Journal of Credit Risk **22**(2), 1–33
DOI: 10.21314/JCR.2025.023

Research Paper

Credit risk meets insurance risk: a unified framework

Guusje Delsing,[1] Michel Mandjes,[2,3] Peter Spreij[3,4,5] and Erik Winands[1,3]

[1] Rabobank, Croeselaan 18, 3521 CB Utrecht, Netherlands;
emails: guusje.delsing@rabobank.nl, erik.winands@rabobank.nl

[2] Mathematical Institute, Leiden University, Niels Bohrweg 1, 2333 CA Leiden, Netherlands; email: m.r.h.mandjes@math.leidenuniv.nl

[3] Korteweg–de Vries Institute for Mathematics, University of Amsterdam, Postbus 94248, 1090 GE Amsterdam, Netherlands; email: p.j.c.spreij@uva.nl

[4] Institute for Mathematics, Astrophysics and Particle Physics, Radboud University, Heyendaalseweg 135, 6525 AJ Nijmegen, Netherlands

[5] ISEG Lisbon School of Economics and Management, University of Lisbon, Rua do Quelhas 6, 1200-781 Lisbon, Portugal

(Received June 10, 2025; revised October 7, 2025; accepted October 20, 2025)

ABSTRACT

This paper introduces a continuous-time extension to the influential CreditRisk+ model for portfolio credit risk modeling. For capital calculations it introduces a risk measure based on the maximum of the loss process of a portfolio over a specified time interval. An extensive numerical study demonstrates that this extension provides an accurate risk assessment. The new framework has many advantages. First, it enables loss evaluation over a continuous time period rather than at a fixed point in time as in the original CreditRisk+ model. Second, the framework offers great flexibility, for example, for the incorporation of collateral risks, interest income and regime shifts. It also accommodates the calibration of almost any observed dependence between counterparties, including dependencies of the form used in structural

Corresponding author: G. Delsing

Print ISSN 1744-6619 | Online ISSN 1755-9723

credit risk models. Finally, the framework establishes a direct connection to insurance industry ruin theory models, allowing the model to leverage exact results and established algorithms from this field to efficiently and accurately determine the credit loss distribution within the continuous-time framework. This extended model therefore provides a comprehensive and theoretically robust approach to assessing credit risk over time and shows the potential for broader use in risk management.

Keywords: ruin theory; CreditRisk+; risk management; portfolio credit risk loss; continuous-time modeling; dependence modeling.

1 INTRODUCTION

In recent decades, financial institutions have developed models to assess credit risk in lending portfolios. Benchmark frameworks such as CreditMetrics (JP Morgan) and CreditRisk+ (Credit Suisse Financial Products) have been pivotal in shaping industry standards and influencing the Basel capital accords. The original CreditRisk+ model draws on insurance-based event-risk models, in which default is the main source of credit risk.[1] Dependence between obligors/instruments is introduced through background macroeconomic risk factors impacting the default probabilities of the obligors/instruments. Our paper enhances CreditRisk+ by introducing a versatile continuous-time framework that incorporates various dependence structures and connects portfolio credit risk models with insurance industry ruin theory models.

A continuous-time framework is particularly relevant in light of regulatory and internal capital management practices. As outlined by Basel Committee on Banking Supervision (2017), banks' capital requirements are based on two complementary perspectives: regulatory capital (RC), reflecting regulatory and supervisory requirements; and economic capital (EC), capturing the bank's internal risk view. Within the RC framework for credit risk, capital calculations are based on the asymptotic single risk factor (ASRF) model (see Basel Committee on Banking Supervision 2005). Whereas the value-at-risk (VaR) calculations in that model assume obligors can default only at the end of the horizon, the possibility of defaults occurring throughout the year is still incorporated in RC, ensuring conservative capital buffers. First, the estimation of the input parameters for the probability of default explicitly accounts for the default behavior of the obligors over the full period. Second, the actual capital held by banks is not based solely on the VaR output of the ASRF model but includes various regulatory and management buffers to, for example, account for model simplifications, such as the ASRF assumption of defaulting only at the end

[1] For a comprehensive overview of CreditRisk+ and its extensions, see Gundlach and Lehrbass (2004).

of the time horizon. As mentioned, the EC perspective complements the RC framework by allowing an institution to apply its own internally developed models and methodologies tailored to its specific risks, portfolios and risk strategy. The results of our paper can play a valuable role in precisely this context by reinforcing the rationale behind the prudent adjustments and buffers within the RC framework. The explicit incorporation of the continuous-time effect allows for the substantiation of these buffers under both baseline and stressed conditions. Further, the economic view provided by our continuous-time model can be used for capital allocation across business lines and portfolios, particularly in areas where the impact of our extension is most pronounced.

The original CreditRisk+ model, by evaluating only end-of-horizon losses, overlooks fluctuations during the period. A discrepancy arises between the maximum loss over the entire horizon and the loss observed at the end if losses are nonmonotonic and reach their peak earlier. Such behavior naturally arises once compensating effects (eg, interest income, prepayments or expected losses) are taken into account.

In response, a growing literature extends portfolio credit risk models beyond the one-period setting. Multiperiod models allow the joint evolution of defaults, migrations and exposures to be captured dynamically. Notable contributions include those of García-Céspedes and Moreno (2017), who extend the Vašíček model to discrete multiperiod losses, and Chongo and Julia-Sala (2024), who simulate correlated defaults and rating migrations over multiple periods. However, these are both still discrete-time approaches. A line of work more closely related to ours is pursued by Reiß (2004), who extends CreditRisk+ by modeling risk factors as dependent geometric Brownian motions, enabling a continuous-time description of credit risk. This framework provides explicit expressions for the first two moments of the loss process. While Reiß (2004) enriches the dynamics of the risk factors, our contribution is to extend the loss process itself to continuous time, with explicit treatment of compensating inflows and outflows. Our framework also accommodates more general dependence structures, though we adopt geometric Brownian motions in our numerical work to highlight the link to Reiß (2004). In addition to these contributions, our paper also explicitly connects portfolio credit risk to ruin theory.

Dependence between obligors in portfolio credit risk models is typically modeled via common risk factors, often interpreted as drivers of the economic environment such as macrofinancial variables. A stream of research has focused on embedding different dependence structures into portfolio credit risk models. For instance, Han and Kang (2014) extend CreditRisk+ by incorporating a generalized common risk factor framework, and Fischer and Dietz (2011) develop the common background vector model, in which sectoral dependencies are linked to multiple background variables reflecting the state of the economy and macroeconomic conditions. In our continuous-time CreditRisk+ model we propose a general framework for dependence between the underlying risk processes, and as a result the dependence between counterparties can be modeled in terms of almost any observed structure.

A continuous-time CreditRisk+ model will share various characteristics with ruin theory models, which study the cash reserves of an institution subject to claims and premium inflows, with a focus on the probability of ruin (reserves falling below zero) (see Asmussen and Albrecher 2010). Ruin-theoretic methods can be adapted to portfolio credit risk to analyze maximum losses and default events in continuous time. Prior work has explored this link: Chen and Panjer (2009) apply ruin theory to individual obligor default by drawing parallels to first-passage structural models; Adékambi and Essiomle (2020) and Yang (2003) instead model institutional cash reserves. Adékambi and Essiomle model loan arrivals as cash outflows, while amortization payments and interest income form the inflows; our model, in line with regulatory capital requirements, abstracts from new loan deals and instead focuses on losses in the existing portfolio, with greater flexibility in default dependence and loss amounts. Yang adopts a discrete-time setup in which reserve fluctuations depend solely on credit ratings, modeled by a Markov chain; unlike our approach, Yang's framework cannot be calibrated to a portfolio with individual default risks and explicit dependence. Collectively, the abovementioned studies illustrate the promise of connecting ruin theory and credit risk, but they do not capture portfolio-level dependencies and continuous-time loss dynamics. Our continuous-time extension of CreditRisk+ addresses these gaps in the literature by focusing on portfolio-level losses while retaining ruin-theoretic tractability.

The main contributions of our paper are the following.

- The original CreditRisk+ model is extended to a continuous-time version that enables loss evaluation over a continuous period of time, rather than at a fixed point in time as in the original CreditRisk+ model. An extensive numerical study shows that the difference between the two models increases when systemic volatility rises or compensating effects increase. Approximating the continuous-time model is achieved by evaluating losses at a finite number of uniformly spaced points over the time horizon.

- The extended model creates an extremely flexible credit risk framework by allowing for different loss distributions, compensating functions and risk factor processes. This enables the inclusion of, for example, collateral risks, interest income and regime shifts. In addition, the dependence between counterparties can be modeled in terms of virtually any observed structure, including the correlation structure in the widely used structural models that form the basis of the Basel capital formula. In our numerical study we also examine the sensitivity of the model to various parameters and correlation structures.

- The proposed framework bridges the gap between portfolio credit risk models and insurance industry ruin theory models. Similarities are drawn between credit and insurance risks, leading to a productive exchange of results and methods. Specifically, the numerical work demonstrates that ruin theory methods can be effectively applied to accurately evaluate the credit loss distribution within the continuous-time CreditRisk+ framework, providing an alternative to simulation-based approaches. While the connection between credit and insurance risk models has been noted previously (as discussed by, for example, Delsing 2022), our work contributes to this line of research and may help further encourage cross-fertilization between these fields.

The remainder of the paper is structured as follows. Section 2 introduces the original CreditRisk+ model and its continuous-time extension. Further, it presents a comparison with a widely used class of structural models. Section 3 establishes a connection between portfolio credit risk models and ruin theory models. Section 4 provides a detailed numerical study of the continuous-time framework. Section 5 states our conclusions and discusses areas for future research.

2 THE MODEL

In this section we introduce our model, an extension of the Credit Suisse Financial Products CreditRisk+ model, and explain its relevance for default and credit risk modeling purposes. The objects of interest are the distribution of the (default) loss in a credit portfolio and capital calculations based on this distribution. We start with a brief introduction to the original CreditRisk+ model of default risk and use similar notation to that of Gordy (2000).

