Comput Manage Sci
DOI 10.1007/s10287-008-0082-3

ORIGINAL PAPER

American option pricing under stochastic volatility: an efficient numerical approach

Farid AitSahlia · Manisha Goswami · Suchandan Guha

Abstract This paper develops a new numerical technique to price an American option written upon an underlying asset that follows a bivariate diffusion process. The technique presented here exploits the supermartingale representation of an American option price together with a coarse approximation of its early exercise surface that is based on an efficient implementation of the least-squares Monte–Carlo algorithm (LSM) of Longstaff and Schwartz (Rev Financ Stud 14:113–147, 2001). Our approach also has the advantage of avoiding two main issues associated with LSM, namely its inherent bias and the basis functions selection problem. Extensive numerical results show that our approach yields very accurate prices in a computationally efficient manner. Finally, the flexibility of our method allows for its extension to a much larger class of optimal stopping problems than addressed in this paper.

Keywords American option pricing · Optimal stopping · Approximate dynamic programming · Stochastic volatility · Doob–Meyer decomposition · Monte–Carlo

1 Introduction

The early exercise feature of an American option, which enables its holder to select the time at which to buy (call option) or sell (put option) the underlying asset, is the main source of its pricing challenge. In contrast, the price of the corresponding European option, which can only be exercised at the option expiration, may be determined in closed-form in a few important instances. One such example is the classical Black–Scholes formula developed in Black and Scholes (1973), which

F. AitSahlia (✉) · M. Goswami · S. Guha
Department of Industrial and Systems Engineering, Weil Hall 303,
University of Florida, Gainesville, FL 32611, USA
e-mail: farid@ise.ufl.edu

results from the assumption of the underlying asset following a geometric Brownian motion with constant coefficients. Ever since the publication of this landmark formula, significant efforts have concentrated on developing computational pricing techniques for models where the underlying asset follows more realistic processes that capture empirical observations. These include leptokurtic and skewed asset returns, the leverage effect, whereby declining asset returns are observed together with larger variations, and volatility clustering, which describes alternating, sustained periods of high and low variability. Among the more popular approaches that have been adopted, by both academics and practitioners, the stochastic volatility model of Heston (1993) stands out particularly as it possesses the features just described. It also has the advantage of resulting in a (nearly) closed-form formula for the price of the corresponding European option. As we shall soon demonstrate, this feature is particularly useful in our approach for American options, which relies on the supermartingale (Doob–Meyer) decomposition of their price into a sum of the associated European option price and the early-exercise premium. In fact, this decomposition formula has been developed for both the classical version for the Black–Scholes model (cf. Kim 1990; Jacka 1991; Carr et al. 1992) and the Heston model (Chiarella and Ziogas 2005). The early-exercise premium involves the early exercise surface that is notoriously difficult to evaluate. In the classical constant-coefficient set-up of Black and Scholes (1973) and Merton (1973), there is now a well-developed efficient procedure to determine the early exercise boundary via spline approximations (cf. AitSahlia and Lai 1999, 2001). In the present paper the boundary surface is first approximated through a coarse and fast simulation based on the Longstaff and Schwartz (2001) algorithm and then used in the American option price decomposition formula. As an application, when the stochastic volatility of the underlying asset follows a mean-reverting process, we show that the approach is fast and accurate through diverse numerical scenarios.

We remark that ever since its publication, the Longstaff and Schwartz algorithm has been widely adopted for its simplicity and its flexibility. However, its actual implementation requires the selection of basis functions and the set of exercisable dates. While its theoretical convergence rests on the number of each becoming infinite (cf. Clement et al. 2002), practical considerations restrict them to be finite obviously. As there is little guidance for these choices (cf. Tianhai and Burrage 2003; Moreno and Navas 2003; Zhou 2004), and given the systematic bias observed by Longstaff and Schwartz (2001) for their algorithm, our approach is practical and avoids precisely these issues.

This paper is organized as follows. In the next section, the stochastic volatility model of Heston (1993) is reviewed. Section 3 develops our approximation approach to price American options under this model. Section 4 contains a systematic numerical evaluation of this approximation and Sect. 5 concludes.

2 Heston pricing model

In this model, the volatility of the underlying asset is assumed to be stochastic and its square (the variance) follows a mean-reverting process that indicates its tendency to return to a long-term average. Specifically, denoting the underlying price process by $\{S(t)\}$ and its return volatility process (hereafter in the variance sense, as is common) by $\{V(t)\}$, we have the bivariate specification:

$$dS(t) = (r-q)S(t)dt + \sqrt{V(t)}S(t)dW_1(t), \tag{1}$$

$$dV(t) = \kappa(\theta - V(t))dt + \sqrt{V(t)}\sigma_v\left(\rho dW_1(t) + \sqrt{1-\rho^2}dW_2(t)\right), \tag{2}$$

where r and q denote the risk-free rate and the dividend yield, respectively. W_1 and W_2 are two independent standard Brownian motions defined on a common underlying complete filtered probability space $\left(\Omega, (\mathcal{F}_t)_t, Q\right)$, where Q is the risk-neutral measure. The volatility $V(t)$ therefore evolves as a mean-reverting square root process with a rate of mean reversion κ, long-term mean θ and volatility of volatility σ_v. This square root process specification is particularly appealing as it guarantees that $V(t)$ remains positive as long as $2\kappa\theta \geq \sigma_v^2$ (cf. Feller 1951).

