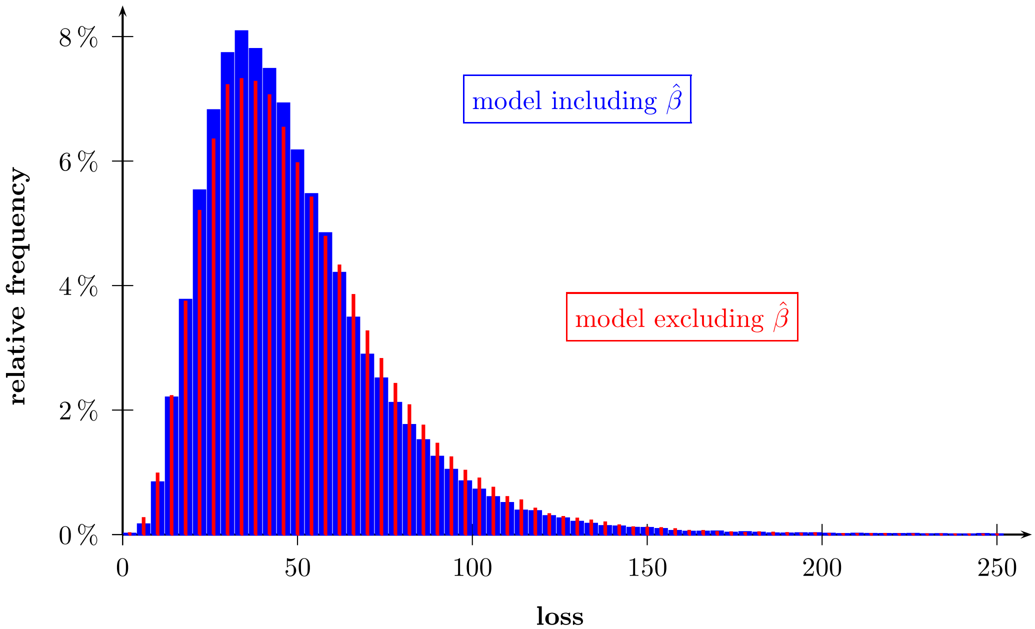

Figure 5: Estimated loss distributions 2003 (empirical data, $\theta = 20\%$)

model	mean	std.dev.	$\text{VaR}_{0.999}$	$\text{CVaR}_{0.999}$
estimated model including contagion factor	48.83	26.42	206.00	245.83
estimated model excluding contagion factor	50.34	27.08	194.00	218.67

Table 14: Descriptive statistics of loss distributions (empirical data)

shows the three resulting loss distributions. By shifting the contagion factor, the mean, the variance, and the weights in the tails increase.

The corresponding descriptive statistics of the loss distributions as well as the risk measures Value-at-Risk and Conditional Value-at-Risk can be taken from table 15. Since increasing the contagion factor also increases the expected loss, both the absolute and the relative risk measures are given. The latter are defined as follows:

Definition 3 (Value-at-Risk$^{\text{rel.}}$) $\text{VaR}_\alpha^{\text{rel.}}(X) = \inf\{x \in \Re | P(X \leq x) \geq \alpha\} - E(X)$.

Definition 4 (Conditional Value-at-Risk$^{\text{rel.}}$) $\text{CVaR}_\alpha^{\text{rel.}}(X) = \text{E}\{X | X \geq \text{VaR}_\alpha(X)\} - E(X)$.

The effect of the increased contagion factor on the quantil-based risk measures is enormous. While the expected loss increases by only 16.36% using $\hat{\beta}_{0.999} = 8.5532$ instead of the point estimate $\hat{\beta} = 3.5532$ the relative $\text{CVaR}_{0.999}$ increases by 62.95%.

It should be noted that only the loss distribution of the infected firms is affected by the factor shift, the loss distribution of the infecting firms remains completely unaffected. For a portfolio consisting of infected firms only the consequences of the factor shift are accordingly more severe. Tables 16 to 18 compare the relative increase of the risk measures for a portfolio consisting of all the infecting firms only (table 16), a portfolio consisting of all the infected firms only (table 17), and the total portfolio (table 18).

