TRUE MARTINGALES FOR UPPER BOUNDS ON BERMUDAN OPTION PRICES UNDER JUMP-DIFFUSION PROCESSES

Helin Zhu
Fan Ye
Enlu Zhou

Industrial & Enterprise Systems Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

ABSTRACT

Fast pricing of American-style options has been a difficult problem since it was first introduced to financial markets in 1970s, especially when the underlying stocks' prices follow some jump-diffusion processes. In this paper, we propose a new algorithm to generate tight upper bounds on the Bermudan option price without nested simulation, under the jump-diffusion setting. By exploiting the martingale representation theorem for jump processes on the dual martingale, we are able to construct a martingale approximation that preserves the martingale property. The resulting upper bound estimator avoids the nested Monte Carlo simulation suffered by the original primal-dual algorithm, therefore significantly improves the computational efficiency. Theoretical analysis is provided to guarantee the quality of the martingale approximation. Numerical experiments are conducted to verify the efficiency of our proposed algorithm.

1 INTRODUCTION

Pricing American-style derivatives (which is essentially an optimal stopping problem) has been an active and challenging problem in the last thirty years, especially when the underlying stocks' prices follow some jump-diffusion processes, as they become more and more critical to investors. To present time, various jump-diffusion models for financial modelling have been proposed to fit the real data in financial markets, including: (i) the normal jump-diffusion model, see Merton (1976); (ii) the jump models based on Levy processes, see Cont and Tankov (2003); (iv) the exponential jump diffusion models, see Kou (2008). All these models are trying to capture some interesting features of the market behaviour that cannot be well explained by the pure-diffusion models, such as the heavy-tail risk suffered by the market. In general, closed-form expressions for the American-style derivatives can hardly be derived under these jump-diffusion models due to the multiple exercise opportunities and the randomness in the underlying asset price caused by both jumps and diffusions. Hence, various numerical methods have been proposed to tackle the American-style option pricing problems under the jump-diffusion models, including: (i) solving the free boundary problems via lattice or differential equation methods, e.g. Feng and Linetsky (2008); (ii) quadratic approximation and piece-wise exponential approximation methods, e.g. Kou and Wang (2004). A thorough study on jump-diffusion models for asset pricing has been done by Kou (2008). In a broader sense, an elegant overview of financial models under jump processes is provided in Cont and Tankov (2003).

Another class of widely-used methods is the Monte Carlo simulation-based method, which has been successfully implemented on option pricing problems under the pure-diffusion models, see Longstaff and Schwartz (2001), Tsitsiklis and van Roy (2001). They are able to approximate the continuation values by regression on certain basis functions sets (called "function bases"), which leads to good suboptimal

exercise strategies and lower bounds on the exact option price. Moreover, their methods bypass "the curse of dimensionality" and scale well with the number of underlying variables, working efficiently for high-dimensional problems under the pure-diffusion models. Though these methods can be naturally adapted to option pricing problems under the jump-diffusion setting, two key questions regarding the effectiveness of these methods remain to be addressed: (i) how to choose the function bases for regression. (ii) how to measure the quality of the lower bounds.

The second question is partially addressed by the dual approach proposed independently by Rogers (2002), Haugh and Kogan (2004), and Anderson and Broadie (2004). They are able to generate the upper bounds on the option price by solving the associated dual problem, which is obtained by subtracting the payoff function by a dual martingale adapted to a proper filtration. In theory, if the dual martingale is the Doob-Meyer martingale part of the option price process, namely the "optimal dual martingale", then the resulting upper bound equals the exact option price. In practice, the optimal dual martingale is not available, but good approximations of it can generate tight upper bounds. With the access to the upper bounds, the quality of suboptimal exercise strategies or lower bounds could be measured empirically by looking at the duality gaps. Glasserman (2004) provides an elegant and thorough overview of the duality theory for option pricing problems.

A lot work has been done following the duality theory. To name a few, Ye and Zhou (2013b) apply the primal-dual approach with particle filtering techniques to optimal stopping problems of partially observable Markov processes. Desai et al. (2012) consider an additional path-wise optimization procedure in constructing the dual martingales for optimal stopping problems. Rogers (2007), Brown et al. (2010) generalize the duality theory to general discrete-time dynamic programming problems and provide a broader interpretation of the dual martingale. From Brown et al. (2010)'s perspective, the dual martingale can be regarded as the penalty for the access to the future information (information relaxation) and different degrees of relaxation result in different levels of upper bounds. In particular, the dual martingales constructed by Haugh and Kogan (2004), Anderson and Broadie (2004) can be interpreted as perfect information relaxation, which means the option holder has access to all the future prices of the underlying assets. Ye and Zhou (2012) consider an additional path-wise optimization technique in constructing the penalties for general dynamic programming problems. Ye and Zhou (2013a) also develop the duality theory for general dynamic programming problems under a continuous-time setting.

