

EFFICIENT SIMULATION FOR RISK MEASUREMENT IN PORTFOLIO OF CDOs

Michael Gordy

Federal Reserve Board
Washington, DC 20551, USA

Sandeep Juneja

Tata Institute of Fundamental Research
Homi Bhabha Road, Mumbai 400 005, India

ABSTRACT

We consider a portfolio containing CDO tranches as well as ordinary bonds. Our interest is in large loss probabilities and risk measures such as value-at-risk. When loss is measured on a mark-to-market basis, estimation via simulation requires a nested procedure: In the outer step one draws realizations of all risk factors up to the horizon, and in the inner step one re-prices each instrument in the portfolio at the horizon conditional on the drawn risk factors. Practitioners perceive the computational burden of such nested schemes to be unacceptable, and adopt a variety of somewhat ad hoc measures to avoid the inner simulation. In this paper, we question whether such short cuts are necessary. We show that a relatively small number of trials in the inner step can yield accurate estimates, and analyze how a fixed computational budget may be allocated to the inner and the outer step to minimize the mean square error of the resultant estimator.

1 INTRODUCTION

A collateralized debt obligation (“CDO”) is a structured financial product in which a set of debt claims serve as collateral for a set of securities that differ in seniority. The juniormost claimant, known as the equity tranche, is the first to absorb default losses in the pool. When default losses fully exhaust the equity tranche, a mezzanine tranche absorbs subsequent losses, and so on. Senior tranches are very unlikely to realize any credit losses and typically are rated AA or AAA by public rating agencies. CDOs have revolutionized bank balance sheet management, and more recently have given rise to a speculative market in “correlation trading” of standardized CDO tranches constructed synthetically on traded credit default swap indexes.

Most existing research related to CDOs concerns pricing. Duffie and Gârleanu (2001) demonstrate how CDO tranches may be priced using the stochastic default intensity methods widely applied to single-name investment grade bonds and credit derivatives. This paper remains influential as a conceptual benchmark, but practitioners generally find the computational burden of this model prohibitive

for real-time trading. Instead, copula models based on Li (2000) are commonly used, despite compelling evidence that calibration of such models to prices can yield misleading comparative statics (Willemann 2005).

Risk-measurement applications introduce additional challenges. When loss is measured on a mark-to-market basis, estimation via simulation of risk-measures such as Value-at-Risk (“VaR”) or large loss probabilities requires a nested procedure: In the outer step one draws realizations of all risk factors up to the horizon, and in the inner step one re-prices each instrument in the portfolio at the horizon conditional on the drawn risk factors. As stochastic default intensity models are perceived as too costly even for one-time pricing calculations, it has been widely assumed that such models are infeasible for the inner step of a VaR simulation, because the inner step must be executed once for each trial in the outer step. Copula pricing methods may sometimes be used in the inner step. However, copula models are inherently static, and so cannot be applied to repricing at a future horizon except by ad hoc means. More typically, practitioners avoid the inner step altogether by decomposing the CDO tranches into credit default swaps (“CDS”) on the individual names in the collateral pool. The exposure assigned to each name is determined by the delta of the tranche price with respect to a small change in the name’s CDS spread. How well this approach works in practice has not been studied carefully, but intuition suggests that it will work poorly. In risk-management applications, one is interested in tail events that generate large portfolio losses. Large losses occur when many *defaults* (and not just small credit spread changes) have occurred. Delta-implied exposures in general will bear little resemblance to tranche price changes in the event of a default in the collateral pool.

In this paper, we question whether inner step simulations must necessarily impose a large computational burden. We show that a relatively small number of trials in the inner step can yield accurate estimates for VaR and large loss probabilities. Since an expectation is replaced by a noisy sample mean, the estimator is biased, and we are able to characterize this bias asymptotically. Our main contribution is to analyze how a fixed computational budget may be allocated to the inner and the outer step to minimize the

mean square error of the resultant estimator. We note that the inner step can be implemented in broadly two different ways, and we analyze the optimal allocation in both the settings. Numerical results testing the efficacy of the proposed algorithms will be presented at the conference.

In Section 2 we develop the modelling framework for portfolio of CDOs. In Section 3, we introduce the nested simulation methodology, and analyze the optimal allocation of computational resources between the two stages that minimizes the mean square error of the resultant estimator. Here, we also characterize the bias in the simulation estimator. We conclude briefly in Section 4. Here we also discuss our related ongoing research.

