VU Research Portal

An analytic approach to credit risk of large corporate bond and loan portfolios

Lucas, A.; Klaassen, P.; Spreij, P.J.C.; Straetmans, S.T.M.

1999

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Lucas, A., Klaassen, P., Spreij, P. J. C., & Straetmans, S. T. M. (1999). *An analytic approach to credit risk of large corporate bond and loan portfolios*. (Research Memorandum; No. 1999-18). FEWEB.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

SERIE RESEARCH MEMORANDA

An Analytic Approach to Credit Risk of Large Corporate Bond and Loan Portfolios

André Lucas
Pieter Klaassen
Peter Spreij
Stefan Straetmans

Research Memorandum 1999-I 8

February 1999

vrije Universiteit *amsterdam*

An Analytic Approach to Credit Risk of Large Corporate Bond and Loan Portfolios*

André Lucas^{‡§} Pieter Klaassen^{†‡}

Peter Spreij[†] Stefan Straetmans[‡]

This version: February 22, 1999

*We thank Laurens de Haan, Bernard Harmon, Herbert Rijken, and Ronald van Dijk for useful comments. Andre Lucas also thanks the Dutch Organization for Scientific Research (N.W.O.) for financial support. Correspondence to: alucas@econ.vu.nl, Pieter.Klaassen@nl.abnamro.com, spreij@wins.uva.nl, or sstraetmans@econ.vu.nl.

[‡]Dept. Finance and Financial Sector Management, Vrije Universiteit, De Boelelaan 1105, NL-1081HV Amsterdam, the Netherlands

[§]Tinbergen Institute Amsterdam, Keizersgracht 482, NL-1017EG Amsterdam, the Netherlands

[†]ABN-AMRO Bank NV, Financial Markets Risk Management, P.O.Box 283, NL-1000EA Amsterdam, the Netherlands

[†]Korteweg-de Vries Institute, University of Amsterdam, Plantage Muidersgracht 24, NL-1018TV Amsterdam, the Netherlands

An Analytic Approach to Credit Risk of Large Corporate Bond and Loan Portfolios

Abstract

We consider portfolio credit loss distributions based on a factor model for individual exposures and establish an analytic characterization of the credit loss distribution if the number of exposures tends to infinity. Using this limiting distribution, we explain how skewness and leptokurtosis of credit loss distributions relate to the underlying factor model and the portfolio composition. A key role is played by the R^2 of the factor model regression. Based on the limiting distribution and empirical data, it appears that the Basle 8% rule is not an unreasonable approximation for high confidence (99.9%) quantiles of credit losses of a typical portfolio of rated corporate bonds. The practical relevance of our results for credit risk management is investigated by checking the applicability of the limiting distribution to portfolios with a finite number of exposures. It appears that for relatively homogeneous portfolios a minimum of 300 exposures is enough, while for relatively heterogeneous portfolios a number of 800 exposures suffices to obtain an adequate approximation. Thus, our approach can be a fast and accurate alternative to the standard Monte-Carlo simulation approach adopted in much of the literature and in practice.

Key words: credit risk; factor model; fat-tailed distributions; skewness; asymptotic analysis.

JEL Codes: G21; G33; G29; C19.

1 Introduction

Increasingly, banks are using portfolio models to quantify the aggregate credit risk they are exposed to through their loan and trading books. These models generate the distribution of potential losses due to credit risk, as well as some summary statistics like standard deviations and percentiles. Loss distributions are used by banks internally to measure the profitability of (subsets of) transactions in relation to the risk they contribute to the portfolio. This information can result in either laying off certain exposures, for example through securitization, or taking on additional exposures. Additionally, the loss distribution can be used to determine the level of capital that the bank

needs in order to protect itself (with a certain level of confidence) against unexpected credit risk losses. Similarly, it is possible to use the portfolio models to analyze portfolios of assets to be securitized.

The increased use of credit risk portfolio models by financial intermediaries potentially has a significant impact on the pricing of credit-risky instruments in financial markets. A parallel may be drawn with the relationship between equity returns and compensation for systematic risk, as established by the Modern Portfolio Theory of Markowitz (1952) and Sharpe's Capital Asset Pricing Model. One can also envisage far-reaching implications of this development for the capital adequacy regime to which banks are subjected. Since the introduction of the Basle Accord in 1988, see Basle Committee on Bank Supervision (1988), capital charges are determined for individual assets. These charges are summed to arrive at the capital required for a bank. The current rules ignore portfolio effects by levying the same capital charge for corporate debtors of varying creditworthiness. As a result, banks have become actively engaged in 'regulatory arbitrage' transactions. These transactions reduce the regulatory capital charge without decreasing the credit risk exposure proportionally. This undermines the effectiveness of the capital adequacy regime. The shortcomings of the current regime also distort price signals in the market; see ISDA (1998) and IIF (1998) for an overview of shortcomings of the current regime.

The general characteristics of the credit risk loss distribution resulting from portfolio models are badly understood. It is often observed that the loss distribution exhibits significant skewness and leptokurtosis, but the prominence of these properties very much depends on the composition of the specific portfolio under consideration. In this paper we derive an efficient analytic approximation to the loss distribution if the portfolio contains a large number of exposures. Our approximation enables us to study the sensitivity of the loss distribution, and in particular the shape of its tails, to specific portfolio characteristics. These include its overall credit quality, the degree of systematic risk, and the maturity profile. It is shown for portfolios with realistic complexity that the approximation is reasonable already for portfolios with 300 to 800 exposures. As compared to using fully-fledged Monte Carlo simulation to generate the loss distribution, which is commonly done in practice, our approximation can be evaluated much more efficiently in most practical instances.

Numerical results in this paper indicate that the shape of the loss distribution is particularly sensitive to the initial credit quality of the portfolio and the extent to which credit events occur simultaneously for different debtors. The credit quality of a portfolio is often measured by assigning ratings to each of the debtors for which exposures are present in the portfolio. These

ratings can come from external rating agencies such as Standard & Poor's and Moody's, as is the case in *CreditMetrics* of J.P. Morgan (1999). Alternatively, the ratings can be assigned by banks internally. Each debtor's rating is associated with a certain probability of default. A different route to estimate the default probabilities is based on the option-theoretic approach pioneered by Merton (1974), and later extended by Black and Cox (1976) and Longstaff and Schwartz (1995). In this approach, the equity of a company is viewed as an option on its assets with the strike price equal to the level of liabilities. The portfolio model of KMV combines this approach with historical default statistics to assign a default probability to each debtor individually, provided it has equity listed on a stock exchange, Kealhofer (1995).

Correlation between credit events of different debtors is induced by the fact that their well-being is influenced by the same (economic) factors. A positive correlation is induced by the fact that default rates are significantly higher in economic recessions than in periods of economic growth, see also Jónsson and Fridson (1996) and Fons (1991). The more a portfolio of exposures is diversified over different countries and industries, the smaller the 'average' correlation will be in the portfolio. We show that this decreases the likelihood of extremal portfolio credit losses. The correlation effect on the shape of the loss distribution also depends on the initial credit quality of the portfolio. As Zhou (1997) shows, for a given correlation between the asset values of two companies, the correlation between default events of both companies is higher when the creditworthiness of both is lower. Hence, correlation has a larger impact on the tails of the distribution when the credit quality of the portfolio is lower. This is confirmed by our numerical results, and in line with the analysis of Carey (1998).

To our knowledge, only Carey (1998) has thus far performed a systematic study of the tails of the credit loss distribution. In his study, he uses historical data on exposures and credit losses stemming from private placements by US life insurers. Carey samples exposures from this large database to obtain portfolios with certain characteristics in terms of initial credit quality. He then analyzes the actual loss experience from these sampled portfolios. His study yields insights into the effect of credit quality, the size of the portfolio, and the state of the economy on the tails of the loss distribution. His conclusions, however, are only valid as far as the exposures in the database, and their aggregation into portfolios, are representative of actual portfolios.

We describe our analysis from the perspective of portfolios with corporate bonds and loans. This is also the perspective taken in the credit risk portfolio models that are available in the market, such as *CreditMetrics* of J.P. Morgan (1999), *CreditRisk+* of Credit Suisse (1999), *PortfolioManager* of KMV (Kealhofer (1995)), and *CreditPortfolio View* of McKinsey (Wilson

(1997a,b)). Although the approach in each of these models appears quite different at first sight, Koyluoglu and Hickman (1998) have outlined the underlying unifying framework. We follow their framework in the set-up of our model.

This set-up is described in Section 2. In Section 3 we derive the asymptotic loss distribution, and point out the salient features of the obtained expression. Section 4 investigates the properties of the tails of the distribution.

In Sections 5 and 6 we study the properties of the asymptotic loss distribution for a large number of portfolios with different characteristics. Section 5 considers stylized portfolios that differ in initial credit quality, the degree of inherent systematic risk, and the maturity of the exposures. The analysis in this section shows the sensitivity of the loss distribution to each of its parameters. The composition of the portfolios in Section 6 more closely approximates actual corporate bond and loan portfolios in two respects: (i) the distribution of exposures over initial credit ratings, and (ii) the level and variability in systematic risk across exposures. For a typical corporate bond portfolio we find that the standard 8% capital charge from the Basle 1988 Accord roughly corresponds to a confidence level of about 99.9%.

Section 7 investigates how large the number of exposures in a portfolio needs to be to render the asymptotic loss distribution derived in the paper a good approximation to the actual loss distribution. In studying the convergence properties, we especially pay attention to the tail behaviour of the distribution. For relatively homogeneous exposures it is shown that the approximation is already quite accurate for portfolios with a few hundred exposures. If a portfolio contains relatively heterogeneous exposures, we find that approximately 800 exposures suffice to get a close fit with the actual loss distribution. Section 8 concludes, while the Appendix gathers the proofs.

2 Theoretical framework

We consider a portfolio containing n exposures. Each exposure j is characterized by a four-dimensional stochastic vector

$$(S_j, k_j, \ell_j, \pi(j, k_j, \ell_j, \psi)). \quad (1)$$

The first element of this vector triggers the mechanism for defaults and rating migrations. A prime candidate for S_j is the company's surplus, i.e., the difference between the market values of liabilities and assets. If this surplus falls below a certain threshold, default occurs. We assume that the portfolio

exposures are driven by a vector of common factors:

$$\mathbf{Sj} = \mu_j + \beta_j^\top f + \varepsilon_j, \quad (2)$$

where $\mu_j \in \mathbb{R}$ is a constant term, $\beta_j \in \mathbb{R}^m$ is a vector of factor loadings, $f \in \mathbb{R}^m$ is the vector of common factors, and $\varepsilon_j \in \mathbb{R}$ is a scalar representing idiosyncratic risk. This set-up follows the model structure of, e.g., J.P. Morgan (1999). For expositional purposes, we set $\mu_j = 0$ for all j . All results remain valid, however, for non-zero intercept terms, see also further below. Furthermore, we assume that

$$f \sim N(0, \Omega_f), \quad (3)$$

and

$$\varepsilon_j \sim N(0, \omega_j), \quad (4)$$

with $E(\varepsilon_j f) = 0$ for all j , Ω_f positive definite and $E(\varepsilon_i \varepsilon_j) = 0$ for all $i \neq j$. The normality assumption for the factors and the idiosyncratic shocks considerably simplifies the proof of our main theorem in Section 3. The proof can, however, easily be generalized to the case of heavy tails and volatility clustering in the companies' surplus variables S_j .

