Journal of Credit Risk **10**(4), 45–70

Analysis of credit portfolio risk using hierarchical multifactor models

Pak-Wing Fok
Ewing Hall 412, Department of Mathematical Sciences, University of Delaware, Newark, DE 19716, USA; email: pakwing@udel.edu

Xiuling Yan
017 Purnell Hall, Department of Accounting and Management Information Systems, University of Delaware, Newark, DE 19716, USA; email: xyan@udel.edu

Guangming Yao
Department of Mathematics, Clarkson University, 8 Clarkson Avenue, Potsdam, NY 13676, USA; email: gyao@clarkson.edu

(Received June 28, 2013; revised August 15, 2014; accepted August 27, 2014)

In this paper, we generalize Vasicek's asymptotic single-risk factor solution to multiple factors organized with a particular hierarchical structure. We use this model to investigate credit portfolio loss. In this hierarchical factor model, the asset returns of a company depend on a global factor, a sector factor and an idiosyncratic risk factor. All companies share the same global factor and all companies within a sector share the same sector factor. Using the central limit theorem, we derive closed-form solutions for the value-at-risk (VaR) and expected shortfall under the assumption that the number of sectors in the portfolio is large and the exposures scale is the reciprocal of the number of sectors. Our results for the VaR agree with Monte Carlo simulations, provided the sector factor loadings and variance of systematic risk are not too large.

1 INTRODUCTION

Credit portfolios are portfolios of fixed-income investment products such as bonds, loans and credit derivatives. Fixed-income investment products provide the investor

The authors wish to thank William Morokoff and Saiyid Islam for many helpful comments on the manuscript and for introducing them to the credit portfolio problem during the Delaware 2012 Mathematical Problems in Industry workshop. They also thank Lou Rossi and David Edwards for organizing the workshop. This work was partially supported by a grant from the Simons Foundation (282579 to Pak-Wing Fok)

with a steady stream of cash inflow (eg, in the form of interest payments) during the lifetime of the product. The trade-off is limited upward potential for portfolio gain. Banks, insurance companies and other financial institutions regularly maintain and manage large numbers of credit portfolios. The main risk associated with such portfolios is a debtor defaulting on its obligation. Although such an event is rare, a single default often means that the entire portfolio goes to loss. Therefore, investors in credit portfolios need systematic methods to analyze the associated risk and create financial instruments to insure against losses should they arise. The management of credit risk is a vital area of research within quantitative finance (see, for example, Bohn and Stein 2009; Denault 2001; De Servigny and Renault 2004; McNeil *et al* 2010).

It is well-known that companies do not default independently from each other (Lucas 1995). One common way of modeling correlated company defaults is through "factor models" (Bluhm *et al* 2010; Burtschell *et al* 2009; Schönbucher 2001). In these models, a representation for the asset return of a company is specified in terms of random variables. When the asset return drops below a given threshold, the company defaults. Correlation between company defaults is included by allowing the random variables to share common factors. Factor models are studied and commonly employed by companies such as Moody's KMV (Crouhy *et al* 2000) and the RiskMetrics Group (Gordy 2003).

The Vasicek (1987) credit model provides a simple analytic solution for a portfolio containing identical companies that are coupled with a single global factor. While Vasicek's asymptotic single-risk factor (ASRF) solution is simple to derive, it can also be easily extended to the heterogeneous case and, importantly, forms the foundation of the Basel accords for bank capital requirements (Basel Committee on Banking Supervision 2004). Other authors have extended the ASRF model to account for uncertainty over loss given default (Kupiec 2008) and multiple global risk factors (Pykhtin 2004; Schönbucher 2001).

In this paper, we propose and validate an analytic formula for the loss distribution of a credit portfolio, assuming a hierarchical multifactor model. In such a model, all companies have exposure to a global risk factor; in addition, companies in a given sector are subject to a local risk factor. Thus, it may be regarded as a simple extension of the ASRF model to an economically intuitive multifactor case. Since our loss formula can be written entirely in terms of elementary and special functions, it is much quicker to evaluate than Monte Carlo simulations, which are often time consuming and computationally expensive. Our derivation involves analyzing the sector loss and then applying the central limit theorem to all the sectors. It is similar to saddlepoint methods (Huang *et al* 2007; Jensen 1995; Lugannani and Rice 1980) in the sense that both methods approximate sums of random variables through asymptotic formulas.

FIGURE 1 Hierarchical factor model of company asset returns within a global economy with $N = 3$ sectors and $n = 3$ companies per sector.

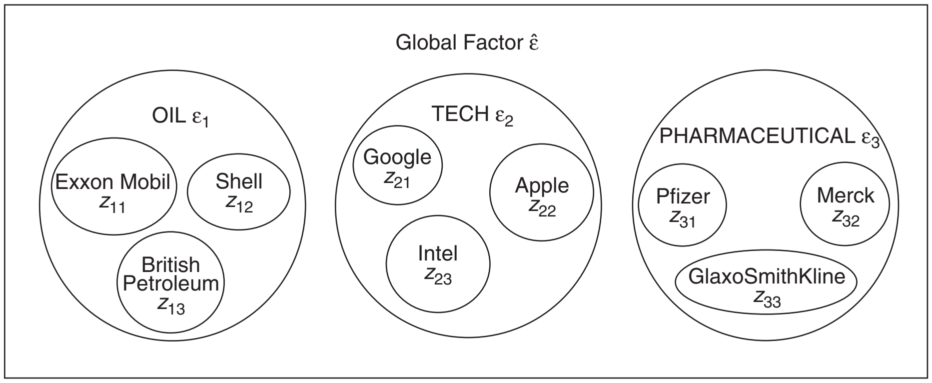

Every company participates in a global economy and belongs to exactly one sector. Each company's asset return z_{ij} is affected by a global factor $\hat{\varepsilon}$, a sector factor ε_i and an idiosyncratic factor ζ_{ij} (not shown). See (2.1).

The layout of this paper is as follows. In Section 2, we introduce the hierarchical factor model for a firm's value and set up the portfolio in terms of individual companies and their default probabilities. In Section 3, we derive the value-at-risk (VaR) for a portfolio coupled with a hierarchical factor model. In Section 4, we compare our solutions for the VaR with Monte Carlo simulations. We conclude the paper in Section 5.

2 PORTFOLIO PRICING IN A LOCALIZED ONE-FACTOR MODEL

One critical issue that determines the value of credit portfolios is the default correlation among companies (Schönbucher 2001). Although the default probability of a company may be very small, defaults between companies are often correlated. Factor models incorporate the correlation among asset returns explicitly by assuming they are driven by M shared "factors" that are modeled as independent random variables. For example, in an $M = 2$ factor model, these shared factors could represent the state of a country's economy (a recession negatively impacts all companies in that country), or the price of a resource (a lower price would lower the expenses of all companies that use the resource). All asset returns in an $M = 2$ factor model would be influenced by the country's economy and the price of the resource.

In this paper, we restrict our attention to hierarchical (or "localized") factor models, which have the advantage of being simple yet economically intuitive. In such models, the asset return for a company depends on a "global" factor that is shared by all companies and exactly one of N other "sector" factors (see Figure 1 on the preceding page). Hence, all companies are correlated through the global factor and all companies in the same sector are further correlated. Although in this paper our model partitions a portfolio into different industrial sectors, our approach can also be applied to partitions of geography and size buckets.

Specifically, using a Merton (1974) model for a firm's value, we consider a special case of an $(N+1)$-factor model where the normalized asset return of the jth company in sector i is given by

$$z_{ij} = \sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i) + \sqrt{1-\rho_{ij}}\zeta_{ij}, \quad i = 1, \dots, N, \; j = 1, \dots, n, \quad (2.1)$$

and $0 < \rho_{ij} < 1$, $0 < \beta_{ij}, \hat{\beta}_{ij} < 1$. In this model, the asset returns are normalized so that $V(z_{ij}) = 1$, ie, $\hat{\beta}_{ij}^2 + \beta_{ij}^2 = 1$. The $\hat{\beta}_{ij}$ and β_{ij} are the global and sector factor loadings; knowledge of β_{ij} determines $\hat{\beta}_{ij}$, and vice versa. Since $\sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i)$ is the systematic risk, ρ_{ij} is the variance of systematic risk, $\hat{\varepsilon} \sim \mathcal{N}(0,1)$ is the global risk factor, $\{\varepsilon_i\} \sim \mathcal{N}(0,1)$ are the N independent sector risk factors and $\{\zeta_{ij}\} \sim \mathcal{N}(0,1)$ are the $N \times n$ independent idiosyncratic risk factors.

Although there are $N+1$ factors altogether, each company only depends on two of them. Any two companies are always correlated at the global level, and possibly also at the sector level. Specifically, we have

$$\operatorname{Corr}(z_{ij}, z_{kl}) = \begin{cases} \sqrt{\rho_{ij}\rho_{kl}}\hat{\beta}_{ij}\hat{\beta}_{kl} & \text{if } i \neq k, \\ \sqrt{\rho_{ij}\rho_{il}}(\hat{\beta}_{ij}\hat{\beta}_{il} + \beta_{ij}\beta_{il}) & \text{if } i = k. \end{cases} \quad (2.2)$$

In other words, the asset returns of any two companies are always correlated through their global factor loadings. If the companies also happen to be in the same sector, they are further correlated through their sector factors.