The numerical effectiveness of the primal-dual approach has been demonstrated in pricing multidimensional American-style options. A possible deficiency of the algorithm is the heavy computation (quadratic in the number of exercisable periods), caused by the nested simulation in constructing the dual martingale. To address the computational issue, Belomestny et al. (2009) propose an alternative algorithm to generate approximations of optimal dual martingale via non-nested simulation under the Wiener process setting. By exploiting the martingale representation theorem on the optimal dual martingale driven by Wiener processes, they are able to approximate the optimal dual martingale through regressing the integrand on some function bases at finite number of time points. The resulting approximation preserves the martingale property and generates a true upper bound. More importantly, their algorithm avoids nested Monte Carlo simulation and is linear in the number of exercisable periods.

In this paper, we will generalize Belomestny et al. (2009)'s idea of "true martingale" to Bermudan option pricing problems under jump-diffusion processes and provide a new perspective in understanding the structure of the optimal dual martingale, which facilitates us to construct good approximations of it. According to our knowledge, we are among the first to ever consider estimating the upper bounds on American-style option price under the jump-diffusion models. In a greater detail, we have made the following contributions.

- Motivated by Belomestny et al. (2009), we propose a new algorithm, which is referred as the "true martingale algorithm" (T-M algorithm), to compute the upper bounds on the Bermudan option price under the jump-diffusion models. The resulting approximation (called "true martingale approximation") preserves the martingale property, therefore generates true upper bounds on the

Bermudan option price. Moreover, compared with the primal-dual algorithm proposed by Anderson and Broadie (2004) (A-B algorithm), our proposed T-M algorithm avoids the nested Monte Carlo simulation and scales linearly with the exercisable periods, and hence achieves a higher computational efficiency.

- We investigate the numerical effectiveness of Longstaff and Schwartz (2001)'s least-squares regression approach (L-S algorithm) for Bermudan option price under the jump-diffusion models. In particular, we find that by incorporating the European option price under the corresponding pure-diffusion model (referred as the "non-jump European option price") into the function basis of the L-S algorithm, the quality of the induced suboptimal exercise strategies and the lower bounds can be significantly improved.

- Motivated by the explicit structure of the optimal dual martingale (Theorem 1), we propose a function basis that can be employed in our proposed algorithm to obtain upper bounds on the option price. This function basis is also derived based on the non-jump European option price, which is critical to the true martingale approximation and hence the quality of the true upper bounds. By implementing our algorithm together with the A-B algorithm on several sets of numerical experiments, the numerical results demonstrate that both methods can generate tight and stable upper bounds on option price of the same quality; however, we observe that our algorithm is much more efficient than the A-B algorithm in practice due to the relief from nested simulation.

To summarize, the rest of this paper will be organized as follows. In section 2, we describe the Bermudan option pricing problem under general jump-diffusion models and review the dual approach. We develop the true martingale approach and provide error analysis and convergence analysis of it in section 3. Section 4 focuses on the detailed T-M algorithm and its numerical advantages. Numerical experiments are conducted in section 5 to verify the computational efficiency of the T-M algorithm. Conclusion and future directions are given in section 6.

2 MODEL FORMULATION

2.1 Primal Problem

In this paper, we consider a special case of asset price models−jump-diffusion processes, i.e., the asset price $X(t)$ satisfies the following stochastic differential equation (SDE):

$$dX(t) = b(t, X(t)) dt + \sigma(t, X(t)) dW(t) + \int_{\mathbb{R}^d} J(t, X(t), y) \mathscr{P}(dt, dy), \tag{1}$$

where $t \in [0, T]$, $X(t) = [X_1(t), ..., X_n(t)]$ is a random process with a given initial deterministic value $X(0) = X_0 \in \mathbb{R}^n$, $W(t) = [W_1(t), ..., W_{n_w}(t)]$ is a standard vector Wiener process, $\mathscr{P}(dt, dy)$ is the Poisson random measure (see Definition 2.18 in Cont and Tankov (2003)) defined on $[0, T] \times \mathbb{R}^d \subset \mathbb{R}^{d+1}$ with the intensity measure $\mu(dt \times dy)$, the coefficients b, σ and J are functions $b : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$, $\sigma : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^{n_w}$ and $J : \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^n$ satisfying mild continuity conditions (such as uniformly Lipschitz continuous or Holder continuous). Throughout $\mathscr{F} = \{\mathscr{F}_t : 0 \leq t \leq T\}$ is the augmented filtration generated by the Wiener process $W(t)$ and the Poisson random measure $\mathscr{P}$.

