Valuation of Forward-Starting Basket Default Swaps

by

Wanhe Zhang

Research Paper

Department of Computer Science
University of Toronto

Abstract

Valuation of Forward-Starting Basket Default Swaps

Wanhe Zhang

Research Paper

Graduate Department of Computer Science

University of Toronto

2007

A basket default swap (BDS) is a credit derivative with contingent payments that are triggered by a combination of default events of the reference entities. A forward-starting BDS is a basket default swap starting at a specified future time. In this paper, we study valuation methods for a forward-starting BDS. We begin by reviewing the popular factor copula model. The widely used Monte Carlo method and associated variance reduction techniques are surveyed. The analytical solution of a recursive algorithm is developed. Conditional on a specified common factor, we explore the possible combination of defaults during the life of the forward contract; under each scenario, we evaluate the object functions; we finish the valuation by computing the expectations of these object functions. The possible combination of defaults results in a large combinatorial problem. In order to overcome the inefficiency of the method outlined above, a more applicable approximation method that omits or interpolates the unimportant scenarios is proposed. Numerical results compare the accuracy and performance of these methods and illustrate the effect of contract parameters.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The credit derivative market has grown explosively during the last 10 years. Among these credit derivatives, the most sophisticated ones are the products associated with a portfolio of underlying assets, such as basket default swaps (BDS) and collateralized debt obligations (CDO). In this chapter, we discuss the mechanics of forward-starting BDS, survey the literature, state the purpose of this paper, and outline the paper structure.

1.1 The mechanics of BDS

A basket default swap (BDS) is a credit derivative, the underlying assets of which are corporate bonds or other assets subject to credit risk. In an mth-to-default BDS, the protection buyer pays a specified rate (known as the premium or spread) on a specified notional principal periodically until the mth default occurs among the reference entities or until the maturity of the contract. If the mth default happens prior to the maturity of the BDS, the protection seller pays the losses caused by the mth default only to the protection buyer. This arrangement is depicted in Figure 1.1.

A forward-starting BDS is a forward contract obligating the holder to buy or sell a BDS at a specified future time. For example, such a contract might obligate the holder to buy five-year protection on a second-to-default BDS with 10 reference entities. Suppose

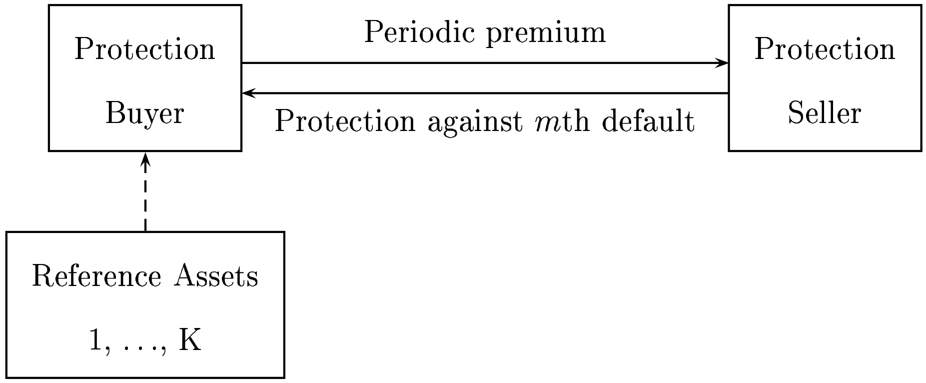

Figure 1.1: Cash flows of an mth-to-default BDS

the contract starts one year later and the premium is 100 basis points per year. During the first year, there is no payment between the buyer and the seller. At the end of the first year, if three reference entities have defaulted, the forward contract obligates the holder to enter a five-year second-to-default BDS on the remaining seven reference entities. The premium is 100 basis points per year on the outstanding notional values.

We denote

T: The maturity date of the forward contract, also the starting date of the BDS;

T^*: The maturity date of the BDS;

T_i: The swap premium dates, for $i = 1, \ldots, n$, and $T = T_0 < T_1 < \ldots < T_n = T^*$.

Figure 1.2 illustrates the cash flows of a forward-starting mth-to-default BDS from the protection seller's point-of-view. Whether the BDS starts or not is determined by the number of entities left in the basket at T: if less than m names survive till T, the contract terminates without any payments as shown in case (a); if at least m entities survive till T, the BDS starts and the cash flows are the same as those in a normal BDS starting at T as shown in cases (b) and (c).

In this paper, if all the underlying names in the reference pool have identical recovery-adjusted notional values, identical default correlations and identical risk-neutral default probabilities, the pool is called *completely homogeneous*. If the recovery-adjusted notional values are the same, the pool is termed *homogeneous*. Otherwise, the pool is named

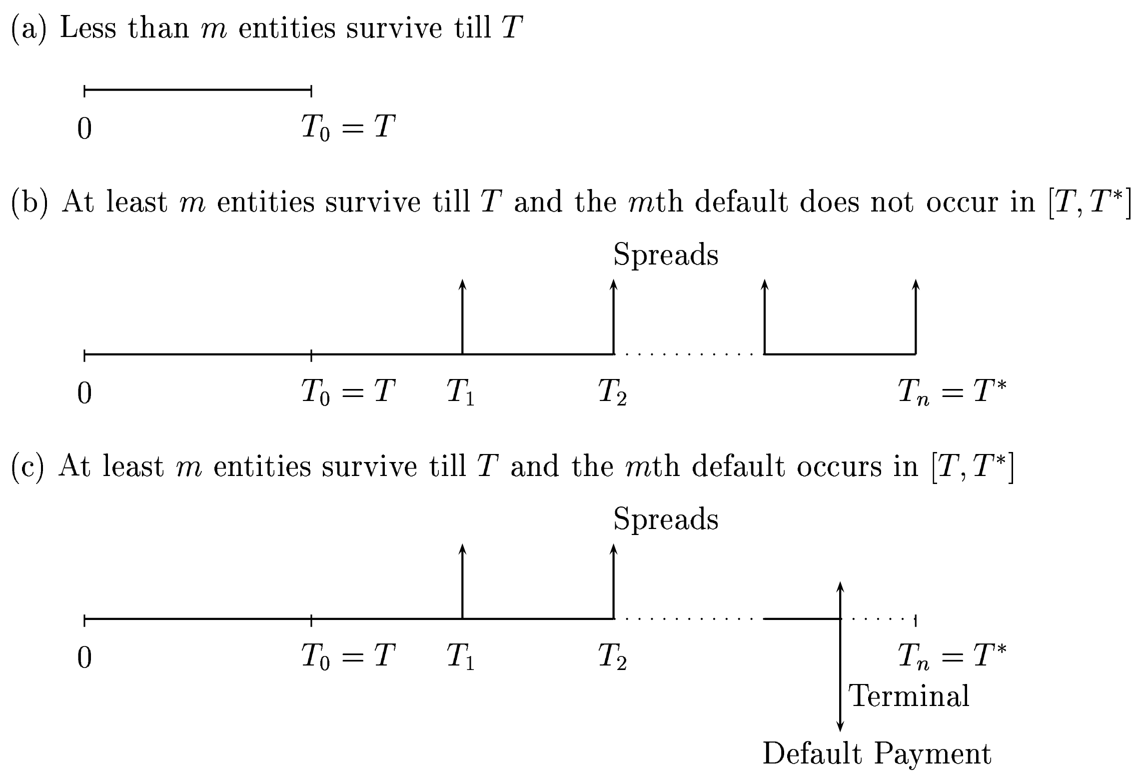

Figure 1.2: Cash flows of a forward-starting mth-to-default BDS

inhomogeneous or *heterogeneous*.

1.2 Literature review

The valuation methods for mth-to-default BDS can be coarsely divided into two classes: Monte Carlo simulations and analytical approaches. The first class consists of Monte Carlo methods (e.g., Andersen et al. [1]) and Monte Carlo methods coupled with variance reduction techniques (e.g., Joshi and Kainth [7] and Chen and Glasserman [3]). Such methods are flexible, but are computationally expensive. Therefore, the second class of methods – analytical approaches – are desirable. The analytical approach computes the expected premium and default payments using probability theory and obtains a closed-form, or semi closed-form, solution.

Factor copula models are widely used in the analytical approach by researchers, e.g., Li [11], Hull and White [5], and Laurent and Gregory [9], to name just a few. The advantage of the factor copula model is that the default probabilities of different names conditional on the constant common factor(s) are assumed to be independent. Under this conditional independence assumption, computing the loss distribution, which is the key step in the valuation of basket credit derivatives, becomes tractable. Among these copula based models, Laurent and Gregory [9] use the fast Fourier transform to calculate the conditional loss distribution as a convolution of the conditional default probabilities of each reference entity. Hull and White [5] derive a recursive relationship to determine the conditional distribution of the number of defaults, assuming the pool is homogeneous. For the inhomogeneous case, they divide the pool into homogeneous buckets, solve them separately and recombine them by the convolution technique. Iscoe and Kreinin [6] use the order statistics of the default times of the names to obtain a recursive relationship for the conditional default distribution. Their method is applicable to any kind of basket. The models based on factor copulas are tractable and suitable for risk management, but they are essentially static and incapable of modeling dynamic evolution.

Recently, another class of analytical or semianalytical models has been introduced by Bennani [2], Schönbucher [13], and Sidenius et al. [14]. Instead of modeling the default distribution of individual names, they model the aggregate portfolio loss as a jump-diffusion process. Although such models allow for dynamic evolution, they cannot easily produce an arbitrage-free loss distribution, and they are unsuitable for risk management.

1.3 The problem

The object of this paper is to explore methods for the valuation of the forward-starting BDS. We study not only the general Monte Carlo simulation, but also analytical methods based on the popular factor copula model.

