Leave-One-Out Least Square Monte Carlo Algorithm for Pricing American Options

Jeechul Woo
Department of Mathematics, University of Illinois Urbana-Champaign, jeechul.woo@gmail.com

Chenru Liu
Department of Management Science and Engineering, Stanford University, liucr@stanford.edu

Jaehyuk Choi
Peking University HSBC Business School, jaehyuk@phbs.pku.edu.cn

The least square Monte Carlo (LSM) algorithm proposed by Longstaff and Schwartz (2001) is widely used for pricing American options. The LSM estimator contains look-ahead bias, and the conventional technique of removing it necessitates doubling simulations. This study proposes a new approach for efficiently eliminating, thereby measuring, look-ahead bias by using the leave-one-out cross-validation (LOOCV). With the leave-one-out LSM (LOOLSM) method, we show that the look-ahead bias is asymptotically proportional to the regressors-to-simulation paths ratio and the LSM estimator have overvalued the multi-asset claims. Analysis and computational evidence support our findings.

Key words: American option, Bermudan option, Least square Monte Carlo, Longstaff–Schwartz algorithm, Look-ahead bias, Foresight bias, Leave-one-out-cross-validation

---

1. Introduction

1.1. Background

Derivatives with early exercise features are popular, with American- and Bermudan-style options being the most common types. Nonetheless, the pricing of these options is a difficult problem in the absence of closed-form solutions, even in the simplest case of valuing American options on a single asset. Researchers have thus developed various numerical methods for pricing that largely fall into two categories: the lattice-based and simulation-based approaches.

In the lattice-based approach, pricing is performed on a dense lattice in the state space by valuing the options at each point of the lattice using suitable boundary conditions and the mathematical relations among neighboring points. Examples include the finite difference scheme (Brennan and Schwartz 1977), binomial tree (Cox et al. 1979), and its multidimensional generalizations (Boyle 1988, Boyle et al. 1989, He 1990). These methods are known to work well in low-dimensional problems. However, they become impractical in higher-dimensional settings, mainly because the lattice size grows exponentially as the number of state variables increases. This phenomenon is commonly referred to as the curse of dimensionality.

In the simulation-based approach, the price is calculated as the average of the option values over simulated paths, each of which represents a future realization of the state variables with respect to the risk-neutral measure. While the methods in this category is not challenged by dimensionality, they entail finding the optimal exercise rules. Several simulation-based methods propose various approaches for estimating the continuation values as conditional expectations. Equipped with stopping time rules, they calculate the option price by solving a dynamic programming problem whose Bellman equation is essentially the comparison between the continuation values and exercise values.

The randomized tree method (Broadie and Glasserman 1997) estimates the continuation value at each node of the tree as the average discounted option values of its children. This non-parametric approach is of the most generic type, but its use is limited in scope because the tree size still grows exponentially in the number of exercise times. The stochastic mesh method (Broadie and Glasserman 2004) overcomes this issue by using the mesh structure in which all the states at the next exercise time are the children of any state at the current exercise time. The conditional expectation is computed as a weighted average of the children, where the weights are determined by likelihood ratios. Regression-based methods (Carriere 1996, Tsitsiklis and Van Roy 2001, Longstaff and Schwartz 2001) use regression techniques to estimate the continuation values from the simulated paths. Those approaches are computationally tractable, as they are linear not only in the number of simulated paths, but also in the number of exercise times. Fu et al. (2001) and Glasserman (2003) provide an excellent review of the implementation and comparison of simulation-based methods.

Among the regression-based methods, the least square Monte Carlo (LSM) algorithm proposed by Longstaff and Schwartz (2001) has been widely used in practice for valuing options with early exercise features for its simplicity and efficiency. Particularly, it is heavily used for pricing callable structured notes whose coupons have complicated dependency on other underlying assets such as equity prices, foreign exchange rates, and benchmark interest swap rates. The financial institutions that issue the notes effectively buy the Bermudan-style right to early redeem the notes. In return, the investors receive a premium in the form of a higher yield compared to the non-callable notes with the same structure. Because multi-factor models are required for the underlying assets as well as for the yield curve term structure, the use of Monte Carlo simulation along with the LSM method is inevitable for pricing such notes. This study has been motivated and developed in the circumstances.

1.2. Biases in the LSM method

In simulation-based methods including the LSM method, there are two main sources for bias that are opposite in directions. The low-side bias is related to suboptimal exercise decision owing to various approximations adopted in the method. In the LSM method, for example, finite basis functions cannot fully represent the conditional payoff function and the price estimated from finite simulation paths contains noise. The resulting exercise policies deviate from the most optimal one and therefore lead to a lower option price. For this reason, it is also called *suboptimal* bias. The high-side bias comes from sharing the simulation results for the exercise policy and payoff valuation. As explained in Broadie and Glasserman (1997), this practice undesirably creates positive correlation between exercise decisions and future payoffs; the algorithm is more likely to continue (exercise) precisely when the *future* payoff in the simulation paths is higher (lower). For this reason, it is called *look-ahead* or *foresight* bias. The LMS estimator has both low- and high-sided biases, hence, Glasserman (2003) calls it an interleaving estimator. Other simulation estimators in the literature are typically either low-biased or high-biased. For example, Broadie and Glasserman (1997) carefully constructs both low- and high-biased estimators to form a confidence interval for the true option price.

We argue that, in buyer-driven option markets, the look-ahead bias is more dangerous than the suboptimal bias, and that the look-ahead bias being mixed with the suboptimal bias is a significant drawback of the LSM estimator. By buyer-driven, we refer to the markets in which option buyers are capable of using quantitative models to value and risk-manage the options. The buyers pay the option premium, less some margin, and realize gain through delta-hedging and possibly early exercising the option. On the contrary, the sellers pocket the premium and hold the option without active risk management. While sellers may not have valuation capability, they can choose the buyer who bid the highest option premium. The structured notes market is a prime example of the buyer-driven option market.

From the buyers' perspective, the look-ahead bias is malicious because it wrongly inflates the option premium they pay. Regardless of delta-hedging, the value attributed from the look-ahead bias shrinks to zero when the position is near the maturity or the early exercise because there is no more *future* to look into by then. The suboptimal bias, on the contrary, is benign to the buyers. Although it deflates the premium, the gain realized through delta-hedging under the suboptimal exercise policy is just as much. In short, the option buyers *get what they pay for*. The only downside from the suboptimal bias is to lose the trade to competitors. In general, option buyers are more concerned with the low-biased estimator or the lower bound of the option value to ensure that the premium they pay is lower than the true option value. However, the look-ahead bias mixed in the LSM method makes the *conservative* estimation difficult for buyers.

A standard technique of eliminating the look-ahead bias is to calculate the exercise criteria by using an additional independent set of Monte Carlo paths, thereby eliminating the correlation between exercise decision and simulated payoff. While this two-pass approach removes the look-ahead bias, it comes at the cost of the computational burden from doubling the simulation effort. It is uncommon that the simulation of stochastic processes requires the time-discretized Euler scheme. The design of the LSM estimator to include the biases in both directions primarily is to retain the computational efficiency rather than to raise accuracy by letting these two biases partially offset. Moreover, Longstaff and Schwartz (2001) claim that the look-ahead bias of the LSM estimator is negligible by presenting a single-asset put option case tested with the two-pass simulation as supporting evidence. In this regard, the LSM estimator has been considered low-biased.

However, there has been concerns about the validity of the assumption. Although the look-ahead bias is small in ideal settings, it may not generalize to broader cases in practice. Carriere (1996) reports that, when the simulation paths are small, the look-ahead bias is statistically significant in the estimator predating, but similar to, the LSM estimator. Fries (2005) remarks the same for the LSM estimator and conjectures that the bias can be greater in high dimensional models without evidence. Indeed, the practitioners in structured notes market observe that, when they include higher order regression variables in an effort to better capture the exercise boundary (i.e., reduce the suboptimal bias), the look-ahead bias also increases. The tendency is more pronounced in multi-asset or multi-factor settings. Given the desire to keep the classical one-pass LSM implementation for computational efficiency, they are reluctant to include higher order terms in the LSM regression. There are methods developed to estimate both lower and upper bounds of American options based on policy iteration (Kolodko and Schoenmakers 2006, Beveridge et al. 2013) and duality representation (Haugh and Kogan 2004, Andersen and Broadie 2004), respectively. However, their computational cost is heavy to be used in day-to-day pricing and risk-management as nested simulations are required. Therefore, it is of significant practical importance to understand the magnitude of the look-ahead bias in the LSM estimator and to develop an efficient algorithm for removing this bias.

