

# Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

Kin Hung (Felix) Kan <sup>\*</sup>    Greg Frank <sup>†</sup>    Victor Mozgin <sup>‡</sup>    R. Mark Reesor <sup>§</sup>

**Abstract.** Building on the least-squares Monte Carlo (LSM) method that was originally proposed by Longstaff and Schwartz (2001) to price American options, we develop a new version of the LSM method, which we term “optimized least-squares Monte Carlo” (OLSM), to measure the counterparty credit exposure of American-style options. In order to enhance its performance, OLSM is integrated with the following three techniques: variance reduction, initial state dispersion and multiple bucketing (piecewise linear regression). Numerical results demonstrate the power of the OLSM method.

**Keywords.** American option, counterparty exposure, Monte Carlo simulation, variance reduction, piecewise linear regression

## 1 Introduction

This paper focuses on the methodology for estimating the counterparty credit exposure of American options. The exposure here means the potential loss to the buyer of an American option resulting from a naked position in the option. Because our goal is to develop an efficient algorithm that is applicable to most financial instruments, including those that have many underlying risk factors, we concentrate exclusively on the Monte Carlo approach as it can overcome the curse of dimensionality.

There is a rich literature in pricing American options by simulation. In particular, [6] proposes the least-squares Monte Carlo (LSM) method which is very popular in practice since it is very easy to implement and readily applicable to various options that contain an early-exercise feature. Thus, LSM serves as a natural starting point to measure the credit exposure of American options.

There are two main components in calculating the credit exposure of an American option: continuation value function (CVF) estimation and exposure valuation. Suppose we are interested in calculating the credit exposures at each of the possible exercise times. Then, we need to know

---

<sup>\*</sup>Department of Applied Mathematics, University of Western Ontario, [kkan8@uwo.ca](mailto:kkan8@uwo.ca)

<sup>†</sup>Capital Markets Risk Management, CIBC, [Greg.Frank@cibc.com](mailto:Greg.Frank@cibc.com)

<sup>‡</sup>Capital Markets Risk Management, CIBC, [Victor.Mozgin@cibc.ca](mailto:Victor.Mozgin@cibc.ca)

<sup>§</sup>Departments of Applied Mathematics and Statistical & Actuarial Sciences, University of Western Ontario, [mreesor@uwo.ca](mailto:mreesor@uwo.ca)

## Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

the risk-neutral option values across a wide range of states of the underlying risk factors at each possible exercise time. Putting this into the LSM framework, we obtain the risk-neutral values of the option by estimating the CVFs through a cross-sectional linear regression at each possible exercise time. The sample paths for this estimation purpose are generated using evolution models of the underlying risk factors under the risk-neutral measure. For valuing credit exposure, however, we simulate the sample paths under the physical measure, since our concern is on the potential future market value of the option. The continuation value is estimated by inserting the realized values of the risk factors into the estimated CVF. The exposure is given by the maximum of the continuation value and the immediate exercise value. The future exposures are set to zero once the option is exercised. We term this version of LSM that is tailored to calculate credit exposure the optimized least-squares Monte Carlo (OLSM) method.

While LSM usually yields a reasonably accurate estimate for the option price today (using a moderate sample size, say 10000), the estimated CVFs at future exercise opportunities are generally less accurate, resulting in possible incorrect exercise decisions, which in turn produces misleading credit exposures. For instance, the future exposures are set to be zero once the option is determined to be exercised by the LSM algorithm. If the option actually should not be exercised, that implies its exposure is grossly underestimated. To minimize such error, we apply variance reduction techniques, initial state dispersion and multiple bucketing (piecewise linear regression) to OLSM.

We test the OLSM method on a simple American call option as there exists an analytical approximation for it in [3]. It is assumed that the underlying stock does not distribute dividends. By no arbitrage theory, it is never optimal to exercise this kind of option early. Thus, its exposure is very sensitive to incorrect exercise decisions as a result of simulation error or regression model error in the stopping time estimators. We do not address the important issue of suboptimal exercise decisions due to other factors such as pricing model error or investor behaviour. Exposure estimates due to these other factors can be above or below the actual exposure. The main reason for this choice is that the exposures are often underestimated in case the option is mistakenly exercised early, which means that the average exposures likely fall under the graph of the analytical option prices. This allows for easy observation of the bias in exposures. On the contrary, for other American options whose optimal stopping times are before maturity, incorrect exercise decisions could lead to over- or under-estimations, causing the average exposures to appear above or below the analytical option prices. Moreover, the positive and negative biases due to incorrect exercise decisions could offset each other. These make it hard to observe the significance of the bias. Therefore, we pick the simple American call option as a benchmark instrument to evaluate the accuracy of the continuation value function approximations. Numerical results show that OLSM works well for the call option even when the underlying's volatility is as high as 80%, but further research on multiple bucketing is required in order to generalize this method to handle various financial instruments.

The rest of this article is organized as follows. The next section sets up the OLSM framework for calculating credit exposures, entailing the two main components aforementioned. Section 3 is

the heart of this paper, which develops several techniques to improve the performance of OLSM. We provide numerical results for an American call option to demonstrate the effectiveness of OLSM in Section 4. Section 5 concludes and suggests future research directions.

## 2 OLSM Framework

OLSM consists of two phases. Phase one estimates the CVF at each exercise opportunity, whereas phase two makes use of the estimated CVFs to calculate credit exposure to the counterparty. Assume that we are working under the constraint that the sample paths are generated externally, and all the sample paths are used to maximize the performance of OLSM. Specifically,  $M_1 = 10000$  paths are simulated under the risk-neutral measure and  $M_2 = 10000$  paths are simulated under the physical measure. This fixed sample size will affect our selection of techniques that accelerate the convergence of the continuation value estimators. The risk-neutral prices and physical prices are used to estimate the CVFs and the credit exposures, respectively.

In this paper, without loss of generality, we assume the lognormal model for the underlying physical prices

$$S(t) = S_0 \exp \left[ \left( \mu - \frac{1}{2} \sigma^2 \right) t + \sigma W(t) \right] \quad (1)$$

and the underlying risk-neutral prices

$$S(t) = S_0 \exp \left[ \left( r - \frac{1}{2} \sigma^2 \right) t + \sigma W^*(t) \right] \quad (2)$$

where  $W(t)$  is a standard Brownian motion under the physical measure,  $W^*(t)$  is a standard Brownian motion under the risk-neutral measure,  $S_0$  is the initial stock price,  $\mu$  is the drift,  $r$  is the risk-free interest rate, and  $\sigma$  is the volatility. Note that the  $S_0$  in (1) could be different from that in (2) as we disperse the initial states in Section 3.2. The ideas developed in this paper are applicable to more general price process models (e.g., stochastic volatility, jump diffusions).

### 2.1 CVF Estimation

[6] develop an LSM interleaving estimator, where the determiner (the estimator used to make the exercise/hold decision) and the propagator (option value passed on to preceding exercise opportunity) are dependent through one overlapping path. This overlap causes the estimator to be biased high for a relatively small sample size, since (a) there exists a Jensen's inequality effect, and (b) it makes use of future information to make exercise decisions, violating the definition of a stopping time. However, it converges to a lower bound for the true option value as the sample size gets large in which case the determiner tends to a constant and defines a valid stopping time.

