Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

EFFICIENT MONTE CARLO CVA ESTIMATION

Samim Ghamami

Federal Reserve Board
Washington, D.C. 20551
Center for Risk Management Research
University of California, Berkeley, CA 94720
USA

Bo Zhang

IBM Thomas J. Watson Research Center
Yorktown Heights, NY 10598
USA

ABSTRACT

This paper presents an overview of the efficient Monte Carlo counterparty credit risk (CCR) estimation framework recently developed by Ghamami and Zhang (2014). We focus on the estimation of credit value adjustment (CVA), one of the most widely used and regulatory-driven counterparty credit risk measures. Our proposed efficient CVA estimators are developed based on novel applications of well-known mean square error (MSE) reduction techniques in the simulation literature. Our numerical examples illustrate that the efficient estimators outperform the existing crude estimators of CVA substantially in terms of MSE.

1 INTRODUCTION

Counterparty credit risk (CCR) is the risk that a party to a derivative contract may default prior to the expiration of the contract and fail to make the required contractual payments. We refer the reader to Canabarro and Duffie (2003) and Gregory (2010) for the basic CCR concepts and definitions. Counterparty credit risk has been widely considered as one of the key drivers of the 2007-08 credit crisis, and it has become one of the main focuses of the major global and U.S. regulatory frameworks (Basel III and the Dodd-Frank Act of 2009-10; see, e.g., Bohme et al. (2011)). It is well known that pricing and measuring counterparty credit risk is computationally extremely intensive; large financial institutions invest large amounts of resources developing and maintaining Monte Carlo simulation “engines” to manage their counterparty risk (Pykhtin and Zhu (2006), Gregory (2010), and Canabarro and Duffie (2003)). While various aspects of counterparty credit risk have been subject of extensive research post 2007-08 financial crisis, statistical efficiency of the CCR estimators has received little attention in the literature. Considering various counterparty credit risk measures, Ghamami and Zhang (2014) are the first to develop an efficient Monte Carlo framework for pricing and measuring CCR. This paper presents an overview of the efficient Monte Carlo counterparty credit risk (CCR) estimation framework recently developed by Ghamami and Zhang (2014). We focus on the estimation of credit value adjustment (CVA). Ghamami and Zhang introduce their efficient Monte Carlo framework by focusing on a different CCR measure, namely, expected positive exposure (EPE). Due to the similarities in the mathematical formulation of EPE and CVA, they address efficient Monte Carlo CVA estimation indirectly through efficient EPE estimation. The attention that CVA has received post financial crisis is unprecedented; this paper gives a detailed treatment of efficient Monte Carlo CVA estimation that can be read independently from Ghamami and Zhang (2014).

CVA is defined as the difference between the risk free value of a derivatives portfolio and the counterparty default risky derivatives portfolio value (Pykhtin and Zhu (2007)), and it has become one of the main focuses of the Basel III; e.g., derivative dealers are required to calculate CVA *capital charges* for each of their counterparties on a frequent basis. Efficiency criteria under consideration are variance, bias, and computing time of the Monte Carlo estimators. Our proposed Monte Carlo estimators of CVA outperform the existing crude estimators substantially in terms of mean square error (MSE).

1.1 Problem formulation

Consider the setting where a financial institution holds a portfolio of derivative contracts with another financial institution, namely, its counterparty. Counterparty credit exposure (Canabarro and Duffie (2003)) denoted by V, of a financial institution against the counterparty, is the larger of zero and the market value of the portfolio of derivatives contracts the financial institution holds with this counterparty.

Contract level credit exposure at time $t > 0$ is the maximum of a contract's market value and zero, $\max\{C_t, 0\}$, where C_t denotes the time-t value of the derivative contract. Consider a financial institution that holds a portfolio of k derivative contracts with its counterparty. Counterparty level credit exposure is

$$V_t = \sum_{i=1}^{k} \max\{C_t^i, 0\}, \tag{1}$$

where C_t^i denotes the time-t value of the i'th derivative contract in the derivatives portfolio. In practice, V_t may need to be valued differently when risk mitigants are employed. For instance, in the presence of netting agreements, credit exposure becomes, (see Gregory (2010) and the references therein for a comprehensive overview of different types of risk mitigants and netting agreements). We emphasize that our methods apply irrespective of the presence of risk mitigants and for simplicity one may assume that (1) holds for the remainder of this paper.

$$V_t = \max\left\{\sum_{i=1}^{k} C_t^i, 0\right\}. \tag{2}$$

Let τ, a positive random variable, denote the default time of the counterparty and $T > 0$ represent the expiration time of the longest maturity derivative contract in the OTC derivatives portfolio. It can be shown that CVA, the price of the counterparty credit risk, is equal to the risk neutral expected discounted loss, i.e.,

$$\text{CVA} \equiv E[(1-R)D_\tau V_\tau \mathbf{1}\{\tau \leq T\}], \tag{3}$$

where $1\{A\}$ is the indicator of the event A, $D_t = B_0/B_t$ is the stochastic discount factor at time t, B_t is the value of the money market account at time t, and R is the financial institution's recovery rate, (see, for instance, Chapter 7 of Gregory (2010) for a derivation of this formula). When V and τ are assumed to be independent, CVA is referred to as independent CVA, which we denote by CVA_I. Wrong (right) way risk are referred to as cases where credit exposures are negatively (positively) correlated with the credit quality of the counterparty (Ghamami and Goldberg (2014), Canabarro and Duffie (2003), and Hull and White (2012)). In this paper we focus on the estimation of CVA_I. To simplify the presentation of the main idea, we suppress the dependence of the CVA on the recovery rate, R, and that on the stochastic discount factor by assuming zero short rate. Let F denote the cumulative distribution function of τ, which is assumed to be known (market observable) from, for instance, credit default swap spreads of the counterparty (e.g., Hull and White (2012)). Independent CVA can be written as follows,

$$\text{CVA}_I = E\left[E[V_\tau \mathbf{1}\{\tau \leq T\}|\tau]\right] = \int_0^T E[V_t]dF_t, \tag{4}$$

where the last equality follows from conditioning on τ and the independence of V and τ. The standard crude Monte Carlo approach, which is being used in practice, estimates the independent CVA based on a time-discretized summation approximation of the integral in (4) and Monte Carlo estimation of expected exposures, $E[V_t]$. We are to introduce efficient Monte Carlo CVA estimation schemes. Specifically, we are interested in efficient Monte Carlo estimation of the Riemann-Stieltjes integral of the mean of a stochastic process which depends on many other stochastic processes in a complicated way; note that under either (1) or (2), V_t is determined by C_t^i, $i = 1, ..., k$, whose dynamics further depend on possibly multiple underlying stochastic processes. The order of k can be more 1000 for each portfolio in practice. Each C_t^i can depend on the evolution of multiple *risk factors* such as interest rates and commodity prices.