For illustrative purposes, we consider an American call option with strike K. Let $C_A(S, v, \tau)$ denote its price when the underlying has price S and spot volatility v, with τ units of time left to expiry. Using standard arbitrage arguments, C_A can be shown to satisfy the following partial differential equation

$$\begin{aligned}\frac{\partial C_A}{\partial \tau} &= \frac{vS^2}{2}\frac{\partial^2 C_A}{\partial S^2} + \rho\sigma v S\frac{\partial^2 C_A}{\partial S \partial v} + \frac{\sigma^2 v}{2}\frac{\partial^2 C_A}{\partial v^2} \\ &\quad + (r-q)S\frac{\partial C_A}{\partial S} + (\kappa[\theta - v] - v\lambda)\frac{\partial C_A}{\partial v} - rC_A\end{aligned} \tag{3}$$

in the region $\mathcal{D} = \{0 \leq \tau \leq T,\ 0 \leq S \leq b(v,\tau),\ 0 \leq v < \infty\}$ along with the boundary conditions

$$\begin{aligned}&C_A(S, v, 0) = \max(S-K, 0), \\ &C_A(b(v,\tau), v, \tau) = b(v,\tau) - K, \\ &\lim_{S \to b(v,\tau)} \frac{\partial C_A}{\partial S} = 1, \\ &\lim_{S \to b(v,\tau)} \frac{\partial C_A}{\partial v} = 0,\end{aligned}$$

where $b(v, t)$ denotes the optimal early exercise price (boundary) at time t for spot volatility v, and λv denotes the corresponding market price of volatility risk, with λ determined empirically. This market price of risk approach is a common way to address the market incompleteness that is inherent in the stochastic volatility formulation (cf. Heston 1993).

Chiarella and Ziogas (2005) use the method of Jamshidian (1992) to convert the homogeneous PDE (3) defined in the region $\mathcal{D}$ above to an inhomogeneous one in an unrestricted domain:

Table 1 American put option prices (constant volatility) with strike = 100, $\sigma = 0.4$ and $\rho = r/\sigma^2 = 1.22$

T	S0	BM	SA	1k5	10k5	10k10	50k10	100k5	100k10	100k25
0.021	80	20.002	20.000	19.951	19.960	19.980	19.980	19.960	19.980	19.992
	90	10.000	10.017	9.845	9.972	9.937	9.958	9.952	9.941	9.927
	100	2.138	2.151	2.122	2.123	2.114	2.112	2.121	2.112	2.106
	120	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
0.083	80	20.002	20.002	19.850	19.897	19.943	19.961	19.895	19.947	19.969
	90	10.354	10.377	10.433	10.523	10.408	10.487	10.502	10.428	10.356
	100	3.942	3.992	4.011	4.008	3.971	3.974	4.007	3.964	3.937
	120	0.190	0.226	0.228	0.227	0.223	0.223	0.227	0.223	0.221
0.146	80	20.002	20.075	19.904	20.009	20.007	20.052	19.980	20.017	20.006
	90	10.842	10.924	11.056	11.205	11.007	11.072	11.150	11.012	10.893
	100	4.979	5.028	5.080	5.132	5.049	5.043	5.116	5.038	4.979
	120	0.651	0.660	0.670	0.673	0.660	0.659	0.671	0.659	0.650

Columns BM and SA are, respectively, the benchmark and spline approximation values from Table 2 of AitSahlia and Lai (1999). MkN denotes prices obtained using $M \times 1{,}000$ sample paths and N time steps in our method

$$\frac{\partial C_A}{\partial \tau} = \frac{v}{2}\frac{\partial^2 C_A}{\partial x^2} + \rho\sigma v\frac{\partial^2 C_A}{\partial x \partial v} + \frac{\sigma^2 v}{2}\frac{\partial^2 C_A}{\partial v^2} + \left(r - q - \frac{v}{2}\right)\frac{\partial C_A}{\partial x}$$
$$+(\alpha - \beta v)\frac{\partial C_A}{\partial v} - H(x - \ln b(v,\tau))\{e^{r\tau}(qe^x - rK)\}, \tag{4}$$

where $\alpha \equiv \kappa\theta$ and $\beta \equiv \kappa + \lambda$, in the unrestricted domain $-\infty < x < \infty$, $0 < v < \infty$, $0 \leq \tau \leq T$, subject to the boundary conditions:

$$C_A(x, v, 0) = \max(e^x - K, 0),$$
$$\lim_{x \to \ln b(v,\tau)} \frac{\partial C_A}{\partial x} = b(v,\tau)e^{r\tau},$$
$$\lim_{x \to \ln b(v,\tau)} \frac{\partial C_A}{\partial v} = 0,$$

where $H(x)$ is the Heaviside step function defined as

$$H(x) = \begin{cases} 1, & x > 0, \\ \frac{1}{2}, & x = 0, \\ 0, & x < 0. \end{cases}$$

To obtain C_A through Eq. (4), one still needs the knowledge of the optimal stopping (early exercise) boundary $b(v, t)$. In the classical context of constant volatility for the underlying asset return, AitSahlia and Lai (1999) have shown that this boundary is well-approximated by linear splines with very few knots, typically 3 or 4. When the

Table 2 American call option with strike = 100, $r = 0.03$, $q = 0.05$, $\sigma = 0.1$, $\kappa = 4.0$, $\rho = 0$, $\theta = 0225$

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
1	90	1.545 (1.536–1.554)	1.637	1.589	1.560	1.558	1.589	1.557	1.550
	95	2.878 (2.865–2.89)	3.055	2.970	2.919	2.918	2.960	2.916	2.903
	100	4.872 (4.855–4.888)	5.123	4.991	4.902	4.904	4.958	4.903	4.879
	105	7.53 (7.511–7.55)	7.875	7.685	7.540	7.543	7.612	7.547	7.509
	110	10.828 (10.805–10.851)	11.290	11.020	10.804	10.804	10.898	10.814	10.775
0.8	90	1.221 (1.213–1.228)	1.264	1.263	1.241	1.242	1.250	1.245	1.232
	95	2.495 (2.484–2.506)	2.557	2.556	2.520	2.519	2.540	2.524	2.506
	100	4.46 (4.445–4.475)	4.537	4.545	4.487	4.484	4.523	4.489	4.466
	105	7.149 (7.131–7.168)	7.251	7.274	7.179	7.183	7.237	7.185	7.154
	110	10.537 (10.515–10.558)	10.660	10.704	10.549	10.569	10.634	10.571	10.531
0.6	90	0.871 (0.865–0.877)	0.901	0.895	0.889	0.890	0.898	0.887	0.883
	95	2.034 (2.025–2.043)	2.077	2.067	2.049	2.053	2.067	2.050	2.040
	100	3.956 (3.943–3.969)	4.028	4.019	3.983	3.989	4.014	3.991	3.970
	105	6.726 (6.709–6.742)	6.815	6.816	6.754	6.761	6.802	6.767	6.737
	110	10.26 (10.241–10.28)	10.369	10.383	10.282	10.297	10.353	10.303	10.269
0.4	90	0.502 (0.498–0.506)	0.515	0.513	0.507	0.507	0.514	0.510	0.505
	95	1.465 (1.458–1.472)	1.488	1.490	1.478	1.480	1.490	1.481	1.473
	100	3.338 (3.327–3.349)	3.374	3.374	3.354	3.358	3.373	3.355	3.344
	105	6.219 (6.205–6.234)	6.286	6.281	6.248	6.250	6.279	6.249	6.232
	110	9.99 (9.973–10.007)	10.068	10.071	10.019	10.022	10.064	10.026	10.000