American option pricing is an optimal stopping time problem. Transforming it into a dynamic programming problem and running it on the  $M_1$  risk-neutral price paths, the recursive equations for the interleaving estimator are

$$\tilde{H}_k^i = x_k^i \tilde{\beta}_k \tag{3}$$

$$\hat{H}_k^i = e^{-r\Delta T} \tilde{B}_{k+1}^i \tag{4}$$

and

$$\tilde{B}_k^i = \begin{cases} \hat{H}_k^i & \text{if } \tilde{H}_k^i > P_k^i, \\ P_k^i & \text{if } \tilde{H}_k^i \leq P_k^i, \end{cases} \tag{5}$$

where  $x_k^i$  is a  $(1 \times p)$  vector of basis functions evaluated at time  $k$  for path  $i$ ,  $\tilde{\beta}_k$  is a  $(p \times 1)$  vector of least-squares regression coefficients,  $p$  is the number of basis functions,  $\tilde{H}_k^i$  (determiner) and  $\hat{H}_k^i$  (propagator) are the time- $k$ , path- $i$  continuation value estimators,  $\tilde{B}_k^i$  is the time- $k$ , path- $i$  option value estimator,  $P_k^i$  is the time- $k$ , path- $i$  exercise value,  $\Delta T = T/N$  is the time interval,  $T$  is the maturity,  $k \in \{1, \dots, N\}$  is shorthand for  $k\Delta T$  which is the possible exercise time, and the terminal conditions are given by  $\tilde{H}_N^i = \hat{H}_N^i = 0$  for all  $i \in \{1, \dots, M_1\}$ .

The basis functions that we will use are monomials up to the  $3^{rd}$  degree, i.e.,  $x = [1, S, S^2, S^3]$ . The only results from dynamic programming that we need to calculate the exposures are  $\tilde{\beta}_k$ , which are equal to  $(\mathbf{X}_k' \mathbf{X}_k)^{-1} \mathbf{X}_k' \hat{\mathbf{H}}_k$ , where  $\mathbf{X}_k = ((x_k^1)', \dots, (x_k^{M_1})')$ ,  $\hat{\mathbf{H}}_k = (\hat{H}_k^1, \dots, \hat{H}_k^{M_1})'$ ,  $^{-1}$  denotes the matrix inverse, and  $'$  denotes the matrix transpose.

## 2.2 Exposure Valuation

The exposures are measured based on the simulated physical prices. At each time on each path, we can calculate the immediate exercise value and the basis functions values using the physical prices, where the product of the latter and the previously estimated regression coefficients gives the continuation value. The exposure is equal to the maximum of the immediate exercise value and the continuation value. Once the option is exercised, the exposures at the later time points are all set to be zero. In practice, however, the exposure should only be set to zero after an appropriate close-out period. Therefore, the actual exposure is slightly underestimated by our algorithm.

There are some subtle issues that are not considered in our exposure estimation approach. For instance, when a firm buys an option with no margining, this can bring about additional exposure due to suboptimal decisions by the firm, possibly resulting in higher amounts of exposures for unexercised options. However, where there is a margining agreement, an unexercised sold option can produce additional counterparty (the writer of the option) exposure arising from unreturned margin. It could lead to an increase in bilateral exposures if the derivative is something like a swaption that can generate profit and loss of either sign after exercise.

## 3 Beyond LSM

We would like to highlight the fact that incorrect exercise decisions will result in misleading exposures due to the practice of setting the exposures at the later points to be zero after exercise.

Thus, it is highly important to estimate the continuation value functions (characterized by  $\beta_k$ ) corresponding to the physical prices accurately. However, the original LSM method does not produce CVFs that are accurate for a wide range of underlying prices due to extrapolation of the fitted regression beyond the data values used to estimate  $\beta_k$ . This might result in incorrect exercise decisions especially at future times in which real-world values of the underlying price are likely to be outside of the range of risk-neutral prices used to estimate  $\beta_k$ . To remedy this problem, several techniques are proposed in this section.

### 3.1 Faster Convergence to the Approximation

Given that the set of basis functions have been fixed, we would like to accelerate the convergence of the continuation value estimators to the true approximation provided by this set of basis functions. In other words, our aim is to reduce the simulation error. This can be achieved by bias and/or variance reduction.

#### 3.1.1 Bias Reduction

The probabilistic nature of simulation makes incorrect exercise decisions (w.r.t. the true approximation) possible, giving rise to estimator bias. [5] develops a bias correction technique for Monte Carlo pricing of early-exercise options. Some numerical results for the LSM high- and low-biased estimators are given in this paper. [4] contains a detailed analysis of the bias of the LSM estimators and provides comprehensive numerical results on these estimators and their bias-corrected versions. Although all the numerical results show that the bias correction technique is effective in reducing the bias of the LSM estimators, we do not apply this technique to calculate the exposures because we work with a sample size of 10000, where there is, in general, little simulation bias to be corrected. This is evident by the numerical results for the pricing of a five-dimensional American max-call option [4].

#### 3.1.2 Variance Reduction

Since bias is not a big issue for a sample size of 10000, we focus on enhancing the convergence rate by reducing the variance. We consider two common variance reduction techniques in the following.

##### (a) Antithetic variates

Antithetic variates can be applied to the sample paths generated under the physical measure. Specifically, here 5000 antithetic pairs of negatively correlated paths are simulated, and exposures for each path are estimated. At each possible exercise time, the average exposure for each antithetic pair is computed, resulting in 5000 values that form an empirical distribution of the exposures. We can then estimate the expected exposure or the value-at-risk (VaR) using this distribution. This method works well when the function to be estimated is monotonic.

**(b) Inner control variates**

Since we would like to reduce the variance of the continuation value estimators in the intermediate time steps instead of time zero, we use control variates in the cross-sectional regressions. To simplify notations, we suppress the subscript  $k$ . Rewriting the determiner in (3) as a weighted average of the propagators gives

$$\tilde{H}^i = x^i \tilde{\beta} = \sum_{j=1}^M \omega(i, j) \hat{H}^j, \quad (6)$$

where  $\omega(i, j)$  denotes the weight connecting path  $i$  at time  $k$  to path  $j$  at time  $k + 1$ . With control variates  $Y^j$ , (6) becomes

$$\tilde{H}^i(\alpha) = \sum_{j=1}^M \omega(i, j) \hat{H}^j - \alpha \left( \sum_{j=1}^M \omega(i, j) Y^j - \sum_{j=1}^M \omega(i, j) \mathbb{E}[Y^j] \right) \quad (7)$$

$$= \sum_{j=1}^M \omega(i, j) \left( \hat{H}^j - \alpha (Y^j - \mathbb{E}[Y^j]) \right), \quad (8)$$

where  $\alpha = \text{cov}(Y^j, \hat{H}^j) / \text{var}(Y^j)$ . As a result, we adjust the propagators with the control variates before running a cross-sectional regression to obtain the variance-reduced determiners.

In our numerical experiments, we will use three martingale control variates, namely  $e^{-r(\tau-k\Delta T)} S_\tau$ ,  $e^{-(2r+\sigma^2)(\tau-k\Delta T)} S_\tau^2$  and  $e^{-(3r+3\sigma^2)(\tau-k\Delta T)} S_\tau^3$ , which are sampled at the estimated exercise times to make them have a higher correlation to the propagators. The higher the correlation the more significant the amount of variance reduction.