Section 2 summarizes the common features of the Monte Carlo CCR framework widely used by financial institutions and introduces the notion of *Marginal Matching*, which enables us to define and differentiate the two widely used CCR sampling methods, *Path Dependent Simulation* (PDS) and *Direct Jump to Simulation* (DJS) date. Practitioners often choose either of the sampling methods arbitrarily. Section 3 introduces an efficient Monte Carlo framework for estimating CVA_I. We illustrate that PDS and DJS-based CCR estimators have drastically different MSE. This is an important result that could have broader applicability. Our numerical examples indicate that employing our Monte Carlo CVA estimation schemes leads to substantial MSE reduction.

2 MONTE CARLO COUNTERPARTY CREIDT RISK ESTIMATION

A typical Monte Carlo counterparty risk engine estimates various CCR measures by sampling from the credit exposure process on a time grid, $0 < t_1 < ... < t_n = T$, where T denotes the maturity of the longest transaction in a portfolio and t_i's are sometimes referred to as *valuation points*. Set $V_i \equiv V_{t_i}$.

In what follows we first summarize the simulation of the credit exposure process. Then, we introduce the notion of Marginal Matching in sampling from the time evolution of the credit exposure process.

2.1 Simulating the Credit Exposure Process

Suppose that credit exposure is a stochastic process $\{V_t \ ; t \geq 0\}$ defined on a given filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{0 \leq t \leq \infty}, P)$. Given (1) and (2), V_t can be viewed as a function of the stochastic processes that drive the values of the derivative contracts, $C_t^1, ..., C_t^k$. In risk management, these underlying stochastic processes are usually referred to as risk factors, e.g., interest rates, commodity prices, and equity prices. To generate a Monte Carlo realization of V_t, for a fixed $t > 0$, first, the underlying risk factors should be sampled from up to time $t > 0$. Next, given the Monte Carlo realization of the risk factors up to time $t > 0$, the derivative contracts should be valued. This two-step procedure leads to a single Monte Carlo realization of V_t. For risk management applications, one often chooses the physical probability measure in the first step and the risk-neutral measure in the second. However, since CVA is often viewed as the market price of counterparty credit risk, risk-neutral measure is usually used in both steps. Depending on the complexity of the payoff function of the derivative contracts, the valuation step could take straightforward Black-Scholes-type analytical calculations, or it could demand approximations that depending on the desired level of accuracy might be computationally intensive. For instance, the valuation step could involve a second layer of Monte Carlo (Gordy and Juneja (2010), Broadie et al. (2011), and Chapter 8 of Glasserman (2004)).

2.2 Marginal Matching

Let $X = (X_1, ..., X_n)$ denote a random vector with distribution function F_X. Let $\omega_X \equiv (E[h_1(X_1)], ..., E[h_n(X_n)])$ for some functions $h_1, ..., h_n$. Let $\theta_X \equiv g(\omega_X)$ for a function g that maps ω_X from $\mathbb{R}^n$ to $\mathbb{R}$. Two simple examples are $\sum_{i=1}^n E[h(X_i)]$ and $\max\{E[h(X_1)], ..., E[h(X_n)]\}$, for which θ_X is defined based on the marginal distribution of (functions of) $X_1, ..., X_n$. Let $Y = (Y_1, ..., Y_n)$ denote another random vector with distribution function F_Y such that, $X \neq^d Y$ and yet $X_i =^d Y_i$ for all $i = 1, ..., n$, where $=^d$ denotes "equal in distribution". Simply note that since the marginal distributions of X and Y *match*, $\theta_X = \theta_Y$. Now, suppose that θ_X is to be estimated with Monte Carlo simulation. Given $X_i =^d Y_i$, samples can be drawn from F_X or F_Y. Let $\hat{\theta}_{X,m}$ and $\hat{\theta}_{Y,m}$ denote Monte Carlo estimators of θ_X based on m simulation runs when samples are drawn from F_X and F_Y, respectively. Obviously, since $X \neq^d Y$, we have $\hat{\theta}_{X,m} \neq^d \hat{\theta}_{Y,m}$, and so between $\hat{\theta}_{X,m}$ and $\hat{\theta}_{Y,m}$, i.e., when deciding on whether to sample from F_X or F_Y, the estimator with a lower MSE should be chosen.

Stochastic Models of the Risk Factors Let $\{S_t\ ;t \geq 0\}$, representing the dynamics of a risk factor, denote a stochastic process defined on a given filtered probability space, $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{0 \leq t \leq \infty}, P)$. Similar to the setting of Ghamami and Zhang (2014), we assume that $\{S_t\ ;t \geq 0\}$ is in the following class: a Gauss-Markov process (Chapter 5 of Karatzas and Shreve (1991) or Chapter 3 of Glasserman (2004)), a Geometric Brownian motion (GBM), or a square-root diffusion. Many of the widely used continuous time stochastic processes in finance and economics are in this class. Consider the finite dimensional distribution of $\{S_t\ ;t \geq 0\}$ on a time grid, $t_1, \ldots, t_n$ and set $S_i \equiv S_{t_i}$. Suppose that $S = (S_1, \ldots, S_n)$ can be sampled from *exactly* in the sense that the distribution of the simulated S is precisely that of the $\{S_t\ ;t \geq 0\}$ process at times $t_1, \ldots, t_n$; examples are Brownian motion, Ornstein-Uhlenbeck processes, GBM, and the square-root diffusion whose simulations involve generating positively correlated normal random variables. Let $\tilde{S} = (\tilde{S}_1, \ldots, \tilde{S}_n)$ denote a random vector for which $\tilde{S} \neq^d S$ but $\tilde{S}_i =^d S_i$ for all $i = 1, \ldots, n$ and $\text{cov}(\tilde{S}_i, \tilde{S}_j) = 0$ for all $i \neq j$. That is, simulation of $\tilde{S}_1, \ldots, \tilde{S}_n$ can be done by generating n uncorrelated or simply independent normal random variables.

