



ELSEVIER

Journal of Banking & Finance 25 (2001) 1635–1664

Journal of
BANKING &
FINANCE

www.elsevier.com/locate/econbase

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

André Lucas ^{a,b}, Pieter Klaassen ^{a,c}, Peter Spreij ^d,
Stefan Straetmans ^{a,*}

^a *Department of Finance and Financial Sector Management, Vrije Universiteit, De Boelelaan 1105, NL-1081 HV Amsterdam, Netherlands*

^b *Tinbergen Institute Amsterdam, Keizersgracht 482, NL-1017 EG Amsterdam, Netherlands*

^c *ABN-AMRO Bank NV, Financial Markets Risk Management, P.O. Box 283, NL-1000 EA Amsterdam, Netherlands*

^d *Korteweg-de Vries Institute, University of Amsterdam, Plantage Muidergracht 24, NL-1018 TV Amsterdam, Netherlands*

Received 29 April 1999; accepted 10 June 2000

Abstract

We derive an analytic approximation to the credit loss distribution of large portfolios by letting the number of exposures tend to infinity. Defaults and rating migrations for individual exposures are driven by a factor model in order to capture co-movements in changing credit quality. The limiting credit loss distribution obeys the empirical stylized facts of skewness and heavy tails. We show how portfolio features like the degree of systematic risk, credit quality and term to maturity affect the distributional shape of portfolio credit losses. Using empirical data, it appears that the Basle 8% rule corresponds to quantiles with confidence levels exceeding 98%. The limit law's relevance for credit risk management is investigated further by checking its applicability to portfolios with a finite number of exposures. Relatively homogeneous portfolios of 300 exposures can be well approximated by the limit law. A minimum of 800 exposures is required if portfolios are relatively heterogeneous. Realistic loan portfolios often contain thousands of exposures implying that our analytic approach can be a fast and accurate

* Corresponding author. Tel.: +31-20-44-46039; fax: +31-20-44-46005.

E-mail addresses: alucas@econ.vu.nl (A. Lucas), pieter.klaassen@nl.abnamro.com (P. Klaassen), spreij@wins.uva.nl (P. Spreij), sstraetmans@econ.vu.nl (S. Straetmans).

alternative to the standard Monte-Carlo simulation techniques adopted in much of the literature. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: G21; G33; G29; C19

Keywords: Credit risk; Factor model; Fat tails; Skewness; Asymptotic analysis; Nonnormal factors

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 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 capital requirements (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 the Capital Asset Pricing Model of Sharpe (1964). 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 Basle rules of 1988 ignore portfolio effects and levy 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. See also the new proposals to revise the 1988 capital accord, Basle Committee on Bank Supervision (1999).

The general characteristics of the credit risk loss distribution resulting from portfolio models are badly understood. It is often observed that this distri-

bution is heavy-tailed and skewed, but the prominence of these properties seems to depend on the composition of the specific portfolio under consideration. The precise mechanism linking portfolio characteristics to skewness and heavy tails is relatively unexplored. In this paper, we derive an efficient analytic approximation to the loss distribution if the portfolio contains a large number of exposures. This approximation enables us to study the sensitivity of the loss distribution, and in particular the credit loss tail, to exposure characteristics like credit quality, the degree of systematic risk, and the maturity profile. The loss profile of portfolios with a realistic diversity in portfolio characteristics can be approximated fairly well, even if they are medium-sized, containing say 300–800 exposures. Compared to Monte-Carlo simulation of portfolio losses, our approximation is much more efficient in practice.

The credit quality of individual exposures is usually expressed by their ratings. These can be assigned either by external rating agencies such as Standard & Poor's and Moody's, as is the case in *CreditMetrics* of J.P. Morgan (1999), or by banks internally. Each debtor's rating is associated with a certain probability of default. An alternative route for estimating default probabilities constitutes the option-theoretic framework 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 in order to assign default probabilities to debtors, provided these have equity quoted on a stock exchange, Kealhofer (1995).

Co-movements in credit quality changes for different debtors may be triggered by some common, underlying macro-factors. For example, default rates tend to rise during recessions, 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 diversification decreases the likelihood of extreme portfolio credit losses. The correlation effect on the shape of the loss distribution also depends on the initial credit quality of the portfolio. Zhou (1997) shows that 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. Our numerical results confirm that for a given correlation a higher portfolio quality lowers extreme credit loss quantiles.

To our knowledge, only Carey (1998) and Gordy (1999) have thus far performed a systematic study of the tails of the credit loss distribution. The former study uses historical data on exposures and credit losses stemming from private placements by US life insurers. Exposures are sampled from this large database to obtain portfolios with certain characteristics in terms of initial credit quality. The actual loss experience is then analyzed for these

sampled portfolios. Carey's study yields insight 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. Gordy (1999) studies credit loss tails for realistic portfolios using two competing portfolio models, namely *CreditMetrics* and *CreditRisk⁺*. He focusses on default losses only. We extend his results by incorporating losses due to credit rating migrations. Moreover, we study the effect of bond maturity on extreme quantiles and present analytical as well as numerical results on tail fatness in relation to portfolio characteristics.

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 *CreditPortfolioView* of McKinsey (Wilson, 1997a,b). Despite the apparent differences between these approaches, they exhibit a common underlying framework, see Koyluoglu and Hickman (1998) and Gordy (1999).

The paper is set up as follows. Our model, which follows the Koyluoglu and Hickman (1998) set-up, is described in Section 2. In Section 3, we derive the asymptotic loss distribution and its salient features. Section 4 investigates the tail shape of the limiting distribution.

In Section 5, we study how the asymptotic loss distribution is altered when the characteristics of the underlying exposures are altered. The 'stylized' portfolios we consider mimic typical bank portfolios in terms of: (i) the distribution of exposures over initial credit ratings, (ii) the sizes of exposures per rating category, and (iii) the level and variability in systematic risk across exposures. For a typical high quality corporate bond portfolio, we find that the standard 8% capital charge from the Basle 1988 Accord corresponds to a 98% or higher confidence quantile. It is also shown that the location of high-confidence quantiles like the 99.9% quantile is quite sensitive to the choice of distribution of the common risk factors.

Section 6 questions the number of exposures required to render the asymptotic loss distribution a good approximation to the actual loss distribution. In studying the convergence properties, we especially pay attention to the tail behavior of the distribution. It is shown that the speed of convergence of the actual credit loss quantiles to their analytic counterparts is influenced by the degree of portfolio heterogeneity. For relatively homogeneous portfolios, the approximation is already accurate for portfolios with 300 exposures or more. Relatively heterogeneous portfolios should contain at least 800 exposures in order to get a similar accuracy. Section 7 concludes the paper. Proofs of all results in this paper are available upon request.

2. Theoretical framework

Consider a credit portfolio that consists of exposures to n companies. Each company j is characterized by a four-dimensional stochastic vector

$$(S_j, k_j, \ell_j, \pi(j, k_j, \ell_j)). \tag{1}$$

The first element S_j represents a latent variable triggering defaults or rating migrations. A prime candidate for S_j is the company’s surplus, i.e., the difference between the market value of assets and liabilities. Default may then be defined as a situation in which the surplus falls below a certain threshold. We assume that the companies’ surplus variables are both driven by economy-wide features (business cycle conditions, stock market fluctuations) and firm-specific circumstances. This justifies the specification of a factor model for the surplus variable

$$S_j = \mu_j + \beta_j^T f + \varepsilon_j, \tag{2}$$

where $\mu_j \in \mathbb{R}$ is a constant term, $\beta_j \in \mathbb{R}^m$ a vector of factor loadings, $f \in \mathbb{R}^m$ the vector of common factors, and $\varepsilon_j \in \mathbb{R}$ is a scalar representing idiosyncratic shocks. This set-up follows the model structure of, e.g., J.P. Morgan (1999). Without loss of generality, we set $\mu_j = 0$ for all j . For the time being, we assume that the factor vector f and innovations ε_j (for all j) are normally distributed with zero-mean. They have covariance matrix Ω_f and variance ω_j , respectively. We further assume that the idiosyncratic shocks are independent across firms and that f and $\{\varepsilon_j\}_{j=1}^\infty$ are independent. The proof of the limit law for portfolio credit losses does not hinge upon the zero mean and normality assumption. Thus, the limit law can be easily generalized. In Section 5, we investigate the consequences of relaxing the normality assumption on the analytic credit loss quantiles.

The above model structure implies that the surplus variables are correlated across firms. As the S_j s also trigger the default mechanism, correlation between the S_j s results in correlated default probabilities. In Sections 3 and 4, we show that this correlation induces the credit loss distribution to be heavy-tailed, notwithstanding the normality assumption for the company’s surplus. A company’s initial and end-of-period ratings are represented in (1) by k_j and ℓ_j , respectively. We assume r rating categories, such that $k_j, \ell_j \in \{1, \dots, r\}$. We further assume that migrations are driven by a Markovian transition matrix P ,

$$P = \begin{pmatrix} p_{11} & \dots & p_{1r} \\ \vdots & \ddots & \vdots \\ p_{r1} & \dots & p_{rr} \end{pmatrix}, \tag{3}$$

where $p_{k\ell}$ denotes the probability that a firm with initial rating k switches to rating ℓ over the time horizon considered (one year in this paper). For sim-

plicity, we assume that the transition matrix is common to all firms. Note that $P\mathbf{1}_r = \mathbf{1}_r$, where $\mathbf{1}_r$ is an r -dimensional vector of ones. By setting $p_{r1} = \dots = p_{r,r-1} = 0$ and $p_{rr} = 1$, we can identify the r th rating category as the state of default. Numerous authors already calculated historical averages for the migration probabilities in (3). We use the estimates provided by J.P. Morgan (1999) on their website, see Section 5 for the numerical values.