### 3.2 Improving the Approximation

The methods in Section 3.1 can only reduce the number of sample paths required to attain the same accuracy for the crude Monte Carlo estimators. They help the estimators converge faster to the true approximation for a given set of basis functions, but they do not produce a better approximation to the true value. In fact, not much work has been done to address the approximation error associated with using a finite set of basis functions in the literature. We provide part of a solution to this problem in the following.

In OLSM, the CVF is estimated by a simulated regression. Initial state dispersion deals with the simulation part, generating a wider support for regression and hence reducing extrapolation error when computing credit exposures. Multiple bucketing improves the fit of the regression model.

#### 3.2.1 Initial State Dispersion

Determining the exercise strategy of an American option does not require the risk-neutral sample paths to have the same initial price. Thus, we disperse the initial state so as to improve the stability of the regressions. Now, the question is how to disperse the initial state.

[7] suggests simulation of the state variables from a fictitious initial time point prior to time zero and the original initial state using the discounted underlying asset prices, a martingale under the risk-neutral measure. The distribution at time 0 will reflect the volatility of the underlying assets while being centered at the original initial state. The problem with this method is that it is difficult to determine the fictitious initial time point. Preliminary numerical results showed that Rasmussen's method is not very stable for the purpose of exposure calculation.

Since our goal is to measure the exposure, we have to estimate the risk-neutral value of the option at the physical prices. The drift of the model for the physical prices is usually greater than that for the risk-neutral prices. This implies a demand for an accuracy of the CVF for a wide support of underlying asset prices. Instead of drawing initial states from a sampling distribution, we avoid introducing one more dimension of variation by deterministically allocating the initial states. This is to help ensure that there exist initial states in the target region from which the risk-neutral prices are simulated to cover a wide range of the physical prices at future times. On the contrary, drawing from a sampling distribution might result in initial states concentrated on a small region, which is not a desirable property because that does not make the support wide enough at future times for accurate estimation of the CVF.

We divide the initial states into three regions,  $[10, 80]$ ,  $[80, 280]$  and  $[280, 510]$ . 4000, 2000 and 4000 initial states are allocated to the three regions, respectively. Within each region, the initial states are chosen uniformly. The ratios 4:2:4 are chosen since we want a very accurate estimated CVF at the beginning and more simulated data points are necessary to capture the shape of the CVF for the widely distributed physical prices near the maturity. The region boundaries are determined by matching the mean of the risk-neutral prices with the mean plus six standard deviations of the physical prices at  $T/N$ ,  $T/2$  and  $T$ , i.e.,

$$S_0^Q e^{rk\Delta T} = S_0 e^{\mu k\Delta T} + 6S_0 e^{\mu k\Delta T} \sqrt{e^{\sigma^2 k\Delta T} - 1}, \quad (9)$$

where  $S_0^Q$  is an initial state for the risk-neutral prices,  $k$  is equal to 1, 20 and 40, respectively. With the parameter values in our numerical study, the  $S_0^Q$ 's are rounded off to 80, 280 and 510, respectively.

There are several remarks on our proposed method of initial state dispersion:

- There is no need to determine a fictitious initial time point.
- The way to choose the initial state regions is not very rigorous, but it yields satisfactory results.
- The ratios 4:2:4 are more like guidelines than rules.
- This method is independent of the financial instrument type. All the dispersion parameter values remain the same if various financial instruments share the same risk factors.

### 3.2.2 Multiple Bucketing

After generating a wide range of values of the regressors, our next target is to improve the fit of the regression model. This could be achieved by running regressions on two buckets rather than using one regression for all the underlying asset prices. This approach is especially effective when it comes close to the maturity date because the true CVF is not like a polynomial there. On the other hand, the CVF is pretty smooth at the beginning. It is expected that a single regression model would provide a good fit of the simulated data points. Following these principles, the bucket boundaries are chosen as follows:

- In the first quarter of the maturity, the bucket boundary for the underlying stock prices is chosen to be 100 (rounded-off), which is calculated using the sum of the expected value and four times the standard deviation of the underlying stock prices,

$$S_0 e^{rk\Delta T} + 4S_0 e^{r\Delta T} \sqrt{e^{\sigma^2 k\Delta T} - 1}, \quad (10)$$

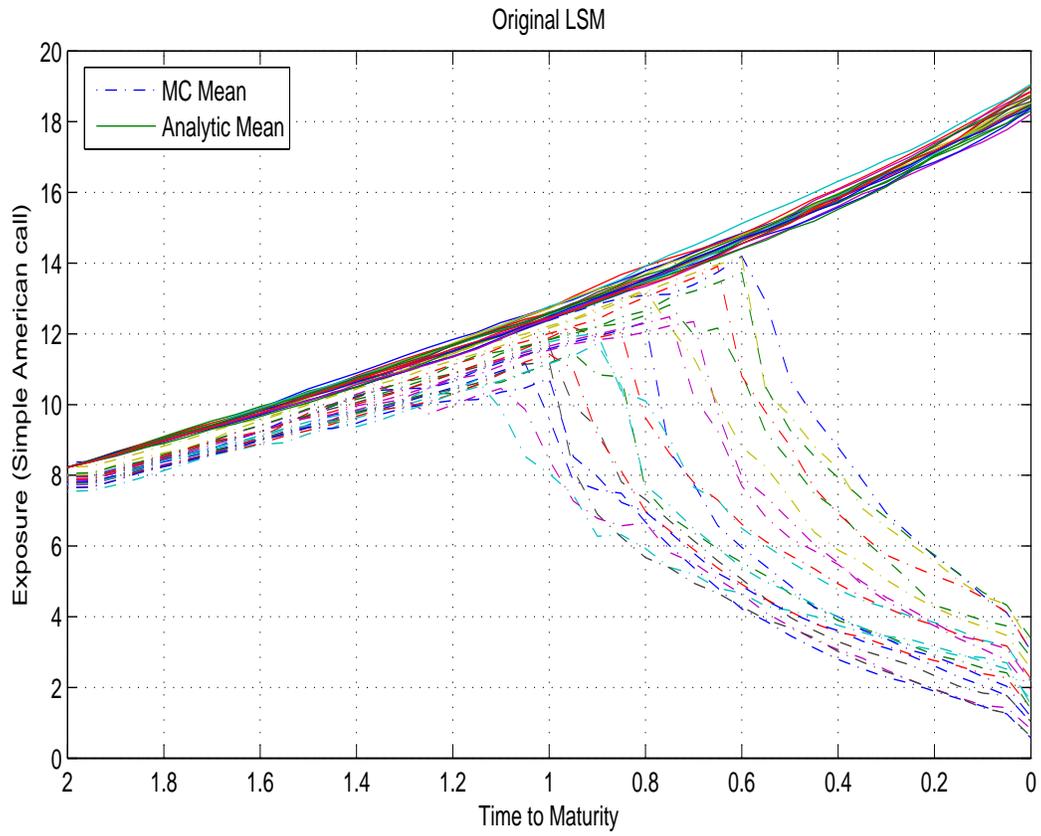
where  $k = 5$ . The interval  $(0, 100)$  would cover most simulated stock prices, thus we essentially run a single bucket regression in the first quarter, although technically speaking, that is a two-bucket regression since another regression is performed over the interval  $(100, \infty)$ .