If the factor structure in (2) holds, the surplus variables of different firms are correlated. Because the S_j 's also trigger the default mechanism, correlation between the S_j 's results in correlated default probabilities. It turns out that this correlation causes the portfolio credit loss distribution to exhibit heavy tails when the number of exposures becomes large, see Sections 3 and 4. It is important to stress that the mere correlation between the credit exposures suffices to induce heavy tails of the credit loss distribution, even if the underlying stochastic variables f and ε_i are thin-tailed, e.g., normal.

The second and third element in (1), k_j and ℓ_j , represent the exposure's initial and its end-of-period rating category, respectively. We assume r rating categories, such that $k_j, \ell_j \in \{1, \dots, r\}$. In this paper we work within a static, one-period framework. We can therefore assume that the migrations are driven by a Markovian transition matrix P ,

$$P = \begin{pmatrix} p_{11} & \dots & p_{1r} \\ \vdots & \ddots & \vdots \\ p_{r1} & \dots & p_{rr} \end{pmatrix}, \quad (5)$$

where $p_{k\ell}$ denotes the probability that a firm with initial rating k switches to rating ℓ over the period considered. Note that $P\iota_r = \iota_r$, where ι_r is an r -dimensional vector with ones. By setting $p_{r1} = \dots = p_{r,r-1} = 0$ and

Table 1: Rating Migration Probability Matrix and Credit Spreads

k	ℓ									1Y	5Y	10Y
	A A A	A A	A	BBB	BB	B	C C C	D				
A A A	9082	826	74	6	11	0	0	0	58	77	125	
A A	65	9088	769	58	5	13	2	0	62	87	145	
A	8	242	9130	523	68	23	1	5	74	102	160	
BBB	3	31	587	8746	496	108	12	17	87	116	180	
BB	2	12	64	771	8116	840	98	98	175	210	350	
B	0	10	24	45	686	8350	392	492	278	475	630	
c c c	21	0	41	124	267	1170	6448	1929	435	585	980	
D	0	0	0	0	0	0	0	10000	—	—	—	
Base yield									429	409	429	

The table contains the probability (in basis points) of a credit rating migration from category k to ℓ over a 1-year period. The category D stands for default. The last three columns of the table contain the credit spreads (in basis points) for firms with initial rating k corresponding to a bond with a maturity of 1, 5, or 10 years. The base yields are also in basis points and imply a U-shaped yield curve. Source: **CreditMetrics**' web site, October 1998.

$p_{rr} = 1$, we can identify the r th rating category as the state of default. J.P. Morgan (1999) presents a transition matrix of the form (5) on its Web site. This transition matrix is based on S&P or Moody's rating categories and an example is presented in Table 1.

For given values of $p_{k\ell}$, one can select constants $c_{k\ell}$, $k = 1, \dots, r$ and $\ell = 0, \dots, r$, such that $c_{k0} = -\infty$ and $c_{kr} = +\infty$ for all k , and

$$\Phi(c_{k\ell}) - \Phi(c_{k,\ell-1}) = p_{k\ell}, \quad (6)$$

for all k and $\ell = 1, \dots, r$, where $\Phi(\cdot)$ is the standard normal cumulative distribution function (c.d.f.). The end-of-period rating of exposure j (with initial rating k_j) is set equal to ℓ_j , where ℓ_j is such that

$$c_{j,k_j,\ell_j-1} \equiv c_{k_j,\ell_j-1} + \sqrt{\omega_j + \beta_j^\top \Omega_f \beta_j} < S_j \leq c_{k_j,\ell_j} + \sqrt{\omega_j + \beta_j^\top \Omega_f \beta_j} \equiv c_{j,k_j,\ell_j} \quad (7)$$

This is illustrated in Figure 1. The support of the normal distribution of S_j is partitioned by means of the constants $c_{k\ell}$ and the standard deviation of S_j . Each bin corresponds to a specific end-of-period rating. Note that the locations of the bins depend on the exposure's initial rating, as for example highly rated exposures are less likely to default than low-graded ones.

As we assumed $E(f) = 0$ and $\mu_j = 0$ in (2), the unconditional distribution of S_j plotted in Figure 1 is centered around 0. Nothing material changes if we relax either of these assumptions. In our present static, one-period set-

Figure 1: Relation of the random variable S_j and the end-of-period rating ℓ_j in an 8 category rating system. The initial rating of the exposure is $k_j = \text{BBB}$. D denotes default, while $c_{jk\ell}$ equals $c_{k\ell}(\omega_j + \beta_j^\top \Omega_f \beta_j)^{1/2}$. Data are as presented in Table 1.

up the default probabilities are always non-stochastic. By an appropriate choice of the matrix P the default and rating migration mechanisms are dictated by the factor model (2) and the bins (7). If we extend our model to a dynamic setting, the default (and rating migration) probabilities become stochastic. For example, assume that the vector of common factors follows an autoregressive process of order 1,

$$f_t = \Psi f_{t-1} + \eta_t,$$

with η_t independent of ε_s for all s, t . The bins are again given by (7). By the dependence of f_t on its own past, the default probabilities vary over different stages of the business cycle. For example, defaults can become more likely in case of recessions, while upgrades prevail in expansionary periods. These effects can be captured by an appropriate choice of f and ψ (see further below). A more extensive discussion on stochastic migration rates can be found in, e.g., Credit Suisse (1999) and Belkin, Suchower, and Forest (1998). See also Koyluoglu and Hickman (1998) for a synthesis. In the present paper, we fully concentrate on the one-period framework. Detailed extensions to multi-period models are left for further research, see also Wilson (1997a,b).

The final characteristic in (1) of exposure j is its credit loss $\pi(\cdot)$. We assume that the amount of credit loss depends on the exposure's initial (k_j) and final (ℓ_j) rating category, as well as on the state of the economy (ψ).

This is expressed by specifying

$$\pi(j, k_j, \ell_j, \psi). \quad (8)$$

A credit loss not only occurs if a firm defaults, but also if the firm's rating deteriorates. The latter is due to differing credit risk spreads across rating categories, maturities, and industries. This explains the dependence of $\pi(\cdot)$ on j , k_j , and ℓ_j . The height of the credit spreads can also be affected by other economic variables. This is captured by the presence of the variable ψ . Note that the elements of ψ and the common factors f may either overlap partially, completely, or not at all. The presence of ψ in the credit loss function also allows for a straightforward link between credit risk and market risk. Up to now, these types of risks have been treated separately, at least from a supervisory point of view, see Basle (1988, 1996). This separation appears artificial. For example, Duffee (1996) argues that in recessions interest rates are typically lower, while defaults are more likely. This implies for example that the credit risk of the receive-fixed end of a swap transaction between two equally rated companies will be higher than the pay-fixed end: the pay-fixed party of the transaction is more likely to default in the recession, while the value of the receive-fixed end of the swap transaction is more valuable for lower interest rates.

Given all above definitions, the credit loss of a portfolio comprising n exposures is simply the sum of the individual credit losses:

$$C_n = \sum_{j=1}^n \pi(j, k_j, \ell_j, \psi). \quad (9)$$

3 The limiting distribution of portfolio credit losses

In this section we establish the distribution of the portfolio credit loss C_n when the number of exposures becomes large. Before we proceed with the main theorem, it proves useful to make some assumptions regarding the boundedness of factor loadings and potential losses.

Assumption 1

- $n^{-1} \sum_{j=1}^n \beta_j \omega_j^{-1} \beta_j^\top$ converges to a finite, positive definite matrix.
- $\beta_n^\top \beta_n / (n \omega_n) \rightarrow 0$.

$\sup_n n^{-1} \sum_{j=1}^n \sum_{\ell=1}^r \pi(j, k_j, \ell, \psi)^2$ is bounded almost surely (a.s.).

The first requirement in Assumption 1 is satisfied if the factor loadings and the idiosyncratic variances are bounded from above and below, respectively. In economic terms, this means that every exposure in the portfolio should exhibit non-negligible idiosyncratic risk. Moreover, the requirement of a positive definite limiting matrix implies that there are no redundant factors in the exposure's factor model (2). The second boundedness requirement implies that the factor loadings divided by the idiosyncratic variances grow less than linearly in the portfolio size n . Though this requirement differs technically from the first condition in Assumption 1, the economic interpretation is similar: the idiosyncratic risk should be non-negligible for all firms in the portfolio. The third part of Assumption 1 requires that potential squared portfolio losses are bounded on average for sufficiently large portfolios. This assumption is typically satisfied for most financial instruments. In particular for bonds we see that the assumption is trivially satisfied, as the maximum loss is the value of the bond, which is finite. For more complicated instruments like derivatives, note that the average squared loss mentioned in this part of Assumption 1 lies close to the *expected* squared loss for sufficiently large n under standard regularity conditions. Part 3 of the assumption is thus very much linked to the requirement that the expected squared losses are bounded (uniformly) for all exposures. This can be the case even though the loss itself may not be bounded, e.g., $\pi(j, k_j, \ell, \psi) = \max(0, \psi - K)$ with ψ log-normally distributed.¹

The following limit law constitutes the heart of the paper and is proved in the Appendix.

Theorem 1 *Define*

$$R_j^2 = \frac{\beta_j^\top \Omega_f \beta_j}{\omega_j + \beta_j^\top \Omega_f \beta_j} \quad (10)$$

as the R^2 of the factor regression model (2), i.e., the correlation between S_j and its 'fit' $\beta_j^\top f$. Moreover, let

$$v_j = \beta_j^\top \Omega_f^{1/2} \sqrt{1 - R_j^2} / \sqrt{\omega_j R_j^2}, \quad (11)$$

¹This corresponds for example to being on the long side of an OTC transaction involving a call option on a stock index with strike price K , where the counterparty can default.

such that $v_j^\top v_j = 1$, and let Y be an m -dimensional standard normal random variate defined by $Y = \Omega_f^{-1/2} f$. Let

$$\hat{\Phi}_{j\ell} = \Phi \left(\frac{c_{k_j, \ell} - \sqrt{R_j^2} v_j^\top Y}{\sqrt{1 - R_j^2}} \right) - \Phi \left(\frac{c_{k_j, \ell-1} - \sqrt{R_j^2} v_j^\top Y}{\sqrt{1 - R_j^2}} \right), \quad (12)$$

denote the conditional (on f) probability of migrating from rating k_j to rating ℓ , and define

$$B_n = \sum_{j=1}^n \sum_{\ell=1}^r \hat{\Phi}_{j\ell} \cdot \pi(j, k_j, \ell, \psi). \quad (13)$$

Then given Assumption 1 and the framework of Section 2, we have

$$n^{-1} C_n - n^{-1} B_n \xrightarrow{a.s.} 0, \quad (14)$$

with C_n the portfolio credit loss as defined in (9) and $\xrightarrow{a.s.}$ denoting almost sure convergence.

