Validating Structural Credit Portfolio Models

Michael Kalkbrener, Akwum Onwunta
Risk Analytics & Instruments, Deutsche Bank AG, Frankfurt, Germany*

1 Introduction

Concentrations in a bank's credit portfolio are key drivers of credit risk. These risk concentrations may be caused by material concentrations of exposure to individual names as well as large exposures to single sectors (geographic regions or industries) or to several highly correlated sectors. The most common approach to introduce sector concentration into a credit portfolio model is through systematic factors affecting multiple borrowers. The specification of the factor model, which determines the dependence structure of rating migrations and defaults, is a central problem in credit portfolio modeling.

Different data sources can be used as input for the calibration of factor models, i.e., equity data, credit spreads or default and rating information. The most frequent approach is based on information from equity markets, mainly due to the good quality and wide coverage of this data source. Since individual firm returns are available, methodologies based on equity information allow for firm-specific pairwise correlations and more granular industry definitions than historical rating data methods. The main disadvantage of equity data is the fact that equity prices do not solely reflect the credit quality of companies but also information that is unrelated to credit risk, e.g. liquidity issues, risk aversion of market participants, etc.

Historical rating and default data, on the other hand, are not distorted by information unrelated to credit quality but require some aggregation, typically by rating class, country or industry. Hence, these data sources are natural candidates for validating the dependence structure of credit portfolio models on an aggregate level. In this paper, we shall review the most frequently used techniques for deriving correlations from default and rating data: moment estimators (Lucas, 1995; Bahar and Nagpal, 2001; de Servigny and Renault, 2002; Frey and McNeil, 2003; Jobst and de Servigny, 2005) and maximum likelihood estimators (Gordy and Heitfield, 2002; Demey et al., 2004). Our empirical results are based on S&P rating data from 1981 to 2009. We group the rated companies into cohorts that correspond to industry segments, rating classes or combinations of both. Intra-cohort and inter-cohort correlations are calculated using maximum likelihood estimation. Our results are broadly in line with previous studies reported in the literature:

(i) Correlations within industry segments are consistently higher than correlations between industry segments. In particular, this observation confirms that correlations are usually underestimated if industry information is ignored in the specification of the cohorts, i.e., if cohorts correspond to rating classes.

(ii) Correlations strongly vary across industries and over time.

(iii) On average, correlations derived from historical default and rating data are lower than equity correlations. In addition, we observed a significant difference between correlations obtained from default data and correlations obtained from rating data.

Most of our calculations are based on the standard assumption that the risk factors follow a multivariate normal distribution. In addition, we perform calculations using t-distributions and skew normal distributions. The best calibration results are obtained with a multivariate t-distribution with a high degree of freedom, e.g. in the range of 65.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of Deutsche Bank AG.

The paper has the following structure. In Section 2, we give a formal specification of the structural credit portfolio model, which is used for the analysis of historical default data, together with a review of the concepts of default and asset correlations. Different approaches for calibrating and validating factor models are discussed in Section 3. The next section is devoted to the derivation of moment estimators and maximum likelihood estimators when applied to historical default data. The actual calibration results based on S&P data are presented in Section 5. We generalize the credit portfolio model to cover rating migration and apply maximum likelihood estimation to S&P rating data in Section 6. Calibration results for non-Gaussian risk factors are presented in Section 7. In addition, we review different concepts for incorporating the credit cycle. In Section 8, we conclude the paper by comparing equity correlations with correlations derived from historical default and rating data.

2 Credit portfolio model for default risk

2.1 Definition of the model

Credit portfolio models can be divided into reduced-form models and structural (or firm-value) models. We refer to Crouhy et al. (2000), Bluhm et al. (2002) and McNeil et al. (2005) for a survey on credit portfolio modeling. In this paper, we use a structural credit portfolio model similar to CreditMetrics (Gupton et al., 1997).

The progenitor of all structural models is the model of Merton (1974), which links the default of a firm to the relationship between the firm's assets and liabilities at the end of a given time period $[0, T]$. More precisely, in a structural credit portfolio model the j-th obligor defaults if at time T its ability-to-pay variable A_j is below a default threshold c_j: the default event is defined as $\{A_j \leq c_j\} \subseteq \Omega$, where A_j is a real-valued random variable on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$ and $c_j \in \mathbb{R}$. The corresponding default indicator is denoted by

$$I_j = \mathbf{1}_{\{A_j \leq c_j\}}. \tag{1}$$

If A_j is standardized and normally distributed then the default threshold c_j is linked to the default probability $p_j \in [0, 1]$ via $c_j = N^{-1}(p_j)$, where N denotes the distribution function of a standardized normal distribution.

For each obligor, a loss variable is defined by $L_j = l_j \cdot I_j$, where $l_j \in \mathbb{R}$ denotes the loss-at-default. In order to aggregate the loss variables L_j of the individual obligors to a loss variable L of the entire portfolio a dependence structure of the A_j is specified. This is done via the introduction of a factor model consisting of systematic and idiosyncratic factors. More precisely, each ability-to-pay variable A_j is decomposed into a sum of systematic factors $\Psi_1, \ldots, \Psi_m$ and an idiosyncratic (or specific) factor ε_j, that is,

$$A_j = \sqrt{R_j^2} \sum_{i=1}^{m} w_{ji} \Psi_i + \sqrt{1 - R_j^2} \varepsilon_j. \tag{2}$$

The idiosyncratic factors are independent of each other as well as independent of the systematic factors. It is usually assumed that the idiosyncratic and systematic factors are centered and follow a multivariate Gaussian distribution. The systematic weights $w_{j1}, \ldots, w_{jm} \in \mathbb{R}$ are scaled such that the systematic component $\phi_j := \sum_{i=1}^{m} w_{ji} \Psi_i$ is a standardized normally distributed variable. Each $R_j^2 \in [0, 1]$ determines the impact of the systematic component on A_j and therefore the correlation between A_j and ϕ_j: it immediately follows from (2) that $R_j^2 = \mathrm{Corr}(A_j, \phi_j)^2$. Finally, the portfolio loss variable $L : \Omega \to \mathbb{R}$ is the sum of the loss variables of the obligors, that is

$$L = \sum_{j=1}^{n} L_j = \sum_{j=1}^{n} l_j \cdot I_j. \tag{3}$$

An important instance of this model class is the homogeneous one-factor model: assume that all obligors have the same loss-at-default $l \in \mathbb{R}$, default probability p with corresponding default threshold $c := N^{-1}(p)$

and the same R^2. In this case, (2) and (3) specialize to

$$A_j = \sqrt{R^2}\Psi + \sqrt{1-R^2}\varepsilon_j, \qquad L = \sum_{j=1}^{n} L_j = l \cdot \sum_{j=1}^{n} I_j, \tag{4}$$

where Ψ is a standardized normally distributed variable.

2.2 Default and asset correlations

Via equation (1), the credit portfolio model specifies a dependence structure for the default indicators I_i. We shall measure dependence by calculating default and asset correlations. This subsection provides a short review of the definitions and main properties.