- After  $T/4$ , we use in-the-money (ITM) and out-of-the-money (OTM) buckets. In other words, the bucket boundary is the strike price.

Different from the initial state dispersion, the selection of the bucket boundaries is sensitive to the type of the financial instrument. While formula (10) still applies in general, it is hard to justify why the bucket boundary should be at the strike price after  $T/4$  for instruments other than the simple American call option. One possible solution to this problem is to choose the “best” boundary out of several candidates according to a certain criterion. This is left for future research.

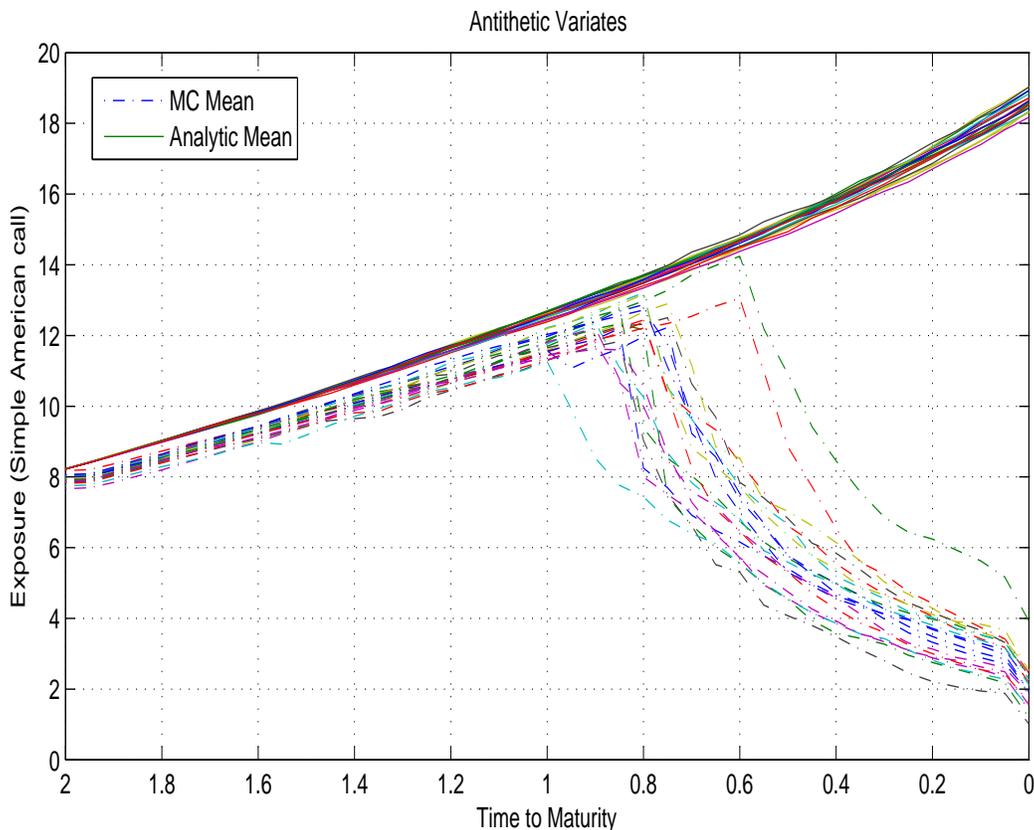
## 4 Numerical Results

The OLSM method is tested on a simple American call option where the underlying stock does not pay dividends. By no arbitrage theory, it can be shown that it is optimal not to exercise this option early. Thus, the estimated exposures of this option are very sensitive to incorrect exercise decisions. If they closely resemble the exposures calculated using the approximate formula given in [3], that indicates the estimated exposures are accurate. We consider a call option that has a maturity of two years, and a strike price of 40. The maturity is divided into 40 time steps, so the time interval is 0.05 year. The risk-free interest rate is 6%. The initial underlying stock price is 36. The stock prices are governed by a geometric Brownian motion with drift 20% and volatility 40% for Figures 1-8, and with volatility 80% for Figures 9-11, respectively. Note that we used the volatility of 80% to derive the initial states regions and the bucket boundaries in Section 3.2. Nonetheless, the same



**Figure 1:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Volatility = 40%.

## Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

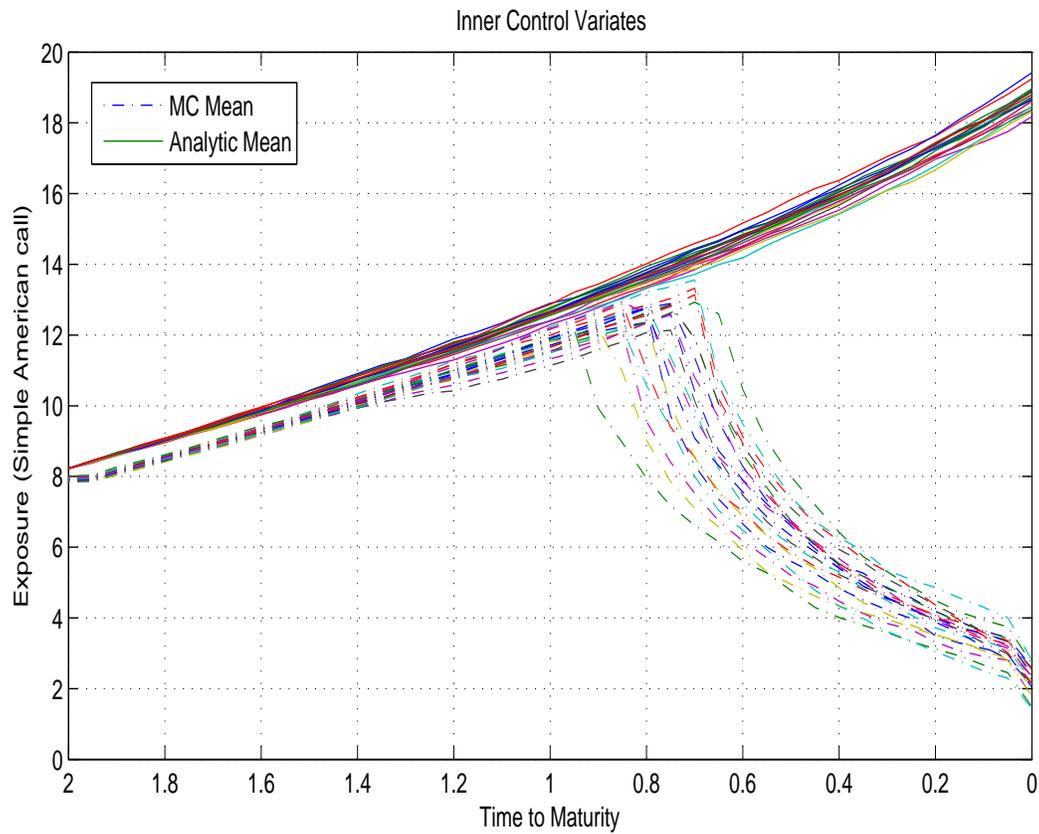


**Figure 2:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Antithetic variates are used on the “exposure” paths. Volatility = 40%.

results are applied to estimate the exposures for the option with the underlying’s volatility equal to 40%.