2.1 The CreditRisk+ model

In the original CreditRisk+ model each obligor in the credit portfolio is considered to have only two possible states: default and nondefault. If obligor i defaults, a loss is suffered of fixed size z_i. The default event of obligor i prior to the fixed maturity time T is denoted by $D_{i,T}$, and its indicator function by $\mathbf{1}_{D_{i,T}}$. The default event can be written in terms of the default time τ_i: $D_{i,T} = \{\tau_i \leqslant T\}$. The probability of default is given by

$$p_{i,T} := \mathbb{P}(D_{i,T}).$$

Consider a portfolio of n obligors with the same potential fixed loss $z_i \equiv z$, constructed by grouping together losses of similar sizes. The portfolio loss up to a fixed time T is then given by

$$L_T := \sum_{i=1}^{n} z_i \mathbf{1}_{D_{i,T}} = z \sum_{i=1}^{n} \mathbf{1}_{D_{i,T}}.$$

Default correlations between obligors are introduced by the risk factor $X := (X_1, \dots, X_K)$. The X_k are often assumed to be nonnegative and (independently) gamma distributed. Conditional on X, the default events of obligors are independent and Bernoulli distributed. The default probability of obligor i now becomes a random variable, which in CreditRisk+ is specified as

$$p_{i,T}^X := p_{i,T} \sum_{k=1}^{K} w_{i,k} X_k,$$

where the $w_{i,k}$ represent factor loadings measuring the sensitivity of obligor i to each of the risk factors, satisfying $\sum_{k=1}^{K} w_{i,k} = 1$. The risk factor and corresponding weights are such that $\mathbb{E}[p_{i,T}^X] = p_{i,T}$ for all obligors. The CreditRisk+ model assumes that conditional on X and under small default probabilities, the total number of obligor defaults can be approximated by a Poisson random variable N_T with parameter $\mu_T^X := \sum_{i=1}^{n} p_{i,T}^X$ (see Gordy 2000). This works when $p_{i,T}$ is small as we can then ignore the constraint that a single obligor can default only once. The portfolio loss is then approximated by

$$\hat{L}_T = \sum_{j=1}^{N_T} z = zN_T.$$

Traditionally, a bank's credit risk is managed by considering its losses over a certain period $[0, T]$ and making sure these do not exceed a certain (high) level with a given (low) probability α. This risk measure is often referred to as the value-at-risk (VaR). More specifically, for losses L_T,

$$\mathrm{VaR}^{\alpha}(T) := \inf\{u \geqslant 0 \mid \mathbb{P}(L_T \geqslant u) \leqslant \alpha\}. \tag{2.1}$$

Typical values for the maturity and threshold are one year ($T = 1$) and $\alpha = 0.01\%$. In practice, the losses are only considered at the end of the time horizon, at time T. Although it is impossible to know a bank's credit losses over a particular time interval in advance, an expectation of the losses can be determined. These losses are referred to as "expected losses". The expected losses over time are often already embedded in the pricing of credit instruments and/or provisioning by the financial firm, as they are viewed as a cost of doing business. A bank's capital is meant to provide a buffer against losses that significantly exceed expected levels. These losses are referred to as "unexpected losses" and are the gap between the expected losses and VaR.

The CreditRisk+ model has been celebrated for its simplicity as it allows for an expression for the distribution of portfolio losses. We refer the reader to Gundlach and Lehrbass (2004) for a more comprehensive overview of the original CreditRisk+ model and various extensions and applications. These extensions of the framework include the integration of migration risk and variable default severity. For the integration of migration risk we refer the reader to Binnenhei (2004).

2.2 A continuous-time CreditRisk+ model

In the original CreditRisk+ model, as well as most generalizations, only the loss at the end of a certain time horizon is studied. We will often refer to this as the "static" CreditRisk+ model. As mentioned by Reiß (2004), it is also relevant to consider credit risk over time. This is especially true when generalizing the model to consider not only the outgoing cashflows due to defaults but also some positive cashflows such as interest rate premiums received or (expected) losses that are already accounted for. The maximum of the loss process over the time horizon, in this case, may occur prior to the end of the time horizon. In this section we propose an extension of CreditRisk+ by adjusting the model on a few fronts, which results in a dynamic, "continuous-time" portfolio loss process. We will elaborate on all the changes in later subsections, but we start with a brief overview of the changes made to the original CreditRisk+ model.

- We allow for random losses Z_i instead of fixed losses z_i. The random losses are independent and identically distributed (iid) and nonnegative.
- We include a compensating nonnegative increasing real function $t \mapsto h(t)$ defined on $\mathbb{R}$.
- We allow for risk factors to be random processes instead of random variables (ie, $X(t) := (X_1(t), \dots, X_K(t))$). Let

  $$\mathcal{X} = (\mathcal{X}(t))_{t\in[0,T]}, \qquad \mathcal{X}(t) := \sigma(X(s) \colon 0 \leqslant s \leqslant t),$$

  denote the filtration generated by the multivariate risk factor process.
- For obligor i, instead of the Poisson random variable $N_{i,t}$, we consider the inhomogeneous conditional Poisson process $N_i(t)$ with nonnegative intensity process $\lambda_i^X(t)$, which is assumed to be $\mathcal{X}(t)$-measurable and whose integral $\int_0^t \lambda_i^X(s)\,\mathrm{d}s$ is well defined and finite for all $t \geqslant 0$ almost surely. The Poisson process is independent of the random losses Z_i. Analogously to the original CreditRisk+ model, for the number of defaults in the credit portfolio up to time t, conditional on $\mathcal{X}(t)$, we take the Poisson process $N(t)$ with intensity parameter $\mu^X(t) := \sum_{i=1}^n \lambda_i^X(t)$.

As a result, the loss over a credit portfolio up to time t is modeled as the random process

$$\hat{L}(t) := \sum_{i=1}^{N(t)} Z_i - h(t). \tag{2.2}$$

In the continuous-time framework, the default time τ_i of obligor i is explicitly modeled as the first jump of a nonhomogeneous Poisson process with (possibly stochastic) intensity $\lambda_i^X(t)$. We denote by $p_i^X(t) := \mathbb{P}(\tau_i \leqslant t \mid \mathcal{X}(t))$ the probability of default by time t, conditional on the risk factor process. Under the continuous-time CreditRisk+ model this is given by $p_i^X(t) = 1 - \exp(-\int_0^t \lambda_i^X(s)\,\mathrm{d}s)$. As in the original CreditRisk+ model, and as is common in credit risk over short time horizons, for a small integrated default intensity $\int_0^t \lambda_i^X(s)\,\mathrm{d}s$ this default probability can be approximated by $p_i^X(t) \approx \int_0^t \lambda_i^X(s)\,\mathrm{d}s$.

The quantity of interest is the maximum of the loss process over a specified time horizon. More specifically, for risk and capital calculations, we are interested in the probability of large losses over the time interval. This is given by

$$\Psi(u, T) = \mathbb{P}\Big(\sup_{0 \leqslant t \leqslant T} \hat{L}(t) > u\Big).$$

In many cases it is often easiest to first determine the conditional (on $\mathcal{X}(T)$) probability and then integrate or simulate out the risk factor processes.

For capital calculations we consider the following risk measure:

$$\rho^{\alpha}(T) := \inf\{u \geqslant 0 \mid \Psi(u, T) \leqslant \alpha\}. \tag{2.3}$$

If the expected losses are accounted for in $h(t)$, the loss process $\hat{L}(t)$ represents the unexpected losses and the corresponding risk measure $\rho^{\alpha}(T)$ can be seen as the capital estimation for credit risk. A similar risk measure is considered by Boudoukh *et al* (2004), who, instead of the value at the end of the time horizon, as is considered in the original VaR measure (2.1), use the maximum or minimum value before the end of the time horizon.

The final model, as given in (2.2), is common in the field of ruin theory. We will elaborate on this in Section 3 and show that results derived in ruin theory can be used to determine the distribution of losses. We now proceed by further elaborating on the above extensions of the model and their relation to the existing literature in separate subsections below.

2.2.1 Random losses

While the current model allows for both constant and random losses, the latter are often applicable in the presence of collateral risk. For the CreditRisk+ model, random losses have been considered by Akkaya *et al* (2004).

2.2.2 Compensating function $h(t)$

By including a nonnegative function $h(t)$, the model allows for compensating effects such as interest income or adjusting for expected losses. While the original CreditRisk+ model and the regulatory ASRF model do not take this compensating aspect into account, many other market practice models do consider it to some extent.

Examples include CreditMetrics and Moody's KMV, both of which have a valuation module. Other examples in the literature include models by Schlottmann *et al* (2004) and Adékambi and Essiomle (2020).

A linear function is an appropriate choice for the compensating function. Premiums or interest rates received are often modeled as a linear function of time. A linear function of time is also a logical choice in the case of the compensating function $h(t)$ representing the expected losses $\mathbb{E}[\sum_{i=1}^{N(t)} Z_i]$ over time when the Poisson process $N(t)$ is homogeneous. When, conditional on $\mathcal{X}(t)$, the loss arrival process $N(s)$ for $s \leqslant t$ is a homogeneous Poisson process with fixed conditional intensity parameter $\lambda^X \in \mathbb{R}_+$ and iid losses Z_i, the corresponding conditional expected losses at time t are given by $\lambda^X t\mathbb{E}[Z_i]$, which is a linear function of time. By setting $h(t)$ as the expected losses, the loss process $\hat{L}(t)$ represents the unexpected loss process.

In practice, the compensating function $h(t)$ can be calibrated using empirical portfolio data or by estimating expected losses. For instance, scheduled loan repayments, interest or fee income and expected recoveries from defaulted exposures can provide a deterministic approximation of cumulative inflows. Constructing $h(t)$ from future cashflows typically involves identifying and aggregating all predictable portfolio inflows over the horizon, including contractual repayments, scheduled interest and other recurring income. This requires detailed portfolio knowledge. Ideally, these cashflows would also be incorporated directly into the loss process Z_i by making them time dependent; this remains an area for future research. Alternatively, $h(t)$ can be defined as the expected cumulative losses, which naturally aligns the model with the unexpected loss process used for capital calculations. This is the approach adopted in our numerical study in Section 4. The deterministic inflows and/or expected losses can be aggregated across the portfolio and, if necessary, smoothed to account for seasonal patterns or discretization, ensuring $h(t)$ captures the baseline around which stochastic deviations occur.