2 MODEL OF PORTFOLIO OF CDO TRANCHES

In this section, we develop a model of a portfolio containing CDO tranches as well as single-name instruments. In our setting, a single-name instrument can be represented as a special case of a CDO, so without loss of generality we assume that all positions in the portfolio are CDO tranches. For simplicity in exposition, these are assumed to be *unfunded, synthetic* CDOs. An unfunded synthetic CDO is nothing more than a basket of single-name credit default swaps (CDS) with prioritized structuring of default payouts by the tranche holders. This greatly simplifies the cashflow rules without limiting the more general applicability of the methodology.¹

The reference names in the CDO collateral pools are drawn from a universe consisting of m obligors. Each CDO tranche is defined by a set of variables $(A_t, \Gamma_t, \Delta_t, s, T)$ where

- A_t is a vector $A_t = (a_{1t}, \dots, a_{mt})$ of dollar exposures in the CDO pool to each name in the underlying universe. Exposure is zero for names not included in a pool.
- Γ_t and Δ_t are the attachment and detachment points, respectively, for a tranche. These too are expressed in dollars. The face value of the tranche at time t is $\Delta_t - \Gamma_t$.
- s is the spread on the tranche. We assume the CDO issuer pays the tranche holder a continuous stochastic premium of $s(\Delta_t - \Gamma_t)$. Note that s is constant over the life of the CDO, but the residual face value of the tranche is nonincreasing over time due to default losses.
- T is the maturity of the CDO.

¹In the US and Europe, there are CDS indices based on the most widely-traded names. One can also trade standard “index tranches” of unfunded synthetic CDOs of the index names. The index tranche market has become quite liquid. Positions in this market fall naturally into our model framework.

We can approximate a “plain vanilla” loan or bond in the portfolio as a single-name credit default swap. That is, the credit risk for the holder of a bond issued by obligor i is equivalent to the credit risk for the seller of a CDS written on i . It is represented by a vector A_t that is equal to the bond exposure a at position i and zeros elsewhere; $\Gamma_t = 0$; $\Delta_t = a$; and T equal to the maturity of the bond. Therefore, without loss of generality, we say that the portfolio consists of K CDO tranches. Position k will be represented by

$$\mathbf{P}_{kt} = (A_{kt}, \Gamma_{kt}, \Delta_{kt}, s_k, T_k).$$

Note that we can have two positions in the pool that are different tranches of the same CDO—no complications arise in this case.

Credit risk at the obligor level is modeled in the reduced-form framework of Duffie and Singleton (1999). Obligor i has a stochastic default time τ_i . When $t < \tau_i$, the obligor has a stochastic default intensity λ_{it} . The vector of default intensities is denoted Λ_t . In models of multi-name derivatives such as CDOs and basket default swaps, cross-sectional dependence in Λ_t is a central concern. We assume that the λ_{it} are correlated via common dependence on a diffusion W_t . Conditional on W_t , changes in the λ_{it} are independent and default events are independent.

When a default occurs in a CDO pool, the tranche holder pays out its contractual protection. Let χ_{jt} be the payout required of the tranche in the event that obligor j defaults at time t . A payout is required only if the loss on this default exhausts all tranches junior to the given CDO tranche, and can never be greater than the remaining face value of the CDO tranche, so

$$\chi_{jt} = \min\{\max\{(1 - R_{jt})a_{jt} - \Gamma_t, 0\}, \Delta_t - \Gamma_t\}$$

where R_{jt} is a recovery rate. The recovery rate may be stochastic but, following standard practice, we assume that any uncertainty is idiosyncratic. Now let X_{jt} represent the cumulative payout by the tranche for default of obligor j . This is

$$X_{jt} = I_{t \geq \tau_j} \cdot \chi_{j, \min\{t, \tau_j\}}$$

Subsequent to the payout on the defaulted obligor, the vector of exposures A and the attachment and detachment points must be updated. Exposure to defaulted obligor j in the CDO pool is set to zero, i.e., $a_{j, \tau_j^+} = 0$. The tranche attachment and detachment are decremented as well via:

$$\begin{aligned} d\Gamma_t &= -\min\{\Gamma_t, \sum_i dX_{it}\} \\ d\Delta_t &= -\min\{\Delta_t, \sum_i dX_{it}\} \end{aligned}$$

To keep the focus on credit risk, we assume that risk-free interest rates are constant at r . In this case, the price of position k at time t can be written as the function

$$V(\Lambda_t, A_{kt}, \Gamma_{kt}, \Delta_{kt}, s_k, T_k - t).$$

Note that V does not depend directly on k because the instruments all take the same contractual form. Similarly, V does not depend directly on time because the effects of the passage of time on the CDO are absorbed in the A , Γ and Δ parameters and the remaining maturity $T_k - t$. Thus, we can write a very generic pricing function.

