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

Jeechul Woo
University of Illinois Urbana-Champaign

Chenru Liu
Stanford University

Jaehyuk Choi*
Peking Univeristy

May 2021

Abstract

The least square Monte Carlo (LSM) algorithm proposed by Longstaff and Schwartz (2001) is widely used for pricing American options. The LSM estimator contains undesirable look-ahead bias, and the conventional technique of removing it necessitates doubling simulations. We present the leave-one-out LSM (LOOLSM) algorithm for efficiently eliminating look-ahead bias. We also show that look-ahead bias is asymptotically proportional to the regressors-to-simulation paths ratio. Our findings are demonstrated with several option examples, including the multi-asset cases that the LSM algorithm significantly overvalues. The LOOLSM method can be extended to other regression-based algorithms that improve the LSM method.

Keywords: American option, Least square Monte Carlo, Look-ahead bias, Leave-one-out-cross-validation
JEL Classification: G12, G17

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

Jeechul Woo[a], Chenru Liu[c], Jaehyuk Choi[b,1,*]

[a]*Department of Mathematics, University of Illinois Urbana-Champaign, Champaign, IL 61820, US*
[b]*Peking University HSBC Business School,*
University Town, Nanshan, Shenzhen 518055, China
[c]*Department of Management Science and Engineering, Stanford University, Stanford, CA 94305, US*

Abstract

The least square Monte Carlo (LSM) algorithm proposed by Longstaff and Schwartz (2001) is widely used for pricing American options. The LSM estimator contains undesirable look-ahead bias, and the conventional technique of removing it necessitates doubling simulations. We present the leave-one-out LSM (LOOLSM) algorithm for efficiently eliminating look-ahead bias. We also show that look-ahead bias is asymptotically proportional to the regressors-to-simulation paths ratio. Our findings are demonstrated with several option examples, including the multi-asset cases that the LSM algorithm significantly overvalues. The LOOLSM method can be extended to other regression-based algorithms that improve the LSM method.

Keywords: American option, Least square Monte Carlo, Look-ahead 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 are 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. Among the regression-based methods, the least square Monte Carlo (LSM) algorithm proposed by Longstaff and Schwartz (2001) is the most popular for its simplicity and efficiency. Fu et al. (2001) and Glasserman (2003) provide comprehensive reviews of the implementation and comparison of simulation-based methods.

The LSM method is essential 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 notes are the effective buyers of the Bermudan option to redeem notes early. On the contrary, investors, be them individual or institutional, are the effective sellers of the option. They receive the premium in the form of an enhanced yield compared with that of non-callable notes with the same structure. Because a multi-factor model is required for the underlying assets as well as the yield curve term structure, the use of Monte Carlo simulation along with the LSM method is inevitable for pricing and risk-managing such notes. Like many previous studies of this topic (Kolodko and Schoenmakers, 2006; Beveridge et al., 2013), this study is motivated and developed in the context of callable structured notes.

1.2. Biases in the LSM Method

In simulation-based methods including the LSM method, there are two main sources of bias, which run in opposite directions. Low-side bias is related to suboptimal exercise decisions owing to various approximations adopted in the method. In the LSM method, for example, finite basis functions cannot fully represent the conditional payoff function. The resulting exercise policy deviates from the most optimal one and therefore leads to a lower option price. For this reason, it is also called *suboptimal* bias. High-side bias comes from using one simulation set for both the exercise decision and the payoff valuation. As explained by Broadie and Glasserman (1997), this practice creates a fictitious positive correlation between exercise decisions and future payoffs; the algorithm is more likely to continue (exercise) precisely when the *future* payoff in the simulation is higher (lower). For this reason, it is called *look-ahead* or *foresight* bias. The LSM estimator has both low- and high-side 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 construct both low- and high-biased estimators to form a confidence interval for the true option price.

In callable note markets, look-ahead bias is more dangerous than suboptimal bias, and look-ahead bias being mixed with suboptimal bias is a significant drawback of the LSM estimator. This is closely related to the fact that buyers (financial institutions) act as market makers and pricing agents. Because buyers have to risk-manage and optimally exercise the option, they typically use the LSM method. Sellers (investors) usually hold the note until maturity without hedging and therefore are less sensitive to accurate valuation. From buyers' perspective, look-ahead bias is malicious because it wrongly inflates the option premium they pay. No matter how well the option is delta-hedged, the option value attributed from 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. Suboptimal bias, on the contrary, is benign. Although it deflates the option value, the gain realized through delta-hedging under the suboptimal exercise policy is just as much as the deflated option value. In short, option buyers *get what they pay for*. The only downside of suboptimal bias is for buyers to lose trades to competitors who bid a higher (more optimal) option premium. Therefore, buyers prefer the low-biased estimator to ensure that the premium they pay is lower than the true value. However, look-ahead bias mixed in the LSM method makes a *conservative* valuation difficult for buyers.

A standard technique for eliminating look-ahead bias is to calculate the exercise decision by using an additional independent set of Monte Carlo paths, thereby eliminating the correlation between the exercise decision and simulated payoff. While this two-pass approach removes look-ahead bias, it comes at the cost of doubling the computational cost, which is already heavy because the simulation of stochastic processes frequently requires the time-discretized Euler scheme. The design of the LSM estimator to include the biases in both directions primarily aims to retain computational efficiency rather than 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 to be low-biased.

However, researchers and practitioners have raised concerns that look-ahead bias may not be small in multi-asset problems, where the simulation has to be the method of last resort. The numerical results of Létourneau and Stentoft (2014); Fabozzi et al. (2017) show that look-ahead bias increases when the simulation size becomes smaller or the polynomial order of the basis becomes higher. Carriere (1996) and Fries (2005) remark the same. Practitioners in the structured notes market also observe that, when higher-order regression variables are used to better capture the exercise boundary (i.e., reduce suboptimal bias), look-ahead bias also increases. Given the desire to keep the one-pass LSM implementation for computational efficiency, they are reluctant to include higher-order terms in the LSM regression in fear of overpricing. It is possible to check the validity of the LSM price against several methods to estimate both the lower and the 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 too 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 look-ahead bias in the LSM estimator and develop an efficient algorithm for removing it.

1.3. Contribution of this Study

In this study, we present an efficient approach for removing look-ahead bias, motivated by the cross-validation practice in statistical learning. Standard practice is to separate the datasets for training and testing to avoid overfitting. In this 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). Similarly, using an independent simulation set for the exercise policy corresponds to the hold-out method, one of the simplest cross-validation techniques.

Among advanced cross-validation techniques, we recognize that leave-one-out cross-validation (LOOCV) fits with 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, it is well known that the corrections from the full regression on all samples can be computed altogether with a simple linear algebra operation (Hastie et al., 2009, § 7.10). Therefore, our new leave-one-out LSM (LOOLSM) algorithm eliminates look-ahead bias from the LSM method without incurring an extra computational cost. The LOOLSM method can thus be understood as an extension of the low-biased estimator of Broadie and Glasserman (1997) in the sense that self-exclusion is conducted on all simulation paths rather than on each state separately. By using the LOOLSM method, practitioners can therefore reliably obtain the low-biased price—even with higher-order regression basis functions. The LOOLSM algorithm can also be applied along with other regression methods proposed to improve

4

least squares regression (Tompaidis and Yang, 2014; Chen et al., 2019; Ibáñez and Velasco, 2018; Fabozzi et al., 2017; Belomestny, 2011; Ludkovski, 2018).

Furthermore, this study contributes to the line of research dealing with the convergence of the LSM algorithm, which is a problem of fundamental importance given the popularity of the method. Several authors theoretically analyze the convergence of the LSM method; Clément et al. (2002) prove the convergence of the LSM price for a fixed set of regressors based on the central limit theorem. Stentoft (2004) analyzes the convergence rate of the continuation value function when the number of regressors M also goes to infinity. Glasserman and Yu (2004) discuss how quickly the simulation size N has to grow relative to M to achieve uniform convergence. Zanger (2018) estimates the stochastic component of the error for a general class of approximation architecture.

Previous studies focus primarily on the convergence of the continuation value functions in the L^2 space. By construction, they do not analyze the convergence rate of look-ahead bias specific to the LSM method, in which the estimated continuation value functions are evaluated for the training samples. We bridge this research gap. In particular, we formulate look-ahead bias as the difference between the LSM and LOOLSM prices, with which we theoretically analyze its convergence rate and derive the upper bounds in Theorem 1. Empirically, the formulation provides a way to measure look-ahead bias that is more robust to Monte Carlo noise. We conduct numerical studies for options whose true prices are known and obtain results consistent with our theoretical findings.

To the best knowledge of the authors, previous works estimating look-ahead bias, theoretically or empirically, are scarce and those correcting such bias are rare. Carriere (1996) predicts that the high-side bias of the estimator asymptotically scales to $1/N + O(1/N^2)$. Our analysis and simulation results not only reaffirm this observation, but also show that any realistic look-ahead bias decays at the rate of M/N at least. Fries (2005, 2008) formulates 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 with these studies, our LOOLSM method does not depend on any model assumption and more accurately targets look-ahead bias in the LSM setting.

Beyond the American option pricing, our new method can be applied to various stochastic control problems in finance where least squares regression is used to approximate the optimal strategy. For examples, 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 convergence rate of look-ahead bias in the LSM method. In § 4, the numerical results are presented. Finally, § 6 concludes.

2. Method

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

- The option can be exercised at a discrete time set $\{0 < t_1 < \cdots < t_I = T\}$. As is customary, we 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. We denote the value at t_i by $S^{[i]} = S(t_i)$.

- $Z^{[i]}(s)$ denotes the expected payout given the option is exercised at time t_i and state $S^{[i]} = s$. It is discounted to the present time $t_0 = 0$. For example, $Z^{[i]}(s) = e^{-r\,t_i}\max(K - s_1, 0)$ for the classical single-stock put option with strike price K and risk-free rate r. In general, the exact payout may depend on the path of $S(t)$ after t_i; hence, the expected payout.

- $V^{[i]}(s)$ and $C^{[i]}(s)$ denote the discounted option values at time t_i and state $S^{[i]} = s$ given that the option was not exercised up to (and at) t_{i-1} and t_i, respectively. Prior studies commonly refer to $C^{[i]}(s)$ as the *continuation value*.

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

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

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

This is equivalent to a dynamic programming problem using the continuation value. Since $C^{[i]}(s)$ and $V^{[i+1]}(s)$ are related by

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

we calculate the option value at t_i by backward induction,

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

This effectively means that the option continues at t_i if $C^{[i]}(s) \ge Z^{[i]}(s)$ and is 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, we express the optimal stopping time τ 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 works in the 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 $S^{[0]}$. We denote the n-th simulation value of $S^{[i]}$ by $S_n^{[i]}$.