1.4 Paper structure

The rest of the paper is organized as follows. Chapter 2 derives the pricing equations for both the normal BDS and the forward-starting BDS and reviews the conditional independence framework. Chapter 3 studies the Monte Carlo method, starting with the naive version and then considers improvements based on one of the variance reduction techniques. Chapter 4 gives an analytical solution for the forward-starting BDS, derives solutions for the completely homogeneous pool and applies them to the homogeneous and inhomogeneous cases. Chapter 5 reports the accuracy and performance of the methods. Chapter 6 provides some conclusions and discusses future work.

Chapter 2

Background Knowledge

Because of their tractability, Gaussian factor copula models have become an industry standard to specify the correlation of the underlying assets in multi-name credit derivatives. In this chapter, we first make assumptions for the forward-starting BDS contract. Based on these assumptions, we derive the pricing equations. Then, we review the popular Gaussian factor copula model.

2.1 Pricing equations

There are several kinds of BDS with different payment structures. The BDS considered here is a common one. Its pricing equations can be easily modified to extend our approach to other cases.

2.1.1 Problem assumptions

We consider a BDS containing K instruments with the recovery-adjusted notional value $N^{(k)}$ for name k in the original basket, where $k = 1, 2, \ldots, K$. Assume that the recovery rates are constant and the interest rate process is deterministic. Let $D(t, T)$ be the discount factor at t for payment at T.

We assume no replacement of the underlying assets in the basket and constant premium[1] s. There are several variants of the BDS contract due to different payment structures. For simplicity, we ignore the accrued interest at default for the premium payments (known as the premium leg); for the protection payment (known as the default leg), if the terminal default occurs during the life of the BDS, the compensation is paid out at the nearest premium date following or equal to the terminal default time.

Let $\tau^{(k)}$ denote the default time of the kth name in the basket ($\tau^{(k)} = +\infty$, if name k never defaults). The terminal default time τ, which triggers the default payment, can be expressed as a function of $\tau^{(k)}$, where $k = 1, 2, \ldots, K$. For example, in a normal first-to-default BDS, $\tau = \min_{1 \le k \le K} \tau^{(k)}$. We also define the index of the premium date that is just before the terminal default time by

$$i(\tau) = \max\{i : T_i < \tau\} \tag{2.1}$$

and the default payment time by

$$\bar{\tau} = \begin{cases} T_{i(\tau)+1}, & \tau < T^* \\ T^*, & \tau \ge T^* \end{cases} \tag{2.2}$$

Denote the loss of the forward-starting BDS at the terminal default time by

$$L = \begin{cases} g(N^{(k)}), & \tau = \tau^{(k)} \in (T, T^*] \\ 0, & \text{otherwise} \end{cases}$$

where $g(\cdot)$ is a payoff function. As mentioned in [6], this payoff function is flexible enough to represent different types of default payments, e.g., mth-to-default, digital mth-to-default, and call option on mth default.

[1]This assumption assists us to compute a fair BDS spread, which balances the expected revenue from the payout of the contract. If we are interested only in computing the value of the forward-starting BDS, the restriction can be relaxed and a nonconstant premium considered.

2.1.2 Pricing equations for forward-starting BDS

The value of a forward-starting BDS at time 0 is given by

$$V_{\text{fwd}}(0) = D(0,T)\mathbb{E}[V_{\text{swap}}(T)] \tag{2.3}$$

where $V_{\text{swap}}(T)$, the value of the BDS at T, is defined by

$$V_{\text{swap}}(T) = V_{\text{def}}(T) - V_{\text{prem}}(T) \tag{2.4}$$

where $V_{\text{def}}(T)$ and $V_{\text{prem}}(T)$ are the value of the default leg and the premium leg at T, respectively. Throughout the paper, $\mathbb{E}$ denotes the risk-neutral expectation with respect to the risk-neutral probability $\mathbb{P}$. In order to compute the risk-neutral expectation numerically, we introduce

$$\Pi_{i|T}^{(k)} = \mathbb{P}(\tau = \tau^{(k)}, \tau \in (T_{i-1}, T_i] \mid \mathcal{B}_T), \quad i = 1, 2, \ldots, n, \quad k \in \mathcal{B}_T$$

where $\mathcal{B}_T$ denotes the set of names in the basket at T. If we let $\mathcal{B}$ be the original basket at time 0, then the set of names that default during the life of the forward contract[2] is $\mathcal{B}\backslash\mathcal{B}_T$. We also denote the number of names in $\mathcal{B}_T$ by $|\mathcal{B}_T|$ and the probability distribution of $\mathcal{B}_T$'s composition by $\mathbb{P}(\mathcal{B}_T)$.

Given the loss information till T (which names default in $[0,T]$, or equivalently which names are in $\mathcal{B}_T$), the value of the default leg at T satisfies

$$V_{\text{def}}(T) = \mathbb{E}[L \cdot D(T,\bar{\tau})] = \sum_{k \in \mathcal{B}_T} g(N^{(k)}) \sum_{i=1}^{n} D(T,T_i)\Pi_{i|T}^{(k)} = \sum_{i=1}^{n} D(T,T_i) \sum_{k \in \mathcal{B}_T} g(N^{(k)})\Pi_{i|T}^{(k)} \tag{2.5}$$

where $\bar{\tau}$ is defined in (2.2). Similarly, given the loss information till T, the value of the premium leg at T satisfies

$$V_{\text{prem}}(T) = \mathbb{E}\Big[sN_T \sum_{i=1}^{i(\tau)} \Delta T_i \cdot D(T,T_i)\Big] = sN_T \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i)\bar{\Pi}_{i|T} \tag{2.6}$$

[2]This is $(0,T]$.

where $i(\tau)$ is defined in (2.1); N_T is the sum of the notational values of all names in $\mathcal{B}_T$; $\Delta T_i = T_i - T_{i-1}$; and the survival probabilities $\bar{\Pi}_{i|T} = \mathbb{P}(\tau > T_i \mid \mathcal{B}_T)$ satisfy the relations

$$
\begin{aligned}
\bar{\Pi}_{0|T} &= 1 \\
\bar{\Pi}_{i|T} &= \bar{\Pi}_{i-1|T} - \Pi_{i|T}, \quad i = 1, 2, \ldots, n \\
\Pi_{i|T} &= \sum_{k \in \mathcal{B}_T} \Pi^{(k)}_{i|T}, \quad i = 1, 2, \ldots, n
\end{aligned}
$$

In a forward-starting mth-to-default BDS, if $|\mathcal{B}_T| < m$, the contract terminates without any payment. Thus we need only to consider the case that $m \le |\mathcal{B}_T| \le K$. Equations (2.3) and (2.4) imply

$$
V_{\text{fwd}}(0) = D(0,T)(\mathbb{E}[V_{\text{def}}(T)] - \mathbb{E}[V_{\text{prem}}(T)])
$$

Therefore we need to evaluate $\mathbb{E}[V_{\text{def}}(T)]$ and $\mathbb{E}[V_{\text{prem}}(T)]$. From (2.5), $\mathbb{E}[V_{\text{def}}(T)]$ satisfies

$$
\begin{aligned}
\mathbb{E}[V_{\text{def}}(T)] &= \mathbb{E}\Big[\sum_{i=1}^{n} D(T,T_i) \sum_{k \in \mathcal{B}_T} g(N^{(k)})\Pi^{(k)}_{i|T}\Big] \\
&= \sum_{i=1}^{n} D(T,T_i)\mathbb{E}\Big[\sum_{k \in \mathcal{B}_T} g(N^{(k)})\Pi^{(k)}_{i|T}\Big] \\
&= \sum_{i=1}^{n} D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}(\mathcal{B}_T) \sum_{k \in \mathcal{B}_T} g(N^{(k)})\Pi^{(k)}_{i|T}
\end{aligned} \tag{2.7}
$$

Similarly, from (2.6), $\mathbb{E}[V_{\text{prem}}(T)]$ satisfies

$$
\begin{aligned}
\mathbb{E}[V_{\text{prem}}(T)] &= \mathbb{E}\Big[sN_T \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i)\bar{\Pi}_{i|T}\Big] \\
&= s\sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i)\mathbb{E}\Big[N_T\bar{\Pi}_{i|T}\Big] \\
&= s\sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}(\mathcal{B}_T) N_T \bar{\Pi}_{i|T}
\end{aligned} \tag{2.8}
$$

The above derivation is straightforward, and hence the proof is omitted.

The normal BDS is a special case of the forward-starting BDS, where $T = 0$. Unlike the unknown starting basket structure in the forward-starting BDS, the basket structure in the normal BDS is certain. More specifically, the normal BDS has only one possible basket structure – the original basket $\mathcal{B}$. Denoting $\Pi_{i|0}^{(k)}$ and $\bar{\Pi}_{i|0}^{(k)}$ by $\Pi_i^{(k)}$ and $\bar{\Pi}_i^{(k)}$, respectively, we can write the pricing equations for the normal BDS as

$$\mathbb{E}[V_{\text{def}}(0)] = V_{\text{def}}(0) = \sum_{k=1}^{K} g(N^{(k)}) \sum_{i=1}^{n} D(0,T_i)\Pi_i^{(k)} = \sum_{i=1}^{n} D(0,T_i) \sum_{k=1}^{K} g(N^{(k)})\Pi_i^{(k)}$$

$$\mathbb{E}[V_{\text{prem}}(0)] = V_{\text{prem}}(0) = sN_0 \sum_{i=1}^{n} \Delta T_i \cdot D(0,T_i)\bar{\Pi}_i$$

Similar results can be found in [8], [9], and [6]; the proof is given by Lando [8].

2.2 Gaussian factor copula model

Due to their tractability, Gaussian factor copula models are widely used to specify a joint distribution for default times consistent with the marginal distributions. In this section, we review the one-factor Gaussian model to illustrate the conditional independence framework.