We point out that the hierarchical factor model for the company's asset return, (2.1), is a special case of a multifactor model in which all companies are influenced (to varying degrees) by multiple systematic risk factors. This general case has been studied by Pykhtin (2004); his approach is to optimally approximate the multifactor model by carefully choosing the factor loadings in the single-factor model. In our approach, we assume a simpler structure for the company's asset return from the very beginning. In return, we obtain an analytic solution that is simple to evaluate and exact in the limit as $N \to \infty$.

We assume that a company defaults if its asset return z_{ij} drops below a threshold value θ_{ij}. In principle, all companies within the global economy could have different thresholds, and details on how to determine them can be found in Crosbie and Bohn (2002). Over a fixed time horizon, we may write the loss on the portfolio as

$$R_\pi = \sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij} R_{ij}(z_{ij}), \quad (2.3)$$

where the exposures $w_{ij} > 0$ satisfy $\sum_{i=1}^{N} \sum_{j=1}^{n} w_{ij} = 1$. More dynamic evolution models that explicitly account for the Brownian nature of a firm's value are studied in Huh and Kolkiewicz (2008), Hurd (2009) and Li (2000). We model the portfolio loss as a mixture of Bernoulli random variables (Joe 1997) by assuming a constant percentage loss given default (LGD):

$$R_{ij}(z_{ij}) = \begin{cases} 0 & \text{if } z_{ij} > \theta_{ij}, \\ c & \text{if } z_{ij} \leqslant \theta_{ij}, \end{cases} \tag{2.4}$$

where $0 < c \leqslant 1$. As a consequence, $0 \leqslant R_\pi \leqslant c$. Because the idiosyncratic factor is unit-normally distributed, the conditional probability of default for company j in sector i is

$$\Pr(z_{ij} < \theta_{ij} \mid \hat{\varepsilon}, \varepsilon_i) = p_{ij}(\varepsilon_i, \hat{\varepsilon}) = \Phi\left[\frac{\theta_{ij} - \sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i)}{\sqrt{1-\rho_{ij}}}\right], \tag{2.5}$$

where $\Phi(z) = [1 + \operatorname{erf}(z/\sqrt{2})/2]$. Throughout this paper, we use $\phi(z) = \Phi'(z) = \exp[-z^2/2]/\sqrt{2\pi}$. As a simple corollary to (2.5), since

$$\int_{-\infty}^{\infty} \Phi(a + b\varepsilon)\phi(\varepsilon)\,\mathrm{d}\varepsilon = \Phi\left(\frac{a}{\sqrt{1+b^2}}\right)$$

for constants a and b, the (unconditional) default probability of the jth company in sector i is

$$\begin{aligned} \mathrm{PD}_{ij} &= \Pr(z_{ij} < \theta_{ij}) \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \Phi\left[\frac{\theta_{ij} - \sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i)}{\sqrt{1-\rho_{ij}}}\right]\phi(\hat{\varepsilon})\phi(\varepsilon_i)\,\mathrm{d}\hat{\varepsilon}\,\mathrm{d}\hat{\varepsilon}_i, \\ &= \Phi(\theta_{ij}). \end{aligned} \tag{2.6}$$

Hence, the default probability is uniquely determined by specifying the threshold value θ_{ij}.

Implementation of the factor model now requires knowledge of the numerical values of β_{ij}, $\hat{\beta}_{ij}$, ρ_{ij} and θ_{ij}. The exposures w_{ij} can be chosen by the user of the model or taken from the call reports of banks, which can be found on the website of the Federal Deposit Insurance Corporation (FDIC). The θ_{ij} are related to the default probability through (2.6), and the default probabilities can be inferred from the credit rating of a firm (Crouhy *et al* 2000). The correlation terms $\hat{\beta}_{ij}$, $\hat{\beta}_{ij}$ and ρ_{ij} are more difficult to obtain, but correlation matrixes can usually be found empirically through historical data (see Andersen *et al* (2003) for more details).

3 LOSS DISTRIBUTION FOR THE PORTFOLIO

The main result in this section is a derivation for the VaR for R_π in (2.3). In our analysis, we assume that the number of sectors is large: $N \gg 1$, and the positive exposures $w_{ij} = O(N^{-1})$ as $N \to \infty$. Although the $N \gg 1$ assumption may be somewhat unrealistic, the value of N does depend on how the loans are grouped together. If historical data and economic intuition can allow a different grouping with a larger N, this would lead to a more accurate model, according to the analysis in this paper. The number of companies per sector n is typically very large (perhaps $n \gtrsim 1000$), but this is not a requirement for our approximations to hold.

We first provide some preliminary results for the moments of R_π before proving the main theorem.

LEMMA 3.1 *Let*

$$R_\pi = \sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij} R_{ij}(z_{ij}),$$

where $R_{ij}(\cdot)$ satisfies (2.4). Furthermore, let $Y_i = \sum_{j=1}^{n} w_{ij} R_{ij}(z_{ij})$, $i = 1, \dots, N$, be the sector losses, so that $R_\pi = \sum_{i=1}^{N} Y_i$, and let the exposures $w_{ij} = O(N^{-1})$ as $N \to \infty$. Define the conditional moments

$$\mu_i(\hat{\varepsilon}, \varepsilon_i) = E[Y_i \mid \hat{\varepsilon}, \varepsilon_i], \tag{3.1}$$

$$\sigma_i^2(\hat{\varepsilon}, \varepsilon_i) = V[Y_i \mid \hat{\varepsilon}, \varepsilon_i], \tag{3.2}$$

$$m_i(\hat{\varepsilon}) = E[Y_i \mid \hat{\varepsilon}], \tag{3.3}$$

$$s_i^2(\hat{\varepsilon}) = V[Y_i \mid \hat{\varepsilon}], \tag{3.4}$$

$$m(\hat{\varepsilon}) = E[R_\pi \mid \hat{\varepsilon}], \tag{3.5}$$

$$s^2(\hat{\varepsilon}) = V[R_\pi \mid \hat{\varepsilon}]. \tag{3.6}$$

Then, as $N \to \infty$, the conditional means satisfy

$$\mu_i(\hat{\varepsilon}, \varepsilon_i) = c \sum_{j=1}^{n} w_{ij} \Phi\biggl[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) - \sqrt{\rho_{ij}}(\hat{\varepsilon}\hat{\beta}_{ij} + \varepsilon_i \beta_{ij})}{\sqrt{1-\rho_{ij}}}\biggr] = O(N^{-1}), \tag{3.7}$$

$$m_i(\hat{\varepsilon}) = c \sum_{j=1}^{n} w_{ij} \Phi\biggl[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) - \sqrt{\rho_{ij}}\hat{\varepsilon}\hat{\beta}_{ij}}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\biggr] = O(N^{-1}), \tag{3.8}$$

$$m(\hat{\varepsilon}) = c \sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij} \Phi\biggl[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) - \sqrt{\rho_{ij}}\hat{\varepsilon}\hat{\beta}_{ij}}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\biggr] = O(1), \tag{3.9}$$

while the conditional variances scale as

$$\sigma_i^2(\varepsilon_i, \hat{\varepsilon}) = O(N^{-2}), \tag{3.10}$$

$$s_i^2(\hat{\varepsilon}) = O(N^{-2}), \tag{3.11}$$

$$s^2(\hat{\varepsilon}) = O(N^{-1}). \tag{3.12}$$

Proof Conditioned on $\hat{\varepsilon}$ and ε_i, the $R_{ij}(z_{ij})$ are independent (scaled) Bernoulli random variables, so (3.7) follows from

$$\begin{aligned}\mu_i(\hat{\varepsilon}, \varepsilon_i) &= \sum_{j=1}^{n} w_{ij} E[R_{ij}(z_{ij})] = c \sum_{j=1}^{n} w_{ij} p_{ij} \\ &= c \sum_{j=1}^{n} w_{ij} \Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) - \sqrt{\rho_{ij}}(\hat{\varepsilon}\hat{\beta}_{ij} + \varepsilon_i \beta_{ij})}{\sqrt{1-\rho_{ij}}}\right] = O(N^{-1}),\end{aligned}$$

using (2.5) and (2.6). Equation (3.8) follows from

$$m_i(\hat{\varepsilon}) = \int_{-\infty}^{\infty} \mu_i(\hat{\varepsilon}, \varepsilon_i) \phi(\varepsilon_i)\, \mathrm{d}\varepsilon_i,$$

and (3.9) follows from $m(\hat{\varepsilon}) = \sum_{i=1}^{N} m_i(\hat{\varepsilon})$. Now we prove (3.10)–(3.12). Conditioned on $\hat{\varepsilon}$ and ε_i, the $R_{ij}(z_{ij})$ are again independent (scaled) Bernoulli random variables. With the shorthand

$$p_{ij}(\hat{\varepsilon}, \varepsilon_i) = \Phi\left[\frac{\theta_{ij} - \sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i)}{\sqrt{1-\rho_{ij}}}\right],$$