- $X^{[i]}(s) = (1, f_1(s), \cdots, f_{M-1}(s))$ denotes the set of M basis functions at time t_i and state $S^{[i]} = s$.

- The N-by-M matrix $\boldsymbol{X}^{[i]}$ is the simulation result of $X^{[i]}(s)$. The n-th row of $\boldsymbol{X}^{[i]}$, denoted by $\boldsymbol{x}_n^{[i]}$, corresponds to $X^{[i]}(S_n^{[i]})$. We assume the basis functions are *diverse* enough to ensure that $\boldsymbol{X}^{[i]}$ has full column rank, M.

- The function $\hat{C}^{[i]}(s)$ is an estimation of $C^{[i]}(s)$ obtained from the simulation set, $\boldsymbol{X}^{[i]}$.

- The length-N column vectors, $\boldsymbol{C}^{[i]}$ and $\boldsymbol{Z}^{[i]}$, are the simulation values of $\hat{C}^{[i]}(s)$ and $Z^{[i]}(s)$, respectively. We denote the n-th elements by $C_n^{[i]} = \hat{C}^{[i]}(S_n^{[i]})$ and $Z_n^{[i]} = Z^{[i]}(S_n^{[i]})$.

- The vector $\boldsymbol{Y}^{[i]}$ is the length-N column vector consisting of the option payout at the stopping time along the simulated paths, conditional on that the option was not exercised before t_i. We denote the n-th element by $Y_n^{[i]}$ and it is equal to $Z_n^{[\tau]}$ for some $i \le \tau \le I$.

- For other variables to be defined later, we use the subscript n and superscript $[i]$ consistently to denote the value of the n-th path at $t = t_i$.

- We use two types of expectation. In the first, we denote the expectation over the N paths in one simulation set by $\mathbb{E}_n$. In the second, we denote the expectation over repeated simulations by $\mathbb{E}_\omega$.

Following the stopping time formulation (1), we compute $\boldsymbol{Y}^{[i]}$ as a path-wise backward induction step: $Y_n^{[I]} = Z_n^{[I]}$ and

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

where $I[\cdot]$ is the indicator function equal to 1 if the condition is satisfied and 0 otherwise. Many authors adopt this backward induction approach, notably Tilley (1993); Carriere (1996); Longstaff and Schwartz (2001). In the final step of the backward induction, we calculate the option price estimate at $t_0 = 0$ as the average option value over the simulated paths:

$$
\hat{V}^{[0]} = \mathbb{E}_n[Y_n^{[0]}] = \frac{1}{N} \sum_{n=1}^{N} Y_n^{[0]}. \tag{4}
$$

The estimation $\hat{V}^{[0]}$, as opposed to the true value $V^{[0]}$, depends on the estimation $\hat{C}^{[i]}(s)$ and the simulation set.

In an alternative backward induction formulation based on (2),

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

which some authors such as Carriere (1996); Tsitsiklis and Van Roy (2001) adopt. However, we do not consider this approach. Carriere (1996), Longstaff and Schwartz (2001), and Stentoft (2014) report that this alternative approach results in a bias significantly higher than the former approach, (3). See Stentoft (2014) for the detailed comparison of the two approaches.

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. This 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 the estimate $\hat{C}_{\text{LSM}}^{[i]}(s)$ as the ordinary least squares (OLS) regression of the next path-wise option values $\boldsymbol{Y}^{[i+1]}$ on the current state $\boldsymbol{X}^{[i]}$:

$$\hat{C}_{\text{LSM}}^{[i]}(s) = X^{[i]}(s)\,\boldsymbol{\beta}^{[i]},$$

where $\boldsymbol{\beta}^{[i]}$ is a length-M column vector of the regression coefficients. Omitting the exercise time superscripts $[i]$ from $\boldsymbol{X}^{[i]}$ for simple notation, $\boldsymbol{C}_{\text{LSM}}^{[i]}$ and $\boldsymbol{\beta}^{[i]}$ are

$$\boldsymbol{C}_{\text{LSM}}^{[i]} = \boldsymbol{X}\boldsymbol{\beta}^{[i]} = \boldsymbol{H}^{[i]}\boldsymbol{Y}^{[i+1]} \quad \text{where} \quad \boldsymbol{\beta}^{[i]} = (\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top \boldsymbol{Y}^{[i+1]}, \quad \boldsymbol{H}^{[i]} = \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}\boldsymbol{X}^\top,$$

where $\boldsymbol{H}^{[i]}$ is the hat matrix. Note that $\boldsymbol{H}^{[i]}$ depends on the current state, $\boldsymbol{X}^{[i]}$, not on the future information, $\boldsymbol{Y}^{[i+1]}$. Using (3), we inductively run the regression for $i = I-1, \cdots, 1$ until we obtain the option price $\hat{V}_{\text{LSM}}^{[0]}$.

To identify how look-ahead bias arises in the LSM algorithm, we focus on the exercise decision at time t_i and state s. For this purpose, we consider only the simulations of size N that have a path passing through the state $S_n^{[i]} = s$ for a dummy path index n. Taking the expectation of (3) over such simulation sets, the option value from the LSM method (with the choice of basis functions) is

$$\mathbb{E}_\omega[\hat{V}_{\text{LSM}}^{[i]}(s)] = \mathbb{E}_\omega[Y_n^{[i]}] = \mathbb{E}_\omega[\,I[C_{n,\text{LSM}}^{[i]} \geq Z_n^{[i]}] \cdot (Y_n^{[i+1]} - Z_n^{[i]}) \,|\, S_n^{[i]} = s\,] + Z^{[i]}(s).$$

Ideally, the exercise decision, $I[C_{n,\text{LSM}}^{[i]} \geq Z^{[i]}(s)]$, and the continuation premium, $Y_n^{[i+1]} - Z_n^{[i]}$, should be independent because the former cannot take advantage of the future information of the simulation path. In the LSM method, however, $C_{n,\text{LSM}}^{[i]}$ depends on $Y_n^{[i+1]}$ via $\boldsymbol{C}^{[i]} = \boldsymbol{H}^{[i]}\boldsymbol{Y}^{[i+1]}$. Therefore, the source of

look-ahead bias is the covariance between the two terms:

$$B^{[i]}(s) = \mathrm{Cov}_\omega \left(I[C_{n,\mathrm{LSM}}^{[i]} \geq Z_n^{[i]}],\, Y_n^{[i+1]} - Z_n^{[i]} \,|\, S_n^{[i]} = s \right). \tag{6}$$

Look-ahead bias is positive because $C_{n,\mathrm{LSM}}^{[i]}$ is always biased toward $Y_n^{[i+1]}$.

We can remove look-ahead bias by de-correlating $\boldsymbol{C}^{[i]}$ from $\boldsymbol{Y}^{[i+1]}$. One method is the standard technique of running an independent simulation set to estimate $\hat{C}^{[i]}(s)$. Applying this type of method, say LSM$'$, to remove look-ahead bias, the option value from the method is suboptimal:

$$\begin{aligned}
\mathbb{E}_\omega[\hat{V}_{\mathrm{LSM}'}^{[i]}(s)] &= \mathbb{E}_\omega[\, I[\, C_{n,\mathrm{LSM}'}^{[i]} \geq Z_n^{[i]}\,]\,|\, S_n^{[i]} = s\,] \cdot \mathbb{E}_\omega[\, Y_n^{[i+1]} - Z_n^{[i]} \,|\, S_n^{[i]} = s\,] + Z^{[i]}(s) \\
&= p_{\mathrm{LSM}'}^{[i]}(s)\, \mathbb{E}_\omega[\, Y^{[i+1]} \,|\, S_n^{[i]} = s\,] + (1 - p_{\mathrm{LSM}'}^{[i]}(s))Z^{[i]}(s) \\
&\leq p_{\mathrm{LSM}'}^{[i]}(s)\, C^{[i]}(s) + (1 - p_{\mathrm{LSM}'}^{[i]}(s))Z^{[i]}(s) \\
&\leq \max(C^{[i]}(s), Z^{[i]}(s)) = V^{[i]}(s).
\end{aligned}$$

Here, $p_{\mathrm{LSM}'}^{[i]}(s) = \mathbb{E}_\omega[\, I[\, C_{n,\mathrm{LSM}'}^{[i]} \geq Z_n^{[i]}\,]\,|\, S_n^{[i]} = s\,]$ is the exercise probability at state s averaged over repeated simulations.

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:

$$B_{\mathrm{Fries}}^{[i]}(s) = \mathrm{Cov}_\omega \left(I[C_{n,\mathrm{LSM}}^{[i]} \geq Z_n^{[i]}],\, C_{n,\mathrm{LSM}}^{[i]} - Z_n^{[i]} \,|\, S_n^{[i]} = s \right).$$

We argue that this definition is inconsistent because it is based on the alternative backward induction (5), even though Fries (2005, 2008) claim to deal with the look-ahead bias in the LSM method.

2.2. The LOOLSM Algorithm

We remove look-ahead bias in the LSM method simply by omitting each simulation path from the regression and making the exercise decision on the path from the self-excluded regression. The bias formulation (6) is free from the correlation because we exclude $Y_n^{[i+1]}$ from the estimation of $C_n^{[i]}$. Figure 1 illustrates this idea with a toy example with three simulation paths.