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

$$\Phi(s_{k,\ell-1}) - \Phi(s_{k\ell}) = p_{k\ell} \quad \text{for all } k \text{ and } \ell = 1, \dots, r, \tag{4}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function (c.d.f.). The end-of-period rating ℓ_j of company j (with initial rating k_j) is such that

$$s_{j,k_j,\ell_j} \equiv s_{k_j,\ell_j} \sqrt{\omega_j + \beta_j^\top \Omega_f \beta_j} < S_j \leq s_{k_j,\ell_j-1} \sqrt{\omega_j + \beta_j^\top \Omega_f \beta_j} \equiv s_{j,k_j,\ell_j-1}. \tag{5}$$

This is illustrated in Fig. 1. The support of the normal distribution of the surplus S_j is partitioned by means of the constants $s_{k\ell}$ 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 company’s initial rating. For example, highly rated companies (AAA) are less likely to default than low-graded ones (CCC), such that from (4) $s_{k,r-1}$ must be higher for bonds with an initial CCC

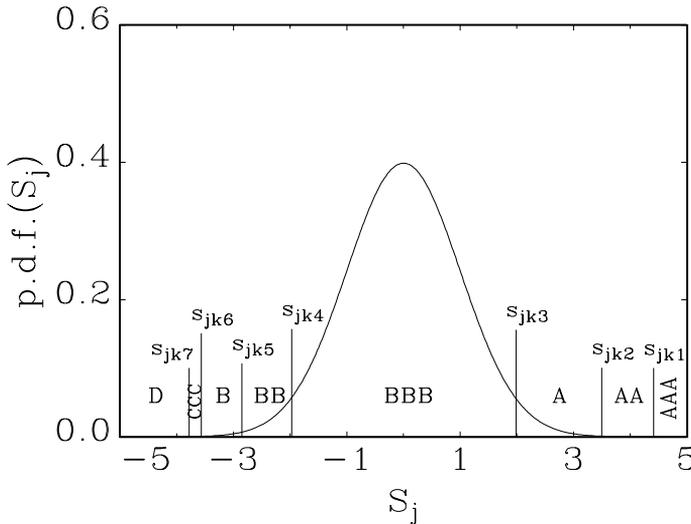


Fig. 1. Relation of the random variable S_j and the end-of-period rating ℓ_j in an eight category rating system. The initial rating of the company is $k_j = \text{BBB}$. D denotes default, while $s_{k\ell}$ equals $s_{k\ell}(\omega_j + \beta_j^\top \Omega_f \beta_j)^{1/2}$.

than with an initial AAA rating. The fact that S_j determines the end-of-period rating ℓ_j makes ℓ_j a stochastic variable in the present set-up.

Note that the described model set-up is static in the sense that the transition matrix is always nonrandom and the rating migration mechanism is solely dictated by the factor model (2) and the bins (5). Extending the model to a dynamic setting requires a random transition matrix. For example, one could assume that the vector of common factors follows an autoregressive process of order 1:

$$f_t = \Psi f_{t-1} + \eta_t, \quad (6)$$

with η_t independent of ε_s for all s, t . By the dependence of f_t on its own past, the credit quality is randomized over time: defaults become more likely during, e.g., recessions whereas upgrades may prevail in expansionary periods. These effects can be captured by an appropriate choice of f_t . In that case, the conditional (given the history of f_t up to time t) expectation of f_{t+1} can be negative in times of recession and positive in expansionary periods. For the sake of clarity, we will not deal with dynamics in this paper. More discussion on stochastic migration rates can be found in, e.g., Credit Suisse (1999) and Belkin et al. (1998b). Detailed simulations for multi-period models are left for further research, see also Wilson (1997a,b).

A firm's credit loss $\pi(\cdot)$ constitutes the final characteristic in (1). We assume that the amount of credit loss depends on the loan's initial (k_j) and final (ℓ_j) rating category. This is expressed by specifying $\pi(j, k_j, \ell_j)$ as a function of k_j and ℓ_j . By definition, a credit loss occurs either if a firm defaults or if the firm's rating deteriorates. The latter is due to differing credit risk spreads across rating categories, maturities, and industries. This justifies the dependence of the credit loss on the initial and end-of-period rating.

Using all the above definitions and symbols, the portfolio credit loss simply is the sum of the n individual credit losses:

$$C_n = \sum_{j=1}^n \pi(j, k_j, \ell_j). \quad (7)$$

Again note that this is a stochastic variable due to the dependence of the ℓ_j s on the random variables S_j .

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, we need to make the following assumption.

Assumption 1. $\sup_{j \geq 1} E[\pi(j, k_j, \ell_j)^2 | f] < \infty$ (a.s.).

Assumption 1 states that the conditional expectations of individual squared losses are bounded uniformly (almost surely). Most financial instruments satisfy this technical condition. For bonds the condition is trivially satisfied because the maximum loss is the (discounted) value of the bond, which is finite. The assumption can also be met under fairly general conditions for more complicated instruments like derivatives with unbounded $\pi(\cdot)$. For example, let $\pi(\cdot)$ be the payoff on the pay-fixed end of a plain vanilla swap. In that case, $\pi(\cdot)$ is unbounded in the interest rate, which we can incorporate in f . When interest rates rise, the pay-fixed end of the swap becomes more valuable. Conditional on f , however, $\pi(\cdot)$ is bounded such that Assumption 1 is met. Note that individual credit losses derive their randomness from the end-of-period ratings $\ell_j \in \{1, \dots, r\}$.

The following limit law constitutes the heart of the paper.

Theorem 1. *Define*

$$R_j^2 = \frac{\beta_j^\top \Omega_f \beta_j}{\omega_j + \beta_j^\top \Omega_f \beta_j} \tag{8}$$

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

$$v_j^\top = \frac{\beta_j^\top \Omega_f^{1/2}}{\sqrt{\beta_j^\top \Omega_f \beta_j}}, \tag{9}$$

such that $v_j^\top v_j = 1$. Define $Y = \Omega_f^{-1/2} f$. Let

$$\Phi_{j\ell} = \Phi \left(\frac{s_{k_j, \ell-1} - \sqrt{R_j^2} v_j^\top Y}{\sqrt{1 - R_j^2}} \right) - \Phi \left(\frac{s_{k_j, \ell} - \sqrt{R_j^2} v_j^\top Y}{\sqrt{1 - R_j^2}} \right) \tag{10}$$

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

$$B_n = E[C_n | f] = \sum_{j=1}^n E[\pi(j, k_j, \ell_j) | f] = \sum_{j=1}^n \sum_{\ell=1}^r \Phi_{j\ell} \pi(j, k_j, \ell). \tag{11}$$

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

$$n^{-1} C_n - n^{-1} B_n \xrightarrow{\text{a.s.}} 0, \tag{12}$$

with C_n the portfolio credit loss as defined in (7) and $\xrightarrow{\text{a.s.}}$ denoting almost sure convergence.

Note that C_n in (12) is the average credit loss, whereas B_n is the average conditional (on f) expectation of credit losses, see (11). Therefore, C_n depends on $(f, \varepsilon_1, \varepsilon_2, \dots)$, but B_n depends on f only, see also Finger (1999). By using B_n rather than C_n , we effectively average out all idiosyncratic risk just as in the case of linear portfolio theory.

To illustrate the limit law, consider a bond portfolio where the migration mechanism is triggered by a single factor ($m = 1$), where systematic risk is equal across the exposures ($R_j^2 \equiv \rho^2, \rho \geq 0$), and where $v_j \equiv 1$. We allow for only two end-of-period ratings $\ell = 1, 2$, the second rating category corresponding to the default state. Moreover, assume that the exposure on each firm has a face value of unity and that everything is lost in case of default. A similar model is studied in Belkin et al. (1998a). Thus, the credit losses in the two possible end-of-period rating states boil down to $\pi(j, k_j, 1) \equiv 0$ and $\pi(j, k_j, 2) \equiv 1$, respectively. Given these assumptions, B_n as defined in (11) simplifies to

$$B_n = \sum_{j=1}^n \Phi\left(\frac{s - \rho Y}{\sqrt{1 - \rho^2}}\right), \tag{13}$$

such that

$$C_n/n - \Phi\left(\frac{s - \rho Y}{\sqrt{1 - \rho^2}}\right) \xrightarrow{\text{a.s.}} 0, \tag{14}$$

where s equals the default threshold for the surplus variable.

