An Efficient Algorithm for Basket Default Swap Valuation

Mi-Hsiu Chiang, Meng-Lan Yueh, and Ming-Hua Hsieh

Mi-Hsiu Chiang
is an assistant professor in the department of money and banking at National Chengchi University in Taiwan.
mhchiang@nccu.edu.tw

Meng-Lan Yueh
is an assistant professor in the department of finance at National Central University in Taiwan.
mlyueh@cc.ncu.edu.tw

Ming-Hua Hsieh
is an associate professor in the department of management information systems at National Chengchi University in Taiwan.
mhsieh@mis.nccu.edu.tw

Importance sampling (IS), as an efficiency improvement technique for Monte Carlo simulations, is particularly well-suited for correlation products which have payoffs contingent on the occurrence of rare events. An appropriate choice of the IS distribution is of absolute necessity to ensure variance reduction. Without being chosen carefully, the use of IS can result in ineffective variance reduction, or at the worse cases, even increase the variance of the Monte Carlo estimators. Hence, it is highly desirable to be able to select the right IS distribution that guarantees variance reduction. In this paper, we propose an effective IS algorithm for the valuation of k-th to default basket default swaps. The algorithm is simple to implement and guarantees variance reduction. We have established a way of ensuring that for every simulation path generated, the desired default events always take place, and this is achieved under an appropriate choice of the IS distribution among the set of possible measures that can force the required events of interest to occur. Our numerical results show that the proposed algorithm is considerably more efficient than the crude simulation method. The gain in variance-reduction efficiency is particularly prominent when the possibility of the k-th default event is remote.

Multi-name credit derivatives, such as basket default swaps, and synthetic collateralized debt obligations, have gained increasing popularity in recent years. For investors and market practitioners seeking to redistribute the credit risk of their underlying portfolios or to release regulatory capital, multi-name credit derivatives provide a feasible way to customize credit protection according to different risk profiles. Recent developments in the credit derivatives market have shifted towards the use of standardized baskets, which are based on diversified sectors of the European iTraxx index, to bring about liquidity and greater price transparency. In particular, the most leveraged trades, standardized first-to-default baskets, are now used by traders as reference benchmarks for their correlation levels.

In contrast to a simple credit default swap which provides credit protection for a single underlying, a basket credit default swap extends the credit protection to a pool of credit instruments under the restriction that compensation is granted for the default of only one underlying obligor. Depending on the ranking of default protection, a basket credit default swap is known as a first-to-default basket, a second-to-default basket, or more generally, a kth-to-default basket. A kth-to-default basket resembles an insurance contract that offers protection against the event of a kth default on a pool of n ($n \geq k$) underlying obligors.

As for credit default swap contracts, the protection buyer pays a premium/spread on the notional of the position at regular time intervals as an insurance premium until the maturity date of the contract or the event of a *k*th default. In return, the buyer of a *k*th-to-default swap contract receives the value of the loss (the notional less the recovery) at the time of a *k*th default. When the designated entity experiences a default event, the credit default swaps on the other reference entities are automatically cancelled at no cost to the protection seller. Such product facilitates the creation of specific exposures to designated default events. For example, a second-to-default basket allows the protection seller to assume the risk on the second reference entity to default within the basket, thereby reducing the seller's risk of loss, while the protection buyer benefits from a lower premium payment as the order of the designated default events progresses.

The rationale behind *k*th-to-default swaps is the alignment of credit risk with the investor's risk tolerance level by taking sequential default risk positions for a basket of reference entities. The resulting structures give rise to different exposures toward default risk, default correlation, and expected loss. Of these exposures, the sensitivity to default correlation, which depends on the order of the designated default events, is the most complex. For example, in a first-to-default basket swap, an increase in default correlation implies a strong possibility that all reference entities default simultaneously. The credit risk that the structure is exposed to, therefore, approaches that of the most risky entity in the reference pool. Thus, the risk of the structure will increase as default correlation decreases due to the protection seller being exposed to all the reference entities.

The characteristic of varying exposures to default correlation in any *k*th-to-default baskets allows traders to trade implied and expected default correlations. By taking a long position in a basket of reference entities and at the same time buying a first-to-default protection on the basket, investors can create synthetic senior positions. Synthetic senior short positions can conversely be constructed by selling a first-to-default protection on the basket and at the same time buying protection on the reference entities. When used as hedging instruments, *k*th -to-default credit default swaps can hedge against risk clustering or exposure to a particular name within a large credit portfolio. The resulting hedging costs in this case will be lower than the sum of hedging individual credits within the portfolio.

The valuation of basket credit default swaps requires specification of the distribution of joint default times; and a pre-specified dependence structure of correlated default times is a vital input. In Li [2000], default times between underlying obligors are assumed to be inhomogeneous Poisson processes and the dependence structure is a Gaussian copula function. The Li [2000] model, commonly known as the Gaussian copula approach, requires only a one-step simulation to generate correlated default times. Even though many have suggested that other copula functions, such as the Student-t copulae (Mashal and Naldi [2002]) and the Clayton copulae (Rogge and Schönbucher [2003]), may better capture the tail dependency of correlated defaults, the Gaussian copula approach remains an industry standard for valuing portfolio credit derivatives because of its ease of implementation via Monte Carlo simulations.