The expression B_n/n in Theorem 1 no longer depends on the idiosyncratic risk factors ε_j , but only on the systematic risk factors f and ψ . This considerably facilitates simulation from the credit loss distribution. A similar result is well known in *linear* portfolio theory. Indeed, within the CAPM model, only the systematic risk matters because it cannot be diversified away by increasing the number of exposures. Theorem 1 generalizes this result to the nonlinear setting of credit losses while simultaneously allowing for richer dynamics in terms of correlated defaults mechanisms and rating migrations.

To see how the simulations from the credit loss distribution can be simplified or even avoided, consider a one-factor model $m = 1$ where $v_j \equiv 1$, and $R_j^2 \equiv \rho^2$. Consider a set of loss functions $\pi(\cdot)$ corresponding to a portfolio corporate bonds or loans, see also the example further below in this section. Given an increasing credit spread for lower rating categories, it can be seen that B_n is a monotonically increasing function of Y . Take $J(Y) = \lim_{n \rightarrow \infty} B_n/n$, then using the usual transformation of variables technique the distribution of credit losses c is given by

$$\phi(\tilde{g}^{-1}(c)) / |\tilde{g}'(\tilde{g}^{-1}(c))|, \quad (15)$$

with $\tilde{g}^{-1}(\cdot)$ and $\tilde{g}'(\cdot)$ denoting the inverse and the first derivative of $\tilde{g}(\cdot)$, respectively. One can now use the trapezoid rule for numerical integration

$$\sum_{i=1}^N \tilde{g}(y_i)^\kappa \phi(y_i) \cdot (y_i - y_{i-1}) \quad (16)$$

for $\kappa = 1, 2$ to obtain an easy approximation of the expected credit loss and its variance, respectively, where $-K = y_0 < y_1 < \dots < y_N = K$ denotes an appropriate partitioning of the interval $[-K, K]$ for a sufficiently large constant $K > 0$. Computing the quantiles of the credit loss in this case is even easier than calculating the moments. Note that

$$P(C \leq c) = \delta \Leftrightarrow c = \tilde{g}(\Phi^{-1}(\delta)), \quad (17)$$

where $C = \lim_{n \rightarrow \infty} C_n/n$, such that the δ -quantile can be obtained by a simple evaluation of $\tilde{g}(\cdot)$ in one point. The present methodology is applicable whenever the credit risk mapping $\tilde{g}(\cdot)$ is a monotonic function of the stochastic variable Y in (12). As long as this is the case, the approach sketched above may be used, even if the notional amounts of the bonds and the R^2 's of the factor model regressions differ across firms. This provides a large computational advantage of our method over simulation based methods such as CreditMetrics, as long as one sticks to one-factor models and long corporate bond portfolios.

A second feature of the expression B_n/n is that the factor loadings and idiosyncratic variances do not enter directly. Only the R_j^2 of the factor model regression and the unit-length vector v_j matter. The R_j^2 determines the magnitude of the impact of systematic risk fluctuations on the j th credit loss. The larger the value of R_j^2 , the higher the influence of the systematic risk factors f on the j th credit loss. In particular, if $R_j^2 \downarrow 0$, the stochastic vector Y does not enter at all into the j th term of B_n in (13). In particular, if $R_j \equiv 0$, B_n becomes non-stochastic. The second exposure specific element entering (12) is the vector v_j . As already noted, this vector has unit length. Therefore, it can be interpreted as a directional vector, indicating which factors matter for a specific exposure. For example, if $m = 1$ such that (2) is a one-factor model, we have $v_j = \pm 1$. The directional vector v_j now indicates whether the systematic risk factor f has a positive or negative impact on S_j . A similar interpretation holds for multi-factor models. Note that the number of elements in (β_j, ω_j) is the same as in (v_j, R_j^2) . However, v_j has the restriction that $v_j^\top v_j = 1$, such that one cannot recover (β_j, ω_j) from (v_j, R_j^2) .

A third feature of (12) through (14) is that by replacing C_n by B_n , we are effectively replacing the actual credit loss due to a rating migration from k_j to ℓ_j by the conditional expectation of the j th credit loss. The conditioning set is given by f . This is essentially the same as in the case of linear portfolio theory.

Finally, it is important to stress that normality of f and the transformed variable Y is not necessary for (14) to hold. Indeed, if the factors f have a clear-cut economic interpretation, the credit risk manager might have some ideas about the future development of f in terms of its forecasting distribution. This distribution can then be used in (14) to obtain simulations from the credit loss distribution that are more relevant from an economic perspective. Alternatively, the credit risk manager might be interested in the effect of specific distributional assumptions for f , e.g., stress scenarios, in which case f places discrete (or unit) mass on certain scenarios. Also in that case, such distributional assumptions can readily be incorporated to obtain simulations for credit losses that are relevant for the purpose at hand.

To get a better feeling for the derived limit law, consider a one-factor version of our model with $m = 1$, $R_j^2 \equiv \rho^2$, $\rho \geq 0$, and $v_j \equiv 1$. So all exposures exhibit the same systematic risk. A similar model is studied in Belkin, Suchower, and Forest (1998a). There are only two rating categories $r = 1, 2$. The second rating category corresponds to a state of default. We also assume that the exposures in our portfolio are ordinary loans that either fully default or not, irrespective of the state of the economy ψ . This is captured by setting $\pi(j, k_j, 1, \psi) \equiv 0$ and $\pi(j, k_j, 2, \psi) \equiv \pi(j)$, where $\pi(j)$ is the size of the j th loan. Given these assumptions, B_n as defined in (13) simplifies to

$$B_n = \sum_{j=1}^n \left[1 - \Phi \left(\frac{c - \rho Y}{\sqrt{1 - \rho^2}} \right) \right] \cdot \pi(j), \quad (18)$$

such that

$$C_n/n - \left[1 - \Phi \left(\frac{c - \rho Y}{\sqrt{1 - \rho^2}} \right) \right] \cdot \bar{\pi}_n \xrightarrow{d} 0, \quad (19)$$

where $\bar{\pi}_n$ is the average size of a loan, $\bar{\pi}_n = n^{-1} \sum_{j=1}^n \pi(j)$, while c is a constant determining the default probability.

Figure 2 plots the c.d.f. of the limiting credit loss C for various values of ρ and c given an average loan size of $\bar{\pi} = 1$. This means that credit losses are expressed as a fraction of the notional amount. The two values of c considered give rise to a default probability of 5% and 1%, respectively.

It is clear from the figure that large credit losses occur much more often if one allows for positive correlations between the underlying ‘surplus’ variables

Figure 2: Asymptotic default loss (C) distributions for constant $R_j^2 \equiv \rho^2$ and $v_j \equiv 1$ in a one-factor model ($m = 1$), see Theorem 1. There are only two rating categories, one of which corresponds to a state of default. The constant c is chosen such that there is either a 5% (left-hand panel) or a 1% (right-hand panel) probability of default.

S_j . The result holds irrespective of the probability of default, as specified through c . We also note that for smaller values of ρ the credit loss distribution becomes more concentrated. For $\rho = 0$, the c.d.f. collapses to a step function, taking the value 0 before, and 1 after the expected credit loss, respectively. In that case, the limiting distribution in Theorem 1 is of no use for credit risk management. A second order limiting result would be needed instead of the presented first order result.

4 Tail behavior of average credit losses

Figure 2 suggests that a higher correlation between default risks increases the likelihood of extreme portfolio credit losses. The increase in probability mass in the tails may partly be due to an increased variance of the credit portfolio. However, we find that the properly rescaled C still exhibits more probability mass in the tails than the normal distribution. Stated otherwise, the tails of the derived limit law seem to decline at a lower than exponential rate. In this section we show that the tail probabilities are polynomially declining functions of the credit loss quantile, i.e., the average credit loss exhibits ‘fat tails’. Moreover we establish the relation between the tail index and the asset correlations through the factor model fit R_j^2 and the directional dependence v_j . A correct assessment of the tail index is important for a proper credit risk assessment, especially if one is interested in the credit loss associated with very small significance levels, i.e., very far into the loss distribution’s tail. In particular, if the credit loss distribution is fat-tailed, common rules of thumb for computing loss quantiles no longer apply. For example, the 99.9%

percentile may lie much more than 3 standard deviations above the expected loss, which is the number one would expect for the normal distribution. In other words, extreme credit losses are much more likely to happen than under the normal distribution.

In order to characterize the tail behavior of credit losses, we first introduce the statistical definition of tail fatness. Let $F(\cdot)$ denote a distribution function. Embrechts, Klüppelberg, and Mikosch (1997) give a necessary and sufficient condition for distributions to exhibit polynomially declining or ‘fat’ tails:

$$\lim_{t \rightarrow \infty} \frac{1 - F(tx)}{1 - F(t)} = x^{-\alpha}, \quad (20)$$

for $x > 0$, where α is called the tail index. The condition is often dubbed the condition of ‘regular variation’. The polynomial’s exponent is the rate at which the tail probabilities decline in x . It can be interpreted as the number of bounded moments that exist (are finite) for a specific distribution. For example, a normal distribution is thin-tailed because all moments exist; its tail index is equal to infinity. The Student-t distribution, by contrast, exhibits a finite tail index equal to the number of degrees of freedom and is, thus, fat-tailed.

For sake of clarity, we start by examining the tail behavior of credit losses within the simplified one-factor setting of Section 3. Because the support of the portfolio credit loss distribution has a finite upper end point equal to 1, the above definition of regular variation is not applicable. Embrechts et al. (1997) also provide a suitable definition for distributions with bounded support. We use this definition to prove the following theorem, see the Appendix.

Theorem 2 *For the one-factor model set out in Section 3, the credit loss distribution has a tail index equal to*

$$\frac{1 - \rho^2}{\rho^2}. \quad (21)$$

Loosely speaking, Theorem 2 implies that the tails of the credit loss distribution are of the form

$$F(C) \approx 1 - (1 - C)^{(1 - \rho^2)/\rho^2}, \quad (22)$$

for sufficiently large credit losses C , i.e., close to 1. Equation (22) clearly demonstrates the polynomially instead of exponentially declining tail shape of the credit loss distribution.

The present results have a direct bearing on empirical findings. First, the empirical detection of fat tails of credit risk distributions can be explained

from a micro-based approach of individual exposures by allowing for common factors. Though the factors and idiosyncratic shocks may be normally distributed (and thus thin-tailed), the credit loss distribution will exhibit heavy tails provided the exposures in the credit portfolio are correlated. Larger correlations imply fatter tails. Second, one should be very careful in using the normal distribution as an approximation to the credit loss distribution in empirical modeling exercises. Of course, situations can be conceived where the approximation provided by this distribution is not too bad. In general, however, we expect fat-tailed and/or skewed distributions to provide better approximations. This holds especially for correlated default probabilities and low probability quantiles, e.g., credit loss realizations in the tail of the distribution such as the 99th or 99.9th percentile, see also Sections 5 through 7. In case one restricts all asset or surplus (and thus default) correlations to zero, one can resort to the normal approximation, see also Section 5. A setting with vanishing asset correlations, however, is highly unrealistic from a practical point of view. It also contradicts the empirical evidence that default correlations are correlated over stages of the business cycle, see Jónsson and Fridson (1996) and Fons (1991). The third implication of the non-normal tail index of the credit loss distribution pertains to the usual way of presenting credit loss quantiles in terms of the number of standard deviations above the expected credit loss. These have to be interpreted with great caution. For different degrees of tail fatness and different portfolio compositions, similar quantiles correspond to widely different numbers of standard deviations in excess of expected loss, see also Section 5. For example, if the tails contain sufficient probability mass ($\rho^2 > 1/3 \Leftrightarrow (1 - \rho^2)/\rho^2 < 2$), estimated second moments can become very high and unstable. As a result, potential credit losses in terms of standard deviations might become very low. By contrast, if ρ is close to zero, the limiting distribution degenerates and the number of standard deviations might become very high.