PDS Sampling versus DJS Sampling In the CCR literature when counterparty risk measures are estimated based on sampling from the finite-dimensional distributions of the underlying risk factors, the sampling is referred to as PDS sampling. When the notion of marginal matching is used, the sampling is referred to as DJS. In Monte Carlo estimation of CCR measures, PDS and DJS sampling have been widely considered equivalent (Pykhtin and Zhu (2006)). We have also observed that practitioners often choose either of the sampling methods arbitrarily. We differentiate DJS and PDS in terms of the mean square error of the estimators of CVA.

3 EFFICIENT MONTE CARLO ESTIMATION OF INDEPENDENT CVA

Consider a time grid, $0 \equiv t_0 < t_1 < \ldots < t_n \equiv T$, with a fixed n. Set $\Delta F_i \equiv F(t_i) - F(t_{i-1})$ and $V_i \equiv V_{t_i}$, $i = 1, \ldots, n$. Let $\hat{\theta}_{b,m,n,k}$ denote a class of Monte Carlo estimators of CVA_I defined as follows,

$$\hat{\theta}_{b,m,n,k} \equiv \sum_{i=1}^{n} \bar{V}_i \Delta F_i,$$

where $\bar{V}_i \equiv \sum_{j=1}^{m} V_{ij}/m$ and $V_{i1}, \ldots, V_{im}$ represent the m simulation samples at valuation point t_i. The subscript b refers to the biased nature of the estimators, and the subscript k could take p and d, referring to PDS and DJS based simulation of the credit exposure process, respectively.

As mentioned in Section 2.1, simulating the credit exposure process involves sampling from the underlying risk factors. Hereafter, PDS and DJS-based simulations of the credit exposure process refer to the cases where the underlying risk factors are sampled from based on their finite dimensional distributions (PDS sampling) and based on the notion of marginal matching (DJS sampling), respectively. Note that,

$$\text{MSE}(\hat{\theta}_{b,m,n,k}) = \text{Var}\left(\sum_{i=1}^{n} \bar{V}_i \Delta F_i\right) + \left(\sum_{i=1}^{n} E[\bar{V}_i] \Delta F_i - \int_0^T E[V_t] dF_t\right)^2.$$

We assume that Monte Carlo realizations of V_i are unbiased estimates of $E[V_i]$, $i = 1, \ldots, n$. This implies that the bias part of the MSE of $\hat{\theta}_{b,m,n,k}$ is not affected by the choice of the sampling method (PDS or DJS) and is only due to the time-discretization of the integral of expected exposures in the definition of CVA. Our framework remains applicable in the absence of this assumption as shown in the Appendix H of Ghamami and Zhang (2014).

In Section 3.1, we assume that n is fixed and compare the efficiency of $\hat{\theta}_{b,m,n,p}$ and $\hat{\theta}_{b,m,n,d}$ in terms of variance and computing time for both *path independent* and *path dependent* derivatives.

3.1 Comparing PDS and DJS-based Estimation of CVA_I

Suppose that the credit exposure process, V, is defined on a given filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{0 \leq t \leq \infty}, P)$, where $(\mathscr{F}_t)_{0 \leq t \leq \infty}$ denote the filtration generated by the underlying risk factors. Consider the setting where

V denotes the contract level exposure and a financial institution takes a position in a maturity-T derivative contract with its counterparty. Let Π_T denote the payoff function of the derivative contract. It is well known from martingale pricing that

$$C_t = M_t E\left[\frac{\Pi_T}{M_T}|\mathscr{F}_t\right], \tag{5}$$

where M is a numeraire. Transactions between the financial institution and its counterparty for which $V_t = \max\{C_t, 0\} = C_t$ for all $0 < t \leq T$ are referred to as *unilateral transactions*, e.g. the financial institution takes a long position in a call option with its counterparty. Transactions for which $V_t = \max\{C_t, 0\} \neq C_t$ for some $0 < t \leq T$ are referred to as *bilateral transactions*, e.g. an interest rate swap between the financial institution and its counterparty.

The following simple example reviews simulation of the exposure process under PDS and DJS. Suppose that $\{S_t \ ; t \geq 0\}$ is a GBM, $S_t = S_0 e^{X_t}$, and $\{X_t \ ; t \geq 0\}$ is a Brownian motion with drift μ and volatility σ. Consider a unilateral transaction. Note that $V_t = C_t = M_t E\left[\frac{\Pi_T}{M_T}|S_t\right] \equiv f(S_t)$. That is, credit exposure is considered as a function of the risk factor.Consider the time grid, $0 \equiv t_0 < t_1 < ... < t_n \equiv T$ and let $V_i \equiv V_{t_i}$. Set $\theta \equiv \sum_{i=1}^n E[V_i]\Delta F_i$. Recall that $\hat{\theta}_{b,m,n,k} = \sum_{i=1}^n \bar{V}_i \Delta F_i$, where $\bar{V}_i$ is the m-simulation-run average of $V_{i1}, ..., V_{im}$. With $V_i = f(S_i)$ and $S_i = S_0 e^{X_i}$, Monte Carlo estimation of θ requires sampling from the multivariate normal random vector, $X = (X_1, ..., X_n)$. This is the so-called PDS sampling method. An alternative sampling method, using the notion of marginal matching, is to sample from the multivariate normal random vector, $Y = (Y_1, ..., Y_n)$, whose components are uncorrelated but marginal distributions match those of X. This is the so-called DJS method. To be more specific, in DJS sampling, S_i is generated from time zero. That is, generate Y_i, a normal random variable with mean μt_i and variance $\sigma^2 t_i$, and set $S_i = S_0 e^{Y_i}$. In PDS sampling, V_i's are sampled based on generating the sample path of the GBM sequentially at $i = 1, ..., n$. That is, to generate a realization of V_i, S_i is generated given the previously generated value of S_{i-1}. More specifically, to sample from S_i generate $\tilde{X}_i$ and set $S_i = S_{i-1} e^{\tilde{X}_i}$, where $\tilde{X}_i$ is a normal random variable with mean $\mu \Delta F_i$ and variance $\sigma^2 \Delta F_i$. Note that since for any given $t > 0$, V_t is a function of $S_t = S_0 e^{X_t}$, DJS-based simulation of the exposure process implies that $\mathrm{cov}(V_i, V_j) = 0$ for any $i \neq j$, $i, j = 1, ..., n$.