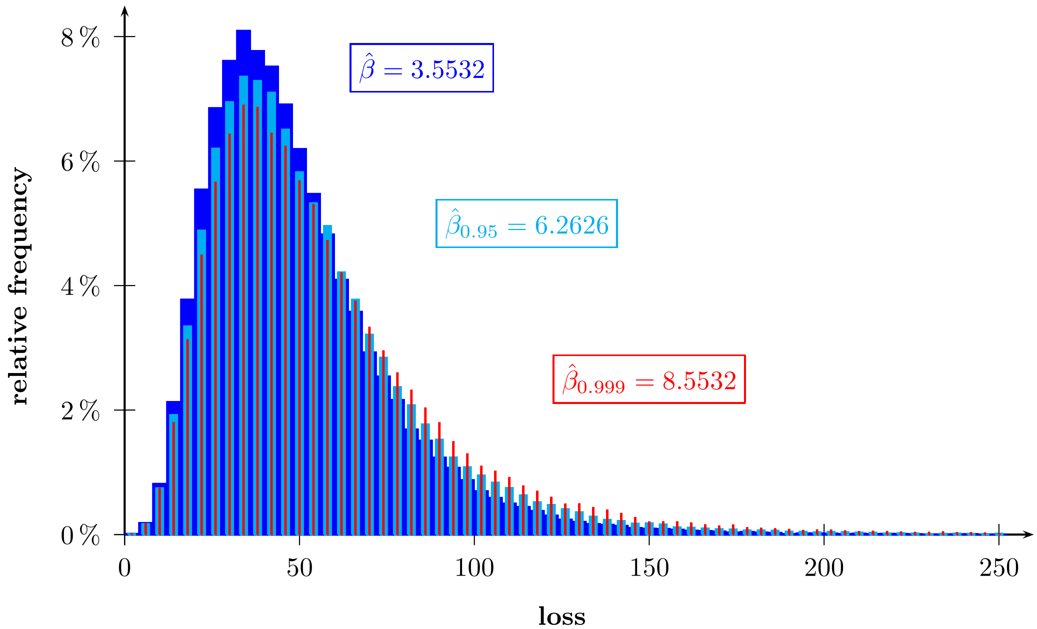

Figure 6: Loss distributions (sensitivity analysis)

	$\hat{\beta} = 3.5532$	$\hat{\beta}_{0.95} = 6.2626$	$\hat{\beta}_{0.999} = 8.5532$
mean	48.83	52.96	56.82
std.dev.	26.42	31.13	35.74
$\mathrm{VaR}_{0.999}$	206.00	251.00	302.00
$\mathrm{CVaR}_{0.999}$	245.83	307.23	377.82
$\mathrm{VaR}^{\mathrm{rel.}}_{0.999}$	157.17	198.04	245.18
$\mathrm{CVaR}^{\mathrm{rel.}}_{0.999}$	197.00	254.27	321.00

Table 15: Descriptive statistics and risk measures of loss distributions (sensitivity analysis)

As expected, the loss of the portfolio consisting of infecting firms is independent from the contagion factor, only random deviations can be observed. The (relatively) greater influence of the contagion factor on the portfolio loss of the infected firms as compared to the total portfolio is apparent.

Instead of conducting sensitivity analyses with respect to the contagion factor β one could also stress the default rate of the infecting firms $DR^I_t = D^I_t/N^I_t$. Since the contagion factor β and the default rate of the infecting firms appear in the model equations in a multiplicative way (see equation 9) the effect of stressing either factor is identical. Stressing the default rate of the infecting firms is similar to the approach by Neu/Kühn (2004) who propose to stress-test their model by setting specific firms into the default state.

	$\hat{\beta} = 3.5532$	$\hat{\beta}_{0.999} = 8.5532$	percentage change
mean	4.89	4.90	+0.04%
std.dev.	3.31	3.32	+0.27%
$\text{VaR}_{0.999}$	23.00	23.00	+0.00%
$\text{CVaR}_{0.999}$	28.28	27.81	−1.67%
$\text{VaR}^{\text{rel.}}_{0.999}$	18.11	18.10	−0.01%
$\text{CVaR}^{\text{rel.}}_{0.999}$	23.39	22.92	−2.02%

Table 16: Percentage increase of risk measures in the portfolio of infecting firms (sensitivity analysis)

	$\hat{\beta} = 3.5532$	$\hat{\beta}_{0.999} = 8.5532$	percentage change
mean	43.93	51.92	+18.18%
std.dev.	23.73	32.88	+38.56%
$\text{VaR}_{0.999}$	184.00	279.00	+51.63%
$\text{CVaR}_{0.999}$	219.35	350.85	+59.95%
$\text{VaR}^{\text{rel.}}_{0.999}$	140.07	227.08	+62.12%
$\text{CVaR}^{\text{rel.}}_{0.999}$	175.42	298.93	+70.41%