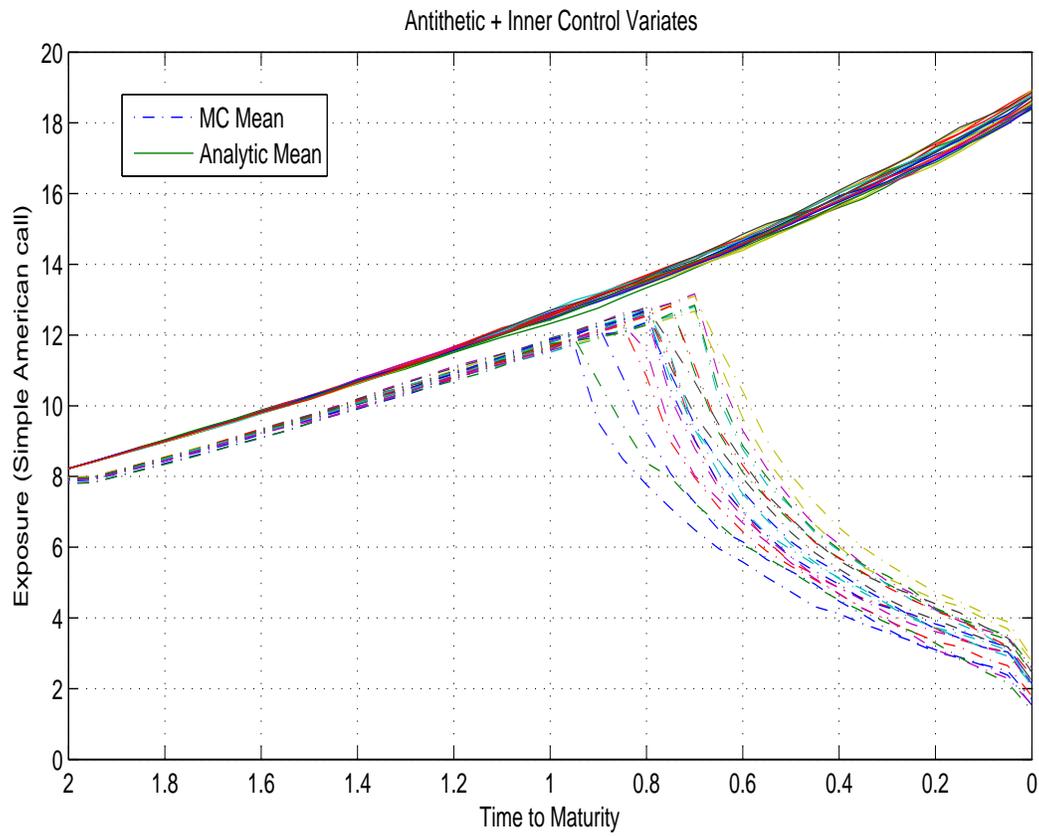
All figures plot 20 average Monte Carlo (MC) exposures and 20 average analytical exposures, where the average is taken over 10000 paths. Figure 1 plots the original LSM exposures. It shows that the original LSM exposures are quite accurate at the beginning, but fall apart from the analytical exposures after half of the maturity. Figures 2-4 plot the Monte Carlo exposures with the use of antithetic variates, inner control variates, and the combination of both, respectively. It appears that either antithetic variates or inner control variates can reduce the variance of the exposures, but using them together gives the best result. However, the MC exposures still fall off after half of the maturity.

In Figure 5, we use two buckets for regressions instead of one bucket, where the bucket boundary is the strike price. It shows that the accuracy of the exposures has been significantly improved —

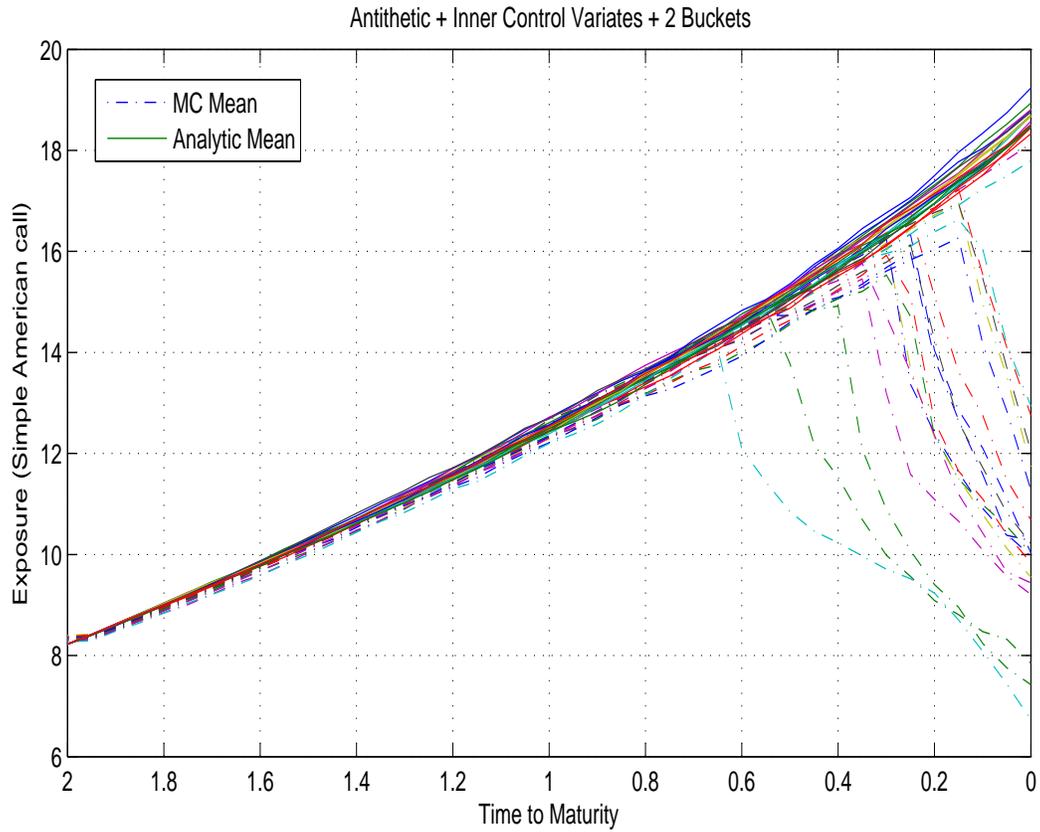


**Figure 3:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Inner control variates are used in estimating continuation value functions. Volatility = 40%.