1.3. Contribution of this study

In this article, we present an efficient approach for removing look-ahead bias, motivated from the cross validation practice in statistical learning. A standard practice is to separate dataset for training and testing purposes to avoid overfitting. In the context of statistical learning, look-ahead bias is an overfitting caused by using the same dataset for both training (i.e., the estimation of the exercise policy) and testing (i.e., the valuation of the options). In this context, using an independent simulation set for exercise policy corresponds to the hold-out method, one of the simplest cross-validation techniques.

Among advanced cross-validation techniques, we recognize that the leave-one-out cross-validation (LOOCV) exactly fits to the LSM method. When making a prediction for a sample, LOOCV trains the model with all samples except the one, thereby separating the dataset for testing in the most minimal way. In linear regression, the correction from the full regression can be analytically computed without actually conducting regressions as many as the number of simulation paths. Our new leave-one-out LSM (LOOLSM) algorithm can be understood as an extension of the low estimator of Broadie and Glasserman (1997) in the sense that the self-exclusion is conducted on all simulation paths rather than on each state separately. Therefore, this simple idea can eliminate the look-ahead bias without extra computation cost and make a genuinely low-biased estimator.

Furthermore, we can easily measure the look-ahead bias as the difference from the LSM method, from which we examine the asymptotic behavior both theoretically and empirically. Numerical experiments are conducted for the cases whose true option values are available. Specifically, we show that the look-ahead bias is proportional to the ratio of the regressors to the simulation paths, unifying the previous observations on the two effects. This result is in contrast to that of Glasserman and Yu (2004) which discusses the numbers of the simulation paths and regressors in the context of Monte Carlo convergence.

To the best knowledge of the authors, previous works on the look-ahead bias correction is limited. Fries (2005, 2008) formulates the look-ahead bias as the price of the option on the Monte Carlo error and derives the analytic correction terms from the Gaussian error assumption. Compared to the study, our method does not depend on any model assumption and the bias to be removed is more consistent with the LSM setting, as we explain in § 2.1 and Appendix B.

Beyond the American option pricing, the our new method can be applied to various stochastic control problems in finance where the least square regression is used to approximate the optimal strategy. For such problems, see (Huang and Kwok 2016) for variable annuities, (Nadarajah et al. 2017) for energy real options, and (Bacinello et al. 2010) for life insurance contracts.

The rest of the paper is organized as follows. In § 2, we describe the LSM pricing framework and introduce the LOOLSM algorithm. In § 3, we analyze the asymptotic behavior of the look-ahead bias in the LSM method. In § 4, the numerical results are presented. Finally, § 5 concludes.

2. Method

In this section, we briefly review the American option pricing primarily to develop our method later. For detailed review, we refer to Glasserman (2003). We first introduce the conventions and notations to be used in the rest of the paper:

- The option can be exercised at a discrete time set $\{0=t_0<t_1<\cdots<t_I=T\}$. It is customary to assume that the present time $t_0=0$ is *not* an exercise time.
- $S(t)=(S_1(t),\cdots,S_J(t))$ denotes the Markovian state vector at time t. The value at t_i is denoted by $S^{[i]}$.

- $Z^{[i]}(S)$ denotes the option payout at t_i in state S. It is discounted to the present time $t_0=0$. For example, $Z^{[i]}(S)=e^{-r\,t_i}\max(S_1^{[i]}-K,0)$ for a single-stock call option struck at K when the risk-free rate is constant as r.
- $V^{[i]}(S)$ and $C^{[i]}(S)$ denotes the discounted option value at t_i in state S given the option was not exercised up to (and including) t_{i-1} and t_i, respectively. In the literature, $C^{[i]}(S)$ is commonly referred to as the *continuation value.*

The exercise time index $[i]$ or the time dependency (t) may be omitted when it is clear from the context.

The valuation of options with early exercise features can be formulated as a maximization problem of the expected future payoffs over all possible choices of discrete stopping times taking values in $\{1,\cdots,I\}$:

$$V^{[0]}(S)=\max_{\tau\in\mathcal{T}}\ \mathbb{E}[Z^{[\tau]}(S^{[\tau]})\,|\,S^{[0]}=S]. \tag{1}$$

It can be alternatively formulated as a dynamic programming problem using the continuation value. Since $C^{[i]}(S)$ and $V^{[i+1]}(S)$ are related by

$$C^{[i]}(S)=\mathbb{E}[V^{[i+1]}(S^{[i+1]})\,|\,S^{[i]}=S]\quad\text{for } 0\le i<I,$$

the option value at t_i is calculated by the backward induction,

$$V^{[i]}(S)=\max(C^{[i]}(S),\,Z^{[i]}(S)). \tag{2}$$

This effectively means that the option is continued at t_i if $C^{[i]}(S)\ge Z^{[i]}(S)$ and exercised otherwise. For consistency, we assume $Z^{[0]}(S)=C^{[I]}(S)=-\infty$ to ensure that $V^{[0]}(S)=C^{[0]}(S)$ (i.e., must continue at $t_0=0$), and $V^{[I]}(S)=Z^{[I]}(S)$ (i.e., must exercise at $t_I=T$ if not before). Therefore, the optimal stopping time τ is expressed in terms of $C^{[i]}(S)$ and $Z^{[i]}(S)$ as

$$\tau=\inf\{0<i\le I: C^{[i]}(S^{[i]})<Z^{[i]}(S^{[i]})\}.$$

To see how the pricing is done in the Monte Carlo simulation setting, we further introduce the following conventions and notations:

- We generate N simulation paths of $S^{[i]}$ $(1\le i\le I)$ with the initial value $S^{[0]}$. The n-th simulation value of $S^{[i]}$ is denoted by $S_n^{[i]}$ (i.e. the subscript n is for the index).
- $X^{[i]}(S)=(1,f_1(S),\cdots,f_{M-1}(S))$ denotes the set of M basis functions at t_i in state S.
- The N by M matrix $\boldsymbol{X}^{[i]}$ is the simulation result of $X^{[i]}(S)$. The n-th row, denoted by $\boldsymbol{x}_n^{[i]}$, corresponds to $X^{[i]}(S_n^{[i]})$.
- The vectors, $\boldsymbol{C}^{[i]}$ and $\boldsymbol{Z}^{[i]}$, are the length-N column vectors consisting of the simulation results of $C^{[i]}(S)$ and $Z^{[i]}(S)$, respectively.

- The vector $\boldsymbol{V}^{[i]}$ is the length-N column vector consisting of the option payout at the stopping time for simulated paths, conditional on that the option was not exercised prior to t_i, i.e. $i \le \tau \le I$.
- For vector $\boldsymbol{A}$, in general, A_n refers to the n-th element (scalar) of $\boldsymbol{A}$. Therefore, $C_n^{[i]}$, $Z_n^{[i]}$, and $V_n^{[i]}$ is understood as the path-wise value of $C^{[i]}(S^{[i]})$, $Z^{[i]}(S^{[i]})$, and $V^{[i]}(S^{[i]})$, respectively.
- $\boldsymbol{\Sigma}^{[i]}$ denotes the sample covariance matrix of $\boldsymbol{X}^{[i]}$ excluding the intercept (henceforth, of size $M-1$ by $M-1$).

Suppose that we have an estimate of the continuation value function as $\hat{C}^{[i]}(S)$, and the corresponding simulation values are $\boldsymbol{C}^{[i]}$. Following the stopping time formulation (1), $\boldsymbol{V}^{[i]}$ is computed as a path-wise backward induction step: $V_n^{[I]} = Z_n^{[I]}$ and

$$V_n^{[i]} = \begin{cases} Z_n^{[i]} & \text{if } Z_n^{[i]} > C_n^{[i]} \\ V_n^{[i+1]} & \text{if } Z_n^{[i]} \le C_n^{[i]} \end{cases} = I[C_n^{[i]} \ge Z_n^{[i]}] \cdot (V_n^{[i+1]} - Z_n^{[i]}) + Z_n^{[i]} \quad \text{for } 0 \le i < I, \tag{3}$$

where $I[\cdot]$ is the indicator function having value 1 if the condition is satisfied or 0 otherwise. In the final step of the backward induction, the option price estimate at $t=0$ is calculated as the average option value over the simulated paths,

$$\hat{V}^{[0]}(S) = \frac{1}{N} \sum_{n=1}^{N} V_n^{[0]}. \tag{4}$$

It should be noted that the estimations, $\hat{C}^{[i]}(S)$ and $\hat{V}^{[0]}(S)$, as opposed to the true value functions $C^{[i]}(S)$ and $V^{[0]}(S)$, depend on the method to estimate $\hat{C}(S)$ (e.g., choice of basis functions) as well as the used simulation set.