Table 2 continued

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
0.2	90	0.139	0.141	0.140	0.139	0.139	0.142	0.140	0.140
		(0.137–0.14)							
	95	0.746	0.753	0.753	0.749	0.749	0.756	0.750	0.747
		(0.742–0.75)							
	100	2.455	2.478	2.477	2.466	2.466	2.477	2.467	2.459
		(2.448–2.463)							
	105	5.579	5.618	5.616	5.598	5.598	5.614	5.597	5.586
		(5.568–5.59)							
	110	9.784	9.841	9.832	9.809	9.810	9.831	9.809	9.797
		(9.77–9.797)							

Columns 1 and 2 refer to maturity and spot asset price, respectively. Column "LSM" refers to the LSM benchmark prices generated as in Sect. 4.2 along with 95% CI. MkN denotes prices obtained using $M \times 1{,}000$ sample paths and N time steps in our method

volatility of the underlying asset itself follows a stochastic process as in (2) above, Broadie et al. (2000) produced empirical evidence to suggest that the corresponding optimal stopping surface can be well-approximated in a log-linear fashion near the long-term variance level; i.e.:

$$\ln b(v,\tau) \approx b_0(\tau) + v b_1(\tau), \qquad \text{for } v \text{ near } \theta,$$

thus reducing the determination of $b(v,\tau)$ to that of $b_0(\tau)$ and $b_1(\tau)$. Under this assumption, Chiarella and Ziogas (2005) then express the solution for the PDE (4) as the following decomposition formula:

$$\begin{aligned} C_A(S,v,\tau) &= Se^{-q\tau} P_1(S,v,\tau,K;0) - Ke^{-r\tau} P_2(S,v,\tau,K;0) \\ &\quad + \int_0^\tau qSe^{-q(\tau-\xi)} P_1(S,v,\tau-\xi,e^{b_0(\xi)};-b_1(\xi))d\xi \\ &\quad - \int_0^\tau rKe^{-r(\tau-\xi)} P_2(S,v,\tau-\xi,e^{b_0(\xi)};-b_1(\xi))d\xi, \end{aligned} \tag{5}$$

where

$$P_j(S,v,\tau-\xi,b;w) = \frac{1}{2} + \frac{1}{\Pi}\int_0^\infty Re\left(f_j(S,v,T-\xi;\phi,w)\frac{e^{-i\phi\ln b}}{i\phi}\right)d\phi \tag{6}$$

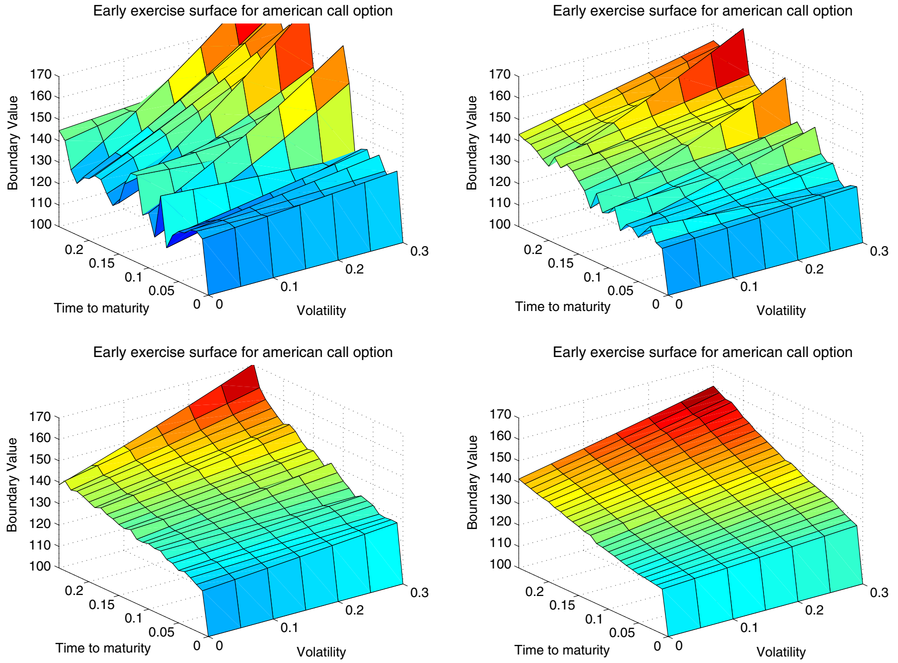

Fig. 1 Early Exercise Boundary for American call option with $T = 0.25$, $\sigma = 0.04$, $\theta = 0.1$, $r = 0.03$, $q = 0.05$ evaluated using 1,000 (*top-left*), 10,000 (*top-right*), 100,000 (*lower-left*) and 1 Million sample paths (*lower-right*), respectively (Number of time steps = 25)

for $j = 1, 2$ and

$$f_1(S, v, \tau-\xi; \phi, w) = e^{-\ln S} e^{-(r-q)(\tau-\xi)} f_2(S, v, T-\xi; \phi, w)$$
$$f_2(x, v, \tau-\xi; \phi, \psi) = \exp\{g_0(\phi, \psi, \tau-\xi) + g_1(\phi, \psi, \tau-\xi)x + g_2(\phi, \psi, \tau-\xi)v\},$$