This idea is well known as LOOCV in statistical learning. This is a special type of the k-fold cross-validation method, where k is equal to the number of data points n. We can obtain the adjusted prediction values analytically without running regressions N times. We express the prediction error with the leave-one-out regression as a correction to that with the full regression (Hastie et al., 2009, § 7.10):

$$\boldsymbol{Y}^{[i+1]} - \boldsymbol{C}_{\mathrm{LOO}}^{[i]} = \frac{\boldsymbol{Y}^{[i+1]} - \boldsymbol{C}_{\mathrm{LSM}}^{[i]}}{\mathbf{1}_N - \boldsymbol{h}^{[i]}},$$

where $\mathbf{1}_N$ is the size-N column vector of 1s, $\boldsymbol{h}^{[i]} = (h_n^{[i]})$ is the diagonal vector of $\boldsymbol{H}^{[i]}$, and the arithmetic

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 $Y_n^{[i+1]} - Z_n^{[i]}$. There are three simulated paths: $n = 1, 2$, and 3. The full regression ($\hat{y} = 1 + x$) under the LSM method indicates that path 2 should be *continued* ($\hat{y}_2 = 1 > 0$). However, this is look-ahead bias influenced by an outlier, path 2 ($y_2 = 4$). Based on the regression without path 2 ($\hat{y}' = -\frac{2}{3} + \frac{5}{6}x$) under the LOOLSM method, it should be *exercised* ($\hat{y}'_2 = -\frac{2}{3} < 0$). The leverage scores, $h_n^{[i]}$, are 13/14, 5/13, and 10/14.
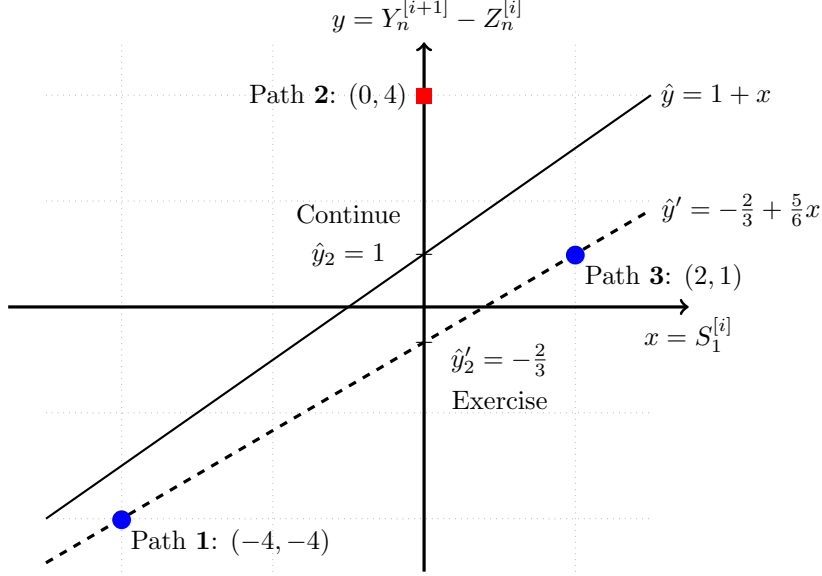
operations between vectors are element-wise. The diagonal element $h_n^{[i]}$, measures the *leverage* of the prediction $C_n^{[i]}$ on the observation $Y_n^{[i+1]}$; that is, $h_n^{[i]} = \partial C_n^{[i]}/\partial Y_n^{[i+1]}$. The value is high when the observation point $\boldsymbol{x}_n^{[i]}$ is far enough away from the others that the regression is more likely fitted close to the observation (see the leverage values in the caption to Figure 1). It also satisfies[2]

$$\frac{1}{n} \leq h_n^{[i]} < 1 \quad \text{and} \quad \sum_{n=1}^{N} h_n^{[i]} = \text{rank}(\boldsymbol{X}^{[i]}) = M. \tag{7}$$

Note that the leave-one-out error is larger in magnitude than the original error because the full regression contains overfitting due to self-influence.

The LOOLSM method we propose is simply to use the corrected continuation value $C_{n,\text{LOO}}^{[i]}$ from the LOOCV in the backward induction step (3):

$$\boldsymbol{C}_{\text{LOO}}^{[i]} = \boldsymbol{C}_{\text{LSM}}^{[i]} - \frac{\boldsymbol{h}^{[i]} \cdot \boldsymbol{e}^{[i]}}{\boldsymbol{1}_N - \boldsymbol{h}^{[i]}} \quad \text{for} \quad \boldsymbol{e}^{[i]} = \boldsymbol{Y}^{[i+1]} - \boldsymbol{C}_{\text{LSM}}^{[i]}. \tag{8}$$

We can compute the whole vector $\boldsymbol{h}^{[i]}$ as the row sum of the element-wise multiplication between $\boldsymbol{X}$ and $\boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}$, which is straightforward from $h_n^{[i]} = \boldsymbol{x}_n(\boldsymbol{X}^\top \boldsymbol{X})^{-1} \boldsymbol{x}_n^\top$. This is much more efficient than

---

[2]The lower bound $1/n$ occurs due to the intercept column in $\boldsymbol{X}^{[i]}$. We exclude the case where $h_n^{[i]} = 1$ because it happens only when $\boldsymbol{X}^{[i]}$ with the n-th observation removed is not full-rank. For the proof and equality condition, see Mohammadi (2016).

obtaining $\boldsymbol{h}^{[i]}$ from the full $\boldsymbol{H}^{[i]}$ matrix. As we must compute the transpose of $\boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X})^{-1}$ for the full regression, we can obtain $\boldsymbol{h}^{[i]}$ with only $O(NM)$ additional operations.

3. Convergence Rate of Look-ahead Bias

3.1. Measuring Look-ahead Bias

In this section, we analyze the convergence rate of look-ahead bias via the LOOLSM method. Given that the LOOCV correction removes the self-influence in the continuation price estimation, it is natural to define look-ahead bias as the difference between the LSM and LOOLSM prices:

$$B_n^{[i]} = Y_n^{[i]} - Y_{n,\text{LOO}}^{[i]} \quad \text{and} \quad \hat{B} = \mathbb{E}_n[B_n^{[0]}] = \hat{V}_{\text{LSM}}^{[0]} - \hat{V}_{\text{LOO}}^{[0]}, \tag{9}$$

where $B_n^{[i]}$ is the path-wise bias and $\hat{B}$ is the final bias in the option value at $t = 0$. To keep the notation simple, we use $Y_n^{[i]}$ instead of $Y_{n,\text{LSM}}^{[i]}$. Measuring look-ahead bias with the LOOLSM method has two advantages over using the two-pass LSM method. First, we eliminate Monte-Carlo error significantly because no extra randomness is required (see Table 2). Moreover, we can analyze the convergence rate of look-ahead bias mathematically thanks to the analytic expression in (8).

There is a subtle difficulty in analyzing $B_n^{[i]}$. The difference between $Y_n^{[i]}$ and $Y_{n,\text{LOO}}^{[i]}$ is easy to measure only when the LSM and LOOLSM methods use the same observations, $\boldsymbol{Y}^{[i+1]} = \boldsymbol{Y}_{\text{LOO}}^{[i+1]}$, for regression. However, this is only guaranteed at the first induction step $i = I - 1$. From the next step, $\boldsymbol{Y}^{[i+1]}$ and $\boldsymbol{Y}_{\text{LOO}}^{[i+1]}$ start deviating. To resolve this difficulty, we introduce a modified LOOLSM method that substitutes $\boldsymbol{Y}^{[i]}$ for $\boldsymbol{Y}_{\text{LOO}}^{[i]}$ in the LOOLSM regression.

$$
\begin{aligned}
Y_n^{[i]} &= I[C_{n,\text{LSM}}^{[i]} \geq Z_n^{[i]}] \cdot (Y_n^{[i+1]} - Z_n^{[i]}) + Z_n^{[i]}, \\
Y_{n,\text{LOO}}^{[i]} &= I[C_{n,\text{LOO}}^{[i]} \geq Z_n^{[i]}] \cdot (Y_{n,\text{LOO}}^{[i+1]} - Z_n^{[i]}) + Z_n^{[i]}, \\
C_{n,\text{LOO}}^{[i]} &= C_{n,\text{LSM}}^{[i]} - (Y_n^{[i+1]} - C_{n,\text{LSM}}^{[i]})\, h_n^{[i]}/(1 - h_n^{[i]})
\end{aligned}
\tag{10}
$$

We numerically verify that the impact of this modification is negligible in pricing because the substitution affects only the estimated continuation value. Therefore, we assume that the LOOLSM price in (9) is measured with the modified method throughout this section.

3.2. Main result

Before stating the main result, we first build an intuition for the convergence rate of look-ahead bias. Suppose the n-th path is an outlier in the LSM regression such that the sample point $Y_n^{[i+1]}$ is much bigger than the prediction $C_{n,\text{LSM}}^{[i]}$.[3] Then, look-ahead bias inverts the exercise decision when $C_{n,\text{LOO}}^{[i]} <$

---

[3] By symmetry, one can also assume that $Y_n^{[i+1]}$ is much smaller than $C_{n,\text{LSM}}^{[i]}$.

$Z_n^{[i]} \le C_{n,\text{LSM}}^{[i]}$. In Lemma 2 in Appendix A, we will show that this is equivalent to

$$0 \le C_{n,\text{LSM}}^{[i]} - Z_n^{[i]} < h_n^{[i]}(Y_n^{[i+1]} - Z_n^{[i]}). \tag{11}$$

We can guess that the probability for above events would decay at the rate of M/N. This is because, as the simulation size N grows larger, $h_n^{[i]}$ becomes smaller as $\mathbb{E}_n[h_n^{[i]}] = M/N$ from (7) while the size of the other terms remain the same. Indeed, this turns out to be the case, as we discuss below in detail.

The main results rely on two technical assumptions common in studies analyzing the convergence of the LSM algorithm (Clément et al., 2002; Stentoft, 2004). First, we work only with realistic payoff functions that grow moderately and well-defined option prices. This is a minimal condition from a practical standpoint.

Assumption 1. *The payout functions, $Z^{[i]}(s)$, are in L^2.*

The second assumption deals with the complication in pricing that arises when the option and continuation values are arbitrarily close with a non-negligible probability. This outcome might lead to a wrong exercise decision at the limit and the LSM algorithm fails to converge to the true price; see Stentoft (2004) for example. Henceforth, we assume that the continuation value is different from the exercise value almost surely.

Assumption 2. *Fix an ordered set of countable basis functions $\{f_m(s) : m = 0, 1, \cdots\}$ in the L^2 space. Let $C_M^{[i]}(s)$ be the continuation value obtained from the LSM method with the first M basis functions at the limit as $N \to \infty$, such that $C_M^{[i]}(s) \to C^{[i]}(s)$ as $M \to \infty$. Furthermore, let*

$$P_M^{[i]}(c) = \mathbb{P}[\,|C_M^{[i]}(S^{[i]}) - Z^{[i]}(S^{[i]})| \le c\,]$$

be the probability of the absolute continuation premium not exceeding c. Then, we assume that

$$\lim_{c \to 0} P_M^{[i]}(c) = 0 \quad \text{for all} \quad M = 1, 2, \cdots.$$

The following theorem is the main result of this section. In analyzing the convergence rate, we regard any derived quantity (e.g., $\hat{B}$) as a random variable, and examine how its expected value behaves as N increases.

Theorem 1. *The following hold under Assumptions 1 and 2.*

(i) $B_n^{[i]} \sim O_p(M/N)$.