The backward induction approach in (3) is adopted by many authors, notably Tilley (1993), Carriere (1996), Longstaff and Schwartz (2001). There is an alternative formulation based on Equation (2),

$$V_n'^{[i]} = \max(C_n^{[i]}, Z_n^{[i]}) = I[C_n^{[i]} \ge Z_n^{[i]}] \cdot (C_n^{[i]} - Z_n^{[i]}) + Z_n^{[i]}. \tag{5}$$

It is also used extensively, see for instance Carriere (1996), Tsitsiklis and Van Roy (2001). However, it is reported that the high bias of the latter approach is significantly higher than that of the former (Carriere 1996, Longstaff and Schwartz 2001).

2.1. The LSM Algorithm

The main difficulty in pricing Bermudan options with simulation methods lies in obtaining $\hat{C}^{[i]}(S)$ (henceforth $\boldsymbol{C}^{[i]}$) from the simulated paths. It is primarily because the Monte Carlo path generation goes *forward* in time, whereas the dynamic programming for pricing works *backward* in time by construction. Longstaff and Schwartz (2001) obtain $\boldsymbol{C}^{[i]}$ as the least squares regression of the next pathwise option values, $\boldsymbol{V}^{[i+1]}$, on the current state, $\boldsymbol{X}^{[i]}$. Therefore, $\boldsymbol{C}^{[i]} = \boldsymbol{X}^{[i]}\boldsymbol{\beta}^{[i]}$, where $\boldsymbol{\beta}^{[i]}$ is the length-M column vector of regression coefficients. Omitting the exercise time superscript $[i]$ for notational simplicity, $\boldsymbol{C}$ and $\boldsymbol{\beta}$ are given as

$$\boldsymbol{C} = \boldsymbol{X}\boldsymbol{\beta} = \boldsymbol{H}\boldsymbol{V} \quad \text{where} \quad \boldsymbol{\beta} = (\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top \boldsymbol{V}, \quad \boldsymbol{H} = \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top,$$

where $\boldsymbol{H}$ is the hat matrix. Note that $\boldsymbol{H}$ depends on the current state, $\boldsymbol{X}$, not on the future information, $\boldsymbol{V}$. Using the Equation (3), we inductively run the regression for $i = I-1, \cdots, 1$ until we obtain the option price.

There are two points of consideration regarding the implementation of the LSM regression. First, we include the payout function $\boldsymbol{Z}^{[i]}$ as a regressor for our implementation in § 4; The regression can alternatively fit the continuation premium (or penalty), $\boldsymbol{V}^{[i+1]} - \boldsymbol{Z}^{[i]}$, instead of the continuation value, $\boldsymbol{V}^{[i+1]}$. It is then used to estimate $\boldsymbol{C}^{[i]} - \boldsymbol{Z}^{[i]}$. While it is difficult to judge which approach is superior, the outcomes become identical when $\boldsymbol{Z}^{[i]}$ is included as a regression variable. Additionally, the payout function is an important indicator improving the optimality of exercise decision as tested in Glasserman (2003). Second, Longstaff and Schwartz (2001) originally suggest to run the regressions with the in-the-money paths only, i.e., $\{n: Z_n^{[i]} > 0\}$. Based on the later observation that such practice can be inferior in some cases (Glasserman 2003), we use all simulation paths in this study. To the best of our knowledge, it is also consistent with the industry practices. However, we make one adjustment; when the payout is always non-negative, typically in the form of $Z^{[i]}(S) = \max(g(S), 0)$ for a function g, we continue the option when $Z_n^{[i]} = 0$ even if $C_n^{[i]} < Z_n^{[i]}$. For such options, the negative continuation value is an artifact caused by simulation noise or imperfect basis functions and there is *nothing to lose* by continuing the option as the path may hit the in-the-money area later.

To identify how the look-ahead bias arises, we take the expectation of Equation (3) over all possible simulation sets of fixed simulation size N, conditional on that $S_n^{[i]} = S$. The intermediate option price, $V^{[i]}(S)$, is estimated as

$$\hat{V}^{[i]}(S) = \mathbb{E}[I[C_n^{[i]} \geq Z_n^{[i]}] \cdot (V_n^{[i+1]} - Z_n^{[i]}) \,|\, S_n^{[i]} = S] + Z^{[i]}(S).$$

In the LSM method, $C_n^{[i]}$ depends on $V_n^{[i+1]}$ via $\boldsymbol{C}^{[i]} = \boldsymbol{H}^{[i]}\boldsymbol{V}^{[i+1]}$. Therefore, the look-ahead bias comes from the covariance between the exercise decision, $I[C_n^{[i]} \geq Z_n^{[i]}]$, and the continuation premium, $V_n^{[i+1]} - Z_n^{[i]}$, which is future information of the path:

$$B^{[i]}(S) = \mathrm{Cov}\Big(I[C_n^{[i]} \geq Z_n^{[i]}],\, V_n^{[i+1]} - Z_n^{[i]} \,\Big|\, S_n^{[i]} = S\Big). \tag{6}$$

The look-ahead bias is positive because $C_n^{[i]}$ is tilted toward $V_n^{[i+1]}$.

Several additional comments on the the look-ahead bias are in order. First, the look-ahead bias is removed when an independent simulation set is used to estimate $\boldsymbol{C}^{[i]}$ because it is no longer correlated with $\boldsymbol{V}^{[i+1]}$. Second, if the look-ahead bias is removed, therefore the covariance is zero, the remaining estimator is suboptimal: This can be shown inductively,

$$\begin{aligned}
\hat{V}^{[i]}(S) &= \mathbb{E}[I[\,C_n^{[i]} \geq Z_n^{[i]}\,] \,|\, S_n^{[i]} = S] \cdot \mathbb{E}[V_n^{[i+1]} - Z_n^{[i]} \,|\, S_n^{[i]} = S] + Z^{[i]}(S) \\
&\leq p(S)\,\mathbb{E}[V^{[i+1]}(S_n^{[i]}) \,|\, S_n^{[i]} = S] + (1 - p(S))Z^{[i]}(S) \\
&= p(S)\,C^{[i]}(S) + (1 - p(S))Z^{[i]}(S) \\
&\leq \max(C^{[i]}(S), Z^{[i]}(S)) = V^{[i]}(S).
\end{aligned}$$

Here, $p(S) = \mathbb{E}[I[\,C_n^{[i]} \geq Z_n^{[i]}\,] \,|\, S_n^{[i]} = S]$ is the exercise probability at the state S. Third, our look-ahead bias expression is subtly different from that of Fries (2005, 2008). He defines it as the value of the option on the Monte Carlo error in the estimation of the continuation values,

$$B_{\text{Fries}}^{[i]}(S) = \mathrm{Cov}\Big(I[C_n^{[i]} \geq Z_n^{[i]}],\, C_n^{[i]} - Z_n^{[i]} \,\Big|\, S_n^{[i]} = S\Big).$$

We argue that this definition is inconsistent because it is based on the alternative backward induction, Equation (5), even though Fries (2005, 2008) investigates the look-ahead bias in the LSM method. In Appendix B, we discuss the difference in detail.

2.2. The LOOLSM Algorithm

The look-ahead bias in the LSM method is removed simply by omitting each simulation path from the regression and making the exercise decision on the path from the self-excluded regression. In the bias formulation, Equation (6), the correlation is eliminated because $V_n^{[i+1]}$ is excluded from estimating $C_n^{[i]}$. Figure 1 illustrates this idea with a toy example with three simulation paths.

This idea is well-known as LOOCV (Hoaglin and Welsch 1978). It is a special type of the k-fold cross-validation method where k is equal to the number of data points. The adjusted prediction values can be analytically obtained without running regressions N times. The prediction value $\boldsymbol{C}'$ with the leave-one-out regression is expressed as a correction to $\boldsymbol{C}$:

$$\boldsymbol{C}' = \boldsymbol{C} - \frac{\boldsymbol{h} \cdot \boldsymbol{e}}{\boldsymbol{1}_M - \boldsymbol{h}} \quad \text{for} \quad \boldsymbol{e} = \boldsymbol{V} - \boldsymbol{C}, \tag{7}$$

where $\boldsymbol{e}$ is the vector of the prediction errors, $\boldsymbol{1}_M$ is the size-M column vector of 1's, $\boldsymbol{h} = (h_n)$ is the diagonal vector of $\boldsymbol{H}$, and the arithmetic operations between vectors are done element-wise. Here, h_n measures the self-sensitivity of the prediction C_n on the observation V_n, i.e., $h_n = \partial C_n / \partial V_n$, and is referred to as *leverage score*. It satisfies $0 \leq h_n \leq 1$ and $\sum_{n=1}^N h_n = \mathrm{rank}(\boldsymbol{X}) \leq M$, because $\boldsymbol{H}$ is symmetric and idempotent. Equation (7) is well-defined if $h_n < 1$, or equivalently, $\boldsymbol{X}^\top \boldsymbol{X} - \boldsymbol{x}_n^\top \boldsymbol{x}_n$ is of full rank. Therefore, the LOOLSM method we propose can be trivially modified from the LSM method by replacing $C_n^{[i]}$ with $C_n'^{[i]}$ in the backward induction step (3).