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

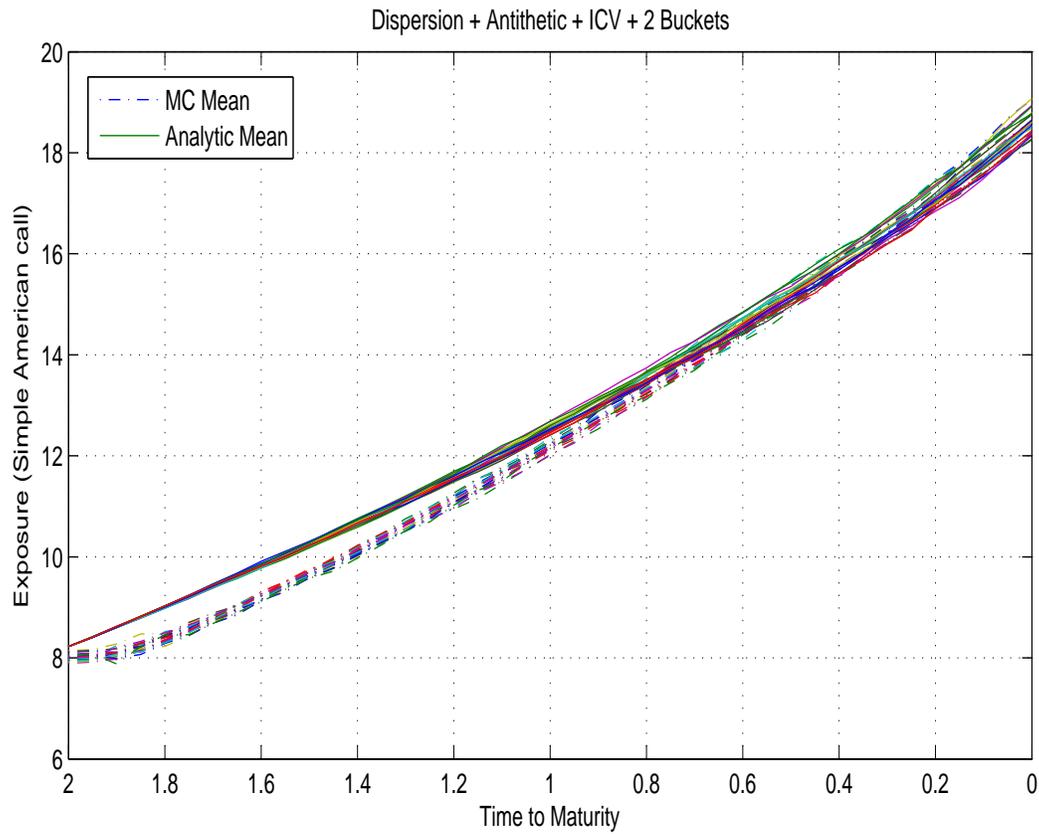


**Figure 4:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Volatility = 40%.

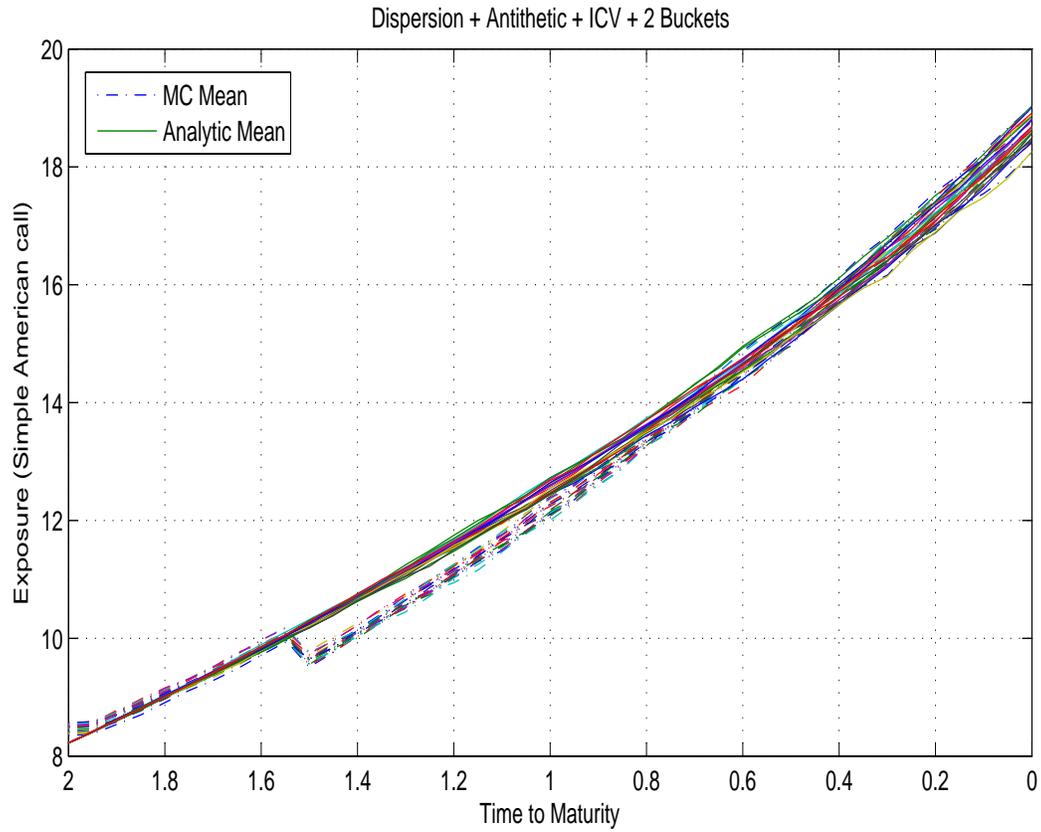


**Figure 5:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is the strike price. Volatility = 40%.

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

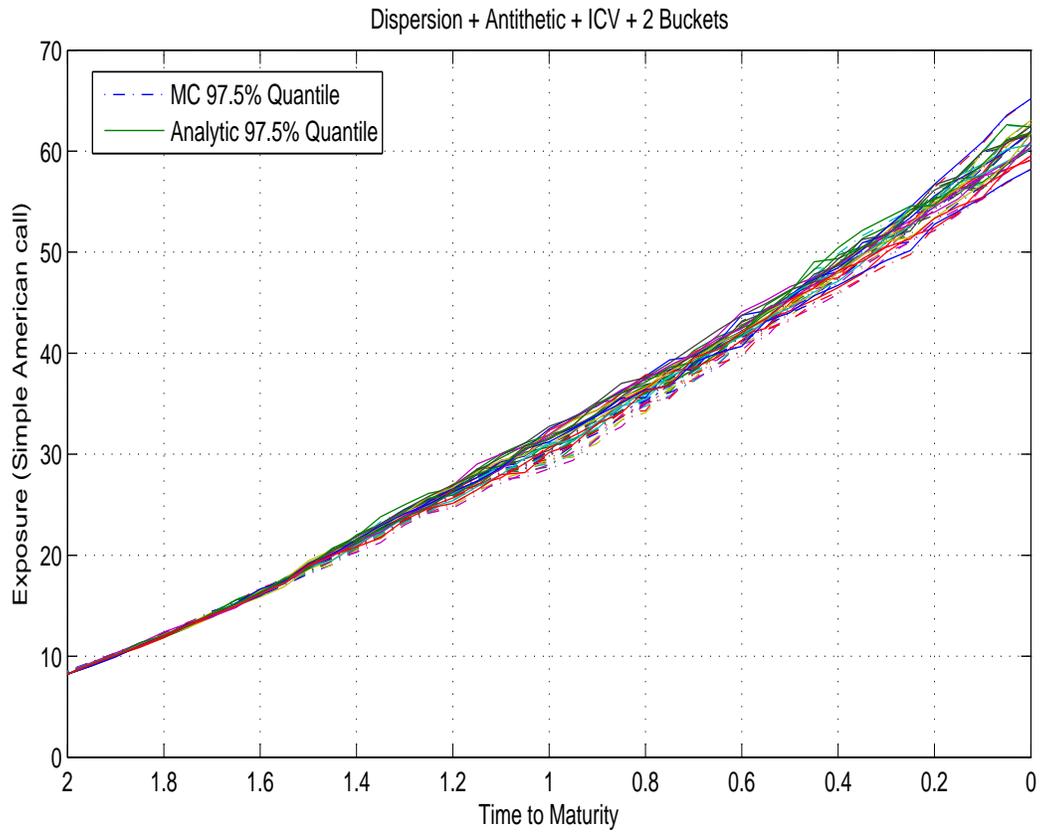


**Figure 6:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is the strike price. Initial states are dispersed based on the ratio 4:2:4 to  $(10, 80)$ ,  $(80, 280)$  and  $(280, 510)$ . Volatility = 40%.