2.2.1 Conditional forward default probabilities

Assume the risk-neutral (cumulative) default probabilities

$$\hat{\pi}^{(k)}(t) = \mathbb{P}(\tau^{(k)} \leq t), \quad k = 1, 2, \ldots, K$$

are known[3]. In order to generate the dependence structure of default times, we introduce random variables U_k that satisfy

$$U_k = \beta_k X + \sigma_k \varepsilon_k, \quad \text{for } k = 1, 2, \ldots, K \tag{2.9}$$

where X is the systematic risk factor reflecting the health of the macroeconomic environment; ε_k are idiosyncratic risk factors which are independent with each other and also

[3] Usually, the risk-neutral default probabilities are implied from the market price of defaultable bonds or credit default swaps. For more details, see [10].

independent with X; the constants β_k and σ_k satisfy

$$\beta_k^2 + \sigma_k^2 = 1, \quad \text{for } k = 1, 2, \ldots, K$$

The random variables X and ε_k follow zero-mean unit-variance distributions, so the correlation between U_i and U_j is $\beta_i \beta_j$.

The default times $\tau^{(k)}$ and the random variables U_k are connected by a percentile-to-percentile transformation, such that

$$\hat{\pi}^{(k)}(t) = \mathbb{P}(\tau^{(k)} \leq t) = \mathbb{P}(U_k \leq u_k(t))$$

Thus the dependence among default times is captured by the common factor X. If we assume X and ε_k follow standard normal distributions, U_k also follows a standard normal distribution, hence we have

$$u_k(t) = \Phi^{-1}(\hat{\pi}^{(k)}(t)). \tag{2.10}$$

where Φ is the standard normal cumulative distribution function.

Conditional on a particular value x of X, the conditional risk-neutral default probabilities are defined as

$$\hat{\pi}^{(k)}(t, x) \equiv \mathbb{P}(\tau_k \leq t \mid X = x) = \mathbb{P}(U_k \leq u_k(t) \mid X = x) \tag{2.11}$$

Substituting (2.9) and (2.10) into (2.11), we have

$$\hat{\pi}^{(k)}(t, x) = \mathbb{P}\big[\beta_k x + \sigma_k \varepsilon_k \leq \Phi^{-1}(\hat{\pi}^{(k)}(t))\big] = \Phi\left[\frac{\Phi^{-1}(\hat{\pi}^{(k)}(t)) - \beta_k x}{\sigma_k}\right]$$

In this framework, the default events of the names are assumed to be conditionally independent. Thus, the problem of correlated names is reduced to the case of independent names. From (2.7), the expected value of the default leg at time T can be evaluated as

$$\begin{aligned}
\mathbb{E}[V_{\text{def}}(T)] &= \int_{-\infty}^{\infty} \mathbb{E}_x[V_{\text{def}}(T)] d\Phi(x) \\
&= \int_{-\infty}^{\infty} \sum_{i=1}^{n} D(T, T_i) \mathbb{E}_x \Big[\sum_{k \in \mathcal{B}_T} g(N^{(k)}) \Pi_{i|T}^{(k)} \Big] d\Phi(x) \\
&= \int_{-\infty}^{\infty} \sum_{i=1}^{n} D(T, T_i) \sum_{|\mathcal{B}_T| = m}^{K} \mathbb{P}_x(\mathcal{B}_T) \sum_{k \in \mathcal{B}_T} g(N^{(k)}) \Pi_{i|T}^{(k)}(x) \, d\Phi(x)
\end{aligned} \tag{2.12}$$

where $\mathbb{E}_x$ denotes the risk-neutral expectation with respect to the risk-neutral probability $\mathbb{P}_x$, conditional on $X = x$. From (2.8), the expected value of the premium leg at time T satisfies

$$\begin{aligned}\mathbb{E}[V_{\text{prem}}(T)] &= \int_{-\infty}^{\infty} \mathbb{E}_x[V_{\text{prem}}(T)]d\Phi(x) \\ &= \int_{-\infty}^{\infty} s \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i)\mathbb{E}_x\left[N_T \bar{\Pi}_{i|T}\right] d\Phi(x) \\ &= \int_{-\infty}^{\infty} s \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) N_T(x) \bar{\Pi}_{i|T}(x)\, d\Phi(x) \qquad (2.13)\end{aligned}$$

Therefore, the main challenge in pricing a forward-starting BDS lies in computing $\mathbb{P}_x(\mathcal{B}_T)$, $N_T(x)$ and $\Pi_{i|T}^{(k)}(x)$. To this end, we introduce the conditional forward default probabilities

$$\hat{\pi}^{(k)}(t|T,x) = \frac{\hat{\pi}^{(k)}(t,x) - \hat{\pi}^{(k)}(T,x)}{1 - \hat{\pi}^{(k)}(T,x)}, \quad \text{for } t > T \qquad (2.14)$$

which represents the risk-neutral probability that, conditional on a specified value x of X and conditional on surviving till T, name k defaults before t.

2.2.2 Conditional forward default intensities

Assume the conditional forward default distribution that name k defaults in $(T, t]$ follows the Cox process

$$\mathbb{P}_x(\tau^{(k)} \le t \,|\, \tau^{(k)} > T) = 1 - \exp\left(-\Lambda^{(k)}(t|T,x)\right) \qquad (2.15)$$

where

$$\Lambda^{(k)}(t|T,x) = \int_T^t \lambda^{(k)}(u|T,x)du \qquad (2.16)$$

and $\lambda^{(k)}(\cdot)$ is the conditional forward default intensity of the kth name. We know

$$\mathbb{P}_x(\tau^{(k)} \le t \,|\, \tau^{(k)} > T) = \hat{\pi}^{(k)}(t|T,x) \qquad (2.17)$$

where $\hat{\pi}^{(k)}(t|T,x)$ is given by (2.14). If we assume $\Lambda^{(k)}(t|T,x)$ is linear between premium dates T_i, then (2.16) implies that $\lambda^{(k)}(t|T,x)$ is a piecewise constant function

$$\lambda^{(k)}(t|T,x) = \lambda_{i|T}^{(k)}(x), \quad \text{for } t \in (T_{i-1}, T_i]$$

where $\lambda_{i|T}^{(k)}(x)$ represent the constant $\lambda^{(k)}(\cdot)$ in $(T_{i-1}, T_i]$. Combining this result with (2.16), we have

$$\Lambda^{(k)}(T_i|T, x) = \Lambda^{(k)}(T_{i-1}|T, x) + \lambda_{i|T}^{(k)}(x)\Delta T_i$$

Furthermore, we obtain

$$\lambda_{i|T}^{(k)}(x) = \frac{1}{\Delta T_i}\Big(\Lambda^{(k)}(T_i|T, x) - \Lambda^{(k)}(T_{i-1}|T, x)\Big) \tag{2.18}$$

From (2.15) and (2.17), we know

$$\Lambda^{(k)}(T_{i-1}|T, x) = -\ln\big(1 - \hat{\pi}_{i-1|T}^{(k)}(x)\big)$$
$$\Lambda^{(k)}(T_i|T, x) = -\ln\big(1 - \hat{\pi}_{i|T}^{(k)}(x)\big)$$

where $\hat{\pi}_{i|T}^{(k)}(x) = \hat{\pi}^{(k)}(T_i|T, x) = \mathbb{P}_x(\tau^{(k)} \leq T_i \mid \tau^{(k)} > T)$. Substituting $\Lambda^{(k)}(T_{i-1}|T, x)$ and $\Lambda^{(k)}(T_i|T, x)$ into (2.18), we obtain

$$\lambda_{i|T}^{(k)}(x) = \frac{1}{\Delta T_i}\ln\Big(\frac{1 - \hat{\pi}_{i-1|T}^{(k)}(x)}{1 - \hat{\pi}_{i|T}^{(k)}(x)}\Big), \quad \text{for } i = 1, 2, \ldots, n \tag{2.19}$$

Chapter 3

Monte Carlo Simulation

Monte Carlo simulation is among the most popular computational tools for the valuation of credit derivatives. As in other application areas, it has the advantage of being flexible, but the disadvantage of being inefficient. This motives the investigation of methods to accelerate simulations through variance reduction. In this chapter, we discuss the naive Monte Carlo method before we describe a more sophisticated method that incorporates stratified sampling – one of the variance reduction techniques.

3.1 Naive method

The main challenge of the Monte Carlo method for basket credit derivatives is to simulate correlated default times with known marginal distributions. The Gaussian factor copula model discussed in Section 2.2 is employed to simulate the correlated default probabilities. Because of the percentile-to-percentile mapping, the default times can be generated from their known marginal distributions. After that, we compute the value of the default leg and the premium leg for this specified simulation. If we replicate the same process many times and average the values, we obtain relatively reliable results. Based on the law of large numbers, the more replications we perform, the more accuracy we obtain. Below we outline Monte Carlo simulation for BDS based on the one-factor Gaussian copula model.