Fig. 2 plots the p.d.f. of the limiting credit loss $C = \lim_{n \rightarrow \infty} C_n/n$ for various values of ρ and s . We consider values for s corresponding with default probabilities of 5% and 1%, respectively. The figure reveals that large credit losses become more likely for higher correlations between the underlying surplus

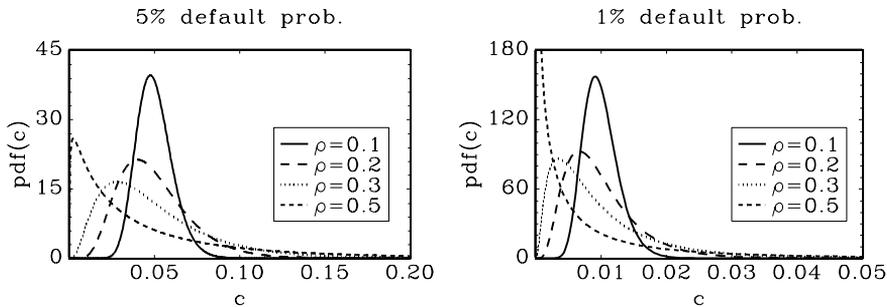


Fig. 2. Asymptotic default loss densities for constant $R_j^2 \equiv \rho^2$ and $v_j \equiv 1$ in a one-factor model ($m = 1$), see Theorem 1. There are only two rating categories, one of which corresponds to a state of default. The constant s is chosen such that there is either a 5% (left-hand panel)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

variables S_j . This feature seems to be independent of the choice of default probability as specified through s . We also note that for smaller values of ρ the credit loss distribution becomes more concentrated and bell-shaped. If $\rho = 0$, the p.d.f. degenerates, i.e., the portfolio credit loss converges almost surely towards the expected credit loss when the number of bonds in the portfolio grows indefinitely.

Knowledge of the limit law’s analytic expression enables the credit risk manager to calculate the loss distribution’s higher moments and quantiles without resorting to simulation, at least for the one-factor case. Table 1 presents the moments of the credit loss distributions shown in Fig. 2. The standard deviation, skewness, and kurtosis are all increasing in the degree of systematic risk ρ . Moreover, skewness and kurtosis increasingly deviate from their Gaussian values of 0 and 3 for higher values of ρ .

In order to calculate quantiles of C for the simplified one-factor example, first note that B_n is monotonically decreasing in Y . Take $g(Y) = \lim_{n \rightarrow \infty} B_n/n$, then using the transformation-of-variables technique, see e.g., Hogg and Craig (1970), the probability density of credit losses c boils down to

$$\frac{\phi(g^{-1}(c))}{|g'(g^{-1}(c))|}, \tag{15}$$

with ϕ denoting the standard normal density and $g^{-1}(\cdot)$ and $g'(\cdot)$ denoting the inverse and the first derivative of $g(\cdot)$, respectively. The trapezoid rule for numerical integration

$$\frac{1}{2} \sum_{j=1}^N [g(y_j)^\kappa \phi(y_j) + g(y_{j-1})^\kappa \phi(y_{j-1})] (y_j - y_{j-1}), \tag{16}$$

provides an easy approximation of, for example, the expected credit loss ($\kappa = 1$) or higher order moments, where $-K = y_0 < y_1 < \dots < y_N = K$ denotes an appropriate partitioning of the interval $[-K, K]$ for a sufficiently large constant $K > 0$. The following chain of equalities shows that the computation

Table 1
Moments of credit loss distributions^a

ρ	$s = \Phi^{-1} (5\%)$				$s = \Phi^{-1} (1\%)$			
	μ	σ	Skew	Kurt	μ	σ	Skew	Kurt
0.1	5	1.0	0.50	3.40	1	0.3	0.71	3.87
0.2	5	2.1	1.01	4.62	1	0.6	1.52	6.98
0.3	5	3.3	1.54	6.73	1	0.9	2.49	13.93
0.5	5	6.1	2.70	13.77	1	1.8	5.10	47.47

^a The table contains the expected (percentage) credit loss (μ) and its standard deviation (σ) as well as the skewness and kurtosis of the credit loss distribution for different default probabilities and different values of ρ . Moments are computed using numerical integration based on (16).

of credit loss quantiles in this simple set-up is even easier than calculating the distributional moments:

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

Hence, the δ -quantile can be obtained by a simple evaluation of $g(\cdot)$ in one point. Analytic quantile calculations for multi-factor models are somewhat more cumbersome, but there is still an advantage over pure simulation. Indeed, the probability $P(C \leq c) = P(g(Y) \leq c)$ can be evaluated by integrating over Y , either by using numerical integration if Y is low-dimensional, or by simulation if Y is high-dimensional. The latter may still be more efficient than full-fledged simulation of the credit loss distribution because we do not need to simulate the idiosyncratic shocks.

Before closing this section we focus upon some additional properties of the limit law. First, note that the expression B_n/n in Theorem 1 no longer depends on the idiosyncratic risk factors ε_j , but only on systematic risk f . A similar result is well known in *linear* portfolio theory. Indeed, within the CAPM model, only systematic risk persists when the number of assets increases. Theorem 1 generalizes this feature to the nonlinear context of credit risk management.

A second point to note is that we have expressed B_n/n in terms of (v_j, R_j^2) rather than (β_j, ω_j) . This mimics the conversion from Cartesian coordinates to polar coordinates. The R_j^2 directly reflect 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. On the other hand, if R_j^2 approaches zero the stochastic vector Y drops out of the j th term of B_n , making the j th term nonrandom. The vector v_j , in contrast, determines the directional sensitivity (rather than the magnitude) of the j th surplus variable with respect to the systematic factors. As already noted, this vector has unit length. 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.

Finally, we like to stress that the proof of Theorem 1 does not hinge upon the normality assumption for Y . Risk managers might have some idea about the future development of Y in terms of its forecast distribution. The latter can then be used in (12) to obtain simulations from the credit loss distribution that are more relevant from an economic perspective. Alternatively, risk managers might be interested in the effect of specific distributional assumptions for Y , e.g., stress scenarios, in which case Y places discrete (or even unit) mass on certain scenarios. Such distributional assumptions can be readily incorporated to obtain simulations for credit losses that are relevant for the purpose at hand.

In Sections 5 and 6, we study the sensitivity of the analytic credit loss quantiles to using the fat-tailed Student t rather than the normal distribution¹ for f .

4. Tail behavior of credit losses

In Section 3, we found that a higher correlation between default risks increases the likelihood of extreme portfolio credit losses, see also Fig. 2. 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 portfolio credit losses still exhibit more probability mass in the tails than a normal distribution with identical mean and variance, see also Table 1. Distributions with this pattern of decay in tail probabilities are also called heavy-tailed or fat-tailed distributions. The tails of the derived limit law decline to zero at a lower than exponential rate. In this section, we investigate this rate of decay using the statistical theory of extremes, see Embrechts et al. (1997) for a nice introduction into extreme value theory with applications in finance. We also reflect on the relationship between the rate of decay and the systematic risk in the portfolio. A correct assessment of the tail behavior is important for risk management. Indeed, if portfolio credit losses exhibit heavy tails, common rules of thumb for computing loss quantiles based on the normal paradigm no longer apply. For example, the 99.9% percentile may lie much more than three standard deviations above the expected loss, which is the number one would expect for the normal distribution, see Section 5.

The remainder of this section is built up as follows. We start in Section 4.1 by commenting on the main result from statistical extreme value theory that we use in this section. We explain its relevance for credit losses. Next, in Section 4.2 we show for models with homogeneous $v_j \equiv v$ (see Theorem 1) that the credit loss distribution has polynomially rather than exponentially declining tails. This limits the applicability of the normal distribution as an approximation to credit loss distributions. We also link the rate of tail decay (the so-called tail index) to the portfolio's characteristics. Exposures with the highest idiosyncratic risk components dominate the extreme tail behavior of credit losses and imply a high rate of tail decay. Less far out in the tails, however, exposures with relatively less idiosyncratic risk may also influence the tail

¹ (i) Note that alternative distributional assumptions for Y require changes in the constants s_{jkl} in order for the unconditional rating migration probabilities to correspond to their long-term averages from Table 2. (ii) The *Creditrisk*⁺ framework uses a Gamma-distribution for f . Checking the sensitivity of our results with respect to the Gamma assumption is difficult in our framework, however, because we consider losses due to *both* defaults and credit rating migrations. It is then much less straightforward to synchronize the frameworks of *CreditMetrics* and *Creditrisk*⁺ as in Gordy (1999).

Table 2
Rating migration probability matrix and credit spreads^a

k	ℓ										
	AAA	AA	A	BBB	BB	B	CCC	D	1Y	5Y	10Y
AAA	9082	826	74	6	11	0	0	0	58	77	125
AA	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
CCC	21	0	41	124	267	1170	6448	1929	435	585	980
D	0	0	0	0	0	0	0	10,000	–	–	–
								Base yield	429	409	429

^aThe table contains the probability (in basis points) of a credit rating migration from category k to ℓ over a one-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* website, October 1998.

behavior significantly. In Section 4.3, we generalize our results to models with heterogeneous v_j . It turns out that such models (*ceteris paribus*) give rise to higher rates of tail decay. Note that the analytic results of the present section only concern the rate of tail decay. As will be explained in Section 4.1, characterizing the complete tail behavior analytically using extreme value theory is beyond the scope of this paper. Due to the detailed nature of the results, the present section is somewhat more technical than the remainder of this paper.

