European Journal of Operational Research 263 (2017) 698–706

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Interfaces with Other Disciplines

An improved least squares Monte Carlo valuation method based on heteroscedasticity

Frank J. Fabozzi[a,*], Tommaso Paletta[b], Radu Tunaru[b]

[a] EDHEC Business School, BP3116, Nice Cedex 3, 06202, France
[b] Kent Business School, University of Kent, Canterbury CT2 7PE, UK

ARTICLE INFO

Article history:
Received 1 September 2016
Accepted 24 May 2017
Available online 31 May 2017

Keywords:
Finance
American options
Heteroscedasticity
Weighted least squares
Least squares Monte Carlo pricing method

ABSTRACT

Longstaff–Schwartz's least squares Monte Carlo method is one of the most applied numerical methods for pricing American-style derivatives. We examine the algorithms regression step, demonstrating that the OLS regression is not the best linear unbiased estimator because of heteroscedasticity. We prove the existence of heteroscedasticity for single-asset and multi-asset payoffs numerically and theoretically, and propose weighted-least squares MC valuation method to correct for it. An extensive numerical study shows that the proposed method produces significantly smaller pricing bias than the Longstaff–Schwartz method under several well-known price dynamics. An empirical pricing exercise using market data confirms the advantages of the improved method.

1. Introduction

The problem of pricing American-style derivatives has been extensively examined over the past 40 years.[1] A long series of papers have focused on the approximation of the conditional expected payoff to the option holder from continuation. While they all use regression methods in a dynamic programming context, they have distinctive features. Carriere (1996) estimates the continuation value along each simulated path by employing spline regressions and regressions with a local polynomial smoother, while Tsitsiklis and Van Roy (2001) and Longstaff and Schwartz (2001) employ the ordinary least squares (OLS) regression.

The regression-based methods for pricing American options are centered on the least squares Monte Carlo (LSMC) method described in Longstaff and Schwartz (2001). Stentoft (2014) justified the widespread use of LSMC by noting that it has the smallest absolute bias and less error accumulation when compared to other regression-based algorithms. Longstaff and Schwartz (2001) proved the convergence for problems with one state variable and only one exercise date (except maturity). Clement, Lamberton, and Protter (2002) showed that, for a given set of basis functions, the error resulting from Monte Carlo simulation goes to zero when the number of paths goes to infinity. Within a multi-dimensional and multi-period setting, Stentoft (2004b) proved convergence as the number of basis functions M and the number of paths n_s go to infinity with $M^3/n_s \to 0$. Studying the LSMC near the beginning of the contract when the time-step size approaches zero, Mostovyi (2013) found that the regression problem is ill-posed, making the LSMC unstable.

On the computational side, Moreno and Navas (2003); Stentoft (2004a) and Areal, Rodrigues, and Armada (2008) assessed the pricing performance of LSMC under different numbers of simulated paths, payoff structures and polynomial families in the regressions. They argued that the performance of LSMC is virtually the same for vanilla options when different polynomial families are employed but that their selection has a major impact in the case of exotic options. Wang and Caflisch (2010) modified LSMC to calculate directly the delta and the gamma parameters.

The literature has shown that the pricing bias in the LSMC is a combination of the downward bias caused by the approximation of this curve by a finite low-dimensional polynomial, and the upward bias caused by using the same paths to estimate the optimal stopping time (see, among others, Létourneau & Stentoft, 2014). Létourneau and Stentoft (2014) employed the linear inequality

* Corresponding author.
E-mail addresses: fabozzi321@aol.com, frank.fabozzi@edhec.edu (F.J. Fabozzi).