(ii) *For any given $\varepsilon > 0$, there exists $r_\varepsilon > 0$ such that the expected look-ahead bias satisfies $\mathbb{E}_\omega[\hat{B}] \le \varepsilon + r_\varepsilon M/N$.*

(iii) *$\hat{B}$ converges to zero in probability.*

Here, the probabilistic asymptotic notation O_p is defined in the probability space of all possible simulation runs of size N. The subscript for path n in (i) is a dummy index because the Monte Carlo paths are drawn independently.

We prove Theorem 1 in Appendix A. Two cases are treated separately in the proof: (a) the contributions to look-ahead bias near the exercise boundary and when the tails of the asset distributions can be

12

made arbitrarily small, and (b) the probability of look-ahead bias occurring elsewhere can be bounded by a constant multiple of leverage $h_n^{[i]}$ with expected value M/N. Since any realistic bias is controlled by (b), its expected value decays at the rate of a constant multiple of M/N at least. Indeed, we report a strong linear relationship between look-ahead bias and M/N with a few examples, see Figures 2, 3, and 4. Although Theorem 1 does not guarantee linearity, we believe this is a direct consequence of $h_n^{[i]}$ primarily determining the convergence rate.

4. Numerical Results

4.1. Overview of Experiments

We price four Bermudan option cases to compare the LSM and LOOLSM methods. We present them in increasing order of the number of underlying assets: single-stock put options, best-of options on two assets, basket options on four assets, and cancellable exotic interest swap under the LIBOR market model. Therefore, the number of regressors, M, also increases in general given the same polynomial orders to include.

We run n_{MC} sets of simulations with N paths each and use the following three estimators for comparison:

- **LSM**: the one-pass (i.e., in-sample) LSM estimator.

- **LSM-2**: the two-pass (i.e., out-of-sample) LSM estimator. We apply the exercise policy computed from an extra set of N paths to the payoff valuation with the original simulation set.

- **LOOLSM**: the LOOLSM estimator.

Using the same N simulation paths for the payoff valuation across the three methods works as a control to reduce the variability of the measured bias (i.e., price difference between methods). We use the antithetic random variate ($N/2 + N/2$) to reduce the variance. We also vary M by selecting different basis sets. From the results of n_{MC} independent simulation sets, we obtain the mean and standard deviation of the option price. If the exact option value $V^{[0]}$ is available, then we report the price offset from $V^{[0]}$:

$$\text{Price Offset} = \mathbb{E}_\omega[\hat{V}^{[0]}] - V^{[0]},$$

where $\hat{V}^{[0]}$ is the price estimate from each simulation. Otherwise, we report the price $\mathbb{E}_\omega[\hat{V}^{[0]}]$. We implemented the numerical experiments in Python (Ver. 3.7, 64-bit) on a personal computer running Windows 10 with an Intel core i7 1.9 GHz CPU and 16 GB RAM.

4.2. Bermudan Options under the Black-Scholes Model

The underlying asset prices $S_j(t)$ of the three examples in this section follow geometric Brownian motions:

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

where r is the risk-free rate, q_j is the dividend yield, σ_j is the volatility, and the $W_j(t)$'s are the standard Brownian motions correlated by $dW_j(t)\, dW_{j'}(t) = \rho_{jj'}\, dt$ ($\rho_{jj} = 1$). The choice of geometric Brownian motion for the price dynamics has several advantages and does not oversimplify the problem. It is easy to implement because an exact simulation is possible. The geometric Brownian motion is a standard choice in the literature and we can take advantage of the exact Bermudan option prices reported previously.

For each example, we run two experiments. The first experiment is to ensure that the LOOLSM method eliminates look-ahead bias in a similar way as the LSM-2 method. We run $n_{\mathrm{MC}} = 100$ sets of simulations with $N = 4 \times 10^4$ paths and price options with the three methods. Additionally, we price the corresponding European options using the Monte Carlo method with the same sets of paths.

The second experiment is to validate the convergence rate of the bias in Theorem 1. We run the LSM and LOOLSM methods with varying N and M. We first generate a pool of 720×10^4 Monte Carlo paths and split them into groups of $N = 0.5, 1.5, 3, 6, 9, 12, 18,$ and 36×10^4 paths. Therefore, the generated groups comprise $n_{\mathrm{MC}} = 720 \times 10^4 / N\ (= 1440, \cdots, 20)$ Monte Carlo runs, and we compute the price offset and standard deviation from the n_{MC} prices. By varying N within 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 measure look-ahead bias. Simultaneously, we vary the number of regressors (M) by including polynomials of higher terms. We thus measure look-ahead bias as a function of M/N.

The three examples are the Bermudan options whose payouts are determined entirely at the time of exercise and are always non-negative. We take advantage of these properties in the implementation of the methods. First, we include the payout function as a regressor, $f_1(s) = Z^{[i]}(s)$. The payout is an important regressor improving the optimality of the exercise decision, as Glasserman (2003) shows. Including the payout also eliminates a specification issue in regression. In an alternative LSM implementation, one may regress the continuation premium, $\boldsymbol{Y}^{[i+1]} - \boldsymbol{Z}^{[i]}$ (instead of $\boldsymbol{Y}^{[i+1]}$), to estimate $\boldsymbol{C}^{[i]} - \boldsymbol{Z}^{[i]}$ (instead of $\boldsymbol{C}^{[i]}$). While it is difficult to determine the superior approach, they become identical when we include $Z^{[i]}(s)$ as a basis function. Second, following Beveridge et al. (2013), we do not exercise the option when $Z_n^{[i]} = 0$, even if $C_n^{[i]} < Z_n^{[i]}$. The negative continuation value is an artifact caused by simulation noise or imperfect basis functions. It is always optimal to continue the option since the future payout is non-negative.[4]

---

[4]On top of the exercise decision override, Longstaff and Schwartz (2001) even suggest running the regressions with the in-the-money paths only; that is, $\{n :\ Z_n^{[i]} > 0\}$. In our experiment, however, this practice makes little difference. Glasserman (2003) even reports that the result can be inferior in some cases. We use all simulation paths in this study.

Case 1: Single-stock Put Option. We start with Bermudan put options on a single stock

$$Z^{[i]}(s) = e^{-rt_i} \max(K - s_1, 0),$$

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

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

We obtain the exact option prices for the strike prices: $K = 80$, 90, 100, 110, and 120 by implementing the binomial tree method. Feng and Lin (2013) reports an exact price for $K = 100$ consistent with our result. For the regressors, we use

$$X^{[i]}(s) = (1, \ Z^{[i]}(s_1), \ s_1, \ s_1^2, \cdots).$$

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 and LSM-2 prices are similar and are slightly lower than the LSM price. This result implies that LOOLSM removes look-ahead bias, although the size is small. To show the statistical significance of the look-ahead bias, Table 2 reports its mean and standard deviation separately. The bias measured with the LOOLSM method has much less deviation than that measured with LSM-2 because the LOOLSM method requires no extra simulation, whereas LSM-2 needs another independent simulation set.

Figure 2 shows the result of the second experiment. The top plot shows the price offset of the LSM and LOOLSM methods as a function of M/N for varying M and N values. It demonstrates how the LSM and LOOLSM prices converge as N increases for a fixed M. The LSM price converges from above and the LOOLSM converges from below, indicating that the LOOLSM price is low-biased compared with the convergent value for a given M. The bottom plot shows 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 convergence rate in Theorem 1. While the figure illustrates one specific option ($K = 80$), the other options in the case exhibit the same pattern.

This single-asset case is similar to the example Longstaff and Schwartz (2001) uses to demonstrate that look-ahead bias is negligible. Indeed, our finding is consistent. In light of the convergence rate analysis, however, the example has the smallest ratio, $M/N = 4/10^5$, among the five examples in Longstaff and Schwartz (2001). The convergence analysis indicates a danger that the LSM price might rise above the true price when using larger basis sets (e.g., $M = 8$ and 12), even with a large simulation size N.

Case 2: Best-of Option on Two Assets. We price best-of (or rainbow) call options on two assets:

$$Z^{[i]}(s) = e^{-rt_i} \max(\max(s_1, s_2) - K, 0),$$

Table 1: Results for the single-stock Bermudan put options (Case 1). We use $N = 40,000$, $M = 5$. The "Exact" columns report the true option prices, while the other columns report the price offset and standard deviation from the $n_{\mathrm{MC}} = 100$ 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 the single-stock Bermudan put options in § 4.2. The columns report the difference between the price of each method and the LSM price, and its error estimate.

K	LSM − LSM-2	LSM − 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

Figure 2: The price offset (top) and look-ahead bias (bottom) as functions of M/N for the single-stock put option with $K = 80$ (Case 1). At the top, given the fixed values of M and N, the higher value corresponds to the LSM method and the lower one to the LOOLSM method.

16
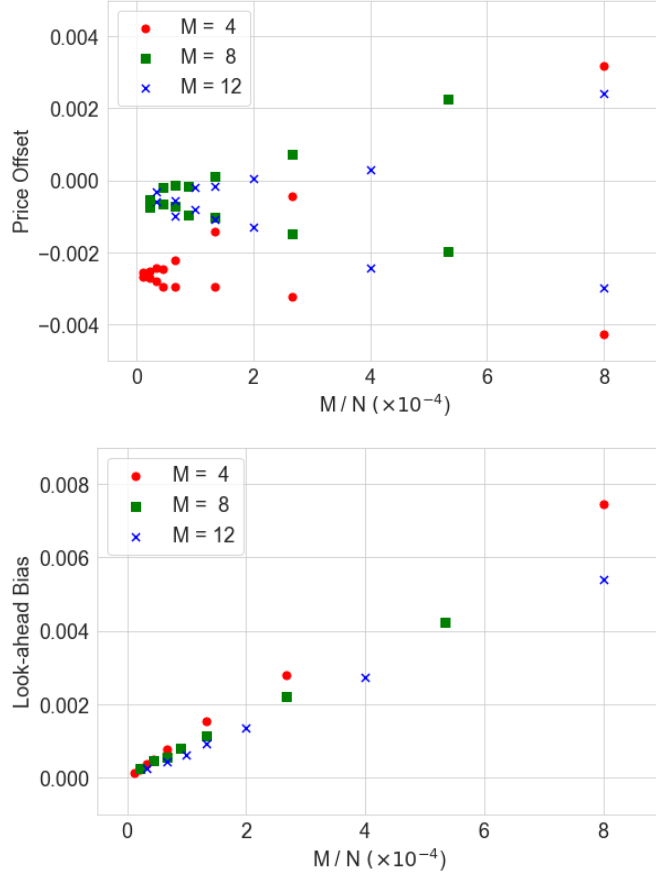
with the parameter set tested by Glasserman (2003) and Andersen and Broadie (2004):

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

The options are priced with three initial asset prices, $S_1(0) = S_2(0) = 90$, 100, and 110. We use the following basis functions (M=11) for the first experiment:

$$X^{[i]}(s) = (1, \ Z^{[i]}(s), \ s_1, \ s_2, \ s_1^2, \ s_1 s_2, \ s_2^2, \ s_1^3, \ s_1^2 s_2, \ s_1 s_2^2, \ s_2^3).$$

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

Table 3 and Figure 3 show the results. Look-ahead bias in the LSM method becomes more pronounced, whereas the LSM price is still lower than the true price, primarily because the exercise boundary of the best-of option is highly non-linear, as Glasserman (2003) observes. As depicted in Figure 3, suboptimal bias quickly decreases as M increases. Nevertheless, look-ahead bias is clearly proportional to M/N, regardless of the suboptimality level.