We now show how the above results on fat tails for credit losses generalize to non-constant R_j^2 and $\pi(j)$ and to multi-factor models. First consider a one-factor model with non-constant R_j^2 and $\pi(j)$. We only consider two different values for R_j^2 and $\pi(j)$. The arguments presented here, however, directly carry over to situations with more than two possible values and even to situations with a continuum of possible values. Let $R_j^2 = \hat{R}_1^2$ and $\pi(j) = \hat{\pi}_1$ for $[\lambda \cdot n]$ of the exposures in the portfolio, with $\lambda \in [0, 1]$ and $[x]$ denoting the integer part of x . For the remaining $n - [\lambda \cdot n]$ exposures, $R_j^2 = \hat{R}_2^2$ and $\pi(j) = \hat{\pi}_2$. Using analogous derivations as in the case of constant R^2 's and $\pi(\cdot)$'s, we establish that the tail index of the credit loss distribution is given

by

$$\max_{j \in \{1,2\}} (1 - \hat{R}_j^2) / \hat{R}_j^2. \quad (23)$$

Note that for $\hat{R}_1^2 \equiv \hat{R}_2^2 \equiv \rho^2$, we recover (21). So the exposure with the smallest correlation with its systematic component, i.e., with the smallest R^2 , dominates the tail behavior of credit losses. This is intuitively clear. For example, consider the case $\hat{R}_1^2 \equiv 0$. Then $[\lambda \cdot n]$ of the firms in the portfolio only display idiosyncratic risk, which can be eliminated through diversification. Because of this diversification argument, the potential credit losses caused by these firms only affect the expected loss and not the tail behavior. In particular, the probability of average credit losses being near the upper bound $\bar{\pi} = \lambda \cdot \hat{\pi}_1 + (1 - \lambda) \cdot \hat{\pi}_2$ is zero (in the limit), implying a thin tail. Now if $R_1^2 > 0$, a similar line of reasoning can be followed to demonstrate that the firms with the smallest systematic risk component dominate the tail behavior.

The above example for non-constant R^2 also illustrates another important property of the tail behavior of credit loss distributions. Again, consider the extreme case $0 = \hat{R}_1^2 < \hat{R}_2^2$. The worst case credit loss outcome is given by $(1 - \lambda)\hat{\pi}_2 < \bar{\pi}$, with $\bar{\pi}$ the maximum average credit loss as defined earlier. So the tail between $(1 - \lambda)\hat{\pi}_2$ and $\bar{\pi}$ has to be flat, i.e., thin. The tail near $C = (1 - \lambda)\hat{\pi}_2$, however, is fat, as can be seen from (21). So if R_j^2 is non-constant, extreme tail behavior is dominated by the largest $(1 - R_j^2)/R_j^2$ values. Quantiles less far out in the tails, however, may also strongly be affected by smaller values of $(1 - R_j^2)/R_j^2$. This mixed tail behavior can result in unconventional combinations of skewness and leptokurtosis. It may therefore be very difficult to devise parametric distributions that capture all salient features of a credit loss distribution for large portfolios. Moreover, it may also be very difficult to employ (semi-nonparametric) extreme-value statistical theory for estimating higher order quantiles of the credit loss distribution, see, e.g., Danielsson and Vries (1997). Such methods presume a certain degree of homogeneity of tail observations, which might be inappropriate given the strongly varying tail behavior of C over different parts of the portfolio.

Deriving the tail behavior of credit losses for multi-factor models is somewhat more complicated. We only present the resulting tail index in case there are only a finite number n^* of different combinations $(\hat{R}_j^2, \hat{v}_j, \hat{\pi}_j)$. The result can be generalized to a continuum of possibilities with a corresponding increase in technical details. First, let $G \subset \{1, \dots, n^*\}$, and define the

limiting credit loss function

$$g(Y, G) = \sum_{j \in G} \lambda_j \cdot \left[1 - \Phi \left(\frac{c - \sqrt{\hat{R}_j^2} \hat{v}_j^\top Y}{\sqrt{1 - \hat{R}_j^2}} \right) \right] \cdot \hat{\pi}_j,$$

with λ_j the fraction of the portfolio with combination $(\hat{R}_j^2, \hat{v}_j, \hat{\pi}_j)$. Let

$$\pi^* = \sup_{y \in \mathbb{R}^m} g(y, (1, \dots, n^*))$$

be the maximum average credit loss of the limiting portfolio. Note that this maximum loss need not coincide with the maximum average loss for a portfolio of finite size n . The latter is given by $\sum_{j=1}^{n^*} \lambda_j \hat{\pi}_j$, because there is always a (small) probability that the idiosyncratic risks push all exposures in the portfolio into default. By contrast, if for example $\hat{R}_j^2 = 0$ and $\hat{\pi}_j > 0$ for some $j \in \{1, \dots, n^*\}$, then $\pi^* < \sum_{j=1}^{n^*} \lambda_j \hat{\pi}_j$. Next, define

$$\mathcal{G} = \{G \subset (1, \dots, n^*) \mid \sup_{y \in \mathbb{R}^m} g(y, G) = \pi^*\}.$$

$\mathcal{G}$ consists of the parts of the portfolio that, when combined, can result in the maximum average limiting credit loss. One can now prove that the upper-tail index in the multi-factor case is given by

$$\min_{G \in \mathcal{G}} \max_{j \in G} \frac{1 - \hat{R}_j^2}{\hat{R}_j^2}, \quad (24)$$

Note that for identical values of the $\hat{R}_j^2$'s, the tail index of the multi-factor model is at most as high as that of the single factor model, see (23). This follows from the fact that (24) takes the minimum of the tail index from (23) over different subsets G . The resulting tail index is, *ceterus paribus*, just as high or higher for multi-factor models compared to one-factor models.² This has a clear intuitive explanation. For the single factor model with $\hat{v}_j \equiv 1$, tail behavior is determined by the exposures with the smallest systematic dependence in terms of $\hat{R}_j^2$. If multiple directions $\hat{v}_j$ are possible and if there are multiple factors, however, it is possible that different realizations of the systematic risk component f in (2) give rise to similar large values of credit

²Note that in practice the $\hat{R}_j^2$'s for multi-factor models will also be higher, thus enforcing the mentioned *ceterus paribus* effect. Also note, however, that in a multi-factor model π^* may turn out to be lower than for the one-factor model, such that the tail index may be higher, while the upper end point of the credit loss distribution is lower.

losses. The systematic dependence (R_j^2), however, of the portfolio exposures in these alternative directions might be very different, resulting in different tail behavior of credit losses in the different directions mentioned. Equation (24) now states that the fatter tail dominates, see Ibragimov and Linnik (1971). This can already be seen in a single factor example ($m = 1$). Take $n^* = 2$, $\lambda_j \equiv 1/2$, $\hat{\pi}_j \equiv 1$, $\hat{v}_1 = 1$, $\hat{v}_2 = -1$, $\hat{R}_1^2 = 0.5$, and $\hat{R}_2^2 = 0.1$. The portfolio now consists of two parts, both of which contain a homogeneous set of exposures and comprise 50% of the portfolio. The first part of the portfolio has a positive ($\hat{v}_1 = 1$) and high ($\hat{R}_1^2 = 0.5$) correlation with the systematic risk factor, while the second part has a negative ($\hat{v}_2 = -1$) and low ($\hat{R}_2^2 = 0.1$) correlation. As the portfolio contains an equal number of type 1 and 2 exposures, the maximum limiting average credit loss ($\pi^* = 0.5$) is attained if the systematic risk factor f tends to either plus or minus infinity. If f tends to plus infinity, the type 1 exposures default, so $f \rightarrow \infty$ results in a tail index of $0.5/0.5 = 1$. Similarly, if f tends to minus infinity, the type 2 exposures default, resulting in a tail index of $0.9/0.1 = 9$. Following standard results for tail behavior from, e.g., Ibragimov and Linnik (1971), the fatter tail dominates, i.e., the one with tail index 1. This also follows from (24), where we take the minimum of 1 and 9.

5 Credit loss quantiles of stylized portfolios

In Sections 3 and 4 we studied the behavior of the limiting *default* loss distribution in the stylized setting of Belkin, Suchower, and Forest (1998a) and somewhat more general factor models. In this section, we investigate the behavior of credit loss distributions in more detail. We generalize the previous set-up by allowing for differences in initial ratings, loan portfolio maturities, and magnitudes of the systematic risk component.

The corporate bond maturities A4 may vary from 1, 5, to 10 years. All bonds are assumed to be of the same type. At the outset, we evaluate the bonds at par using the yields taken from the CreditMetrics site on October 6, 1998. The base yield on the bond is increased by the credit spread. This credit spread depends on the initial rating of the firm (see further below). The spreads were also downloaded from the CreditMetrics site, see Table 1. We assume constant base yields and credit spreads over the credit risk evaluation period in order to focus entirely on the effect of credit risk without incorporating market risk.

The transition probability matrix used is presented in Table 1. We use a classification with 7 categories: AAA, AA, A, BBB, BB, B, and CCC. In addition, we have the default category D, see also Figure 1. The tran-

sition probabilities are used to determine the binning constants $c_{k\ell}$ through (6). Default and rating migration probabilities are then equal to long-term historical averages.

We consider three different quality levels for the initial portfolio. The high-quality portfolio constitutes of (equal) positions in AAA, AA, and A rated companies while the medium quality portfolio has (equal) positions in BBB, BB, and B rated firms. Finally, the low-quality portfolio only contains CCC rated firms.

To complete the comparison, we consider three different values for the degree of systematic risk: $R^2 = 0.2, 0.4, 0.6$. We assume that these R^2 's correspond to a one-factor model where all β_j 's are positive, such that $v_j \equiv 1$. All firms have the same R^2 , i.e., $R_j^2 \equiv R^2$. The chosen values of R^2 imply tail indices as defined in (23) between 4 for $R^2 = 0.2$ and $2/3$ for $R^2 = 0.6$.