Equation (3.10) follows from

$$\sigma_i^2(\hat{\varepsilon}, \varepsilon_i) = c^2 \sum_{j=1}^{n} w_{ij}^2 p_{ij}(\hat{\varepsilon}, \varepsilon_i)[1 - p_{ij}(\hat{\varepsilon}, \varepsilon_i)] = O(N^{-2}),$$

while (3.11) follows from

$$\begin{aligned}s_i^2(\hat{\varepsilon}) &= E[Y_i^2 \mid \hat{\varepsilon}] - E[Y_i \mid \hat{\varepsilon}]^2 \\ &= \int_{-\infty}^{\infty} E[Y_i^2 \mid \hat{\varepsilon}, \varepsilon_i] \phi(\varepsilon_i)\, \mathrm{d}\varepsilon_i - \left(\int_{-\infty}^{\infty} E[Y_i \mid \hat{\varepsilon}, \varepsilon_i] \phi(\varepsilon_i)\, \mathrm{d}\varepsilon_i\right)^2 \\ &= c^2 \sum_{k=1}^{n} \sum_{j=1}^{n} w_{ij} w_{ik} \int_{-\infty}^{\infty} p_{ij}(\hat{\varepsilon}, \varepsilon_i) p_{ik}(\hat{\varepsilon}, \varepsilon_i) \phi(\varepsilon_i)\, \mathrm{d}\varepsilon_i \\ &\quad - \left\{c \sum_{j=1}^{n} w_{ij} \int_{-\infty}^{\infty} p_{ij}(\hat{\varepsilon}, \varepsilon_i) \phi(\varepsilon_i)\, \mathrm{d}\varepsilon_i\right\}^2 \\ &= O(N^{-2}),\end{aligned}$$

and, therefore, (3.12) immediately follows, since $s^2(\hat{\varepsilon}) = \sum_{i=1}^{N} s_i^2(\hat{\varepsilon})$. □

THEOREM 3.2 *Let the normalized asset return of the j th company in sector i follow a hierarchical multifactor model so that the correlated random variables z_{ij} satisfy*

$$z_{ij} = \sqrt{\rho_{ij}}(\hat{\beta}_{ij}\hat{\varepsilon} + \beta_{ij}\varepsilon_i) + \sqrt{1-\rho_{ij}}\zeta_{ij}, \quad i = 1, \dots, N, \ j = 1, \dots, n, \quad (3.13)$$

where $0 < \rho_{ij} < 1$, $0 < \beta_{ij}, \hat{\beta}_{ij} < 1$, $\beta_{ij}^2 + \hat{\beta}_{ij}^2 = 1$ *and* $\varepsilon_i, \hat{\varepsilon} \sim \mathcal{N}(0, 1)$. *Consider a portfolio*

$$R_\pi = \sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij} R_{ij}(z_{ij}), \quad (3.14)$$

with exposures w_{ij} such that $\sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij} = 1$ and $w_{ij} = O(N^{-1})$ as $N \to \infty$, where $R_{ij}(z_{ij})$ follows (2.4): ie, for some $-\infty < \theta_{ij} < \infty$, the company with an asset return that follows (3.13) defaults when $z_{ij} < \theta_{ij}$, incurring a loss c. Then the VaR (VaR_q) of the portfolio R_π at risk level $0 \leqslant q \leqslant 1$ satisfies the asymptotic relation

$$\mathrm{VaR}_q \sim c\sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij}\Phi\biggl[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q)}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\biggr] \quad (3.15)$$

as $N \to \infty$, where

$$\mathrm{VaR}_q = \inf\{x : q \leqslant F_{R_\pi}(x)\}, \quad (3.16)$$

and $F_{R_\pi}(\cdot)$ is the cumulative density function of R_π.

PROOF We write the loss of the portfolio as

$$R_\pi = \sum_{i=1}^{N} Y_i, \quad Y_i = \sum_{j=1}^{n} w_{ij} R_{ij}(z_{ij}), \quad (3.17)$$

so that Y_i is the total loss of all companies in sector i, and R_π is the loss summed over all sectors. Conditioned on the global risk $\hat{\varepsilon}$, Y_i, $i = 1, \dots, N$, are independent random variables and, from the central limit theorem, their sum follows a normal distribution when $N \gg 1$:

$$R_\pi \equiv \sum_{i=1}^{N} Y_i \sim N[m(\hat{\varepsilon}), s^2(\hat{\varepsilon})],$$

where $m(\cdot)$ is given by (3.9) and s^2 obeys the scaling (3.12) (its explicit form is not required to derive (3.15)). The density for the loss of the entire portfolio is given by the law of total probability:

$$f_{R_\pi}(L) \sim \sqrt{N}\int_{-\infty}^{\infty}\frac{\phi(\hat{\varepsilon})}{\sqrt{2\pi\hat{s}^2(\hat{\varepsilon})}}\exp\biggl\{\frac{-N[L-m(\hat{\varepsilon})]^2}{2\hat{s}^2(\hat{\varepsilon})}\biggr\}\,\mathrm{d}\hat{\varepsilon} \quad (3.18)$$

as $N \to \infty$, where we have set $s^2 = \hat{s}^2/N$ in light of (3.12).

We now apply Laplace's method to the integral in (3.18). Laplace's method provides a way to approximate integrals that contain a large parameter by analyzing the stationary point of the integrand. An overview of the method is given in Appendix A. In particular, we refer to (A.2), which approximates integrals of the form

$$\int_a^b g(t)\exp[-k\psi(t)]\,\mathrm{d}t$$

for functions g and ψ, with $a, b \in \mathbb{R}$ and $k \gg 1$. Taking $\psi(\hat{\varepsilon}) = (L - m(\hat{\varepsilon}))^2/2s^2(\hat{\varepsilon})$ and $g(\hat{\varepsilon}) = \phi(\hat{\varepsilon})/\sqrt{2\pi s^2(\hat{\varepsilon})}$, the stationary point $\hat{\varepsilon}^*$ satisfies $\hat{\varepsilon}^* = m^{-1}(L)$ and $\psi''(\hat{\varepsilon}^*) = [m'(\hat{\varepsilon}^*)/\hat{s}(\hat{\varepsilon}^*)]^2$, so that

$$f_{R_\pi}(L) \sim \frac{\phi(\hat{\varepsilon}^*)}{|m'(\hat{\varepsilon}^*)|}, \quad N \to \infty. \tag{3.19}$$

The cumulative density is

$$F_{R_\pi}(L) \sim \int_0^L \frac{\phi[\hat{\varepsilon}^*(L')]}{-m'[\hat{\varepsilon}^*(L')]}\,\mathrm{d}L' = \Phi[-m^{-1}(L)]. \tag{3.20}$$

From (3.16), VaR_q is just the inverse of the cumulative density function; ie, for a given confidence level $0 \leqslant q \leqslant 1$, the VaR_q is found by setting $F_{R_\pi}(L) = q$ and solving for L:

$$L = m[-\Phi^{-1}(q)] \Rightarrow \mathrm{VaR}_q \sim c\sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij}\Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q)}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\right] \tag{3.21}$$

as $N \to \infty$. □

Equation (3.21) is the main contribution of this paper. What are the errors associated with this asymptotic approximation? There are actually two contributions. One is associated with approximating the density of R_π, conditioned on $\hat{\varepsilon}$, with a Gaussian through the central limit theorem. The other is associated with applying Laplace's method to the integral in (3.18): see the higher-order term in (A.2). The first error appears as an extra term under the integral in (3.18): for large but finite N, $f_{R_\pi|\hat{\varepsilon}}$ would actually take the form of a Gaussian plus a small correction (Berry 1941; Esseen 1942). The second gives rise to an additive $O(N^{-1})$ term in (3.19)–(3.21). By comparing our results for VaR with Monte Carlo simulations in Figure 2 on page 55, we find that the dominant error term is, in fact, $O(N^{-1})$:

$$\mathrm{VaR}_q = c\sum_{i=1}^{N}\sum_{j=1}^{n} w_{ij}\Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q)}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\right] + O(N^{-1}) \tag{3.22}$$

as $N \to \infty$. Therefore, either the two contributions are of the same order or the error incurred by using Laplace's method is dominant. Equation (3.22) is confirmed numerically by comparing the double sum with Monte Carlo simulations of the VaR, $\mathrm{VaR}_q^{(\mathrm{num})}$. The error was measured using the infinity norm over $0 \leqslant q \leqslant 0.99$:

$$\text{error} = \max_{0 \leqslant q \leqslant 0.99} \left| c \sum_{i=1}^{N} \sum_{j=1}^{n} w_{ij} \Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q)}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\right] - \mathrm{VaR}_q^{(\mathrm{num})} \right|. \tag{3.23}$$

The portfolios compared in Figure 2 are completely homogeneous (and, therefore, somewhat artificial), but their purpose is to provide a benchmark result to deduce the scaling of the error. Note there is no scaling with n: the errors in part (b) of Figure 2 on the facing page arise from N being finite.