with

$$g_0(\phi, \psi, \tau-\xi) = (r-q)i\phi(\tau-\xi) + \frac{\alpha}{\sigma^2}\left\{(\beta - \rho\sigma i\phi + D_2)(\tau-\xi) - 2\ln\left[\frac{1 - G_2(\psi)e^{D_2(\tau-\xi)}}{1 - G_2(\psi)}\right]\right\},$$
$$g_1(\phi, \psi, \tau-\xi) = i\phi,$$
$$g_2(\phi, \psi, \tau-\xi) = i\psi + \frac{\beta - \rho\sigma i\phi - \sigma^2 i\psi + D_2}{\sigma^2}\left[\frac{1 - e^{D_2(\tau-\xi)}}{1 - G_2(\psi)e^{D_2(\tau-\xi)}}\right],$$
$$D_2^2 \equiv (\rho\sigma i\phi - \beta)^2 + \sigma^2\phi(\phi + i),$$
$$G_2(\psi) \equiv \frac{\beta - \rho\sigma i\phi - \sigma^2 i\psi + D_2}{\beta - \rho\sigma i\phi - \sigma^2 i\psi - D_2}.$$

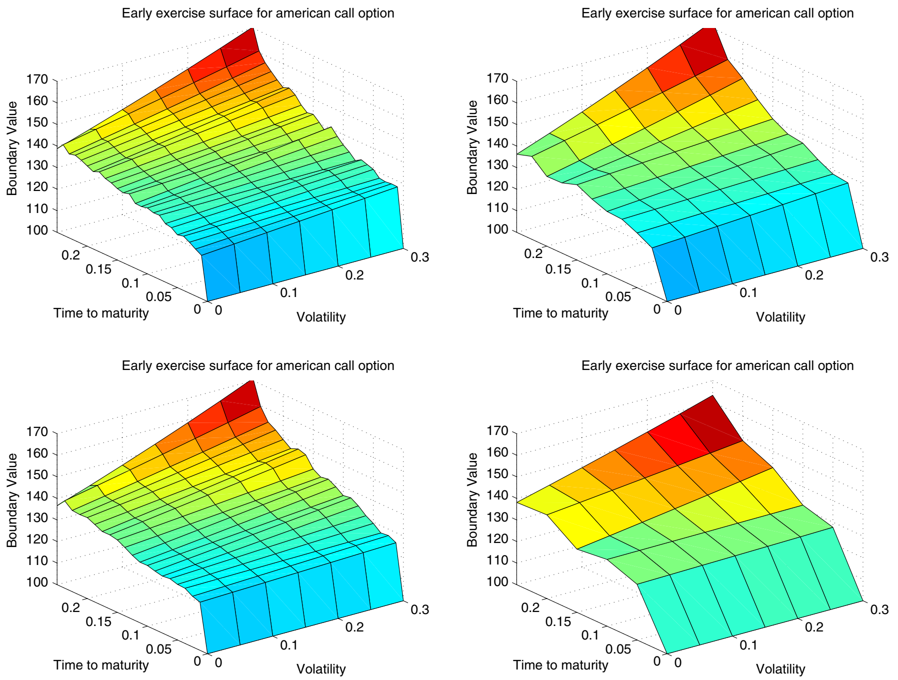

Fig. 2 Early Exercise Boundary for American call option with $T = 0.25$, $\sigma = 0.04$, $\theta = 0.1$, $r = 0.03$, $q = 0.05$ evaluated using 25 (*top-left*), 20 (*top-right*), 10 (*lower-left*) and 5 time steps (*lower-right*), respectively (Number of sample paths = 10,000)

In order to determine approximately the terms $b_0(\tau)$ and $b_1(\tau)$ for every τ, we shall rely on a coarse implementation of the least-squares Monte-Carlo (LSM) algorithm of Longstaff and Schwartz (2001). This flexible method uses a combination of Monte-Carlo simulation with least-squares regression to evaluate American option prices. In our adaptation of this approach, we shall estimate $b_0(\tau)$ and $b_1(\tau)$ over a finite subset of discrete dates $\tau_1, \tau_2, \ldots, \tau_N$, with N very small, typically between 5 and 10, and a small number M of sample paths, just a few thousands, compared to simulation runs several orders of magnitude larger required for accurate results through Monte-Carlo. In fact, our approach is motivated in part by the constant volatility results, which indicate that the exact knowledge of the boundary for the integral representation expression is not critical for the accuracy of the option price calculation (cf. AitSahlia and Lai 2001). In addition, Glasserman (2004) also shows that a simulation-based valuation of American option prices does not critically depend on an accurate evaluation of the optimal exercise strategy.

3 Boundary evaluation

In this section we detail the steps of the approach described above. In particular, we show how the early exercise surface is approximated numerically. Then in the next section we use this approximation through numerical integration in the American option pricing decomposition formula. For our purposes, the early exercise surface will be approximately determined by the LSM method. As a Monte-Carlo based technique, it will generate discrete sample values $\hat{S}_i$ and $\hat{V}_i$ of the stock and its variance, respectively, by discretizing the associated stochastic differential Eqs (1)–(2). A natural choice for this purpose is the Euler scheme:

$$\hat{S}_{i+1} = \hat{S}_i + \mu \hat{S}_i \Delta t + \sqrt{\hat{V}_i} S_t \Delta W_1,$$
$$\hat{V}_{i+1} = \hat{V}_i + \kappa(\theta - \hat{V}_i)h + \sqrt{\hat{V}_i}\sigma_v \left(\rho \Delta W_1 + \sqrt{1-\rho^2}\Delta W_2\right),$$

where ΔW_1 and ΔW_2 are independent standard normal random variables with variance h, which is defined as the time mesh-size. However, we follow Glasserman (2004) who suggests that the second-order scheme of Milstein (1978) and Talay (1982) given below has a better convergence (less bias) for option pricing applications:

$$\begin{aligned}
\hat{S}_{i+1} &= \hat{S}_i \left(1 + rh + \sqrt{\hat{V}_i}\Delta W_1\right) + \frac{1}{2} r^2 \hat{S}_i h^2 \\
&\quad + \left(\left[r + \frac{\rho - \kappa}{4}\right]\hat{S}_i \sqrt{\hat{V}_i} + \left[\frac{\kappa\theta}{4} - \frac{1}{16}\right]\frac{\hat{S}_i}{\sqrt{\hat{V}_i}}\right)\Delta W_1 h \\
&\quad + \frac{1}{2}\hat{S}_i \left(\hat{V}_i + \frac{\rho}{2}\right)(\Delta W_1^2 - h) + \frac{1}{4}\sqrt{1-\rho^2}\,\hat{S}_i\,(\Delta W_2 \Delta W_1 + \xi), \\
V_{i+1} &= \kappa\theta h + (1 - \kappa h)\hat{V}_i + \sqrt{\hat{V}_i}\left(\rho\Delta W_1 + \sqrt{1-\rho^2}\Delta W_2\right) - \frac{1}{2}\kappa^2(\theta - \hat{V}_i)h^2 \\
&\quad + \left(\left[\frac{\kappa\theta}{4} - \frac{1}{16}\right]\frac{1}{\sqrt{\hat{V}_i}} - \frac{3\kappa}{2}\sqrt{\hat{V}_i}\right)\left(\rho\Delta W_1 + \sqrt{1-\rho^2}\Delta W_2\right) h \\
&\quad + \frac{1}{4}\rho^2(\Delta W_1^2 - h) + \frac{1}{4}(1-\rho^2)(\Delta W_2^2 - h) + \frac{1}{2}\rho\sqrt{1-\rho^2}\Delta W_1 \Delta W_2,
\end{aligned}$$

where ξ is a random variable independent of ΔW_1 and ΔW_2 such that $P\{\xi = h\} = P\{\xi = -h\} = 1/2$.

3.1 Boundary evaluation through LSM

Once stock and volatility paths are generated, the LSM method consists of approximating the expected value of continuation by least squares. Specifically, LSM assumes that the option can be exercised at one of N dates, $t_1, t_2, \ldots, t_N = T$, along each of M sample paths generated by Monte-Carlo. In order to calculate the boundary, we perform the following steps for every date (starting from $t_N = T$ and going backwards through $t_{N-1}, t_{N-2}, \ldots t_1$).

Table 3 American call option with strike = 100, $r = 0.03$, $q = 0.05$, $\sigma = 0.1$, $\kappa = 4.0$, $\rho = 0$, $\theta = 0.09$

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
1	90	5.484 (5.477–5.491)	5.599	5.590	5.531	5.492	5.523	5.497	5.476
	95	7.514 (7.506–7.523)	7.652	7.634	7.575	7.519	7.558	7.517	7.497
	100	9.911 (9.902–9.920)	10.097	10.080	9.994	9.911	9.953	9.915	9.894
	105	12.674 (12.664–12.685)	12.888	12.902	12.796	12.669	12.725	12.686	12.656
	110	15.778 (15.767–15.789)	16.038	16.042	15.952	15.782	15.863	15.792	15.765
0.8	90	4.505 (4.498–4.511)	4.570	4.573	4.536	4.518	4.537	4.519	4.504
	95	6.445 (6.438–6.452)	6.537	6.544	6.480	6.456	6.483	6.456	6.440
	100	8.809 (8.800–8.817)	8.911	8.926	8.857	8.821	8.861	8.817	8.798
	105	11.583 (11.574–11.593)	11.724	11.737	11.653	11.595	11.642	11.593	11.569
	110	14.743 (14.732–14.753)	14.890	14.956	14.833	14.746	14.800	14.753	14.729
0.6	90	3.397 (3.392–3.402)	3.433	3.442	3.410	3.412	3.429	3.412	3.400
	95	5.212 (5.206–5.218)	5.263	5.276	5.228	5.229	5.251	5.226	5.210
	100	7.519 (7.512–7.527)	7.585	7.598	7.531	7.535	7.558	7.532	7.512
	105	10.309 (10.300–10.317)	10.374	10.410	10.319	10.323	10.344	10.320	10.299
	110	13.551 (13.541–13.560)	13.633	13.692	13.569	13.565	13.589	13.560	13.538
0.4	90	2.142 (2.138–2.146)	2.158	2.161	2.147	2.150	2.162	2.148	2.143
	95	3.743 (3.738–3.747)	3.769	3.771	3.752	3.754	3.771	3.754	3.745
	100	5.961 (5.955–5.967)	6.002	6.006	5.972	5.973	5.992	5.947	5.962
	105	8.792 (8.785–8.799)	8.837	8.853	8.801	8.807	8.827	8.808	8.792

Table 3 continued

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
	110	12.182	12.250	12.254	12.187	12.204	12.227	12.201	12.183
		(12.174–12.190)							
0.2	90	0.768	0.776	0.775	0.772	0.773	0.777	0.772	0.771
		(0.767–0.770)							
	95	1.936	1.951	1.951	1.944	1.945	1.952	1.944	1.941
		(1.933–1.939)							
	100	3.977	3.993	3.998	3.983	3.985	3.997	3.985	3.978
		(3.973–3.981)							
	105	6.938	6.964	6.968	6.948	6.949	6.964	6.950	6.941
		(6.933–6.943)							
	110	10.693	10.722	10.733	10.705	10.709	10.725	10.710	10.700
		(10.687–10.699)							

Columns 1 and 2 refer to maturity and spot asset price, respectively. Column "LSM" refers to the LSM benchmark prices generated as in Sect. 4.2 along with 95% CI. MkN denotes prices obtained using $M \times 1,000$ sample paths and N time steps in our method