The default or event correlation ρ_{ij}^D of obligors i and j is defined as the correlation of the respective default indicators. Because

$$\mathrm{Var}(I_i) = \mathrm{E}(I_i^2) - (\mathrm{E}(I_i))^2 = \mathrm{E}(I_i) - (\mathrm{E}(I_i))^2 = p_i - p_i^2,$$

the default correlation equals

$$\rho_{ij}^D = \mathrm{Corr}(I_i, I_j) = \frac{\mathrm{E}(I_i I_j) - p_i p_j}{\sqrt{(p_i - p_i^2)(p_j - p_j^2)}}. \tag{5}$$

If the vector of default indicators $I = (I_1, \ldots, I_n)$ is exchangeable[1] the default correlation ρ_{ij}^D has a particularly simple form (equality (7) below). To show that, we first introduce a notation for joint default probabilities:

$$\pi_k := \mathbb{P}(I_{i_1} = 1, \ldots, I_{i_k} = 1), \qquad \{i_1, \ldots, i_k\} \subseteq \{1, \ldots, n\}, k \leq n, \tag{6}$$

i.e., π_k is the probability that an arbitrarily selected subgroup of k obligors defaults. Note that for all $i, j \in \{1, \ldots, n\}$ with $i \neq j$,

$$\mathrm{E}(I_i) = \mathbb{P}(I_i = 1) = p_i = \pi_1,$$

$$\mathrm{E}(I_i I_j) = \mathbb{P}(I_i = 1, I_j = 1) = \pi_2,$$

so that $\mathrm{Cov}(I_i, I_j) = \pi_2 - \pi_1^2$. This implies that for all i, j with $i \neq j$ the default correlation in (5) is given by

$$\rho^D := \rho_{ij}^D = \frac{\pi_2 - \pi_1^2}{\pi_1 - \pi_1^2}, \tag{7}$$

which is a simple function of the first- and second-order default probabilities.

In the credit portfolio model introduced in Section 2.1, the indicator variables I_i are defined in terms of ability-to-pay variables A_i, which are often interpreted as log-returns of asset values. The correlation $\mathrm{Corr}(A_i, A_j)$ is therefore called the asset correlation ρ_{ij}^A of obligors $i \neq j$. As an immediate consequence of (2), the correlation as well as the covariance of the ability-to-pay variables of the counterparties i and j are given by

$$\mathrm{Corr}(A_i, A_j) = \mathrm{Cov}(A_i, A_j) = \sqrt{R_i^2}\sqrt{R_j^2} \sum_{k,l=1}^{m} w_{ik} w_{jl} \mathrm{Cov}(\Psi_k, \Psi_l). \tag{8}$$

Note that to specify the covariance structure of all n ability-to-pay variables it suffices to define the covariance matrix of the m systematic factors and the systematic weights and R^2s of all counterparties. Since m is usually much smaller than n, the factor model approach leads to a significant reduction in the amount of input data, i.e., the number of parameters is of the order of $m(m+n)$ instead of n^2.

[1] A vector of random variables $X = (X_1, \ldots, X_n)$ is called exchangeable if for any permutation $\Pi : \{1, \ldots, n\} \to \{1, \ldots, n\}$ the vectors $(X_1, \ldots, X_n)$ and $(X_{\Pi(1)}, \ldots, X_{\Pi(n)})$ have the same distributions.

There exists an obvious link between default and asset correlations. For given default probabilities, the default correlations ρ_{ij}^D are determined by $\mathrm{E}(I_i I_j)$ according to (5). It follows from (1) that

$$\mathrm{E}(I_i I_j) = \mathbb{P}(A_i \le c_i, A_j \le c_j) = \int_{-\infty}^{c_i} \int_{-\infty}^{c_j} f_{ij}(u,v)dudv,$$

where $f_{ij}(u,v)$ is the joint density function of A_i and A_j. Hence, default correlations depend on the joint distribution of A_i and A_j. If (A_i, A_j) is bivariate normal the correlation of A_i and A_j determines the copula of their joint distribution and hence the default correlation:

$$\mathrm{E}(I_i I_j) = \frac{1}{2\pi\sqrt{1-\rho_{ij}^{A\,2}}} \int_{-\infty}^{c_i} \int_{-\infty}^{c_j} \exp\left(-\frac{u^2 - 2\rho_{ij}^A uv + v^2}{2(1-\rho_{ij}^{A\,2})}\right) dudv. \tag{9}$$

Note that for general ability-to-pay variables outside the multivariate normal class, the asset correlations do not fully determine the default correlations.

2.3 Counting defaults

Let $D : \Omega \to \mathbb{R}$ specify the number of defaults, i.e., $D := \sum_{j=1}^n I_j$. We shall now derive a formula for the distribution of D, which will be used in the calibration techniques presented later. The derivation is based on the fact that the vector of default indicators $I = (I_1, \dots, I_n)$ satisfies the definition of a Bernoulli mixture model if $m < n$ (McNeil et al., 2005). More precisely, there exist functions $\bar{p}_j : \mathbb{R}^m \to [0,1]$, $1 \le j \le n$, such that, conditional on $\Psi = (\Psi_1, \dots, \Psi_m)$, the default indicator $I = (I_1, \dots, I_n)$ is a vector of independent Bernoulli random variables with

$$\mathbb{P}(I_j = 1 | \Psi = \psi) = \bar{p}_j(\psi), \qquad \text{for } \psi \in \mathbb{R}^m.$$

For a given value ψ of the systematic factors, the conditional independence of the I_j follows from the independence of the ε_j. The functions $\bar{p}_j$ have the form

$$\bar{p}_j(\psi) \;=\; \mathbb{P}\left(A_j \le c_j \mid \Psi = \psi\right) \;=\; \mathbb{P}\left(\varepsilon_j \le \frac{c_j - \sqrt{R_j^2}\sum_{i=1}^m w_{ji}\psi_i}{\sqrt{1-R_j^2}}\right) \;=\; N\left(\frac{c_j - \sqrt{R_j^2}\sum_{i=1}^m w_{ji}\psi_i}{\sqrt{1-R_j^2}}\right),$$

since the ε_j are $N(0,1)$-distributed variables.

For the rest of this subsection we make the simplifying assumption that the vector of default indicators $I = (I_1, \dots, I_n)$ is exchangeable. As a consequence, the functions $\bar{p}_1, \dots, \bar{p}_n$ are identical and we therefore drop the subscript. Conditional on $\Psi = \psi$, the number of defaults D is the sum of n independent Bernoulli variables with parameter $\bar{p}(\psi)$. Hence, D has a binomial distribution with parameters $\bar{p}(\psi)$ and n, i.e.,

$$\mathbb{P}(D = j \mid \Psi = \psi) = \binom{n}{j} \bar{p}(\psi)^j (1-\bar{p}(\psi))^{n-j}. \tag{10}$$

The unconditional distribution of D is obtained by integrating over the distribution F of the systematic factors Ψ:

$$\mathbb{P}(D = j) \;=\; \int_{\mathbb{R}^m} \mathbb{P}(D = j \mid \Psi = \psi)\, dF(\psi) \;=\; \binom{n}{j} \int_{\mathbb{R}^m} \bar{p}(\psi)^j (1-\bar{p}(\psi))^{n-j}\, dF(\psi).$$

In particular, for the homogeneous one-factor model defined in (4) the vector of default indicators $I = (I_1, \dots, I_n)$ is exchangeable and the unconditional distribution of D is

$$\begin{aligned} \mathbb{P}(D = j) \;&=\; \binom{n}{j} \int_{\mathbb{R}} \bar{p}(\psi)^j (1-\bar{p}(\psi))^{n-j} dN(\psi) \\ &=\; \frac{1}{(2\pi)^{1/2}} \binom{n}{j} \int_{\mathbb{R}} \bar{p}(\psi)^j (1-\bar{p}(\psi))^{n-j} e^{-\psi^2/2} d\psi, \end{aligned} \tag{11}$$

where $\bar{p}$ is specified by

$$\bar{p}(\psi) = N\left(\frac{c - \sqrt{R^2}\psi}{\sqrt{1-R^2}}\right). \tag{12}$$

3 Calibration and validation of factor models

3.1 The calibration problem

The central problem in credit portfolio modeling is the specification of the dependence structure of defaults. In the quantitative framework introduced in the previous section the probability of joint defaults is specified through a factor model consisting of systematic and idiosyncratic factors. The calibration of the factor model consists of the following steps:

(i) identification of appropriate systematic factors $\Psi_1, \dots, \Psi_m$,

(ii) estimation of the correlations of $\Psi_1, \dots, \Psi_m$,

(iii) estimation of the systematic weights w_{ji} and R_j^2, where $j = 1, \dots, n$ and $i = 1, \dots, m$.