Although conceptually simple and easy to implement, Monte Carlo simulations can be unstable and exhibit slow convergence when dealing with rare events. This is mainly due to the integrand being locally peaked. In the case of pricing basket default swaps, it is only when a *kth* default has taken place before the maturity date that a simulated path can result in a default payoff. Thus, of all paths generated, even with a high hazard rate, the percentage of paths that actually result in default over a contract's lifetime can be very small and their relevant contributions to the value of the default leg are uninformative and insignificant. The convergence rate decreases as the number of underlying obligors increases, especially for contracts of short maturities.

When a reference entity is highly creditworthy, default is a rare event that makes a crude Monte Carlo algorithm inefficient in estimating the point of interest. Importance sampling (IS) is an appealing technique for improving the efficiency of rare-event Monte Carlo simulations. But if the IS distribution is not properly selected, the use of IS can result in ineffective variance reduction and even increase the variance of the estimator under certain circumstances. The IS procedure introduced by Joshi and Kainth [2004] for the valuation of BDS contracts is one case where the variance of an estimator can be increased. When the default times are generated in sequential order and the choice of the likelihood ratio is not subject to any pre-specified minimization constraint for variance reduction, the resulting probability distribution can fail to be an "equivalent measure." Thus, events that take place under the original measure may not be assigned a positive probability of occurring under the IS distribution. Therefore, it is vital to have a method for selecting IS distributions that guarantees variance reduction. In other words, if chosen appropriately, an IS distribution will ensure that the resulting likelihood ratio is always less than one, and that for every path generated, a default event will always take place. In this article, we propose an effective IS algorithm that achieves such objectives, yet whose implementation requires no more complexity than that of a crude Monte Carlo algorithm. Our numerical results confirm that our proposed algorithm achieves effective variance reduction even when the probability of a *k*th default event is remote.

This article is organized as follows: The next section begins with a brief review of credit risk modeling and ends with a description of the Li [2000] model that our approach is based upon. Problem formulation for the valuation of *k*th-to-default basket default swaps is given in the third section. The fourth section proposes an IS algorithm that is simple to implement and that guarantees variance reduction. In the fifth section, we present estimates of fair spreads for *k*th-to-default basket default swaps with underlying obligors in the automobile, industrial, energy and finance sectors. The last section concludes.

CHARACTERIZATION OF DEFAULT TIMES

Credit risk modeling, in particular the modeling of default risk, divides the literature into two mainstreams: one stream considers default events as being reflected in a company's capital structure and the other considers default events as pure surprises. They are known, respectively, as the structural and reduced-form approaches.

The characterization of default risk as a structural approach first-passage time problem was initiated in Merton [1974]. The default event is defined as the value of a company's assets falling below the face value of its debt at any future date. Credit risk models that follow this line of thinking include Black and Cox [1976]; Geske [1977]; Leland [1994]; Longstaff and Schwartz [1995], and Zhou [2001]. All tackle the credit risk of a single entity with the exception of Zhou [2001] which considers the default correlations of two entities as two correlated Brownian motions and finds a closed-form solution for the probability of joint defaults. While Zhou's results cannot readily be extended to cover the case of many issuers, Hull and White [2000] demonstrates the use of Monte Carlo simulations to value a basket credit default swap. And Hull and White [2001] describes a default correlation model that can handle many issuers, but requires the use of computationally intensive numerical procedures.

Reduced-form models make direct use of the market information available (e.g., credit spreads and rating-transition probabilities) to arrive at default risk estimates. In contrast to structural models, firm values are bypassed and defaults are treated as exogenous. Schönbucher [2003]; Duffie [2003], and Lando [2004] provide a detailed review of the reduced-form approach to credit risk modeling.

Until now the industry standard for the joint default probability of multiple underlying obligors has relied on a model of joint default times. The copula approach for specifying the dependence structure among default times was introduced by Li [2000]. Laurent and Gregory [2005] introduced the factor-copula which conceptually coincides with the conditional independence assumption among default events (i.e., conditional on a systematic factor, default events are mutually independent). The conditional independence assumption results in a semi-analytic representation of the loss distribution. Models that follow this assumption include Andersen, Sidenius, and Basu [2003], in which a recursion technique is used to construct the loss distribution, and Gregory and Laurent [2004], in which the loss distribution is derived via the convolution of characteristic functions of individual losses. They apply the Fast Fourier Transform to numerically calculate the integral.

The following is a brief account of the joint-default time model that our approach is based upon.

Let τ_i denote the default time (or equivalently, the time to default) for an underlying obligor i, where $i = 1, \ldots, n$. τ_i is a positive random variable and its distribution is characterized by the hazard rate function $h_i(\cdot)$,

$$\text{prob}(\tau_i > t) = e^{-\int_0^t h_i(u)du} \tag{1}$$

Let $F_i(t)$ denote the cumulative default probability before time t for obligor i (the marginal distribution of the time to default for obligor i),

$$F_i(t) = \text{prob}(\tau_i \le t) = 1 - e^{-\int_0^t h_i(u)du}$$

and let $S_i(t)$ be the survival function of obligor i, so that $S_i(t)$ is expressed as $S_i(t) = 1 - F_i(t)$.