1. For each sample path ω and at every date t_k, we calculate the exercise cash flow $\left\{C(\omega, t_j; t_k, T) : k+1 \le j \le N\right\}$ as in Longstaff and Schwartz (2001). However, in our context of stochastic volatility, we approximate the conditional expected value of continuation $F(S(t_k;\omega), V(t_k;\omega), t_k)$ as

$$F(S(t_k;\omega), V(t_k;\omega), t_k) \approx C_0^{(k)} + C_1^{(k)} L_0\left(\frac{S}{K}\right) + C_2^{(k)} L_1\left(\frac{S}{K}\right) + C_3^{(k)} L_0\left(\frac{S}{K}\frac{V}{\theta}\right), \tag{7}$$

where the argument $(t_k;\omega)$ is omitted from S and V on the right hand side, and K and θ are, respectively, the strike and the long-run mean of the variance in Heston's model. The basis functions for the above regression estimate are Laguerre polynomials $L_0(X) = \exp\left(-\frac{X}{2}\right)$ and $L_1(X) = \exp\left(-\frac{X}{2}\right) \times (1 - X)$. The coefficients $C_0^{(k)}, C_1^{(k)}, C_2^{(k)}$ and $C_3^{(k)}$ in (7) are obtained by regressing the continuation cash flows against the basis functions, with the superscript (k) indicating that the regressions are effected at each time step. As noted earlier, the choice of the basis functions may be debatable for the full-blown implementation of the LSM algorithm, but it is a non-issue for our approach as our goal is to find an approximation to the early exercise surface, the exact determination of which has been shown to matter little.

For each sample path at a given time step, the expected cash flow from continuation is calculated and compared with the current payoff. The option is then exercised if the current payoff is greater than the expected payoff from continuation. This is done for all M stock prices (one for each sample path) at this time step.

Table 4 American call option with strike = 100, $r = 0.03$, $q = 0.05$, $\sigma = 0.1$, $\kappa = 4.0$, $\rho = 0$, $\theta = 02$

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
1	90	11.1305 (11.094–11.167)	11.655	11.579	11.420	11.305	11.311	11.307	11.326
	95	13.536 (13.496–13.576)	14.181	14.115	13.922	13.774	13.788	13.770	13.794
	100	16.192 (16.149–16.235)	16.950	16.903	16.670	16.477	16.493	16.481	16.501
	105	19.019 (18.973–19.065)	19.947	19.925	19.654	19.396	19.410	19.418	19.436
	110	22.0765 (22.028–22.125)	23.148	23.157	22.863	22.539	22.571	22.562	22.587
0.8	90	9.749 (9.715–9.783)	10.039	10.052	9.891	9.832	9.824	9.845	9.837
	95	12.1195 (12.082–12.157)	12.455	12.512	12.311	12.224	12.217	12.242	12.234
	100	14.7255 (14.685–14.766)	15.151	15.247	15.003	14.884	14.888	14.904	14.898
	105	17.5975 (17.554–17.641)	18.098	18.240	17.952	17.800	17.808	17.818	17.816
	110	20.682 (20.636–20.728)	21.250	21.468	21.145	20.953	20.957	20.977	20.976
0.6	90	8.0765 (8.047–8.106)	8.223	8.241	8.127	8.132	8.140	8.141	8.131
	95	10.3565 (10.323–10.39)	10.536	10.581	10.425	10.433	10.445	10.444	10.432
	100	12.943 (12.906–12.98)	13.158	13.235	13.033	13.040	13.061	13.051	13.041
	105	15.813 (15.773–15.853)	16.075	16.186	15.939	15.941	15.966	15.948	15.943
	110	18.9825 (18.94–19.025)	19.250	19.409	19.122	19.115	19.141	19.120	19.119
0.4	90	7.1375 (6.069–6.13)	6.116	6.139	6.087	6.093	6.106	6.100	6.089
	95	8.2345 (8.206–8.263)	8.286	8.319	8.241	8.252	8.268	8.261	8.247
	100	10.7425 (10.711–10.774)	10.832	10.878	10.766	10.784	10.802	10.794	10.778
	105	13.623 (13.588–13.658)	13.737	13.798	13.650	13.673	13.690	13.681	13.664
	110	16.82 (16.782–16.858)	16.974	17.052	16.867	16.891	16.911	16.897	16.882

Table 4 continued

T	S0	LSM	1k5	10k5	10k10	50k10	100k5	100k10	100k25
0.2	90	3.451	3.547	3.469	3.458	3.460	3.465	3.463	3.458
		(3.435–3.467)							
	95	5.3295	5.459	5.356	5.338	5.342	5.349	5.345	5.339
		(5.31–5.349)							
	100	7.7095	7.889	7.759	7.732	7.738	7.745	7.742	7.735
		(7.686–7.733)							
	105	10.614	10.818	10.663	10.623	10.632	10.640	10.636	10.628
		(10.587–10.641)							
	110	13.95	14.199	14.024	13.968	13.981	13.991	13.985	13.977
		(13.92–13.98)							

Columns 1 and 2 refer to maturity and spot asset price, respectively. Column "LSM" refers to the LSM benchmark prices generated as in Sect. 4.2 along with 95% CI. MkN denotes prices obtained using $M \times 1,000$ sample paths and N time steps in our method

(In practice, however, Longstaff and Schwartz (2001) suggest that only in-the-money paths be considered for computational efficiency, which we do as well.)

2. As volatility can theoretically attain any value on the set of positive reals, computational considerations lead to partition them into classes $C_1, C_2, \ldots, C_{n_v}$, with n_v to be fixed. In this way, at every time step t_k, $(S(t_k), V(t_k)) \in C_j$ for some $1 \leq j \leq n_v$ if $V(t_k) \in (V_j, V_{j+1})$, where V_j $(j = 1, \ldots, n_v)$ are equidistant values of volatility, with $V_0 \equiv$ vMin and $V_{n_v} \equiv$ vMax,

$$V_j = V_0 + j\frac{V_{n_v} - V_0}{n_v} \quad \forall \;\; j = 1, \ldots n_v,$$

vMin and vMax are given approximate values according to the distribution of volatility in Eq. (1) for Heston's process. The values n_v, vMin, and vMax clearly depend on the stochastic volatility model and can be set in advance in a number of different ways. In our particular case, we ran numerical simulations of the volatility process alone ahead of the pricing calculations and determined that vMin could be set to its natural value of 0, as it is a variance, and vMax was set to 0.70, which is a conservative estimate as observed variance values are overwhelmingly less 0.06. Numerical experiments were performed as "dry runs" for the fine-tuning of n_v which was then set to 10. We again observe that the fine selection of such parameters does not materially affect the option value thanks to the empirically observed phenomenon of its insensitivity to the fine accuracy of the corresponding exercise surface.