Our objective is to model the distribution of the discounted value of portfolio profit and loss at a fixed horizon. Profit and loss (“P/L”) is measured on a mark-to-market basis and includes the cashflows received and payed out. We normalize the present time to 0. The model horizon is H . For convenience, we assume $H \leq \min\{T_k\}$. Realized loss on position k is denoted Y_k and is calculated as

$$Y_k = V(\Lambda_0, A_{k0}, \Gamma_{k0}, \Delta_{k0}, s_k, T_k) - e^{-rH}V(\Lambda_H, A_{kH}, \Gamma_{kH}, \Delta_{kH}, s_k, T_k - H) + \int_0^H e^{-rt} dX_{kt} - s_k \cdot \int_0^H e^{-rt} (\Delta_{kt} - \Gamma_{kt}) dt.$$

The first term is current value, the second term is the discounted value at the horizon, the third term is the discounted value of the payouts made, and the last term is the discounted value of the stream of premium payments received. Note that Y_k is negative whenever the position generates a profit. Portfolio loss is simply $Y = \sum_{k=1}^K Y_k$.

Whenever P/L is modeled on a mark-to-market basis, we need to consider the relationship between the empirical and risk-neutral properties. The empirical measure governs the evolution of the portfolio over the horizon, whereas the risk-neutral measure governs the repricing of the portfolio at the horizon. In a large and efficient market, it is reasonable to assume that investors demand a premium only on non-diversifiable sources of risk. In our model, W_t represents the only such source of risk, so we assume that a change in drift in W_t determines the relationship between λ_{it} and its counterpart $\tilde{\lambda}_{it}$ under the risk-neutral measure Q . Note that the default event itself (i.e., conditional on λ_{it}) attracts no risk-premium, so at each moment in time we have $\tilde{\lambda}_{it} = \lambda_{it}$. That is, the equality holds instantaneously at each point in time, but the two evolve into the future by different processes (Lando, Jarrow, and Yu 2005).

3 SIMULATION FRAMEWORK

We now develop notation related to the simulation process. The simulation is nested: There is an “outer step” in which we draw histories up to the horizon H . For each trial in

the outer step, there is an “inner step” simulation needed for repricing at the horizon.

Let L be the number of trials in the outer step. In each of these trials, we

1. Draw a single path W_t under the empirical measure for $t \in (0, H]$. Here, the interval $(0, H]$ may be discretized into a large number of equally spaced times and W_t may be sampled at these times (see, e.g., Glasserman 2004).
2. Conditional on W_t , draw paths for the default intensities.
3. Conditional on the $\{\lambda_{it}\}$, draw default events and (when default occurs) recovery rates.
4. Evaluate the cashflow components of Y_k for each k .
5. Evaluate the price of each position at H .

It is the last step that calls for the “inner step” simulation. We consider two distinct methods to conduct this inner step simulation:

- A. We draw N independent trials of joint paths for the m obligors on $(H, \max\{T_k\}]$, and in each path calculate the price of each CDO tranche as a discounted expected value of future cashflows.
- B. For position k , draw $p_k N$ trials of joint paths for the obligors with positive exposure in CDO k . For each path, we calculate the price of position k as a discounted expected value of future cashflows. Here $p_k \geq 0$ and $\sum_{k \leq K} p_k = 1$. The key thing here is that we price each position independently of other positions.

In either case, inner-step simulations are governed by the risk-neutral measure Q .

We assume that the CDO prices at time 0 are already known and can be taken as constants in our algorithm. Of course, this can be relaxed.

Let the random vector ξ denote all the economic information available at the model horizon H that is required to reprice in the inner step simulations and thereby complete a draw of portfolio loss Y . A sufficient representation of ξ includes

- Default times τ_i and recoveries R_i for obligors that default in $(0, H]$.
- The time- H default intensity vector Λ_H for surviving borrowers.
- The time- H common factor W_H .

3.1 Estimating the Probability of Large Losses

We first consider the problem of efficient estimation of $\alpha = P(Y(\xi) > u)$ via simulation for a given u . If for

each generated ξ , the mark-to-market values of each position were known, the associated $Y(\xi)$ would be known and simulation would involve generating i.i.d. samples $Y_1(\xi_1), Y_2(\xi_2), \dots, Y_L(\xi_L)$ and taking the average

$$\frac{1}{L} \sum_{i=1}^L I_i(Y_i(\xi_i) > u)$$

as an estimator of α . However, the mark-to-market value of each CDO is not known and is estimated via the inner step simulations.