Using the quoted market spread on corporate bond or credit default swaps, one can extract the marginal distribution of the time to default, $F_i(t)$, for each obligor i. However, the cashflows of portfolio credit derivatives are a function of a sequence of random default times $(\tau_1, \ldots, \tau_n)$.

Therefore, in order to evaluate multi-name credit derivatives, the modeling challenge is to characterize the dependence structure for the default times, τ_i.

By sampling a set of correlated uniform variates $(U_1, U_2, \ldots, U_n)$, one can then specify a copula function $C(u_1, u_2, \ldots, u_n)$ which defines the dependence structure among default times to link univariate marginals into a multivariate joint distribution,

$$C(u_1, u_2, \ldots, u_n, \rho) = \text{prob}(U_1 \le u_1, U_2 \le u_2, \ldots, U_n \le u_n) \tag{2}$$

The copula-based approach involves the following steps: First, one must generate correlated random numbers X_i, where $i = 1, \ldots, n$. Second, uniformly distributed random variates, $U_i = \Phi(X_i)$, must be obtained from the cumulative normal distribution function, $\Phi(\cdot)$. The final step is the computation of default times for each individual obligor via an inverse mapping of their marginal distributions, $\tau_i = F_i^{-1}(U_i)$.

For the Gaussian copula, the sampling procedure in the first step is very simple. Based on a correlation matrix Σ for the asset returns, a set of correlated standard normal variates $(X_1, X_2, \ldots, X_n)$ are generated either by performing a Cholesky decomposition for the correlation matrix[1] or by using a factor-form representation such as in Laurent and Gregory [2005],

$$X_i = \rho_i Z_0 + \sqrt{1-\rho_i^2 Z_i}$$

where Z_0, Z_i, with $i = 1, \ldots, n$ are independent standard normal variables. The variable Z_0 is the common factor upon which default events of all underlying obligors are dependent, and Z_i represents the firm-specific or industry-specific risk factors. The parameter ρ_i determines how strongly X_i is correlated to the evolution of the common factor Z_0. The correlation between two credits i and j is $Corr(X_i, X_j) = \rho_i \rho_j$. The second and final steps require that the correlated uniform random variables are translated into default times via the inverse of the marginal distribution function for each underlying obligor.

VALUATION OF kTH-TO-DEFAULT BDSs

In this section, we briefly review the valuation procedure of a generic kth-to-default credit default swap relative to a portfolio of n $(n \ge k)$ reference risky obligors. Throughout the section we shall adapt the following notations: t_j is the time the jth premium payment takes place; $\Delta_{j-1,j}$ is the time increment between premium payments at the $(j-1)$th and the jth time points in years; $P(0, t_j)$ is the discount factor for a basket default swap contract with time to maturity t_j; R_i is the recovery rate for the ith obligor when default happens assuming that R_i is equal to a constant R for all i; τ is the k-th default time among the underlying obligors; and M_i is the notional amount for credit i assuming M_i is equal to a constant amount M for all i. The time to maturity of the basket default swap is T, where $T = t_N$; $F_i(t_j)$ is the probability that an underlying credit i defaults before or at time t_j; hence, by definition, $F_i(t_j) = \Pr(\tau_i \le t_j)$. The value $S(t)$ denotes the survival function of the k-th default time, $S(t) = \Pr(\tau > t)$. The distribution function of τ is therefore given by $F(t) = \Pr(\tau \le t) = 1 - S(t)$. Finally, Q denotes the risk-neutral probability measure, s denotes the credit spread, and $I(\cdot)$ is the indicator function.

If we define $PV^{premium}$ as the present value of the premium leg for the k-th-to-default basket swap, then

$$PV^{premium} = E^Q\left[\sum_{j=1}^{N} \Delta_{j-1,j} \times s \times P(0, t_j) \times M \times (\tau > t_j)\right] \tag{3}$$

where the expectation is a function of the risk-neutral pricing measure Q. If accrued premiums are considered, then the present value of the accrued premium $PV^{AccruedPremium}$ can be computed as

$$\begin{aligned} PV^{AccruedPremium} = E^Q\Bigg[&\sum_{j=1}^{N}\left(\tfrac{\tau - t_{j-1}}{t_j - t_{j-1}} \Delta_{j-1,j}\right) \times s \\ &\times P(0,\tau) \times M \times I(t_{j-1} < \tau \le t_j)\Bigg] \end{aligned} \tag{4}$$

And if we define $PV^{default}$ as the present value of the default leg for the k-th-to-default basket swap, then

$$PV^{default} = E^Q[(1 - R) \times M \times P(0, \tau) \times I(\tau \le T)] \tag{5}$$

We can therefore derive the fair spread for the kth-to-default basket swap as follows

$$s = \frac{E^Q[(1-R) \times P(0,\tau) \times I(\tau \le T)]}{E^Q\left[\Sigma_{j=1}^{N} \Delta_{j-1,j} \times P(0,t_j) \times I(\tau > t_j)\right]} \tag{6}$$

The crude Monte Carlo procedure for computing the value of the default leg can be described as follows:

1. Generate independent samples of standard normal variates $Z_0, Z_1, \ldots, Z_n$.
2. Generate correlated normal variates by setting $X_i = \rho_i Z_0 + \sqrt{1-\rho_i^2} Z_i$, $1 \le i \le n$.
3. Set $U_i = \Phi(X_i)$, $1 \le i \le n$.
4. Set $\tau_i = F_i^{-1}(U_i)$, $1 \le i \le n$.
5. Set τ = the k-th order statistic of $(\tau_1, \ldots, \tau_n)$.
6. Set the discounted payoff = $(1 - R)I(\tau \le T)P(0,\tau)$.
7. Repeat Steps 1 to 6 m times; the confidence interval of the DL can then be constructed using the m values of the discounted payoff.