4.1. Extreme value theory

Extreme value theory is concerned with the extreme tail behavior of statistical distributions. The main result we use in this paper is that the credit loss distribution has a tail expansion of the form

$$1 - F(c) = (\bar{c} - c)^\alpha L[(\bar{c} - c)^{-1}], \quad (18)$$

where \bar{c} is the maximum credit loss, $F(\cdot)$ the credit loss distribution, and $L(\cdot)$ is a slowly varying function, i.e., $\lim_{x \rightarrow \infty} L(tx)/L(x) = 1$ for $t > 0$. Typical examples of slowly varying functions are $L(x) = k$ for some constant k , or $L(x) = \ln(x)$. Eq. (18) shows that near the maximum credit loss \bar{c} the dominant factor of $F(\cdot)$ is $(\bar{c} - c)^\alpha$. The larger α , the faster the tail decays to zero. By inspecting α , we thus get an indication about the tail shape of the credit loss distribution for large credit losses. In this section, we show how the tail index α is related directly to the exposures' characteristics. Note, however, that α determines the tail behavior only in part, albeit the dominant part. In order to derive the complete tail behavior of $F(c)$ by means of extreme value theory, we also need an explicit characterization of the slowly varying function $L(\cdot)$. This, however, is very technical and beyond the scope of the present paper. In this section, we only focus on the rate of decay of the tails, i.e., on α . The reader should bear in mind that for two credit loss distributions $F_1(\cdot)$ and $F_2(\cdot)$ with tail indices α_1 and α_2 , respectively, we may well have for some fixed value c that $1 - F_1(c) < 1 - F_2(c)$ even though $\alpha_1 < \alpha_2$. This can be due to a larger value of the slowly varying function $L((\bar{c} - c)^{-1})$ for distribution 2 at c .

4.2. Portfolios with homogeneous v_j

First, we assume that the exposures are characterized by the same directional sensitivity vector $v_j \equiv v$, but possibly different R_j s. As $\|v\| = 1$, note that $v'Y$ in (10) is also standard normally distributed for all j . Therefore, without loss of generality we consider a one-factor model only ($m = 1$) and set $v = 1$. Assume that the portfolio can be divided into two homogeneous groups. Each group is characterized by a different value of R_j^2 , namely \hat{R}_i^2 for group $i = 1, 2$. The arguments presented here carry over directly to situations with an arbitrary number of groups. For expositional purposes, we restrict attention here

and in the following section to a setting with two rating categories only: defaulted or not defaulted. The results are equally valid, however, in settings with multiple rating categories. The following theorem summarizes the tail behavior for the sketched setting:

Theorem 2. *Consider the one-factor model with $v_j \equiv 1$. Assume that a fraction $\lambda \in (0, 1)$ of the firms has $R_j^2 = \hat{R}_1^2$, while the remaining firms have $R_j^2 = \hat{R}_2^2$. Then the credit loss distribution has a tail expansion as in (18) with tail index*

$$\max_{i \in \{1,2\}} (1 - \hat{R}_i^2) / \hat{R}_i^2. \quad (19)$$

Clearly, the credit loss distribution has algebraically rather than exponentially declining tails, as $\alpha < \infty$ for $\hat{R}_i^2 > 0$, see (18). For illustration, first consider the special case of a completely homogeneous portfolio. In that case, $\hat{R}_1^2 = \hat{R}_2^2 = \rho^2$, such that $\alpha = (1 - \rho^2) / \rho^2$. It is then easily seen that increases in the degree of systematic risk (ρ^2) cause a lower tail index, thus decreasing the rate of tail decay. Using (19), we can explain the empirical stylized fact of fat-tailed credit loss distributions from a micro-based approach to individual exposures by allowing for common factors. Even if these exposures are driven by normally distributed (and thus thin-tailed) systematic and idiosyncratic shocks, the portfolio credit loss distribution exhibits polynomially declining tails. Consequently, one should be very careful in using the normal distribution as an approximation to the credit loss distribution. This holds especially if one is interested in extreme quantiles, e.g., the 99.9th percentile, see also Section 5.

Theorem 2 also reveals that the exposures with the highest idiosyncratic risk components, i.e., the smallest \hat{R}_i^2 , dominate the extreme tail behavior and imply a high rate of tail decay. The intuition for this is as follows. The limiting credit loss distribution from Theorem 1 only uses systematic risk, as the idiosyncratic risk can be diversified in large portfolios. If the portfolio now contains a significant number of exposures with a dominant idiosyncratic risk component, i.e., low \hat{R}_i^2 , these exposures are less easily jointly pushed into default using movements in systematic risk only. Consequently, credit losses near the maximum credit loss, which need to be triggered by both groups entering default simultaneously, become less likely. The tail close to this maximum credit loss will therefore be thinner.

4.3. Portfolios with heterogeneous v_j

We now turn to models with varying directional factor sensitivities v_j , including both multi-factor and one-factor models. We restrict our attention to portfolios that can be divided into a finite number n^* of homogeneous groups, cf. the previous subsection. Group i is characterized by a combination $(\hat{R}_i^2, \hat{v}_i, \hat{\pi}_i)$, $i = 1, \dots, n^*$.

To state the theorem, we need some additional notation. Let $G \subset \{1, \dots, n^*\}$ be a set of indices characterizing which groups in the portfolio are to be considered. Define

$$g(Y, G) = \sum_{i \in G} \lambda_i \Phi \left(\frac{s - \sqrt{\hat{R}_i^2} \hat{v}_i^\top Y}{\sqrt{1 - \hat{R}_i^2}} \right) \hat{\pi}_i,$$

with λ_i the proportion of exposures in group i within G . The function $g(Y, G)$ gives the (limiting) credit loss of Theorem 1 for this special setting considering the groups included in G only. We define the maximum possible credit loss as

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

In the previous subsection, π^* was equal to the sum of all exposures due to the homogeneity assumption for v_j . We now consider the tail behavior near π^* . To do this, we need to characterize the parts of the portfolio that can result in the maximum credit loss π^* . Define \mathcal{G} to be the family of subsets G of $\{1, \dots, n^*\}$ such that $\sup_y g(y, G) = \pi^*$ and such that for all proper subsets G^s of G one has $\sup_y g(y, G^s) < \pi^*$. The collection \mathcal{G} then defines parts of the credit portfolio that satisfy two criteria: (i) considered in isolation part G can result in the maximum credit loss, and (ii) there is no strict subset of this original part of the portfolio that can result in the (same) maximum loss. For each part G , we define Θ_G as a subset of the unit sphere $\{\theta : \|\theta\| = 1\}$, consisting of the directions for factor realizations f that push part G of the portfolio to the maximum credit loss π^* . So Θ_G contains the most unfavorable realizations of f for part G . Every such f can be written as $a\theta$ for some $\theta \in \Theta_G$ and $a \in \mathbb{R}^+$. Given all these definitions, we can prove the following theorem.

Theorem 3. *In the factor model where each loan is characterized by a vector $(\hat{R}_i^2, \hat{v}_i, \hat{\pi}_i)$ with $i \in \{1, \dots, n^*\}$, the credit loss distribution has a tail expansion as in (18) with tail index*

$$\min_{G \in \mathcal{G}} \min_{\theta \in \Theta_G} \max_{i \in G} \frac{1 - \hat{R}_i^2}{\hat{R}_i^2 (\hat{v}_i^\top \theta)^2}. \tag{20}$$

Theorem 3 is a clear extension of Theorem 2. The main difference is that we take the smallest tail index in the sense of Theorem 2 over all parts of the portfolio and corresponding unfavorable directions for f that can result in the maximum credit loss. The intuition is as follows. If different realizations of the common risk factor Y give rise to the same maximum credit loss, albeit for different subsets of the portfolio, the tail behavior of portfolio credit losses is a mixture of the tail behavior for each of these different subsets. The tail behavior for each subset is determined by Theorem 2, which explains the $\max_{i \in G} (1 - \hat{R}_i^2) / \hat{R}_i^2$ part in (20). Standard results on tail behavior then reveal

that for a mixture of tails, the fattest tail dominates the extreme tail behavior, see Ibragimov and Linnik (1971). In our context, the fattest tail has the lowest tail index, explaining why the minimum over $G \in \mathcal{G}$ and $\theta \in \Theta_G$ is reached in Theorem 3.

The tail decay rate α of the credit loss distribution for models with heterogeneous v_j is generally not below that of corresponding models with homogeneous v_j . This stems from the fact that both $\|\theta\| = 1$ and $\|v_j\| = 1$, and therefore $(\theta^\top v_j)^2 \leq 1$. The intuition is that with heterogeneous v_j we only consider particularly unfavorable realizations f dictated by Θ_G and relevant for a subset G of the portfolio only. This restriction further reduces the likelihood of ending up near the extreme credit loss π^* and therefore makes the extreme tail thinner, i.e., increases α .