Figure 1 Illustration of the look-ahead bias correction via LOOCV. The x-axis is the current state variable, $S_n^{[i]}$, and the y-axis is the continuation premium, $V_n^{[i+1]} - Z_n^{[i]}$. There are three simulated paths: **1**, **2**, and **3**. The full regression result ($\hat{y} = 1 + x$) indicates that the path **2** should be *continued* ($\hat{y}_2 = 1 > 0$). However, this is influenced from the high value $y_2 = 4$ at the path **2**. Based on the regression without the path **2** ($\hat{y}' = -\frac{2}{3} + \frac{5}{6}x$), it should be *exercised* ($\hat{y}_2' = -\frac{2}{3} < 0$). Therefore, the path **2** contributes to the look-ahead bias.

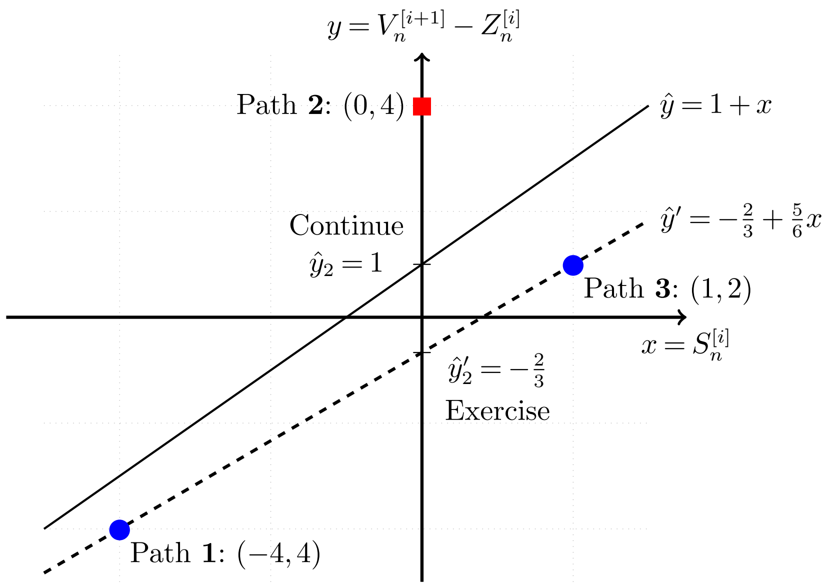

The whole vector $\boldsymbol{h}$ can be computed as the row sum of the element-wise multiplication $\cdot$ of the two matrices,

$$\boldsymbol{h} = \sum_{\text{row}} \boldsymbol{X} \cdot \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1},$$

which is a straightforward calculation from the definition $h_n = \boldsymbol{x}_n(\boldsymbol{X}^\top \boldsymbol{X})^{-1} \boldsymbol{x}_n^\top$. As the transpose of $\boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}$ is already computed for the full regression, LOOLSM can be implemented with only $O(NM)$ additional operations. This is much more efficient than obtaining $\boldsymbol{h}$ from the full $\boldsymbol{H}$ matrix.

The LOOCV correction can be understood more intuitively using the following equivalent form to Equation (7),

$$\boldsymbol{e}' = \boldsymbol{V} - \boldsymbol{C}' = \frac{\boldsymbol{e}}{\boldsymbol{1}_M - \boldsymbol{h}},$$

which shows that the leave-one-out error is always bigger in magnitude. This is a direct consequence of the full regression overfitting due to the self-influence.

When the regression includes an intercept term, h_n satisfies

$$h_n = \frac{1}{N} + \frac{1}{N-1} D_{\boldsymbol{\Sigma}}(\boldsymbol{x}_n, \bar{\boldsymbol{x}})^2,$$

where $D_{\boldsymbol{\Sigma}}(\boldsymbol{x},\boldsymbol{y}) = (\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{y})$ is the Mahalanobis distance associated to the sample covariance matrix $\boldsymbol{\Sigma}$.

From Equation (7), we guess that the average correction of the continuation value is roughly proportional to M/N. Indeed, it turns out to be case as discussed in the next section.

3. Asymptotic Analysis of Look-ahead Bias

We analyze the asymptotic behavior of the look-ahead bias. The purpose of the analysis in this section is not to provide a correction term to the LSM price. In order to obtain the look-ahead bias-removed price, the LOOLSM method can be directly implemented since the modification from the LSM method and the additional computation are very light. The analysis is provided to retrospectively investigate when and how much the look-ahead bias have contributed to the LSM price.

The look-ahead bias is measured as the price difference between the LSM and LOOLSM methods,

$$\hat{B} = \hat{V}^{[0]}_{\text{LSM}} - \hat{V}^{[0]}_{\text{LOOLSM}}.$$

The set of basis functions $\{f_m : m = 0, 1, \cdots, M-1\}$ is fixed in the L^2 space, and the M basis functions always means $\{f_0, \cdots, f_{M-1}\}$. The approximation of $C^{[i]}(S)$ with the M basis functions is denoted by $C^{[i]}_M(S)$.

We make some assumptions before stating the main result, which are mostly standard ones in the literature. First, we only work with realistic payoff functions that grow moderately and the price of the option is well defined. This is a minimal condition from a practical standpoint, nevertheless we state it for completeness.

Assumption 1. *The payout functions $Z^{[i]}(S)$ are continuous and have finite norms.*

Second, $\hat{C}^{[i]}$, when estimated with the M basis functions, converges to $C^{[i]}_M$ as N goes to infinity. Here, N comes in through the use of N i.i.d. samples for estimating the continuation value function. It is essentially based on the law of large numbers, which only requires weak moment assumptions.

Assumption 2. *$\hat{C}$ converges to C_M as N approaches to infinity.*

The last assumption is more subtle. A complication in pricing arises when the option and continuation values can be arbitrarily close with non-negligible probability, in which case the backward induction becomes numerically unstable. Therefore, we impose conditions to control the behavior around the exercise boundary.

Assumption 3. *Let $P^{[i]}_M(k) = \mathbb{E}[I(\,|C^{[i]}_M(S) - Z^{[i]}(S)| \le k\,)]$ be the probability that the absolute continuation premium does not exceed k. Then, $\lim_{k\to 0} P^{[i]}_M(k) = 0$ and $P^{[i]}_M$ converges to $P^{[i]}$ uniformly.*

It is usually assumed in the literature analyzing the convergence that $P^{[i]}(0) = 0$. Stentoft (2004) explains the condition is necessary for the algorithm to identify the right exercise policy at the limit. Assumption 3 extends it to all Ms as we are interested not only in M at the limit, but in small Ms for practical importance.

THEOREM 1. *Suppose $M \ll N$ and the option can be exercised at t_1 and t_2. Under Assumptions 1, 2 and 3, the look-ahead bias $\hat{B}$ satisfies*

$$\mathbb{E}[\hat{B}] \sim \epsilon + O(\frac{M}{N})$$

for any $\epsilon > 0$.

We prove Theorem 1 in Appendix A. Essentially, the proof shows that look-ahead bias mainly comes from a small neighborhood of the exercise boundary whose volume is $O(M/N)$, whereas the expected bias size is bounded above. The same idea may potentially work for multi-period cases (i.e. $I > 2$), but require stronger assumptions as the effects of propagation is difficult to incorporate. By choosing a very small ϵ, the theorem implies that any realistic bias should decay at least at the rate of M/N. On the contrary, the result should be interpreted with caution as the proof does not provide much clue on how small M/N has to be until we observe the asymptotic behavior in the theorem.

Based on the estimator with Equation (5), Carriere (1996) predicts that their high bias decays at the rate of $1/N$ as N increases:

$$\mathbb{E}[\hat{B}'] = \frac{a}{N} + O(\frac{1}{N^2}) \quad \text{for a constant } a.$$

Indeed, we find a similar pattern when we estimate look-ahead bias as the difference between the LSM and LOOLSM estimators. While we present the empirical results in § 4, here we attempt to provide a theoretical justification.