2.2.3 Risk factor processes

The common risk factors representing the systematic component upon which obligors are dependent are often associated with the state of the environment/economy, which changes over time. Reiß (2004) also considers a continuous-time version of the original CreditRisk+ model and introduces a geometric Brownian motion to model the risk factor processes. Our model deviates from the work of Reiß (2004) as we not only consider a different risk factor process but also factor in a compensating function, something that turns out to be very relevant when considering the difference between the original and continuous-time CreditRisk+ models. Further, we present a more general framework for the risk factor processes, presenting multiple examples in addition to the geometric Brownian motion case.

Another common way to introduce dependence is via regime switching or Markov-modulation; that is, by introducing a Markov environmental process with a finite state space that influences the parameters of the model (ie, the intensity of the Poisson process). Further, note that our model setup still allows for piecewise constant processes such that the risk factors can also be modeled via random variables.

2.2.4 Inhomogeneous Poisson process

The original CreditRisk+ model only studies the loss at the end of a certain time horizon, but in many cases a continuous-time model is more appropriate. The inhomogeneous Poisson process is the natural extension of the original CreditRisk+ model to continuous time.

2.2.5 Dependence structures

Our continuous-time CreditRisk+ model provides a general framework for the dependence between the underlying risk processes. We will now elaborate on potential options to define the default intensity $\lambda_i^X(t)$. The most straightforward model for the default intensity is an extension of the original CreditRisk+ model in the form

$$\lambda_i^X(t) = p_i(1) \sum_{k=1}^{K} w_{i,k} X_k(t),$$

where $p_i(1)$ denotes the one-year default probability of obligor i and the $w_{i,k} \geqslant 0$ are the factor loadings as in the original CreditRisk+ model (see, for example, Gordy 2000; Credit Suisse Financial Products 1997, Section A12.3). The risk factors $X_k(t)$ are assumed to be nonnegative and satisfy $\mathbb{E}[\int_0^1 X_k(t)] = 1$, ensuring consistency with the default probability $p_i(1)$. Recall that the default probability conditional on $\mathcal{X}$, $p_i^X(t)$, can be approximated by $\int_0^t \lambda_i^X(s)\,\mathrm{d}s$ in the continuous-time CreditRisk+ model, which now gives $\mathbb{E}[p_i^X(1)] \approx p_i(1)$.

Below, we briefly introduce some dependence structures (ie, options for the intensity of the Poisson process $\lambda^X(t)$) based on the literature, and we reflect on their advantages and disadvantages. Each dependence structure results in a different distribution of the loss process $\hat{L}(t)$ and the corresponding maximum loss distribution $\Psi(u, T)$. In Section 4 all of the following correlation structures are implemented for the continuous-time CreditRisk+ model.

(1) In the industry, credit risk modeling often employs credit ratings, commonly represented as a finite-state continuous-time Markov process (see Berd 2005). When considering a two-state model (ie, default and nondefault), with an absorbing default state, the survival probability of obligor i up to time t is

given by $e^{-c_i t}$ (and the corresponding default probability by $1 - e^{-c_i t}$) for some constant $c_i > 0$. This aligns with the survival probability over a fixed time t when default is the first jump of a homogeneous Poisson process with constant intensity rate c_i. Thus, a Poisson process with constant intensity rate c_i is a suitable choice for modeling default/nondefault environments with absorbing defaults in a continuous-time Markov model. Using the approximation $1 - e^{-c_i} \approx c_i$ for $c_i \approx 0$, the intensity constant c_i can be set to the one-year probability of default $p_i(1)$ when this probability is small. Dependence between obligors can be introduced as in the regular CreditRisk+ framework by assuming the common risk factor is constant for all t (ie, $X_k(t) = X_k$, where X_k is a random variable). More specifically,

$$\lambda_i^X(t) = c_i \sum_{k=1}^{K} w_{i,k} X_k, \quad \text{with } \mathbb{E}\bigg[\sum_{k=1}^{K} w_{i,k} X_k\bigg] = 1.$$

As a result, conditional on the common risk factors X_k, the process $N(t)$ is a homogeneous Poisson process. Popular choices for X_k include independent gamma variables or a discrete distribution with finite states. While conditioning on X_k yields a simpler homogeneous Poisson process, it does not capture changes in risk factors over time, which are often seen in practice.

(2) Reiß (2004, Section 13.2) suggests a continuous-time extension of the original CreditRisk+ by considering a Poisson process in which the default intensity of obligor i is set as the one-year probability of default $p_i(1)$, assumed to be small. Dependence between obligors is introduced through dependent geometric Brownian motion factor processes $X_k(t)$, $k \in \{1, \dots, K\}$, and the assumption that $X_k(0) = 1$ and $\mathbb{E}[X_k(t)] = 1$. The default intensity of obligor i for Reiß (2004) is then a stochastic process given by

$$\lambda_i^X(t) = p_i(1) \sum_{k=1}^{K} w_{i,k} X_k(t),$$

where the factor weights ensure $\sum_{k=1}^{K} w_{i,k} = 1$ for each obligor i. The advantage of this dependence structure is that the risk factor changes over time.

(3) Regime switching is a popular method for introducing uncertainty or dependence in models, including ruin theory models. We consider a single finite-state Markov process $X(t)$ that influences the intensity of the Poisson process. This gives default arrival intensity $\lambda_i^X(t) = \lambda_{i,j}$ for obligor i when $X(t) = j$, where $\lambda_{i,j}$ is fixed (see, for example, Asmussen and Albrecher 2010, Chapter VII). The advantage of this dependence structure is that it allows for the

risk process to change over time but reduces the complexity by limiting the number of potential outcomes. Since analytical expressions over a finite time interval are not available, numerical approximations must be resorted to.

(4) We pose for known and small $p_i^X(t)$ the approximation

$$\lambda_i^X(t) \approx \frac{\partial}{\partial t} p_i^X(t)$$

to obtain the already established approximation $p_i^X(t) \approx \int_0^t \lambda_i^X(s)\,\mathrm{d}s$. This is in line with the original CreditRisk+ model methodology. An example where $p_i^X(t)$ is known from a continuous-time structural model is given in Section 2.3. By imposing very few conditions on the intensity $\lambda_i^X(t)$, this dependence structure allows for a lot of flexibility, and dependence may change over time. On the other hand, calculations can become complicated and time-consuming.

The functions $h(t)$ and $\mu^X(t)$ are time dependent but not influenced by default levels. This assumption is reasonable when defaulted credit assets/obligors are replenished by similar ones with the same credit quality and/or when the number of defaults is small relative to the total number of obligors in the portfolio. Delsing and Mandjes (2021) investigate a model in which the income rate and the (default) arrival rate are dependent on the number of surviving obligors in the system.

2.3 Comparison with structural models

In the continuous-time CreditRisk+ model, default is modeled with a reduced-form/ intensity model (ie, default is described through an exogenous jump process). In this section we consider a continuous-time structural default model (ie, default is triggered by the value of the firm dropping below the debt level). We will compare it with the continuous-time CreditRisk+ model.

In the continuous-time structural model, the loss over a credit portfolio up to time t is a random process given by

$$L(t) := \sum_{i=1}^{n} Z_i \mathbf{1}_{D_{i,t}} - h(t),$$

where we have used the same notation as above; ie, Z_i and $h(t)$ denote the random losses and nonnegative compensating function, respectively. The default is given by $D_{i,t} := \{\tau_i \leqslant t\}$, where the first passage time is $\tau_i := \inf\{t \geqslant 0 \colon A_i(t) < B_i\}$. The asset value process $A_i(t)$ is a geometric Brownian motion such that the log return of the asset value is a Brownian motion with constant drift μ_i and volatility σ_i. The unconditional probability of default of obligor i up to time t in the structural model is given by

$$\begin{aligned} p_i(t) &= \mathbb{P}(\mathbf{1}_{D_{i,t}}) = \mathbb{P}(\mathbf{1}_{\tau_i \leqslant t}) = \mathbb{P}\Big(\inf_{0<s\leqslant t} A_i(s) < B_i\Big) \\ &= \mathbb{P}\bigg(\inf_{0<s\leqslant t} \mu_i s + \sigma_i W_i(s) < \log \frac{B_i}{A_i(0)}\bigg) \\ &= \Phi\bigg(\frac{\log(B_i/A_i(0)) - \mu_i t}{\sigma_i \sqrt{t}}\bigg) \\ &\quad + \bigg(\frac{A_i(0)}{B_i}\bigg)^{-2\mu_i/\sigma_i^2} \Phi\bigg(\frac{\log(B_i/A_i(0)) + \mu_i t}{\sigma_i \sqrt{t}}\bigg), \end{aligned} \tag{2.4}$$

where $W_i(t)$ denotes a standard Brownian motion. The final expression is derived in a straightforward manner from the distribution of the maximum of a Brownian motion with drift as given by Dębicki and Mandjes (2015, (4.6)). As in the CreditMetrics and KMV structural models, we introduce dependence through common risk factor processes $X(t) := (X_1(t), \dots, X_K(t))$ (with factor loadings $w_i := (w_{i,1}, \dots, w_{i,K})$) driving the asset values of the obligor. More concretely,

$$W_i(t) = \sum_{k=1}^{K} w_{i,k} X_k(t) + \eta_i \varepsilon_i(t),$$

where $\varepsilon_i(t)$ (with weight η_i) denotes an idiosyncratic process independent of the common risk factor processes. The process $X(t)$ is a K-dimensional Brownian motion with mean zero and covariance matrix $\mathbb{E}[X(t)X(t)^{\mathrm{T}}] = \Sigma t$ with unit diagonal. The idiosyncratic risk processes $\varepsilon_1(t), \dots, \varepsilon_K(t)$ are mutually independent standard Brownian motions. Without loss of generality, we impose the condition that $\mathbb{V}[W_i(t)] = t$, where $\mathbb{V}[X]$ denotes the variance of the random variable X. We adopt the same notation as in Section 2, using $\mathcal{X}$ to denote the filtration generated by the multivariate risk factor process.