**Figure 7:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is 100 for the first quarter of the maturity, and the strike price thereafter. Initial states are dispersed based on the ratio 4:2:4 to (10, 80), (80, 280) and (280, 510). Volatility = 40%.

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options



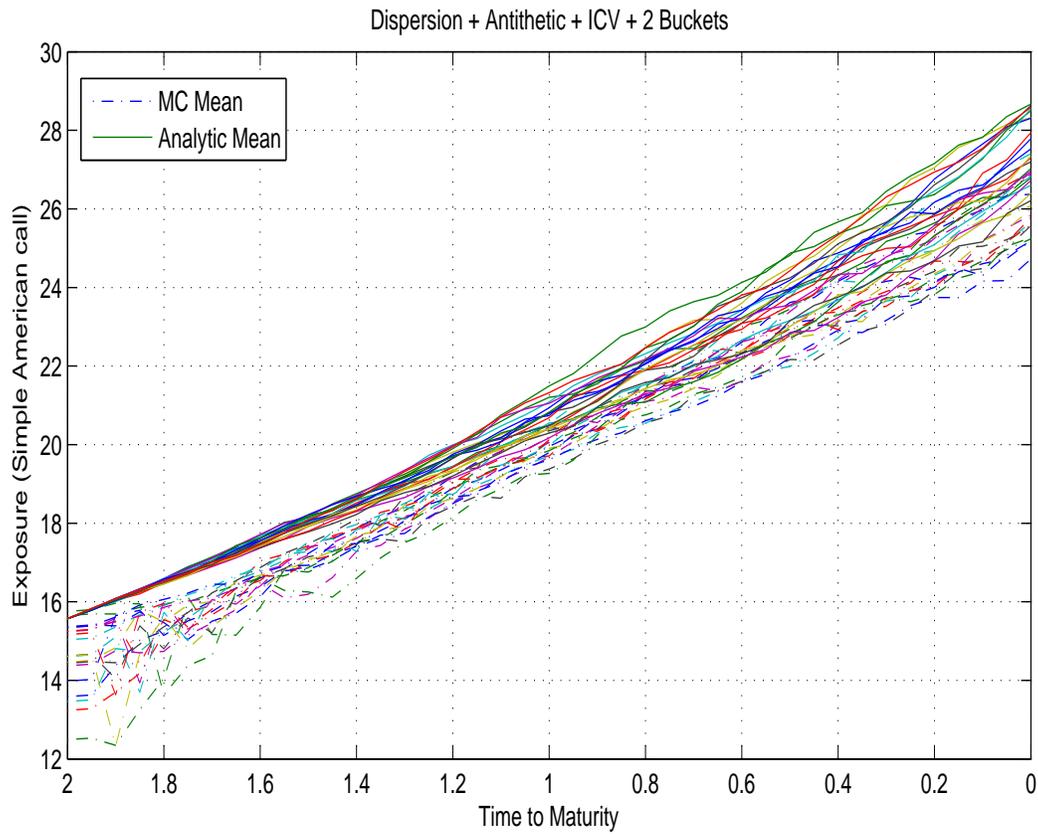
**Figure 8:** 20 97.5% quantile Monte Carlo and 20 97.5% quantile analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is 100 for the first quarter of the maturity, and the strike price thereafter. Initial states are dispersed based on the ratio 4:2:4 to (10, 80), (80, 280) and (280, 510). Volatility = 40%.

the MC exposures are “exact” in the first half and fall apart in the last quarter of the maturity. We incorporate the initial state dispersion in Figure 6. Amazingly, the MC exposures no longer fall apart from the analytical exposures, which indicates that very few or no incorrect exercise decisions have been made. However, this comes with the cost that the errors of the MC exposures at the beginning become significant. To get around this problem, the bucket boundary in the first quarter is changed to 100. Figure 7 shows that the MC exposures are very accurate after the change. Figure 8 demonstrates that the OLSM method works even better for the 97.5% quantiles than the average of the MC exposures.

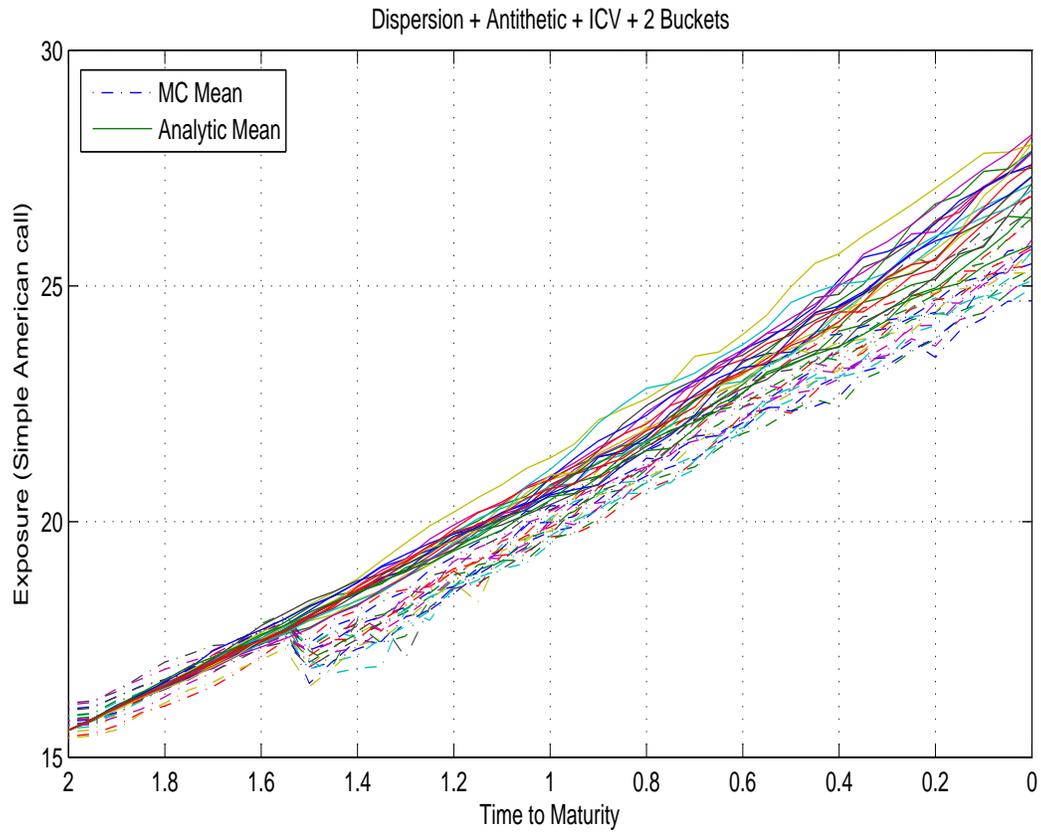
We also plot the estimated exposures corresponding to the volatility of 80% in Figures 9-11. The results are similar to those with the volatility equal to 40%, qualitatively. The variance of the MC exposures is bigger as expected.