4. Numerical Results

4.1. Design of Experiments

Three Bermudan option cases are priced to compare the LSM and LOOLSM methods. We present them in an increasing order of the number of the underlying assets: single-stock put options, best-of options on two assets, and basket options on four assets. Therefore, the number of the regressors, M, also increases given the same polynomial orders to include. In all examples, the underlying asset prices $S_j(t)$ follow geometric Brownian motions (GBMs):

$$\frac{dS_j(t)}{S_j(t)} = (r - q_j)dt + \sigma_j\, dW_j(t),$$

where r is the risk-free rate, q_i is the dividend yield, σ_i is the volatility, and $W_j(t)$'s are the standard Brownian motions correlated by $dW_j(t)\,dW_{j'}(t) = \rho_{jj'}\,dt \quad (\rho_{jj}=1)$. The choice of the GBM for the price dynamics has several advantages while it does not over-simplify the problem. It can be easily implemented because exact simulation is possible. More complicated stochastic processes require the Euler scheme, which may exhibit another kind of Monte Carlo bias resulting from the time discretization. The GBM is a standard choice in the literature and we can take advantages of the exact Bermudan option prices reported previously. In order for easy comparison to the true price, we typically report the price offset from the exact value and standard deviation from the results of n_{MC} independent simulation sets:

$$\text{Price Offset} = \mathbb{E}[\hat{V}^{[0]}] - V^{[0]}, \quad \text{Standard Deviation} = \sqrt{\frac{n_{\text{MC}}}{n_{\text{MC}}-1}(E[(\hat{V}^{[0]})^2] - E[\hat{V}^{[0]}]^2)},$$

where $\hat{V}^{[0]}$ is the price from each simulation and $V^{[0]}$ is the exact option price. We use antithetic random variate $(N/2+N/2)$ to reduce variance.

For each case, two experiments are run. The first experiment is to ensure that the LOOLSM method eliminates the look-ahead bias in a similar way the conventional two-simulation-set method does. We run $n_{\text{MC}} = 100$ sets of simulations with $N = 40,000$ paths each and use the following three estimators for comparison:

- **LSM**: the classic LSM estimator.
- **LSM-2**: the two-pass LSM estimator. The exercise policy computed from independent $N = 40,000$ paths is applied to the payoff valuation with the original simulation set.
- **LOOLSM**: the LOOLSM estimator.

Using the same simulation paths for the payoff valuation across the three methods works as a control variate to reduce the variability of the measured bias, i.e., price difference between methods. Additionally, we price the corresponding European options with Monte Carlo method using the simulation set.

The second experiment is to validate the asymptotic behavior of the bias in Theorem 1. Here, we run the LSM and LOOLSM methods with varying N and M. We first generate a pool of 1.44×10^6 Monte Carlo paths and split them into n_{MC} chunks for $n_{\text{MC}} = 10$, 20, 30, 40, 60, 120, 240, and 720. Therefore, one simulation set comprises $N = 1.44 \times 10^6 / n_{\text{MC}}$ paths and the price offset is computed from the n_{MC} prices. By varying N among the same path pool, we control the Monte Carlo variance as much as possible and make the simulation size N the most important factor to determine the look-ahead bias. At the same time, the number of regressors, M, is also varied by including polynomials of higher terms. In this way, we measure the look-ahead bias as a function of M/N. For the second experiment, we use the European option price as a control variate in order to make better estimates of the prices when n_{MC} is small. However, the measured look-ahead bias is not affected by the control variate because the same amount is adjusted for both LSM and LOOLSM prices.

4.2. Case 1: Single-stock Put Option

We start with Bermudan put options on single stock:

$$Z^{[i]}(S) = e^{-rt_i}\max(K - S_1, 0),$$

with the parameter set tested in Feng and Lin (2013):

$$S_1(0) = 100,\ \sigma_1 = 20\%,\ r = 5\%,\ q_1 = 2\%,\ t_i = \frac{i}{5},\ I = 5\ (T = 1).$$

The option prices are obtained for the strikes, $K = 80$, 90, 100, 110, and 120. We implement the binomial tree method to calculate exact prices and validate the accuracy by comparing the price for $K = 100$ to that in Feng and Lin (2013). For the regressors, we use the following set:

$$X^{[i]}(S) = (1,\ Z^{[i]}(S),\ S_1,\ S_1^2, \cdots).$$

We use the first $M = 5$ functions (up to S_1^3) for the first experiment and $M = 4$, 8, and 12 for the second.

Table 1 reports the result of the first experiment. As expected, the LOOLSM produces similar results to those of the LSM-2 method. The LSM results are slightly higher than the other two methods, implying that look-ahead bias is small. To show the statistical significance of the look-ahead bias, we separately report the statistics of the bias in Table 2. It should be noted that the bias measured against the LOOLSM method has much smaller deviation compared to that measured against the LSM-2. It is intuitive because the LOOLSM method requires no extra random number whereas the LSM-2 method requires another independent simulation set.

The result of the second experiment is shown in Figure 2. The top plot shows the price offset of the LSM and LOOLSM methods as a function of M/N for different M values. It demonstrates how the LSM and LOOLSM prices converge to each other as N is increased for a fixed M. The LSM price converges from above and the LOOLSM converges from below, indicating that the LOOLSM is a low-biased compared to the convergent value for a given M. The bottom plot show the look-ahead bias as a function of M/N. Notably, the data from the three M values form a clear linear pattern, confirming the asymptotic behavior in Theorem 1. While the figure is for one specific option ($K = 80$), other options in the case exhibit the same patterns.

This single-asset case is similar to the test case used in Longstaff and Schwartz (2001) to demonstrate that the look-ahead bias is negligible. Indeed, our finding is consistent. In light of the asymptotic behavior, however, the case used in Longstaff and Schwartz (2001) has very small ratio, $M/N = 4/10^5$, which is the smallest among the five cases discussed in the paper. The asymptotic analysis indicates that the LSM price possibly goes above the true price when larger basis set, e.g., $M = 8$ or 12, is used even with big simulation size N.

Table 1 Results for single-stock Bermudan put options in § 4.2. We use $N = 40,000$, $M = 5$. The 'Exact' columns report true option prices while the other columns report the price offset and standard deviation from $n_{\text{MC}} = 100$ simulation results. All values are rounded to three decimal places.

	Bermudan				European	
K	Exact	LSM	LSM-2	LOOLSM	Exact	MC
80	0.856	-0.002 ± 0.014	-0.003 ± 0.014	-0.003 ± 0.014	0.843	-0.002 ± 0.015
90	2.786	-0.002 ± 0.019	-0.004 ± 0.019	-0.003 ± 0.018	2.714	-0.002 ± 0.024
100	6.585	-0.001 ± 0.020	-0.003 ± 0.020	-0.003 ± 0.020	6.330	-0.000 ± 0.029
110	12.486	-0.009 ± 0.024	-0.011 ± 0.023	-0.012 ± 0.024	11.804	-0.001 ± 0.026
120	20.278	-0.014 ± 0.033	-0.014 ± 0.033	-0.016 ± 0.033	18.839	-0.003 ± 0.018

Table 2 Results for the first experiment on single-stock Bermudan put options in § 4.2. The columns report the difference between the price of each method and LSM price, and its error estimate. The negative values are because high-bias is removed. All values are rounded to four decimal places.

K	LSM-2	LOOLSM
80	-0.0013 ± 0.0026	-0.0011 ± 0.0005
90	-0.0017 ± 0.0035	-0.0014 ± 0.0007
100	-0.0025 ± 0.0072	-0.0024 ± 0.0014
110	-0.0021 ± 0.0088	-0.0024 ± 0.0011
120	-0.0003 ± 0.0086	-0.0022 ± 0.0013

4.3. Case 2: Best-of Option on Two Assets

We price best-of (or rainbow) call options on two assets:

$$Z^{[i]}(S) = e^{-rt_i} \max(\max(S_1, S_2) - K, 0),$$

with the parameter set tested in Glasserman (2003) and Andersen and Broadie (2004),

$$K = 100,\ \sigma_j = 20\%,\ r = 5\%,\ q_j = 10\%,\ \rho_{j\neq j'} = 0,\ t_i = \frac{i}{3},\ I = 9\ (T = 3).$$

We price the options with three initial asset prices, $S_1(0) = S_2(0) = 90$, 100, and 110. We use the following basis functions for the first experiment:

$$X^{[i]}(S) = (1,\ Z^{[i]}(S),\ S_1,\ S_2,\ S_1^2,\ S_1S_2,\ S_2^2,\ S_1^3,\ S_1^2S_2,\ S_1S_2^2,\ S_2^3) \quad (M = 11),$$

For the second experiment, $M = 4$, 7, and 11 are used, which correspond to the terms including up to linear, quadratic, and cubic polynomial terms, respectively. We use the exact Bermudan option

Figure 2 The price offset (top) and look-ahead bias (bottom) as functions of M/N for the single-stock put option with $K = 80$ in § 4.2. On the top, given the fixed value of M/N, the higher value corresponds to the LSM method and the lower to the LOOLSM.