Historical time-series of default indicators or of ability-to-pay variables are typically used as input for the calibration procedure. For example, observations $(a_1(1), \dots, a_n(1)), \dots, (a_1(s), \dots, a_n(s)) \in \mathbb{R}^n$ of equity or asset returns at different time points are often interpreted as realizations of the ability-to-pay variables $A = (A_1, \dots, A_n)$.

Several statistical techniques can be applied to the calibration of the factor model. These techniques estimate the parameters of the factor model in order to replicate the dependence of the input time-series as closely as possible. Additional requirements arise from the application of credit portfolio models in risk management. In particular, the systematic factors should have a clear economic interpretation, such as representing individual countries or industries. Furthermore, all obligors should have plausible factor weights. These features are important for the economic interpretation of risk concentrations identified by the portfolio model. They are also necessary for the implementation of economic stress scenarios in the model. These non-statistical model requirements are one of the reasons why statistical techniques for factor model calibration are usually complemented by qualitative assessments. For example, country and industry weights may be derived from balance sheet and legal information to avoid counterintuitive results from regression on market data. Note, however, that there is typically a trade-off between the accurate replication of the dependence structure of the input data and the economic interpretation of the model.

3.2 Calibration with equity data

The calibration of factor models can be classified according to the data source used as input, i.e., equity data, credit spreads or default and rating information. The most frequent approach is based on information from equity markets. The main advantage is that equity information has wider coverage both in terms of industry and geography. The information comes from liquid markets, is of good quality, and is readily available. Thus, unlike the ratings-based approach, one can empirically observe nuances between firms on a global basis. Since individual firm returns are available, methodologies based on equity information allow for firm-specific pairwise correlations and more granular industry definitions than historical rating data methods. With these methodologies, however, there is a concern that equity prices reflect - in addition to credit related information - information that is unrelated to credit risk, such as liquidity issues, risk aversion of market participants, etc.

Another important issue is the question of whether equity or asset returns can be used for the calibration of ability-to-pay variables. Following Merton's approach, a company defaults if the value of its assets is no longer sufficient to cover its liabilities. Default correlations are therefore determined by the corresponding asset correlations and default probabilities. Since asset processes cannot be observed in the market and are difficult to determine, an alternative approach is to derive default correlations from equity correlations instead of asset correlations. However, this raises the question of whether equity correlations are reliable proxies for asset correlations. The relationship between asset and equity correlations is still a controversial issue in the literature since authors come to rather different conclusions, see, for example, Zeng and Zhang

(2002) and de Servigny and Renault (2002). In Section 8, we will compare equity correlations to correlations derived from historical default and rating data.

3.3 Validation with historical default information

The main objective of the factor model in (2) is the specification of the dependence structure of default indicators $I_1, \dots, I_n$. It is therefore natural to use historical default data for model calibration. Compared to other data sources such as equity data the use of historical defaults has the important advantage that the calibration results are not distorted by information unrelated to credit quality. Time series of historical defaults are provided by rating agencies for a large pool of companies with a particularly good coverage in the US. Despite these advantages it is uncommon in the industry to rely on historical default data for formal statistical estimation of model parameters. There are good reasons for this, the main one being that, particularly for higher-rated companies, defaults are rare events. In other words, there is simply not enough relevant data on historical defaults to obtain reliable parameter estimates for individual companies by formal inference alone. As a consequence, methodologies based on historical default data require some aggregation, typically by rating class, country or industry. Hence, the main application of default data is the validation of portfolio models calibrated to equity data or credit spreads.

4 Estimators applied to default data

4.1 Homogeneous cohorts

In the following, we analyze different techniques for estimating dependence in time-series of default data. The methods we describe are motivated by the format of the data we consider, which comprises observations of the default or non-default of groups of monitored companies in a number of time periods $t = 1, \dots, s$. This kind of data is readily available from rating agencies. In each time period $t = 1, \dots, s$ we group the monitored companies into cohorts. Let $n(t)$ denote the number of obligors in a given cohort at the start of period t and $d(t)$ the number of observed defaults during that period. The companies within a cohort are assumed to be homogeneous in the following sense: conditionally on systematic factors $\Psi(t) = (\Psi_1(t), \dots, \Psi_m(t))$ their default indicators $I_1(t), \dots, I_{n(t)}(t)$ are independent and identically distributed. Hence, in each time period the defaults of a cohort are generated by an exchangeable Bernoulli mixture model. Furthermore, the random variable $D(t)$ that counts the number of defaults is conditionally binomially distributed and its unconditional distribution has the form (11). Note that $d(1), \dots, d(s)$ are observations of the random variables $D(1), \dots, D(s)$. Finally, we assume that the vectors of systematic factors $\Psi(1), \dots, \Psi(s)$ as well as the random variables $D(1), \dots, D(s)$ are independent.

We consider two methods for estimating the dependence structure of default indicators $I_1(t), \dots, I_{n(t)}(t)$, $t = 1, \dots, s$, from the time-series $d(1), \dots, d(s)$ of default counts. The method of moments calculates the joint default probabilities π_k specified in (6) within a cohort (*intra-correlations*) and between different cohorts (*inter-correlations*). The maximum likelihood method uses the parametric form of the conditional probability function $\bar{p}$ specified in (12) to calibrate the parameters of the factor model.

4.2 Moment estimator

The application of moment estimators is based on the following generalization of the default count $D(t)$, see McNeil et al. (2005). Define the random variable

$$D^{(k)}(t) := \sum_{\{i_1,\dots,i_k\} \subseteq \{1,\dots,n(t)\}} I_{i_1}(t) \cdots I_{i_k}(t) = \binom{D(t)}{k} \tag{13}$$

representing the number of possible subgroups of k obligors among the defaulting obligors in period t. Obviously, $D^{(1)}(t)$ equals $D(t)$. By taking expectations in (13) we get

$$\mathrm{E}(D^{(k)}(t)) = \binom{n(t)}{k} \pi_k,$$

where π_k specifies the joint default probability of an arbitrarily selected subgroup of k exchangeable obligors, see (6). Hence, the unknown default probability π_k can be easily estimated from the first moment of the generalized default count $D^{(k)}(t)$ by taking the empirical average of s years of data. More precisely, the estimator for π_k is defined by

$$\hat{\pi}_k := \frac{1}{s}\sum_{t=1}^{s} D^{(k)}(t) \Big/ \binom{n(t)}{k} = \frac{1}{s}\sum_{t=1}^{s} \binom{D(t)}{k} \Big/ \binom{n(t)}{k}.$$

For a given time-series of default counts $d(1), \ldots, d(s)$ we obtain

$$\hat{\pi}_k = \frac{1}{s}\sum_{t=1}^{s} \frac{d(t)(d(t)-1)\cdots(d(t)-k+1)}{n(t)(n(t)-1)\cdots(n(t)-k+1)}.$$