Table 3: Results for the best-of Bermudan call options (Case 2). We use $N = 40,000$ and $M = 11$. The "Exact" columns report the 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.

$S_j(0)$	Bermudan				European	
	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

**Case 3: Basket Option on Four Assets**. Next, we price Bermudan call options on a basket of four stocks:

$$Z^{[i]}(s) = e^{-rt_i} \max\left(\frac{s_1 + s_2 + s_3 + s_4}{4} - K, 0\right).$$

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

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

The options are priced for a range of strikes, $K = 60, 80, 100, 120$, and 140. Because the underlying assets do not pay dividends, it is optimal not to exercise the option until maturity; hence, the European option price is equal to the Bermudan price. Therefore, we refer to Choi (2018) for the exact prices. For

Figure 3:   The price offset (top) and look-ahead bias (bottom) as functions of M/N for the best-of call option with $S_j(0) = 100$ (Case 2). At the top, given the fixed values of M and N, the higher value corresponds to the LSM method and the lower one to the LOOLSM method.
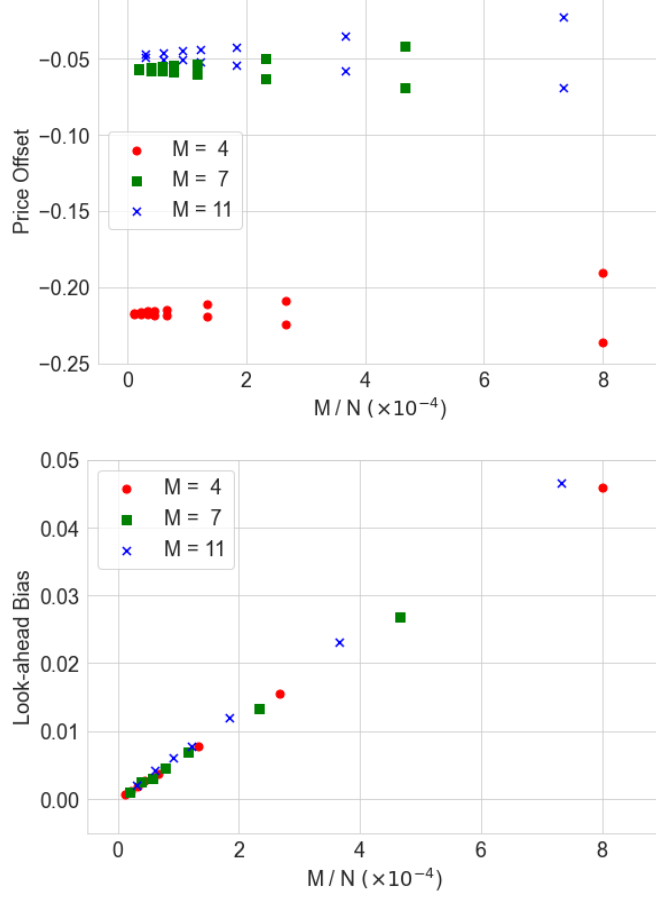
the regressors, we use polynomials up to degree 2 ($M = 16$) for the first experiment:

$$X^{[i]}(s) = (1,\, Z^{[i]}(s),\, s_j, \cdots,\, s_j^2, \cdots,\, s_j s_{j'}, \cdots) \quad \text{for} \quad 1 \leq j < j' \leq 4,$$

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

Table 4 and Figure 4 report the results. In this example, the LSM method noticeably overprices the option for all strike prices, whereas the LOOLSM and LSM-2 prices are consistently low-biased. 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, therefore only the linear basis functions ($M = 6$) capture the exercise boundary accurately. The look-ahead biases for the different M's collapse into a function of M/N, although linear convergence clearly appears when M/N is very small (see the inset of the bottom plot of Figure 4).

Table 4: Results for the four-asset basket options (Case 3). We use $N = 40,000$ and $M = 16$. The "Exact" columns report the true option prices, while the other columns report the price offset and standard deviation from $n_{\mathrm{MC}} = 100$ 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

Figure 4: The price offset (top) and look-ahead bias (bottom) as functions of M/N for the four-asset basket option with $K = 100$ (Case 3). At the top, given the fixed values of M and N, the higher value corresponds to the LSM method and the lower one to the LOOLSM method.
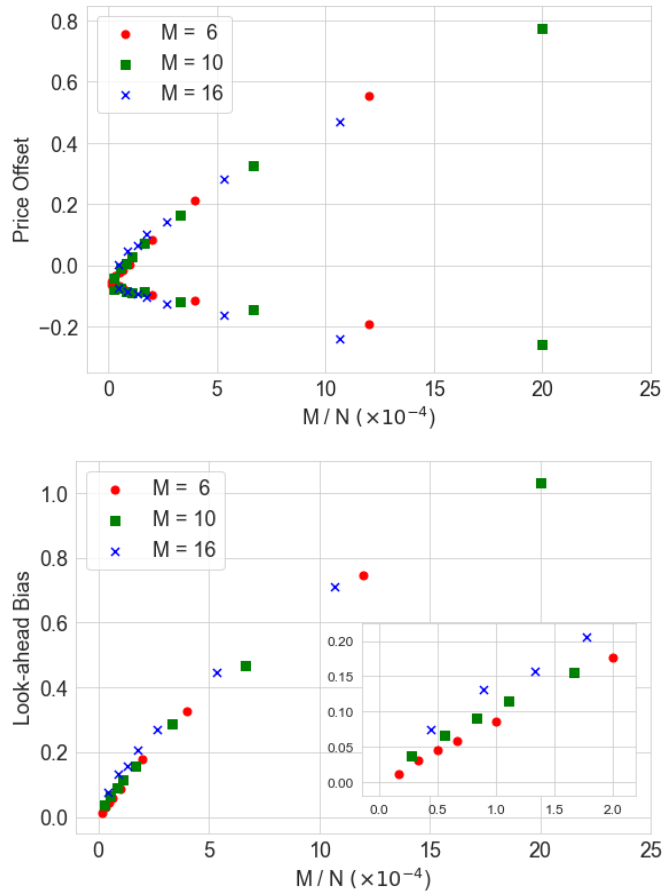
4.3. Cancellable Exotic Interest Rate Swap under the LIBOR Market Model

Finally, we apply the LOOLSM method to a cancellable exotic interest rate derivative under the LIBOR market model (Brace et al., 1997; Jamshidian, 1997). This last example is different from the previous three examples in that it is closer to the structured product traded in the market and has a much higher computational cost than the previous examples. The exact option price is not available. We use this example to demonstrate the computational advantage of the LOOLSM algorithm.

We briefly introduce the LIBOR market model first before introducing the payout. Let $P(t, T)$ denote

the time t price of the zero-coupon bond paying \$1 at T. For a set of equally spaced dates, $T_j = j\delta$ for a tenor δ, the forward rate between T_j and T_{j+1} seen at time $t \leq T_j$ is given by

$$F_j(t) = \frac{1}{\delta}\left(\frac{P(t, T_j)}{P(t, T_{j+1})} - 1\right).$$

We also denote the *spot* rate by $F_j = F_j(T_j)$. The LIBOR market model evolves $\{F_j(t)\}$, and then yields the discount curve $P(t, \cdot)$. Among the various model specifications, we follow the displaced-diffusion stochastic volatility implementation of Joshi and Rebonato (2003); the forward rates follow displaced geometric Brownian motions

$$\frac{d\,F_j(t)}{F_j(t) + \alpha} = \mu_j(t)\,dt + \sigma_j(t)\,dW_j(t) \quad \text{for} \quad 0 \leq t \leq T_j,$$

where the $W_j(t)$s are the correlated standard Brownian motions, and the volatility takes the time-homogeneous form

$$\sigma_j(t) = \big(a + b(T_j - t)\big)e^{-c(T_j - t)} + d \quad \text{for} \quad 0 \leq t \leq T_j.$$

This *abcd* volatility structure is popular in the literature (Joshi and Tang, 2014; Beveridge et al., 2013). Joshi and Rebonato (2003) further makes volatility stochastic by letting a, b, $\log c$, and $\log d$ evolve over time following the Ornstein–Uhlenbeck process

$$d\,h(t) = \lambda(\,h_\infty - h(t)) + \sigma_h dW_h(t), \tag{12}$$

where $W_h(t)$ is a standard Brownian motion independent of the $W_j(t)$s. Displaced diffusion and stochastic volatility enable the model to exhibit the swaption volatility smile observed in the market. We choose a spot measure where the numeraire asset is a discretely compounded money market account with \$1 invested at $t = 0$. The value of the numeraire asset at $t = T_j$ is

$$P_j^* = \prod_{k=0}^{j-1}(1 + \delta F_j)$$

The drift, $\mu_j(t)$, is determined by the arbitrage condition depending the numeraire choice. The predictor–corrector method (Hunter et al., 2001) provides an efficient approximation of the integrated drift required for the simulation of $F_j(t)$.