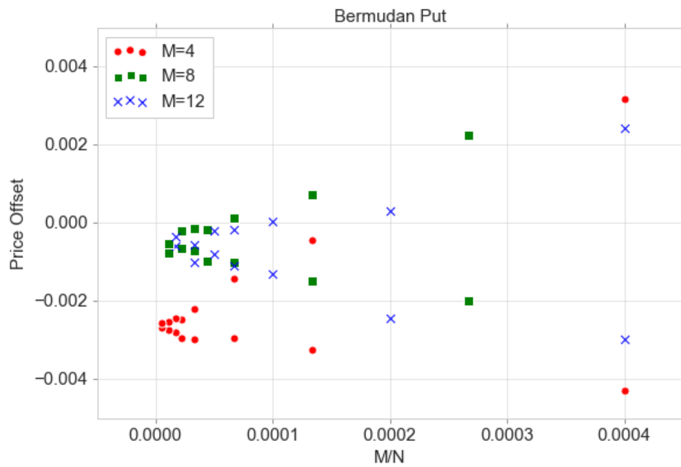

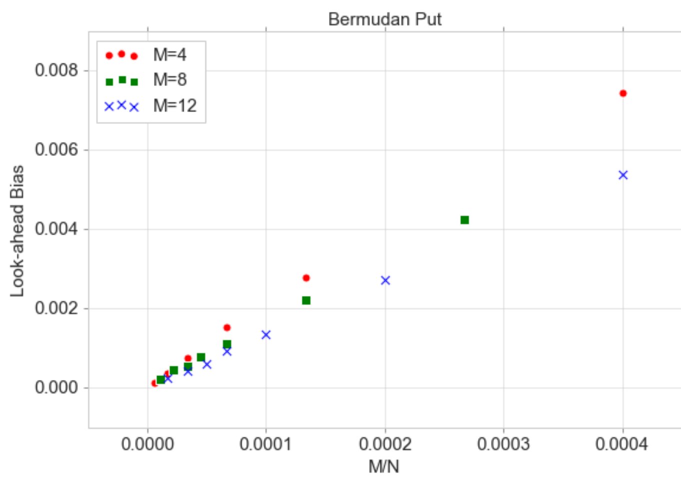

prices from Andersen and Broadie (2004) and compute the exact European option prices from the analytic solutions expressed in terms of the bivariate cumulative normal distribution (Rubinstein 1991).

The results are shown in Table 3 and Figure 3. The look-ahead bias in the LSM method becomes more pronounced, but the LSM price is still lower than the true price. This is primarily because the exercise boundary of the best-of option is highly non-linear as observed in Glasserman (2003). As seen in Figure 3, the suboptimal bias quickly decrease as M is increased. Nevertheless, the look-ahead bias is clearly proportional to M/N regardless of the sub-optimality level.

Table 3 Results for best-of the Bermudan best-of call options in § 4.3. We use $N = 40,000$ and $M = 11$. The 'Exact' columns report true option prices while the other columns report the price offset and standard deviation from $n_{\text{MC}} = 100$ simulation results. All values are rounded to three decimal places.

	Bermudan				European	
$S_j(0)$	Exact	LSM	LSM-2	LOOLSM	Exact	MC
90	8.075	-0.020 ± 0.055	-0.036 ± 0.056	-0.035 ± 0.054	6.655	0.011 ± 0.062
100	13.902	-0.036 ± 0.060	-0.052 ± 0.062	-0.054 ± 0.058	11.196	0.011 ± 0.078
110	21.345	-0.040 ± 0.065	-0.062 ± 0.068	-0.059 ± 0.064	16.929	0.013 ± 0.096

Figure 3 The price offset (top) and look-ahead bias (bottom) as functions of M/N for the best-of call option with $S_j(0) = 100$ in § 4.3. On the top, given the fixed value of M/N, the higher value corresponds to the LSM method and the lower to the LOOLSM.

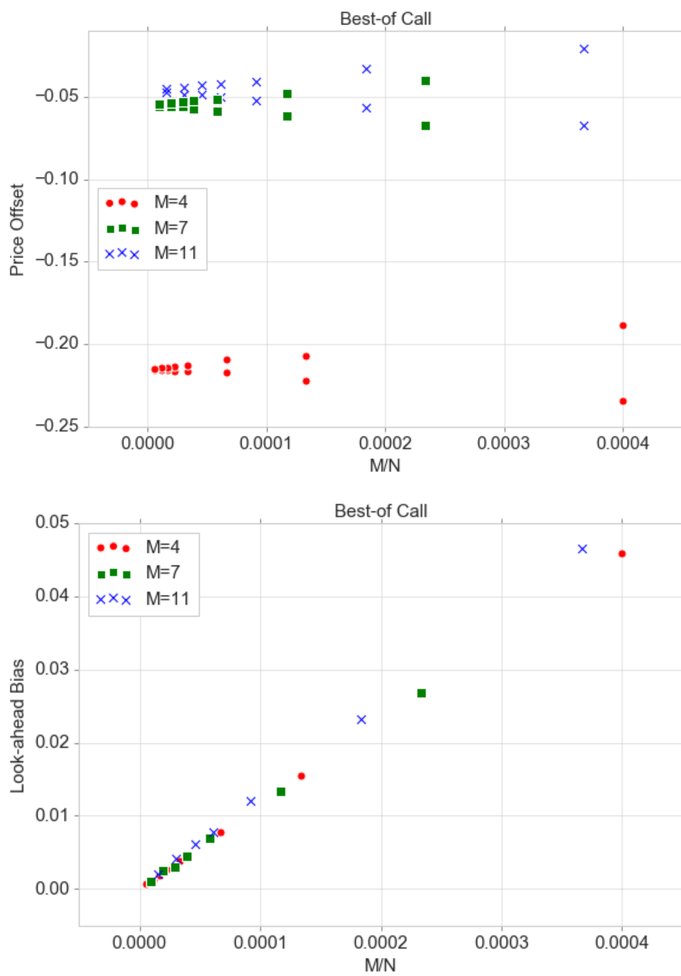

4.4. Case 3: Basket Option on Four Assets

Finally, we price Bermudan calls option on a basket of four stocks,

$$Z^{[i]}(S) = e^{-rt_i} \max\left(\frac{S_1 + S_2 + S_3 + S_4}{4} - K, 0\right).$$

Table 4 Results for four-asset basket options in § 4.4. We use $N = 40,000$ and $M = 16$. The 'Exact' columns report true option prices while the other columns report the price offset and standard deviation from $n_{\text{MC}} = 100$ simulation results. All values are rounded to three decimal places.

K	Exact	LSM	LSM-2	LOOLSM	European
60	47.481	0.233 ± 0.223	-0.205 ± 0.213	-0.209 ± 0.196	0.012 ± 0.309
80	36.352	0.230 ± 0.255	-0.174 ± 0.244	-0.158 ± 0.235	0.012 ± 0.316
100	28.007	0.235 ± 0.237	-0.117 ± 0.238	-0.109 ± 0.231	0.012 ± 0.309
120	21.763	0.226 ± 0.236	-0.084 ± 0.245	-0.080 ± 0.229	0.013 ± 0.293
140	17.066	0.213 ± 0.224	-0.086 ± 0.222	-0.075 ± 0.223	0.015 ± 0.275

with the parameter set tested by Krekel et al. (2004) and Choi (2018) in the context of the European payoff,

$$S_j(0) = 100,\ \sigma_j = 40\%,\ r = q_j = 0,\ \rho_{j \neq j'} = 0.5,\ t_i = \frac{i}{2},\ I = 10\,(T = 5).$$

We price the options with a range of strikes, $K = 60, 80, 100, 120$, and 140. Because the underlying assets are not paying dividend, the optimal exercise policy for the option holder is not to exercise the option until maturity, so the European option price is equal to the Bermudan's. Therefore, we use Choi (2018) for the exact prices. For the regressors, the polynomials up to degree 2 are used for the first experiment:

$$X^{[i]}(S) = (1,\, Z^{[i]}(S),\, S_j, \cdots, S_j^2, \cdots, S_j S_{j'}, \cdots) \quad \text{for} \quad 1 \le j < j' \le 4 \quad (M = 16),$$

and the subsets, $M = 6, 10$, and 16, are used for the second experiment.

Table 4 and Figure 4 report the results. In this example, the LSM method noticeably overprice the option for all strike prices whereas the LSM-2 and LOOLSM prices are consistent. Unlike the best-of option case, the suboptimal level is unchanged for increasing M because the payoff function is a linear combination of the asset prices and, therefore, the exercise boundary is accurately captured with the linear basis functions only ($M = 6$). The look-ahead bias collapses into a function of M/N although the asymptotic behavior clearly appears when N is very large (see the inset of the bottom plot of Figure 4).