ONE-FACTOR-GAUSSIAN-COPULA($\beta, \hat{\pi}$)

Draw $X \sim N(0,1)$ $\quad \triangleright$ $N(0,1)$ is the standard normal distribution
for $k \leftarrow 1$ **to** K
  do Draw $\varepsilon_k \sim N(0,1)$
    $u_k \leftarrow \beta_k X + \sqrt{1-\beta_k^2}\,\varepsilon_k$
    $\tau_k \leftarrow t$ s.t. $\hat{\pi}^{(k)}(t) = \Phi(u_k)$

MONTE-CARLO-BDS($\beta, \hat{\pi}, M$)

1 **for** $i \leftarrow 1$ **to** M $\quad \triangleright$ M is number of trials
2   **do** $\tau = (\tau_1 \ldots \tau_K) \leftarrow$ ONE-FACTOR-GAUSSIAN-COPULA($\beta, \hat{\pi}$)
3   Generate path present value $\tilde{V}_{\text{prem}}[i]$ and $\tilde{V}_{\text{def}}[i]$
4 $\hat{V}_{\text{prem}} \leftarrow \sum_{i=1}^{M} \tilde{V}_{\text{prem}}[i]/M; \quad \hat{V}_{\text{def}} \leftarrow \sum_{i=1}^{M} \tilde{V}_{\text{def}}[i]/M$

The algorithm is applicable to both normal BDS and forward-starting BDS by specifying the payment structures in Line 3 of MONTE-CARLO-BDS(). In the following, we estimate Monte Carlo simulation error of the default leg; a similar estimate applies to the premium leg. The $\hat{V}_{\text{def}}$ is an approximation to the real quantity V_{def}, so the error of Monte Carlo simulation is $\hat{V}_{\text{def}} - V_{\text{def}}$. The law of large numbers tells us that

$$\hat{V}_{\text{def}} \xrightarrow{a.s.} V_{\text{def}}, \text{ as } M \to \infty$$

and the central limit theorem states

$$\frac{\hat{V}_{\text{def}} - V_{\text{def}}}{\sigma/\sqrt{M}} \xrightarrow{D} N(0,1), \text{ as } M \to \infty \tag{3.1}$$

where σ is the standard deviation of V_{def}. Elementary statistics tells us that the Monte Carlo standard error

$$\hat{\sigma} = \sqrt{\frac{1}{M-1} \sum_{i=1}^{M} \left(\tilde{V}_{\text{def}}[i] - \hat{V}_{\text{def}}\right)^2} \tag{3.2}$$

is an unbiased estimator for σ. Substituting $\hat{\sigma}$ for σ in (3.1), we obtain

$$\frac{\hat{V}_{\text{def}} - V_{\text{def}}}{\hat{\sigma}/\sqrt{M}} \xrightarrow{D} N(0,1), \text{ as } M \to \infty$$

or informally

$$\hat{V}_{\text{def}} \sim N(V_{\text{def}}, \hat{\sigma}^2/M)$$

The 95% confidence interval for the actual V_{def} is

$$\hat{V}_{\text{def}} - 1.96\frac{\hat{\sigma}}{\sqrt{M}} \leq V_{\text{def}} \leq \hat{V}_{\text{def}} + 1.96\frac{\hat{\sigma}}{\sqrt{M}}$$

Therefore, the Monte Carlo standard error is $O(1/\sqrt{M})$. So, for example, to reduce the error by a factor of 10, the number of trials (i.e., M) must be increased by a factor of 100. This explains why Monte Carlo simulation is so computationally expensive.

3.2 Stratified Monte Carlo simulation

While increasing the sample size is one technique for reducing the standard error of a Monte Carlo simulation, a better solution is to employ some variance reduction technique. Standard techniques of variance reduction include antithetic variates, control variates, importance sampling, and stratified sampling. For normal BDS, Glasserman and Li [4], Joshi and Kainth [7], and Chen and Glasserman [3] use importance sampling to accelerate the simulation. Here we present the stratified sampling technique, which can be applied to more general problems, including the forward-starting BDS. In our illustration of the method, we use the one-factor Gaussian copula model; the method can be easily extended to a multi-factor version as shown in [3].

Recall that the one-factor Gaussian copula model is of the form

$$u_k = \beta_k X + \sqrt{1-\beta_k^2}\varepsilon_k, \quad k = 1, \ldots, K$$

We apply stratified sampling to the common factor X, which follows the standard normal distribution. We partition the real line $[-\infty, \infty]$ into m disjoint subintervals $D_1, \ldots, D_m$ with equal probability $1/m$ by setting

$$D_i = \left[\Phi^{-1}\left(\frac{i-1}{m}\right), \Phi^{-1}\left(\frac{i}{m}\right)\right], \quad i = 1, \ldots, m$$

Note $\bigcup_{i=1}^{m} D_i = [-\infty, \infty]$. Instead of drawing X randomly over the entire real line M times, the stratified sampling method draws X in each subinterval D_i M/m times. For simplicity, we assume M is divisible by the number of strata m. In order to generate a sample of X in D_i, we first draw a unformly distributed random variable U_i over $[\frac{i-1}{m}, \frac{i}{m}]$, then we set $X = \Phi^{-1}(U_i) \in D_i$, which follows the standard normal distribution. The algorithm for BDS is stated below.

Monte-Carlo-Stratified-Sampling$(\beta, \hat{\pi}, m, M)$

for $j \leftarrow 1$ **to** M/m
    do for $i \leftarrow 1$ **to** m
        do Draw $U_i \sim U\left(\frac{i-1}{m}, \frac{i}{m}\right)$ $\triangleright$ $U(\cdot)$ follows uniform distribution
            $X_i \leftarrow \Phi^{-1}(U_i)$
            for $k \leftarrow 1$ **to** K
                do Draw $\varepsilon_k \sim N(0,1)$
                    $u_k \leftarrow \beta_k X_i + \sqrt{1-\beta_k^2}\varepsilon_k$
                    $\tau_k \leftarrow t$ s.t. $\hat{\pi}^{(k)}(t) = \Phi(u_k)$
            Generate path present value $\bar{V}_{\text{prem}}[i][j]$ and $\bar{V}_{\text{def}}[i][j]$
        $\tilde{V}_{\text{prem}}[j] \leftarrow \sum_{i=1}^{m} \bar{V}_{\text{prem}}[i][j]/m$; $\tilde{V}_{\text{def}}[j] \leftarrow \sum_{i=1}^{m} \bar{V}_{\text{def}}[i][j]/m$
$\hat{V}_{\text{prem}} \leftarrow \sum_{j=1}^{M/m} \tilde{V}_{\text{prem}}[j]/(M/m)$; $\hat{V}_{\text{def}} \leftarrow \sum_{j=1}^{M/m} \tilde{V}_{\text{def}}[j]/(M/m)$

The stratified sampling method described above draws an equal number of samples from equiprobable strata (known as proportional allocation). From Madras [12], we know that the estimators are unbiased and the variance of the stratified sampling is smaller than that of the naive Monte Carlo method.

Chapter 4

Analytical Solution

Recall that, conditional on a specified common factor x, the pricing equations satisfy

$$\mathbb{E}_x[V_{\text{def}}(T)] = \sum_{i=1}^{n} D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) \sum_{k\in\mathcal{B}_T} g(N^{(k)})\Pi_{i|T}^{(k)}(x) \tag{4.1}$$

$$\mathbb{E}_x[V_{\text{prem}}(T)] = s\sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) N_T(x)\bar{\Pi}_{i|T}(x) \tag{4.2}$$

In order to evaluate the right sides of the equations above, we need to compute $\mathbb{P}_x(\mathcal{B}_T)$, the conditional probability distribution of $\mathcal{B}_T$'s composition; $N_T(x)$, the total notional values of $\mathcal{B}_T$; and $\Pi_{i|T}^{(k)}(x)$, the conditional forward probabilities that name k's default triggers the terminal default payment.

We apply the recursive method for normal BDS proposed by Iscoe and Kreinin [6] to compute $\Pi_{i|T}^{(k)}(x)$. For $\mathbb{P}_x(\mathcal{B}_T)$ and $N_T(x)$, we consider completely homogeneous, homogeneous and heterogeneous cases separately. Furthermore, we complete the evaluation by considering them separately.

4.1 Computing terminal default probabilities

Suppose the original basket $\mathcal{B}$ contains names $1,\ldots,K$. Conditional on a specified scenario x, the basket may experience different numbers of defaults in $(0,T]$. Assume names

$1', \ldots, K'$ survive till T, so $\mathcal{B}_T$ is composed of these names. We need to explore only the cases for which $m \leq K' \leq K$, as the contract terminates at T without any payments if $K' < m$.

In a forward-starting first-to-default contract, the conditional probabilities $\Pi_{i|T}^{(k)}(x) = \mathbb{P}_x(\tau = \tau^{(k)}, \tau \in (T_{i-1}, T_i] \mid \mathcal{B}_T)$ satisfy

$$\Pi_{i|T}^{(k)}(x) = \frac{\lambda_{i|T}^{(k)}(x)}{\lambda_{i|T}(x)} \big(\bar{\Pi}_{i-1|T}(x) - \bar{\Pi}_{i|T}(x)\big), \quad i = 1, \ldots, n;\ k = 1', \ldots, K' \tag{4.3}$$

where $\lambda_{i|T}^{(k)}(x)$ is the conditional forward default intensities defined in (2.19); $\lambda_{i|T}(x) = \sum_{k=1'}^{K'} \lambda_{i|T}^{(k)}(x)$; and $\bar{\Pi}_{i|T}(x) = \prod_{k=1'}^{K'} \big(1 - \hat{\pi}_{i|T}^{(k)}(x)\big)$. The proof of (4.3) for normal BDS is given by Iscoe and Kreinin [6]. The proof is also valid for forward-starting BDS. Here we explain it informally. The term $\bar{\Pi}_{i|T}(x)$ in (4.3) is the conditional probability that no name defaults in $(T, T_i]$, therefore, $\bar{\Pi}_{i-1|T}(x) - \bar{\Pi}_{i|T}(x)$ is the conditional probability that at least one name defaults in $(T_{i-1}, T_i]$, which is also the probability that the terminal default (first-to-default) occurs in $(T_{i-1}, T_i]$. The term $\Pi_{i|T}^{(k)}(x)$ is name k's contribution to the probability of the terminal default occurring in $(T_{i-1}, T_i]$. Since we assume the forward default intensity $\lambda_{i|T}^{(k)}(x)$ is piecewise constant, the probability that name k's default occurs in $(T_{i-1}, T_i]$ follows an exponential distribution. Thus, the first-to-default event in $(T_{i-1}, T_i]$ can be treated as an arrival problem for the names in the pool, and, from a general result for the arrival problem, the probability that name k defaults first out of the pool is $\lambda_{i|T}^{(k)}(x)/\lambda_{i|T}(x)$. Thus, (4.3) is true.