For $k = 1, 2$ we get the estimators

$$\hat{\pi}_1 = \frac{1}{s}\sum_{t=1}^{s} \frac{d(t)}{n(t)}, \qquad \hat{\pi}_2 = \frac{1}{s}\sum_{t=1}^{s} \frac{d(t)(d(t)-1)}{n(t)(n(t)-1)}. \tag{14}$$

The default correlation ρ^D can be estimated by (7). The estimator $\hat{\pi}_k$ is unbiased and consistent, see Frey and McNeil (2001). However, de Servigny and Renault (2002) have argued that for low-default cohorts the estimator $\hat{\pi}_2$ might lead to spurious negative correlations and proposed the following (biased) estimator instead

$$\tilde{\pi}_2 = \frac{1}{s}\sum_{t=1}^{s} \frac{d(t)^2}{n(t)^2}. \tag{15}$$

The estimators $\hat{\pi}_2$ and $\tilde{\pi}_2$ are used for calculating the default correlation of two obligors in a homogeneous cohort. The methodology can be easily extended to default correlations of obligors in different cohorts. More precisely, let $n_1(t)$ and $n_2(t)$ be the number of observed companies at the start of period t in cohort 1 and cohort 2, respectively. The corresponding numbers of defaults are denoted by $d_1(t)$ and $d_2(t)$. Assume that obligor i is in the first cohort and obligor j is in the second cohort. Their joint default probability $\mathrm{E}(I_i I_j)$ is estimated by

$$\frac{1}{s}\sum_{t=1}^{s} \frac{d_1(t)d_2(t)}{n_1(t)n_2(t)} \tag{16}$$

and the corresponding default correlation is specified by (5).

4.3 Maximum likelihood estimator

To implement a maximum likelihood procedure we assume a parametric form for the function $\bar{p}_b : \mathbb{R}^m \to [0,1]$ with parameters $b = (b_1, \ldots, b_j)$, where $\bar{p}_b(\psi)$ determines the binomial distribution (10) of the default count $D(t)$ conditionally on $\Psi(t) = \psi$. The likelihood function $L(b \mid d)$ is the product of the unconditional probabilities of the default counts $D(1), \ldots, D(s)$ evaluated at observed defaults $d = (d(1), \ldots, d(s))$:

$$\begin{aligned} L(b \mid d) &= \prod_{t=1}^{s} \mathbb{P}(D(t) = d(t)) \\ &= \prod_{t=1}^{s} \binom{n(t)}{d(t)} \int_{\mathbb{R}^m} \bar{p}_b(\psi)^{d(t)} (1-\bar{p}_b(\psi))^{n(t)-d(t)}\, dF(\psi), \end{aligned} \tag{17}$$

where F is the distribution function of the systematic factors ψ. The parameters $b = (b_1, \ldots, b_j)$ can be calculated by maximizing the likelihood function $L(b \mid d)$ or the log-likelihood function $\log(L(b \mid d))$ with respect to b. The log-likelihood function

$$\log(L(b \mid d)) = \sum_{t=1}^{s} \log(\mathbb{P}(D(t) = d(t))) \tag{18}$$

has the same maxima as $L(b \mid d)$ but its maximization is usually easier from a numerical point of view.

We shall now apply the maximum likelihood estimator to the homogeneous one-factor model (4). More precisely, we assume that the homogeneous default probability p has already been specified. Hence, the default threshold is determined by $c = N^{-1}(p)$. We will apply maximum likelihood estimation to the remaining model parameter, the R^2. It follows from the unconditional distribution (11) of $D(t)$ that the likelihood function has the following form:

$$L(R^2 \mid d) = \prod_{t=1}^{s} \binom{n(t)}{d(t)} \int_{\mathbb{R}} \bar{p}_{R^2}(\psi)^{d(t)} (1 - \bar{p}_{R^2}(\psi))^{n(t)-d(t)} \, dN(\psi), \tag{19}$$

where

$$\bar{p}_{R^2} = N\left(\frac{c - \sqrt{R^2}\psi}{\sqrt{1 - R^2}}\right).$$

The integrals in (19) are easily evaluated numerically.

Again, the methodology can be extended to default correlations of obligors in different cohorts. For example, consider two homogeneous groups of obligors, e.g. cohorts 1 and 2. We assume that all obligors in cohort i depend only on the systematic factor Ψ_i and have the same default threshold c_i and R_i^2. The systematic factors Ψ_1 and Ψ_2 are standardized and follow a bivariate normal distribution N_ρ with correlation ρ.

We assume that observations for s time periods are given: the cohort size and the number of defaults in the first and second cohort are denoted by $n_1(t), n_2(t)$ and $d_1(t), d_2(t)$ respectively, the corresponding count variables are specified by $D_1(t)$ and $D_2(t)$. We apply maximum likelihood estimation to the parameters R_1^2, R_2^2, ρ and obtain the likelihood function

$$\begin{aligned} L(R_1^2, R_2^2, \rho \mid d_1, d_2) &= \prod_{t=1}^{s} \int_{\mathbb{R}^2} \prod_{i=1}^{2} \mathbb{P}(D_i(t) = d_i(t) \mid \Psi_i = \psi_i) \, dN_\rho(\psi_1, \psi_2) \qquad (20) \\ &= \prod_{t=1}^{s} \int_{\mathbb{R}^2} \prod_{i=1}^{2} \binom{n_i(t)}{d_i(t)} \bar{p}_{R_i^2}(\psi_i)^{d_i(t)} (1 - \bar{p}_{R_i^2}(\psi_i))^{n_i(t)-d_i(t)} \, dN_\rho(\psi_1, \psi_2). \end{aligned}$$

5 Model validation based on default data

5.1 Correlations of rating cohorts

We now apply the estimators presented in the previous section to Standard & Poor's data set CreditPro consisting of data from July 1981 to June 2009. Each of the companies included in this data set is rated according to the eight rating categories AAA, AA, A, BBB, BB, B, CCC and Default and classified with respect to 13 industry segments, which are listed in Table 2.

In this subsection, we ignore industry information and partition the set of rated companies into 7 cohorts that correspond to rating classes AAA to CCC. We model each cohort as a homogeneous one-factor model: the ability-to-pay variable of each obligor j in rating class $k = \text{AAA}, \ldots, \text{CCC}$ has the form

$$A_j = \sqrt{R_k^2}\Psi_k + \sqrt{1 - R_k^2}\varepsilon_j, \tag{21}$$

where $\Psi_{\text{AAA}}, \ldots, \Psi_{\text{CCC}}$ are the 7 systematic factors of the model. The default threshold c_k is derived from the average default rate of the companies in rating class k.

Twelve time-series with non-overlapping 1y observation periods are constructed for each rating class, say BB. The time-series cover 27 years, their starting dates are shifted by one month: the first time-series covers July 1981 to July 2008, the second starts in August 1981 and the last time-series covers June 1982 to June 2009. Each of these BB time-series $d(1), \ldots, d(27)$ counts the number of defaults $d(t)$ of those companies that had a BB rating at the beginning of period t and defaulted within the next 12 months. Correlation estimates are derived separately from every BB time-series and then averaged. This procedure is motivated by the observation that the correlation estimates are rather sensitive to the precise bucketing of defaults. For example, correlation estimates derived from data sets July 1981 to July 2008 and January 1982 to January 2009 might differ significantly.