5. Conclusion

This article study a new efficient approach for removing the look-ahead bias of the LSM algorithm (Longstaff and Schwartz 2001). It is natural to apply the LOOCV in this context, a well-known cross-validation technique in statistical learning. The resulting LOOLSM estimator can be implemented with little extra computational cost. We validate this approach with several examples. In particular, we demonstrate that the LSM price can be biased high for multi-asset options and that the LOOLSM algorithm can effectively eliminate look-ahead bias. Finally, we discuss the asymptotic behavior of look-ahead bias, measured as the difference between the LSM and LOOLSM estimators. We uncover interesting connection of the bias decay not only to the number of Monte Carlo paths, but also to the number of regressors.

Figure 4 The price offset (top) and look-ahead bias (bottom) as functions of M/N for the best-of call option with $K = 100$ in § 4.4. On the top, given the fixed value of M/N, the higher value corresponds to the LSM method and the lower to the LOOLSM.

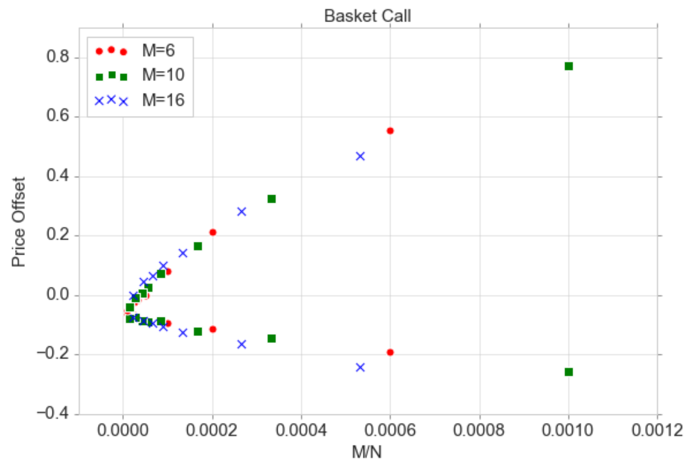

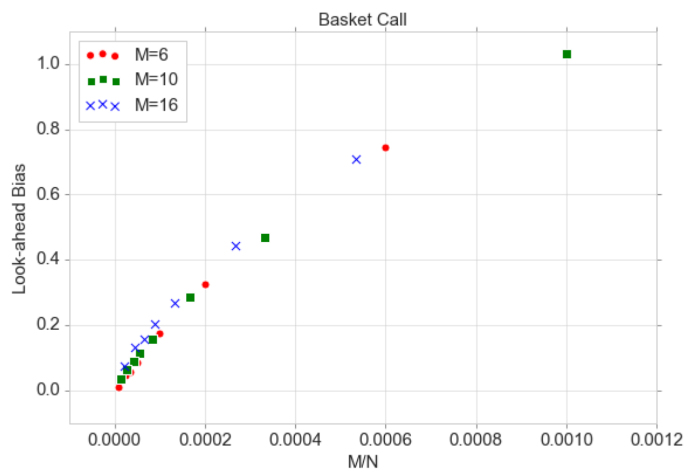

Appendix A: Proof of Theorem 1

THEOREM 1. *Suppose $M \ll N$ and the option can be exercised at t_1 and t_2. Under Assumptions 1, 2 and 3, the look-ahead bias $\hat{B}$ satisfies*

$$\mathbb{E}[\hat{B}] \sim \epsilon + O(\frac{M}{N})$$

for any $\epsilon > 0$.

From Equation (3),

$$\hat{V}^{[1]}_{n,LSM} = I[Z_n \le \hat{C}_n](V_n - Z_n) + Z_n,$$
$$\hat{V}^{[1]}_{n,LOOLSM} = I[Z_n \le \hat{C}'_n](V_n - Z_n) + Z_n.$$

Here, we omitted the superscript $[i]$ to simply write as $Z_n = Z_n^{[1]}$, $C_n = C_n^{[1]}$ and $V_n = V_n^{[2]}$. In the signle period case, the look-ahead bias can only arise at time t_1. Therefore,

$$\begin{aligned}\hat{B}_n &= \hat{V}_{n,LSM}^{[1]} - \hat{V}_{n,LOOLSM}^{[1]} \\ &= (I[Z_n \le \hat{C}_n] - I[Z_n \le \hat{C}'_n])(V_n - Z_n) \\ &= (I[D_n^+] - I[D_n^-])(V_n - Z_n),\end{aligned}$$

where $D_n^+ = \hat{C}'_n < Z_n \le \hat{C}_n$ and $D_n^- = \hat{C}'_n \ge Z_n > \hat{C}_n$. In other words, D_n^+ and D_n^- corresponds to the situation where the LSM algorithm continues and exercises due to the foresight bias when it should have exercised and continued, respectively. The second term is the price impact as a result. By applying Equation (7), we get

$$\begin{aligned}D_n^+ &\iff 0 \le \hat{C}_n - Z_n < \hat{C}_n - \hat{C}'_n \\ &\iff 0 \le \hat{C}_n - Z_n < \frac{h_n}{1-h_n}(V_n - \hat{C}_n) \\ &\iff 0 \le \hat{C}_n - Z_n < h_n(V_n - Z_n).\end{aligned}$$

Similarly, $D_n^- \iff 0 > \hat{C}_n - Z_n \ge h_n(V_n - Z_n)$.

From Assumption 1, there exists a compact set $\mathcal{U}$ in the space of state vectors such that

$$\mathbb{E}[I[S \in \mathcal{U}^c] \cdot |V(S) - Z(S)|] < \frac{\epsilon}{2} \tag{8}$$

for a given $\epsilon > 0$. Since $\mathcal{U}$ is compact, there exist c_0 such that $|V(S) - Z(S)| < c_0$ for $S \in \mathcal{U}$. Note that $\mathcal{U}$ and c_0 depend only on ϵ, but not on the simulation parameters N and M. Then,

$$\mathcal{A}_n = \left\{ S \,\middle|\, |\hat{C}_n - Z_n| < c_0 h_n \right\}$$

is a superset of D_n^+ and D_n^- in $\mathcal{U}$, i.e. $((D_n^+ \cup D_n^-) \cap \mathcal{U}) \subset (\mathcal{A}_n \cap \mathcal{U})$.

From Assumption 2 and 3, there exists $c_1 > 0$ such that

$$\max(P_M(\frac{2}{c_1}), \mathbb{E}\left[I[|\hat{C}(S) - C_M(S)| \ge \frac{1}{c_1}]\right]) < \frac{\epsilon}{4c_0}$$

for previous chosen c_0, ϵ, and for any M and sufficiently large N. Then,

$$\begin{aligned}\mathbb{E}\left[I[|\hat{C}_n - Z_n| < \frac{1}{c_1}]\right] &\le \mathbb{E}\left[I[\max(|\hat{C}_n - Z_n|, |\hat{C}_n - C_{M,n}|) < \frac{1}{c_1}|]\right] + \mathbb{E}\left[I[|\hat{C}_n - C_{M,n}| \ge \frac{1}{c_1}|\right] \\ &< \mathbb{E}\left[I[|C_{M,n} - Z_n| < \frac{2}{c_1}\right] + \frac{\epsilon}{4c_0} < \frac{\epsilon}{2c_0}.\end{aligned} \tag{9}$$

Putting these all together, we have

$$\begin{aligned}
\mathbb{E}[\hat{B}] = \mathbb{E}[\hat{B}_n] &= \mathbb{E}[I[S_n \in \mathcal{U}^c]\hat{B}_n] + \mathbb{E}[I[S_n \in \mathcal{U}]\hat{B}_n] \\
&< \mathbb{E}[I[S_1 \in \mathcal{U}^c]|V_n - Z_n|] + \mathbb{E}[I[S_n \in \mathcal{A} \cap \mathcal{U}]|V_n - Z_n|] \\
&< \frac{\epsilon}{2} + c_0\mathbb{E}[I[S_n \in \mathcal{A} \cap \mathcal{U}]] \qquad \text{from Equation (8)} \\
&< \frac{\epsilon}{2} + c_0(\mathbb{E}[I[|\hat{C}_n - Z_n| < \frac{1}{c_1}]] + \mathbb{E}[I[\frac{1}{c_1} \le |\hat{C}_n - Z_n| < c_0 h_n]]) \\
&< \frac{\epsilon}{2} + c_0(\frac{\epsilon}{2c_0} + \mathbb{E}[c_0 c_1 h_n]) \qquad \text{from Equation (9)} \\
&\le \epsilon + c_0^2 c_1 \frac{M}{N}.
\end{aligned}$$

The last inequality follows from $\mathbb{E}[h_n] \le M/N$. This completes the proof.

Appendix B: Comparison with look-ahead bias in Fries (2005, 2008)

In Fries (2005, 2008), the look-ahead bias at t_i is defined as the option value on the Monte Carlo error $\xi^{[i]} = \hat{C}^{[i]} - C^{[i]}$ in the estimation of the continuation value.