This procedure provides a point estimate for the default leg, denoted by α, such that

$$\alpha = \frac{1}{m}\sum_{i=1}^{m}(1-R)I(\tau^{(i)} \le T)P(0,\tau^{(i)}) \tag{7}$$

where $\tau^{(i)}$ is the i-th independent sample of τ.

THE PROPOSED ALGORITHM

The efficiency of Monte Carlo simulation is measured by the variance of the estimator and the ease of generating a replication (Glynn and Whitt [1992]). Our goal is twofold: First, we need to establish a way of selecting an appropriate IS distribution among the set of possible measures in order to ensure k default events of interest on every generated path. Second, we need to ensure that the gain in variance reduction of the proposed algorithm always outweighs that of a crude Monte Carlo algorithm.

Let f be the density function of τ and suppose that $g(x)$ is also a density function, such that $g(x) > 0$ for $0 \le x \le T$. Let $L(x) = f(x)/g(x)$, which is the likelihood ratio, and $\tau^{(i)}$ be independent values sampled from $g(x)$. Then, we see that the point estimator

$$\beta = \frac{1}{m}\sum_{i=1}^{m}(1-R)I(\tau^{(i)} \le T)P(0,\tau^{(i)})L(\tau^{(i)}) \tag{8}$$

is an alternative estimator for the default leg. The point estimator β is known as an IS estimator (Glynn and Iglehart [1989], and Glasserman [2004]), and its variance depends on the choice of the IS density $g(x)$. In order to select an appropriate $g(x)$, we first need a simple alternative characterization for the default event of the basket $\{\tau \le T\}$.

Proposition 1. *The default event of the basket $\{\tau \le T\}$ is equivalent to the event $\{Z_0 \le b\}$ if b is the $(n-k+1)$-th order statistic of $(b_1, \ldots, b_n)$, where $b_i = (\Phi^{-1}(F_i(T)) - \sqrt{1-\rho_i^2}Z_i)/\rho_i$; that is, conditional on $Z_1, \ldots, Z_n$, the event of interest can be determined solely by the common factor Z_0.*

Proof. Since

$$\begin{aligned}
\{\tau_i \le T\} &= \{F_i^{-1}(\Phi(X_i)) \le T\} \\
&= \{\Phi(X_i) \le F_i(T)\} \\
&= \{X_i \le \Phi^{-1}(F_i(T))\} \\
&= \{\rho_i Z_0 + \sqrt{1-\rho_i^2} Z_i \le \Phi^{-1}(F_i(T))\} \\
&= \left\{ Z_0 \le \frac{\Phi^{-1}(F_i(T)) - \sqrt{1-\rho_i^2} Z_i}{\rho_i} \right\} \\
&= \{Z_0 \le b_i\}
\end{aligned}$$

we arrive at an alternative characterization for the default event $\{\tau_i \le T\}$ (i.e., the event $\{\tau_i \le T\}$ is equivalent to $\{Z_0 \le b_i\}$).

Now, consider the event of interest $\{\tau = T\}$. Note that

$$\begin{aligned}
I(\tau \le T) = 1 &\Leftrightarrow \sum_{i=1}^{n} I(\tau_i \le T) \ge k \\
&\Leftrightarrow \sum_{i=1}^{n} I(Z_0 \le b_i) \ge k
\end{aligned}$$

By setting $b = (n-k+1)$-th order statistic of $(b_1, \ldots, b_n)$, we see that

$$\sum_{i=1}^{n} I(Z_0 \le b_i) \ge k \Leftrightarrow I(Z_0 \le b) = 1$$

Thus, $\{\tau \le T\}$ if and only if $\{Z_0 \le b\}$.

Proposition 1 is the key idea behind our proposed algorithm. By carefully choosing the probability measure for Z_0, we can ensure that $\{Z_0 \le b\}$ occurs with a 100% probability. Therefore, Proposition 1 provides a simple way to ensure that for every path generated, k default events always take place. However, there are an infinite number of measures which prohibit the event $\{Z_0 \leq b\}$ from occurring with a 100% probability. Thus, we need an additional guideline for selecting an appropriate IS distribution for Z_0 among the set of all possible measures that force k default events to take place on every generated path. The following lemma provides a characterization for such guideline.