Since the default indicators of each rating cohort form an exchangeable Bernoulli mixture model we can use the moment estimators $\hat{\pi}_2$ and $\tilde{\pi}_2$ defined in (14) and (15) for calculating the joint default probabilities within each rating cohort. The corresponding asset correlations are derived from equality (9). Asset correlations between different cohorts are calculated with estimator (16) and equality (9). Intra-correlations for the rating classes[2] A, BBB, BB, B and CCC as well as the average inter-correlation are given in the first three rows of Table 1.

We now turn to maximum likelihood estimation. Note that each rating cohort specifies a likelihood function of the form (19). The ML estimate of each R_k^2 in (21) is calculated by maximizing this likelihood function.[3] In a second step, the correlation between two systematic factors is estimated by maximizing (20). The results obtained with both ML estimators are shown in rows 4 and 5 of Table 1.

	Intra						Inter
in %	A	BBB	BB	B	CCC	Avg	Avg
Moment estimator (14)	7.8	2.8	7.8	8.9	7.2	6.9	
Moment estimator (15)	13.6	8.2	10.5	9.9	12.3	10.9	
Moment estimator (16)							4.7
ML estimator (19)	11.8	6.2	9.5	11.1	7.9	9.3	
ML estimator (20)							5.9

Table 1: Asset correlations of rating cohorts

In the Basel II Accord (Basel Committee on Banking Supervision, 2006) asset correlations for corporates, sovereigns and banks are modeled as a decreasing function of default probabilities. The correlation estimates in Table 1 do not support a systematic dependence between asset correlations and credit quality. Furthermore, the estimates for intra-correlations are smaller than the Basel II asset correlations for large corporates, which range between 12% and 24%. The difference is even more pronounced for inter-correlations. In the following subsection we will show that these differences can be partly reconciled by taking industry information into account.

Note that the estimates calculated with the unbiased moment estimator (14) are always lower than the ML estimates. The biased moment estimator (15) produces the highest estimates for all but one rating class. These results seem to confirm that the first moment estimator tends to underestimate correlations in small or low-default cohorts whereas the second moment estimator overestimates correlations for these cohorts.[4] These observations are in line with a Monte Carlo study in Gordy and Heitfield (2002), where ML estimators outperform moment estimators. Hence, in the rest of the paper we will focus on maximum likelihood estimation.

[2]Rating classes AAA and AA are not displayed since no reliable estimates of their intra-correlations could be calculated due to the low number of defaults.

[3]As mentioned above, the correlation estimates are very sensitive to the precise bucketing of defaults. For example, the intra-correlation of 9.5% for BB has been obtained as the average of ML estimates that vary between 7.5% (Nov 81 - Nov 08) and 11.9% (Sep 81 - Sep 08).

[4]A comprehensive theoretical analysis of the different estimators will be presented in the PhD thesis of the second author.

5.2 Correlations of industry cohorts

The correlation estimates in Table 1 were calculated without using the S&P industry classification. We now apply the same MLE methodology to 13 cohorts that correspond to the S&P industry segments. The estimates for the 13 intra-correlations are presented in the first column of Table 2. Demey et al. (2004) and Demey and Roncalli (2004) apply the same setup to S&P data covering 22 years from 1981 to 2002. In addition to the one-factor MLE model, they also present an ML estimator based on two systematic factors for each cohort. Both sets of results are given in columns 4 and 5 of Table 2 for comparison.

Note that each of these industry cohorts consists of companies with different ratings. In order to ensure that the rating profiles of the cohorts are homogeneous we partition each industry cohort into 7 subcohorts corresponding to the different rating classes. More precisely, the improved model has the following form:

(i) 91 cohorts are specified by taking all combinations of 7 rating classes and 13 industries.

(ii) The model has 13 systematic factors Ψ_i that correspond to the 13 industries.

(iii) All companies in an industry depend only on the systematic industry factor and have the same R^2. More formally, if a company j is in industry i then the ability-to-pay variable A_j has the form

$$A_j = \sqrt{R_i^2}\Psi_i + \sqrt{1 - R_i^2}\varepsilon_j.$$

(iv) The default threshold of a cohort is specified by the average default rate of the companies in the cohort.

Conditions (ii) and (iii) ensure that only 13 R^2s and 78 correlation parameters have to be estimated. The ML estimate of each R_i^2 is derived from the 7 time-series corresponding to the different rating cohorts of the i-th industry. In order to apply the ML estimator simultaneously to 7 cohorts the maximum likelihood function (19) is generalized to

$$L(R_i^2 \mid d_{i1}, \dots, d_{i7}) = \prod_{t=1}^{s} \int_{\mathbb{R}} \prod_{k=1}^{7} \binom{n_{ik}(t)}{d_{ik}(t)} \bar{p}_{R_i^2}(\psi)^{d_{ik}(t)} (1 - \bar{p}_{R_i^2}(\psi))^{n_{ik}(t) - d_{ik}(t)} \, dN(\psi), \tag{22}$$

where d_{ik} and n_{ik} are the time-series of observed defaults and cohorts size for $i = 1, \dots, 13$ and $k = 1, \dots, 7$. The maximum likelihood function (20) for estimating the correlations of the systematic factors is generalized in the same way. We will use this MLE model in all subsequent calculations.

The second column in Table 2 displays the R^2 estimates obtained by maximizing likelihood function (22). As in all previous calculations, the underlying time-series consist of 27 default counts for 1y observation periods. It is obvious that the defaults are concentrated in the low rating classes since each observation period covers only one year. In contrast, the results in the third column are obtained by applying the ML estimator (22) to time-series with 3y observation periods.

A comparison of the results in Tables 1 and 2 clearly shows that the average intra-correlation of the industry cohorts in Table 2 is significantly higher than the average intra-correlation of the rating cohorts in Table 1. The increased correlations seem to be caused by the higher homogenity of the industry cohorts. This explanation is also supported by a comparison of the first and the second columns in Table 2: for all but one industry segment intra-correlations increased if the estimates are derived from cohorts that correspond to combinations of industries and ratings. The third column shows that the intra-correlation are relatively stable if the same ML estimator is applied to 3y observation periods.

On average, the intra-correlations calculated in Demey et al. (2004) and Demey and Roncalli (2004) are lower, see also Jobst and de Servigny (2005). The main reason is that correlations increased for most industries in the period 2003 to 2009, which is not covered in these papers. In fact, correlations vary significantly over time as illustrated in Table 3. The correlation estimates in Table 3 were computed for time-series of non-overlapping 1y observation periods that cover 18 years. The 18y periods are shifted by

in %	Intra Correlations				
	Industry 1y	Ind / Rat 1y	Ind / Rat 3y	Demey 1F	Demey 2F
Aerospace / Automotive	8.2	8.9	11.5	11.6	11.2
Consumer / Service Sector	6.4	9.7	7.5	7.5	8.7
Energy & Natural Resources	16.5	26.2	23.9	11.5	21.3
Financial Institutions	12.6	15.8	19.9	12.2	15.7
Forest & Building Products	10.9	17.1	10.5	14.5	6.8
Health Care / Chemicals	8.5	11.3	12.3	9.2	8.3
High Technology	14.5	17.7	16.6	4.7	6.8
Insurance	9.2	10.3	11.4	7.6	12.2
Leisure Time / Media	10.8	18.0	20.6	7.0	7.0
Real Estate	44.7	51.0	59.1	27.7	35.9
Telecommunications	34.9	31.1	27.5	34.3	27.1
Transportation	7.8	11.2	9.5	8.3	6.8
Utility	27.1	29.8	26.8	21.2	18.3
Average	16.3	19.8	19.8	13.6	14.3