To conclude this section, we stress once more that the analytic expressions obtained for the tail index only partly characterize the actual tail behavior. For a full characterization, one also needs expressions for the slowly varying functions in (18). These are much more difficult to derive analytically. The tail index, however, clearly reveals the rate of tail decay and therefore contains useful information. For a complete characterization, one can also use the limit law from Theorem 1 directly. This is done in the empirical illustrations of the following section.

5. Examples based on empirical data

In Sections 3 and 4, we studied the behavior of the limiting *default* loss distribution in the stylized setting of Belkin et al. (1998a) and somewhat more general factor models. In this section, we investigate the behavior of *credit* loss distributions in more detail. Credit losses comprise both losses due to default and losses due to changes in the credit quality or credit rating of a firm. The latter occur in a mark-to-market framework where credit spreads differ over credit ratings (and possibly industries and maturities). We again focus on a one-factor model for rating migrations but generalize the previous set-up by allowing for differences in initial ratings, factor model fits, exposure sizes, and bond maturities. We also consider thin-tailed and fat-tailed distributions for the common risk factors f and show that minor changes in distributional assumptions can have large effects on extreme credit loss quantiles. The sensitivity analyses conducted in this section would be time-consuming with a standard simulation approach. Using the efficient computation techniques presented in this paper, however, we reduce this burden significantly.

In order to calculate empirical credit quantiles we need to know the rating migration probabilities, the yields and yield spreads, the initial ratings of the exposures in the portfolio, the credit loss functions $\pi_j(\cdot)$, and the R_j^2 s of the factor model (2). We use a classification with seven rating categories: AAA,

AA, A, BBB, BB, B, and CCC. In addition, default (D) can occur, see also Fig. 1. The rating migration probabilities, yields, and yield spreads were downloaded from the *CreditMetrics* website. The transition probability matrix used is presented in Table 2. Default and rating migration probabilities are then equal to long-term historical averages. The transition probabilities are used to determine the binning constants s_{kl} through (4).

The table's right-hand panel contains yield spreads for corporate bond maturities of 1, 5, and 10 years. We also have observations (not displayed) for maturities of 2, 3, and 7 years, to be used later on. The maturities are assumed to be equal for all exposures within a given portfolio. This is because we were unable to obtain detailed and reliable information on the distribution of maturities across initial ratings, exposure sizes, and degrees of systematic risk. At the outset, we start with bonds that are priced at par using the base yields and credit spreads in the table. The spreads depend on the initial rating of the firm. For the sake of convenience we assume that the term structure of yields is constant over time. In this way, we fully concentrate on credit risk without the interference of market risk. Required yields that are not observed directly are computed by linear interpolation based on observed yields.

We also vary the portfolio credit quality. In order to have a realistic distribution of initial credit ratings for empirical work, we use the data in Gordy (1999). His study provides the percentage exposure and number of obligors in specific rating categories for typical bank portfolios. These data are based on a bank survey by the Federal Reserve on the credit quality composition of bank portfolios as well as on the private placements report by the Society of Actuaries (1998). We combine the numbers provided in Gordy (1999) to determine the average size of an exposure for each rating category. The precise numbers used in our study are provided in Table 3. We consider three portfolio quality levels: high quality, average or benchmark quality, and low quality.

We also need a specification for the loss functions $\pi_j(\cdot)$. As mentioned, we abstract from changes in the yield curve over time in order to abstract from

Table 3
Typical bank portfolio characteristics^a

	Rating						
	AAA	AA	A	BBB	BB	B	CCC
High quality	4	6	29	36	21	3	1
Average quality	3	5	13	29	35	12	3
Low quality	1	1	4	15	40	34	5
Rel. exposure size	109	93	96	93	107	107	69

^a The table contains the percentage of obligors in each rating category for a typical bank portfolio of high, average, or low quality. The bottom line gives the mean size of the loan in each rating category. The numbers are based on Gordy (1999).

market risk. We assume that the portfolio consists of plain vanilla loans or bonds only. Changes in the market value of a bond at any time are then caused by changes in the yields only. As there is no time-variation in the yields, such a change can only arise due to a credit rating migration and a shortening of the maturity of the bond. Let $y_{k,m}$ denote the yield on a bond with initial rating class k and maturity m . If the bond is initially evaluated at par, the bond value for a unit principal at the planning horizon of one year is given by

$$\begin{aligned} &(1 + y_{\ell,m-1})^{-(m-1)} + \sum_{i=0}^{m-1} (1 + y_{\ell,m-1})^{-i} y_{k,m} \\ &= 1 + y_{k,m} - \left(1 - \frac{y_{k,m}}{y_{\ell,m-1}}\right) \left(1 - (1 + y_{\ell,m-1})^{-(m-1)}\right), \end{aligned}$$

which follows by a straightforward discounting of present and future cash flows. This only holds if there is no default. If the bond defaults, a fraction γ of the (unit) principal amount, the so-called recovery rate, can be recovered. Summarizing, the individual credit loss function boils down to

$$\begin{aligned} &\pi(j, k, \ell) \\ &= \begin{cases} \left[\left(1 - \frac{y_{k,m}}{y_{\ell,m-1}}\right) \left(1 - (1 + y_{\ell,m-1})^{-(m-1)}\right) - y_{k,m} \right] e_j & \text{for } \ell = 1, \dots, r-1, \\ (1 - \gamma) e_j & \text{for } \ell = r, \end{cases} \end{aligned} \tag{21}$$

with e_j the size of the j th principal amount. The values of e_j depend on k_j and can be found in the bottom line in Table 3. Using recovery data at emergence from S&P (1998), we set the recovery rate to 0.6.

Finally, in order to calculate the degrees of systematic risk or R_j^2 we ran CAPM regressions with the S&P500 as market index for the 1662 stocks in the Research Top 2000 company list from DATASTREAM for which at least 5 years of data were available. Using regressions at a monthly, quarterly, and annual frequency, we obtain three distributions of R^2 depicted in Fig. 3. The averages of these distributions are equal to 0.20, 0.27, and 0.35, respectively.

Given our available data set, we did not find any significant relationship between R^2 values and firm ratings. We therefore impose the same distribution of R^2 s per rating category. We also inspected the values of the directional vectors v_j of Theorem 1. In our one-factor set-up, we have $v_j = 1$ or $v_j = -1$ 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 $v_j \equiv 1$ for all exposures in the sample.

Using Theorem 1 and expression (17), we can now compute the credit loss quantiles without resorting to simulations. The mean and variance of portfolio credit losses can moreover be computed by numerical integration, see (16). We consider 27 credit loss distributions corresponding to three maturities, three

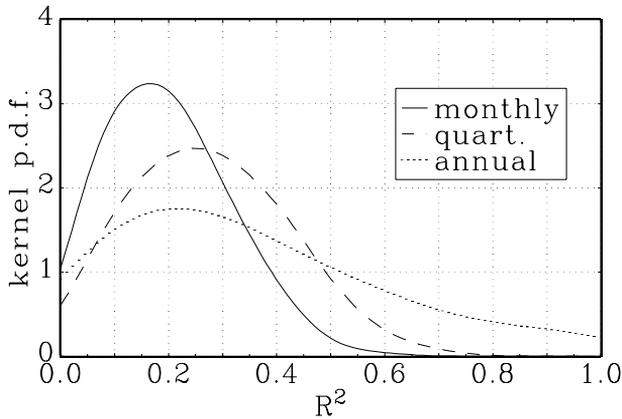


Fig. 3. R^2 values of factor model regressions using monthly, quarterly, and annual data. The data are obtained from the Research Top 2000 list of DATASTREAM Inc. and comprise 1662 firms observed over the period January 1980–December 1998. The factor model explains total return of the firm by a constant and the total return on the S&P500. A minimum of 5 years of data is used for the factor model regressions.

distributions for R^2 , and three portfolio quality levels. The boxplots in Fig. 4 summarize the results for normally distributed systematic (f) and idiosyncratic (ε_j) risk factors.

First, consider the effect of changing the degree of systematic risk (R^2) in Fig. 4. 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. This stems from the fact that credit rating upgrades will also be more correlated over the different exposures in the portfolio. Especially for high-risk quantiles, the effects of a higher average R^2 are substantial: for the 1 year bond and the average quality (B), the 99.9th percentile of credit losses shifts from about 5% ($M - R^2$) of the notional amount via 7% ($Q - R^2$) to 10% ($Y - R^2$).

The right-hand panels in Fig. 4 express the loss quantiles in terms of standard deviations in excess of the expected loss. If the R^2 distribution is kept fixed and the portfolio's initial quality is changed, the distributions with the most prolonged tails in the left-hand panels appear to have the thinnest tails in the right-hand panels. By contrast, if the portfolio's initial quality is fixed and the distribution of R^2 is changed, the ordering of tails in the left-hand and right-hand panels remains unaltered. The 99.9th percentiles for all portfolios (indicated by the top of the whisker) in the right-hand panels of the figure are significantly larger than 3, the number one expects for a normal distribution ($\Phi(3) \approx 99.9\%$). This illustrates that the normal approximation to credit loss quantiles may be completely inappropriate. Instead of relying on the normal, it is more appropriate to use the limiting distribution of Theorem 1 directly, see

the left-hand panels in Fig. 4. In any case, one should be careful when interpreting credit loss quantiles expressed in terms of standard deviations in excess of expected loss.