Lemma 2. *If X is a random variable with density function $f_X(x)$ and distribution function $F_X(x)$, and we are interested in estimating $F_X(a)$, if we define the likelihood ratio $L(x)$ as*

$$L(x) := f_X(x)/f_X^*(x)$$

for $x \leq a$, where the density function $f_X^*(x)$ for X is chosen to be

$$f_x^*(x) = \begin{cases} f_X(x)/F_X(a), & x \leq a \\ 0, & \text{otherwise} \end{cases}$$

then, $Y \triangleq I(X \leq a)\ L(x)$ is an unbiased estimator for $F_X(a)$ and $Var^*[Y] = 0$.

Proof. It is easy to verify that $f_X^*(x)$ is an optimal IS density for this simple problem, because it produces an estimator with zero variance. By definition, $L(x) = f_X(x)/f_X^*(x) = F_X(a)$ for $x \leq a$, therefore

$$\begin{aligned} E^*[I(X \leq a)L(x)] &= \int_{-\infty}^{a} F_X(a) f_X^*(x)dx \\ &= F_X(a)\int_{-\infty}^{a} F_X^*(x)dx = F_X(a) \end{aligned} \tag{9}$$

where $E^*(\cdot)$ denotes the expectation operator under the density function $f_X^*(x)$. Hence, Y is an unbiased estimator for $F_X(a)$.

Now, if $Var^*(\cdot)$ denotes the variance operator under the density function $f_X^*(x)$, then

$$\begin{aligned} Var^*[Y] &= E^*[I^2(X \leq a)L^2(x)] - (E^*[I(X \leq a)L(x)])^2 \\ &= \int_{-\infty}^{a} F_X^2(a) f_X^*(x)dx - F_X^2(a) \\ &= F_x^2(a)\int_{-\infty}^{a} f_X^*(x)dx - F_X^2(a) = 0 \end{aligned} \tag{10}$$

The above analysis provides us with a guideline for choosing the new probability measure for Z_0. In particular, for a given b, when the common factor Z_0 is assumed to be Gaussian, we can set

$$f^*(z) = \begin{cases} \phi(z)/\Phi(b), & z \leq b; \\ 0, & \text{otherwise} \end{cases}$$

where $\phi(z)$ is a standard Gaussian density. With this specific IS density, the proposed algorithm is as follows:

1. Generate independent samples of standard normal variates $Z_1, \ldots, Z_n$.
2. Set
$$b_i = \frac{\Phi^{-1}(F_i(T)) - \sqrt{1-\rho_i^2}\,Z_i}{\rho_i}, \quad 1 \leq i \leq n$$
and let $b = (n-k+1)$-th order statistic of $(b_1, \ldots, b_n)$.
3. Set $L = \Phi(b)$ and generate Z_0 according to the formula $\Phi^{-1}(LU)$, where U is a uniform(0, 1) random variate.
4. Set $X_i = \rho_i Z_0 + \sqrt{1-\rho_i^2}\,Z_i, \quad 1 \leq i \leq n.$
5. Set $U_i = \Phi(X_i), 1 \leq i \leq n.$
6. Set $\tau_i = F_i^{-1}(U_i), 1 \leq i \leq n.$
7. Set τ = the k-th order statistic of $(\tau_1, \ldots, \tau_n)$.
8. Set the discounted payoff = $(1 - R)\mathrm{P}(0, \tau)L$.
9. Repeat Steps 1 to 8 m times; the confidence interval of the default leg can then be constructed by the m copies of the discounted payoff.

Note that our proposed algorithm does not require the common factor Z_0 to be restricted to the Gaussian case. If Z_0 is not distributed normally, but has a continuous distribution function $F(\cdot)$, and the associated inverse function $F^{-1}(\cdot)$ can be identified, only Step 3 requires further modification by replacing $\Phi(\cdot)$ and $\Phi^{-1}(\cdot)$ with $F(\cdot)$ and $F^{-1}(\cdot)$, respectively. The algorithm also allows for more generalized problem setting. When the specific factors $Z_1, \ldots, Z_n$ are of arbitrary distributions, we need only to modify Steps 1, 2, and 5 where $\Phi(\cdot)$ is replaced by the distribution function associated with X_i. Finally, the recovery rate R in Step 8 can be assumed to be random and correlated to the default times τ_i.

In addition to the advantages listed above, the strength of our proposed lies in its variance reduction efficiency established by the following proposition:

Proposition 3. If β^* denotes the estimator generated by the proposed algorithm and α denotes the estimator produced by a crude Monte Carlo algorithm, then the following always holds:

$$Var^*(\beta^*) < Var(\alpha)$$

where $Var^*(\cdot)$ is the variance operator under the density function f^*.