In what follows we compare the efficiency of $\hat{\theta}_{b,m,n,p}$ and $\hat{\theta}_{b,m,n,d}$ in terms of variance and computing time for *path independent* and *path dependent* derivatives. We consider unilateral and bilateral transactions in both single risk-factor and multi-risk factor settings. That is, we consider two cases: a stylized setting where $(\mathscr{F}_t)_{0 \leq t \leq \infty}$ is the filtration generated by a single risk factor; we also consider the more general multi-risk factor settings.

3.1.1 Path Independent Case

The above mentioned example shows that under DJS, $\mathrm{cov}(V_u, V_t) = 0$ for any $0 < u < t < T$. Proposition 1 and Proposition 2 consider this covariance function of the contract level credit exposure process under the PDS method for unilateral and bilateral transactions, respectively, and identify conditions under which $\mathrm{cov}(V_u, V_t) > 0$ for any $0 < u < t < T$. Condition 2 of Proposition 1 below uses the well known changes of numeraire techniques of Geman et al. (1995) for option type contracts with at most three distinct sources of randomness: stochastic short rate and a maximum of two risky assets. Well known examples of these contracts are options written on stocks or bonds, e.g. European options and exchange options.

Proposition 1 (Ghamami and Zhang (2014)) Consider the credit exposure process, $\{V_t \ ; t \geq 0\}$, defined on a given filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{0 \leq t \leq \infty}, P)$, and a T-maturity transaction between the financial institution and its counterparty that is unilateral, i.e. the credit exposure process is the price process, $V_t = C_t > 0$ for all $0 \leq t \leq T$, where C_t denotes the time-t value of the derivative contract with payoff Π_T. Then,

$$\mathrm{cov}(V_u, V_t) > 0,$$

for any $0 < u < t < T$ under any of the following conditions:

Condition 1: Numeraire is the money market account, B, with deterministic short rate, r, and Π_T is a function of $N \geq 1$ exogenously given risky assets;

Condition 2: Short rate is stochastic and the T-payoff function is a function of at most two risky assets as follows $\Pi_T = (\alpha_1 S_1(T) + \alpha_2 S_2(T))^+$, where α_1 and α_2 are any real numbers, and S_1 and/or S_2 are risky assets.

In the case of bilateral transactions for which the exposure process satisfies $V_t = \max\{C_t, 0\} \neq C_t$ for some $0 < t \leq T$, where C_t denotes the time-t value of the derivative contract with payoff function Π_T stronger assumptions are required to analytically show that $\text{cov}(V_u, V_t) > 0$ for any $0 < u < t < T$. This is shown in Proposition 2 below.

Proposition 2 (Ghamami and Zhang (2014)) Consider the credit exposure process, $\{V_t\ ; t \geq 0\}$, defined on a given filtered probability space $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{0 \leq t \leq \infty}, P)$, and a T-maturity transaction between the financial institution and its counterparty that is bilateral, i.e. the credit exposure process is the price process, $V_t = \max\{C_t, 0\} \neq C_t$ for some $0 < t \leq T$, where C_t denotes the time-t value of the derivative contract with payoff function Π_T. Then,

$$\text{cov}(V_u, V_t) > 0$$

for any $0 < u < t < T$ under the following condition:

Numeraire is the money market account, B, with deterministic short rate, r, and Π_T is a monotone function of a single risky asset whose dynamics is modeled by a GBM, a Gauss-Markov process or a square-root diffusion.

Propositions 1 and 2 identify conditions for unilateral and bilateral transactions under which the credit exposure process satisfies $\text{cov}(V_u, V_t) > 0$ for any $0 < u < t < T$. This, then, implies that

$$\text{Var}(\hat{\theta}_{b,m,n,d}) \leq \text{Var}(\hat{\theta}_{b,m,n,p}). \tag{6}$$

Note that the above inequality holds since

$$\text{Var}(\hat{\theta}_{b,m,n,d}) = \sum_{i=1}^{n} \frac{\text{Var}(V_i)\Delta F_i^2}{m} \leq \sum_{i=1}^{n} \frac{\text{Var}(V_i)\Delta F_i^2}{m} + \frac{2}{m} \sum_{i<j}\sum \text{cov}(V_i, V_j)\Delta F_i \Delta F_j = \text{Var}(\hat{\theta}_{b,m,n,p}). \tag{7}$$

3.1.2 Path Dependent Case

Suppose that V_t is time t value of a maturity-T contract, where the payoff at the time T is a function of $S_1, ...S_n$, (for instance, an arithmetic Asian option). That is, $V_i = g(S_1, ..., S_i)$, where g is a function from $\mathbb{R}^i$ to $\mathbb{R}$. The DJS sampling method is to make $V_i = g(S_1, ..., S_i)$ and $V_j = g(S_1, ..., S_j)$, $i < j$, uncorrelated random variables. That is, sample from $S_1, ..., S_i$ to generate a single realization of V_i. To generate V_j, start again from time zero, and sample from $S_1, ..., S_i, ...S_j$. Under this DJS-type sampling method, V_i and V_j become uncorrelated, $\text{cov}(V_i, V_j) = 0$. In the PDS-type sampling, given the Monte Carlo realization of V_i, to generate V_j, one uses the previously generated $S_1, ..., S_i$ and only samples from $S_{i+1}, ..., S_j$. In this case V_i and V_j are dependent.

Using conditional covariance formula and arguments similar to the ones used in the path independent case, it can be shown that $\text{cov}(V_i, V_j) > 0$, $i \neq j$. More specifically, it can be shown that $\text{cov}(V_i, V_j) > 0$ for unilateral and bilateral transactions under the first condition of Proposition 1 and Proposition 2's condition, respectively. That is, for the above mentioned covariance function to be positive, we need the numeraire money market account with deterministic short rate in the unilateral case. The bilateral case, additionally, requires monotonicity of the payoff function and its dependence on a single risk factor.