Next, Fig. 4 clearly shows that low-quality portfolios have a worse credit loss performance as upper credit loss quantiles shift to the right. Lower quantiles may shift to the left as positive returns (through correlated upgrades) also become more likely.

Finally, the bonds' maturities have the expected effect on the credit loss quantiles. For 1-year bonds, the only credit losses are those due to default. For 5-year and 10-year bonds, however, the effect of rating migrations and differing credit spreads is also taken into account. These have a higher impact the longer the maturity, i.e., the higher the duration or interest elasticity of the bond portfolio. Therefore, upper quantiles are higher for longer maturities.

In order to check the robustness of the results somewhat further, we conduct the following experiment.² Gordy (1999) notes that the distribution of the systematic risk component f may be important for the location of extreme credit loss quantiles. To check this, we deviate from the normality assumption by allowing for fat tails in the common factor f . More specifically, we fit a Student t distribution to the S&P return series, which is our common factor in the CAPM regressions. Maximum likelihood estimation reveals that the index return can best be described by a Student t distribution with 5 degrees of freedom. Fig. 5 presents the normal and the (unit variance) Student $t(5)$ distributions in one graph. It is seen that although the two distributions superficially look alike, their extreme quantiles differ quite substantially. Using the Student $t(5)$ instead of the normal as the distribution of f , we recomputed all credit loss quantiles. The results are presented in Fig. 6.

As expected, given the large shift in extreme (99.9%) quantiles of f in Fig. 5, Fig. 6 reveals much larger credit loss quantiles than in the case of a normally distributed f . For the 99.9th percentile, the increase varies from about 500 to 1000 basis points of the invested notional amount, see the top of the whiskers in the boxplots. Also note that quantiles less far out in the tail are affected to a much lesser extent. An important conclusion from this experiment, therefore, is that one should pay more attention to the distributional specifications of the risk factors f and ε_j in (2) if one is interested in quantiles extremely far out in the tail (e.g. 99.9%). Superficially minor changes in the distribution may result in large shifts of the quantiles. A reliable estimate of the distribution of f may, however, not be easily obtained given the data that are typically available. Appropriate sensitivity analyses such as the one presented here are therefore a prerequisite before any of these credit risk quantiles can be used to formulate

² We are grateful to one of the referees for suggesting this type of experiment.

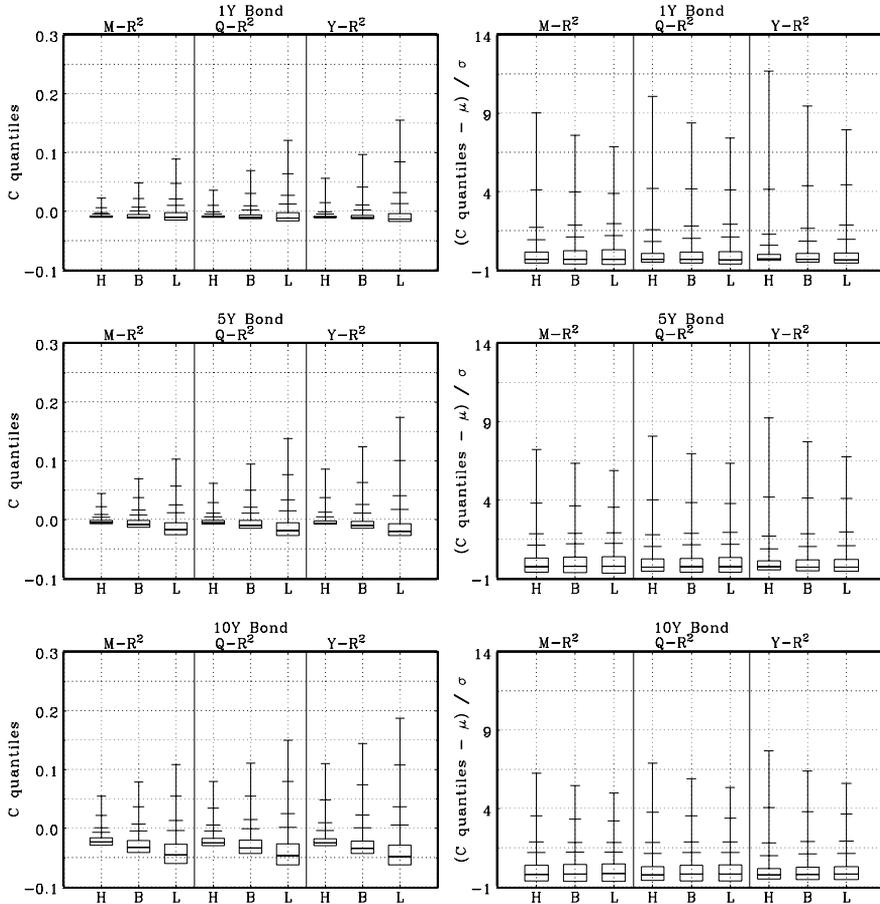


Fig. 4. The figure summarizes the effect of differing maturities, systematic risk and portfolio quality on the limiting credit loss distribution by means of boxplots. The plots express the credit loss either as a fraction of the notional amount (left three panels), or in terms of the number of standard deviations in excess of the expected loss (right three panels). Each row of two plots relates to corporate bond portfolios of a given maturity. Each plot contains three panels for three different distributions of systematic risk R^2 . The left, middle, and right panels are based on the distribution of R^2 s using regressions with monthly (M- R^2), quarterly (Q- R^2), and yearly (Y- R^2) data, respectively. See also Fig. 3. Within each of these panels, three boxplots are presented for portfolios of high (H), average or benchmark (B), and low (L) quality, respectively. The initial rating distributions corresponding to these different levels of quality can be found in Table 3. Each box represents the interquartile range of credit losses whereas the middle line indicates the median. The whisker of the boxplot has four markings, relating to the 0.9, 0.95, 0.99, and 0.999 quantile of credit losses. A recovery rate of $\gamma = 0.6$ is used for all bonds. The risk factors f and ε_j are normally distributed.

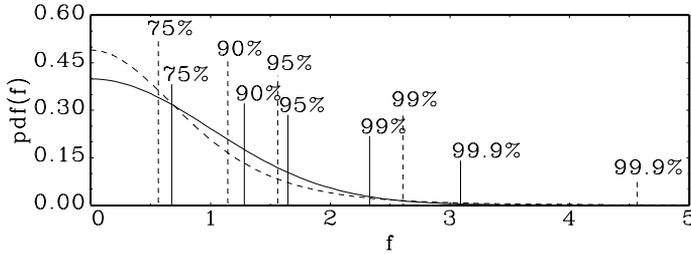


Fig. 5. Normal (solid) and unit variance Student $t(5)$ (dashed) density with quantiles. Quantiles are indicated by solid vertical lines for the normal, and by dashed lines for the Student t . As both distributions are symmetric around zero, we only plot the distributions on the positive halfline.

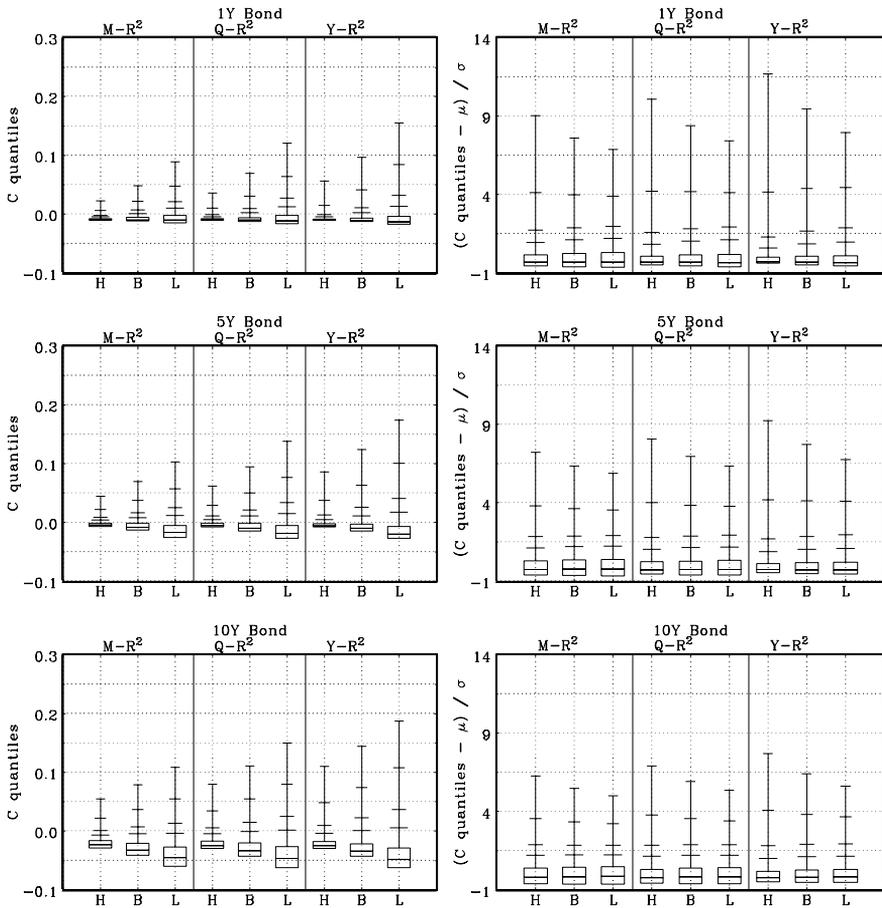


Fig. 6. The figure contains the same information as Fig. 4, the difference being that f is now assumed to follow a (unit variance) Student $t(5)$ distribution instead of a normal.

precise capital requirements. Such sensitivity analyses can be speeded up using the analytic approach set out in the present paper.