Once the approximation to the VaR has been obtained, related quantities such as the expected shortfall (ES) and ES contribution are easily approximated as

$$\begin{aligned} \mathrm{ES}_q &= \frac{1}{1-q}\int_q^1 \mathrm{VaR}_{q'}\, \mathrm{d}q' \\ &\sim \sum_{i=1}^{N}\sum_{j=1}^{n} \frac{c w_{ij}}{1-q}\int_q^1 \Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q')}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\right]\mathrm{d}q', \qquad (3.24) \\ \mathrm{ES}_{q,ij} &= w_{ij}\frac{\partial \mathrm{ES}_q}{\partial w_{ij}} \\ &\sim \frac{c w_{ij}}{1-q}\int_q^1 \Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_{ij}) + \sqrt{\rho_{ij}}\hat{\beta}_{ij}\Phi^{-1}(q')}{\sqrt{1-\rho_{ij}+\beta_{ij}^2\rho_{ij}}}\right]\mathrm{d}q' \qquad (3.25) \end{aligned}$$

as $N \to \infty$.

Now consider the case where the parameters β_{ij}, ρ_{ij} and θ_{ij} are only sector dependent: $\beta_{ij} = \beta_i$, $\rho_{ij} = \rho_i$ and $\theta_{ij} = \theta_i$ for the sector index i. Then all companies within a sector are statistically identical and we can consider a distribution of exposures that only depends on the sector. By defining $\hat{w}_i = \sum_{j=1}^{n} w_{ij}$, (3.22) reduces to

$$\mathrm{VaR}_q = c\sum_{i=1}^{N}\hat{w}_i\Phi\left[\frac{\Phi^{-1}(\mathrm{PD}_i) + \sqrt{\rho_i}\hat{\beta}_i\Phi^{-1}(q)}{\sqrt{1-\rho_i\hat{\beta}_i^2}}\right] + O(N^{-1}), \tag{3.26}$$

where $\sum_{i=1}^{N}\hat{w}_i = 1$ and $\mathrm{PD}_i = \Phi(\theta_i)$ (cf. (2.6)). It is easy to show that when $n = 1$ (recall that n does not have to be large for our approximations to be valid), (3.26) is identical to the VaR for N firms coupled with a single factor, with rescaled systematic and idiosyncratic risks. In fact, (2.1) implies that the asset return of each firm is given by

$$z_{i1} = \sqrt{r_i}\hat{\varepsilon} + \sqrt{1-r_i}\eta_i, \quad r_i = \rho_{i1}\hat{\beta}_{i1}^2,\ i = 1,\dots,N, \tag{3.27}$$

FIGURE 2 Error of analytic approximation (3.21), as defined by (3.23), scales as $O(N^{-1})$ when compared with Monte Carlo simulation and is independent of the number of companies per sector n.

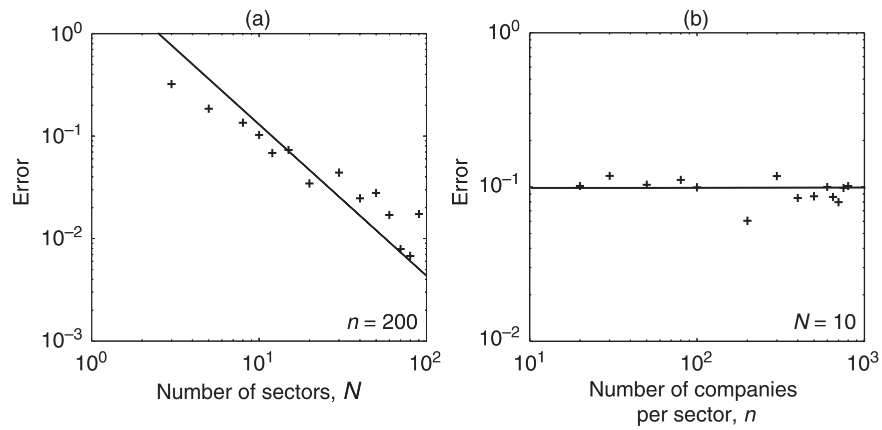

Solid line has slope −1 in (a) and 0 in (b). Parameters used were $c = 2$, $w_{ij} = 1/(Nn)$, $\beta_{ij} = 0.8$, $\theta_{ij} = -1.3$ and $\rho_{ij} = 0.7$. Monte Carlo VaRs are found from 20 000 realizations.

with $\sqrt{1 - r_i}\eta_i \equiv \sqrt{\rho_{i1}}\beta_{i1}\varepsilon_1 + \sqrt{1 - \rho_{i1}}\zeta_{i1}$ and $\eta_i \sim N(0, 1)$. Application of Vasicek's formula to (3.27) then yields (3.26). The general case with n identical companies in each sector also collapses to the single-factor case because the first firm is representative of all firms in that sector, and its asset return is described by (3.27). When the sectors are homogeneous, the hierarchical multifactor model can, in fact, be treated as a single-factor model.

4 RESULTS AND DISCUSSION

To validate our analytic approximation (3.21), we construct a proxy portfolio using exposure data taken from the call report of a large bank (JP Morgan 2014) with $N = 17$ "sectors" (see Table 1 on the next page). In this example, the sectors correspond to different types of institution that borrow from the bank. For each of the sectors, we estimate sector default probabilities $\overline{\mathrm{PD}}_i$ and variances of systematic risk $\bar{\rho}_i$. In this example, we also assume that the factor loadings are constant for every company. The $\overline{\mathrm{PD}}_i$ are estimated by assuming that large government institutions and corporations are less likely to default than small companies and consumers. More accurate values could come from the credit rating of these entities (Crouhy *et al* 2000).

TABLE 1 Proxy credit portfolio motivated by exposures taken from a JP Morgan call report.

i	Loan sector	Exposure (US$)	$\bar{w}_i$ (%)	$\overline{\mathrm{PD}}_i$	$\bar{\rho}_i$
1	Construction, land development	3 815	0.60	10^{-2}	0.23
2	Farmland	211	0.03	10^{-2}	0.22
3	**1–4 residential properties**	203 246	**32.17**	10^{-2}	0.25
4	5+ residential properties	45 090	7.14	10^{-2}	0.21
5	Nonfarm, nonresidential	27 153	4.30	10^{-2}	0.27
6	Commercial US banks	3 157	0.50	10^{-3}	0.15
7	Banks in foreign countries	18 933	3.00	10^{-3}	0.13
8	Agricultural loans	788	0.12	10^{-2}	0.12
9	**US commercial/industrial**	90 879	**14.39**	10^{-2}	0.18
10	Non-US commericial/industrial	33 624	5.32	10^{-2}	0.28
11	Credit cards	26 189	4.15	0.05	0.15
12	Other revolving credit plans	2 584	0.41	0.05	0.17
13	Automobile loans	41 517	6.57	0.05	0.21
14	Other consumer loans	19 837	3.14	0.05	0.16
15	Foreign governments	1 031	0.16	10^{-4}	0.12
16	US states and subdivisions	12 680	2.01	10^{-4}	0.17
17	**Nondepository/other institutions**	101 000	**15.99**	10^{-2}	0.15
	Total	631 734	100.00		

The three largest obligors are highlighted in bold face. Exposure dollar amounts are in millions.

From the sector parameters $\bar{\rho}_i$, $\overline{\mathrm{PD}}_i$ and $\bar{w}_i$, we generate firm-level parameters by

$$\rho_{ij} = \bar{\rho}_i + \delta\rho_{ij}, \tag{4.1}$$
$$\theta_{ij} = \Phi^{-1}(\overline{\mathrm{PD}}_i) + \delta\theta_{ij}, \tag{4.2}$$
$$\beta_{ij} = \text{constant}, \tag{4.3}$$
$$w_{ij} = \bar{w}_i/n, \tag{4.4}$$

for $i = 1, \dots, N$, $j = 1, \dots, n$, with $\delta\rho_{ij} \sim N(0, 10^{-4})$ and $\delta\theta_{ij} \sim N(0, 10^{-4})$. Therefore, each of the seventeen loans in Table 1 is subdivided equally into $n = 1000$ subloans, with corresponding exposure $\bar{w}_i/n$. At the sector level, this portfolio is quite typical in the sense that the distribution of exposures is "lumpy", with the portfolio being dominated by a few large loans – in this case, to residential real estate (32%), US commercial and industrial companies (14%) and nondepository/other institutions (16%).

Because the firm-level parameters are small perturbations of the sector-level parameters in (4.1)–(4.4), the correlation between two companies in sector i is approximately given by $\bar{\rho}_i$, and the correlation between companies in sector i and j $(i \neq j)$ is approximately $\sqrt{\bar{\rho}_i \bar{\rho}_j}\,\hat{\beta}_i \hat{\beta}_j$. The loading factors $\hat{\beta}_i$ essentially control cross-sector correlations. For the parameters in Table 1 on the facing page, when $\hat{\beta}_i = 0.95$, firms in different sectors are correlated at between 11% and 25%; when $\hat{\beta}_i = 0.87$, they are correlated at between 9% and 21%; and when $\hat{\beta}_i = 0.6$, they are correlated at between 4% and 10%.