Proof. Let

$$V(\tau, R) = (1 - R)P(0, \tau)I(\tau \leq T)$$

then

$$Var^*(\beta^*) = Var^*(V(\tau, R)L]/m^2$$

and

$$Var(\alpha) = Var(V(\tau, R)L]/m^2$$

Therefore,

$$Var^*(V(\tau, R)L) < Var V(\tau, R))$$

Since $L = \Phi(b)$ is always <1, we have

$$\begin{aligned} Var^*(V(\tau,R)L) &= E^*[V(\tau,R)^2 L^2] - E^*[V(\tau,R)L]^2 \\ &< E^*[V(\tau,R)^2 L] - E[V(\tau,R)]^2 \\ &= E[V(\tau,R)^2] - E[V(\tau,R)]^2 \\ &= Var[V(\tau,R)] \end{aligned}$$

Proposition 3 ensures that the gain in variance reduction of the proposed algorithm always outweighs that of a crude Monte Carlo algorithm. In practice, the actual gain in computational efficiency depends on a few driving factors; we will discuss the two most important ones. First, the analysis in the proof of Proposition 3 shows that the magnitude of $Var^*(\beta^*)$ depends on the variance of the likelihood ratio L. As ρ_i approaches 1, b_i becomes deterministic. Therefore, if ρ_i is near 1, the variance of L is small. This implies that the proposed algorithm is noticeably more effective when the factor loadings, ρ_i, for the common factor Z_0 are high. The numerical results in the next section confirm this observation. Second, the IS technique is well known to be most effective for rare-event simulations. Thus, the performance gain should further increase when the default probability is low. Our numerical results in the next section also confirm this expectation.

NUMERICAL RESULTS

Based on the IS algorithm proposed in the previous section, we can present numerical results for the estimation of fair spreads for kth- to-default basket default swaps.

The market quotes for standardized 5-year *k*th-to-default basket default swaps are obtained from J. P. Morgan and are defined using five underlying reference entities. The risk-neutral default probabilities of the reference entities over the 5-year period are extracted directly from market-quoted credit default swap spreads. All five underlying references are of equal nominal value and assumed to have a recovery rate of 40%. We study four sector *k*th-to-default basket default swaps with underlyings specific to each of the following sectors: automobiles, industrials, energy and finance.[2]

Our focus is on estimating the expected values of the default leg and premium leg, which are vital inputs in calculating fair spreads, *s*, for the *k*th-to-default basket default swaps. In order to avoid cancellation effects, in each exhibit we present only the results for the default legs.

In our simulations, the values for k are 1, 2, and 3. We use the risk free rate $r = 5\%$. To understand how correlation affects the variance reduction efficiency of our IS algorithm, we repeat the same calculation with the constant factor loadings of 0.5, 0.7, 0.9, and 0.99.

We report the standard errors (s.e.) of the simulation with and without importance sampling in the fourth and sixth columns of each exhibit, where the standard error for estimator $\beta^* = \frac{1}{m}\Sigma_{i=1}^m Y_i$ is defined as

$$s.e(\beta^*) = \frac{\sqrt{\frac{1}{m-1}\Sigma_{i=1}^m (Y_i - \beta^*)^2}}{\sqrt{m}}$$

By comparing the two columns, one readily observes the significant gain in efficiency achieved by our proposed IS algorithm in terms of the overall variance reduction.

In the last column of each exhibit, we list the sample variance ratio $R = \frac{[s.e.(\alpha)]^2}{[s.e.(\beta^*)]^2} \times 10$, where α refers to the crude simulation estimator and β^* refers to the IS estimator based on our IS technique.

As is evident in each exhibit, the estimator for the default leg, based on our IS algorithm, is substantially more efficient than that estimated by crude Monte Carlo simulations, for all cases considered. The value in the last column of each exhibit shows that the IS estimator, based on our proposed algorithm, shows a marked decrease in variance. The efficiency gain is more prominent when the correlation ρ is higher. Moreover, the efficiency of our proposed IS algorithm increases as the value of k increases, which indicates that our algorithm reduces the variance of the estimator especially when the possibility of the k-th default event is remote.

CONCLUSION

In this article, we have proposed an importance sampling algorithm for the valuation of basket default swaps. The algorithm is simple to implement and guarantees variance reduction via the appropriate choice of the IS distribution.

In Proposition 1, we established a way of ensuring that for every path generated, k default events always take place. We also established an additional guideline for selecting an appropriate IS distribution among the set of possible measures that can force the required events of interest to take place; Lemma 2 characterizes this guideline. Although importance sampling is prone to ineffective variance reduction, and in some cases to an increase in the variance of the point estimators, the strength of our proposed algorithm lies in its unconditional variance reduction efficiency; this assertion is validated in Proposition 3.

Our numerical results indicate that the proposed algorithm is considerably more efficient than the crude simulation method. The efficiency gain is more prominent for underlying credits with a high correlation and especially for basket default swaps with a remote probability of a kth default. We therefore conclude that the

EXHIBIT 1

Estimated Values for First (k = 1), Second (k = 2), and Third (k = 3)-to-Default BDS: Automobile Sector

First-to-Default Basket Default Swap ($k = 1$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default Probability	Default Leg (with 10^6 runs)	Standard Error s.e. (α)	Default Leg (with 10^5 runs)	Standard Error s.e. (β^*)	Reduction R
0.5	0.171066	0.091444	0.000202	0.092286	0.000404	2.5
0.7	0.144736	0.077525	0.000189	0.077637	0.000261	5.2
0.9	0.099715	0.053398	0.000161	0.053508	0.000106	23.2
0.99	0.063024	0.033581	0.000130	0.033660	2.57E-05	254.4