We first focus on the Method A described earlier in this section for generating the inner step simulations. Method B is similar and is discussed later. Suppose that through outer step simulation a sample ξ is generated. For CDO k , suppose that $(\Lambda_H, A_{kH}, \Gamma_{kH}, \Delta_{kH}, s_k, T_k - H)$ denotes the information associated with ξ . Then, in the inner step, N trials are conducted, and at each trial, a noisy sample of $\sum_{k \leq K} V(\Lambda_H, A_{kH}, \Gamma_{kH}, \Delta_{kH}, s_k, T_k - H)$ of mark-to-market CDO prices is generated. Let $Z_i(\xi)$ denote the zero-mean pricing error associated with the i^{th} such sample, and set

$$\bar{Z}^N(\xi) = \frac{1}{N} \sum_{i=1}^N Z_i(\xi).$$

In place of $Y(\xi)$, we take as its surrogate $Y(\xi) + \bar{Z}^N(\xi)$ as an estimate of loss in the portfolio. By the law of large numbers,

$$\bar{Z}^N(\xi) \rightarrow 0 \quad a.s.$$

as $N \rightarrow \infty$. The estimator for $P(Y(\xi) > u)$ then involves generating i.i.d. samples

$$(Y_1(\xi_1) + \bar{Z}_1^N(\xi_1)), \dots, (Y_L(\xi_L) + \bar{Z}_L^N(\xi_L))$$

via outer and inner step simulation and taking the average

$$\hat{\alpha}_{LN} = \frac{1}{L} \sum_{i=1}^L I_i[Y_i(\xi_i) + \bar{Z}_i^N(\xi_i) > u].$$

We now develop an expression for the mean square error of the estimator represented by $\hat{\alpha}_{LN}$ for both Methods A and B and discuss the parameters that minimize this.

3.1.1 Inner step simulation: Method A

Let α_N denote $P(Y(\xi) + \bar{Z}^N(\xi) > u)$. The mean square error of the estimator $\hat{\alpha}_{LN}$ separates into

$$\begin{aligned} E(\hat{\alpha}_{LN} - \alpha)^2 &= E(\hat{\alpha}_{LN} - \alpha_N + \alpha_N - \alpha)^2 \\ &= E(\hat{\alpha}_{LN} - \alpha_N)^2 + (\alpha_N - \alpha)^2. \end{aligned}$$

Further, note that

$$E(\hat{\alpha}_{LN} - \alpha_N)^2 = \frac{\alpha_N(1 - \alpha_N)}{L}.$$

Suppose that the computational effort to generate one sample of ξ (outer step simulation) on an average equals $\beta > 0$, and to generate an inner step simulation sample on the average equals γ . Then, average effort per iteration of simulation equals $N\gamma + \beta$ and average effort for L iterations equals $L(N\gamma + \beta)$. By law of large numbers this is close to the actual effort when L is large. We suppose that the overall computational budget is fixed C . We analyze the problem of determining best N, L that minimize the mean square error of the estimator $\alpha_{L,N}$ when $L(N\gamma + \beta) = C$ for large C .

Thus we consider the optimization problem

$$\min_{N,L \geq 0} \frac{\alpha_N(1 - \alpha_N)}{L} + (\alpha_N - \alpha)^2$$

subject to $L(N\gamma + \beta) = C$, as $C \rightarrow \infty$.

We address this problem by first showing that

$$\alpha_N = \alpha + c/N + o(1/N)$$

for a specified c . The objective function then reduces to finding N that minimizes

$$\frac{N\alpha(1 - \alpha)\gamma + o(N)}{C} + \frac{c^2}{N^2} + o(1/N^2).$$

It is easy to see that an optimal N for this has the form

$$N^* = \left(\frac{2c^2}{\alpha(1 - \alpha)\gamma} \right)^{1/3} C^{1/3} + o(C^{1/3}).$$

Therefore optimal L has the form

$$L^* = \left(\frac{\alpha(1 - \alpha)}{2\gamma^2 c^2} \right)^{1/3} C^{2/3} + o(C^{2/3}),$$

and the mean square error at optimal N^* equals

$$\left(\frac{c\alpha(1 - \alpha)\gamma}{2C} \right)^{2/3} + o(C^{-2/3}).$$

We now evaluate the expression $\alpha_N - \alpha$ to show that c above equals

$$\frac{1}{2} \frac{d}{du} f_Y(u) E[\sigma_\xi^2 | Y = u],$$

where σ_ξ^2 denotes the conditional variance of Z_i (conditioned on ξ) and $f_Y(\cdot)$ denotes the marginal pdf of Y (we suppress the dependence on ξ in our notation wherever it is visually convenient). Our methods here are drawn from the literature on granularity adjustment (Martin and Wilde 2002, Gordy 2004) which in turn is based on sensitivity analysis of VaR (Gourieroux, Laurent, and Scaillet 2000).