Similar to the work in Gordy (2000) on the original CreditRisk+ model, we will map the continuous-time structural model onto the mathematical framework of the continuous-time CreditRisk+ model. As the compensating function $h(t)$ and the loss size distribution in both models can be set equal, we focus on the default events and number of defaults in the portfolio. We first derive the implied default probability function (up to time t) of obligor i conditional on the risk factor process (ie, $p_i^X(t)$) in the structural model:

$$p_i^X(t) = \mathbb{P}\bigg(\inf_{0<s\leqslant t} \mu_i s + \sigma_i \sum_{k=1}^{K} w_{i,k} X_k(s) + \sigma_i \eta_i \varepsilon_i(s) < \log \frac{B_i}{A_i(0)} \;\bigg|\; \mathcal{X}(t)\bigg).$$

The unconditional default probability up to time t is then given by $p_i(t) = \mathbb{E}[p_i^X(t)]$. As mentioned in Section 2.2.5, the default intensity of the Poisson process conditional on the risk factor process $\lambda_i^X(t)$ can now be determined as

$$\sum_{i=1}^{n} \frac{p_i^X(t)}{\partial t},$$

where we have assumed that $p_i^X(t)$ is given by the structural model as above and is differentiable with a positive derivative. In Section 4 we implement this as the "Structural" dependence structure. Finally, we still need to simulate the risk factor process $X(t)$.

3 RUIN THEORY

The continuous-time CreditRisk+ model (2.2), conditional on $\mathcal{X}(t)$, has previously been considered in ruin theory for modeling an insurer's cash reserves. Consequently, all applicable results from ruin theory can be used to determine the loss distribution. For example, many of the ruin theory results presented by Asmussen and Albrecher (2010) can be used in this context. In this section we discuss how and which results from ruin theory can be used to calculate the (conditional) distribution of losses and capital reserves for the continuous-time CreditRisk+ model.

In the classical Cramér–Lundberg model, the evolution of the cash reserves of an insurance firm experiences fluctuations due to the claim amounts (Z_i), the arrival of claims $(N(t))$ and the incoming premiums $(h(t))$. Ruin theory primarily focuses on determining the probability of ruin: the probability that the supremum of aggregated losses (due to claims minus the received income) over time exceeds the initial capital reserves. More concretely, the probability of ruin over infinite and finite time horizons is given by

$$\psi(u) := \mathbb{P}(\tau(u) < \infty) = \mathbb{P}\Big(\sup_{0 \leqslant t < \infty} S(t) > u\Big),$$

$$\psi(u, T) := \mathbb{P}(\tau(u) \leqslant T) = \mathbb{P}\Big(\sup_{0 \leqslant t \leqslant T} S(t) > u\Big),$$

respectively, where $S(t)$ denotes the aggregate loss (or risk) process, u the initial reserve level, and $\tau(u) := \inf\{t \geqslant 0 \colon S(t) > u\}$ the time of ruin. Note that the finite-time ruin probability coincides with the probabilities of interest, $\Psi(u, T)$, given in Section 2 when $S(t) = \hat{L}(t)$.

Without wishing to provide a full overview of the ruin theory literature, we proceed by describing some results that are applicable in determining the distribution of $\hat{L}(t)$ and $\psi(u, T)$. As previously mentioned, in most cases it is easiest to first condition on the common risk factor process $X(t)$ and then derive the conditional ruin probabilities.

When conditioned on the risk factor process $X(t)$, if the loss arrival $N(t)$ is a homogeneous Poisson process and the compensating function is linear in time (ie, $h(t) = rt$ for fixed $r \geqslant 0$), then the loss process $\hat{L}(t)$, as defined in (2.2), corresponds to the classical Cramér–Lundberg model from ruin theory. The classical Cramér–Lundberg model is a well-studied standard model, and analytic formulas for the finite- and infinite-time ruin probability have been found for a few special cases depending on the distribution of the loss sizes (or claim sizes in the context of insurance modeling). One such case is when losses/claims are assumed to be exponentially distributed. In most cases, however, numerical or analytic approximations must be resorted to in order to calculate the probability of ruin, especially in finite time. These approximations are validated and in general very accurate.

Another special instance occurs when the risk factor process is a Markov process that drives the parameters (intensity and, possibly, the loss size distribution) of the classical Cramér–Lundberg model (see Section 2.2.4). This regime-switching model is fairly common in ruin theory (see, for example, Asmussen and Albrecher 2010, Chapter VII; Dickson and Qazvini 2018). To determine ruin probabilities, especially in finite time, numerical or analytic approximations are typically resorted to.

For more general inhomogeneous premium and arrival mechanisms, we refer the reader to Lefèvre and Loisel (2009) or to Asmussen and Albrecher (2010, Section VII.6), which predominantly focuses on periodic risk processes. Those references consider the classical Cramér–Lundberg model, with nonlinear premium process $h(t)$ and inhomogeneous Poisson arrival process $N(t)$ as specified in (2.2). To our knowledge, no explicit analytic formulas exist for the finite-time ruin probability of this model. Analytic and numerical approximations include an averaged model, as studied by Asmussen and Albrecher (2010) (who apply the classical Cramér–Lundberg framework with averaged parameters), and recursion formulas, as studied by Lefèvre and Loisel (2009). We draw the reader's attention to the recursion formulas from Lefèvre and Loisel (2009), which accommodate a general time-dependent compensating function $h(t)$, making them highly valuable for practical implementation.

In this section and throughout the remainder of the paper we focus on the results of ruin theory, which can be used to analyze the credit risk loss distribution. However, we would like to emphasize that the established connection between portfolio credit risk models and ruin theory facilitates a much broader exchange of results and methods. Further details are given by Delsing (2022). One notable result in ruin theory, which is of significant interest for credit risk management, is the existence of several generic methods for optimally allocating capital reserves across various business lines and subportfolios (see Delsing *et al* 2022).

4 NUMERICAL WORK

This section studies the continuous-time CreditRisk+ model and its relation to the original CreditRisk+ model and continuous-time structural model using comparative simulations. The main aim of this section is fourfold:

(1) to show the effect of time in the CreditRisk+ model;

(2) to compare the continuous-time CreditRisk+ and structural models;

(3) to show the sensitivity of the results to various parameters and correlation structures; and

(4) to demonstrate the efficiency and higher accuracy of using ruin theory results instead of simulations.

We consider a homogeneous portfolio consisting of 500 obligors with similar risk characteristics and with equal contributions to the portfolio, as it allows us to best demonstrate the various features of the models. The same exercise can also be carried out for a portfolio that is heterogeneous (in terms of the probability of default and size distribution), but for ease of computation and to better demonstrate the various features of the models, we have chosen to consider homogeneous portfolios in our experiments.

The credit quality of these obligors is characterized by the Standard & Poor's rating grade BBB, with a corresponding unconditional annual default probability $p_i(1)$ of 0.18% (see Table 6.9 of the CreditMetrics technical document by Gupton *et al* (2007)). Historically, the normalized volatility of the default probability of a single BBB obligor, $\sqrt{\mathbb{V}[p_i^X(1)]}/p_i(1)$, has a value of 0.4 according to Gordy (2000, Section 3.2, Table 2). Here we have used the notation $\mathbb{V}[X]$ to denote the variance of the random variable X. Similarly to the numerical work of Gordy (2000), we assume that the loss given default is a fixed proportion (30%) of the book value. The losses in the CreditRisk+ model (original/static and continuous-time) are then given by $z_i = Z_i \equiv 0.3(1/500)$ for a portfolio with a total book value of 1. Except for Section 4.3, where we examine the sensitivity of the results to other values of r, we assume the compensating function is given by $h(t) = rt$ with $r = 0.3 \times 0.18\% = 0.00054$. Note that this corresponds to the expected loss of the portfolio over one year (ie, $\mathbb{E}[\hat{L}(1)]$). Further, we consider a time horizon of one year in all numerical experiments, in line with regulatory banking requirements.

In Section 2.2.4 we introduced several dependence structures, all of which will be considered in our numerical (simulation) experiments, leading to different loss distributions. To calibrate these dependence structures, we follow the approach of Gordy (2000), matching the implied model and historical values for the annual default probability and its volatility.

All dependence structures considered in this numerical study are based on a single stochastic risk factor process, $X(t)$. For notational convenience, we define the systemic volatility ξ in the continuous-time CreditRisk+ model as

$$\xi^2 := \mathbb{V}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg].$$

Following Gordy (2000), who examined systemic volatility values of 1.0, 1.5 and 4, we set $\xi = 1.5$ in our analysis. The sensitivity of the results to variations in ξ is explored in Section 4.3.

Except in the case of the structural model, we further assume that the intensity process for obligor i is given by $\lambda_i^X(t) = p_i(1)(1 - w + wX(t))$. This specification corresponds to a two-factor version of the model mentioned in Section 2.2.5, where the first risk factor is 1 (ie, has zero volatility). To ensure that the expected default probability aligns with the one-year default probability, we impose the condition $\mathbb{E}[p_i^X(1)] = p_i(1)$, as in Section 2.2. This condition implies $\mathbb{E}[\int_0^1 X(t)\,\mathrm{d}t] = 1$.

The weight w is determined by matching the normalized variance (or volatility) of the default probability from the model, $\mathbb{V}[p_i^X(1)]/p_i(1)^2$, to historical observations. In our examples, this normalized variance is given by

$$\frac{\mathbb{V}[\int_0^1 \lambda_i^X(t)\,\mathrm{d}t]}{p_i(1)^2} = w^2\mathbb{V}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg] = w^2\xi^2.$$

Given $\xi^2 = 1.5^2$ and an observed normalized variance of 0.4^2, the weight w is uniquely determined as 0.267.

We now present the different dependence structures and briefly describe the calibration of their risk factor processes.