We now compare the VaRs predicted by (3.21) with portfolio losses generated by drawing random variables defined by (2.1). Our Monte Carlo simulations use $n = 1000$ companies per sector and 1000 trials to simulate the value of the portfolio at some risk level q. In Figure 3 on the next page, we see that the agreement is good, provided the β_{ij} are not too large. As β_{ij} increases from 0.3 via 0.5 to 0.8, our analytic approximation (3.21) becomes worse, particularly for smaller values of q. The error bars for the Monte Carlo simulated VaR represent 99% confidence intervals. We see that for $\beta_{ij} \equiv 0.3$ and 0.5, the analytic solution is within the intervals for $q = 0.2, 0.4, 0.6$ and 0.8. For large factor loading $\beta_{ij} \equiv 0.8$, there is a significant departure from the Monte Carlo simulations, especially when $q \lesssim 0.6$.

The portfolio in Table 1 on the facing page is fairly homogeneous in terms of the systematic risk variance ρ_i. Many portfolios of interest are more inhomogeneous in that they contain a few companies or sectors with defaults that are very strongly correlated, while the defaults of most companies are only weakly correlated. In Table 2 on page 59, we sharply increase the value of $\bar{\rho}_i$ for large institutions in the portfolio. Again, we test the analytic VaR of (3.21) against Monte Carlo simulations when $\beta_{ij} = 0.3, 0.5$ and 0.8: see Figure 4 on page 61 (note the β_{ij} in Table 2 on page 59 are not used for these results; they are used in the next set of simulations described below). As in the first portfolio, the agreement is good when $\beta_{ij} = 0.3$ or 0.5. In this portfolio, the firm–firm correlations have a wide range, spanning between 14% and 97% when $\beta_{ij} = 0.3$, between 12% and 97% when $\beta_{ij} = 0.5$ and between 6% and 97% when $\beta_{ij} = 0.8$. The strongest correlations are, of course, between companies in the same sector, and obligors belonging to US states and subdivisions are the most strongly correlated in this portfolio. These institutions are responsible for the upper bound of approximately 97% in the correlation matrix since $\rho_{16} = 0.97$ in Table 2 on page 59.

In part (a) of Figure 5 on page 62, we plot the VaR for the portfolio in Table 2 on page 59, relax the constant β_{ij} assumption and instead generate them at the firm level through

$$\beta_{ij} = \bar{\beta}_i + \delta\beta_{ij}, \tag{4.5}$$

FIGURE 3 Comparison of VaRs generated through Monte Carlo simulation (dashed gray) and analytic approximation (solid black) as determined by (3.21).

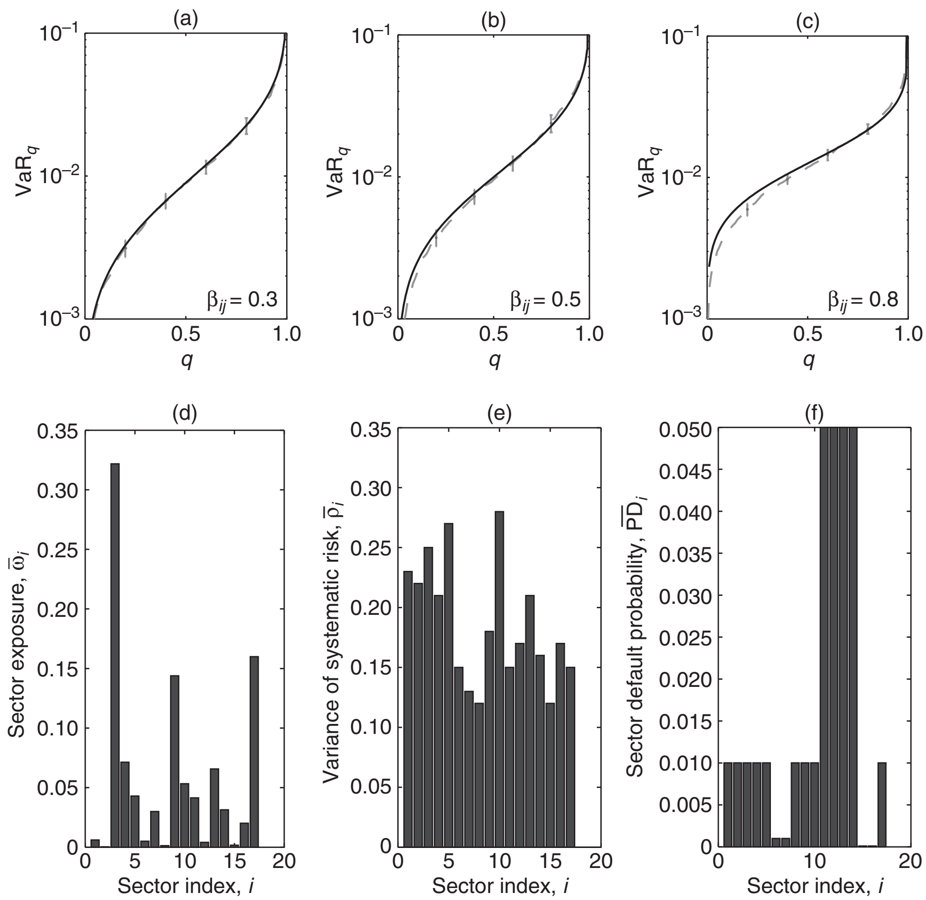

Error bars for the Monte Carlo simulations are 99% confidence intervals derived from 50 000 bootstrap samples. Parameters are taken from Table 1 on page 56 with intrasector correlations given by $\bar{\rho}_i$. Cross-sector correlations are 11–25% for (a), 9–21% for (b) and 4–10% for (c). The bottom panels show the distribution of (d) sector exposures, (e) systematic risk variances and (f) default probabilities.

where $\bar{\beta}_i$ are taken from Table 2 on the facing page and $\delta\beta_{ij} \sim \mathcal{N}(0, 10^{-4})$. We choose the loading factors $\bar{\beta}_i$ to be closer to zero for loans with larger exposures. Hence, the three largest loans are assigned $\bar{\beta}_i = 0.05$, loans with exposures between about 2% and 7% are assigned $\bar{\beta}_i = 0.1$ and the remaining loans are assigned values from 0.8 to 0.9. The rationale is that we wish to mimic a portfolio containing a few large loans with defaults that may be highly correlated with respect to global risk. Therefore, we

TABLE 2 A more strongly correlated proxy credit portfolio motivated by exposures taken from a JP Morgan call report.

i	Loan sector	Exposure (US$)	$\bar{w}_i$ (%)	$\overline{\mathrm{PD}}_i$	$\bar{\rho}_i$	$\bar{\beta}_i$
1	Construction, land development	3 815	0.60	10^{-2}	0.33	0.8
2	Farmland	211	0.03	10^{-2}	0.32	0.9
3	**1–4 residential properties**	203 246	**32.17**	10^{-2}	0.35	0.05
4	5+ residential properties	45 090	7.14	10^{-2}	0.31	0.1
5	Nonfarm, nonresidential	27 153	4.30	10^{-2}	0.37	0.1
6	Commercial US banks	3 157	0.50	10^{-3}	0.95	0.9
7	Banks in foreign countries	18 933	3.00	10^{-3}	0.93	0.1
8	Agricultural loans	788	0.12	10^{-2}	0.52	0.7
9	**US commercial/ industrial**	90 879	**14.39**	10^{-2}	0.48	0.05
10	Non-US commercial/ industrial	33 624	5.32	10^{-2}	5×10^{-2}	0.1
11	Credit cards	26 189	4.15	5×10^{-2}	0.15	0.1
12	Other revolving credit plans	2 584	0.41	5×10^{-2}	0.17	0.8
13	Automobile loans	41 517	6.57	0.05	0.21	0.1
14	Other consumer loans	19 837	3.14	0.05	0.16	0.1
15	Foreign governments	1 031	0.16	10^{-4}	0.92	0.8
16	US states and subdivisions	12 680	2.01	10^{-4}	0.97	0.1
17	**Nondepository/ other institutions**	101 000	**15.99**	10^{-2}	0.85	0.05
	Total	631 734	100.00			

The three largest obligors are highlighted in bold face. Exposure dollar amounts are in millions.

choose the loading factors so that the corresponding firms are more tightly coupled with $\hat{\varepsilon}$. Now, with β_{ij} stochastically generated through (4.5), the agreement with the Monte Carlo simulated VaR is excellent for a large range of q values. The reason for this could be that, even though there are six sectors (construction, farmland, commercial US banks, agriculture, revolving credit plans and foreign governments) that have large $\bar{\beta}_i$, the average factor loading across all companies in the portfolio is

TABLE 3 Portfolio containing $N = 8$ sectors, based on grouping together loans from Table 2 on the preceding page.

i	Loan sector	$\bar{w}_i$ (%)	$\overline{\mathrm{PD}}_i$	$\bar{\rho}_i$	$\bar{\beta}_i$
1	Real estate	44.25	10^{-2}	0.33	0.39
2	Depository institutions and banks	3.5	10^{-3}	0.94	0.5
3	Agricultural loans	0.12	10^{-2}	0.52	0.7
4	Commercial/industrial loans	19.71	10^{-2}	0.49	0.075
5	Consumer loans	14.27	5×10^{-2}	0.32	0.45
6	Foreign governments	0.16	10^{-4}	0.92	0.8
7	US states and subdivisions	2.00	10^{-4}	0.97	0.1
8	Nondepository/other institutions	15.99	10^{-2}	0.85	0.05
	Total	100.00			

small. We have seen from Figure 3 on page 58 and Figure 4 on the facing page that our analytic solution performs better when firm factor loadings β_{ij} are small.