Table 2: Intra-correlations of industry cohorts

one year, i.e., the first period covers 1982 to 1999, the last period covers 1991 to 2008. The calculations are performed for all industry segments. For example, Table 3 shows that, for High Technology, the initial intra-correlation of 6.3% in the period 1982 - 1999 increases to 17.7% in the last period 1991-2008.

in %	Time period									
	82-99	83-00	84-01	85-02	86-03	87-04	88-05	89-06	90-07	91-08
Aerospace / Automotive	4.5	5.3	10.0	9.3	9.0	8.2	8.2	7.8	9.2	10.2
Consumer / Service Sector	5.2	6.1	7.9	7.5	6.9	6.1	6.8	9.0	10.3	10.2
Energy & Natural Resources	21.2	21.4	20.7	19.0	18.2	6.4	7.2	10.3	14.1	14.4
Financial Institutions	14.1	17.7	14.9	14.3	16.1	19.6	20.8	21.8	17.6	14.2
Forest & Building Products	16.6	15.1	14.5	15.0	14.5	15.1	18.3	19.7	19.4	17.4
Health Care / Chemicals	6.4	9.2	7.8	5.7	4.8	7.4	8.8	12.5	14.7	13.5
High Technology	6.3	6.6	8.5	8.7	9.8	13.9	14.7	13.6	16.4	17.7
Insurance	14.0	13.3	11.3	9.8	8.8	7.2	9.3	9.3	11.5	10.4
Leisure Time / Media	13.4	12.6	10.9	8.9	10.3	14.0	14.9	17.4	17.8	13.5
Real Estate	44.9	52.6	54.0	54.1	57.4	60.1	62.3	66.4	66.3	51.5
Telecommunications	11.7	7.6	19.7	22.9	23.7	25.6	29.4	30.0	32.0	38.3
Transportation	9.4	9.6	8.5	10.4	10.6	12.2	10.6	10.2	13.2	12.1
Utility	42.8	30.2	18.9	25.6	22.3	25.0	24.6	27.0	28.8	29.8

Table 3: Variation of intra-correlations over time

In addition to intra-correlations, we calculated inter-correlations for each pair of the 13 industries using the same model as in the second column of Table 2. The average inter-correlation is 10.1%, which is approximately half the size of the average intra-correlation of 19.8%.

6 Validation based on rating data

6.1 Credit portfolio model for default and migration risk

The portfolio model defined in Section 2.1 distinguishes only between two states: default and non-default. We shall now extend the model to cover rating migrations, which provides the appropriate framework for deriving correlation estimates from time-series of rating data. More precisely, instead of focusing on defaults within a given cohort we shall now consider time-series that specify the entire rating distribution of a cohort at the end of each observation period.

Let r be the number of rating classes including default and define thresholds

$$-\infty = c_{r,j} \leq c_{r-1,j} \leq \ldots \leq c_{1,j} \leq c_{0,j} = \infty, \qquad j = 1, \ldots, n$$

for each of the n obligors. The event $\{c_{k,j} < A_j \leq c_{k-1,j}\}$ is interpreted as the j-th obligor being in rating class k at the end of the time period $[0, T]$. The corresponding migration probabilities

$$p_{k,j} := \mathbb{P}(c_{k,j} < A_j \leq c_{k-1,j}) \tag{23}$$

are usually taken from a rating migration matrix. Note that the default probability is defined by $p_{r,j} = \mathbb{P}(A_j \leq c_{r-1,j})$. The vector $(l_{1,j}, \ldots, l_{r,j}) \in \mathbb{R}^r$ with $l_{1,j} \leq \ldots \leq l_{r,j}$ specifies the loss $l_{k,j}$ if the j-th obligor migrates to rating class k.[5] Hence, the loss variable $L_j : \Omega \to \mathbb{R}$ of the j-th obligor is generalized to

$$L_j := \sum_{k=1}^{r} l_{k,j} \cdot \mathbf{1}_{\{c_{k,j} < A_j \leq c_{k-1,j}\}}.$$

The definition of portfolio loss remains unchanged, i.e., $L = \sum_{j=1}^{n} L_j$. Note that if $r = 2$, then the multi-state model is just the two-state model from Section 2.1.

6.2 Counting rating migrations

The objective of this subsection is the generalization of the one-dimensional random variable D that counts defaults to an r-dimensional variable M that counts rating migrations. First of all, we define for $k = 1, \ldots, r$ and $j = 1, \ldots, n$ a rating indicator

$$I_{k,j} := \mathbf{1}_{\{c_{k,j} < A_j \leq c_{k-1,j}\}}. \tag{24}$$

Note that $I_{r,j}$ is the default indicator. Let $M_k := \sum_{j=1}^{n} I_{k,j}$ specify the number of obligors in rating class k at time T.

For the rest of this subsection, we assume that the n-dimensional vector

$$((I_{1,1}, \ldots, I_{r,1}), \ldots, (I_{1,n}, \ldots, I_{r,n})) \tag{25}$$

is exchangeable. As a consequence, each obligor has the same vector of migration thresholds $(c_0, \ldots, c_r)$ and there exist functions $\bar{p}_k : \mathbb{R}^m \to [0, 1]$, $1 \leq k \leq r$, which specify the conditional migration probabilities for all obligors:

$$\bar{p}_k(\psi) = \mathbb{P}(c_{k,j} < A_j \leq c_{k-1,j} \mid \Psi = \psi)$$

for all $\psi \in \mathbb{R}^m$ and $j = 1, \ldots, n$.

We will now compute conditional and unconditional distributions of $M = (M_1, \ldots, M_r)$. Conditional on $\Psi = \psi$, the probability that M equals $j = (j_1, \ldots, j_r)$ with $\sum_{k=1}^{r} j_k = n$ is given by

$$\mathbb{P}(M = j \mid \Psi = \psi) = \frac{n!}{\prod_{k=1}^{r} j_k!} \prod_{k=1}^{r} \bar{p}_k(\psi)^{j_k}.$$

[5] Note that $l_{k,j}$ can also be negative, i.e., $l_{k,j}$ represents a gain.

The unconditional distribution of M is obtained by integrating over the distribution F of the systematic factors Ψ:

$$\mathbb{P}(M = j) \;=\; \int_{\mathbb{R}^m} \mathbb{P}(M = j \mid \Psi = \psi)\, dF(\psi) \;=\; \frac{n!}{\prod_{k=1}^{r} j_k!} \int_{\mathbb{R}^m} \prod_{k=1}^{r} \bar{p}_k(\psi)^{j_k}\, dF(\psi).$$

As in the two-state case we consider the example of a homogeneous one-factor model: assume that all obligors have the same migration probabilities $p_1, \ldots, p_r$ with corresponding thresholds $c_0, \ldots, c_r$ specified by (23) and the same R^2. For this model class the vector of rating indicators (25) is exchangeable and the unconditional distribution of M equals

$$\begin{aligned}
\mathbb{P}(M = j) &= \frac{n!}{\prod_{k=1}^{r} j_k!} \int_{\mathbb{R}} \prod_{k=1}^{r} \bar{p}_k(\psi)^{j_k}\; dN(\psi) \\
&= \frac{1}{(2\pi)^{1/2}} \frac{n!}{\prod_{k=1}^{r} j_k!} \int_{\mathbb{R}} \prod_{k=1}^{r} \bar{p}_k(\psi)^{j_k} e^{-\psi^2/2} d\psi, \qquad (26)
\end{aligned}$$

where $\bar{p}_k$ is specified by

$$\bar{p}_k(\psi) = N\left(\frac{c_{k-1} - \sqrt{R^2}\psi}{\sqrt{1 - R^2}}\right) - N\left(\frac{c_k - \sqrt{R^2}\psi}{\sqrt{1 - R^2}}\right). \qquad (27)$$

6.3 Maximum likelihood estimator for rating data

In Section 3.3, moment estimators and MLE techniques were presented to derive parameters of the factor model from default data. We shall now generalize maximum likelihood estimation to rating data. The replacement of default data by rating data significantly increases the amount of information available for calibration of the model: we now infer dependence not only from joint defaults but also from rating upgrades and downgrades. This is an important improvement, particularly for good rating classes, which are characterized by a low number of defaults.