Discrete. In line with (1) in Section 2.2.5, we make use of a discrete and constant common risk variable (ie, $X(t) = X$ for all t). The common risk variable X is a discrete random variable that has two possible states, x_1 and x_2, with corresponding probabilities q and $1 - q$. These parameters are selected arbitrarily, subject to the following conditions: $0 < q < 1$, $\mathbb{E}[X] = 1$ and $\mathbb{V}[X] = \xi^2 = 1.5^2$. We take $q = 121/130$, $x_1 = 13/22$ and $x_2 = 13/2$.

Gamma. As another example of (1) in Section 2.2.5, we again consider $X(t) = X$ constant over time. Aligned with the original CreditRisk+ model, the risk factor X is now assumed to be gamma distributed with shape parameter $k > 0$ and scale parameter $\theta > 0$. The model is calibrated in the same way as the discrete example above by setting $\mathbb{E}[X] = 1$ and $\xi^2 := \mathbb{V}[X] = 1.5^2$. This gives $\theta = 2.25$ and $k = 1/\theta$.

Geometric Brownian motion. We now examine the dependence structure given in (2) in Section 2.2.5 by considering the common factor process $X(t)$ given by a geometric Brownian motion:

$$X(t) = X(0)\exp((\mu_X - \tfrac{1}{2}\sigma_X^2)t + \sigma_X W(t)),$$

where $W(t)$ is a standard Brownian motion and μ_X, σ_X denote the drift and volatility parameters, respectively. Further, $X(0) = 1$ and $\mathbb{E}[X(t)] = 1$ for all $t > 0$, as described in Section 2.2.4. These restrictions give $\mu_X = 0$. Note that this also gives $\mathbb{E}[\int_0^1 X(t)\,\mathrm{d}t] = 1$. We calibrate the remaining model parameter in a similar way to the discrete and gamma cases by setting $\xi^2 := \mathbb{V}[\int_0^1 X(t)\,\mathrm{d}t] = 1.5^2$. This gives $\sigma_X = 1.6777$ due to the fact that

$$\mathbb{V}\bigg[\int_0^t X(s)\,\mathrm{d}s\bigg] = \frac{2(-\sigma_X^2 t + \mathrm{e}^{\sigma_X^2 t} - 1)}{\sigma_X^4} - t^2.$$

Regime switching. We now assume that the common factor process $X(t)$ is driven by a Markov environmental process $J := \{J(t) \colon t \geqslant 0\}$ with support $\{1, 2\}$. This leads to the scenario described in (3) of Section 2.2.5. More concretely, if at time t the environmental process J is in state j, the common factor takes on value $x_j > 0$. The transition rate matrix governing J is denoted by $Q = (q_{k,l})_{k,l\in\{1,2\}}$. Inspired by Asmussen (1989), we assume that transitions of J from 1 to 2 occur at rate $q_{1,2} = \rho$, and those from 2 to 1 at rate $q_{2,1} = 2\rho$. We further assume that the initial environment is in state 1. The parameters of the model (ie, $x_1, x_2 > 0$ and ρ) are arbitrarily chosen such that

$$\mathbb{E}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg] = 1, \qquad \xi^2 := \mathbb{V}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg] = 1.5^2.$$

Let T_1 denote the occupation time of environmental state 1 (that is, $T_1 := \int_0^1 \mathbf{1}_{J(t)=1}\,\mathrm{d}t$). We may now write the expectation and variance of the integral over the factor process as

$$\mathbb{E}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg] = (x_1 - x_2)\mathbb{E}[T_1] + x_2,$$
$$\mathbb{V}\bigg[\int_0^1 X(t)\,\mathrm{d}t\bigg] = (x_1 - x_2)^2\mathbb{V}[T_1].$$

This gives

$$x_1 = x_2 + (1 - x_2)/\mathbb{E}[T_1],$$
$$x_2 = \xi\mathbb{E}[T_1]\sqrt{1/\mathbb{V}[T_1]} + 1.$$

Expressions for $\mathbb{E}[T_1]$ and $\mathbb{V}[T_1]$ can be obtained from the probability distribution function, given by Yoon *et al* (2011, (6)). Finally, we set $\rho = 1/2$, $x_1 = 0.0848$ and $x_2 = 5.7801$.

Structural. To implement structural dependence, we follow the approach in (4) from Section 2.2.5. The intensity function of the Poisson process, conditioned on $\mathcal{X}(t)$, is given by $\partial p_i^X(t)/\partial t$, where $p_i^X(t)$ is defined in Section 2.3. We approximate this intensity $\lambda_i^X(t)$ numerically by solving the Volterra equation from Peskir and Shiryaev (2006, Theorem 14.3) using the discretization method of McLeish and Metzler (2011, Section 4.1). The factor process $X(t)$ follows a standard Brownian motion, with drift and volatility set at $\mu_i = 0.05$ and $\sigma_i = 0.3$ for all i, consistent with Zhou (2001). The default threshold $\log B_i/A_i(0)$ is calibrated to match the unconditional observed annual default probability $p_i(1) = 0.18\%$, yielding $\log B_i/A_i(0) = -0.8907$, where $p_i(1)$ is determined using (2.4). Since we consider a single risk factor process $X(t)$ and a homogeneous portfolio, the dependence parameter w is the same for each obligor and represents their correlation (ie, $\mathbb{E}[W_i(t)W_j(t)] = w^2 t$ for $i \neq j$). The value of w is determined as before by matching the model's normalized variance for the default probability to historical observations. Following Gordy (2000), the variance of the default probability in this structural model is given by

$$\mathbb{P}\bigg(\inf_{t\in[0,1]} \mu_1 t + \sigma_1 W_1(t) < \log \frac{B_1}{A_1(0)},\\ \inf_{t\in[0,1]} \mu_2 t + \sigma_2 W_2(t) < \log \frac{B_2}{A_2(0)}\bigg) - p_i(1)^2.$$

An expression for this probability, evaluated numerically, is given by He *et al* (1998). By equating the model and observed variances, we find $w = 0.1227$, which closely aligns with the systemic risk weight in the static structural model (Gordy 2000, Table 2).

By considering multiple dependence structures and various parameters (see Section 4.3), we ensure that our findings are general and not dependent on any one specific instance of the model.

In the following subsections, all the evaluations are retrieved by simulation using PYTHON. The quantiles of a distribution are determined using linear interpolation where necessary. To ensure accuracy we performed 2 million simulations (1 million for the structural dependence due to numerical complexity) using a time discretization step of 0.005 years.

We note that current industry practices commonly rely on simulation approaches and often implement a multiperiod discretized time grid to capture cashflows and

FIGURE 1 Continuous-time CreditRisk+ sample paths.

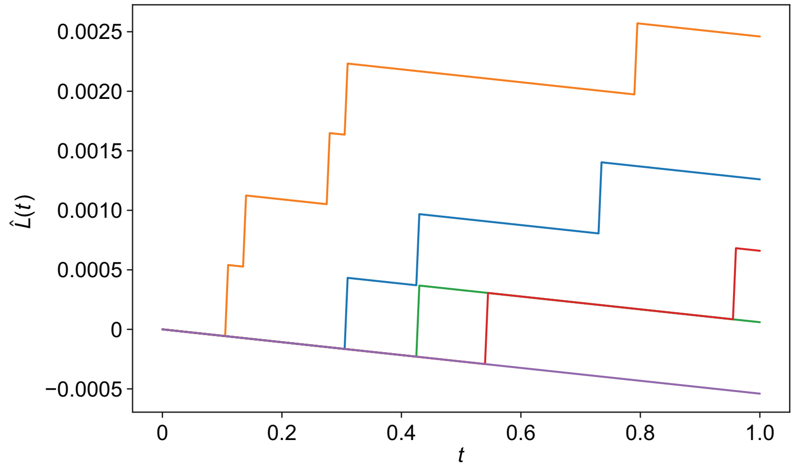

migration risk. Consequently, the need for simulation should not be seen as a significant drawback of the model, especially given its solid intuitive foundation. For large-scale portfolios, computational efficiency can be greatly improved through parallelization and variance-reduction techniques, such as importance sampling.

4.1 The effect of time in the CreditRisk+ model

As mentioned, the original or static CreditRisk+ model only considers the loss distribution at a certain point in time (ie, $\hat{L}(1)$). In practice, however, the maximum of the loss process is often attained not at the end of the interval but prior to that: see the sample paths of $\hat{L}(t)$ in the CreditRisk+ model in Figure 1. In the continuous-time CreditRisk+ model we therefore consider the distribution of $\sup_{t\in[0,1]} \hat{L}(t)$ instead. In this subsection we investigate the difference between these two distributions and the resulting capital estimates (based on quantiles as in (2.3) and (2.1)).

In Figure 2 the cumulative distribution functions of the two variables $\hat{L}(1)$ (static) and $\sup_{t\in[0,1]} \hat{L}(t)$ (continuous time) are compared assuming a gamma dependence structure. The static distribution is piecewise linear, whereas the continuous-time CreditRisk+ model gives a smoother cumulative distribution function due to the effect of the compensating function $h(t)$, which is considered not only at the end of the time interval but also at all times prior to that. The continuous-time CreditRisk+ cumulative distribution lies to the right of its static counterpart, illustrating that the continuous-time CreditRisk+ distribution obtains higher loss values than

FIGURE 2 Comparison of static and continuous-time distributions with gamma dependence.