To compare the efficiency of the DJS and PDS-based estimators of θ in the path dependent case, computing time is also to be considered in parallel with variance of the estimators. In the path independent case computing time of DJS and PDS-based estimators of θ are roughly equal. More specifically, the estimator with the lower ($variance\ per\ replication \times expected\ computing\ time$) should be selected (see Glynn and Whitt (1992) for the formal formulation of this useful criterion in comparing alternative Monte Carlo estimators). Consider, for instance, arithmetic Asian options. Suppose that the computational time to calculate $\hat{\theta}_{b,m,n,k}$ is proportional to the number of random variables that are to be generated. Let $\mathrm{ct}(\hat{\theta}_{b,m,n,k})$ denote the computational effort associated with $\hat{\theta}_{b,m,n,k}$. Note that,

$$\frac{\mathrm{ct}(\hat{\theta}_{b,1,n,d})}{\mathrm{ct}(\hat{\theta}_{b,1,n,p})} \approx n \quad \text{and} \quad \frac{\mathrm{Var}(\hat{\theta}_{b,1,n,p})}{\mathrm{Var}(\hat{\theta}_{b,1,n,d})} \approx n. \tag{8}$$

To see why (8) holds note that to calculate $\hat{\theta}_{b,1,n,d}$, $\frac{n(n+1)}{2}$ random variables are to be generated while $\hat{\theta}_{b,1,n,p}$ requires generating n random variables. Also, note that as can be seen from (7), the variance of the PDS-based estimator is of order n^2 because of the covariance terms while the DJS-based estimator has a variance of order n . So, $\hat{\theta}_{b,m,n,d}$ and $\hat{\theta}_{b,m,n,p}$ have a similar performance for fixed and sufficiently large n. PDS and DJS-based estimators of other derivatives whose payoff depends on the path in a different form can be compared similarly.

3.1.3 Summary of Section 3.1

To compare the DJS and PDS-based estimators CVA_I (viewed as weighted sums of expected exposures) variance and computing time of the Monte Carlo estimators are considered. The DJS method induces zero covariance between any two distinct time points of the simulated credit exposure process. So, it remains to look at this covariance function for the credit exposure process under the PDS method. When the dynamics of the risk factors are modeled by the class of continuous time stochastic processes considered in this paper, the covariance function of the credit exposure process under the PDS method becomes positive under conditions of Proposition 1 and 2 for unilateral and bilateral path independent derivatives transactions, respectively. Similar results hold for path dependent derivatives. That is, under conditions of Proposition 1 and 2, DJS-based estimators of CVA outperform the PDS-based estimators in terms of variance. For path independent derivatives PDS and DJS-based computing times are roughly equal. So, we recommend that the counterparty credit risk modeler uses DJS for path independent derivatives. For path dependent derivatives, DJS-based estimators usually have larger computing times. The criterion introduced above considers the computing time in parallel with variance. There are widely traded path dependent derivatives for which PDS and DJS-based estimators of CVA perform approximately equally. For instance, for arithmetic Asian options the DJS and PDS-based estimators of CVA perform similarly. There are contracts whose payoff function does not exactly match the mathematical conditions of Proposition 1 and 2. For those contracts, a small simulation study could compare the variance of the DJS and PDS-based estimators of CVA.

Hereafter, we assume that the credit exposure process V satisfies $cov(V_u,V_t) = 0$ and $cov(V_u,V_t) > 0$ when simulated under the DJS and PDS methods, respectively, for any $0 < u < t$.

3.2 Efficient Monte Carlo CVA_I Estimation: Biased Estimators

In this subsection, we are interested in finding the approximately optimal number of valuation points, n, and the number of simulation runs at each valuation point, m, that minimize $\mathrm{MSE}(\hat{\theta}_{b,m,n,k})$,

$$\mathrm{MSE}(\hat{\theta}_{b,m,n,k}) = \mathrm{Var}(\hat{\theta}_{b,m,n,k}) + (E[\hat{\theta}_{b,m,n,k}] - \mathrm{CVA}_I)^2.$$

given a fixed computational budget, denoted by s, that is proportional to, mn. Also, $k = p$, and d refer to PDS and DJS-based simulation of the credit exposure process on a time grid $0 \equiv t_0 < t_1 < ... < t_n \equiv T$. That is, as shown in the previous section, under PDS sampling and DJS sampling, $\mathrm{cov}(V_i,V_j) > 0$ and $\mathrm{cov}(V_i,V_j) = 0$, respectively, for any $i \neq j$, $i, j = 1,...,n$.

To formulate and solve this optimization problem, we specify the order of the variance and bias of the Monte Carlo estimator of CVA_I, $\hat{\theta}_{b,m,n,k}$. Note that from basic results on endpoint Reimann sum approximation of integrals, time-discretization bias is of order $1/n$. We are not concerned with deriving sharp estimates of the orders of variance. In fact, our numerical examples indicate that choosing approximately optimal m and n using even very rough approximates for the orders of variance and bias leads to substantial MSE reduction compared to industry practice.

Consider the equidistant time grid, $0 \equiv t_0 < t_1 < \ldots < t_n \equiv T$, where $t_i - t_{i-1} = \Delta_i \equiv \Delta = \frac{1}{n}$. Note that $\Delta F_i \leq \Delta_i \sup_{t_{i-1} \leq x \leq t_i} f(x)$, where f denotes the density of the counterparty's default time, τ. That is, we can set $\Delta F_i = O(\frac{1}{n})$. We assume that $E[V_t^2] < \infty$ for all $t \in [0,T]$. First, we note that

$$\text{Var}(\hat{\theta}_{m,n,d}) = O(\frac{1}{mn}). \tag{9}$$

To see this, consider $M > 0$ such that $E[V_t^2] \leq M$ for $t \in (0,T]$. Note that,

$$\text{Var}(\hat{\theta}_{m,n,d}) = \Delta^2 \sum_{i=1}^{n} \frac{\text{Var}(V_i)}{m} \leq (\frac{T}{n})^2 \sum_{i=1}^{n} \frac{E(V_i^2)}{m} \leq \frac{MT^2}{mn}.$$

Now, consider the variance of the PDS-based estimator, $\hat{\theta}_{m,n,p}$,

$$\text{Var}(\hat{\theta}_{m,n,p}) = \Delta^2 \sum_{i=1}^{n} \frac{\text{Var}(V_i)}{m} + \Delta^2 \frac{2}{m} \sum_{i<j} \sum \text{cov}(V_i, V_j).$$