So far, our results have concentrated on a portfolio with $N = 17$ sectors. As discussed before, the value of N depends on how loans are classified. Is the approximation (3.21) still accurate when N is reduced? In Table 3, we reduce N by grouping together loans that are economically similar. For example, we group commercial US bank loans and loans to foreign banks into a single "depository institutions and banks" sector. This results in a new portfolio with $N = 8$ sectors, with some of the new sectors encompassing several of the old sectors in the $N = 17$ portfolio. The new sector exposures $\bar{w}_i$ are sums of the exposures in the old portfolio and the new $\overline{\mathrm{PD}}_i$, $\bar{\rho}_i$ and $\bar{\beta}_i$ are averages of the parameters in the old portfolio. In part (b) of Figure 5 on page 62, we see that, although the analytic approximation (3.21) becomes worse for smaller values of q, it still lies within the 99% confidence intervals for $q \gtrsim 0.4$. In particular, even though $N = 8$ is not "large", the agreement between the analytic solution and simulation results is still excellent for values of q close to 1.

We now compare the ES for Monte Carlo simulated portfolios and the analytic approximation (3.24). The analytic approximation is computed using a compound trapezoid rule with 2501, 5001 and 10 001 abscissae for $q = 0.95, 0.90$ and 0.80, respectively. The Monte Carlo ES, ES_q, is computed by finding the mean loss conditioned on the loss being larger than VaR_q:

$$\mathrm{ES}_q = E[R_\pi \mid R_\pi > \mathrm{VaR}_q]. \tag{4.6}$$

From Table 4 on page 63, we see that our analytic approximation generally does a reasonable job of predicting the ES for portfolios with small β_{ij}. For portfolios 1,

FIGURE 4 Comparison of VaRs generated through Monte Carlo simulation (dashed gray) and analytic approximation (solid black) as determined by (3.21).

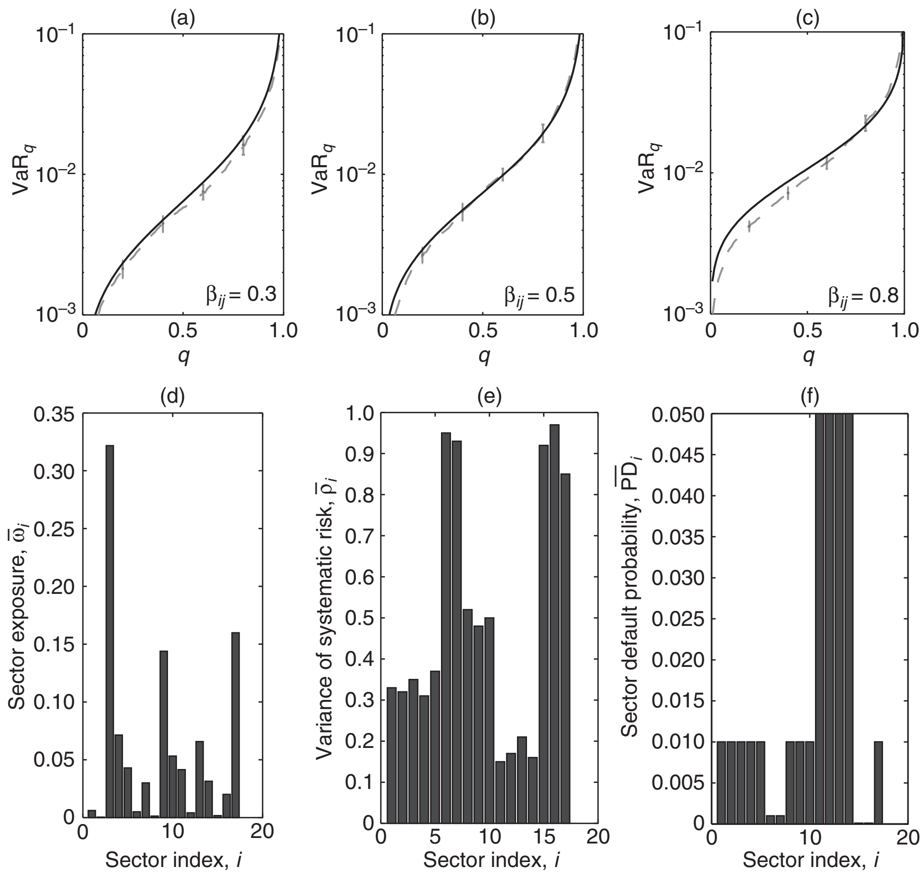

Error bars for the Monte Carlo simulations are 99% confidence intervals derived from 50 000 bootstrap samples. Parameters are taken from Table 2 on page 59 with intrasector correlations given by $\bar{\rho}_i$. Cross-sector correlations are 14–87% for (a), 12–72% for (b) and 8–49% for (c). The bottom panels show the distribution of (d) sector exposures, (e) systematic risk variances and (f) default probabilities.

2 and 4, the relative error in ES is between 5% and 10%. When $\beta_{ij} = 0.8$, the ES disagree by up to 25%.

Finally, we try to formulate some guidelines for when our analytic approximation (3.21) is valid; see Table 5 on page 64. This table provides a range of factor loadings β_{ij} and risk variances ρ_{ij}, within which the agreement between the analytic solution (3.21) and Monte Carlo simulated VaR is good. Specifically, a check mark in Table 5 on page 64 indicates that the analytic solution lies within the 99% confidence intervals

FIGURE 5 VaR for portfolios containing (a) $N = 17$ and (b) $N = 8$ sectors, corresponding to Table 2 on page 59 and Table 3 on page 60.

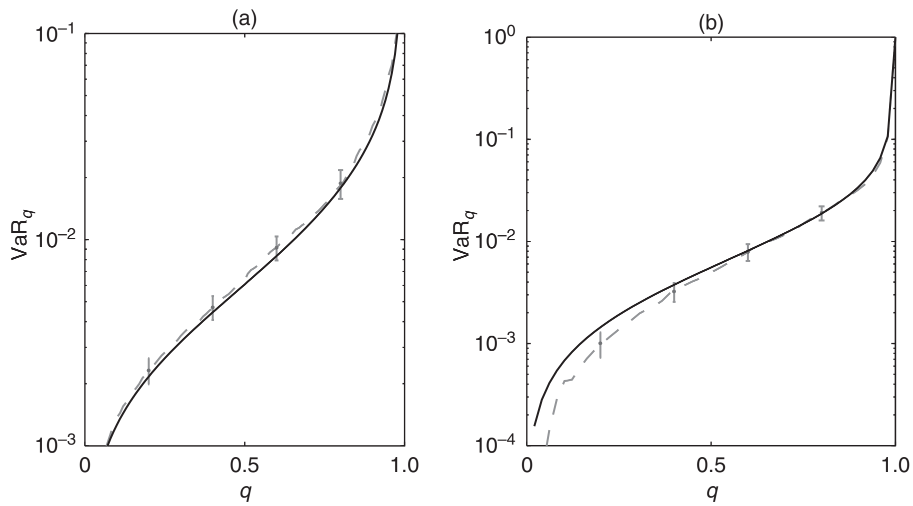

Solid black curve indicates analytic solution (3.21), dashed gray curve indicates Monte Carlo simulated values and error bars are 95% confidence intervals generated using 50 000 bootstrap samples.

of the Monte Carlo solution at $q = 0.1, 0.3, 0.5$ and 0.7. The parameter values were generated using (4.1)–(4.4), with $\bar{\rho}_i = \bar{\rho}_0 + 0.05Z_i$, where $Z_i \sim \mathcal{N}(0, 1)$. We set β_{ij} to be constant across all companies and sectors and the sector default probabilities $\overline{\mathrm{PD}}_i = 0.0102$ for all i. The sector exposures $\bar{w}_i$ in (4.4) are randomly generated, and there are $n = 1000$ companies per sector. We find that, for portfolios with seventeen sectors or more, our analytic approximation agrees with the Monte Carlo simulations, provided the asset returns are no larger than about 0.3 and the loading factors β_{ij} are no larger than about 0.7.