Case 4: Cancellable CMS Spread and LIBOR Range Accrual. Using the LIBOR market model, we price a callable structure note with an exotic coupon rate. In the equivalent swap form, the note issuer (option buyer) pays an exotic coupon with annual rate R_j and the investor (option seller) pays the market rate F_{j-1} at the end of each period, $t = T_j$. The exotic coupon is paid only when the spread of

two constant maturity swaps (CMS) and LIBOR rates are within certain ranges. Specifically, we assume

$$R_j = \begin{cases} 0.095 & \text{in the first one year} \quad (j = 1, \ldots, 1/\delta) \\ 0.095 \cdot I[\, \mathrm{Sw}_j^{2\mathrm{y}} \leq \mathrm{Sw}_j^{10\mathrm{y}} \,] \cdot I[\, 0 \leq F_j \leq 0.03 \,] & \text{afterwards,} \end{cases}$$

where $\mathrm{Sw}_j^{n\mathrm{y}}$ is the n-year swap rates implied from the forward rates at the time of cashflow exchange, $\{F_k(T_j) : k \geq j\}$. The two conditions embedded in R_j have been popular among investors since the financial crisis in 2008; the inversion of the swap curve (i.e., $\mathrm{Sw}_j^{2\mathrm{y}} > \mathrm{Sw}_j^{10\mathrm{y}}$) is historically rare and it is expected that the Federal Reserve will maintain a low realized short-term rate (i.e., F_j). However, the risk-neutral probability of the conditions implied from the option market is lower. Therefore, the coupon rate (i.e., 9.5%) can be set high to balance the present values of the two parties. While most market trades use daily range accrual to mitigate the fixing risk, our example uses a single observation per coupon period to simplify pricing.

We assume that the forward rate tenor is 6 months ($\delta = 0.5$) and that the swap matures in 20 years. Therefore, we simulate 60 forward rates, $\{F_j(t) : 1 \leq j \leq 60\}$, until $t = 20$. The issuer can cancel the swap every year at $t_i = T_{2i} = i$ for $i = 1, \ldots, 20$ ($I = 20$). Following the market convention, cancellation does not apply to the cashflow exchange at the same time. To price the trade in a Bermudan option form, we decompose the cancellable swap into two trades: the (non-cancellable) underlying swap and the Bermudan swaption, where the holder has the right to enter the swap from t_i to $t = 20$. The payout of the Bermudan swaption to the option holder is

$$Z^{[i]} = \sum_{j=2i+1}^{2I} \frac{\delta}{P_j^*}(R_j - F_{j-1}) \quad (Z^{[I]} = 0).$$

We obtain the price of the cancellable swap as the sum of the prices for the two trades accordingly.

Table 5: The Ornstein–Uhlenbeck parameters of the stochastic *abcd* volatility for cancellable exotic interest rate swap (Case 4). The parameter values are from Joshi and Rebonato (2003, Table 4)

$h(t)$	$h(0) (= h_\infty)$	σ_h	λ
a	-0.020	0.05	0.5
b	0.108	0.1	0.3
$\log c$	$\log(0.800)$	0.1	0.5
$\log d$	$\log(0.114)$	0.2	0.4268

We use the model parameters in Joshi and Rebonato (2003, § 8.4) for the simulation. The displacement is $\alpha = 0.025$ and the correlation between forward rates decays exponentially, $dW_j(t)dW_{j'}(t) = e^{-\theta\delta|j-j'|} dt$ with $\theta = 0.1$. The Ornstein–Uhlenbeck parameters for the *abcd* volatility are given in Table 5. See Joshi and Rebonato (2003, § 8) for the swaption volatilities implied from the parameter set. We simulate the stochastic volatility with the Euler scheme with time step, $\Delta t = 0.25$. The initial forward rates are $F_j(0) = 0.045 - 0.0425\, e^{-j\delta/4}$ such that F_j increases from 0.25% to 4.5% as j increases.

Because the payout $Z^{[i]}$ is not determined at the time of exercise, we require a different regression implementation from the previous cases. We do not include $Z^{[i]}$ in the basis function. We apply the regression to the continuation premium (or penalty), $\boldsymbol{Y}^{[i+1]} - \boldsymbol{Z}^{[i]}$ instead.[5] For the basis set $X^{[i]}$, we consider two groups of variables: (i) the 6 variables related to interest rates, F_{2i}, $\mathrm{Sw}_{2i}^{2\mathrm{y}}$, $\mathrm{Sw}_{2i}^{10\mathrm{y}}$, $\mathrm{Sw}_{2i}^{(20-i)\mathrm{y}}$ (co-terminal swap rate), $I[\,\mathrm{Sw}_{2i}^{2\mathrm{y}} \leq \mathrm{Sw}_{2i}^{10\mathrm{y}}\,]$, and $I[\,0 \leq F_{2i} \leq 0.03\,]$); and (ii) the 4 volatility parameters, a, b, c, and d at t_i. With these terms, we construct the 4 basis sets in increasing order of M:[6]

- $M = 11$: linear terms of both groups.

- $M = 28$: linear and quadratic terms of the interest rate group.

- $M = 32$: linear and quadratic terms of the interest rate group and linear volatility terms.

- $M = 66$: linear and quadratic terms of all variables in both groups.

In Table 6, we present the pricing result. Similar to the previous examples, the LOOLSM prices are closer to the LSM-2 prices, indicating that the LOOLSM method removes look-ahead bias efficiently. The look-ahead bias, measured as the difference between the LSM and LOOLSM prices, proportionally increases as the number of basis M increases. In particular, the LOOLSM (and LSM-2) price no longer increases after $M = 28$, whereas the LSM prices keeps increasing due to the look-ahead bias. Therefore, it is safe to use higher-order basis functions under the LOOLSM algorithm.

Table 6: Results for the cancellable exotic swap (Case 4). The prices are for the notional value of 100. The prices are the average of the $n_{\mathrm{MC}} = 200$ simulation results with $N = 10,000$ paths. The standard deviations are constant around 0.22. The price of the underlying swap is -4.19 ± 0.20. All values are rounded to two decimal places.

	LSM	LSM-2	LOOLSM	LSM $-$ LOOLSM
$M = 11$	9.09	9.04	9.06	0.04 ± 0.01
$M = 28$	9.45	9.38	9.38	0.07 ± 0.02
$M = 32$	9.48	9.39	9.39	0.09 ± 0.02
$M = 66$	9.53	9.40	9.36	0.18 ± 0.03

Table 7: Average computation time in seconds for pricing the cancellable exotic swap (Case 4) with $N = 10,000$ paths ($2N$ paths for LSM-2) and $M = 66$ basis functions. For smaller basis sets ($M = 11$, 28, and 32), the "Regression and Pricing" time is smaller in proportion to M, while the "Path Generation" time is unchanged.

	LSM	LSM-2	LOOLSM
Path Generation	65.19	130.38	65.19
Regression and Pricing	0.35	0.64	0.42
Total	65.54	131.02	65.60

In Table 7, we compare the computation time of the three methods. We separately measure the time to generate paths and perform the regression and valuation. Note that path generation takes the

---

[5]In general, we can separately regress $\boldsymbol{Y}^{[i+1]}$ and $\boldsymbol{Z}^{[i]}$ (Piterbarg, 2003).

[6]All basis sets include 1 for the intercept.

majority of the total pricing time due to the complexity of the stochastic LIBOR market implementation. Since LSM-2 requires another set of simulation paths, it takes twice as much as time to generate paths. Regarding the time for regression and valuation, the increment from LSM to LOOLSM is marginal, as in § 2.2, but LSM-2 takes longer than LOOLSM because we must evaluate the regression on the two simulation sets. Overall, the computational gain of the LOOLSM method is significant compared to the LSM-2 method, while they achieve the same goal of removing look-ahead bias.

5. The extension of LOOLSM to the other regression estimators

Many studies have aimed to improve the LSM method using advanced regression methods, such as ridge regression (Tompaidis and Yang, 2014), least absolute shrinkage and selection operator (LASSO) (Tompaidis and Yang, 2014; Chen et al., 2019), weighted least squares regression (Fabozzi et al., 2017; Ibáñez and Velasco, 2018), and non-parametric kernel regression (Belomestny, 2011; Ludkovski, 2018). The LOOLSM method can be flexibly extended to these alternatives to the LSM method because they are essentially linear projections via the hat matrix. As long as $\boldsymbol{h}^{[i]}$ is available, the corrections from the original regression are obtained from (8) in the same manner and the two-pass approach becomes unnecessary. Below, we present their hat matrices with discussions.

Ridge regression and LASSO are linear regressions with L^2 and L^1-regularization, respectively (Hastie et al., 2009, § 3.4). These methods outperform the LSM method in small simulation paths (Tompaidis and Yang, 2014) and provide stable estimates of the value-at-risk (Chen et al., 2019). The hat matrix of ridge regression is

$$\boldsymbol{H}^{[i]} = \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{X} + \lambda \boldsymbol{I}_M)^{-1} \boldsymbol{X}^\top,$$

where λ is the regularization strength. From its diagonal, we can see the implication of regularization for look-ahead bias. The effective degree of freedom, defined as the sum of the leverages (see Hastie et al. (2009, (3.50))), is less than that of the OLS regression in (7):

$$\sum_{n=1}^{N} h_n^{[i]} = \sum_{j=1}^{M} \frac{d_j^2}{d_j^2 + \lambda} < M,$$

where d_j is the j-th singular value of $\boldsymbol{X}^{[i]}$. Therefore, we expect the look-ahead bias to decrease as λ increases, following Section 3. However, regularization alone cannot remove look-ahead bias completely, and we still need an additional method such as the LOOLSM method. The hat matrix of LASSO is not analytically available, not to mention that the method is not exactly a linear projection. Given the selected regressors from shrinkage, however, it is a linear projection. Therefore, an approximation of the hat matrix can be obtained accordingly. The effective degree of freedom is equal to the number of the selected regressors under the approximation, which also indicates that LASSO has an effect of reducing look-ahead bias to some extent.

In the weighted linear regression, the hat matrix is

$$\boldsymbol{H}^{[i]} = \boldsymbol{X}(\boldsymbol{X}^\top \boldsymbol{W} \boldsymbol{X})^{-1} \boldsymbol{X}^\top \boldsymbol{W},$$

where $\boldsymbol{W}$ is an N-by-N diagonal weight matrix. Fabozzi et al. (2017) adopt this approach to deal with heteroscedasticity and Ibáñez and Velasco (2018) to give higher weights to the paths near the exercise boundary.

Despite heavy computation, the non-parametric kernel regression is an alternative to the OLS regression (Belomestny, 2011; Ludkovski, 2018). For kernel function $\mathcal{K}(x,y)$, the hat matrix is the normalized kernel value between sample points. The (n,n') element of $\boldsymbol{H}^{[i]}$ is

$$H_{n,n'}^{[i]} = \frac{\mathcal{K}(X_n^{[i+1]}, X_{n'}^{[i+1]})}{\sum_{k=1}^{N} K(X_n^{[i+1]}, X_k^{[i+1]})}.$$

The adjusted prediction value from (8) is simply the self-excluded kernel estimate:

$$C_{n,\text{LOO}}^{[i]} = \frac{\sum_{k \neq n} \mathcal{K}(X_n^{[i+1]}, X_k^{[i+1]}) Y_k^{[i+1]}}{\sum_{k \neq n} \mathcal{K}(X_n^{[i]}, X_k^{[i]})}.$$

In the kernel regression, the LOOLSM algorithm not only saves out-of-sample path generation, but also saves costly kernel evaluations.