For the mth-to-default BDS, Iscoe and Kreinin [6] derive the recursive relation between the mth-to-default and the $(m-1)$st-to-default contract:

$$(m-1)\mathcal{P}_m(\mathcal{B}_T) = \sum_{j \neq k} \mathcal{P}_{m-1}(B_T^{[j]}) - (K' - m + 1)\mathcal{P}_{m-1}(\mathcal{B}_T)$$

where $\mathcal{P}_m(\mathcal{B}_T) = \mathbb{P}(\tau = \tau^{(k)}, \tau \in (T_{i-1}, T_i] \mid \mathcal{B}_T)$ for the mth-to-default BDS; and $\mathcal{B}_T^{[j]}$ is the set of names obtained by excluding name j from $\mathcal{B}_T$. Naive implementation of this recursion causes the recalculation of the same probabilities, thus reducing the

mth-to-default problem to the first-to-default problem as

$$\mathcal{P}_m(\mathcal{B}_T) = \sum_{v=0}^{m-1}(-1)^{m-v-1}\binom{K'-v-1}{m-v-1}\sum_{\mathcal{J}\subset\mathcal{B}_T:|\mathcal{J}|=v}\mathcal{P}_1(\mathcal{B}_T^{[\mathcal{J}]})$$

where $\mathcal{J}$ is a subset of $\mathcal{B}_T$ and $\mathcal{B}_T^{[\mathcal{J}]} = \mathcal{B}_T \setminus \mathcal{J}$. Here, for simplicity, we demonstrate the recursion for the unconditional probabilities, but it is also valid for the conditional probabilities.

4.2 Completely homogeneous case

A completely homogeneous pool has identical recovery-adjusted notional values, identical risk-neutral default probabilities and identical default correlations. Hence, given a scenario x, the number of defaults in $(0,T]$ follows a binomial distribution with probability $\hat{\pi}^{(1)}(T,x)$, which is defined in (2.11). With identical recovery-adjusted notional values, all equal to $N^{(1)}$, the pool losses till T are $L_T = vN^{(1)}$, where v is the number of defaults in $(0,T]$. Therefore, the outstanding notional at T is $N_T(x) = N - L_T = (K-v)N^{(1)}$, where $0 \le v \le K - m$. The conditional distribution of the basket composition at T satisfies

$$\mathbb{P}_x(|\mathcal{B}_T| = K - v) = \mathbb{P}_x(L_T = vN^{(1)}) = \mathrm{Bin}(v; K, \hat{\pi}^{(1)}(T,x)) \tag{4.4}$$

We present two methods to evaluate the pricing equations (4.1) and (4.2) for a completely homogeneous pool: one approach is to compute the conditional expectations exactly; the other is to evaluate them approximately.

4.2.1 Exact method

For the exact method, we explore the payout in all possible cases and compute the associated expectations. From (4.4) and $N^{(k)} = N^{(1)}$, the conditional expectation of the default leg at T satisfies

$$\begin{aligned}
\mathbb{E}_x[V_{\text{def}}(T)] &= \sum_{i=1}^{n} D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) \sum_{k\in\mathcal{B}_T} g(N^{(k)})\Pi_{i|T}^{(k)}(x) \\
&= g(N^{(1)}) \sum_{i=1}^{n} D(T,T_i) \sum_{v=0}^{K-m} \mathbb{P}_x(L_T = vN^{(1)}) \sum_{k\in\mathcal{B}_T} \Pi_{i|T}^{(k)}(x) \\
&= g(N^{(1)}) \sum_{i=1}^{n} D(T,T_i) \sum_{v=0}^{K-m} \text{Bin}(v; K, \hat{\pi}^{(1)}(T,x))(K-v)\Pi_{i|T}^{(1)}(x) \qquad (4.5)
\end{aligned}$$

In deriving the last equality of (4.5), we use the fact that all the names in $\mathcal{B}_T$ are identical, so their terminal default probabilities are equal, i.e., $\Pi_{i|T}^{(k)}(x) = \Pi_{i|T}^{(1)}(x)$. Similarly, the conditional expectation of the premium leg at T satisfies

$$\begin{aligned}
\mathbb{E}_x[V_{\text{prem}}(T)] &= s \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) N_T(x) \bar{\Pi}_{i|T}(x) \\
&= s \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{v=0}^{K-m} \mathbb{P}_x(L_T = vN^{(1)})(K-v)N^{(1)}\bar{\Pi}_{i|T}(x) \\
&= sN^{(1)} \sum_{i=1}^{n} \Delta T_i \cdot D(T,T_i) \sum_{v=0}^{K-m} \text{Bin}(v; K, \hat{\pi}^{(1)}(T,x))(K-v)\bar{\Pi}_{i|T}(x) \qquad (4.6)
\end{aligned}$$

Conditional on a scenario x, we need to evaluate $n(K-m+1)$ terms in the sums (4.5) and (4.6). If we assume the integration over X can be computed numerically using an effective quadrature rule by evaluating M different x over the real line, the running time of the whole method for the valuation of a forward-starting BDS for a completely homogeneous pool is $O(M \cdot n \cdot (K-m))$.

4.2.2 Approximation method

From the previous subsection, we know $\mathbb{E}_x[V_{\text{prem}}(T)]$ can be computed as the summation of functions $h(v)$:

$$\mathbb{E}_x[V_{\text{prem}}(T)] = \sum_{v=0}^{K-m} h(v)$$

where

$$h(v) = sN^{(1)} \sum_{i=1}^{n} \Delta T_i \, D(T,T_i) \, \text{Bin}(v; K, \hat{\pi}^{(1)}(T,x)) \, (K-v) \, \bar{\Pi}_{i|T}(x)$$

We employ an effective quadrature rule, such as the M-point Gauss-Legendre or Hermite rule, to approximate $\mathbb{E}[V_{\text{prem}}(T)]$:

$$\begin{aligned}
\mathbb{E}[V_{\text{prem}}(T)] &= \int_{-\infty}^{\infty} \mathbb{E}_x[V_{\text{prem}}(T)]d\Phi(x) \\
&= \int_{-\infty}^{\infty} \mathbb{E}_x[V_{\text{prem}}(T)]\varphi(x)dx \\
&\approx \sum_{j=1}^{M} w_j \Big(\mathbb{E}_{x_j}[V_{\text{prem}}(T)]\varphi(x_j)\Big) \\
&= \sum_{j=1}^{M} w_j \sum_{v=0}^{K-m} f(x_j, v)
\end{aligned}$$

where $\varphi(\cdot)$ is the probability density function of the standard normal distribution; w_j is the weight factor; and x_j is the associated node for the M-point quadrature rule. Consider the second line of the equation above. The density function $\varphi(x)$ is very small outside of the range of $[-5, 5]$, whereas $\mathbb{E}_x[V_{\text{prem}}(T)]$ is not unduly large outside of $[-5, 5]$. Therefore, the integral is generally approximated by a quadrature rule over $[-5, 5]$. A similar quadrature rule is also applied to $\mathbb{E}[V_{\text{def}}(T)]$ over the truncated range $[-5, 5]$.

Figure 4.1 shows representative shapes of $\mathbb{E}[V_{\text{def}}(T)]$ and $\mathbb{E}[V_{\text{prem}}(T)]$ as functions of the common factor X and the number of defaults v. The plot is based on a forward-starting second-to-default BDS with 10 names in the original pool. In order to compute $\mathbb{E}[V_{\text{prem}}(T)]$ and $\mathbb{E}[V_{\text{def}}(T)]$, we need to sum up all the heights over the coordinate.

Figure 4.1 demonstrates that most of the points (x, v) contribute little to the final value of $\mathbb{E}[V_{\text{prem}}(T)]$, so we can omit these nodes or approximate them to reduce the computational workload. If the unimportant points are dropped, the computed values for both the default leg and premium leg are less than the actual values; if they are approximated, the computed values are almost the same as the actual ones. In either case, the error in the computed values for the default leg and premium leg are likely correlated. Since the value of BDS is the difference of these two legs, the error likely cancels to some extend.

To implement the approximation method, we need a heuristic to decide which points

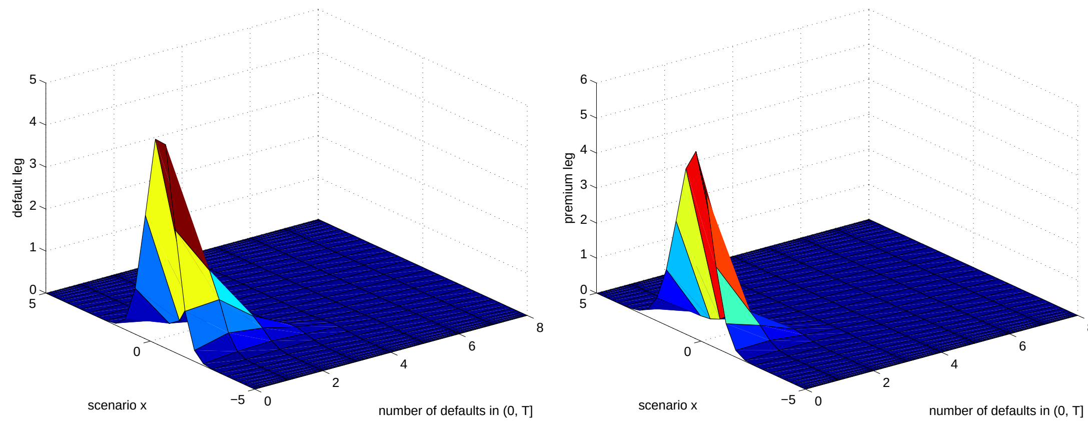

Figure 4.1: Plot of $\mathbb{E}[V_{\text{prem}}(T)]$ and $\mathbb{E}[V_{\text{def}}(T)]$ as a function of X and v

to drop or approximate. To this end, we choose a tolerance, which can be either absolute or relative. If the object value at the current point is below the tolerance, we omit or interpolate it. With this heuristic, we need an algorithm to determine these ignorable points cheaply. From the pricing equations (4.5) and (4.6), we find that the values of $\mathbb{E}_x[V_{\text{prem}}(T)]$ and $\mathbb{E}_x[V_{\text{def}}(T)]$ are determined mostly by $\mathbb{P}_x(\mathcal{B}_T)$ or equivalently the binomial distribution $\text{Bin}(v; K, \hat{\pi}^{(1)}(T,x))$. Figure 4.2 plots the value of the default leg versus the number of defaults under different scenario values of X on the left side and the relevant binomial distribution of the number of defaults in $(0,T]$ versus the number of defaults under the same scenario values of X on the right side. The figure clearly demonstrates that the binomial distribution determines the shape of $\mathbb{E}_x[V_{\text{def}}(T)]$. This observation underpins the following algorithm. A similar relationship holds for the premium leg.