Second-to-Default Basket Default Swap ($k = 2$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default Probability	Default Leg (with 10^6 runs)	Standard Error s.e. (α)	Default Leg (with 10^5 runs)	Standard Error s.e. (β^*)	Reduction R
0.5	0.031144	0.016136	9.02E-05	0.016085	0.000126	5.1
0.7	0.043805	0.022966	0.000108	0.022808	9.55E-05	12.7
0.9	0.052184	0.027674	0.000118	0.027518	5.14E-05	52.9
0.99	0.047831	0.025420	0.000114	0.025269	1.59E-05	514.0

Third-to-Default Basket Default Swap ($k = 3$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default Probability	Default Leg (with 10^6 runs)	Standard Error s.e. (α)	Default Leg (with 10^5 runs)	Standard Error s.e. (β^*)	Reduction R
0.5	0.005457	0.002781	3.76E-05	0.002808	3.54E-05	11.3
0.7	0.014542	0.007541	6.22E-05	0.007565	3.87E-05	25.9
0.9	0.030589	0.016124	9.10E-05	0.016240	3.19E-05	81.6
0.99	0.039179	0.020808	0.000103	0.020892	1.28E-05	654.5

EXHIBIT 2

Estimated Values for First ($k = 1$), Second ($k = 2$), and Third ($k = 3$)-to-Default BDS: Industrial Sector

First-to-Default Basket Default Swap ($k = 1$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.29367	0.157668	0.000245	0.158459	0.000502	2.4
0.7	0.246693	0.132590	0.000232	0.132324	0.000339	4.7
0.9	0.17272	0.092680	0.000203	0.092382	0.000149	18.7
0.99	0.110533	0.058994	0.000168	0.059096	3.61E-05	216.6

Second-to-Default Basket Default Swap ($k = 2$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.079521	0.041325	0.000141	0.041264	0.000222	4.0
0.7	0.095089	0.049969	0.000155	0.050101	0.000164	8.9
0.9	0.100733	0.053511	0.00016	0.053364	8.39E-05	36.5
0.99	0.091452	0.048669	0.000154	0.048749	2.65E-05	336.6

Third-to-Default Basket Default Swap ($k = 3$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.019114	0.009763	7.01E-05	0.009912	8.38E-05	7.0
0.7	0.037817	0.019633	9.93E-05	0.019777	7.89E-05	15.8
0.9	0.064	0.033760	0.000129	0.033863	5.66E-05	52.3
0.99	0.078513	0.041700	0.000143	0.041746	2.26E-05	401.4

proposed algorithm provides a suitable basis for efficient estimation of the default leg when calculating the fair spreads of kth-to-default basket default swaps.

ENDNOTES

[1]To generate correlated standard normal variates $\underline{X} = (X_1, X_2, \ldots, X_n)$, we first apply the Cholesky decomposition to the correlation matrix $\Sigma = AA^T$, then generate n i.i.d. standard normal random numbers $z_1, \ldots, z_n \in (0, 1)$. Finally, we set $\underline{X} = A_z$, where $\underline{X} = (X_1, \ldots, X_n)$ are the correlated random numbers fitting the joint normal distribution.

[2]The five underlyings of the kth-to-default basket default swap for the auto sector are Volkswagen AG, Peugoet, Renault, Continental AG, and Valeo. The reference names of the kth-to-default basket default swap for the industrial sector are Bayer AG, Rolls-Royce, Lafarge, Compagnie de Saint-Gobain, and BAE systems. The five underlyings of the kth-to-default basket default swap for the financial sector are Bayerische Hypo-Und Vereinsbank, AXA, Allianz, Commerzbank, and Abbey National. The reference names of the kth-to-default basket default swap for the energy sector are Suez, Veolia Environment, RWE, Endessa, and Repsol YPF.

EXHIBIT 3

Estimated Values for First (k = 1), Second (k = 2), and Third (k = 3)-to-Default BDS: Financial Sector

First-to-Default Basket Default Swap ($k = 1$)

ρ	Default Probability	Naive Simulation: Default Leg (with 10^6 runs)	Naive Simulation: Standard Error s.e. (α)	Importance Sampling: Default Leg (with 10^5 runs)	Importance Sampling: Standard Error s.e. (β^*)	Variance Reduction R
0.5	0.300344	0.161309	0.000247	0.160582	0.000504	2.4
0.7	0.25196	0.135468	0.000234	0.135485	0.000343	4.7
0.9	0.177454	0.095189	0.000206	0.095386	0.000153	18.1
0.99	0.11644	0.062120	0.000172	0.062090	3.84E-05	199.8

Second-to-Default Basket Default Swap ($k = 2$)

ρ	Default Probability	Naive Simulation: Default Leg (with 10^6 runs)	Naive Simulation: Standard Error s.e. (α)	Importance Sampling: Default Leg (with 10^5 runs)	Importance Sampling: Standard Error s.e. (β^*)	Variance Reduction R
0.5	0.082375	0.042802	0.000143	0.042949	0.000229	3.9
0.7	0.0976	0.051283	0.000156	0.051491	0.000166	8.8
0.9	0.1037	0.055064	0.000162	0.054873	8.64E-05	35.3
0.99	0.095534	0.050837	0.000157	0.050749	2.84E-05	305.4

Third-to-Default Basket Default Swap ($k = 3$)