6. Conclusion

This study shows that it is possible to eliminate undesirable look-ahead bias in the LSM method (Longstaff and Schwartz, 2001) using LOOCV without extra simulations. By measuring look-ahead bias with the LOOLSM method, we also find that the bias size is asymptotically proportional to the regressors-to-simulation paths ratio. With numerical examples, we demonstrate that the LOOLSM method effectively prevents the possible overvaluation of multi-asset American options without extra computation.

Appendix A. Proof of Theorem 1

We first introduce two technical Lemmas 2 and 3, then prove the main Theorem 1 in Section 3. For ease of notation, we omit the exercise time superscripts $[i]$ from $Z_n^{[i]}$, $C_n^{[i]}$, $h_n^{[i]}$, and $B_n^{[i]}$ when the context is clear. However, we preserve the superscript in $Y_n^{[i]}$ and $Y_n^{[i+1]}$ to avoid ambiguity.

Lemma 2. *Let $I[\cdot]$ be the indicator function. Then,*

$$B_n^{[i]} \le I\left[\,|C_{n,\text{LSM}} - Z_n| \le h_n\,|Y_n^{[i+1]} - Z_n|\,\right]\,|Y_n^{[i+1]} - Z_n| + B_n^{[i+1]}.$$

PROOF (PROOF OF LEMMA 2.).
We first formulate the conditions under which look-ahead bias changes the exercise decision. From (10),

$$
\begin{aligned}
B_n^{[i]} &= Y_n^{[i]} - Y_{n,\text{LOO}}^{[i]} \\
&= (I[Z_n \le C_{n,\text{LSM}}] - I[Z_n \le C_{n,\text{LOO}}])(Y_n^{[i+1]} - Z_n) + I[Z_n \le C_{n,\text{LOO}}]\,B_n^{[i+1]} \\
&\le (I[D_n^+] - I[D_n^-])(Y_n^{[i+1]} - Z_n) + B_n^{[i+1]}.
\end{aligned}
$$

where $D_n^+ \iff \{C_{n,\text{LOO}}^{[i]} < Z_n \le C_{n,\text{LSM}}^{[i]}\}$ and $D_n^- \iff \{C_{n,\text{LSM}}^{[i]} < Z_n \le C_{n,\text{LOO}}^{[i]}\}$. Here, D_n^+ (D_n^-) is the condition in which the LSM algorithm incorrectly continues (exercises) due to look-ahead bias, but the LOOLSM algorithm exercises (continues), and the term $(Y_n^{[i+1]} - Z_n)$, is the price change caused by the inverted exercise decision. From (8), we obtain the following equivalence:

$$
\begin{aligned}
D_n^+ &\iff 0 \le C_{n,\text{LSM}} - Z_n < C_{n,\text{LSM}} - C_{n,\text{LOO}} \\
&\iff 0 \le C_{n,\text{LSM}} - Z_n < \frac{h_n}{1 - h_n}(Y_n^{[i+1]} - C_{n,\text{LSM}}) \\
&\iff 0 \le C_{n,\text{LSM}} - Z_n < h_n(Y_n^{[i+1]} - Z_n);
\end{aligned}
$$

similarly,

$$D_n^- \iff h_n(Y_n^{[i+1]} - Z_n) \le C_{n,\text{LSM}} - Z_n < 0.$$

Since D_n^+ and D_n^- are mutually exclusive events, we get

$$
\begin{aligned}
B_n^{[i]} &\le I[D_n^+ \cup D_n^-] \cdot |Y_n^{[i+1]} - Z_n| + B_n^{[i+1]} \\
&\le I\left[\,|C_{n,\text{LSM}} - Z_n| \le h_n|Y_n^{[i+1]} - Z_n|\,\right] \cdot |Y_n^{[i+1]} - Z_n| + B_n^{[i+1]}. \quad \square
\end{aligned}
$$

Lemma 3. *The following hold under Assumptions 1 and 2.*

(i) $Y_n^{[i+1]} - Z_n \sim O_p(1)$ and $\mathbb{E}_\omega\left[(Y_n^{[i+1]} - Z_n)^2\right]$ are finite.

(ii) $h_n \sim O_p(M/N)$.

(iii) $|C_{n,\text{LSM}} - Z_n|^{-1} \sim O_p(1)$.

PROOF (PROOF OF LEMMA 3.). (i) Since $Y_n^{[i+1]} = Z_n^{[\tau]}$ for some $i+1 \le \tau \le I$, we obtain

$$\left(Y_n^{[i+1]} - Z_n^{[i]}\right)^2 \le \left(\sum_{j=i}^{I} |Z_n^{[j]}|\right)^2 \le (I - i + 1)\sum_{j=i}^{I}\left(Z_n^{[j]}\right)^2$$

for any N. Then, the statements follow from Assumption 1 that $Z^{[i]}(s)$ is in L^2.

(ii) For any given $\varepsilon > 0$,

$$\mathbb{P}_\omega\left[h_n > \frac{1}{\varepsilon}\frac{M}{N}\right] \le \frac{\varepsilon N}{M}\,\mathbb{E}_\omega\left[h_n\right] < \varepsilon$$

25

using Markov's inequality and $\mathbb{E}_n[h_n] = M/N$.

(iii) It is sufficient to prove that for any given $\varepsilon > 0$, there exists $c > 0$ such that $\mathbb{P}_\omega\left[\,|C_{n,\text{LSM}} - Z_n| < c\,\right] < \varepsilon$ for any sufficiently large N. From Assumption 2, we can choose c so that

$$\mathbb{P}_\omega\left[\,|C_{n,M} - Z_n| < 2c\,\right] < \frac{\varepsilon}{2}\,,$$

From Clément et al. (2002, Lemma 3.2), which is also based on Assumption 2, $\hat{C}^{[i]}_{\text{LSM}}(s)$ almost surely converges to $C_M(s)$ for a large N:

$$\mathbb{P}_\omega\left[\,|C_{n,\text{LSM}} - C_{n,M}| \geq c\,\right] < \frac{\varepsilon}{2}.$$

Here, the choice of c does not depend on M because Assumption 2 holds uniformly on M. Then,

$$
\begin{aligned}
\mathbb{P}_\omega\left[\,|C_{n,\text{LSM}} - Z_n| < c\,\right] &\leq \mathbb{P}_\omega\left[\,\{|C_{n,\text{LSM}} - Z_n| < c\} \cap \{|C_{n,\text{LSM}} - C_{n,M}| < c\}\,\right] + \mathbb{P}_\omega\left[\,|C_{n,\text{LSM}} - C_{n,M}| \geq c\,\right] \\
&< \mathbb{P}_\omega\left[\,|C_{n,M} - Z_n| < 2c\,\right] + \mathbb{P}_\omega\left[\,|C_{n,\text{LSM}} - C_{n,M}| \geq c\,\right] \\
&= \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon. \quad \square
\end{aligned}
$$

Theorem 1. *The following hold under Assumptions 1 and 2.*

(i) $B_n^{[i]} \sim O_p(M/N)$.

(ii) For any given $\varepsilon > 0$, there exists $r_\varepsilon > 0$ such that the expected look-ahead bias satisfies $\mathbb{E}_\omega[\hat{B}] \leq \varepsilon + r_\varepsilon M/N$.

(iii) $\hat{B}$ converges to zero in probability.

Here, the probabilistic asymptotic notation O_p is defined in the probability space of all possible simulation runs of size N. The subscript for path n in (i) is a dummy index because the Monte Carlo paths are drawn independently.

PROOF. Proof of Theorem 1.

(i) We prove this Theorem inductively. First, $B_n^{[I]} = 0$, and we assume that $B_n^{[i+1]} \sim O_p(M/N)$ for $i < I$. We define E_n, F_n and G_n as

$$
\begin{aligned}
E_n &\iff \{h_n > kM/N\} \cup \{|C_{n,\text{LSM}} - Z_n| < c\}, \\
F_n &\iff \{|Y_n^{[i+1]} - Z_n| > l\}, \\
G_n &\iff \{B_n^{[i+1]} > sM/N\}.
\end{aligned}
$$

For $\varepsilon > 0$, we choose c, k, l, and s such that $\mathbb{P}_\omega[E_n \cup F_n \cup G_n] < \varepsilon$ by Lemma 3 and the induction assumption. If $I[E_n \cup F_n \cup G_n] = 0$, then

$$
\begin{aligned}
B_n^{[i]} &\leq I\left[\,|C_{n,\text{LSM}} - Z_n| \leq h_n|Y_n^{[i+1]} - Z_n|\,\right] |Y_n^{[i+1]} - Z_n| + B_n^{[i+1]} \\
&\leq I\left[c \leq \frac{kM}{N}|Y_n^{[i+1]} - Z_n|\right] |Y_n^{[i+1]} - Z_n| + B_n^{[i+1]} \leq \left(\frac{kl^2}{c} + s\right)\frac{M}{N}
\end{aligned}
$$

by Lemma 2 and Markov's inequality. Finally,

$$\mathbb{P}_\omega\left[B_n > \left(\frac{kl^2}{c} + s\right)\frac{M}{N}\right] \leq \mathbb{E}_\omega\left[I[E_n \cup F_n \cup G_n]\right] < \varepsilon.$$

Therefore, $B_n^{[i]} \sim O_p(M/N)$.