Let $d^{[i]} = Z^{[i]} - C^{[i]}$. The value of the option $E^{[i]}$ can be decomposed into two sources of bias,

$$\begin{aligned}
E^{[i]} &= \mathbb{E}[\max(Z^{[i]}, C^{[i]} + \xi^{[i]})] - \max(Z^{[i]}, C^{[i]}) \\
&= Z^{[i]} + \mathbb{E}[I[\xi^{[i]} \ge d^{[i]}](\xi^{[i]} - d^{[i]})] - \max(Z^{[i]}, C^{[i]}) \\
&= \underbrace{\mathrm{Cov}(I[\xi^{[i]} \ge d^{[i]}], \xi^{[i]})}_{A} + \underbrace{Z_i - d^{[i]}\mathbb{E}[I[\xi^{[i]} \ge d^{[i]}]] - \max(Z^{[i]}, C^{[i]})}_{B}.
\end{aligned}$$

The covariance term A is the look-ahead bias, while the term B is the suboptimal bias. Under the Gaussian error assumption $\xi^{[i]} \sim N(0, \sigma^2)$, the look-ahead bias term $B^{[i]}_{\text{Fries}}$ has a closed-form expression, $B^{[i]}_{\text{Fries}} = \sigma\,\phi(d^{[i]}/\sigma)$, where ϕ is the probability density function of the standard normal distribution.

This is closely related to the definition in Equation (6), but with a crucial difference. By using the same notations, Equation (6) can be rewritten as

$$B_i = \mathrm{Cov}(I[\xi^{[i]} \ge d^{[i]}], \zeta^{[i]}),$$

where $\zeta^{[i]} = \hat{V}^{[i+1]} - C^{[i]}$. This error is different from $\xi^{[i]}$ in that $\zeta^{[i]}$ is the deviation at the sample level, whereas $\xi^{[i]}$ is at the resulting estimator level. If we denote the Monte Carlo error of the LOOLSM estimator by $\delta^{[i]} = \hat{C}'^{[i]} - C^{[i]}$, we get

$$\begin{aligned}
\xi^{[i]} &= \hat{C}^{[i]} - C^{[i]} \\
&= \delta^{[i]} + h^{[i]}(\hat{V}^{[i+1]} - (C^{[i]} + \delta^{[i]})) \quad \text{(from Equation (??))} \\
&= (1 - h^{[i]})\delta^{[i]} + h^{[i]}\zeta^{[i]}.
\end{aligned}$$

Therefore, the total Monte Carlo error $\xi^{[i]}$ is a weighted average of the two independent error terms, $\delta^{[i]}$ and $\zeta^{[i]}$. As the LSM method follows Equation (3), our look-ahead bias term correctly captures the contribution of $\zeta^{[i]}$ to the exercise decision. On the contrary, the look-ahead bias in Fries (2008) includes the contribution from $\delta^{[i]}$ as well since it defines the bias term by using the convexity of the maximum function of the total Monte Carlo error $\epsilon^{[i]}$. Such term is not a source of the high bias in the original LSM formulation.

Acknowledgments

The research of Jaehyuk Choi was supported by Bridge Trust Asset Management Research Fund (2016-2017, Deep learning applied to finance).

References

Andersen L, Broadie M (2004) Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options. *Management Science* 50(9):1222–1234, ISSN 0025-1909, URL http://dx.doi.org/10.1287/mnsc.1040.0258.

Bacinello AR, Biffis E, Millossovich P (2010) Regression-based algorithms for life insurance contracts with surrender guarantees. *Quantitative Finance* 10(9):1077–1090, ISSN 1469-7688, URL http://dx.doi.org/10.1080/14697680902960242.

Beveridge C, Joshi M, Tang R (2013) Practical policy iteration: Generic methods for obtaining rapid and tight bounds for bermudan exotic derivatives using Monte Carlo simulation. *Journal of Economic Dynamics and Control* 37(7):1342–1361, URL http://dx.doi.org/10.1016/j.jedc.2013.03.004.

Boyle PP (1988) A Lattice Framework for Option Pricing with Two State Variables. *The Journal of Financial and Quantitative Analysis* 23(1):1–12, URL http://dx.doi.org/10.2307/2331019.

Boyle PP, Evnine J, Gibbs S (1989) Numerical Evaluation of Multivariate Contingent Claims. *The Review of Financial Studies* 2(2):241–250.

Brennan MJ, Schwartz ES (1977) The Valuation of American Put Options. *The Journal of Finance* 32(2):449–462, URL http://dx.doi.org/10.2307/2326779.

Broadie M, Glasserman P (1997) Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control* 21(8-9):1323–1352, URL http://dx.doi.org/10.1016/S0165-1889(97)00029-8.

Broadie M, Glasserman P (2004) A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance* 7(4):35–72, URL http://dx.doi.org/10.21314/JCF.2004.117.

Carriere JF (1996) Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19(1):19–30, URL http://dx.doi.org/10.1016/S0167-6687(96)00004-2.

Choi J (2018) Sum of all Black-Scholes-Merton models: An efficient pricing method for spread, basket, and Asian options. *Journal of Futures Markets* 38(6):627–644, URL http://dx.doi.org/10.1002/fut.21909.

Cox JC, Ross SA, Rubinstein M (1979) Option pricing: A simplified approach. *Journal of Financial Economics* 7(3):229–263, URL http://dx.doi.org/10.1016/0304-405X(79)90015-1.

Feng L, Lin X (2013) Pricing Bermudan Options in Lévy Process Models. *SIAM Journal on Financial Mathematics* 4(1):474–493, URL http://dx.doi.org/10.1137/120881063.

Fries CP (2005) Foresight Bias and Suboptimality Correction in Monte-Carlo Pricing of Options with Early Exercise: Classification, Calculation and Removal. *Available at SSRN* URL http://christian-fries.de/finmath/foresightbias/.

Fries CP (2008) Foresight Bias and Suboptimality Correction in Monte-Carlo Pricing of Options with Early Exercise. *Progress in Industrial Mathematics at ECMI 2006*, 645–649 (Springer), URL http://dx.doi.org/10.1007/978-3-540-71992-2_107.

Fu MC, Laprise SB, Madan DB, Su Y, Wu R (2001) Pricing American options: a comparison of Monte Carlo simulation approaches. *Journal of Computational Finance* 4(3):39–88, URL http://dx.doi.org/10.21314/JCF.2001.066.

Glasserman P (2003) Chapter 8. Pricing American Options. *Monte Carlo Methods in Financial Engineering*, 421–478 (New York), 2003 edition edition.

Glasserman P, Yu B (2004) Number of paths versus number of basis functions in American option pricing. *The Annals of Applied Probability* 14(4):2090–2119, URL http://dx.doi.org/10.1214/105051604000000846.

Haugh MB, Kogan L (2004) Pricing American options: A duality approach. *Operations Research* 52(2):258–270.

He H (1990) Convergence from Discrete- to Continuous-Time Contingent Claims Prices. *The Review of Financial Studies* 3(4):523–546.

Hoaglin DC, Welsch RE (1978) The Hat Matrix in Regression and ANOVA. *The American Statistician* 32(1):17–22, ISSN 0003-1305, URL http://dx.doi.org/10.1080/00031305.1978.10479237.

Huang YT, Kwok YK (2016) Regression-based Monte Carlo methods for stochastic control models: Variable annuities with lifelong guarantees. *Quantitative Finance* 16(6):905–928, ISSN 1469-7688, URL http://dx.doi.org/10.1080/14697688.2015.1088962.

Kolodko A, Schoenmakers J (2006) Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics* 10(1):27–49, ISSN 1432-1122, URL http://dx.doi.org/10.1007/s00780-005-0168-5.

Krekel M, de Kock J, Korn R, Man TK (2004) An analysis of pricing methods for basket options. *Wilmott Magazine* 2004(7):82–89.

Longstaff FA, Schwartz ES (2001) Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies* 14(1):113–147, URL http://dx.doi.org/10.1093/rfs/14.1.113.

Nadarajah S, Margot F, Secomandi N (2017) Comparison of least squares Monte Carlo methods with applications to energy real options. *European Journal of Operational Research* 256(1):196–204, ISSN 0377-2217, URL http://dx.doi.org/10.1016/j.ejor.2016.06.020.

Rubinstein M (1991) Somewhere Over the Rainbow. *Risk* 1991(11):63–66.

Stentoft L (2004) Convergence of the Least Squares Monte Carlo Approach to American Option Valuation. *Management Science* 50(9):1193–1203, URL http://dx.doi.org/10.1287/mnsc.1030.0155.

Tilley JA (1993) Valuing American options in a path simulation model. *Transactions of the Society of Actuaries* 45:499–550.

Tsitsiklis JN, Van Roy B (2001) Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12(4):694–703, URL http://dx.doi.org/10.1109/72.935083.