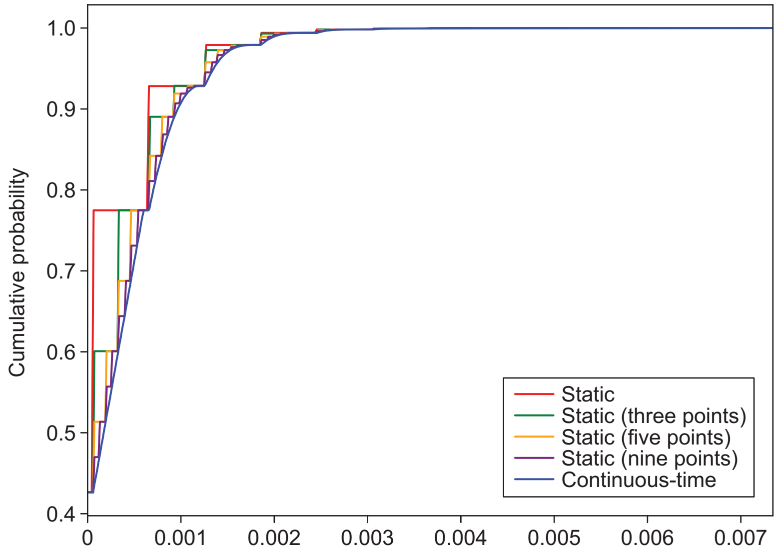

its static counterpart with the same probability. This is what we would expect. The difference between the static and continuous-time CreditRisk+ distributions is less pronounced for large tail percentiles/quantiles of the distribution. As the number of defaults/jumps increases, the difference between the maximum value and the value at the one-year horizon decreases. This phenomenon is evident in the sample paths depicted in Figure 1 and can also be observed from certain distributional statistics in Tables 1 and 2. Considering the gamma dependence structure, the 75% quantile differs by a factor of 10, whereas the 99% quantile only differs by 8% and the 99.9% quantile differs only by 2.6%. In capital calculations it is often the higher quantiles that are of interest, while provisions are often calculated based on expected losses. For the lower quantiles we note that the static $\hat{L}(1)$ can also obtain negative values, corresponding to losses smaller than the expected losses. The supremum function, however, is naturally floored at zero. This effect is most pronounced in the (relatively) lower tail percentiles of the distribution, where there are fewer defaults/jumps.

To bridge the gap between the static and continuous-time models, we could consider taking the maximum of a finite number of points in the interval. More specifically, instead of considering only $\hat{L}(1)$, we could consider two points in time and take

TABLE 1 Static CreditRisk+ distribution $\hat{L}(1)$ with different dependence structures.

	Discrete	Gamma	GBM	RS	Structural
Mean	−1.5510e−7	−1.5510e−7	−3.9000e−9	−2.4633e−6	−3.0420e−7
SD	6.0845e−4	6.0882e−4	6.0776e−4	6.0670e−4	6.0794e−4
Skew	1.3851	1.3689	1.7437	1.3002	1.2790
Kurtosis	5.7517	5.7343	13.6025	5.1048	5.0151
0.5	6.0000e−5	6.0000e−5	6.0000e−5	6.0000e−5	6.0000e−5
0.75	6.0000e−5	6.0000e−5	6.0000e−5	6.0000e−5	6.0000e−5
0.99	1.8600e−3	1.8600e−3	1.8600e−3	1.8600e−3	1.8600e−3
0.999	3.0600e−3	3.0600e−3	3.0600e−3	3.0600e−3	3.0600e−3

SD, standard deviation. GBM, geometric Brownian motion. RS, regime switching.

TABLE 2 Continuous-time CreditRisk+ distribution $\sup_{t\in[0,1]} \hat{L}(t)$ with different dependence structures.

	Discrete	Gamma	GBM	RS	Structural
Mean	3.6020e−4	3.6000e−4	3.6129e−4	3.5184e−4	2.6908e−4
SD	4.8264e−4	4.8172e−4	4.8370e−4	4.7518e−4	4.4165e−4
Skew	1.9733	1.9589	2.6152	1.8873	2.3773
Kurtosis	8.5744	8.6248	25.8038	7.6605	10.0109
0.5	1.7340e−4	1.7340e−4	1.8420e−4	1.5720e−4	8.4300e−5
0.75	5.5950e−4	5.5950e−4	5.6220e−4	5.5410e−4	2.5710e−4
0.99	2.0193e−3	2.0112e−3	1.9815e−3	1.9896e−3	1.8897e−3
0.999	3.1518e−3	3.1410e−3	3.1950e−3	3.0924e−3	3.0654e−3

SD, standard deviation. GBM, geometric Brownian motion. RS, regime switching.

their maximum, $\max(\hat{L}(0), \hat{L}(1))$, or three points in time, with $\max(\hat{L}(0), \hat{L}(0.5), \hat{L}(1))$. We refer to these as the static case with multiple evaluation points. In Figure 2 we plot the static case with three, five and nine points in time. When considering m uniformly spaced evaluation points between 0 and t, we will denote this case by $\hat{L}^{(m)}(t)$. When the number of evaluation points increases, the distribution comes closer to the continuous-time distribution, as we would expect. In Figure 3 the maximum of the absolute error $|\hat{L}^{(m)}(1) - \sup_{t\in[0,1]} \hat{L}(t)|$ is given as a function of the number of evaluation points m. The error is computed per simulation and the maximum over all simulation runs is shown. A small number of evaluation points is sufficient to achieve a substantial reduction in approximation error; for instance, with just 13 uniformly spaced points (including the starting point $t = 0$), the error has already dropped significantly, and further increases in resolution yield only marginal improvements. In other words, assessing losses on a monthly basis and taking the

FIGURE 3 Maximum absolute error across simulations between the static approximation with multiple evaluation points $\hat{L}^{(m)}(1)$, based on m uniformly spaced time points, and the continuous-time maximum $\sup_{t\in[0,1]}\hat{L}(t)$.

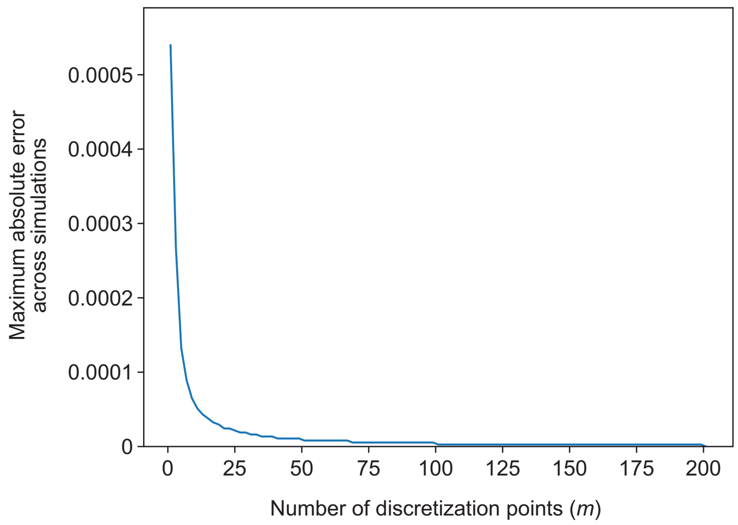

maximum across all months yields results much closer to the continuous-time case than those of the static approach of evaluating losses only at maturity.

4.2 Comparison between the structural and CreditRisk+ models

Gordy (2000) shows that the mapping of a static structural model to the original (static) CreditRisk+ model produces $\hat{L}(1)$ distributions that are roughly similar. The same can be observed for the continuous-time distribution $\sup_{t\in[0,1]}\hat{L}(t)$: see Table 3.

As the derivation of the CreditRisk+ model assumes small probabilities of default, we include a similar comparison for a homogeneous CCC rated portfolio in the same table. The annual default probability of these CCC rated obligors is $p_i(1) = 19.14\%$, and the expected one-year loss is given by $r = 0.3 \times 19.14\% = 0.05742$. Further, the other parameters are calibrated using the same values as before, giving $\log B_i/A_i(0) = -0.3561$ and $w = 0.2852$. We observe, in line with Gordy (2000), that as the credit quality deteriorates, the extreme percentile values in the continuous-time CreditRisk+ increase more quickly than in the continuous-time structural model. The difference, however, still remains below 2%.

TABLE 3 Comparison between the continuous-time structural model and CreditRisk+ with structural dependence for a BBB rated and CCC rated portfolio.

	BBB		CCC	
	CreditRisk+	**Structural**	**CreditRisk+**	**Structural**
Mean	2.6908e−4	2.6835e−4	1.0229e−2	1.0180e−2
SD	4.4416e−4	4.4034e−4	1.5625e−2	1.5470e−2
Skew	2.3773	2.3814	1.9509	1.9239
Kurtosis	10.0109	10.0567	7.2140	7.0403
0.5	8.4300e−5	8.16e−5	7.9650e−4	8.6730e−4
0.75	2.5710e−4	2.5440e−4	1.6077e−2	1.6077e−2
0.99	1.8897e−3	1.8897e−3	6.5841e−2	6.4952e−2
0.999	3.0654e−3	3.0654e−3	9.2241e−2	9.0780e−2

SD, standard deviation.

4.3 The sensitivity of the results to various parameters

In this subsection we demonstrate the effects of some of the parameters in the models, such as the systemic volatility ξ, the rate r and the number of obligors in the portfolio. While the numerical experiments are based on the gamma dependence structure, similar results are obtained for the other dependence structures (discrete, geometric Brownian motion and regime switching).

The systemic volatility ξ controls the shape of the distribution of $\int_0^1 X(t)\,\mathrm{d}t$, and tail probabilities are sensitive to the choice of ξ. There is, however, no obvious additional information available to determine the value of this parameter. As mentioned by Gordy (2000), the difficulty of calibrating ξ should not be interpreted as a disadvantage of the CreditRisk+ model compared with the structural model, because the structural model simply does not allow for this additional flexibility; the structural model imposes very strong restrictions on the shape of the distribution tail by assuming a Brownian motion risk factor process.

In the original CreditRisk+ a single-factor model calibration of the model (ie, $w = 1$) is used to determine the value of ξ as 1 (see Credit Suisse Financial Products 1997, Section A7.3). For our homogeneous portfolio of BBB rated obligors, this would yield $\xi = 0.4$ and $w = 1$. To illustrate the sensitivity of the systemic volatility ξ, simulation results of the distribution with gamma dependence are presented in Table 4. For $\xi = 0.4$, the skew and kurtosis of the static $\hat{L}(1)$ and continuous-time $\sup_{t\in[0,1]} \hat{L}(t)$ distributions are closer together than under all other considered values, resulting in tail distributions that are also the closest together. As ξ increases, both the static and continuous-time CreditRisk+ distributions exhibit heavier tails. While standard deviations remain relatively unchanged, 99.9% tail per-

TABLE 4 The effect of the systemic volatility on the tail distribution of static and continuous-time (CT) CreditRisk+ with gamma dependence.