The setup is similar to that of Section 3.3. In each time period $t = 1, \ldots, s$ we group the monitored companies into cohorts. Let $n(t)$ denote the number of obligors in a given cohort at the start of period t and $m(t) = (m_1(t), \ldots, m_r(t))$ the rating vector of the cohort. More precisely, $m_k(t)$ specifies the number of obligors in rating class k at the end of period t. The companies within a cohort are assumed to be homogeneous, i.e., conditionally on systematic factors their vectors of rating indicators $(I_{1,1}, \ldots, I_{r,1}), \ldots, (I_{1,n}, \ldots, I_{r,n})$ are independent and identically distributed.

The specification of the likelihood function follows a similar approach as in Section 4.3. The main difference is that the variables $D(t)$, which count defaults, are replaced by the more general variables $M(t) = (M_1(t), \ldots, M_r(t))$, which specify the number of companies in the different rating classes at the end of period t.

We shall now apply the maximum likelihood estimator to the homogeneous one-factor model. Assuming that the homogeneous migration probabilities $p_1, \ldots, p_r$ with corresponding thresholds $c_0, \ldots, c_r$ have already been specified it remains to apply maximum likelihood estimation to the homogeneous R^2. It follows from the unconditional distribution (26) of $M(t)$ that the likelihood function has the following form:

$$L(R^2 \mid m) = \prod_{t=1}^{s} \frac{n(t)!}{\prod_{k=1}^{r} m_k(t)!} \int_{\mathbb{R}} \prod_{k=1}^{r} \bar{p}_{k,R^2}(\psi)^{m_k(t)}\, dN(\psi), \qquad (28)$$

where $\bar{p}_{k,R^2}$ is specified by

$$\bar{p}_{k,R^2}(\psi) = N\left(\frac{c_{k-1} - \sqrt{R^2}\psi}{\sqrt{1 - R^2}}\right) - N\left(\frac{c_k - \sqrt{R^2}\psi}{\sqrt{1 - R^2}}\right). \qquad (29)$$

Finally we derive the likelihood function for the model consisting of two homogeneous cohorts that was specified in Section 4.3. We assume that we have observations for s time periods: the rating vector for cohort

$i = 1, 2$ and time period $t = 1, \ldots, s$ is denoted by $m_i(t) = (m_{1i}(t), \ldots, m_{ri}(t))$. The likelihood function $L(R_1^2, R_2^2, \rho \mid m_1, m_2)$ for the parameters R_1^2, R_2^2, ρ equals

$$\prod_{t=1}^{s} \int_{\mathbb{R}^2} \prod_{i=1}^{2} \frac{n_i(t)!}{\prod_{k=1}^{r} m_{k,i}(t)!} \prod_{k=1}^{r} \bar{p}_{k,R_i^2}(\psi_i)^{m_{k,i}(t)} \, dN_\rho(\psi_1, \psi_2),$$

where $\bar{p}_{k,R_i^2}$ is specified by (29).

6.4 Model validation with rating data

We repeat the calculations presented in the second column of Table 2 but replace the time-series of defaults by time-series of rating migration vectors. The ML estimation is based on the likelihood function (28). The results are displayed in the first column of Table 4. The second column is given for comparison and presents the corresponding estimates derived from default data (column 2, Table 2).

in %	Intra Correlations	
	Rating	Default
Aerospace / Automotive	5.5	8.9
Consumer / Service Sector	3.0	9.7
Energy & Natural Resources	11.4	26.2
Financial Institutions	13.4	15.8
Forest & Building Products	5.9	17.1
Health Care / Chemicals	5.0	11.3
High Technology	5.4	17.7
Insurance	13.9	10.3
Leisure Time / Media	9.8	18.0
Real Estate	29.0	51.0
Telecommunications	10.3	31.1
Transportation	7.6	11.2
Utility	6.1	29.8
Average	9.7	19.8

Table 4: Intra correlations using rating data

The intra-correlations derived from the rating time-series are lower for all but one industry. The average intra-correlation is reduced by more than 50% from 19.8% to 9.7%. Inter-correlations show the same behavior: the average inter-correlation is reduced from 10.1% to 4.5%. These estimates are in line with the results in the literature, see for example the correlation estimates obtained with the Directional Rating Transition Matrix in Akhavein et al. (2005). We shall discuss potential explanations for the striking difference between correlations obtained from rating data and correlations obtained from default data in Section 7.2.

7 Model extensions

7.1 Relaxing the normal distribution assumption

In all of the previous calculations we have assumed that the risk factors and, consequently, the ability-to-pay variables follow a multivariate normal distribution. In the following, we relax the normal distribution assumption to allow for fat tails and skews. More precisely, we calculate estimates for intra-correlations under three different distribution assumptions for each of the 13 industry cohorts (with 7 rating subcohorts):

(i) The systematic factor follows a t-distribution, the idiosyncratic factors are normally distributed.

(ii) The systematic factor and the idiosyncratic factors follow a multivariate t-distribution.[6]

(iii) The systematic factor follows a skew normal distribution, the idiosyncratic factors are normally distributed.[7]

Compared to the Gaussian model, the additional parameter in each of these models improves the fit to the data. Overall, the best calibration results are obtained in model (ii) or, more precisely, in a model that follows a multivariate t-distribution with a high degree of freedom. Table 5 illustrates this result for rating time-series. The first row shows the average intra-correlations (R^2s) for selected degrees of freedom. The last column labeled by ∞ corresponds to the normally distributed model. The average intra-correlation of 9.7% is consistent with the average correlation shown in Table 4. Note that intra-correlations decrease for lower degrees of freedom, which is counterbalanced by higher tail dependencies in these models. The second row in Table 5 displays the average value of the log-likelihood function at the 13 R^2s for each of the selected degrees of freedom. The maximum is obtained for a degree of freedom that is close to 65. A similar result holds for default data.

Degree of freedom	4	10	30	50	60	65	70	90	∞
Intra-correlation	4.8	5.2	6.1	6.7	6.9	7.0	7.0	7.3	9.7
Log-likelihood	-847.71	-733.35	-712.80	-710.89	-710.73	-710.72	-710.76	-711.11	-726.15

Table 5: Variation of intra-correlation (in %) and log-likelihood function with degrees of freedom

Further analysis is required to decide whether it is justified to replace the Gaussian distribution by a more complex distribution class.[8]

7.2 Modeling the credit cycle

Rating systems are usually characterized as point-in-time or through-the-cycle, see Basel Committee on Banking Supervision (2000). In a point-in-time process, a rating reflects an assessment of the borrower's current condition and most likely future condition over the course of the chosen time horizon, typically one year. As a consequence, the point-in-time rating changes as the borrower's condition changes over the course of the credit cycle. In contrast, a through-the-cycle process requires assessment of the borrower's riskiness based on a bottom-of-the-cycle scenario. In this case, a borrower's rating would tend to stay the same over the course of the credit cycle. Ratings from external rating agencies are typically characterized as through-the-cycle. It is obvious, however, that the S&P rating system is a combination of both rating philosophies, see Cornaglia and Morone (2009). This is highlighted by the facts that the proportion of downgrades to upgrades as well as the actual default ratios of the different rating classes are strongly correlated with the credit cycle.