ρ	Default Probability	Naive Simulation: Default Leg (with 10^6 runs)	Naive Simulation: Standard Error s.e. (α)	Importance Sampling: Default Leg (with 10^5 runs)	Importance Sampling: Standard Error s.e. (β^*)	Variance Reduction R
0.5	0.019848	0.010134	7.14E-05	0.010289	8.59E-05	6.9
0.7	0.039201	0.020348	0.000101	0.020336	8.15E-05	15.3
0.9	0.065671	0.034654	0.000131	0.034762	5.81E-05	50.8
0.99	0.079967	0.042483	0.000145	0.042530	2.39E-05	364.3

REFERENCES

Andersen, L., J. Sindenius, and S. Basu. "All Your Hedges in One Basket." *Risk* (November 2003), pp. 67–72.

Black, F., and J. Cox. "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions." *Journal of Finance*, 31 (1976), pp. 351–367.

Duffie, D., and K.J. Singleton. *Credit Risk Modeling: Pricing, Measurement and Management*. Princeton, NJ: Princeton University Press, 2003.

Geske, R. "The Valuation of Corporate Liabilities as Compound Option." *Journal of Financial and Quantitative Analysis*, 5 (1977), pp. 541–552.

Glasserman, P. *Monte Carlo Methods in Financial Engineering*. New York: Springer Verlag, 2004.

Glynn, P.W., and D.L. Iglehart. "Importance Sampling for Stochastic Simulations." *Management Science*, 35 (1989), pp. 1367–1392.

Glynn, P.W., and W. Whitt. "The Asymptotic Efficiency of Simulation Estimators." *Operations Research*, 40 (1992), pp. 505–520.

Gregory, J., and J.P. Laurent. "In the Core of Correlation." *Risk* (October 2004), pp. 87–91.

Hull, J., and A. White. "Valuing Credit Default Swaps I: No Counterparty Default Risk." *Journal of Derivatives*, Vol. 8, No.1 (2000), pp. 29–40.

EXHIBIT 4

Estimated Values for First (k = 1), Second (k = 2), and Third (k = 3)-to-Default BDS: Energy Sector

First-to-Default Basket Default Swap ($k = 1$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.225387	0.120701	0.000225	0.120296	0.000452	2.5
0.7	0.189166	0.101458	0.000211	0.102356	0.000301	4.9
0.9	0.130722	0.070057	0.000181	0.070316	0.000126	20.7
0.99	0.083519	0.044520	0.000148	0.044550	3.09E-05	228.6

Second-to-Default Basket Default Swap ($k = 2$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.049933	0.025911	0.000113	0.026161	0.00017	4.4
0.7	0.064506	0.033847	0.000129	0.033785	0.000126	10.5
0.9	0.072649	0.038539	0.000138	0.038397	6.60E-05	43.8
0.99	0.065248	0.034722	0.000132	0.034940	2.06E-05	408.7

Third-to-Default Basket Default Swap ($k = 3$)						
ρ		Naive Simulation		Importance Sampling		Variance
	Default	Default Leg	Standard Error	Default Leg	Standard Error	Reduction
	Probability	(with 10^6 runs)	s.e. (α)	(with 10^5 runs)	s.e. (β^*)	R
0.5	0.010395	0.005302	5.18E-05	0.005304	5.44E-05	9.1
0.7	0.023691	0.012287	7.91E-05	0.012265	5.53E-05	20.4
0.9	0.044602	0.023524	0.000109	0.023483	4.25E-05	66.0
0.99	0.054828	0.029117	0.000121	0.029240	1.68E-05	522.1

——— "Valuing Credit Default Swaps II: Modeling Default Correlations." *Journal of Derivatives*, Vol. 8, No. 3 (2001), pp. 12–22.

Joshi, M.S., and D. Kainth. "Rapid and Accurate Development of Prices and Greeks for Nth to Default Credit Swaps in the Li Model." *Quantitative Finance*, 4 (2004), pp. 266–275.

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

Laurent, J.P., and J. Gregory. "Basket Default Swaps, CDOs and Factor Copulas." *Journal of Risk*, 7 (2005), pp. 103–122.

Leland, H.E. "Risky Debt, Bond Covenants and Optimal Capital Structure." *Journal of Finance*, 49 (1994), pp. 1213–1252.

Li, D.X. "On Default Correlation: A Copula Function Approach." *Journal of Fixed Income*, 9 (2000), pp. 43–54.

Longstaff, F., and E.S. Schwartz. "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt." *Journal of Finance*, 50 (1995), pp. 789–819.

Mashal, R., and M. Naldi. "Pricing Multiname Credit Derivatives: Heavy Tailed Hybrid Approach." Working paper, 2002.

Merton, R. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 29 (1974), pp. 449–470.

Rogge, E., and P.J. Schönbucher. "Modelling Dynamic Portfolio Credit Risk." Working paper, 2003.

Schönbucher, P. *Credit Derivatives Pricing Models: Model, Pricing and Implementation*. Chichester, England: Wiley, 2003.

Zhou, C. "An Analysis of Default Correlations and Multiple Defaults." *Review of Financial Studies*, Vol. 14, No. 2 (2001), pp. 555–576.