Using the present setting, we derive the credit loss distribution. Note that the right-hand side of (12) reduces to

$$\sum_{k=1}^{r-1} \sum_{\ell=1}^r \lambda_k \left[\Phi \left(\frac{c_{k\ell} - \sqrt{R^2} Y}{\sqrt{1 - R^2}} \right) - \Phi \left(\frac{c_{k,\ell-1} - \sqrt{R^2} Y}{\sqrt{1 - R^2}} \right) \right] \cdot \tilde{\pi}(k, \ell), \quad (25)$$

where $r = 8$, λ_k denotes the fraction of firms in the portfolio with initial rating k , e.g., 33% for AAA ($r = 1$) in the high-quality portfolio, and $\tilde{\pi}(k, \ell)$ is the credit loss/gain on the corporate bond when a firm migrates from rating k to rating ℓ , i.e.,

$$\tilde{\pi}(k, \ell) = \begin{cases} (1 - y_k/y_\ell) \cdot (1 - (1 + y_\ell)^{-(M-1)}) & \text{for } \ell = 1, \dots, r-1, \\ 1 + y_k & \text{for } \ell = r, \end{cases} \quad (26)$$

with y_k the yield on an M -year corporate bond with initial rating k .

Using the methodology mentioned in the first comment to Theorem 1, we can now compute the credit loss quantiles without resorting to simulations. The expected loss and its variance can moreover be computed by numerical integration, see (16). The boxplots in Figure 3 summarize the 27 credit loss distributions resulting from our experiments.

First consider the effect of changing the degree of systematic risk, R^2 . It is clear that more systematic risk leads to more prolonged tails. The upper credit risk quantiles all shift to the right. Note, however, that the lower quantiles may shift in the opposite direction, see for example the medians of the C rated portfolios. This stems from the fact that credit rating upgrades will also be more correlated over the different exposures in the portfolio. The effects are substantial. For certain settings, the 99.9th percentile of

Figure 3: The figure summarizes the effect of varying maturities, systematic risk, and portfolio quality on the limiting credit loss distribution by means of boxplots. The 3 left plots express the credit quantiles as a fraction of the notional whereas the three right plots express the credit quantiles in terms of the numbers of standard deviations in excess of the expected loss. Each row of two plots corresponds with a given maturity. Each plot contains 3 panels corresponding with different degrees of systematic risk R^2 . Each panel comprises 3 boxplots representing credit loss distributions of a high-quality portfolio (A) with equal positions in AAA, AA, and A rated firms, a medium quality portfolio (B) with equal positions in BBB, BB, and B rated firms, and a low-quality portfolio (C) consisting of CCC rated firms only. Each box represents the interquartile range of credit losses, the middle line indicating the median. The whisker of the boxplot has 4 markings, relating to the 0.9, 0.95, 0.99, and 0.999 quantile of credit losses.

credit losses may shift by more than 20% of the invested notional when R^2 is increased from 0.4 to 0.6.

If we express the quantiles in terms of standard deviations in excess of the expected loss (right panels in Figure 3), we see a remarkable feature of credit loss distributions. The distributions with the most prolonged tails appear to have the smallest ‘surprise’ element in terms of standard deviations. Though the C rated portfolios have the highest 99.9% quantiles, these quantiles are only about 3 to 4 standard deviations in excess of the expected loss. This is the number one would expect when using a normal distribution. By contrast, the 99.9% quantiles of the A rated portfolios appear almost negligible, but they are between 9 and 13 standard deviations in excess of the expected loss. These large numbers illustrate that the use of the normal distribution for approximating credit loss quantiles may be completely inappropriate in this case: the 99.9% quantile of the normal is only about 3 standard deviations in excess of its mean. The log-normal distribution, by contrast, has its 99.9% quantile about 9.4 standard deviations in excess of its mean. As such, the log-normal may be more suited for the highly rated portfolios, but it will prove inappropriate for the portfolios with many low-rated exposures. Instead of relying on the normal or the log-normal, it appears more appropriate to use the limiting distribution of Theorem 1 directly, either using the analytical method of the present paper, or the more traditional simulation based methods, see, e.g., J.P. Morgan (1999). In any case, great care has to be taken in interpreting credit loss quantiles represented in terms of standard deviations in excess of expected loss.

Next, Figure 3 clearly shows that low-quality portfolios have a worse credit risk performance. Indeed, upper credit loss quantiles shift to the right if one includes more poorly rated companies in the portfolio. The lower quantiles also shift to the right most of the time. If the maturity of the bond is as long as 10 years, however, the lower quantiles shift to the left for a high degree of systematic risk ($R^2 = 0.6$). If we express all quantiles in terms of standard deviations, we again note the reverse in the magnitudes compared to the raw quantiles.

Finally, the bonds’ maturities appear to have only minor effects on the credit loss quantiles. Of course, longer maturities give rise to higher upper credit loss quantiles. For 1 year bonds, the only credit losses are those due to default. For 5 and 10 year bonds, we also have to take the effect of rating migrations and differing credit spreads into account. These effects will have a higher impact the longer the maturity, i.e., the higher the duration or interest elasticity of the bond. The effect can work both ways, because upgrades as well as downgrades have a larger effect for longer maturities.

Table 2: Percentage of firms in the portfolio with specific initial ratings

Rating	AAA	AA	A	BBB	BB	B	CCC
%	3	13	30	22	17	14	1

6 Examples based on empirical data

In the previous section, we studied the behavior of the limiting credit loss distribution for stylized portfolios. The advantage of using such portfolios is that one can isolate the effects of parameter changes on the limiting distribution. On the other hand, the credit loss quantiles have limited practical applicability for credit portfolios in practice. In this section we study the limiting distribution on a more realistic portfolio of corporate bonds.

We need 5 types of data: the rating migration probabilities, the yields and yield spreads, the initial ratings of the exposures in the portfolio, the credit loss functions $\pi(\cdot)$, and the R^2 's of the factor model (2). The rating migration probabilities, yields, and yield spreads used in this section are the same as in the previous section and were reported in Table 1.

Because proprietary data on the typical quality of a bank's portfolio are difficult to obtain, we consider a portfolio with initial rating distribution equal to the distribution implied by the total number of S&P rated companies in 1997, see S&P (1998). The ratings' empirical distribution is given in Table 2 and comprises a wide variety of firms. The reported distribution forms the benchmark case. For sake of comparison, we also consider a low-initial-rating and a high-initial-rating portfolio. The high-rating portfolio is obtained by upgrading exposures in the benchmark portfolio (except the AAA ones) by one category. Similarly, the low-rating portfolio is constructed by downgrading the exposures in the benchmark portfolio (except the CCC rated firms) by one category.

Suitable empirical values for the R^2 are obtained as follows. We take equity returns as a proxy for the surplus variables S_j in (2), see also J.P. Morgan (1999). We downloaded the Research Top 2000 company list (comprising monthly total returns of 1762 listed companies) from DATASTREAM Inc. The sample period runs from January 1980 until December 1998. We regress each return on a constant and on the total return of the S&P500 index. If less than 5 years of data are available, the firm is omitted from the sample. This eventually results in a total of 1645 R^2 's. Figure 4 presents kernel estimates of the R^2 frequency distributions for three different data frequencies: low R^2 (monthly data), middle R^2 (quarterly data), and high R^2 (annual data).

Given our available data set, we did not find any significant relationship between R^2 values and firm ratings. We therefore impose the same distri-

Figure 4: R^2 values of factor model regressions using monthly, quarterly, and annual data. The data are obtained from the Research Top 2000 list of DATASTREAM Inc. and comprise 1762 firms observed over the period January 1980–December 1998. The factor model explains total return of the firm by a constant and the total return on the S&P500. A minimum of 5 years of data is used for the factor model regressions.

bution of R^2 's per rating categories. We also inspected the values of the directional vectors v_j of Theorem 1. In our one-factor set-up we have $v_j = \mathbf{1}$ or $v_j = -\mathbf{1}$ depending on whether $\beta_j > 0$ or $\beta_j < 0$, respectively. For the vast majority of firms in our sample, the estimated β_j was positive and none was found to be significantly negative. Thus we can safely set $v_j \equiv \mathbf{1}$ for all exposures in the sample.

As the structure of a typical bank portfolio is difficult to determine due to data (non)-availability, we only consider a fixed maturity loan for each firm in the portfolio. This can, however, easily be altered. In our present study, the maturities and sizes of the loans are identical for all firms. This may seem overly simplistic. Note, however, that the size and maturity of the loan can be interpreted as a kind of portfolio average, see (19), such that the maturity of the bond can be interpreted as the maturity of the bank portfolio. As a result, only considering an identical fixed maturity loan for each firm is a viable approach from a limiting portfolio point of view. Of course, differing loan sizes across firms may be very important for the adequacy of the limiting distribution as an approximation to the credit loss distribution for portfolios of finite size, but this is deferred until Section 7. We consider the same three cases as in Section 5, namely a unit loan size with a maturity of 1, 5, and 10 years, respectively.

A final important difference with the set-up of Section 5 is that we consider a non-zero recovery rate. Based on S&P (1998), a realistic setting is obtained if we set the recovery rate to 0.5 for all bonds and firms. This means that the lower line in (26) is replaced by $(1 + y_k)/2$.

The results are presented in Figure 5. Altering the parameters of the credit loss distribution seems to shift the quantiles in roughly the same direction as Figure 3. Longer maturities, lower portfolio quality, and a higher degree of systematic risk all make extreme credit losses realizations more likely. For example, using the quarterly instead of the monthly R^2 for the benchmark portfolio results in a 99.9% credit loss quantile that is higher by about 2% of the notional. A 4% increase in terms of the notional is established for the monthly R^2 values if the maturity of the portfolio is increased from 1 year to 10 years. The differences in the upper credit loss quantiles are extremely large if the initial rating distribution of the portfolio is varied from high-quality (H) to low-quality (L). By contrast, switching from the quarterly to the annual R^2 's has virtually no effect.

Concerning the size of the upper quantiles, we note that for the benchmark portfolio with the R^2 's based on monthly data, the 99.9% quantile lies between 5% and 13% of the invested notional. For a typical bank portfolio, we expect a duration of 5 rather than 1 or 10 years. For the 5 year maturity bond, the 99.9% quantile lies between 7 and 9 per cent of the notional. This comes close to the number of 8% prescribed by the Basle proposals, Basle Committee on Bank Supervision (1988). For high-quality portfolios, the typical 99.9% lies significantly below the Basle guideline of 8%, while the reverse holds for the low-quality portfolios. Also note that the normal approximation is not applicable in the present empirical setting. The 99.9% quantiles all lie far more than 3 standard deviations above the expected loss, which would be the appropriate number for the normal distribution. Also note the inapplicability of the log-normal distribution, as the 99.9% quantile lies substantially above or below 9.4 for most portfolios considered. In any case, it appears that proper credit risk management should allow for more diversified capital requirements depending on the maturity, composition, quality, and systematic risk of the portfolio. Similar focal points apply to supervisory institutions for credit risk.