3. For a given time step t_k, the boundary values $B_1(t_k), B_2(t_k), \ldots, B_{n_v}(t_k)$ associated with volatilities $\bar{v}_1, \bar{v}_1, \ldots, \bar{v}_{n_v}$, where $\bar{v}_j = \frac{V_j + V_{j+1}}{2}$), are obtained as

$$B_j(t_k) = \begin{cases} \max\{s_i | (s_i, v_i) \in C_j\} & \text{for put option,} \\ \min\{s_i | (s_i, v_i) \in C_j\} & \text{for call option.} \end{cases}$$

In this way, we are able to obtain boundary value estimates for different volatility ranges at each time step.

4. Finally, given the empirically supported approximation $lnb(v, t) \approx b_0(t)+vb_1(t)$, for v near θ (cf. Broadie et al. 2000), our objective now is to determine $b_0(t_k)$ and $b_1(t_k)$ in the equation

$$\ln B_j(t_k) \approx b_0(t_k) + \bar{v}_j b_1(t_k). \tag{8}$$

For this purpose, $B_1(t_k), B_2(t_k), \dots, B_{n_v}(t_k)$ are then regressed against the mid-points of $\bar{v}_1, \bar{v}_1, \dots, \bar{v}_{n_v}$.

4 Numerical implementation

With given parameters for the Heston model, the price of an American call option is computed by using in (5) the values of $b_0(\tau)$ and $b_1(\tau)$ obtained from the previous section. As can be seen in this expression, in order to evaluate the price of the option, we need to proceed with some numerical integration. The outer integrals for the early exercise premium need to account for the fixed times where the boundary is estimated and are therefore computed using Simpson's rule. However, within these integrals and for the calculations of the European option, which involves Fourier inversion, additional numerical integration is required, over unbounded intervals. We proceed with the use of Gauss-Laguerre quadrature, with our code based on readily available routines from the QuantLib library (http://quantlib.org).

4.1 A constant volatility test

Before we proceed with the validation of our proposed approach on the stochastic volatility model of Heston (1993), we first evaluate it on the well-known constant volatility for which efficient and accurate numerical techniques are available. In this case, our approach is simplified as there is no need to partition the volatility dimension. Table 1 refers to results in Table 2 of AitSahlia and Lai (1999), particularly their last three rows, which correspond to an entire class of American option pricing problems such that $r/\sigma^2 = 1.22$ and $q = 0$, with their canonical time scale s defined as $s = -\sigma^2 T$. For illustrative purposes, we set $\sigma = 0.4$ and the strike $K = 100$, thus capturing their columns with spot prices equal to 80, 90, 100, and 120, respectively. The benchmark values (BM) in Table 1 are theirs re-expressed for our spot prices and are obtained via their converging Bernoulli algorithm. As indicated through this table, our approximation is well within acceptable accuracy for the constant volatility test case.

4.2 Stochastic volatility test

In this case, the benchmark values against which to compare those generated through the proposed numerical approach are computed according to the following two steps:

1. Apply LSM algorithm to estimate the early exercise surface.
2. Stop simulated finer paths at this boundary to obtain option price.

In step 1, we first simulate the stock price paths following the second-order scheme described in Sect. 3. LSM is then applied to these simulated paths to calculate the boundary using the procedure given in 3.1. The calculated values of the boundary are stored for the next step. Step 2 generates a fresh set of sample paths and stops them at the boundary to get the price for each path. In order to check if a particular path can be stopped at a specific time, we must first determine the volatility class of the corresponding stock price and volatility. The benchmark boundary is then generated using 1,000,000 sample paths and time increments of 0.01 year each. Prices are subsequently calculated by generating a fresh set of 10,000,000 sample paths and stopping them at the boundary obtained in the previous step, with the same number of time steps. We should note that the benchmark values thus obtained are not necessarily the most accurate prices for the American option prices but are good (tight) lower bounds. As shown in Clement et al. (2002), in order to obtain complete convergence of the LSM algorithm to the correct American option price, the number of basis functions, that of the sample paths, and that of exercisable dates all have to become very large.

In our numerical experimentation, the parameter values are set as $T = 0.25$, $\rho = 0.0$, $\kappa = 1.0$, $\theta = 0.09$, $\sigma = 0.1$ and various combinations of sample sizes with number of time steps are considered. In addition, different maturities $\tau \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ and spot prices are used. For each value of τ, spot prices s_0 are taken from the set $\{90, 95, 100, 105, 110\}$.

In order to study the effect of changing the number of sample paths, we evaluate the boundary surface using 10 K, 50 K, 100 K and 1,000 K sample paths, keeping the number of time steps constant at 25. As shown in Fig. 1, starting with 1,000 sample paths, the approximated boundary surface becomes progressively better. In fact, there is little noticeable difference between the one generated with 100,000 sample paths and that which is generated with 1 million.

To study the effect of changing the number of time steps, the boundary surface is again evaluated using 5, 10, 20, 50 time steps, keeping the number of sample paths constant at 100 K and the set of parameters as in the previous setting. The obtained boundaries are plotted in Fig. 2, which shows that piecewise linear approximations with only a few time steps and sample paths are warranted.

Next, we compute option prices according to our method. The column heading "MkN" in Table 2 indicates that prices are calculated on the basis of an approximate boundary obtained with N time steps and $M \times 1,000$ sample paths. The column labeled "LSM" refers to prices computed according to the two steps described in Sect. 4.2 above, together with the corresponding 95% confidence intervals.