5 CONCLUSIONS

In this paper, we computed the VaRs and ESs for a bond portfolio under a hierarchical multifactor model for the asset returns. Our main results are (3.21) and (3.24), which are analytic approximations to the portfolio's VaR and ES, given a risk level $0 \leqslant q \leqslant 1$. Our approximation to the VaR is written entirely in terms of easy-to-compute special and elementary functions and represents an economically intuitive extension of Vasicek's ASRF result to multiple sectors. It is much quicker to compute than predicting the VaR_q through many trials of a Monte Carlo simulation. Our formulas for the VaR_q give good approximations to the loss as predicted by Monte

TABLE 4 Comparison of ES at level q (average loss in the worst $100(1-q)\%$ of cases).

Portfolio 1: $\beta_{ij} = 0.3$			
	q **= 0.95**	q **= 0.90**	q **= 0.80**
ES_q (analytic)	0.1172	0.0795	0.0519
ES_q (MC)	0.1237	0.0848	0.0545
Portfolio 2: $\beta_{ij} = 0.5$			
	q **= 0.95**	q **= 0.90**	q **= 0.80**
ES_q (analytic)	0.1019	0.0720	0.0487
ES_q (MC)	0.1123	0.0800	0.0527
Portfolio 3: $\beta_{ij} = 0.8$			
	q **= 0.95**	q **= 0.90**	q **= 0.80**
ES_q (analytic)	0.0631	0.0497	0.0377
ES_q (MC)	0.0841	0.0611	0.0433
Portfolio 4: random β_{ij}			
	q **= 0.95**	q **= 0.90**	q **= 0.80**
ES_q (analytic)	0.1241	0.0825	0.0531
ES_q (MC)	0.1396	0.0934	0.0591

Portfolios 1, 2 and 3 have $\beta_{ij} \equiv 0.3, 0.5, 0.8$, respectively, and other firm parameters are generated through (4.1), (4.2) and (4.4). Portfolio 4 has β_{ij} generated through (4.5). All values are correct to four decimal places.

Carlo simulations when the sector factor loadings and systematic risk variances are not too close to 1. When the sector factor loadings are increased, we find that our approximations deviate from Monte Carlo simulated VaR_q for small values of risk level q.

Our approximations are able to account for asset-return correlations among companies at a global and sector level. We derived the formulas by using the central limit theorem and Laplace's method to approximate the loss distribution conditioned on the global risk when the number of sectors N is large, and then integrating over the global risk. The final analytic approximations have a similar mathematical structure to the ASRF but explicitly feature local and global factor loadings.

Although we have given some guidelines, in terms of model parameters, for when the analytic solution (3.21) may be accurate, quantifying and understanding its accu-

TABLE 5 Comparison of analytic solution with Monte Carlo simulated VaR for $N = 13$ (top panel), $N = 17$ (middle panel) and $N = 23$ (bottom panel) sector portfolios.

	β_{ij}						
	0.1	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**
$\bar{\rho}_0 = 0.1$	✓		✓	✓		✓	
$\bar{\rho}_0 = 0.2$	✓	✓	✓	✓			
$\bar{\rho}_0 = 0.3$	✓		✓	✓			
$\bar{\rho}_0 = 0.4$				✓			
$\bar{\rho}_0 = 0.5$							

	β_{ij}						
	0.1	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**
$\bar{\rho}_0 = 0.1$	✓	✓	✓	✓	✓	✓	✓
$\bar{\rho}_0 = 0.2$	✓	✓	✓		✓	✓	✓
$\bar{\rho}_0 = 0.3$	✓	✓	✓	✓		✓	
$\bar{\rho}_0 = 0.4$					✓		
$\bar{\rho}_0 = 0.5$							

	β_{ij}						
	0.1	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**
$\bar{\rho}_0 = 0.1$	✓	✓	✓	✓	✓	✓	
$\bar{\rho}_0 = 0.2$	✓	✓	✓	✓	✓	✓	
$\bar{\rho}_0 = 0.3$	✓	✓	✓			✓	✓
$\bar{\rho}_0 = 0.4$							

A check mark indicates that the analytic solution falls within the 99% confidence intervals for the simulated VaR at $q = 0.1, 0.3, 0.5, 0.7$.

racy in terms of the correlation among firms is still an open question (the correlation matrix is easily found from (2.1)). Certainly, our result is exact in the limit as $N \to \infty$, but for finite values of N we do not know which correlation matrixes give good agreement between (3.21) and simulation results, and which do not. We have found instances where the agreement is good and there is a wide range of correlation among firms, and instances where the agreement is poor and there is a narrow range of correlation among firms.

In summary, our work builds on current research in developing analytic and numerical tools to study credit portfolios. We hope that it will motivate further studies to build analytic approximations that could be accurate even for a very small number of sectors (say $N = 3$ or 4) and extreme parameter values.

APPENDIX A. LAPLACE'S METHOD FOR EVALUATING INTEGRALS

In our calculation, we employ Laplace's method to approximate integrals. Here, we briefly review this method, which is a technique for asymptotically evaluating integrals of the form

$$I(k) = \int_a^b g(t) \exp[-k\psi(t)] \, \mathrm{d}t, \tag{A.1}$$

when $k \gg 1$. The method relies on the important fact that when k is large, most of the mass of the integrand will be located around a stationary point t^*, where $\exp[-k\psi(t)]$ is maximal or, equivalently, where $\psi(t)$ is minimal: $\psi'(t^*) = 0$, $\psi''(t^*) > 0$. When $t^* \in (a, b)$, we make the approximations $g(t) \approx g(t^*) + g'(t^*)(t - t^*) + \cdots$ and $\psi(t) \approx \psi(t^*) + \frac{1}{2}\psi''(t^*)(t - t^*)^2 + \cdots$ to find that

$$I(k) \sim \mathrm{e}^{-k\psi(t^*)} \sqrt{\frac{2\pi}{k\psi''(t^*)}} \left\{ g(t^*) + O\left(\frac{1}{k}\right) \right\}, \quad k \to \infty. \tag{A.2}$$

The first term on the right-hand side of (A.2) is Laplace's approximation to $I(k)$. The second term in the series can be used to give an error estimate of the first term. A full account of Laplace's method and smoothness conditions required for g and ψ can be found in many texts, such as Ablowitz and Fokas (1997), Olver (1997) and Erdelyi (1956). An explicit form for the $O(k^{-1})$ term in (A.2) can be found in Bender and Orszag (2010).