(ii) Set $\varepsilon > 0$. From Lemma 3, we can choose c and k for E_n such that

$$\mathbb{P}_\omega[E_n]\,\mathbb{E}_\omega[(Y_n^{[i+1]} - Z_n)^2] < \frac{\varepsilon^2}{(I-1)^2}$$

for any sufficiently large N. Then,

$$
\begin{aligned}
\mathbb{E}_\omega[\,B_n^{[i]} - B_n^{[i+1]}] &\leq \mathbb{E}_\omega\left[[\,|C_{n,\mathrm{LSM}} - Z_n| \leq h_n|Y_n^{[i+1]} - Z_n|\,] \cdot |Y_n^{[i+1]} - Z_n|\right] \qquad \text{(by Lemma 2)} \\
&\leq \mathbb{E}_\omega\left[I[E_n] \cdot |Y_n^{[i+1]} - Z_n|\right] \\
&\quad + \mathbb{E}_\omega\left[I[E_n^c \cap (|C_{n,\mathrm{LSM}}^{[i]} - Z_n| \leq h_n|Y_n^{[i+1]} - Z_n|)] \cdot |Y_n^{[i+1]} - Z_n|\right] \\
&\leq \mathbb{E}_\omega\left[I[E_n] \cdot |Y_n^{[i+1]} - Z_n|\right] + \mathbb{E}_\omega\left[I\left[c \leq \frac{kM}{N}|Y_n^{[i+1]} - Z_n|\right] \cdot |Y_n^{[i+1]} - Z_n|\right] \\
&\leq \mathbb{E}_\omega\left[I[E_n] \cdot |Y_n^{[i+1]} - Z_n|\right] + \frac{kM}{cN}\mathbb{E}_\omega\left[(Y_n^{[i+1]} - Z_n)^2\right] \qquad \text{(by Markov's inequality)} \\
&\leq \left(\mathbb{P}_\omega[E_n]\,\mathbb{E}_\omega\left[(Y_n^{[i+1]} - Z_n)^2\right]\right)^{1/2} + \frac{kM}{cN}\mathbb{E}_\omega\left[(Y_n^{[i+1]} - Z_n)^2\right] \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by the Cauchy–Schwarz inequality)} \\
&< \frac{\varepsilon}{I-1} + \frac{M}{N}r_\varepsilon^{[i]}. \qquad\qquad\qquad\qquad\qquad \text{(for some } r_\varepsilon^{[i]}, \text{ by Lemma 3(i))}
\end{aligned}
$$

Finally, we can aggregate the step-wise bounds for the incremental bias to obtain an upper bound for the overall bias:

$$\mathbb{E}_\omega[\hat{B}] = \mathbb{E}_\omega[\,\mathbb{E}_n[B_n^{[1]}]] = \mathbb{E}_\omega[B_n^{[1]}] = \sum_{i=1}^{I-1}\mathbb{E}_\omega[B_n^{[i]} - B_n^{[i+1]}] < \varepsilon + \frac{M}{N}\sum_{i=1}^{I-1}r_\varepsilon^{[i]}.$$

This completes the proof.

(iii) For any given δ and ε, we can choose ε_1 and ε_2 such that $\mathbb{E}_\omega[B_n^{[1]}] \leq \varepsilon_1 + r_{\varepsilon_1}M/N < \delta\varepsilon$ for any $M/N < \varepsilon_2$. Then,

$$\mathbb{P}_\omega[\hat{B} > \delta] \leq \frac{1}{\delta}\mathbb{E}_\omega[\hat{B}] = \frac{1}{\delta}\mathbb{E}_\omega[B_n^{[1]}] < \varepsilon.$$

Therefore, $\hat{B} \sim o_p(1)$. $\square$

References

Andersen, L., Broadie, M., 2004. Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options. Management Science 50, 1222–1234. doi:10.1287/mnsc.1040.0258.

Bacinello, A.R., Biffis, E., Millossovich, P., 2010. Regression-based algorithms for life insurance contracts with surrender guarantees. Quantitative Finance 10, 1077–1090. doi:10.1080/14697680902960242.

Belomestny, D., 2011. Pricing Bermudan options by nonparametric regression: Optimal rates of convergence for lower estimates. Finance and Stochastics 15, 655–683. doi:10.1007/s00780-010-0132-x.

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, 1342–1361. doi:10.1016/j.jedc.2013.03.004.

Boyle, P.P., 1988. A Lattice Framework for Option Pricing with Two State Variables. The Journal of Financial and Quantitative Analysis 23, 1–12. doi:10.2307/2331019.

Boyle, P.P., Evnine, J., Gibbs, S., 1989. Numerical evaluation of multivariate contingent claims. Review of Financial Studies 2, 241–250.

Brace, A., Gatarek, D., Musiela, M., 1997. The Market Model of Interest Rate Dynamics. Mathematical Finance 7, 127–155. doi:10.1111/1467-9965.00028.

Brennan, M.J., Schwartz, E.S., 1977. The Valuation of American Put Options. The Journal of Finance 32, 449–462. doi:10.2307/2326779.

Broadie, M., Glasserman, P., 1997. Pricing American-style securities using simulation. Journal of Economic Dynamics and Control 21, 1323–1352. doi: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, 35–72. doi:10.21314/JCF.2004.117.

Carriere, J.F., 1996. Valuation of the early-exercise price for options using simulations and nonparametric regression. Insurance: Mathematics and Economics 19, 19–30. doi:10.1016/S0167-6687(96)00004-2.

Chen, J., Sit, T., Wong, H.Y., 2019. Simulation-based Value-at-Risk for nonlinear portfolios. Quantitative Finance 19, 1639–1658. doi:10.1080/14697688.2019.1598568.

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, 627–644. doi:10.1002/fut.21909, arXiv:1805.03172.

Clément, E., Lamberton, D., Protter, P., 2002. An analysis of a least squares regression method for American option pricing. Finance and Stochastics 6, 449–471. doi:10.1007/s007800200071.

Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option pricing: A simplified approach. Journal of Financial Economics 7, 229–263. doi:10.1016/0304-405X(79)90015-1.

Fabozzi, F.J., Paletta, T., Tunaru, R., 2017. An improved least squares Monte Carlo valuation method based on heteroscedasticity. European Journal of Operational Research 263, 698–706. doi:10.1016/j.ejor.2017.05.048.

Feng, L., Lin, X., 2013. Pricing Bermudan Options in Lévy Process Models. SIAM Journal on Financial Mathematics 4, 474–493. doi:10.1137/120881063.

Fries, C.P., 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, C.P., 2008. Foresight Bias and Suboptimality Correction in Monte-Carlo Pricing of Options with Early Exercise, in: Bonilla, L.L., Moscoso, M., Platero, G., Vega, J.M. (Eds.), Progress in Industrial Mathematics at ECMI 2006. 2008 edition ed., pp. 645–649. doi:10.1007/978-3-540-71992-2_107.

Fu, M.C., Laprise, S.B., Madan, D.B., Su, Y., Wu, R., 2001. Pricing American options: A comparison of Monte Carlo simulation approaches. Journal of Computational Finance 4, 39–88. doi:10.21314/JCF.2001.066.

Glasserman, P., 2003. Chapter 8. Pricing American Options, in: Monte Carlo Methods in Financial Engineering. 2003 edition ed.. New York, pp. 421–478.

Glasserman, P., Yu, B., 2004. Number of paths versus number of basis functions in American option pricing. The Annals of Applied Probability 14, 2090–2119. doi:10.1214/105051604000000846.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition. 2nd edition ed., New York, NY. URL: https://web.stanford.edu/~hastie/ElemStatLearn/.

Haugh, M.B., Kogan, L., 2004. Pricing American Options: A Duality Approach. Operations Research 52, 258–270. doi:10.1287/opre.1030.0070.

He, H., 1990. Convergence from Discrete- to Continuous-Time Contingent Claims Prices. The Review of Financial Studies 3, 523–546.

Huang, Y.T., Kwok, Y.K., 2016. Regression-based Monte Carlo methods for stochastic control models: Variable annuities with lifelong guarantees. Quantitative Finance 16, 905–928. doi:10.1080/14697688.2015.1088962.

Hunter, C.J., Jäckel, P., Joshi, M.S., 2001. Drift Approximations in a Forward-Rate-Based LIBOR Market Model. Risk , 10.

Ibáñez, A., Velasco, C., 2018. The optimal method for pricing Bermudan options by simulation. Mathematical Finance 28, 1143–1180. doi:10.1111/mafi.12158.

Jamshidian, F., 1997. LIBOR and swap market models and measures. Finance and Stochastics 1, 293–330. doi:10.1007/s007800050026.

Joshi, M., Tang, R., 2014. Effective sub-simulation-free upper bounds for the Monte Carlo pricing of callable derivatives and various improvements to existing methodologies. Journal of Economic Dynamics and Control 40, 25–45. doi:10.1016/j.jedc.2013.12.001.

Joshi, M.S., Rebonato, R., 2003. A displaced-diffusion stochastic volatility LIBOR market model: Motivation, definition and implementation. Quantitative Finance 3, 458–469. doi:10.1088/1469-7688/3/6/305.

Kolodko, A., Schoenmakers, J., 2006. Iterative construction of the optimal Bermudan stopping time. Finance and Stochastics 10, 27–49. doi:10.1007/s00780-005-0168-5.

Krekel, M., de Kock, J., Korn, R., Man, T.K., 2004. An analysis of pricing methods for basket options. Wilmott Magazine 2004, 82–89.

Létourneau, P., Stentoft, L., 2014. Refining the least squares Monte Carlo method by imposing structure. Quantitative Finance 14, 495–507. doi:10.1080/14697688.2013.787543.

Longstaff, F.A., Schwartz, E.S., 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. The Review of Financial Studies 14, 113–147. doi:10.1093/rfs/14.1.113.

Ludkovski, M., 2018. Kriging metamodels and experimental design for Bermudan option pricing. Journal of Computational Finance 22, 37–77. doi:10.21314/JCF.2018.347.

Mohammadi, M., 2016. On the bounds for diagonal and off-diagonal elements of the hat matrix in the linear regression model. Revstat–Statistical Journal 14, 75–87.

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, 196–204. doi:10.1016/j.ejor.2016.06.020.

Piterbarg, V., 2003. A Practitioner's Guide to Pricing and Hedging Callable Libor Exotics in Forward Libor Models. SSRN Electronic Journal doi:10.2139/ssrn.427084.

Rubinstein, M., 1991. Somewhere Over the Rainbow. Risk 1991, 63–66.

Stentoft, L., 2004. Convergence of the Least Squares Monte Carlo Approach to American Option Valuation. Management Science 50, 1193–1203. doi:10.1287/mnsc.1030.0155.

Stentoft, L., 2014. Value function approximation or stopping time approximation: A comparison of two recent numerical methods for American option pricing using simulation and regression. Journal of Computational Finance 18, 65–120. doi:10.21314/JCF.2014.281.

Tilley, J.A., 1993. Valuing American options in a path simulation model. Transactions of the Society of Actuaries 45, 499–550.

Tompaidis, S., Yang, C., 2014. Pricing American-style options by Monte Carlo simulation: Alternatives to ordinary least squares. Journal of Computational Finance 18, 121–143. doi:10.21314/JCF.2014.279.

Tsitsiklis, J.N., Van Roy, B., 2001. Regression methods for pricing complex American-style options. IEEE Transactions on Neural Networks 12, 694–703. doi:10.1109/72.935083.

Zanger, D.Z., 2018. Convergence of a Least-Squares Monte Carlo Algorithm for American Option Pricing with Dependent Sample Data. Mathematical Finance 28, 447–479. doi:10.1111/mafi.12125.