Figure 5: The figure summarizes the effect of differing maturities, systematic risk and portfolio quality on the limiting credit loss distribution by means of boxplots. The plots express the credit loss either as a fraction of the notional (left 3 panels), or in terms of the number of standard deviations in excess of the expected loss (right 3 panels). Each row of 2 plots relates to corporate bond portfolios of a given maturity. Each plot contains three panels for three different degrees of systematic risk R^2 . The left panel is based on the distribution of R^2 's using regressions with monthly data ($M-R^2$), while the middle and right panel use quarterly ($Q-R^2$) and yearly ($Y-R^2$) data, respectively. See also Figure 4. Within each of these panels, 3 boxplots are presented for a high-quality portfolio (H) with initial rating distribution as in Table 2, but with all firms upgraded by one rating category; a benchmark portfolio (B) with initial rating distribution as in Table 2, and a low-quality portfolio (L) with all firms downgraded one category with respect to the benchmark situation. Each box represents the interquartile range of credit losses whereas the middle line indicates the median. The whisker of the boxplot has 4 markings, relating to the 0.9, 0.95, 0.99, and 0.999 quantile of credit losses. The R^2 distribution is generated using data extracted from DATASTREAM Inc. A recovery rate of 0.5 is used for all bonds.

7 Speed of convergence to the limiting distribution

So far, we have concentrated on calculating credit loss quantiles if the number of exposures in the portfolio gets very large. For credit risk management, however, quantiles corresponding to a limited number of exposures in the portfolio may be more relevant. In this section we investigate for which portfolio size n the upper quantiles for the limiting distribution provide a reasonable approximation for the finite sample quantiles. As in the previous sections, we concentrate on the one-factor model.

The previous section constitutes the theoretical framework for studying the speed of convergence towards the limit law. We only consider loan maturities of 1 year and 5 years. Note that credit risk equals default risk for a one year maturity. In contrast, for a maturity of 5 years, credit risk both encompasses default risk and risk due to credit rating migrations. For each of the 2 maturities, we conduct 9 different experiments, relating to 3 different distributions of R^2 (Monthly, Quarterly, Yearly) and 3 initial rating distributions (High, Benchmark, Low).

For a finite number of exposures n in the portfolio, the distribution of R^2 's over the portfolio has to be discretized. Let $F_{R^2}^{-1}(\cdot)$ denote the inverse c.d.f. of the R^2 's corresponding to the p.d.f.'s provided in Figure 4. Assume n_k exposures with initial rating k in a portfolio of size n . For these exposures, we set the R^2 value equal to

$$F_{R^2}^{-1}(i/(n_k + 1)), \quad (27)$$

for $i = 1, \dots, n_k$. This discretization implies an identical distribution of R^2 's across rating categories when the number of exposures becomes large. For finite n , the R^2 's are spread evenly over the interval $[0,1]$ using the inverse c.d.f. For example for the R^2 's based on monthly data this implies that there will be relatively more low R^2 values than high ones for every n . In the limit $n \rightarrow \infty$, the postulated distribution $F_{R^2}(\cdot)$ of R^2 is recovered for every rating category.

We extend the above simulation exercise by introducing a measure of portfolio 'heterogeneity' or 'dispersion'. Portfolio dispersion may slow down the convergence towards the limiting distribution. In order to create dispersion in the portfolio, we only vary the loan sizes. Let $\nu \in \mathbb{N} \cup \{0\}$ be our measure of portfolio dispersion. If $\nu = 0$, we set all loan sizes to unity such that there is no dispersion. If $\nu > 0$, we set

$$\pi(j, k_j, \ell_j) = \varsigma_j \tilde{\pi}(k_j, \ell_j), \quad (28)$$

where $\tilde{\pi}(\cdot)$ denotes the loss function of a unit loan size, and ς_j denotes the size of loan j . Let $\lfloor x \rfloor$ again denote the integer part of x . Then we let $\lfloor n/\nu \rfloor$ of the exposures have loan sizes that are uniformly distributed over the interval $[0.5 + \nu, 1.5 + \nu]$. The remaining exposures have loan sizes uniformly distributed over the interval $[0.5, 1.5]$. Broadly speaking, we now have a portfolio with both large and small loans. The large loans comprise a fraction of about $1/\nu$ of the portfolio, while the remaining fraction of $(\nu-1)/\nu$ consists of small loans. By increasing ν , we can increase the loan portfolio's degree of dispersion or heterogeneity. For example, for $\nu = 1$, loan sizes are distributed uniformly (per rating category) over the interval $[1.5, 2.5]$, such that the degree of heterogeneity is relatively modest. For $\nu = 10$, 10% of the portfolio consists of loans that have an approximately ten times larger than the median portfolio loan size. This type of dispersion may significantly distort the applicability of the limiting distribution for finite n .

Instead of introducing the different loan sizes directly at the portfolio level, we introduce them at the initial rating level. In particular, for each rating category k with n_k exposures, we consider a fraction of $1/\nu$ of larger loans and $(\nu - 1)/\nu$ of smaller loans. We adapt the procedure for assigning R^2 -values to individual exposures, accordingly. In particular, we use the approach sketched in (27) per rating category and segment of loan size (large/small) instead of per rating category only.

We are now ready to calculate credit loss quantiles of the limiting credit loss distribution under different scenarios for maturities, portfolio dispersion, R^2 distribution, initial rating distribution, and portfolio size. These quantiles are used as the benchmark in checking the convergence speed of the credit loss distribution. They can be computed using (17) without the need for simulations.

For finite portfolios, computing quantiles is much more difficult. Generally, we have to resort to simulations. Estimated quantiles of credit losses based on simulations can be very unstable. We tackle this problem as follows. For portfolios of size $n = 100, 200, \dots, 1000$, we first generate 20,000 simulations from the factor model (2) using 10,000 pairwise antithetic draws. These simulations are used to obtain estimates of the 50th, 75th, 90th, 95th, 99th, and 99.9th percentile of credit losses. In order to further reduce the variability of the simulated quantiles, this process is repeated 10 times. The final estimates of the quantiles are the averages over the 10 replications. The discrepancies between these averages and the limiting distribution's quantiles in percentages of the notional are presented in Figures 6 through 13.

For the homogeneous portfolios ($\nu = 0$), the results in Figures 6 and 7 reveal that the quantiles of the limiting credit loss distribution generally lie very close to those of the distribution for finite portfolio sizes. The difference

Figure 6: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 1 year for all firms in the portfolio. The degree of portfolio heterogeneity ν is 0, so all firms have the same bond size. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 8: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 1 year for all firms in the portfolio. The degree of portfolio heterogeneity ν is 1, so the bond size is uniformly distributed on the interval $[1.5, 2.5]$. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 9: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 5 years for all firms in the portfolio. The degree of portfolio heterogeneity ν is 1, so the bond size is uniformly distributed on the interval $[1.5, 2.5]$. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 10: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 1 year for all firms in the portfolio. The degree of portfolio heterogeneity ν is 5, so 20% of the firms have a loan size that is approximately 5 times the loan size of the remaining 80% of firms. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 11: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 5 years for all firms in the portfolio. The degree of portfolio heterogeneity ν is 5, so 20% of the firms have a loan size that is approximately 5 times the loan size of the remaining 80% of firms. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 12: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 1 year for all firms in the portfolio. The degree of portfolio heterogeneity ν is 10, so 10% of the firms have a loan size that is approximately 10 times the loan size of the remaining 90% of firms. The discrepancy is given in percentage terms of the notional. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

Figure 13: The figure presents the estimated discrepancy between the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) for a portfolio consisting of n firms and those of the corresponding limiting distribution. The figure presents the results for a bond maturity of 5 years for all firms in the portfolio. The degree of portfolio heterogeneity ν is 10, so 10% of the firms have a loan size that is approximately 10 times the loan size of the remaining 90% of firms. The figure contains 9 plots. The columns contain the results for given degree of systematic risk (R^2) based on monthly (M), quarterly (Q), or annual (Y) factor model regressions, while the rows contain the results for a given distribution of initial ratings (see Section 5 for the details). Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles. The quantiles are estimated on samples of size 20,000.

appears to be in an acceptable range for portfolios with at least $n = 300$ exposures. This certainly holds if we account for the sampling uncertainty in the quantile estimates for finite portfolio sizes, see for example the instability of the estimated 99.9% quantile in the bottom-middle and lower-right panels of Figure 7. For homogeneous portfolios, the convergence behavior is very similar for different initial rating distributions and values of systematic risk (R^2). The maturity does not have a significant effect on convergence either.

Figures 8 and 9 show that small deviations from portfolio homogeneity leave the finite portfolio convergence behavior virtually unaffected. A number of 300 to 400 exposures in the portfolio suffices to get a good match between the finite sample quantiles and the limiting quantiles. Hence the limiting distribution provides a good approximation for most practical purposes.

If the portfolio heterogeneity is more pronounced, the convergence behavior is affected significantly. This is shown in Figures 10 and 11 for $\nu = 5$, and in Figures 12 and 13 for $\nu = 10$. As expected, much larger portfolio sizes are needed to obtain a similar accuracy as in the case of homogeneous portfolios, i.e., n should preferably lie within the range 500-800. Thus the slowdown in convergence speed is more pronounced for larger values of ν . For portfolios comprising more than 1,000 exposures, however, the figures also reveal that the limiting distribution still provides a useful approximation for practical risk management problems.

Summarizing, the limiting distribution fits the distribution for finite portfolios very closely for reasonably sized homogeneous portfolios ($n \geq 300$). The fit decreases if the homogeneity of the portfolio decreases, i.e., if ν increases. In that case, larger portfolio sizes are needed, n between 500 and 800. It is well-known that there is a generic uncertainty surrounding some of the input parameters of credit risk models, e.g., precise default probabilities and recovery rates, as well as some of the output, e.g., simulated upper (99.9%) quantiles. Given this uncertainty, the discrepancies reported in Figures 6 through 13 seem acceptable once one is willing to adopt a simulation based credit risk management model. The additional error caused by the use of a limiting distribution to approximate quantiles of finite portfolios appears limited for all practical purposes if the portfolio size is sufficiently large.

8 Conclusions

In this paper we studied the credit loss distribution of portfolios comprising a large number of exposures. We concentrated on the study of corporate bond or loan portfolios, but our results can be generalized towards more

complicated financial instruments. The proposed approach builds further upon the factor model approach to credit risk as laid out in J.P. Morgan (1999). Using asymptotic distribution theory, we formally derived the credit loss distribution if the number of exposures in the portfolio gets large. The limiting distribution reveals that tail behavior of credit losses is highly influenced by the fit of the factor model regressions, where the fit is measured in terms of R^2 . Higher values of R^2 imply more systematic risk in the credit loss portfolio and, therefore, fatter tails of the credit loss distribution.

Using an extreme value theory perspective, we derived the tail indices of the credit loss distribution in one-factor and multi-factor models. Again we found that higher values of R^2 lead to smaller tail indices, i.e., fatter tails. If, however, one allows for heterogeneity in the fits of the factor model regressions, i.e., differing R^2 's over the exposures in the portfolio, we find that the smallest R^2 values ultimately dominate tail behavior. The exposures with the smallest R^2 have the relatively highest idiosyncratic risk components. As idiosyncratic risk can be diversified, the contribution of these exposures to the limiting distribution of credit losses is small. As a result, it becomes less likely that the limiting credit loss will hit the maximum loss, which in turn implies that the tail of the distribution is thinner. For multi-factor models, we found that the tail index is *ceterus paribus* at most as high as that of the one-factor model, implying a fatter tail for multi-factor models. These effects may be counterweighted by a drop in the maximum possible limiting loss for multi-factor models vis-a-vis one-factor models.