To assess the accuracy of our approach, for each column we determine whether a price falls in the 95% benchmark confidence interval (CI). If it misses the CI, we record the amount by which it misses and an average is taken for the corresponding column. Tables 2, 3, 4 display the prices obtained for values of $\theta = 0.0225$, 0.09, and 0.2, respectively, with all the other parameters remaining the same as above. Tables 5, 6, 7 list the corresponding summary results. It is important to note that all the prices generated by our method are higher than those labeled LSM. This is a good indication

Table 5 Summary results for $\theta = 0.0225$

	1k5	10k5	10k10	50k10	100k5	100k10	100k25
Max distance from CI of LSM	0.439	0.169	0.029	0.028	0.076	0.026	0.030
Average distance from CI of LSM	0.084	0.058	0.008	0.010	0.040	0.011	0.002
Percentage in CI of LSM	0	4	16	12	0	12	68
Average computation time (s)	0.08	0.32	0.56	2.98	3.04	6.6	16.8

Table 6 Summary results for $\theta = 0.09$

	1k5	10k5	10k10	50k10	100k5	100k10	100k25
Max distance from CI of LSM	0.784	0.432	0.099	0.011	0.035	0.015	0.002
Average distance from CI of LSM	0.174	0.115	0.012	0.004	0.019	0.007	0.000
Percentage in CI of LSM	0	0	24	32	8	16	76
Average computation time (s)	0.08	0.32	0.56	2.98	3.04	6.6	16.8

Table 7 Summary results for $\theta = 0.2$

	1k5	10k5	10k10	50k10	100k5	100k10	100k25
Max distance from CI of LSM	1.023	1.032	0.738	0.414	0.446	0.437	0.462
Average distance from CI of LSM	0.301	0.314	0.156	0.092	0.103	0.101	0.100
Percentage in CI of LSM	4	0	10	12	16	20	24
Average computation time (s)	0.08	0.32	0.57	2.97	3.08	6.63	16.77

of the accuracy of our method as LSM-generated prices are tight lower bounds of the correct American option prices. Thus, it can be observed that the proposed method works generally well, and even better for low standard volatility values (i.e. when the long-run mean of the volatility θ is low). Therefore, the results for $\theta = 0.0225$ are better than those with $\theta = 0.09$, which are in turn better that the ones with $\theta = 0.2$. However, $\theta = 0.09$ corresponds to a volatility of 30% which by itself is a fairly common market volatility value. For each instance of θ, we observe that the price estimate improves as the number of sample paths and the number of time steps increase, as expected.

Also reported in Tables 5, 6, 7 are the average computation times for each column. All computations were performed using a Pentium M processor (1.60 GHz) and 512 MB of RAM. It is clear that a reduction in computation time comes at the cost of increasing error. However it is evident that good computation speed is achieved with very little loss in accuracy, even for the small values of sample paths and time steps.

5 Conclusion

In this paper we presented a novel and efficient numerical technique to price American options when the underlying asset follows a diffusion process indexed by a volatility that itself follows a stochastic process. Extensive numerical tests indicate that this approach is very efficient and accurate. This method significantly improves upon the LSM algorithm of Longstaff and Schwartz (2001) by reducing both computational time and pricing bias. Further work is under way to extend it to stochastic volatility models with random jumps, where the underlying price process also is subject to random shocks (as a result, for example, of unforseen economic developments.) In fact, since the method relies on the combination of a very general result regarding optimal stopping (the Doob–Meyer decomposition of Snell envelopes) together with a very flexible Monte-Carlo approach, further expansions into larger classes of problems is envisioned.

References

AitSahlia F, Lai T (2001) Exercise boundaries and efficient approximations to American option prices and hedge parameters. J Comput Finance 4:85–103

AitSahlia F, Lai TL (1999) A canonical optimal stopping problem for American options and its numerical solution. J Comput Finance 3:33–52

Black F, Scholes M (1973) The pricing of options and corporate liabilities. J Polit Econ 81:637–659

Broadie M, Detemple J, Ghysels E, Torres O (2000) American options with stochastic dividends and volatility: a nonparametric investigation. J Econom 94:53–92

Carr P, Jarrow R, Myneni R (1992) Alternative characterizations of American put options. Math Finance 2:87–106

Chiarella C, Ziogas A (2005) Pricing American options under stochastic volatility. In: Computing in Economics and Finance, 77, Society for Computational Economics

Clement E, Lamberton D, Protter P (2002) An analysis of a least squares regression method for American option pricing. Finance Stochastics 6:449–471

Feller W (1951) Two singular diffusion problems. Ann Math 54:173–182

Glasserman P (2004) Monte-Carlo methods in financial engineering, 1st edn. Springer, Berlin

Heston S (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. Rev Financ Stud 6(2):327–343

Jacka SD (1991) Optimal stopping and the American put. Math Finance 1:1–14

Jamshidian F (1992) An analysis of American options. Rev Futures Mark 11:72–80

Kim IJ (1990) The analytical approximation for the American options. Rev Financ Stud 3:547–572

Longstaff FA, Schwartz ES (2001) Valuing American options by simulation: A simple least-squares approach. Rev Financ Stud 14:113–147

Merton RC (1973) Theory of rational option pricing. Bell J Econ Manage Sci 4:141–183

Milstein GN (1978) A method of second-order accuracy integration of stochastic differential equations. Theory Probab Appl 19:557–562

Moreno M, Navas J (2003) On the robustness of least-squares Monte-Carlo (lsm) for pricing American derivatives. Rev Derivatives Res 6:107–128

Talay D (1982) How to discretize stochastic differential equations, pp. 276–292. Lecture Notes in Mathematics, vol 972. Springer, Berlin

Tianhai T, Burrage K (2003) Accuracy issues of Monte-Carlo methods for valuing American options. ANZIAM J Aust Math Soc 44:C739–C758

Zhou Y (2004) On the existence of an optimal regression complexity in the least-squares Monte-Carlo (lsm) framework for option pricing. In: Proceedings, 39th Actuarial Research Conference, Society of Actuaries