Some notation is needed for this. Let $\tilde{Z}_N = \bar{Z}^N \cdot \sqrt{N}$. Then $\alpha_N = P(Y + \tilde{Z}_N/\sqrt{N} > u)$. Let $f_\xi(\cdot)$ denote the pdf of random vector ξ . Let $f_{Y,\tilde{Z}_N}(\cdot, \cdot)$ denote the joint pdf of Y and \tilde{Z}_N . Then,

$$\alpha_N = \int_{\mathbb{R}} \int_{u-z/\sqrt{N}}^{\infty} f_{Y,\tilde{Z}_N}(y, z) dy dz,$$

so that

$$\alpha_N - \alpha = \int_{\mathbb{R}} \int_{u-z/\sqrt{N}}^u f_{Y,\tilde{Z}_N}(y, z) dy dz. \quad (1)$$

Using the Taylor's series expansion of the pdf in (1),

$$f_{Y,\tilde{Z}_N}(y, z) = f_{Y,\tilde{Z}_N}(u, z) + (y - u)f'_{Y,\tilde{Z}_N}(u, z) + o(y - u),$$

where $f'_{Y,\tilde{Z}_N}(y, z)$ is the derivative of the bivariate density with respect to y , we get

$$\begin{aligned} \alpha_N - \alpha &= \int_{\mathbb{R}} \frac{z}{\sqrt{N}} f_{Y,\tilde{Z}_N}(u, z) dz \\ &\quad - \int_{\mathbb{R}} \frac{z^2}{2N} f'_{Y,\tilde{Z}_N}(u, z) dz + o(1/N). \end{aligned}$$

Technical conditions on the smoothness of f are imposed to ensure that the remainder term has the form $o(1/N)$.

Note that the first term above equals

$$\frac{f_Y(u)}{\sqrt{N}} E[\tilde{Z}_N | Y = u]$$

which equals zero since

$$E[\tilde{Z}_N | Y = u] = E[E[\tilde{Z}_N(\xi) | Y(\xi) = u, \xi]] = 0.$$

We now show that $\int_{\mathbb{R}} y^2 f'_{Y,\tilde{Z}_N}(u, z) dz$ equals

$$\frac{d}{du} f_Y(u) E[\sigma_\xi^2 | Y = u].$$

This follows as:

$$\begin{aligned} \int_{\mathbb{R}} z^2 f'_{Y,\tilde{Z}_N}(u, y) dy &= \frac{d}{du} \int_{\mathbb{R}} z^2 f_{Y,\tilde{Z}_N}(u, z) dz \\ &= \frac{d}{du} f_Y(u) E[\tilde{Z}_N^2 | Y = u] \\ &= \frac{d}{du} f_Y(u) E[E[\tilde{Z}_N^2 | \xi] | Y = u] \\ &= \frac{d}{du} f_Y(u) E[\sigma_\xi^2 | Y = u]. \end{aligned}$$

It can also be seen that

$$\begin{aligned} \int_{\mathbb{R}} z^2 f'_{Y,\tilde{Z}_N}(u, z) dz &= f_Y(u) \int_{\mathbb{R}} z^2 \frac{f'_{Y,\tilde{Z}_N}(u, z)}{f_{Y,\tilde{Z}_N}(u, z)} \frac{f_{Y,\tilde{Z}_N}(u, z)}{f_Y(u)} dz \\ &= f_Y(u) E \left[\tilde{Z}_N^2 \frac{f'_{Y,\tilde{Z}_N}(Y, \tilde{Z}_N)}{f_{Y,\tilde{Z}_N}(Y, \tilde{Z}_N)} \Big| Y = u \right] \\ &= f_Y(u) E \left[\sigma_\xi^2 \frac{f'_{Y,\tilde{Z}_N}(Y, \tilde{Z}_N)}{f_{Y,\tilde{Z}_N}(Y, \tilde{Z}_N)} \Big| Y = u \right]. \end{aligned}$$

The latter two representations have the advantage that they may be estimated via Monte-Carlo simulation whenever the expressions for f_{Y,\tilde{Z}_N} and f'_{Y,\tilde{Z}_N} are available or can be cheaply estimated.