APPENDIX B. MATLAB CODES FOR VALUE-AT-RISK EVALUATION

The MATLAB code `GenerateMatrices.m` generates the model parameters w_{ij}, β_{ij}, θ_{ij} and ρ_{ij} for the hierarchical multifactor portfolios in this paper. The code `MCsimulatedVaR.m` is used to simulate the VaRs using Monte Carlo simulation after calling `GenerateMatrices.m` at the command prompt. Finally, `AnalyticVaR.m` is used to compute the analytic solution as given by (3.21). The codes can also be found on Pak-Wing Fok's website, http://udel.edu/~pakwing/MATLAB_codes/JCRcodes.txt.

```
function [w,beta,theta,rho,wvec,betavec,thetavec,rhovec] = GenerateMatrices(n)
% generate "reasonable" model parameters w_{ij}, \beta_{ij}, theta_{ij}, \rho_{ij}
% for a large bank like JP Morgan for N=17 sectors and
% specified n (# companies per sector). The w_{ij} are based on call report
% data

% N     Loan type/sector
% 1     construction and land development
% 2     farmland
% 3     1-4 residential properties
% 4     5+ residential properties
% 5     nonfarm, non-residential
% 6     commercial US banks
% 7     banks in foreign countries
```

```
% 8     agricultural loans
% 9     US commercial/industrial loans
% 10 non-US commercial/industrial loans
% 11 Credit Cards
% 12 Other revolving credit plans
% 13 automobile loans
% 14 other consumer loans
% 15 foreign governments
% 16 US states and subdivisions
% 17 non-depository financial institutions & other

wvec = [
   0.006038934108343
   0.000334001336005
   0.321727182643328
   0.071374977443038
   0.042981697993143
   0.004997356482317
   0.029969892391418
   0.001247360439679
   0.143856433245638
   0.053224933278880
   0.041455739282673
   0.004090329157525
   0.065719115957033
   0.031400874418664
   0.001632016006737
   0.020071739054729
   0.159877416760852];

rhovec = [
    0.3300
    0.3200
    0.3500
    0.3100
    0.3700
    0.9500
    0.9300
    0.5200
    0.4800
    0.5000
    0.1500
    0.1700
    0.2100
    0.1600
    0.9200
    0.9700
    0.8500];

% rhovec = [
%     0.23
%     0.22
%     0.25
%     0.21
%     0.27
%     0.15
%     0.13
%     0.12
%     0.18
%     0.28
%     0.15
%     0.17
%     0.21
%     0.16
%     0.12
%     0.17
%     0.15
%     ];
```

```
thetavec = [
  -2.326347874040841
  -2.326347874040841
  -2.326347874040841
  -2.326347874040841
  -2.326347874040841
  -3.090232306167814
  -3.090232306167814
  -2.326347874040841
  -2.326347874040841
  -2.326347874040841
  -1.644853626951473
  -1.644853626951473
  -1.644853626951473
  -1.644853626951473
  -3.719016485455709
  -3.719016485455709
  -2.326347874040841];

betavec = [
    0.8000
    0.9000
    0.0500
    0.1000
    0.1000
    0.9000
    0.1000
    0.7000
    0.0500
    0.1000
    0.1000
    0.8000
    0.1000
    0.1000
    0.8000
    0.1000
    0.0500];

N=length(wvec);
w = zeros(N,n);
rho = zeros(N,n);
theta = zeros(N,n);
for i=1:N
    w(i,:) = wvec(i)/n;
    rho(i,:) = rhovec(i) + 0.01*randn(1,n);
    theta(i,:) = thetavec(i) + 0.01*randn(1,n);
    beta(i,:) = betavec(i) + 0.01*rand(1,n);
end

function [VaR,VaR_upper,VaR_lower,R_pi] = MCsimulatedVaR(rho,beta,theta,w,...
c,q,num_trials,confidence,qs)
%
% This function produces a Monte Carlo simulated VaR_q plot for the
% hierarchical multi-factor credit
% portfolio problem and generates error bars at values of q specified
% in the qs vector
%
% rho:     matrix of asset return variances
% beta:    matrix of factor loadings
% theta:   matrix of threshold default values s.t. prob default =
%          Phi(theta_{ij})
% w:       matrix of exposures
% c:       loss given default (LGD). It's a scalar e.g. c=1.
% q:       vector of risk-values at which to evaluate VaR. Must be between
%          0 and 1 e.g. q = linspace(0,1,50)
%
% num_trials:            number of portfolios to generate e.g. num_trials = 5000
```

```
% confidence:            confidence values to create error bars. It's a
%                        number between 0 (no confidence) and 1 (complete confidence).
% VaR:                   Values-at-Risk corresponding to q
% [VaR_upper,VaR_lower]: The upper and lower error bars for VaR,
%                        calculated using the bootstrap method
% qs:                    Values of q at which to calculate error bars e.g.
%                        qs = [0.2 0.4 0.6 0.8]
% R_pi:                  Portfolio loss for each of the num_trials trials

NNN = 50000; % number of bootstrap samples
beta_hat = sqrt(1-beta.^2);

% parameter values
[N,n] = size(rho); % N: number of sectors, n: number of companies per sector
R_pi = zeros(1,num_trials);
for i=1:num_trials
    R_pi(i) = one_draw(rho,beta,beta_hat,theta,w,c,N,n);
end

P = zeros(NNN,length(qs));
for i=1:NNN
    if mod(i,2000) == 0
        sprintf('Generating bootstrap samples: %d/%d',i,NNN)
    end
    % generate num_trials random integers from [1:num_trials]
    k = randi(num_trials,1,num_trials);
    bootstrap_sample = R_pi(k);
    P(i,:) = quantile(bootstrap_sample,qs);
end
alpha = (1-confidence)/2; % e.g. alpha = 0.05 for confidence_percentage = 0.9
VaR_lower = quantile(P,alpha);
VaR_upper = quantile(P,1-alpha);
VaR = quantile(R_pi,q); VaR(end) = 1;
semilogy(q,VaR,'r--','LineWidth',2);
hold on;

for i=1:length(qs)
  val(i) = (VaR_lower(i) + VaR_upper(i))/2;
  errorbar(qs(i),val(i),val(i)-VaR_lower(i),VaR_upper(i)-val(i),'r-','LineWidth',2);
end

axis([0 1 1e-3 1e-1]);
h=gca;
set(h,'FontSize',14,'FontName','Times');
end

function R_pi = one_draw(rho,beta,beta_hat,theta,w,c,N,n)
    % different rows of rho, beta etc. correspond to different sectors,
    % following the convention in the paper
    epsilon_hat = randn(1,1);
    epsilon = kron(randn(N,1),ones(1,n));
    Z1 = sqrt(rho) .* beta_hat*epsilon_hat;
    Z2 = sqrt(rho) .* beta .* epsilon;
    Z3 = sqrt(1-rho) .* randn(N,n);
    Z = Z1+Z2+Z3;
    LOSSES = c*(Z <= theta);
    R_pi = sum(sum( w.*LOSSES ));
end

function VaR = AnalyticVaR(rho,beta,theta,w,c,q)
% provides analytic solution for N sectors and n companies per sector
% rho = A(:,1); beta = A(:,2); theta = A(:,3); w = A(:,4);
beta_hat = sqrt(1-beta.^2);
[N,n] = size(rho);
for i=1:length(q)
    arg = ( theta + sqrt(rho).*beta_hat.*invPhi(q(i)) )./ (sqrt(1-rho+beta.^2.*rho));
    VaR(i) = c*sum(sum(w.* Phi(arg)));
end
```

```
semilogy(q,VaR,'b-','LineWidth',2); hold on;
xlabel('q','FontSize',14,'FontName','Times','FontAngle','italic');
ylabel('VaR_q','FontSize',14,'FontName','Times');
end


function out = Phi(x)
out = 1/2*(1+erf(x/sqrt(2)));
end

function out = invPhi(x)
out = sqrt(2)*erfinv(2*x-1);
end
```

REFERENCES

Ablowitz, M. J., and Fokas, A. (1997). *Complex Variables*. Cambridge University Press.

Andersen, L., Sidenius, J., and Basu, S. (2003). All your hedges in one basket. *Risk* **16**, 67–72.

Basel Committee on Banking Supervision (2004). Basel II: international convergence of capital measurement and capital standards: a revised framework. Report, June, Bank for International Settlements.

Bender, C. M., and Orszag, S. A. (2010). *Advanced Mathematical Methods for Scientists and Engineers: Asymptotic Methods and Perturbation Theory*. Springer.

Berry, A. C. (1941). The accuracy of the Gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society* **49**, 122–136.

Bluhm, C., Overbeck, L., and Wagner, C. (2010). *Introduction to Credit Risk Modeling*, 2nd edn. Chapman & Hall/CRC.

Bohn, J. R., and Stein, R. M. (2009). *Active Credit Portfolio Management in Practice*, Vol. 384. Wiley.

Burtschell, X., Gregory, J., and Laurent, J. (2009). A comparative analysis of CDO pricing models. *Journal of Derivatives* **16**(4), 9–37.

Crosbie, P., and Bohn, J. (2002). Modeling default risk. Report, January, KMV LLC, San Francisco, CA.

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

Denault, M. (2001). Coherent allocation of risk capital. *The Journal of Risk* **4**(1), 7–21,

De Servigny, A., and Renault, O. (2004). *Measuring and Managing Credit Risk*. McGraw-Hill, New York.

Erdelyi, A. (1956). *Asymptotic Expansions*. Dover, New York.

Esseen, C. G. (1942). On the Liapunoff limit of error in the theory of probability. *Arkiv för Matematik, Astronomi Och Fysik* A **28**, 1–19.

Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation* **12**(3), 199–232.

Huang, X., Oosterlee, C., and Van der Weide, J. (2007). Higher order saddlepoint approximations in the Vasicek portfolio credit loss model. *The Journal of Computational Finance* **11**(1), 93–113.

Huh, J., and Kolkiewicz, A. (2008). Computation of multivariate barrier crossing probability and its applications in credit risk models. *North American Actuarial Journal* **12**(3), 263–291.

Hurd, T. R. (2009). Credit risk modeling using time-changed Brownian motion. *International Journal of Theoretical and Applied Finance* **12**, 1213–1230.

Jensen, J. (1995). *Saddlepoint Approximations*. Oxford University Press.

Joe, H. (1997). *Multivariate Models and Dependence Concepts*. Chapman & Hall/CRC.

JP Morgan (2014). Consolidated reports of condition and income for a bank with domestic and foreign offices – FFIEC 031. JPMorgan Chase Bank, National Association, Columbus, OH. URL: www.fdic.gov/bank/statistical.

Kupiec, P. H. (2008). A generalized single common factor model of portfolio credit risk. *Journal of Derivatives* **15**(3), 25–40.

Li, D. X. (2000). On default correlation: a copula function approach. *Journal of Fixed Income* **9**, 43–54.

Lucas, D. J. (1995). Default correlation and credit analysis. *Journal of Fixed Income* **4**, 76–87.

Lugannani, R., and Rice, S. (1980). Saddlepoint approximations for the distribution of the sum of independent random variables. *Advances in Applied Probability* **12**, 475–490.

McNeil, A. J., Frey, R., and Embrechts, P. (2010). *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton University Press.

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

Olver, F. W. J. (1997). *Asymptotics and Special Functions*. A. K. Peters/CRC Press.

Pykhtin, M. (2004). Multi-factor adjustment. *Risk* **17**(3), 85–90.

Schönbucher, P. J. (2001). Factor models: portfolio credit risks when defaults are correlated. *Journal of Risk Finance* **3**, 45–56.

Vasicek, O. A. (1987). Probability of loss on loan portfolio. Internal Report, KMV.