When applying our limiting result to stylized and empirically oriented portfolios, we found that perturbations in factor model fits (R^2) and portfolio quality (rating distribution) matter a great deal for credit risk. Quantiles of the limiting credit loss distribution can vary substantially depending on the choices made for these portfolio characteristics. Moreover, expressing portfolio credit risk in terms of standard deviations in excess of the expected credit loss produced a reversed ordering of riskiness compared to the direct measurement of credit loss quantiles. Thus, one should be very careful when interpreting the standard way of reporting credit losses in terms of standard deviations. From calculations based on representative empirical portfolios it further proved that the Basle 8% capital requirement seems a reasonable approximation to the 99.9% quantile of the credit loss distributions for typical portfolios of corporate rated bonds.

The derived limiting distribution is only interesting as far as it provides a good approximation to the credit loss distribution of a finite sized portfolio. We therefore studied the speed of convergence of the credit loss distribution for an increasing number of exposures and found that this convergence is rapid enough for most practical circumstances. The limiting distribution ap-

pears applicable if the portfolio comprises more than say 300 exposures. The degree of portfolio heterogeneity, however, plays a key role in the convergence behavior. If the portfolio contains a highly heterogeneous set of exposures (e.g., strongly differing loan sizes), the required number of exposures per portfolio might be higher for the limiting distribution to become applicable.

The paper suggests several important topics for future research. First, our approach can be extended to a dynamic setting. Second, the model can be used to obtain a complete assessment of portfolio risk, comprising both market risk and credit risk. Such an integrated credit-risk/market-risk management perspective would provide a valuable contribution to the current literature.

A Proofs

Proof of Theorem 1: For simplicity, we only provide the proof for $\pi(j, k_j, \ell_j, \psi) \equiv \pi(j, k_j, \ell_j)$. The proof for non-redundant ψ can be established similarly. Let T_n be the column vector with elements $S_1, \dots, S_n$, and summarize the regression equations for the $\mathbf{Sj}$ ($j = 1, \dots, n$) as $T_n = B_n f + \xi_n$ with $\xi_n = (\varepsilon_1, \dots, \varepsilon_n)^\top$, and $B_n^\top = (\beta_1, \dots, \beta_n)$. ξ_n has diagonal covariance matrix $\Sigma_n = \text{diag}(\omega_1, \dots, \omega_n)$.

Let h_n be defined as $h_n = E(f|T_n)$. Then the sequence (h_n) is a uniformly integrable martingale, which by the Martingale Convergence Theorem has an a.s. limit $h_\infty = E(f|S_1, S_2, \dots)$, see Doob (1953). Since for all $p \geq 1$ all moments $E|h_n|^p$ are bounded by $E|f|^p$, we also have that $E|h_n - h_\infty|^p \rightarrow 0$ (for $p \geq 1$). Moreover, since $\text{Cov}(f - h_\infty, h_\infty) = 0$ we also have the identity

$$\text{Cov}(f) = \text{Cov}(f - h_\infty) + \text{Cov}(h_\infty).$$

We now compute $\text{Cov}(h_n)$. As a first step we have $\text{Cov}(h_n) = \text{Cov}(f) - E \text{Cov}(f|T_n)$. Since $f, S_1, S_2, \dots$ is a Gaussian sequence, we have that

$$E \text{Cov}(f|T_n) = \text{Cov}(f|T_n) = \text{Cov}(f) - \text{Cov}(f, T_n) \text{Cov}(T_n)^{-1} \text{Cov}(T_n, f) = \\ \Omega_f - \Omega_f B_n^\top (B_n \Omega_f B_n^\top + \Sigma_n)^{-1} B_n \Omega_f = (\Omega_f^{-1} + B_n^\top \Sigma_n^{-1} B_n)^{-1}.$$

Hence $\text{Cov}(h_n) = \Omega_f - (\Omega_f^{-1} + B_n^\top \Sigma_n^{-1} B_n)^{-1}$. Notice that this gives an increasing sequence. Given the first part of Assumption 1, we obtain that $\text{Cov}(h_\infty) = \Omega_f$. But then $\text{Cov}(f - h_\infty) = 0$, and so $h_\infty = f$ with probability (w.p.) 1, because f and h_∞ have the same expectation.

Next we look at the conditional distribution of S_{n+1} given T_n . The distribution is Gaussian with mean $m_{n+1} = \beta_{n+1}^\top h_n$ and variance

$$s_{n+1}^2 = (\omega_{n+1} + \beta_{n+1}^\top \Omega_f \beta_{n+1}) - \beta_{n+1}^\top \Omega_f B_n^\top (B_n \Omega_f B_n^\top + \Sigma_n)^{-1} B_n \Omega_f \beta_{n+1} = \\ \omega_{n+1} + \beta_{n+1}^\top (\Omega_f^{-1} + B_n^\top \Sigma_n^{-1} B_n)^{-1} \beta_{n+1}.$$

Introduce the following notation. $\tilde{P}_n = \Omega_f^{-1} + B_n^\top \Sigma_n^{-1} B_n$ and $q_n = \beta_n^\top \tilde{P}_n^{-1} \beta_n$. Then

$$s_{n+1}^2 = \omega_{n+1} + \beta_{n+1}^\top (\tilde{P}_{n+1} - \frac{1}{\omega_{n+1}} \beta_{n+1} \beta_{n+1}^\top)^{-1} \beta_{n+1} = \omega_{n+1} \frac{1}{1 - \frac{q_{n+1}}{\omega_{n+1}}} > \omega_{n+1}.$$

Define $H_n = \tilde{P}_n - \Omega_f^{-1}$. We have $q_n/\omega_n = \beta_n^\top (n^{-1}\tilde{P}_n)^{-1}\beta_n/(n\omega_n)$. From parts 1 and 2 of Assumption 1, we then obtain $\frac{q_n}{\omega_n} - \frac{\beta_n^\top (n^{-1}H_n)^{-1}\beta_n}{n\omega_n} \rightarrow 0$ and $\frac{\beta_n^\top (n^{-1}H_n)^{-1}\beta_n}{n\omega_n} \rightarrow 0$. The latter follows from the convergence of H_n/n and $\beta_n^\top \beta/(n\omega_n)$ to a positive definite matrix and to zero, respectively. As a result, we establish that $q_n/\omega_n \rightarrow 0$ and $s_{n+1}^2/\omega_{n+1} \rightarrow 1$.

We now turn to the analysis of C_n . The first idea is to compensate C_n to form a martingale. So we look at

$$M_n = C_n - \sum_{j=1}^n \mathbb{E}(\pi(j, k_j, \ell_j) | T_{j-1}). \quad (\text{A1})$$

Let $1_{k,\ell}$ stand for the indicator of the event that a transition of firm j from its initial state (rating) k_j to ℓ takes place. Then we can write $\pi(j, k_j, \ell_j) = \sum_{\ell=1}^r \pi(j, k_j, \ell) 1_{k,\ell}$, and the conditional expectation $\sum_{j=1}^n \mathbb{E}(\pi(j, k_j, \ell_j) | T_{j-1})$ in (A1) becomes $\sum_{j=1}^n \sum_{\ell=1}^r \pi(j, k_j, \ell) \mathbb{E}(1_{k,\ell} | T_{j-1})$. By definition of the numbers $c_{k\ell}$ and the fact that (unconditionally) S_j has a normal distribution with mean zero and variance $\sigma_j^2 = \beta_j \Omega_f \beta_j + \omega_j$, such a transition takes place if S_j falls into the interval $(c_{k_j, \ell-1}\sigma_j, c_{k_j, \ell}\sigma_j)$, See Section 2. Since, conditional On T_{j-1} , S_j has a $N(m_j, s_j^2)$ distribution, we get the conditional probability

$$\begin{aligned} P_{j\ell} &:= \mathbb{E}[1_{k,\ell} | T_{j-1}] = P(c_{k_j, \ell-1}\sigma_j < S_j < c_{k_j, \ell}\sigma_j) \\ &= \Phi\left(\frac{c_{k_j, \ell}\sigma_j - m_j}{s_j}\right) - \Phi\left(\frac{c_{k_j, \ell-1}\sigma_j - m_j}{s_j}\right). \end{aligned} \quad (\text{A2})$$

The next step is to show that the martingale central limit theorem holds for M_n , or rather, that (A1) is of order $\sqrt{n}$ in probability. We compute the conditional variance $\Delta\langle M \rangle_j = \mathbb{E}[(M_j - M_{j-1})^2 | T_{j-1}]$. In order to get a decent and compact expression for this, it is convenient to introduce some additional notation. let $\pi_j = (\pi(j, k_j, 1), \dots, \pi(j, k_j, r))^\top$ and $P_j = (P_{j1}, \dots, P_{jr})^\top$. Then it is straightforward to show using elementary calculations that $\Delta\langle M \rangle_j = \pi_j^\top (\text{diag}(P_j) - P_j P_j^\top) \pi_j$. This is obviously bounded by $\pi_j^\top \pi_j$. From Assumption 1, we have that $\frac{1}{n} \sum_{j=1}^n \pi_j^\top \pi_j$ is bounded. Hence we obtain that M_n is of stochastic order at most $\sqrt{n}$, see Doob (1953).

The next object to study is the compensating term in (A1). We will show that it is of order n , so therefore (in view of the asymptotics of M_n) its asymptotic behavior coincides with that of C_n . We aim at showing that we can replace m_j and s_j in (A2) by $\beta_j^\top f$ and σ_j . Consider

$$\Phi\left(\frac{c_{k_j, \ell}\sigma_j - m_j}{s_j}\right). \quad (\text{A3})$$

Use the notation $\gamma_{j\ell} = c_{k_j, \ell}\sigma_j/\omega_j^{1/2}$, $\xi_j = \beta_j/\omega_j^{1/2}$, and $v_j = s_j/\omega_j^{1/2}$ to rewrite (A3) as

$$\Phi((\gamma_{j\ell} - \xi_j^\top h_{j-1})/v_j). \quad (\text{A4})$$

Notice that $v_j \geq 1$ and that $v_j \rightarrow 1$. We want to approximate (A4) by $\Phi(\gamma_{j\ell} - \xi_j^\top f)$. Therefore, we consider $\Phi((\gamma_{j\ell} - \xi_j^\top h_{j-1})/v_j) - \Phi(\gamma_{j\ell} - \xi_j^\top f)$. Let (h_j^*, v_j^*) denote a random point between (h_{j-1}, v_j) and $(f, 1)$, and let $t_j = (\gamma_{j\ell} - \xi_j^\top h_j^*)$. Moreover, let $\phi(\cdot)$ denote the density of the standard normal distribution. We now apply the mean value theorem, and obtain

$$\Phi((\gamma_{j\ell} - \xi_j^\top h_{j-1})/v_j) - \Phi(\gamma_{j\ell} - \xi_j^\top f) =$$

$$\phi(t_j/v_j^*)(-\frac{\xi_j^\top}{v_j^*}, -(\gamma_{j\ell} - \xi_j^\top h^*)/(v_j^*)^2) \begin{pmatrix} h_{j-1} - f \\ v_j - 1 \end{pmatrix}.$$

Rewrite this with $\Delta h_j = h_{j-1} - f$ and $\Delta v_j = v_j - 1$ as

$$-\phi(t_j/v_j^*)\frac{\xi_j^\top}{v_j^*}\Delta h_j - \phi(t_j/v_j^*)\frac{t_j}{v_j^*}\frac{1}{v_j^*}\Delta v_j.$$

Now consider the absolute value of each of the two summands. Since ϕ is bounded by 1 and $v_j^* \geq 1$, we obtain that the first summand is bounded by $|\xi_j^\top \Delta h_j|$. As $u\phi(u)$ is also bounded by 1, we moreover obtain that the second summand is bounded by $|\Delta v_j|$. Therefore, the absolute value of $\Phi((\gamma_{j\ell} - \xi_j^\top h_{j-1})/v_j) - \Phi(\gamma_{j\ell} - \xi_j^\top f)$ is bounded by $|\xi_j^\top \Delta h_j| + |\Delta v_j|$. Notice that this is independent of ℓ .