In the conference, we present simulation methods for estimating c and the numerical results demonstrating the effectiveness of such estimates in minimizing the mean square error of resultant performance measures.

3.2 Inner Step Simulation: Method B

We now discuss Method B to generate inner step simulations. Here, each CDO's mark-to-market price at the portfolio horizon is estimated by conducting independent simulations. As mentioned earlier, we conduct $p_i N$ simulations for CDO i where each $p_i \geq 0$ and $\sum_{i \leq K} p_i = 1$. Suppose that the average effort to generate a single such simulation is γ_i . Then, total inner step simulation effort equals $N\gamma$ where

$$\gamma = \sum_{i \leq K} p_i \gamma_i.$$

The resulting sample from an iteration of outer and inner step simulation is

$$Y(\xi) + \sum_{i=1}^K \frac{1}{p_i N} \sum_{j=1}^{p_i N} Z_{ij}(\xi)$$

where $Z_{ij}(\xi)$ denotes the sample from simulation j for CDO position i , again normalized to zero mean. (Here we ignore minor technicalities associated with $p_i N$ not being an integer.)

The analysis to compute the mean square error proceeds exactly as in Section 3.1.1. The resultant c in this setting can be seen to equal

$$\frac{1}{2} \frac{d}{du} f_Y(u) \sum_{i=1}^K \frac{1}{p_i} E[\sigma_{i,\xi}^2 | Y = u] \quad (2)$$

where $\sigma_{i,\xi}^2$ denotes the variance of $Z_{ij}(\xi)$ conditioned on ξ . Recall that the mean square error at optimal N^* equals

$$\left(\frac{c\alpha(1-\alpha)\gamma}{2C} \right)^{2/3} + o\left(\frac{1}{C^{2/3}}\right).$$

We now consider the problem of determining approximation to optimal $(p_i \geq 0 : i \leq K)$. Note that u and hence α and C are fixed, therefore it is reasonable to search for $(p_i \geq 0 : i \leq K)$ that minimize the product $c\gamma$. Recall that $\gamma = \sum_{i \leq K} p_i \gamma_i$. Typically, the exact expression for

$$\frac{d}{du} E[\sigma_{i,\xi}^2 | Y = u]$$

may not be known. However, in many portfolio credit risk models, for a reasonable class of parameters, it is possible to express $E[\sigma_{i,\xi}^2 | Y = u]$ as approximately $\delta_i g(u)$ for $\delta_i \geq 0$ and a smooth non-negative function $g(u)$. In CreditRisk+, for example, the conditional variance $E[\sigma_{i,\xi}^2 | Y = u]$ is equal to a constant times u plus terms proportional to u^2 (Gordy 2004). When such a relationship holds, (2) may be re-expressed as

$$\sum_{i=1}^K \frac{1}{p_i} E[\sigma_{i,\xi}^2 | Y = u] \cdot \frac{d}{du} \left(f_Y(u) + \frac{g'(u)}{g(u)} \right).$$

Therefore, under these circumstances a good solution to minimizing the mean-square error corresponds to finding $(p_i \geq 0 : i \leq K)$ that minimize

$$\sum_{i \leq K} p_i \gamma_i \sum_{j=1}^K \frac{1}{p_j} E[\sigma_{j,\xi}^2 | Y = u],$$

subject to $\sum_{i \leq K} p_i = 1$.

Since the terms p_i and p_j appear as ratios in the objective, the constraint $\sum_{i \leq K} p_i = 1$ simply involves normalizing any solution of the unconstrained problem. From the first order conditions, it is easy to see that the solution to this

problem $(p_i^* : i \leq K)$ corresponds to:

$$p_i^* = \frac{\sqrt{E[\sigma_{i,\xi}^2 | Y = u] / \gamma_i}}{\sum_{j \leq K} \sqrt{E[\sigma_{j,\xi}^2 | Y = u] / \gamma_j}}.$$

This is intuitive as one expects that higher computation resources be allocated to a position with higher variability and lower computational effort. This is captured by $E[\sigma_{j,\xi}^2 | Y = u]$ in the numerator and γ_i in the denominator.

3.3 Estimating Value-At-Risk

We now consider the problem of efficient estimation of VaR for Y . For a target insolvency probability α , VaR is the value y_α given by

$$y_\alpha = \inf\{y : P(Y \leq y) \geq 1 - \alpha\}$$

We consider Method A for the inner step of simulation. The extension to Method B is straightforward and is omitted.