Table 17: Percentage increase of risk measures in the portfolio of infected firms (sensitivity analysis)

	$\hat{\beta} = 3.5532$	$\hat{\beta}_{0.999} = 8.5532$	percentage change
mean	48.83	56.82	+16.36%
std.dev.	26.42	35.74	+35.30%
$\text{VaR}_{0.999}$	206.00	302.00	+46.60%
$\text{CVaR}_{0.999}$	245.83	377.82	+53.69%
$\text{VaR}^{\text{rel.}}_{0.999}$	157.17	245.18	+56.00%
$\text{CVaR}^{\text{rel.}}_{0.999}$	197.00	321.00	+62.95%

Table 18: Percentage increase of risk measures in the total portfolio (sensitivity analysis)

5 Conclusion

In this paper, we have extended the standard factor model to allow for credit contagion. Contagion is assumed to occur within business sectors. Every sector consists of infecting and contaminated firms. We assume a one-way dependency structure with the probabilities of default of the contaminated firms depending on the default rate of the infecting firms in the respective industry sector.

Our key findings can be summarized as follows: We have shown in a simulation study that contagion factors can be estimated accurately from historical data using a maximum-likelihood approach. Parameter estimation from empirical default data has revealed the existence of significant contagion effects in real-world portfolios. Concluding, we have conducted sensitivity analyses for the empirical portfolio with respect to the contagion factor, demonstrating the dramatic impact of a shift of the contagion factor on various risk measures. Our findings are not only relevant for financial institutions but also for regulatory purposes. To not underestimate credit risk, contagion effects should be accounted for in credit portfolio models.

A possible extension of our model is the inclusion of sector and/or rating grade specific systematic factors. At the moment, the available historical time series of empirical default data are too scarce to estimate the necessary additional parameters. However, if more data becomes available the estimation of the additional parameters can be easily implemented.

References

Basel Committee on Banking Supervision 2004. International Convergence of Capital Measurement and Capital Standards – A Revised Framework, Bank for International Settlements, June 2004.

Bluhm, C., Overbeck, L., Wagner, C., 2003. An Introduction to Credit Risk Modeling, Chapman & Hall, London.

Credit Suisse Financial Products, 1997. CreditRisk+ – A Credit Risk Management Framework, London.

Davis, M., Lo, V., 2001. Infectious defaults, Quantitative Finance 1, 382-387.

Egloff, D., Leippold, M., Vanini, P., 2004. A simple model of credit contagion, Working Paper, University of Zurich.

Finger, C., 1998. Sticks and Stones, RiskMetrics Group Working Paper, October.

Frey, R., McNeil, A., Nyfeler, M., 2001. Copulas and Credit Models, Risk 14, October, 111-114.

Frey, R., McNeil, A., 2003. Dependent Defaults in Models of Portfolio Credit Risk, Journal of Risk 6, 59-92.

Giesecke, K., Weber, S., 2002. Credit contagion and aggregate losses, Working Paper, Cornell University and Technische Universität Berlin.

Gordy, M.B., 2000. A Comparative Anatomy of Credit Risk Models, Journal of Banking and Finance 24, 119-149.

Gordy, M.B., Heitfield, E.A. 2000. Estimating Factor Loadings When Ratings Performance Data Are Scarce, Board of Governors of the Federal Reserve System, Division of Research and Statistics, Working Paper.

Hamerle, A., Rösch, D., 2005. Misspecified Copulas in Credit Risk Models: How Good Is Gaussian?, Journal of Risk 8(1), 41-58.

Hamerle, A., Rösch, D., 2007. Parameterizing Credit Risk Models, Journal of Credit Risk 2(4), 101-122.

Jarrow, R.A., Yu, F., 2001. Counterparty Risk and the Pricing of Defaultable Securities, Journal of Finance 56, 1765-1799.

Merton, R., 1974. On the pricing of corporate debt: the risk structure of interest rates, Journal of Finance 29, 449-470.

Neu, P., Kühn, R., 2004. Credit risk enhancement in a network of interdependent firms, Physica A 342, 639-655.

Pinheiro, J.C., Bates, D.M., 1995. Approximations to the Log-likelihood Function in the Nonlinear Mixed-effects Model, Journal of Computational and Graphical Statistics 4, 12-35.