	$\xi = 0.4$		$\xi = 1$		$\xi = 2$		$\xi = 4$	
	Static	**CT**	**Static**	**CT**	**Static**	**CT**	**Static**	**CT**
Mean	3.5070e−7	3.6111e−4	4.1130e−7	3.6106e−4	5.6640e−7	3.6039e−4	8.0550e−7	3.5896e−4
SD	6.0869e−4	4.8067e−4	6.0871e−4	4.8195e−4	6.0893e−4	4.8394e−4	6.0860e−4	4.8439e−4
Skew	1.2694	1.8082	1.3171	1.8938	1.4037	2.0474	1.5708	2.3221
Kurtosis	4.9387	7.2058	5.2925	7.9592	6.0540	9.6494	8.0444	13.7018
0.5	6.0000e−5	1.6800e−4	6.0000e−5	1.7070e−4	6.0000e−5	1.7340e−4	6.0000e−5	1.7880e−4
0.75	6.0000e−5	5.6490e−4	6.0000e−5	5.6220e−4	6.0000e−5	5.5950e−4	6.0000e−5	5.5680e−4
0.99	1.8600e−3	1.9977e−3	1.8600e−3	2.0058e−3	1.8600e−3	2.0112e−3	1.8600e−3	2.0031e−3
0.999	3.0600e−3	3.0843e−3	3.0600e−3	3.1194e−3	3.0600e−3	3.1869e−3	3.6600e−3	3.6816e−3

SD, standard deviation.

FIGURE 4 The 99% and 99.9% quantiles of the static and continuous-time (CT) CreditRisk+ models with $h(t) = rt$ as a function of r.

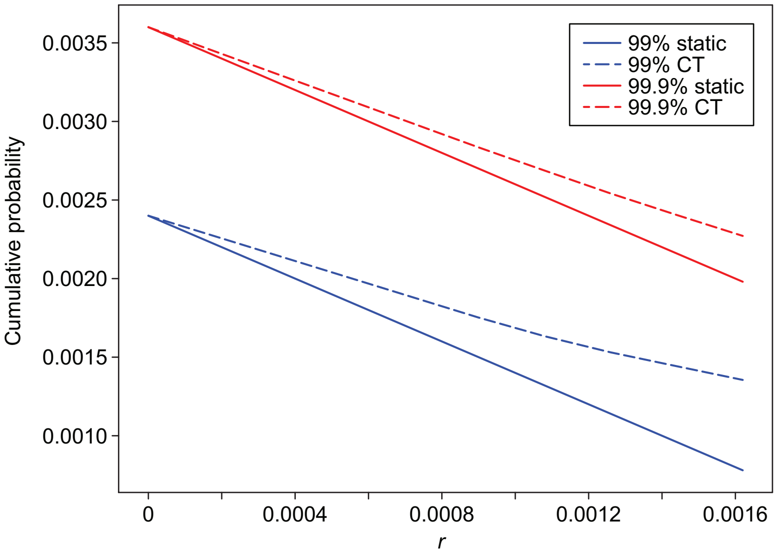

centiles rise significantly, consistent with findings of Gordy (2000) on the original CreditRisk+ model. As the static distribution remains rather stable over time, due to its piecewise linear form as observed in Figure 2, the difference between the static and continuous-time distribution also generally increases if the systemic volatility increases.

If the negative slope of the sample paths increases (ie, a higher value of r), the difference between the supremum $\sup_{t\in[0,1]} \hat{L}(t)$ and the final value $\hat{L}(1)$ also grows. Figure 4 illustrates the 99% and 99.9% quantiles for the static $\hat{L}(1)$ and continuous-time $\sup_{t\in[0,1]} \hat{L}(t)$ CreditRisk+ model under a gamma dependence structure, showing that higher r leads to diverging tail percentages. In this specific instance the static model reflects a lower level of risk than the continuous-time CreditRisk+ model as r increases, with a 99% quantile difference of 8% for $r = \mathbb{E}[\hat{L}(1)] = 0.00054$, 25% for $r = 2\mathbb{E}[\hat{L}(1)]$ and 75% for $r = 3\mathbb{E}[\hat{L}(1)]$. The difference is less pronounced for higher tail percentiles, as we concluded before.

Finally, we also change the number of obligors in the portfolio. The original portfolio contained only 500 obligors and, as a result, the static cumulative distribution function is discrete with a relatively low number of possible values in practice. Due to this discrete distribution for the static model, the difference between the static

TABLE 5 The effect of the number of obligors on the distribution of static and continuous-time (CT) CreditRisk+ with gamma dependence.

	No. of obligors = 1000		**No. of obligors = 5000**		**No. of obligors = 10 000**	
	Static	**CT**	**Static**	**CT**	**Static**	**CT**
Mean	2.7105e−7	2.7970e−4	5.5800e−7	1.5360e−4	−3.5175e−7	1.2287e−4
SD	4.5715e−4	3.4689e−4	2.8138e−4	2.1163e−4	2.5068e−4	1.9353e−4
Skew	1.2660	2.0864	1.8300	3.1757	2.2116	3.6280
Kurtosis	5.9022	10.0676	9.6803	19.2324	11.8204	23.2423
0.5	6.0000e−5	1.8120e−4	−6.0000e−5	8.4600e−5	−6.0000e−5	5.5200e−5
0.75	3.6000e−4	4.2450e−4	1.2000e−4	1.9950e−4	9.0000e−5	1.4430e−4
0.99	1.5600e−3	1.5681e−3	1.0200e−3	1.0227e−3	9.3000e−4	9.4620e−4
0.999	2.4600e−3	2.4708e−3	1.7400e−3	1.7508e−3	1.6500e−3	1.6527e−3

SD, standard deviation.

FIGURE 5 Difference between the simulated and exact cumulative density functions for continuous-time CreditRisk+ with discrete dependence.

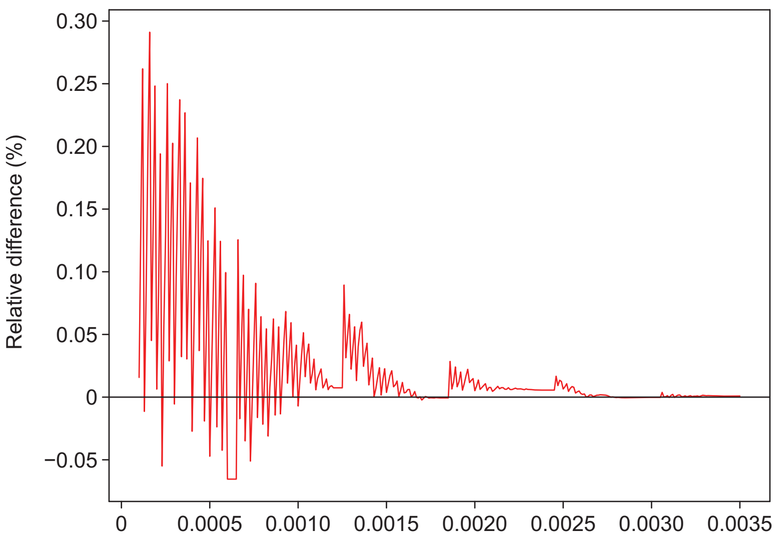

and continuous-time CreditRisk+ models is fairly large. In Table 5 we present the results for the gamma dependence CreditRisk+ model with portfolio sizes of 1 000, 5 000 and 10 000 obligors. The distribution of these portfolios has a larger number of attainable values for the static case, and as a result the static and continuous-time distributions are closer together, especially for the high tail percentiles such as 99.9%. In other words, a small portfolio size may lead to a large difference between the static and continuous-time CreditRisk+ distributions. Further, as observed by Gordy (2000), risk is reduced significantly when increasing the number of obligors. Specifically, the standard deviation of the portfolio with 10 000 obligors is roughly 80% smaller than that of the portfolio with 1000 obligors, and the 99.9% quantile values drop by about 50%.

4.4 Using ruin theory results

Section 3 elaborated on the use of ruin theory methods for determining the distribution of losses in the continuous-time CreditRisk+ model. Here, we implement some of these methods to showcase their workings and their accuracy in comparison with the simulation approaches applied so far. We apply general ruin theory methods to derive the distribution of the supremum of the loss process, $\sup_{t\in[0,1]} \hat{L}(t)$, using recursive formulas instead of relying on simulation. When conditioning on the factor

TABLE 6 Quantile comparison of simulation and exact results for continuous-time CreditRisk+ with a discrete dependence structure.

	Simulation	Exact
0.5	1.73e−4	1.75e−4
0.75	5.60e−4	5.59e−4
0.99	2.019e−3	2.024e−3
0.999	3.152e−3	3.156e−3

process $X(t)$, the distribution of $\sup_{t\in[0,1]} \hat{L}(t)$ can be derived from ruin probabilities. In fact, for discrete and gamma dependence, conditional on the factor process, the losses $\hat{L}(t)$ take on the form of a classical Cramér–Lundberg model with constant claims. A recursive formula for calculating the finite-time ruin probability in the classical Cramér–Lundberg model with constant claims is provided by Rullière and Loisel (2004, Theorem 2.6). We will apply this formula to the continuous-time CreditRisk+ model incorporating a discrete dependence structure. In Figure 5 we present the difference between the cumulative distribution function of simulation-driven results and the distribution obtained by using the exact recursive formula for ruin probabilities. Table 6 shows that the results are near-exact. In addition to being highly accurate, this recursive approach is also computationally more efficient: in our implementation on the same system, the simulation-based method is almost 10 times slower than the exact ruin theory approach.

4.5 Summary

The numerical results obtained in this section are summarized below.

- The continuous-time CreditRisk+ model offers an accurate representation of risk over the full time horizon. In specific instances, the static (final time point) CreditRisk+ distribution may indicate lower risk and required capital than are captured by the continuous-time model. However, differences between the static and continuous-time approaches tend to diminish in the upper-tail percentiles of the distribution, regardless of the dependence structure. A practical approximation of the continuous-time model can be obtained by evaluating losses at a finite number of uniformly spaced points throughout the time horizon. Importantly, even a limited number of such points yields a maximum of the loss process that closely aligns with the continuous-time outcome.