The peak probability of the binomial distribution occurs around its expected value $K\hat{\pi}^{(1)}(T,x)$. Therefore, when evaluating the object function, we compute its value at the point $(x, \lfloor K\hat{\pi}^{(1)}(T,x)\rfloor)$ first. We continue the evaluation process on either side of that point until the object function value is below the tolerance or the evaluation reaches the boundary of the possible number of defaults. The algorithm POINT-EVALUATION() on page 25 gives the evaluation process under a specified scenario x. A similar algorithm

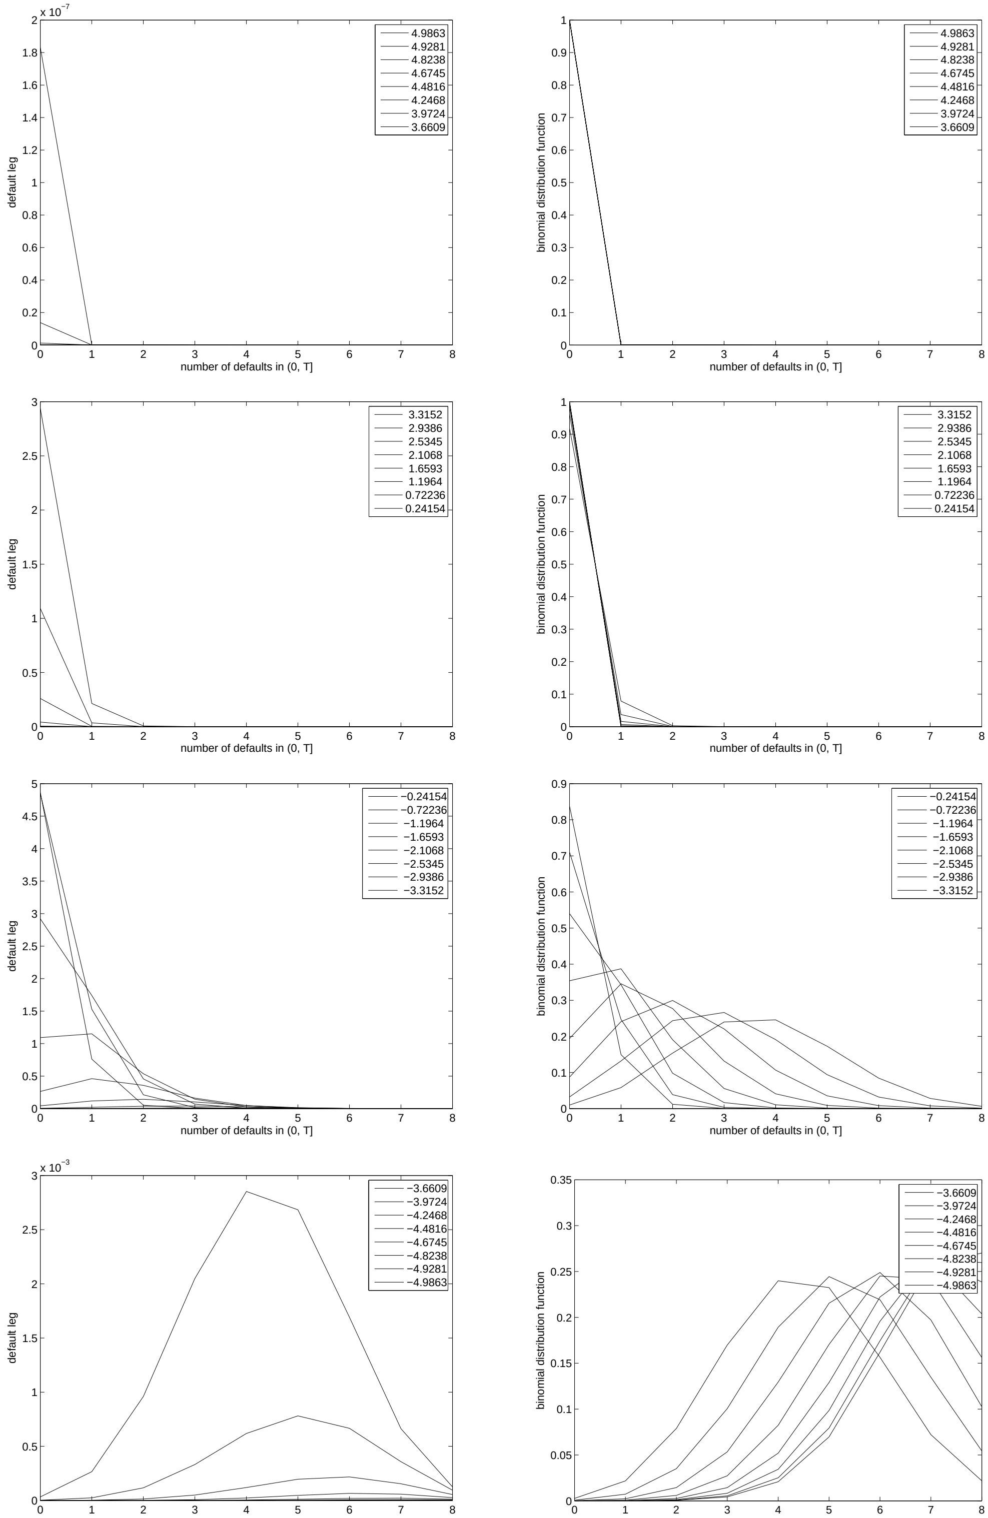

Figure 4.2: The relation between the default leg value and the binomial distribution

POINT-EVALUATION($m, K, \hat{\pi}(x), tol$)

```
Init ← ⌊K π̂^(1)(T, x)⌋            ▷ Initial evaluation point
Evaluate the value of object function V(x, Init)
if V(x, Init) < tol
   then V(x) ← 0
   else V(x) ← V(x, Init)
        Left ← Init − 1                       ▷ Evaluate the left hand side
        while Left ≥ 0
            do Evaluate the value of object function V(x, Left)
               if V(x, Left) < tol
                  then Break
                  else V(x) ← V(x) + V(x, Left)
        Right ← Init + 1                      ▷ Evaluate the right hand side
        while Right ≤ K − m
            do Evaluate the value of object function V(x, Right)
               if V(x, Right) < tol
                  then Break
                  else V(x) ← V(x) + V(x, Right)
```

for the interpolation method can be implemented by modifying lines 4, 10 and 16.

4.3 Homogeneous and inhomogeneous cases

Unlike a completely homogeneous pool, a homogenous or inhomogeneous pool has different default probabilities and correlations. Therefore, the default leg and the premium leg depend on the composition of the basket at T, not just the number of names in the basket. Therefore, we need to explore not only the number of defaults in $(0, T]$, but also, for each number of defaults v, the different subcases due to the actual names in $\mathcal{B}_T$. Thus, the total number of subcases is proportional to 2^K. The evaluation process inside of each subcase is the same as that described above for the completely homogeneous case.

4.3.1 Exact method

Based on the assumption that names $1', \ldots, K'$ survive till T, we have

$$\mathbb{P}_x(\mathcal{B}_T) = \prod_{k \in \mathcal{B} \backslash \mathcal{B}_T} \hat{\pi}^{(k)}(T, x) \prod_{k \in \mathcal{B}_T} (1 - \hat{\pi}^{(k)}(T, x))$$

and

$$N_T(x) = \sum_{k \in \mathcal{B}_T} N^{(k)}$$

Therefore, conditional on a scenario x, the expectation of the default leg and premium leg at T are

$$\mathbb{E}_x[V_{\text{def}}(T)] = \sum_{i=1}^{n} D(T, T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) \sum_{k \in \mathcal{B}_T} g(N^{(k)}) \Pi_{i|T}^{(k)}(x)$$

$$\mathbb{E}_x[V_{\text{prem}}(T)] = s \sum_{i=1}^{n} \Delta T_i \cdot D(T, T_i) \sum_{|\mathcal{B}_T|=m}^{K} \mathbb{P}_x(\mathcal{B}_T) N_T(x) \bar{\Pi}_{i|T}(x)$$

These expressions can be computed by brute force. If we explore M different scenarios of X in the quadrature, the running time of the whole algorithm is proportional to $M \cdot n \cdot 2^K$.

4.3.2 Approximation method

Similar to the completely homogenous case, we omit or approximate the unimportant points. Due to the different marginal default probabilities and correlations associated with a homogeneous or inhomogeneous pool, the conditional default probabilities $\hat{\pi}^{(k)}(T,x)$ are different. Therefore the point evaluation algorithm POINT-EVALUATION() for the completely homogeneous pool is not valid here. However, the idea underpinning the algorithm for the completely homogeneous case is still useful, i.e., whether a point should be omitted is determined by whether it is below the tolerance.