In our analysis we have ignored these characteristics of S&P data. In fact, we made the assumptions that risk factors do not exhibit serial dependence over time and default probabilities of rating classes are constant over time. The application of this oversimplifying estimation model to rating data with rather complex characteristics might explain the significant differences between estimates derived from default data and estimates derived from rating data.

[6] Note that in this setup the idiosyncratic factors are no longer independent, see McNeil et al. (2005).

[7] The skew-normal distribution is an extension of the normal distribution allowing for the presence of skewness, see Arellano-Valle and Azzalini (2006) for more information.

[8] Hamerle and Rösch (2005) analyze the impact of different distribution assumptions for the risk factors on the loss distribution of the credit portfolio model. They come to the conclusion that even if the correct distribution is specified by a t-copula, a misspecified Gaussian model does not necessarily underestimate risk.

The challenge of having to cope with different rating philosophies has already been addressed in Hamerle et al. (2003) and Rösch (2005). They also use S&P default data but mimic a point-in-time rating by including information about the credit cycle. As a consequence, correlation estimates are significantly reduced.

McNeil and Wendin (2006, 2007) propose Generalized Linear Mixed Models (GLMM) as framework for modeling dependent defaults and rating migration. The GLMM-setting allows for the inclusion of observable factors (e.g. macroeconomic indicators as proxies for the credit cycle) and serial dependence of latent risk factors. A drawback of this technique is that serially correlated latent risk factors yield joint migration distributions in terms of high-dimensional integrals, which are difficult to calculate with standard ML techniques. Alternative calibration methodologies are simulated ML (Koopman et al., 2005) and Bayesian techniques. The latter approach is used in McNeil and Wendin (2006, 2007), who report intra-correlations (inter-correlations) of 10.9% (6.5%) for S&P default data and 8.7% (2.6%) for rating data.

8 Conclusion

Let us return to our initial question and compare the correlation estimates derived from historical default and rating data to correlations of equity time-series. Akhavein et al. (2005) apply Fitch's Vector Model 2.0, which is an equity-based approach, and obtain an average intra-correlation of 24.1% and an inter-correlation of 20.9% for the 25 Fitch industry segments. We used a different set of equity time-series covering a 10y period from 1998 to 2007. Based on the S&P industry classification we computed rather homogeneous intra-correlations for the 13 industry segments with an average value of 24.7%. The average inter-correlation is 20.4%. These equity correlations are higher than all the estimates derived from historical default and rating data that were presented in this paper. The difference is particularly pronounced for correlations between industry segments. Hence, we do not see any indication that credit portfolio models underestimate joint default probabilities if calibrated with equity data. On the contrary, there seems to be the danger that equity correlations might underestimate diversification effects between different industries.

The difference between intra-correlations and inter-correlations also has an important consequence for the design of a credit portfolio model. Portfolio models that use a single systematic factor across all industries do not provide sufficient flexibility to capture the complex dependence structure exhibited in rating data. This fact clearly highlights an important weakness of the Basel II framework for credit risk.

Historical ratings are an indispensable data source for calibrating and validating credit portfolio models. However, our analysis provides another confirmation that it is a notoriously difficult task to calibrate a realistic dependence structure for joint defaults and rating migrations. The impact of the credit cycle together with the complex characteristics of the historical rating data make it difficult to specify an appropriate estimation model. The significant differences in the correlation estimates obtained from default and rating data further highlight this difficulty. Another problem is the volatility observed in correlations: the estimates strongly vary across industries as well as over time, which makes it difficult to develop reliable estimation procedures. More work is needed to better understand the dynamics of correlation changes, in particular under stressed market conditions.

Acknowledgement

Financial support from the EU Commission through MRTN-CT-2006-034270 COMISEF is gratefully acknowledged.

REFERENCES

Akhavein, J.D., Kocagil, A.E. and Neugebauer, M. (2005). A Comparative Empirical Study of Asset Correlations. Fitch Ratings, Quantitative Financial Research Special Report.

Arellano-Valle, R.B. and Azzalini, A. (2006). On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics* 33, 561-574.

Basel Committee on Banking Supervision (2000). Range of Practice on Banks' Internal Ratings Systems. Bank for International Settlements, January 2000.

Basel Committee on Banking Supervision (2006). International Convergence of Capital Measurement and Capital Standards: A Revised Framework – Comprehensive Version. Bank for International Settlements, June 2006.

Bahar, R. and Nagpal, K. (2001). Measuring default correlation. *Risk* 14(3), 129-132.

Bluhm, C., Overbeck, L. and Wagner, C. (2002). *An Introduction to Credit Risk Modeling.* CRC Press/Chapman & Hall.

Cornaglia, A. and Morone, M. (2009). Rating philosophy and dynamic properties of internal rating systems: a general framework and an application to backtesting. Working Paper, Intesa Sanpaolo.

Crouhy, M., Galai D. and Mark, R. (2000). A comparative analysis of current credit risk models. *Journal of Banking and Finance* 24, 59-117.

Demey, P., Jouanin, J.F., Roget, C. and Roncalli, T. (2004). Maximum likelihood estimate of default correlations. *Risk* 17(11), 104-108.

Demey, P. and Roncalli, T. (2004). A correction note on "Maximum likelihood estimate of default correlations". Working Paper, Credit Agricole.

De Servigny, A. and Renault, O. (2003). Correlation evidence. *Risk* 16(7), 90-94.

Frey, R. and McNeil, A.J. (2001). Modelling dependent defaults. Working Paper, Department of Mathematics, ETH Zürich.

Frey, R. and McNeil, A.J. (2003). Dependence modeling, model risk and model calibration in models of portfolio credit risk. *Journal of Risk* 6, 59-92.

Gordy, M. and Heitfield, E. (2002). Estimating default correlations from short panels of credit rating performance data. Technical report, Federal Reserve Board.

Gupton, G., Finger, C. and Bhatia, M. (1997). CreditMetrics Technical Document. J.P. Morgan, New York.

Hamerle, A., Liebig, T. and Rösch, D. (2003). Benchmarking Asset Correlations, *Risk* 16(11), 77-81.

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

Jobst, N.J. and De Servigny, A. (2005). An empirical analysis of equity default swaps II: multivariate insights. Working Paper, S&P's Risk Solutions.

Koopman, S.J., Lucas, A., and Daniels, R.J. (2005). A non-Gaussian panel time series model for estimating and decomposing default risk. Working paper, Tinbergen Institute.

Lucas, D. (1995). Default correlations and credit analysis. *Journal of Fixed Income* 9(4), 76-87.

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

McNeil, A.J. and Wendin, J.P. (2007). Bayesian inference for generalized linear mixed models of portfolio credit risk. *Journal of Empirical Finance* 14, 131-149.

McNeil, A.J. and Wendin, J.P. (2007). Dependent credit migrations. *Journal of Credit Risk* 2(3), 87-114.

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

Rösch, D. (2005). An empirical comparison of default risk forecasts from alternative credit rating philosophies. *International Journal of Forecasting* 21(1), 37-51.

Zeng, B. and Zhang, J. (2002). Measuring credit correlations: equity correlations are not enough!. Working Paper, KMV corporation.