As before, our simulation generates samples

$$(Y_1(\xi_1) + \bar{Z}_1^N(\xi_1), \dots, (Y_L(\xi_L) + \bar{Z}_L^N(\xi_L))). \quad (3)$$

Let U_1, \dots, U_L denote the sequence in (3) in descending order. Then, $\hat{y}_\alpha = U_{\lceil \alpha L \rceil}$ provides an estimate of y_α , where $\lceil a \rceil$ denotes the integer ceiling of the real number a . Our interest is in characterizing the mean square error $E[(\hat{y}_\alpha - y_\alpha)^2]$ and then minimizing it.

Let \tilde{y}_α denote the α quantile corresponding to the random variable $Y(\xi) + \bar{Z}^N(\xi)$. We decompose the mean square error as

$$\begin{aligned} E[(\hat{y}_\alpha - y_\alpha)^2] &= E[(\hat{y}_\alpha - E[\hat{y}_\alpha] + E[\hat{y}_\alpha] - \tilde{y}_\alpha + \tilde{y}_\alpha - y_\alpha)^2] \\ &= \sigma^2(\hat{y}_\alpha) + (E[\hat{y}_\alpha] - \tilde{y}_\alpha)^2 + (\tilde{y}_\alpha - y_\alpha)^2 \\ &\quad - 2(E[\hat{y}_\alpha] - \tilde{y}_\alpha)(\tilde{y}_\alpha - y_\alpha). \end{aligned} \quad (4)$$

Treating \hat{y}_α as an estimator of \tilde{y}_α , we have the well-known results

$$E[\hat{y}_\alpha] = \tilde{y}_\alpha - \frac{\alpha(1-\alpha)f'_{Y+\bar{Z}^N}(\tilde{y}_\alpha)}{2(L+2)f_{Y+\bar{Z}^N}(\tilde{y}_\alpha)^3} + O(1/L^2) \quad (5)$$

$$\sigma^2(\hat{y}_\alpha) = \frac{\alpha(1-\alpha)}{(L+2)f_{Y+\bar{Z}^N}(\tilde{y}_\alpha)^2} + O(1/L^2), \quad (6)$$

where $f_{Y+\bar{Z}^N}(\cdot)$ denotes the pdf of rv $Y + \bar{Z}^N$ (see Chen 2002, Law and Kelton 2000).

Note that \tilde{y}_α is a function of N and $\tilde{y}_\alpha \rightarrow y_\alpha$ as $N \rightarrow \infty$. We now evaluate $\tilde{y}_\alpha - y_\alpha$ more precisely by way

of the Taylor's series expansion,

$$\alpha = P(Y + \bar{Z}^N > \tilde{y}_\alpha) = P(Y + \bar{Z}^N > y_\alpha) - (\tilde{y}_\alpha - y_\alpha)f_{Y+\bar{Z}^N}(y_\alpha) + o(\tilde{y}_\alpha - y_\alpha).$$

From our analysis earlier on $\alpha_N - \alpha$, we have

$$P(Y + \bar{Z}^N > y_\alpha) = P(Y > y_\alpha) + \tilde{c}/N + o(1/N)$$

where $\tilde{c} = \frac{1}{2} \frac{d}{du} f_Y(u) E[\sigma_\xi^2 | Y = u] |_{u=y_\alpha}$, and since $P(Y > y_\alpha) = \alpha$, it follows that

$$\tilde{y}_\alpha - y_\alpha = \frac{\tilde{c}}{N f_{Y+\bar{Z}^N}(y_\alpha)} + o(\tilde{y}_\alpha - y_\alpha).$$

Under mild technical conditions, $f_{Y+\bar{Z}^N}(y_\alpha) \rightarrow f_Y(y_\alpha)$, therefore it can be seen that

$$\tilde{y}_\alpha - y_\alpha = \frac{\tilde{c}}{N f_Y(y_\alpha)} + o(1/N). \tag{7}$$

Recall that C denotes our overall computational budget, and $L(N\gamma + \beta) = C$. Plugging this into (5) and imposing mild regularity conditions, we find $(E[\hat{y}_\alpha] - \tilde{y}_\alpha)^2$ is $O(N^2/C^2)$. Similarly, we find $(E[\hat{y}_\alpha] - \tilde{y}_\alpha)(\tilde{y}_\alpha - y_\alpha)$ is $O(1/C)$, and hence the mean square error $E[(\hat{y}_\alpha - y_\alpha)^2]$ equals

$$\frac{\alpha(1-\alpha)(N\gamma + \beta)}{C f_Y(y_\alpha)^2} + O(N^2/C^2) + \frac{\tilde{c}^2}{N^2 f_Y(y_\alpha)^2} + O(1/C).$$