To extend the algorithm for the completely homogeneous case to the homogeneous or inhomogeneous case, we change the initial evaluation point and consider the subcases inside of each point. If we approximate the distribution of the number of defaults in $(0,T]$ by a Poisson distribution with $\hat{\lambda} = \sum_{k=1}^{K} \hat{\pi}^{(k)}(T,x)$, the relationship between the object function and the distribution of the number of defaults are similar to that shown in Figure 4.2. The peak probability of the Poisson distribution occurs around $\lfloor \hat{\lambda} T \rfloor$. Therefore, we start the evaluation at point $(x, \lfloor \hat{\lambda} T \rfloor)$ and explore its subcases, then extend the process to both sides, exploring all subcases for each point, until the terminal condition is satisfied.

Chapter 5

Numerical Results

In this chapter, we compare the naive Monte Carlo method and Monte Carlo with stratified sampling for pricing a forward-starting BDS associated with a completely homogeneous pool. For the same contract, we exam the accuracy and performance of the analytical method described in Chapter 4. The effect of BDS parameters is also studied for the same contract. The approximation method for a forward-starting inhomogeneous BDS is also tested.

5.1 Comparison of the Monte Carlo methods

The first basket is a completely homogeneous pool containing 10 names. The parameters of each name are shown in Table 5.1. The contract is a 5-year BDS starting one year later, so $T = 1$, $T^* = 6$, and $T_i = i + 1$, for $i = 0, \dots, 5$. The recovery rate of each name is 15%. The risk-neutral cumulative default probabilities are listed in Table 5.2. The

Notional	100
Credit Rating	C4
Correlation	0.5

Table 5.1: Parameters of the BDS for a completely homogeneous pool

Credit rating	Time					
	1Y	2Y	3Y	4Y	5Y	6Y
C1	0.0041	0.0052	0.0069	0.0217	0.0288	0.0323
C2	0.0071	0.0185	0.0328	0.0495	0.0682	0.0801
C3	0.0072	0.0225	0.0439	0.0692	0.0967	0.1262
C4	0.0258	0.0575	0.0930	0.1304	0.1683	0.1852
C5	0.0305	0.0616	0.0936	0.1464	0.1702	0.1945
C6	0.0420	0.0713	0.0953	0.1661	0.1950	0.2210
C7	0.0501	0.0802	0.1062	0.2171	0.3030	0.3911
C8	0.0571	0.0872	0.1132	0.2241	0.3100	0.4032

Table 5.2: Risk-neutral cumulative default probabilities

continuously compounded interest rates are listed in Table 5.3.

time	1Y	2Y	3Y	4Y	5Y	6Y
Rate	0.046	0.050	0.056	0.058	0.06	0.061

Table 5.3: Risk-free interest rates

Figure 5.1 illustrates the standard error comparison between the naive Monte Carlo method and the stratified Monte Carlo simulation, where the standard error is computed by (3.2). We use the total number of 10^6 replications in the naive method and 10 strata in the stratified sampling method with 10^5 replications within each stratum. We compute the standard errors for forward-starting mth-to-default contracts (m $= 1, \ldots, 4$) and plot them against m. The left panel is the standard error of the default leg, while the right panel is the standard error of the premium coefficient, which is the value of the premium leg with premium $s = 1$. From the description of the algorithms in Chapter 3, we can see that the complexity of each method is proportional to MK, so the running time of each method is almost the same. However, as shown in Figure 5.1, the stratified method achieves a smaller standard error.

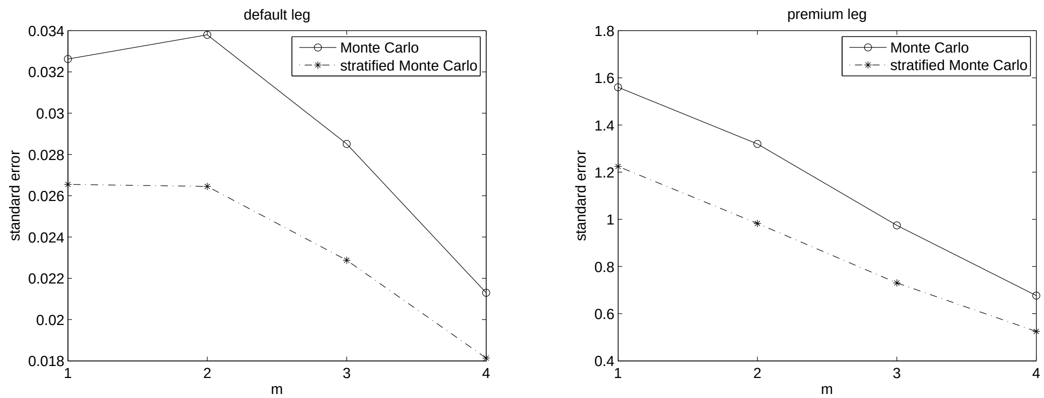

Figure 5.1: Comparison of naive and stratified Monte Carlo methods

5.2 Accuracy of the exact solution

We compare the accuracy of the analytical solution with stratified Monte Carlo simulation using the same contract as above. We compute the premium for the forward-starting mth-to-default BDS ($m = 1, \ldots, 4$) and show the 95% confidence interval of the BDS premium. The 95% confidence interval is computed as follows: we repeat each of the Monte Carlo experiments 500 times; then, we compute the 95% confidence interval from the empirical distribution of those 500 samples. Table 5.4 presents the results for different numbers of scenarios for the stratified Monte Carlo method and the analytical technique. From the table, we can conclude that the premium computed by the analytic method is reliable; in the stratified Monte Carlo simulation, if the number of scenarios is less then 10^6, the relative error of the premium estimation may be greater than 0.5%.

5.3 Effect of parameters

Figure 5.2 shows the error and CPU time versus M for an M-point Gauss-Legendre quadrature rule. The reference solution is computed by a 100-point Gauss-Legendre quadrature rule. The left panel plots the maximum relative error of an mth-to-default contract ($m = 1, \ldots, K$) against M; the right panel plots the CPU time for evaluating

No. of Scenarios	$m = 1$	$m = 2$	$m = 3$	$m = 4$
1×10^5	[256.43, 261.70]	[100.74, 103.17]	[46.74, 48.39]	[21.26, 22.17]
1×10^6	[258.13, 259.74]	[101.60, 102.28]	[47.23, 47.62]	[21.54, 21.84]
2×10^6	[258.36, 259.53]	[101.69, 102.18]	[47.31, 47.60]	[21.59, 21.81]
Analytic	258.97	101.92	47.45	21.70

Table 5.4: Risk premium (bps) for a forward-starting BDS computed by the stratified Monte Carlo (first three rows) and the analytic method (last row)

m from 1 to K against M. As expected, the CPU time is roughly linearly increasing with M and the error is roughly exponentially decreasing with M. When $M > 32$, no more accuracy is obtained as M increases, probably because round-off dominates the truncation error from this point on.

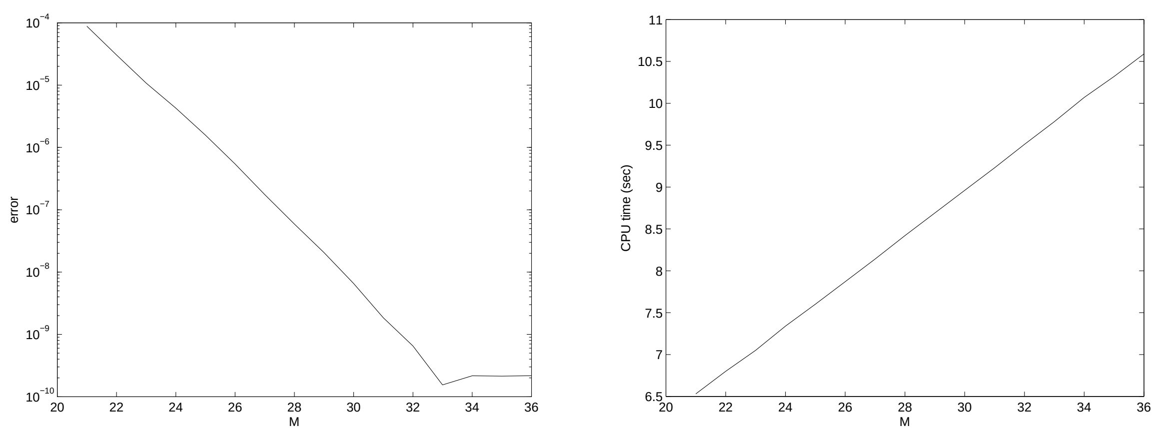

Figure 5.2: Error and CPU time versus M for an M-point Gauss-Legendre quadrature rule

Figure 5.3 plots the premium of a forward-starting mth-to-default contract as a function of m. As the correlation of the names is 0.5, the premium is a decreasing function of m.

In Figure 5.4, we plot the credit risk premium versus the correlation. The correlation factor ranges from 0 (totally uncorrelated) to 0.9 (highly correlated).