As shown before, the first term above is $O(\frac{1}{mn})$. Also, under PDS sampling, the credit exposure process is simulated according to its finite dimensional distributions for which the covariance terms are positive. So, the second term is $O(\frac{1}{m})$. This gives,

$$\text{Var}(\hat{\theta}_{m,n,p}) = O(\frac{1}{mn} + \frac{1}{m}). \tag{10}$$

PDS-Based Biased Efficient Estimator of CVA_I We choose the number of valuation points, n, and number of simulation runs at each valuation point, m, to minimize the mean square error of the PDS-based estimator, $\hat{\theta}_{m,n,p}$, under a fixed computational budget proportional to mn. Approximating the variance of $\hat{\theta}_{m,n,p}$ using (10) leads to the following optimization problems,

$$\min_{m,n} \left(\frac{c_{p,1}}{mn} + \frac{c_{p,2}}{m} + \frac{c_2}{n^2} \right) \quad \text{subject to} \quad s = c_3 mn, \tag{11}$$

for some constants, $c_{p,1}, c_{p,2}, c_2$, and c_3. MSE of $\hat{\theta}_{m,n,p}$ is minimized at,

$$m = cs^{\frac{2}{3}} \quad \text{and} \quad n = \tilde{c}s^{\frac{1}{3}}, \tag{12}$$

for constants c and $\tilde{c}$.

DJS-Based Biased Efficient Estimator of CVA_I Let c_d denote a constant. Given (9), we approximate $\text{Var}(\hat{\theta}_{m,n,d})$ with $\frac{c_d}{mn}$ in the MSE minimization problem for the DJS-based estimator,

$$\min_{m,n} \left(\frac{c_d}{mn} + \frac{c_2}{n^2} \right) \quad \text{subject to} \quad s = c_3 mn,$$

to which the trivial optimal solution is $m = 1$ and $n = \hat{c}s$ fo some constant $\hat{c}$. We note that estimating the various constant parameters appearing in all the above mentioned MSE minimization problems is not possible in practice. In our numerical examples we simply set all these constant parameters equal to 1.

The MSE minimization setup has appeared in various contexts before; e.g., Duffie and Glynn (1995)) and Chapter 6 of Glasserman (2004) and the references therein. However, this has neither appeared in the CCR literature nor been applied by practitioners. Also, it has never been studied in the DJS setting and our result that the efficient DJS-based estimator requires all its computational budget allocated to the number of valuation points is surprising and new.

3.3 Efficient Monte Carlo CVA$_I$ Estimation: Unbiased Estimators

In this section we derive unbiased estimators of CVA$_I$. Specifically, we eliminate the time discretization bias at the expense of introducing additional randomness. To control the variance that would be increased as the result of this new source of randomness, we use stratified sampling. Now, consider the following identity,

$$E[V_\tau \mathbf{1}\{\tau \leq T\}] = \sum_{i=1}^{n} E[V_\tau | \tau \in A_i] P(\tau \in A_i), \tag{13}$$

where stratum i is $A_i = [t_{i-1}, t_i)$. Let m_i, $i = 1, 2, ..., n$ denote the number of simulation runs used to estimate $E[V_i]$, where $V_i \equiv V_{t_i}$, $t_0 \equiv 0$, and $t_n = T$. Also, $N = \sum_{i=1}^{n} m_i$ denotes the total number of simulation runs used in estimating CVA$_I$. Using τ as the stratification variable and the identity (13), the stratified sampling estimator of CVA$_I$ is

$$\hat{\theta}_{u,m,n,k} = \sum_{i=1}^{n} \bar{V}_{\tau_i} p_i, \tag{14}$$

where $p_i \equiv P(\tau \in A_i) = \Delta F_i$, $\tau_i \equiv \tau | \tau \in A_i$, $\bar{V}_{\tau_i} = \sum_{j=1}^{m_i} V_{\tau_{ij}} / m_i$, and $k = p, d$ denotes PDS and DJS sampling, respectively. That is, to draw a single realization of V_{τ_i}, we first sample from τ conditional on $\tau \in A_i$; next, given this realization of τ_i, we generate V_{τ_i}. PDS-based simulation in calculating $\hat{\theta}_{u,m,n,p}$ implies that $\text{cov}(V_{\tau_i}, V_{\tau_j}) > 0$ for $i \neq j$, $i, j = 1, ..., n$, and DJS-based simulation in calculating $\hat{\theta}_{u,m,n,d}$ implies that $\text{cov}(V_{\tau_i}, V_{\tau_j}) = 0$ for $i \neq j$. This immediately implies $\text{Var}(\hat{\theta}_{u,m,n,d}) \leq \text{Var}(\hat{\theta}_{u,m,n,p})$.

In terms of computing time, a biased estimator requires generating N realizations of V_i and an unbiased one requires N additional samples from the truncated τ based on the strata defined above. Note that since generating V_i is computationally much more intensive than the truncated τ, a biased estimator outperforms an unbiased one merely marginally in terms of the computational time.

Proportional stratified sampling sets $m_i = NP(\tau \in A_i)$ and if we further assume $P(\tau \in A_i) = 1/n$, then all m_i's are equal to $m \equiv N/n$. In this paper we do not address further possible improvements of our unbiased stratified sampling-based estimators by attempting to find optimal $m_1, ..., m_n$ and n under fixed computational budgets. Our numerical examples indicate that using our unbiased stratified sampling-based estimators by setting $m_i \equiv m$ and choosing m and n as specified in subsection 3.2 leads to substantial MSE reduction when compared to crude biased Monte Carlo estimators of CVA$_I$.

$\hat{\theta}_{u,m,n,d}$ and the biased DJS-based estimator of CVA$_I$, $\hat{\theta}_{b,m,n,d}$, are asymptotically equivalent in terms of MSE (Proposition 4 in Ghamami and Zhang (2014)). This equivalence is further confirmed by our numerical experiments (see the next subsection) in practical settings with fixed and finite computational budgets proportional to mn. Analytically comparing the $\text{MSE}(\hat{\theta}_{b,m,n,p})$ and $\text{Var}(\hat{\theta}_{u,m,n,p})$ is quite difficult due to the presence of the covariance terms. Our numerical examples show that the unbiased PDS-based estimator of CVA$_I$, $\hat{\theta}_{u,m,n,p}$, outperforms the efficient biased PDS-estimator, $\hat{\theta}_{b,m,n,p}$.