It is worth mentioning that although the initial states regions and the bucket boundaries were not designed for the case where volatility equals 0.4, its results are excellent. This implies that the estimation of the CVF is not very sensitive to the initial states regions and the bucket boundaries when the volatility of the underlying is low. Hence, we should use a high volatility to test the robustness of a new method to measure the exposures.

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options

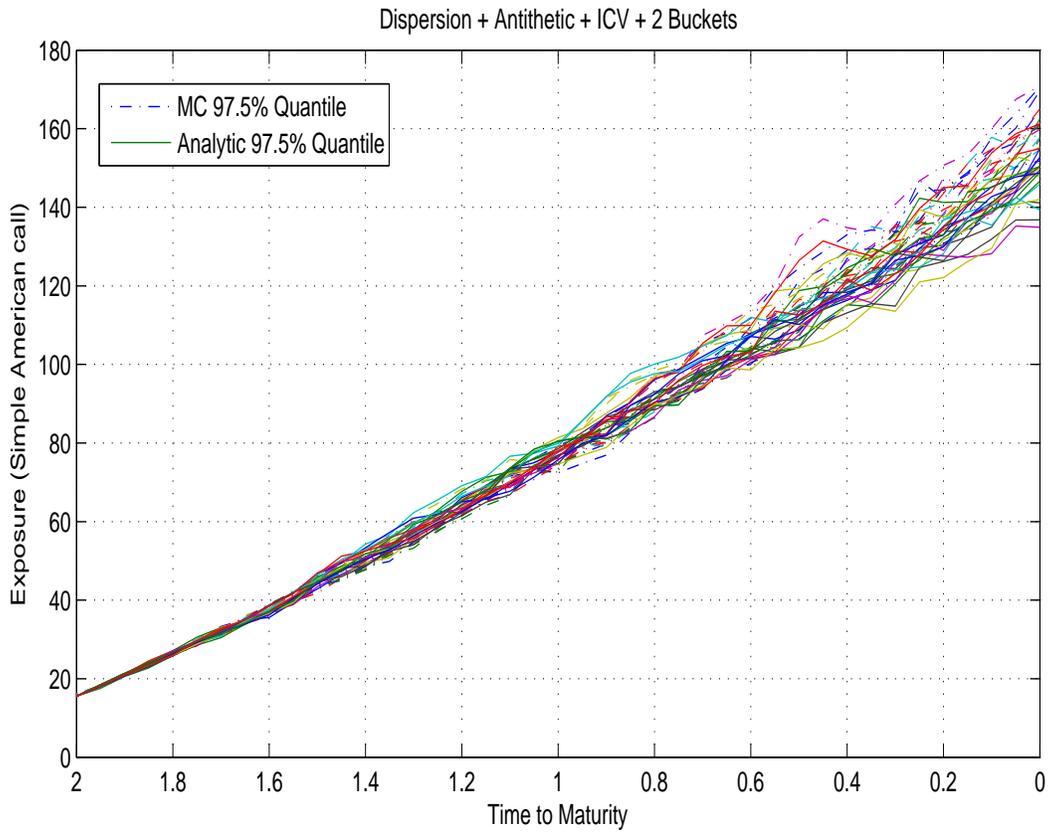


**Figure 9:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is the strike price. Initial states are dispersed based on the ratio 4:2:4 to  $(10, 80)$ ,  $(80, 280)$  and  $(280, 510)$ . Volatility = 80%.



**Figure 10:** 20 average Monte Carlo and 20 average analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is 100 for the first quarter of the maturity, and the strike price thereafter. Initial states are dispersed based on the ratio 4:2:4 to  $(10, 80)$ ,  $(80, 280)$  and  $(280, 510)$ . Volatility = 80%.

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of American-style Options



**Figure 11:** 20 97.5% quantile Monte Carlo and 20 97.5% quantile analytic exposures versus time-to-maturity. Both antithetic and inner control variates are used. Two buckets are used in the regression, where the bucket boundary is 100 for the first quarter of the maturity, and the strike price thereafter. Initial states are dispersed based on the ratio 4:2:4 to  $(10, 80)$ ,  $(80, 280)$  and  $(280, 510)$ . Volatility = 80%.

## 5 Conclusions and Future Work

The popular least-squares Monte Carlo (LSM) method is extended for the purpose of measuring counterparty credit exposure of American-style options. We optimized the LSM method using variance reduction techniques, initial state dispersion and multiple bucketing, hence the name OLSM. Numerical results for a simple American call option indicate that OLSM can reduce the absolute relative error of the average or the 97.5% quantile of the exposures to less than 10%. All of the above techniques can be applied to other American options. In particular, the same initial states can be recycled for multiple options as long as they have some common risk factors. However, the way to determine the bucket boundaries is not very general. Thus, one direction of future work is to develop a systematic approach to pick the bucket boundaries. Furthermore, we are interested in investigating the performance of the OLSM algorithm when the underlying of the option is governed by a volatility model, which allows for time-varying and state-dependent volatilities, including implied volatilities. This can result in additional exposure, or changes in early exercise decisions.

### Acknowledgments

We thank CIBC, MITACS, and UWO for providing financial support for this research project. Thanks also go to Matt Davison, Adam Metzler and the Financial Mathematics Research Group at UWO for helpful comments, leading to improvement of this work. We appreciate the insightful comments from the two anonymous referees. All remaining errors are our responsibility.

### Disclaimer

The methods and opinions expressed in this paper are those of the authors and do not necessarily represent those of CIBC. The exposure estimation approach presented is not necessarily used by CIBC.

## References

- [1] M. Broadie and P. Glasserman, *A stochastic mesh method for pricing high-dimensional American options*, The Journal of Computational Finance **7** (2004), no. 4, 35–72.
- [2] P. Glasserman, *Monte Carlo methods in financial engineering*, Springer-Verlag, New York, 2004.
- [3] N. Ju and R. Zhong, *An approximate formula for pricing American options*, Journal of Derivatives **7** (1999), 31–40. [65](#), [71](#)

Optimized Least-squares Monte Carlo for Measuring Counterparty Credit Exposure of  
American-style Options

- [4] K.H.F. Kan and R.M. Reesor, *Bias reduction for pricing American options by least-squares Monte Carlo*, Working paper, Department of Applied Mathematics, University of Western Ontario, 2009. 68
- [5] K.H.F. Kan, R.M. Reesor, T. Whitehead, and M. Davison, *Correcting the bias in Monte Carlo estimators of American-style option values*, Springer-Verlag, 2009. 68
- [6] F.A. Longstaff and E.S. Schwartz, *Valuing American options by simulation: A simple least-squares approach*, Review of Financial Studies **14** (2001), 113–147. 64, 66
- [7] N.S. Rasmussen, *Control variates for Monte Carlo valuation of American options*, Journal of Computational Finance **9** (2005). 70
- [8] T. Whitehead, *Correcting the Monte Carlo optimal-stopping bias*, Ph.D. thesis, University of Western Ontario, 2008.
- [9] T. Whitehead, R.M. Reesor, and M. Davison, *A bias-reduction technique for Monte Carlo pricing of early-exercise options*, submitted.
- [10] ———, *Correcting the Monte Carlo optimal-stopping bias*, submitted.