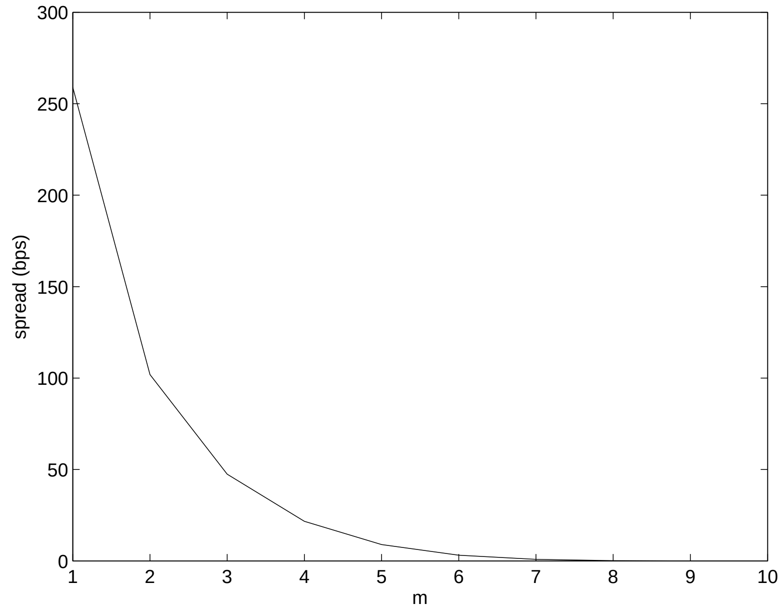

Figure 5.3: Premium of the forward-starting mth-to-default contract versus m

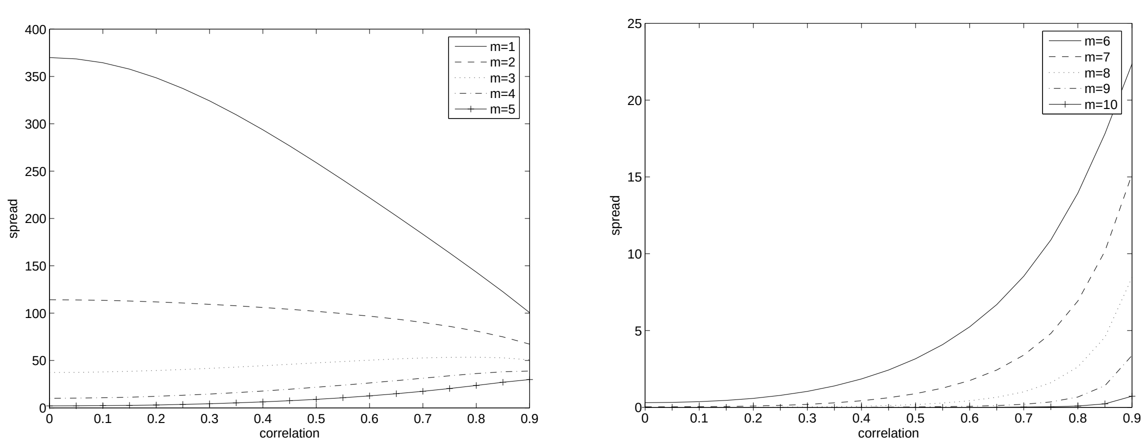

Figure 5.4: Premium of the mth-to-default contract versus correlation

5.4 Accuracy and performance of the approximation method

We test our approximation method by computing the risk premium for forward-starting mth-to-default contracts ($m = 1, \ldots, 4$). The contract is a 5 year BDS which starts at the end of the first year. The basket is a heterogeneous pool containing 10 names with the parameters shown in Table 5.5. The risk-neutral default probabilities and the interest rates are listed in Tables 5.2 and 5.3, respectively. The recovery rate is 15%.

Name	Notional	Credit Rating	Correlation
1	190	C7	0.5
2	80	C2	0.6
3	70	C4	0.9
4	360	C6	0.6
5	100	C5	0.5
6	200	C6	0.4
7	150	C6	0.7
8	123	C5	0.64
9	95	C6	0.55
10	107	C3	0.22

Table 5.5: Parameters of the BDS for a heterogeneous pool

We compute the risk premium for the contract by the exact solution and the approximation method. The reference solution is computed by the exact method with a 32-point Gauss-Legendre quadrature rule. Table 5.6 reports the reference solution and the absolute error for the approximation method with tolerance $= 10^{-4}$.

Figure 5.5 compares the performance of the exact method, the approximation method, and the stratified Monte Carlo method. We use the heterogeneous pool listed in Table

Method	$m = 1$	$m = 2$	$m = 3$	$m = 4$
Exact	274.28	114.35	57.10	28.58
Approximation	4.70×10^{-4}	5.87×10^{-5}	1.01×10^{-5}	5.95×10^{-6}

Table 5.6: Accuracy of the approximation method

5.5 and compute the risk premium for the forward-starting mth-to-default BDS ($m = 1, \dots, K$). The exact solution and the approximation method are described above; the Stratified Monte Carlo method uses 10 strata and 10^5 trials in each stratum. The running time of the exact solution T_{EX} is set as the benchmark; for each of the other two methods, Perf is the ratio of that method's running time to T_{EX}.

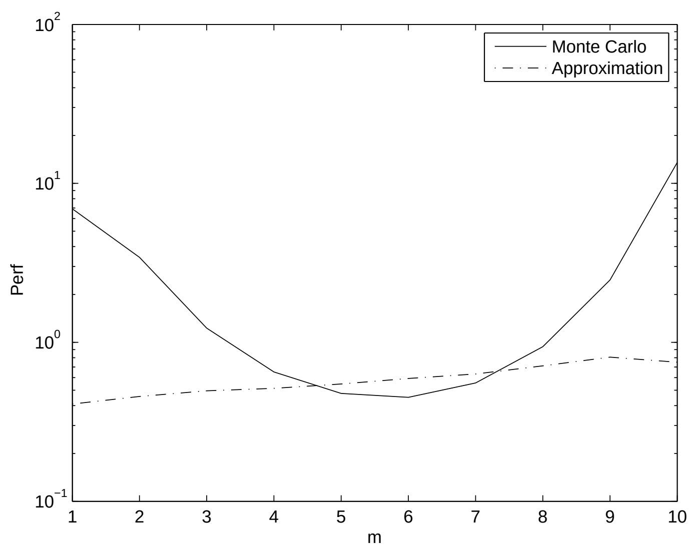

Figure 5.5: Performance comparison

From Table 5.6 and Figure 5.5, we can conclude that the approximation method obtains sufficient accuracy for many applications in about half of the running time used by the exact solution. Thus it is an effective alternative to the exact method. The Monte Carlo method is more efficient than the exact method when m is between about 4 and 7 in this example. However, the complexity of the exact method and the approximation method increases exponentially with the size of the basket, while the complexity of the

Monte Carlo method increases linearly with the size of basket. Therefore, when the pool is large, Monte Carlo simulation becomes relatively more effective.

Chapter 6

Conclusions

In this paper, we study valuation methods for forward-starting BDS in the conditional independence framework. For normal BDS, we transfer the correlated marginal default probabilities to the conditionally independent ones by the factor copula model; for forward-starting BDS, we use a similar mapping process and then convert the conditional default probabilities to the conditional forward ones, which are the default probabilities conditional on the name surviving till the maturity of the forward contract.

We review the flexible but inefficient Monte Carlo simulation and analyze its efficiency. To reduce the computational expensive, we propose the stratified sampling technique. Numerical results verify that the stratified Monte Carlo method performs better than the naive one.

Based on the recursive method proposed by Iscoe and Kreinin [6] for normal BDS, we develop an analytical solution for forward-starting BDS. Conditional on a specified common factor x, we first explore the possible combination of defaults during the life of the forward contract. For each scenario, we employ the recursive method to value the default leg and the premium leg. The risk premium or the value of forward contract is computed from the expectation of these two legs. To evaluate the expectations, we first analysis the exact method for a completely homogeneous pool. Due to the combinatorial nature of the problem, the exact method is computationally expensive for homogeneous and inhomogeneous pools. We develop an approximation method, which omits or interpolates unimportant values, to accelerate the computation.

Our numerical results compare the accuracy and computational time of these methods. The stratified Monte Carlo method is more efficient than the naive one. The analytical solution based on the exact method is reliable. Compared with the exact method, the approximation method is about twice as fast without losing much accuracy, thus it is a more effective method. The influence of different contract parameters is also studied.

In the future, we hope to explore other variance reduction techniques for Monte Carlo simulation, more efficient approximation algorithms, other approaches for the analysis and valuation of BDS options, and extension of our work to stochastic recovery rate and interest rate processes.

Bibliography

[1] L. Andersen, J. Sidenius, and S. Basu. All your hedges in one basket. *Risk*, 16(11):67–72, 2003.

[2] N. Bennani. The forward loss model: A dynamic term structure approach for the pricing of portfolio credit derivatives. Working paper, 2005.

[3] Z. Chen and P. Glasserman. Fast pricing of basket default swaps. Working paper, March 2006.

[4] P. Glasserman and J. Li. Importantce sampling for portfolio credit risk. *Management Science*, 51(11):1642–1656, November 2005.

[5] J. Hull and A. White. Valuation of a CDO and an n^{th} to default CDS without Monte Carlo simulation. *Journal of Derivatives*, 12(2):8–23, Winter 2004.

[6] I. Iscoe and A. Kreinin. Recursive valuation of basket default swaps. *Journal of Computational Finance*, 9, 2006.

[7] M. Joshi and D. Kainth. Rapid and accurate development of prices and Greeks for nth to default credit swaps in the Li model. *Quantitative Finance*, 4:266–275, 2004.

[8] D. Lando. *Credit Risk Modeling: Theory and Applications.* Princeton University Press, 2004.

[9] J. Laurent and J. Gregory. Basket default swaps, CDO's and factor copulas. *Journal of Risk*, 7:103–122, 2005.

[10] D. Li. Constructing a credit curve. *Credit Risk*, pages 40–44, 1998.

[11] D. Li. On default correlation: A copula approach. *Journal of Fixed Income*, 9:43–54, 2000.

[12] N. Madras. *Lectures on Monte Carlo Methods.* American Mathematical Society, 2001.

[13] P. Schönbucher. Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives. Working paper, September 2005.

[14] J. Sidenius, V. Piterbarg, and L. Andersen. A new framework for dynamic credit portfolio loss modeling. Working paper, 2005.