3.4 Numerical Examples

In this section we use simple numerical examples to illustrate the efficiency of our proposed Monte Carlo estimators of CVA$_I$. We consider contract level exposure in a simple setting where $V_t \equiv S_t$ denotes the value of a geometric Brownian motion with drift μ and volatility σ at time $t > 0$. This stylized example enables us to calculate the MSE exactly. We let $T = 1$ and τ be uniformly distributed on $[0, T]$ for simplicity.

We consider six different Monte Carlo estimators of CVA_I. Let $\hat{\theta}_{c,p}$ and $\hat{\theta}_{c,d}$ denote the "crude" and biased Monte Carlo estimators of CVA_I under PDS and DJS sampling, respectively. Let $\hat{\theta}_{e,b,p}$ and $\hat{\theta}_{e,b,d}$ denote the efficient and biased Monte Carlo estimators of CVA_I under PDS and DJS sampling, respectively. In particular, their statistical efficiency is a result of solving the MSE minimization problems in Section 3.2 to derive the (approximately) optimal number of points on the time grid, n, and simulation runs at each of these time points, m, given a fixed computational budget proportional to mn. Let $\hat{\theta}_{u,p}$ and $\hat{\theta}_{u,d}$ denote the unbiased stratified sampling-based Monte Carlo estimators of CVA_I under PDS and DJS sampling, respectively.

The crude estimators of CVA_I are calculated based on 12 valuation points, $n = 12$, at 1, 2, 3, 4, 8, 12, 18, 21, 24, 36, 49 weeks and 1 year. We note that one year with the number of valuation points fixed at 12 is a setting widely used by financial institutions. There is no mathematical basis for this arrangement of valuation points. It is believed that since some trades have "short" expiration times, having more valuation points earlier would increase the accuracy of the estimators of CCR measures. The time grid used to calculate our efficient estimators of CVA_I is equidistant, i.e., $\Delta F_i \equiv \Delta = T/n$. Computational budget, s, is fixed at $12,000$ and $120,000$, respectively. To calculate $\hat{\theta}_{e,b,p}$ under these fixed computational budgets, the solution, (12) with both c and $\tilde{c}$ set to 1, to the MSE minimization problem of Section 3.2 is used. This gives, $n = 23$ and $m = 524$ for $s = 12,000$, and $n = 50$, and $m = 2433$ for $s = 120,000$. Similarly, to calculate $\hat{\theta}_{e,b,d}$, we use the solution to the MSE minimization problem, (3.2). That is, we set $n = 12,000$ and $m = 1$ for $s = 12,000$, and $n = 120,000$ and $m = 1$ for $s = 120,000$. In calculating the stratified sampling estimators of CVA_I, $\hat{\theta}_{u,p}$ and $\hat{\theta}_{u,d}$, we simply use the same setting of (m,n) as $\hat{\theta}_{e,b,p}$ and $\hat{\theta}_{e,b,d}$, respectively.

Table 1 on the facing page illustrates that our proposed estimators of CVA_I lead to substantial MSE reduction when compared to the "crude" Monte Carlo estimators. Comparing the MSE of the PDS-based estimators, $\hat{\theta}_{c,p}$, $\hat{\theta}_{e,b,p}$, and $\hat{\theta}_{u,p}$, we find that our proposed stratified sampling-based estimator leads to an MSE reduction by a factor of up to 100; this unbiased estimator also dominates the efficient biased estimator, in some cases quite substantially (see the third and fourth sections of Table 1). Comparing MSE of the DJS-based Monte Carlo estimators, $\hat{\theta}_{c,d}$, $\hat{\theta}_{e,b,d}$, and $\hat{\theta}_{u,d}$, we observe that the stratified sampling-based estimator and our efficient biased estimator perform similarly. Both efficient DJS estimators lead to substantial MSE reduction when compared to the corresponding crude estimator. Finally, we note that the variance and MSE for the crude estimators do not change much as the computational budget increases from 12,000 to 120,000, whereas those of efficient estimators reduce by up to an order of ten. This contrast yields the simple, yet useful insight that the number of valuation points should vary as the computational budget varies.

4 CONCLUSION

We propose an efficient two-step framework for CVA_I estimation. The counterparty credit risk modeler first needs to choose between the two credit exposure sampling methods: PDS or DJS. Using the notion of marginal matching, we identify conditions under which the PDS method leads to estimators whose variance is substantially larger than the variance of the DJS-based estimators. Taking into account the computational time in parallel with the MSE, we demonstrate that DJS sampling is preferable to PDS sampling for path independent derivatives. For path dependent derivatives since the computational time of the DJS-based estimator usually exceeds that of the PDS-based estimator, the two sampling methods could become approximately equivalent in some cases. Next, in the second step, the modeler needs to choose the number of valuation points and simulation runs at each valuation point. We show that the MSE of the crude Monte Carlo CVA estimators can be substantially reduced by solving approximate MSE minimization problems that specify how to achieve an approximately optimal balance between bias squared and variance. These MSE minimization problems can be easily solved after approximate orders of variance and bias under the PDS and DJS methods are derived. If the PDS method has been chosen in the first step above, we recommend employing our unbiased stratified sampling-based estimator. The unbiased estimator uses stratified sampling with the number of strata and simulation runs allocated to each stratum being chosen

Table 1: Monte Carlo CVA_I estimates for different parameter settings.