Write P_j^0 for the column vector with elements

$$P_{jl}^0 = \Phi\left(\frac{ck_{j,\ell}\sigma_j - \beta_j^\top f}{\omega_j^{1/2}}\right) - \Phi\left(\frac{ck_{j,\ell-1}\sigma_j - \beta_j^\top f}{\omega_j^{1/2}}\right).$$

Then

$$\begin{aligned} (P_{jl} - P_{jl}^0)^2 &\leq 4(|\xi_j^\top \Delta h_j| + |\Delta v_j|)^2 \leq 8(|\xi_j^\top \Delta h_j|^2 + |\Delta v_j|^2) \\ &\leq 8(\xi_j^\top \xi_j \Delta h_j^\top \Delta h_j + |\Delta v_j|^2). \end{aligned}$$

Hence

$$(P_j - P_j^0)^\top (P_j - P_j^0) \leq 8r(\xi_j^\top \xi_j \Delta h_j^\top \Delta h_j + |\Delta v_j|^2). \quad (445)$$

Now consider the difference of

$$n^{-1} \sum_{j=1}^n \mathbb{E}[\pi(j, k_j, l_j) | T_{j-1}] = n^{-1} \sum_{j=1}^n \pi_j^\top P_j$$

and $n^{-1} \sum_{j=1}^n \pi_j^\top P_j^0$, i.e., $n^{-1} \sum_{j=1}^n \pi_j^\top (P_j - P_j^0)$. Its absolute value is bounded by $[n^{-1} \sum_{j=1}^n \pi_j^\top \pi_j]^{1/2} [n^{-1} \sum_{j=1}^n (P_j - P_j^0)^\top (P_j - P_j^0)]^{1/2}$. The first factor is bounded by Assumption 1. We henceforth concentrate on the second factor. Since we know that $\Delta h_j \rightarrow 0$ w.p. 1 and $\Delta v_j \rightarrow 0$, we have that for any given $\eta > 0$ there is an a.s. finite random number N such that $|\Delta h_j| < \eta$ and $|\Delta v_j| < \eta$ if $j > N$ on a subset of the basic probability space with probability one. For each element of this subset we split $n^{-1} \sum_{j=1}^n (P_j - P_j^0)^\top (P_j - P_j^0)$ as $n^{-1} \sum_{j=1}^N (P_j - P_j^0)^\top (P_j - P_j^0) + n^{-1} \sum_{j=N+1}^n (P_j - P_j^0)^\top (P_j - P_j^0)$, because we can take n bigger than N . The first term obviously tends to zero if $n \rightarrow \infty$. Because of (A5), the second term is bounded by $n^{-1} \sum_{j=1}^n 8r(\xi_j^\top \xi_j + 1)\eta^2$, which is less than a constant times η^2 . The latter follows from Assumption 1 by noting that $\xi_j = \beta_j/\omega_j^{1/2}$. Because η is arbitrary, we can now conclude that $n^{-1} \sum_{j=1}^n \pi_j^\top (P_j - P_j^0) \rightarrow 0$ a.s. Now use the final part of Assumption 1.

Finally, note that

$$R_j^2 = \frac{\text{Cov}(S_j, \beta_j^\top f)^2}{\text{Cov}(S_j) \cdot \text{Cov}(\beta_j^\top f)} = \frac{\beta_j^\top \Omega_f \beta_j}{w_j + \beta_j^\top \Omega_f \beta_j},$$

such that

$$\frac{\beta_j^\top \Omega_f \beta_j}{\omega_j} = \frac{1}{1 - R_j^2} - 1 = \frac{R_j^2}{1 - R_j^2}.$$

This proves the theorem. $\blacksquare$

Proof of Theorem 2: Using Corollary 3.3.13 of Embrechts et al. (1997), it suffices to prove that

$$\lim_{C \uparrow 1} \frac{(1 - C) \cdot f(C)}{1 - F(C)} = \frac{1 - \rho^2}{\rho^2}, \quad (\text{A6})$$

with $F(\cdot)$ and $f(\cdot)$ the c.d.f. and p.d.f. of credit losses, respectively. Define $u_1 = 1 - C$. We can now rewrite (A6) as

$$\lim_{u_1 \downarrow 0} \frac{u_1 \cdot f(1 - u_1)}{1 - F(1 - u_1)}. \quad (\text{A7})$$

Using (19), we can rewrite (A7) as

$$\lim_{u_1 \downarrow 0} \frac{u_1 \cdot \sqrt{1 - \rho^2} \cdot \phi\left(\frac{\Phi^{-1}(u_1)\sqrt{1 - \rho^2} - c}{\rho}\right)}{\rho \cdot \Phi\left(\frac{\Phi^{-1}(u_1)\sqrt{1 - \rho^2} - c}{\rho}\right) \cdot \phi(\Phi^{-1}(u_1))}, \quad (\text{A8})$$

with $\Phi(\cdot)$ and $\phi(\cdot)$ the standard normal c.d.f. and p.d.f., respectively. Using the substitution $u_2 = \Phi^{-1}(u_1) \Leftrightarrow u_1 = \Phi(u_2)$, (A8) transforms into

$$\lim_{u_2 \rightarrow -\infty} \frac{\Phi(u_2) \cdot \sqrt{1 - \rho^2} \cdot \phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2} - c}{\rho}\right)}{\rho \cdot \Phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2} - c}{\rho}\right) \cdot \phi(u_2)}. \quad (\text{A9})$$

Now from equation (26.2.13) of Abramowitz and Stegun (1970), we have that for large negative u_2

$$\Phi(u_2) = \frac{\phi(u_2)}{|u_2|} (1 + o(u_2^{-1})). \quad (\text{A10})$$

Applying this result to (A9), we establish

$$\begin{aligned} \lim_{u_2 \rightarrow -\infty} \frac{\Phi(u_2) \cdot \sqrt{1 - \rho^2} \cdot \phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2} - c}{\rho}\right)}{\rho \cdot \Phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2} - c}{\rho}\right) \cdot \phi(u_2)} &= \\ \lim_{u_2 \rightarrow -\infty} \frac{\phi(u_2) \cdot \sqrt{1 - \rho^2} \cdot \phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2}}{\rho}\right) \left|\frac{u_2 \cdot \sqrt{1 - \rho^2}}{\rho}\right|}{|u_2| \cdot \rho \cdot \phi\left(\frac{u_2 \cdot \sqrt{1 - \rho^2}}{\rho}\right) \cdot \phi(u_2)} &= \frac{1 - \rho^2}{\rho^2}, \end{aligned} \quad (\text{A11})$$

which proves the theorem. $\blacksquare$

References

- Abramowitz, M. and I. Stegun (1970). **Handbook of Mathematical Functions**. New York: Dover.
- Basle Committee on Bank Supervision (1988). International convergence of capital measurement and capital standards. Report 4, Bank of International Settlements, Basle.
- Basle Committee on Bank Supervision (1996). Amendment to the capital accord to incorporate market risks. Report 24, Bank of International Settlements, Basle.
- Belkin, B., S. Suchower, and L. Forest (1998a). The effect of systematic credit risk on loan portfolio value-at-risk and loan pricing. *CreditMetrics Monitor* (1st Quarter), 17-28. <http://www.creditmetrics.com>.
- Belkin, B., S. Suchower, and L. Forest (1998b). A one-parameter representation of credit risk transition matrices. *CreditMetrics Monitor* (3rd Quarter), 46-56. <http://www.creditmetrics.com>.
- Black, F. and J. Cox (1976). Valuing corporate securities: some effects of bond indenture provisions. *Journal of Finance* **31**, 351-367.
- Carey, M. (1998). Credit risk in private debt portfolios. *Journal of Finance* **53**(4), 1363-1387.
- Credit Suisse (1999). *CreditRisk⁺*. Downloadable: <http://www.csfp.csh.com>.
- Danielsson, J. and C. D. Vries (1997). Value-at-risk and extreme returns. mimeo.
- Doob, J. (1953). **Stochastic Processes**. New York: Wiley.
- Duffee, G. (1996). On measuring credit risks of derivative instruments. *Journal of Banking & Finance* **20**, 805-833.
- Embrechts, P., C. Klippelberg, and T. Mikosch (1997). **Modeling Extremal Events**, Volume **33** of **Applications of Mathematics; Stochastic Modelling and Applied Probability**. Heidelberg: Springer Verlag.
- Fons, J. (1991, August). An approach to forecasting default rates. Moody's Special Report.
- Ibragimov, I. and Y. Linnik (1971). **Independent and stationary sequences of random variables**. Groningen: Wolters Noordhof.
- IIF (1998, March). Recommendation of revising the regulatory capital rules for credit risk. Report of the Working Group on Capital Adequacy of the Institute of International Finance.
- ISDA (1998, March). Credit risk and regulatory capital. International Swaps and Derivatives Association.
- Jónsson, J. and M. Fridson (1996, June). Forecasting default rates on high-yield bonds. *Journal of Fixed Income*, 69-77.
- J.P. Morgan (1999). *CreditMetrics* (4th ed.). <http://www.creditmetrics.com>.
- Kealhofer, S. (1995). Managing default risk in derivative portfolios. In **Derivative Credit Risk: Advances in Measurement and Management**. London: Risk Publications.
- Koyluoglu, H. and A. Hickman (1998). Reconcilable differences. *Risk* **56**(October), 56-62.
- Longstaff, F. and E. Schwartz (1995). A simple approach to valuing risky and floating rate debt. *Journal of Finance* **50**, 789-819.

- Markowitz, H. (1952). Portfolio selection. ***Journal of Finance* 7**, 77-91.
- Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates. ***Journal Of Finance* 29**, 449-470.
- Sharpe, W. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance* **19**, 429-442.
- S&P (1998). Ratings performance 1997: stability and transition. Report.
- Wilson, T. (1997a, September). Portfolio credit risk, part I. ***Risk***, 111-117.
- Wilson, T. (1997b, October). Portfolio credit risk, part II. ***Risk***, 56-61.
- Zhou, C. (1997). Default correlation: an analytical result. Technical report, Federal Reserve Board.