- Having high values of systemic volatility ξ and/or a high positive compensating function $h(t)$ results in a large difference between the static and continuous-time CreditRisk+ model outcomes. Moreover, a small number of obligors in the portfolio will also have a more pronounced effect.

- Using a structural dependence structure in the continuous-time CreditRisk+ model results in a loss distribution very close to the loss distribution of the actual continuous-time structural model itself. This is the case even for large default portfolios.

- General ruin theory methods should be employed where possible to obtain more accurate loss distributions in the (continuous-time) CreditRisk+ model.

5 DISCUSSION AND CONCLUDING REMARKS

This paper extended the original CreditRisk+ model to a continuous-time version that models credit/default losses continuously rather than only at the end of a time interval. The extension not only accounts for outgoing cashflows due to defaults but also includes a positive compensating function representing, for example, interest rate premiums or (expected) losses that may already be accounted for. As a result, the maximum of the loss process over the time interval may occur prior to the end of the interval. The continuous-time CreditRisk+ model offers an accurate representation of risk over the full time horizon. The difference between the static/original CreditRisk+ model and the continuous-time model was demonstrated through a series of numerical experiments and is strongly influenced by the level of systemic volatility and the compensating function. Further, the introduced continuous-time model is a flexible framework for portfolio credit risk assessment, accommodating various risk factors and dependence structures, including the dependence used in the widely used structural models that are the basis for the current credit risk capital requirements for banks. Finally, the portfolio credit risk models from the banking industry were successfully connected to ruin theory models used in the insurance sector. Notably, it was demonstrated that ruin theory methods can replace simulation methods for credit risk capital calculations in the continuous-time CreditRisk+ model. The analogies established between these model classes pave the way for broader cross-fertilization of results and methods, which will be a focus of future research.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

The research of Michel Mandjes was partly funded by the Netherlands Organisation for Scientific Research (NWO) Gravitation Project "Networks" under grant 024.002.003.

REFERENCES

Adékambi, F., and Essiomle, K. (2020). Ruin probability for stochastic flows of financial contract under phase-type distribution. *Risks* **8**(2), 1–21 (https://doi.org/10.3390/risks8020053).

Akkaya, N., Kurth, A., and Wagner, A. (2004). Incorporating default correlations and severity variations. In *CreditRisk+ in the Banking Industry*, Gundlach, M., and Lehrbass, F. (eds), pp. 129–152. Springer (https://doi.org/10.1007/978-3-662-06427-6_9).

Asmussen, S. (1989). Risk theory in a Markovian environment. *Scandinavian Actuarial Journal* **1989**(2), 69–100 (https://doi.org/10.1080/03461238.1989.10413858).

Asmussen, S., and Albrecher, H. (2010). *Ruin Probabilities*, 2nd edn. World Scientific (https://doi.org/10.1142/7431).

Basel Committee on Banking Supervision (2005). An explanatory note on the Basel II IRB risk weight functions. Working Paper, July, Bank for International Settlements, Basel. URL: www.bis.org/bcbs/irbriskweight.pdf.

Basel Committee on Banking Supervision (2017). Basel III: finalising post-crisis reforms. Standards Document, Bank for International Settlements, Basel. URL: www.bis.org/bcbs/publ/d424.pdf.

Berd, A. (2005). Dynamic estimation of credit rating transition probabilities. Preprint, arXiv (https://doi.org/10.48550/arXiv.0912.4621).

Binnenhei, C. (2004). An analytic approach to rating transitions. In *CreditRisk+ in the Banking Industry*, Gundlach, M., and Lehrbass, F. (eds), pp. 167–185. Springer (https://doi.org/10.1007/978-3-662-06427-6_12).

Black, F., and Cox, J. (1976). Valuing corporate securities: some effects of bond indenture provisions. *Journal of Finance* **31**(2), 351–367 (https://doi.org/10.1111/j.1540-6261.1976.tb01891.x).

Boudoukh, J., Richardson, M., Stanton, R., and Whitelaw, R. (2004). MaxVaR: long-horizon value at risk in a mark-to-market environment. *Journal of Investment Management* **2**(3), 1–6 (https://doi.org/10.2139/ssrn.520805).

Chen, C., and Panjer, H. (2009). A bridge from ruin theory to credit risk. *Review of Quantitative Finance and Accounting* **32**, 373–403 (https://doi.org/10.1007/s11156-008-0100-0).

Chongo, A., and Julia-Sala, F. (2024). A new correlated multi-period multi-state credit risk simulation model. Working Paper, Social Science Research Network (https://doi.org/10.2139/ssrn.4863510).

Credit Suisse Financial Products (1997). CreditRisk+: a credit risk management framework. Technical Document, Credit Suisse Financial Products.

Dębicki, K., and Mandjes, M. (2015). *Queues and Lévy Fluctuation Theory*. Springer (https://doi.org/10.1007/978-3-319-20693-6).

Delsing, G. (2022). Ruin theory for portfolio risk modeling in banking and insurance. PhD Thesis, University of Amsterdam. URL: https://hdl.handle.net/11245.1/caefaa8d-8faf-4e42-bc21-4567e6e76959.

Delsing, G., and Mandjes, M. (2021). A transient Cramér–Lundberg model with applications to credit risk. *Journal of Applied Probability* **58**(3), 721–745 (https://doi.org/10.1017/jpr.2020.114).

Delsing, G., Mandjes, M., Spreij, P., and Winands, E. (2022). On capital allocation for a risk measure derived from ruin theory. *Insurance: Mathematics and Economics* **104**, 76–98 (https://doi.org/10.1016/j.insmatheco.2022.02.001).

Dickson, D., and Qazvini, M. (2018). Ruin problems in Markov-modulated risk models. *Annals of Actuarial Science* **12**(1), 23–48 (https://doi.org/10.1017/S1748499517000124).

Fischer, M., and Dietz, C. (2011). Modeling sector correlations with CreditRisk+: the common background vector model. *The Journal of Credit Risk* **7**(4), 23–43 (https://doi.org/10.21314/JCR.2011.134).

García-Céspedes, R., and Moreno, M. (2017). An approximate multi-period Vasicek credit risk model. *Journal of Banking and Finance* **81**, 105–113 (https://doi.org/10.1016/j.jbankfin.2017.05.002).

Gordy, M. (2000). A comparative anatomy of credit risk models. *Journal of Banking and Finance* **24**(1), 119–149 (https://doi.org/10.1016/S0378-4266(99)00054-0).

Gundlach, M., and Lehrbass, F. (eds) (2004). *CreditRisk+ in the Banking Industry*. Springer (https://doi.org/10.1007/978-3-662-06427-6).

Gupton, G. M., Finger, C. C., and Bhatia, M. (1997). CreditMetrics$^{\text{TM}}$: technical document. Morgan Guaranty Trust Company, New York. URL: https://bit.ly/3YOcXAy.

Han, C., and Kang, J. (2014). An extended CreditRisk+ framework for portfolio credit risk management. Working Paper, Social Science Research Network (https://doi.org/10.2139/ssrn.2427489).

He, H., Keirstead, W., and Rebholz, J. (1998). Double lookbacks. *Mathematical Finance* **8**(3), 201–228 (https://doi.org/10.1111/1467-9965.00053).

Huang, Z., and Kwok, Y. (2021). Efficient risk measures calculations for generalized CreditRisk+ models. *International Journal of Theoretical and Applied Finance* **24**(2), Paper 2150012 (https://doi.org/10.1142/S0219024921500126).

Lefèvre, C., and Loisel, S. (2009). Finite-time ruin probabilities for discrete, possibly dependent, claim severities. *Methodology and Computing in Applied Probability* **11**(3), 425–441 (https://doi.org/10.1007/s11009-009-9123-9).

McLeish, D., and Metzler, A. (2011). A multiname first-passage model for credit risk. *The Journal of Credit Risk* **7**(1), 35–64 (https://doi.org/10.21314/JCR.2011.122).

Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* **29**(2), 449–470 (https://doi.org/10.2307/2978814).

Peskir, G., and Shiryaev, A. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, Basel (https://doi.org/10.1007/978-3-7643-7390-0).

Picard, P., and Lefèvre, C. (1997). The probability of ruin in finite-time with discrete claim size distribution. *Scandinavian Actuarial Journal* **1997**(1), 58–69 (https://doi.org/10.1080/03461238.1997.10413978).

Reiß, O. (2004). Dependent sectors and an extension to incorporate market risk. In *CreditRisk+ in the Banking Industry*, Gundlach, M., and Lehrbass, F. (eds), pp. 215–230. Springer (https://doi.org/10.1007/978-3-662-06427-6_13).

Rullière, D., and Loisel, S. (2004). Another look at the Picard–Lefèvre formula for finite-time ruin probabilities. *Insurance: Mathematics and Economics* **35**(2), 187–203 (https://doi.org/10.1016/j.insmatheco.2004.07.001).

Schlottmann, F., Seese, D., Lesko, M., and Vorgrimler, S. (2004). Risk–return analysis of credit portfolios. In *CreditRisk+ in the Banking Industry*, Gundlach, M., and Lehrbass, F. (eds), pp. 259–278. Springer (https://doi.org/10.1007/978-3-662-06427-6_16).

Vašíček, O. (2002). Loan portfolio value. *Risk* **15**(12), 160–162. URL: www.risk.net/1528714/.

Yang, H. (2003). Ruin theory in a financial corporation model with credit risk. *Insurance: Mathematics and Economics* **33**(1), 135–145 (https://doi.org/10.1016/S0167-6687(03)00149-5).

Yoon, J., Jang, B., and Roh, K. (2011). An analytic valuation method for multivariate contingent claims with regime-switching volatilities. *Operations Research Letters* **39**(3), 180–187 (https://doi.org/10.1016/j.orl.2011.02.010).

Zhou, C. (2001). An analysis of default correlations and multiple defaults. *Review of Financial Studies* **14**(2), 555–576 (https://doi.org/10.1093/rfs/14.2.555).