This can be seen to be minimized at

$$N^* = \left(\frac{2\tilde{c}^2}{\alpha(1-\alpha)\gamma} \right)^{1/3} C^{1/3} + o(C^{1/3}).$$

Therefore optimal L has the form

$$L^* = \left(\frac{\alpha(1-\alpha)}{2\gamma^2\tilde{c}^2} \right)^{1/3} C^{2/3} + o(C^{2/3}).$$

These values are identical up to terms of size $o(C^{2/3})$ to the optimal values for estimating $P(Y > u)$ derived in the previous section when $u = y_\alpha$.

4 CONCLUSION AND ONGOING RESEARCH

In this paper we considered the problem of efficient estimation of the probability of large losses and the value-at-risk for a portfolio comprising many CDO tranches. We proposed a two stage simulation procedure and discussed computational resource allocation to each stage so as to minimize the resulting mean square error. Our analysis was conducted for the reduced form models increasingly popular

for modeling credit risk. It is evident from our analysis that with minor adjustments, it remains valid more generally for other approaches to credit risk (e.g., structural models) as well as market risk modeling. As we noted in the introduction, the simulation results describing the efficacy of the proposed resource allocation strategies would be presented in the conference.

We have not considered importance sampling techniques in this paper. Combining our computational resource allocation methodology with existing IS techniques should be entirely straightforward and effective.

Our ongoing research considers dynamic allocation of N in the simulation inner step. Note that in estimating $P(Y > u)$ as in Section 3, in the inner simulation step if the average of the generated samples $Y + Z_1, \dots, Y + Z_N$, for small value of N , is seen to be much smaller than u or much larger than u , it may be waste of effort to generate many more samples in the inner simulation step. However, if this average is close to u , it makes sense to generate many samples to reduce the bias in the estimator. Another topic of ongoing research concerns the asymptotic behavior of our optimal allocation as the number of positions K or the number of obligors m grow large.

ACKNOWLEDGMENTS

This work was conducted while both authors were visiting the Indian School of Business. The opinions expressed here are our own, and do not reflect the views of the Board of Governors or its staff.

REFERENCES

Chen, E. 2002. Two phase quantile estimation. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yucesan, C. H. Chen, J. L. Snowden, and J. M. Charnes. Piscataway, NJ: IEEE Press.

Duffie, D., and N. Gârleanu. 2001, January-February. Risk and valuation of collateralized debt obligations. *Financial Analysts Journal* 57 (1): 41-59.

Duffie, D., and K. J. Singleton. 1999. Modeling term structures of defaultable bonds. *Review of Financial Studies* 12:687-720.

Glasserman, P. 2004. *Monte carlo methods in financial engineering*. New York: Springer-Verlag.

Gordy, M. B. 2004. Granularity adjustment in portfolio credit risk measurement. In *Risk Measures for the 21st Century*, G. P. Szegö eds. John Wiley & Sons.

Gourieroux, C., J. Laurent, and O. Scaillet. 2000. Sensitivity analysis of values at risk. *Journal of Empirical Finance* 7:225-245.

Lando, D., R. A. Jarrow, and F. Yu. 2005. Default risk and diversification: Theory and empirical implications. *Mathematical Finance* 15 (1): 1-26.

- Law, A., and W. Kelton. 2000. *Simulation modelling and analysis*. 3rd ed. New York: McGraw-Hill.
- Li, D. X. 2000, March. On default correlation: A copula approach. *Journal of Fixed Income* 9:43-54.
- Martin, R., and T. Wilde. 2002, November. Unsystematic credit risk. *Risk* 15 (11): 123-128.
- Willemann, S. 2005, Fall. An evaluation of the base correlation framework for synthetic CDOs. *Journal of Credit Risk* 1 (4).

AUTHOR BIOGRAPHIES

MICHAEL GORDY is a Senior Economist at the Federal Reserve Board in Washington, DC. His primary area of research concerns the design, calibration and computation of models of portfolio credit risk, and on the adaptation of these models to banking regulation. His e-mail address is <michael.gordy@frb.gov> and his web page is <<http://www.federalreserve.gov/research/staff/gordymichaelx.htm>>.

SANDEEP JUNEJA is an Associate Professor in the School of Technology and Computer Science at the Tata Institute of Fundamental Research in India. His research interests include analysis of Monte Carlo methods in finance. His e-mail address is <juneja@tifr.res.in> and his web page is <www.tcs.tifr.res.in/~sandeepj>.