[1] Recently, several important applications emerged that have successfully employed the LSMC algorithm in fields other than the American option pricing problem. Jarrow, Li, Liu, and Wu (2010) priced callable bonds via the LSMC method, showing that the same technique can be applied to mortgage-backed securities. Carmona and Ludkovski (2010) utilised the LSMC for optimal switching models with inventory to evaluate energy storage facilities. Broadie and Detemple (1996); 2004); Glasserman (2003) and Detemple (2005) provide reviews of the literature related to this issue.

http://dx.doi.org/10.1016/j.ejor.2017.05.048
0377-2217/

constrained least squares (ICLS) method to impose monotonicity and convexity properties on the continuation value curve, as the theoretical results suggest. They showed that the ICLS algorithm is less prone to curve-overfitting compared to LSMC and thus the upward pricing bias is substantially reduced.

Another improvement for the Monte Carlo regression used in American option pricing has been described in Belomestny (2011) and Belomestny, Dickmann, and Nagapetyan (2015) where they apply local polynomial kernel regression to the problem of pricing Bermudan options. The idea is to generate an additional independent set of Monte Carlo sample paths to the sample already used in the prior regression step and then average the payoffs stopped according to a simple rule that, although suboptimal, is capable of producing a low-biased estimate for the option price that has improved convergence properties, as discussed in Zanger (2017). Other authors that used methods for generating a new set of independent random paths corresponding to the underlying process at each exercise time increment, independent of all the other sets of paths generated at all other time-steps, were Egloff, Kohler, and Todorovic (2007); Glasserman and Yu (2004) and Zanger (2013). A survey of regression-based Monte Carlo methods for pricing American options can be found in Kohler (2010) and an excellent discussion of the convergence of various algorithms proposed for pricing American options is contained in Zanger (2013).

The calibration and parameter estimation processes can be sometimes impossible to separate and model misspecification may be difficult to disentangle given the information available in the options market, as pointed out by Jarrow and Kwok (2015). Even when the data generating process is fully known, Monte Carlo pricing methods coupled with least-squares algorithms may be subject to inefficient parameter estimation. This seems to be the case in the literature employing LSMC, as we highlight here, and this phenomenon goes beyond the geometric Brownian motion standard assumption widely utilized when pricing American options.

In this paper, we propose an improved pricing method which we refer to as the weighted least squares Monte Carlo (wLSMC) for American put option pricing. The wLSMC, similar in structure to the LSMC, employs the weighted least squares regression (WLS) method instead of the OLS method. We proceed by proving that the homoscedasticity of the errors, one of the assumptions underpinning the OLS method, does not hold for the regressions in the LSMC. Consequently, the errors of the regressions of LSMC are heteroscedastic, a condition which makes the OLS estimators not the best linear unbiased estimators (BLUE). We show that in LSMC, the OLS estimators tend to exhibit large pricing bias because they are more prone to overfitting the continuation value curve. Our analysis extends to the multi-asset American option pricing case, where similar results are valid. Here we also emphasize the importance of an improved estimation and we provide numerical evidence that correcting for heteroscedasticity in our proposed wLSMC method also improves the option price estimators.

The outline of this paper is as follows. In Section 2 we review the American option pricing problem and the LSMC method in Longstaff and Schwartz (2001). Section 3 provides substantial evidence on the existence of heteroscedasticity in each regression step of the LSMC method. After introducing the wLSMC method in Section 4, we compare the pricing performances of the LSMC, ICLS and wLSMC methods under several price dynamics and show how the wLSMC reduces significantly the upward pricing bias of LSMC and ICLS. Section 5 expands the results in the previous sections to multi-asset payoffs. Section 6 highlights the application of our method for stochastic volatility models. A detailed numerical and empirical analysis based of the performance of the new method is provided in the Online Appendix. Section 7 concludes our paper.