	CVA_I	Variance	MSE	CPU Time
Parameters: $S_0 = 30, \mu = .2, \sigma = .3,\ s = 12,000$				
$\hat{\theta}_{c,p}$	34.6559	0.047219	0.48478	0.00380
$\hat{\theta}_{c,d}$	34.6522	0.005028	0.43768	0.00162
$\hat{\theta}_{e,b,p}$	34.1802	0.077212	0.1117	0.00253
$\hat{\theta}_{e,b,d}$	33.9955	0.004785	0.004786	0.00174
$\hat{\theta}_{u,p}$	33.9964	0.072068	0.072064	0.00518
$\hat{\theta}_{u,d}$	33.9956	0.004865	0.004866	0.00335
Parameters: $S_0 = 30, \mu = .2, \sigma = .3,\ s = 120,000$				
$\hat{\theta}_{c,p}$	34.652	0.004791	0.4372	0.03887
$\hat{\theta}_{c,d}$	34.6521	0.000501	0.43303	0.01564
$\hat{\theta}_{e,b,p}$	34.0798	0.016741	0.024026	0.02299
$\hat{\theta}_{e,b,d}$	33.9948	0.000483	0.000483	0.02409
$\hat{\theta}_{u,p}$	33.9957	0.015533	0.015533	0.04420
$\hat{\theta}_{u,d}$	33.9945	0.000486	0.000486	0.03426
Parameters: $S_0 = 30, \mu = 1, \sigma = .3,\ s = 12,000$				
$\hat{\theta}_{c,p}$	57.7556	0.16106	23.5389	0.00389
$\hat{\theta}_{c,d}$	57.7598	0.01628	23.4351	0.00189
$\hat{\theta}_{e,b,p}$	54.1296	0.23369	1.6954	0.00270
$\hat{\theta}_{e,b,d}$	52.9238	0.015853	0.015862	0.00189
$\hat{\theta}_{u,p}$	52.9226	0.217	0.21698	0.00516
$\hat{\theta}_{u,d}$	52.9198	0.015796	0.015796	0.00390
Parameters: $S_0 = 30, \mu = 1, \sigma = .3,\ s = 120,000$				
$\hat{\theta}_{c,p}$	57.7579	0.016112	23.4159	0.03891
$\hat{\theta}_{c,d}$	57.7591	0.001616	23.4136	0.01661
$\hat{\theta}_{e,b,p}$	53.4783	0.047841	0.35899	0.02412
$\hat{\theta}_{e,b,d}$	52.9212	0.001563	0.001564	0.02627
$\hat{\theta}_{u,p}$	52.9189	0.045783	0.045781	0.04657
$\hat{\theta}_{u,d}$	52.9203	0.001565	0.001565	0.03598

based on the solution to the above mentioned MSE minimization problems, namely, under computational budget s, allocating approximately $s^{2/3}$ runs to each of the $s^{1/3}$ valuation points. An interesting case arises when the CCR modeler chooses the DJS method. In this case, our proposed efficient estimator uses 1 simulation run at each valuation point and the total computational budget is allocated to making the discrete time grid (the set of valuation points) as fine as possible.

Finally, we would like to emphasize that our results are to be assessed and applied with the understanding that evaluating the derivatives portfolio value V, consisting of possibly thousands of derivatives contracts that depend on various risk factors, even at a single time point is computationally intensive. Seeking a precisely MSE optimal setup, where the computing time is also to be taken in to account, is not practically possible even in cases where the modeler has a good grasp of the time evolution of the value process of each portfolio constituent, C_t^i. Similar situations could arise in other complex practical settings where developing simple approximate guiding principles would lead to substantial efficiency improvements. When employing our proposed framework, the modeler encounters three choices: PDS versus DJS, biased versus unbiased estimation, number of valuation points versus the number of simulation runs.

ACKNOWLEDGEMENTS

Samim Ghamami is grateful to Robert Anderson, Lisa Goldberg, and Travis Nesmith for their comments.

REFERENCES

Bohme, M., D. Chiarella, P. Harle, M. Neukirchen, T. Poppensieker, and A. Raufuss. 2011, September. "Day of reckoning: new regulation and its impact on capital market businesses". *McKinsey Working Papers on Risk*.

Broadie, M., Y. Du, and M. C. C. 2011. "Risk estimation via regression". Working paper, Columbia University.

Canabarro, E., and D. Duffie. 2003. *Measuring and Marking Counterparty Risk*. London: Asset/Liability Management for Financial Institutions.

Duffie, D., and P. Glynn. 1995. "Efficient Monte Carlo Simulation of Security Prices". *The Annals of Applied Probability* 5 (4): 897–905.

Geman, H., N. E. Karoui, and J. Rochet. 1995. "Changes of Numeraire, changes of probability measure and option pricing". *Journal of Applied Probability* 32:443–458.

Ghamami, S., and L. Goldberg. 2014. "Stochastic Intensity Models of Wrong Way Risk: Wrong Way CVA Need Not Exceed Independent CVA". *Journal of Derivatives* 21:24–35.

Ghamami, S., and B. Zhang. 2014. "Efficient Monte Carlo counterparty credit risk pricing and measurement". *The Journal of Credit Risk*. Forthcoming.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer.

Glynn, P., and W. Whitt. 1992. "The Asymptotic Efficiency Of Simulation Estimators". *Operations Research* 40:505–520.

Gordy, M. B., and S. Juneja. 2010. "Nested Simulation in Portfolio Risk Measurement". *Management Science* 56 (10): 1833–1848.

Gregory, J. 2010. *Counterparty credit risk*. Wiley Finance.

Hull, J., and A. White. 2012. "CVA and Wrong-Way Risk". *Financial Analysts Journal* 68 (5): 58–69.

Karatzas, I., and S. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. Springer.

Pykhtin, M., and S. Zhu. 2006. *Measuring counterparty credit risk for trading products under Basel II*. London: Risk Books.

Pykhtin, M., and S. Zhu. 2007, July/August. "A guide to modeling counterparty credit risk". *GARP Risk Review*:16–22.

AUTHOR BIOGRAPHIES

SAMIM GHAMAMI is an economist at the Board of Governors of the Federal Reserve System and a senior researcher at the Center for Risk Management Research at UC Berkeley. His research has broadly focused on finance and stochastic modeling. He has been a post-doctoral researcher at the CREATE Homeland Security Center, a quantitative analyst at Barclays, an adjunct faculty member of University of Southern California, a visiting scholar at the Department of Economics at UC Berkeley, and a technical advisor to the Basel Committee on Banking Supervision. His email address is samim.ghamami@frb.gov and his webpage is at http://www.federalreserve.gov/econresdata/samim-ghamami.htm.

BO ZHANG is a Research Staff Member in the Business Solutions and Mathematical Sciences Department of the IBM T.J. Watson Research Center in New York. His current research focuses on three application areas: quantitative financial risk management, resource management for information technology and systems and decision-making for emerging healthcare practices. He applies probability, statistics, dynamical systems, and optimization to solve practical problems and also works on theoretical problems in these fields. His email address is bozhang@gatech.edu and his webpage is at http://www2.isye.gatech.edu/~bzhang34/.