To conclude, we comment on the Basle Committee on Bank Supervision (1988) capital adequacy directives, i.e., the 8% capital requirement for corporate bonds and loans. Table 4 gives the probability in basis points (bp) of the credit loss exceeding 8% of the notional amount over a one-year period for the portfolios considered earlier. If the 8% rule would correspond to a 99.9% confidence quantile, the entry in the table would thus be 10 (bp).

Clearly, the 8% guideline corresponds to a quantile with a confidence level above 98% or higher, i.e., exceedance probabilities below 200bp. When the 8% rule is expressed in terms of standard deviations in excess of the expected loss, we obtain numbers between 2.8 and 25.9 (excess) standard deviations. This again illustrates that expressing the 8% guideline in terms of the number of standard deviations in excess of expected loss can be very misleading, as it highly depends on the portfolio composition. Note that the 8% guideline and the credit loss quantiles are not comparable directly. The 8% rule concerns a requirement at the *start* of the planning period, while the quantiles relate to the potential loss at the *end* of the same period. The two can be compared if one agrees upon a discount rate for the credit loss quantiles in Fig. 4, e.g., the appropriate risk free rate or an internal rate of return. Discounting the potential credit loss quantiles makes the 8% guideline more conservative, i.e., decreases the exceedance probabilities in Fig. 4. Also note the correspondence of the exceedance probabilities between normally and Student t distributed f in the upper and lower panels of the table, respectively. In most cases the lower panel probabilities are somewhat lower than those in the upper panel. Resuming, the

Table 4
Probabilities of credit loss exceeding 8% over 1 year^a

Maturity	Credit quality								
	M- R^2			Q- R^2			Y- R^2		
	H	B	L	H	B	L	H	B	L
Normal common factors f									
1 year	0.04	0.8	16	0.5	6	53	4	20	113
5 years	0.4	5	32	3	21	90	13	54	174
10 years	2	9	35	10	36	99	31	85	193
Student $t(5)$ common factors f									
1 year	0.01	0.4	15	0.08	3	47	0.3	9	96
5 years	0.2	4	31	1	15	87	4	38	169
10 years	1	8	34	6	31	99	18	71	196

^aThe table contains the probabilities in basis points of credit losses exceeding 8%. Results are presented for high (H), medium (B), and low (L) portfolio quality, different bond maturities (1, 5, and 10 years), and different distributions for systematic risk R^2 (monthly (M), quarterly (Q), and yearly (Y) frequency), see also the note to Fig. 4.

results in this section suggest that banks and supervisory institutions should take portfolio features such as the maturity, composition, quality, and systematic risk of the portfolio into account when setting capital requirements.

6. Speed of convergence

So far, we concentrated on calculating credit loss quantiles if the number of exposures grows indefinitely. Actual portfolios, however, contain a finite (but possibly large) number of exposures. In this section, we establish the required portfolio size n such that the limit law's upper quantiles are good approximations to the loss quantiles of finite portfolios. We obtain the latter by simulation, which is the dominant approach in the literature. We investigate to what extent portfolio characteristics such as the degree of systematic risk or the average portfolio quality slow down the speed of convergence of simulated quantiles towards their analytic counterparts. A slower convergence may hamper the applicability of our approach to actual portfolios.

The framework is the same as in the previous section. Preliminary simulation evidence revealed that the bond's maturity does not significantly alter the convergence results. Therefore, we only present results for a maturity of 3 years. We simulated 36 finite credit portfolio distributions, i.e., for three distributions of R^2 (monthly, quarterly, yearly), three initial rating distributions (high, benchmark, low), two distributions of f (normal or Student $t(5)$) and 2 degrees of portfolio 'dispersion' (high and low). The portfolio dispersion, labeled as an integer $v \geq 0$, deserves some further clarification. If $v = 0$, we set e_j in (21) equal to the mean loan size for rating k_j , see Table 3. If $v > 0$, we split the portfolio segment with initial rating k into a fraction $1 - 1/v$ with loan sizes equal to the mean loan size, and a fraction $1/v$ with loans of size v times the mean loan size. As a result, the large loans comprise a fraction of about $1/v$ of the portfolio, while the remaining fraction of $(v - 1)/v$ consists of regularly sized loans. The loan portfolio's degree of dispersion v therefore relates to heterogeneity in loan sizes. The heterogeneity increases in v . In our simulations, we use $v = 0$ and $v = 10$. Thus, the portfolio with the highest dispersion ($v = 10$) consists for 10% of bonds with a face value that is 10 times larger than the other portfolio loans.

Note that 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 Fig. 3. 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

$$F_{R^2}^{-1}\left(\frac{i}{n_k + 1}\right) \quad \text{for } i = 1, \dots, n_k. \quad (22)$$

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 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 . For positive degrees of portfolio dispersion ($v > 0$) we slightly adapt the above procedure by assigning R^2 -values using (22) per rating category and segment of loan size (large/small) instead of per rating category only. In that case, n_k has to be interpreted as the number of either small or large loans with initial rating k , respectively.

We can now compute quantiles of the limiting credit loss distribution under different scenarios for maturities, portfolio dispersion (v), R^2 distribution, initial rating distribution, and number of exposures (n). These quantiles are used as the benchmark in checking the convergence speed of the credit loss distribution. For portfolios with a limited number of exposures we need to resort to Monte-Carlo simulation. The resulting credit loss quantiles, however, can be very unstable. We remedy this problem in the following way. 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 quantile estimates are obtained by averaging the 10 replications. The discrepancies between these averages and the limiting distribution's quantiles are presented in Fig. 7 for different parameter configurations. Results for other levels of portfolio quality, distributions of R^2 , etc., appeared to be very similar and are therefore omitted.

For the homogeneous portfolios ($v = 0$) with Gaussian f , the quantiles of the limiting credit loss distribution generally lie very close to those of the finite portfolio distribution, see Fig. 7. For portfolios with at least $n = 300$ exposures, the difference lies in a 0.25% band of the invested notional amount. 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 99.9th percentile in several panels of Fig. 7. The convergence is enhanced if annual instead of monthly R^2 s are used. Note that the former has a higher mean, i.e., a higher average degree of systematic risk. Finite portfolios with lower R^2 s, by contrast, have more idiosyncratic risk. Such idiosyncratic risk is not captured by the limiting result of our paper. This explains why the quantiles for small n generally lie above the limiting quantiles (except for estimation error): the simulation approach accounts for both the idiosyncratic and the systematic risks, whereas the limiting distribution uses the latter only. The importance of idiosyncratic risk, however, is negligible for empirically relevant portfolio sizes.

A higher loan size dispersion ($v = 10$) lowers the convergence speed, see the second row in Fig. 7. Larger portfolio sizes are needed to obtain a similar