2. American options and the LSMC method

Consider the filtered probability space $\left(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{Q}\right)$ associated with a financial market consisting of three assets: a bank account $dM_t = rM_t dt$, where the risk-free interest rate r is assumed constant over time, a risky asset with the dynamics $\{S_t\}_{t\,\geq\,0}$ given under the risk-neutral measure $\mathbb{Q}$ as $S_t = S_0 e^{s_t}$, where $S_0 > 0$ and $\{s_t\}_{t\,\geq\,0}$ is a Markovian process with $s_0 = 0$, and an American put option written on the risky asset (usually referred to as the underlying asset) with strike price K and maturity date T. The pricing problem for the American put option can be formulated as the problem to find the optimal expected discounted payoff given by $\sup_{\tau\in\Gamma} E_0[h(S_\tau)|S_0]$, where $h(S_t) = e^{-rt}\max\{0, K - S_t\}$ is the payoff in time-0 dollars to the option holder from exercise at time t and Γ is the class of admissible stopping times in (0, T]. The numerical applications we carry out in this paper are for the four different dynamics outlined in the Online Appendix: geometric Brownian motion, exponential Ornstein–Uhlenbeck process, log-normal jump-diffusion process and double exponential jump diffusion process. In addition, we also investigate two stochastic volatility models.

Numerical methods usually restrict the pricing of American options to contracts that can be exercised only at a fixed set of exercise opportunities $t_1 < t_2 < \cdots < t_m = T$ and $t_0 = 0$, the time of evaluation, is not usually part of this set. Without loss of generality, we can assume that $\Delta_{t_i} = t_{i+1} - t_i = T/m = \Delta_t$, for any $i = 0, \ldots, m-1$. Henceforth, to simplify the notation under the discrete-time settings, we denote the underlying asset price at the ith exercise opportunity (the one at time t_i), simply as S_i so the logarithmic return over the period (t_i, t_{i+1}) will be $s_{i+1} - s_i$; the payoff function in time-0 dollars for exercise at time t_i when current state is $S_i = \mathcal{X}$ as

$$h_i(\mathcal{X}) = r_{0,i}\max\{0, K - \mathcal{X}\} \tag{1}$$

where $r_{0,i} = e^{-r_i\Delta_t}$ and $V_i(\mathcal{X})$ is the value in time-0 dollars of the option at time t_i given $S_i = \mathcal{X}$, which is calculated with the dynamic programming:[2]

$$\begin{cases} V_m(\mathcal{X}) = h_m(\mathcal{X}) & (2)\\ V_i(\mathcal{X}) = \max\{h_i(\mathcal{X}), C_i(\mathcal{X})\},\ i = 0, \ldots, m-1 & (3)\end{cases}$$

where

$$C_i(\mathcal{X}) = E_i[V_{i+1}(S_{i+1})|S_i = \mathcal{X}] \tag{4}$$

is the continuation value of the American put option measured in time-0 dollars conditioned on the current state $\mathcal{X}$ and $E_i[\cdot]$ is the expectation operator under the risk-neutral measure $\mathbb{Q}$. One is ultimately interested in $V_0(S_0)$. Furthermore, let us define S_{f_i} as the underling asset price such that $h_i\left(S_{f_i}\right) = C_i\left(S_{f_i}\right)$ which is commonly referred to as the optimal exercise price (OEP). By employing the optimal exercise price,[3] an equivalent formulation of problem (2) and (3) for American put options is

$$\begin{cases} V_m(\mathcal{X}) = h_m(\mathcal{X}) & (5)\\ V_i(\mathcal{X}) = \begin{cases} h_i(\mathcal{X}) & \text{if } \mathcal{X} \leq S_{f_i}\\ C_i(\mathcal{X}) & \text{if } \mathcal{X} > S_{f_i}\end{cases}, \quad i = 0, \ldots, m-1. & (6)\end{cases}$$

[2] Note that time t_0 is excluded from the set of exercise opportunities by choosing $h_0(S_0) = 0$.

[3] Chockalingam and Muthuraman (2015) introduced the approximate moving boundaries method which iteratively finds an approximation of the OEP while Chockalingam and Feng (2015) extended Ibanez and Paraskevopoulos (2011) and investigated the cost of using suboptimal OEP. For long-term American options Fabozzi, Paletta, Stanescu, and Tunaru (2016) designed a construction method for the OEP based on an approximation of the optimal exercise price near the beginning of the contract combined with existing quasi-analytical pricing approaches for the remaining part.