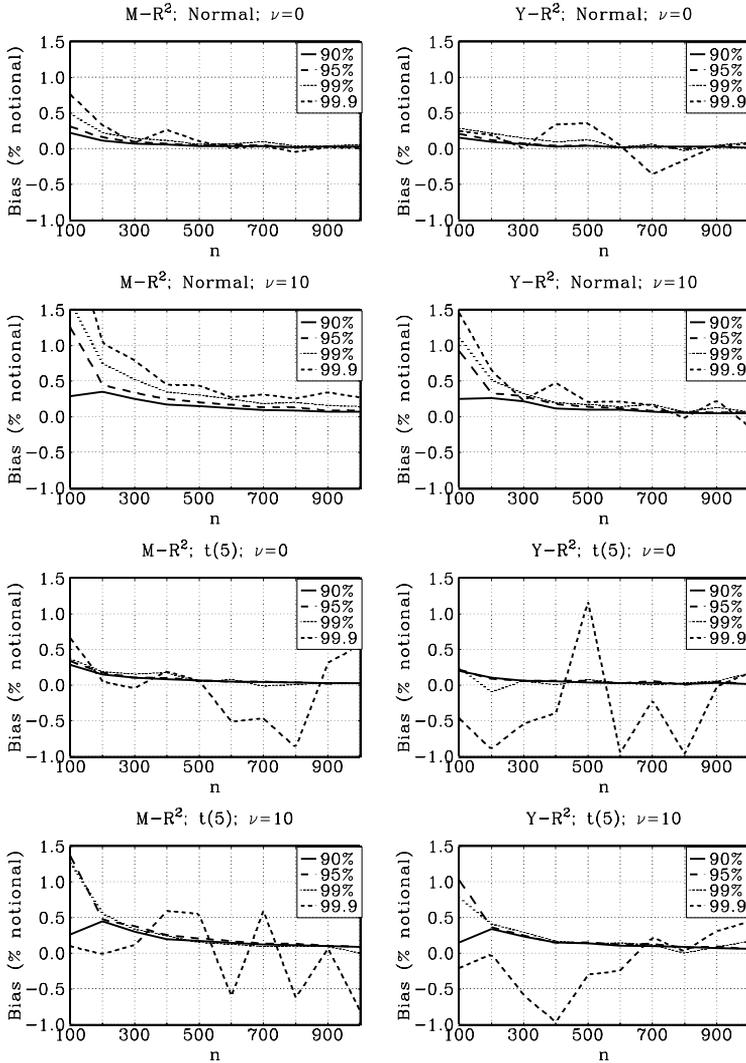


Fig. 7. The figure presents the upper credit loss quantiles (90%, 95%, 99%, and 99.9%) as a percentage of the notional value for a portfolio consisting of n firms. Simulated quantiles are expressed in deviation of the limiting distribution's quantiles. The results are for a bond maturity of 3 years. For $\nu = 0$ all exposures for a fixed rating category have the same size, while for $\nu = 10$, 10% of the exposures (per rating category) have a size 10 times larger than that of the remaining 90% exposures. The left and right columns contain the results for given degree of systematic risk (R^2) based on monthly (M) and annual (Y) CAPM regressions, respectively. The top four graphs are for normally distributed systematic risk factors f , while the bottom four graphs are for f following a Student $t(5)$ distribution. Quantiles of the finite portfolios are based on averages over 10 estimates of the appropriate quantiles, where each of these estimates is based on a Monte-Carlo experiment of size 20,000.

accuracy as in the first row of the figure, i.e., n should be at least 500 to make the discrepancies smaller than 0.25% of the notional amount. Again we see that the convergence sets in earlier for the annual compared to the monthly R^2 distribution.

The convergence results for heavy-tailed (Student $t(5)$) systematic risk factors f are very similar to those for Gaussian f , compare the bottom four graphs with the top four graphs. However, estimating extreme quantiles (like 99.9%) becomes more complicated for finite portfolios if f has fat tails. The quantiles emerging from our limiting approach are much more stable and, therefore, preferable from a robustness point of view.

Summarizing, the limiting distribution fits the distribution for finite portfolios very closely for reasonably sized homogeneous portfolios ($n \geq 300$). The fit decreases if loan sizes are more dispersed. In that case, larger portfolio sizes are needed, i.e., $n > 500$.

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, the precise distribution of f and ε_j , as well as some of the output, e.g., the estimated extreme quantiles like the 99.9th percentile. Given this uncertainty, the discrepancies reported in Fig. 7 seem acceptable once one is willing to adopt a simulation-based model for portfolio credit risk management. 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.

7. Conclusions

In this paper, we studied the credit loss distribution of portfolios comprising a large number of corporate bonds or loans. Our results, however, are also applicable for more complicated financial instruments. The proposed approach builds further upon the factor model approach to credit risk, see e.g., J.P. Morgan (1999). We formally derived the portfolio credit loss distribution if the number of exposures grows indefinitely. This distribution reveals that the tail behavior of credit losses is highly influenced by the fit of the factor model regressions, measured in terms of R^2 .

Using extreme value theory, we showed that the rate of tail decline of credit losses is much lower than that of a normal distribution. We derived how the tail decay rate evolves as a function of portfolio characteristics, both for one-factor and multi-factor models. For the case of identical systematic risk across exposures, we found that higher values of R^2 imply a slower rate of tail decay, i.e., fatter tails. If, however, one allows for differing R^2 s over the exposures, exposures with the highest idiosyncratic risk component ultimately dominate the tail behavior. As idiosyncratic risk can be diversified, the contribution of

these exposures to the limiting distribution of credit losses is small, making it less likely that the maximum credit loss will be hit. This, in turn, implies that the tail of the distribution is thinner.

Applying our limiting result to empirically realistic portfolios, we showed that both increases in systematic risk and declines in portfolio quality (initial rating distribution) shift the credit quantiles to the right. However, if one expresses portfolio credit risk in terms of standard deviations in excess of the expected credit loss, the ordering of riskiness compared to the direct measurement of credit loss quantiles is reversed when considering the sensitivity to changes in the overall portfolio credit quality. Thus, one should be careful when expressing and interpreting credit losses in terms of standard deviations.

The location of extreme credit loss quantiles is particularly sensitive to altering the distributional assumptions for the common risk factors. Replacing the normal distribution by the heavy-tailed Student $t(5)$, the 99.9th percentiles shift to the right by about 5–10% of the notional value.

Our quantile calculations also suggest that the Basle 8% capital requirement corresponds to confidence levels above 98%, or even above 99% for portfolios of average or high quality. If this exercise is repeated in terms of standard deviations in excess of the expected loss, however, the resulting numbers vary considerably with the portfolio characteristics. This illustrates that the latter way of stating capital requirements should not be used if a certain confidence level needs to be attained.

As shown by simulation, the derived limiting distribution may be a useful tool for credit risk managers because it provides an approximation to the ‘true’ (finite portfolio) credit loss distribution. For medium-sized portfolios with homogeneous loan sizes containing at least 300 bonds, the appropriate quantiles of the limiting distribution depart by less than 0.25% of the notional amount from the empirical credit loss quantiles. For highly dispersed loan sizes, the required number of exposures per portfolio should be higher in order to obtain a similar accuracy.

The paper suggests several strands of future research. First, our approach can be extended to a dynamic setting along the lines suggested in Section 2. Also, the model can be used to obtain a complete characterization of portfolio risk, encompassing both market risk and credit risk.

Acknowledgements

We thank Laurens de Haan, Bernard Hanzon, Herbert Rijken, Ronald van Dijk, and two anonymous referees for useful comments and suggestions. André Lucas also thanks the Dutch Organization for Scientific Research (N.W.O.) for financial support.

References

- 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 1999. A new capital adequacy framework. Report 50, Bank of International Settlements, Basle (June).
- Belkin, B., Suchower, S., Forest, L., 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., Suchower, S., Forest, L., 1998b. A one-parameter representation of credit risk transition matrices. *CreditMetrics Monitor* (3rd Quarter), 46–56. <http://www.creditmetrics.com>.
- Black, F., Cox, J., 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⁺*. <http://www.csfp.csh.com>.
- Embrechts, P., Klüppelberg, C., Mikosch, T., 1997. Modeling Extremal Events. Applications of Mathematics: Stochastic Modelling and Applied Probability. Vol. 33, Springer, Heidelberg.
- Finger, C., 1999. Conditional approaches for CreditMetrics portfolio distributions. *CreditMetrics Monitor* 1 (April), 14–33.
- Fons, J., 1991. An approach to forecasting default rates. Moody's Special Report (August).
- Gordy, M., 1999. A comparative anatomy of credit risk models. *Journal of Banking and Finance*, forthcoming.
- Hogg, R., Craig, A., 1970. Introduction to Mathematical Statistics. MacMillan, New York.
- Ibragimov, I., Linnik, Y., 1971. Independent and Stationary Sequences of Random Variables. Wolters-Noordhoff, Groningen.
- IIF, 1998. Recommendation of revising the regulatory capital rules for credit risk. Report of the Working Group on Capital Adequacy of the Institute of International Finance (March).
- ISDA, 1998. Credit risk and regulatory capital. International Swaps and Derivatives Association (March).
- Jónsson, J., Fridson, M., 1996. Forecasting default rates on high-yield bonds. *Journal of Fixed Income* (June), 69–77.
- J.P. Morgan, 1999. *CreditMetrics* 4th ed. <http://www.creditmetrics.com>.
- Kealhofer, S., 1995. Managing default risk in derivative portfolios. *Derivative Credit Risk: Advances in Measurement and Management*. Risk Publications, London.
- Koyluoglu, H., Hickman, A., 1998. Reconcilable differences. *Risk* 56 (October), 56–62.
- Longstaff, F., Schwartz, E., 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.
- Society of Actuaries, 1998. 1986–94 Credit Risk Loss Experience Study: Private Placement Bonds. Society of Actuaries, Schaumburg (IL).
- S&P, 1998. Ratings Performance 1997: Stability and Transition Report.
- Wilson, T., 1997a. Portfolio credit risk: Part I. *Risk* (September), 111–117.
- Wilson, T., 1997b. Portfolio credit risk: Part II. *Risk* (October), 56–61.
- Zhou, C., 1997. Default correlation: An analytical result. Technical Report 1997–27, Board of Governors of the Federal Reserve System. Finance and Economics Discussion Series.