We consider a Bermudan option based on $X(t)$ that can be exercised at any date from the time set $\Xi = \{T_0, T_1, ..., T_{\mathscr{J}}\}$, with $T_0 = 0$ and $T_{\mathscr{J}} = T$. Given a pricing measure $\mathbb{Q}$ and the filtration $\mathscr{F}$, when exercising at time $T_j \in \Xi$, the holder of the option will receive a discounted payoff $H_{T_j} := h(T_j, X(T_j))$, where $h(T_j, \cdot)$ is a Lipschitz continuous function. Our problem is to evaluate the price of the Bermudan option, that is, to find

$$Primal : \qquad V_0^* = \sup_{\tau \in \Xi} E\left[h(\tau, X(\tau)) | X(0) = X_0\right], \tag{2}$$

where τ is an exercise strategy (in this case, a stopping time adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$) taking values in Ξ, V_0^* represents the Bermudan option price at time T_0 given the initial asset price X_0.

2.2 Review of Dual Approach

Let $M = \{M_{T_j} : j = 0, ..., \mathscr{J}\}$ with $M_0 = 0$ be a martingale adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$ and $\mathscr{M}$ represents the set of all such martingales. Anderson and Broadie (2004), Haugh and Kogan (2004) show that the dual problem

$$Dual: \quad \inf_{M \in \mathscr{M}} \left(E \left[\max_{0 \leq j \leq \mathscr{J}} \left(H_{T_j} - M_{T_j} \right) | X(0) = X_0 \right] \right) \tag{3}$$

yields the exact option price V_0^*. Moreover, if we let M_{T_j} in (4) be the Doob-Meyer martingale part of the discounted Bermudan price process $V_{T_j}^*$, denoted by $M_{T_j}^*$, then the infimum in (4) is achieved. Precisely, we have:

$$V_0^* = E \left[\max_{0 \leq j \leq \mathscr{J}} \left(H_{T_j} - M_{T_j}^* \right) | X(0) = X_0 \right]. \tag{4}$$

In practice, the optimal dual martingale is not accessible to us. Nevertheless, we can still obtain an upper bound with an arbitrary $M \in \mathscr{M}$ via

$$V_0^{up}(M) = E \left[\max_{0 \leq j \leq \mathscr{J}} \left(H_{T_j} - M_{T_j} \right) | X(0) = X_0 \right]. \tag{5}$$

It is reasonable to expect that, if M_{T_j} is the martingale induced by a good approximation, V_{T_j}, of the option price process $V_{T_j}^*$, then M_{T_j} is close to the optimal dual martingale $M_{T_j}^*$ and the resulting upper bound $V_0^{up}(M)$ should be close to the exact option price V_0^*. Specifically, suppose $V = \{V_{T_j} : j = 0, ..., \mathscr{J}\}$ is some approximation of $V^* = \{V_{T_j}^* : j = 0, ..., \mathscr{J}\}$. Consider the following Doob-Meyer decomposition:

$$V_{T_j} = V_0 + M_{T_j} + U_{T_j}, \quad j = 0, ..., \mathscr{J}, \tag{6}$$

where V_0 is the approximation of the Bermudan option price at time T_0 and U_{T_j} is the residual predictable process. Then we can obtain the martingale component M_{T_j} in principle via the following recursion:

$$M_{T_{j+1}} = M_{T_j} + V_{T_{j+1}} - E_{T_j} \left[V_{T_{j+1}} \right], \quad \text{with} \quad M_0 = 0, \tag{7}$$

where $E_{T_j}[\cdot]$ means the conditional expectation is taken with respect to the filtration $\mathscr{F}_{T_j}$, i.e., $E_{T_j}[\cdot] = E_{T_j}\left[\cdot | \mathscr{F}_{T_j}\right]$.

Haugh and Kogan (2004), Anderson and Broadie (2004) both use the above theoretical result as the starting point of their algorithms to the upper bounds. The difference between their approaches lies in the ways of generating dual martingales. Nevertheless, due to the nested simulation in approximating the conditional expectation in (7), both of their algorithms lose the martingale property. Thus the resulting upper bounds are not guaranteed to be true upper bounds. Furthermore, the nested simulation requires huge computational effort. Under limited computational resources, this approach might not be realistic. In next section, we will develop an alternative approach to address these issues.

3 TRUE MARTINGALE APPROACH VIA NON-NESTED SIMULATION

In this section, we will develop an approach that is fundamentally different from previous approaches by Haugh and Kogan (2004), Anderson and Broadie (2004). By exploiting the special structure of martingales jointly driven by the Wiener measure and the Poisson random measure, we are able to construct an approximation of M without nested simulation, and thus preserves the martingale property. The following generalized martingale representation theorem provides the intuitive idea in understanding the unique structure of such martingales.

Theorem 1 (***Martingale Representation Theorem***) Fix $T > 0$. Let $\{W(t) : 0 \le t \le T\}$ be a n_w-dimensional Wiener process and $\mathscr{P}$ be a Poisson random measure on $[0, T] \times \mathbb{R}^d$ with intensity $\mu(dt \times dy)$, independent from $W(t)$. If $M = \{M_{T_j} : j = 0, ..., \mathscr{J}\}$ is a locally square-integrable (real-valued) martingale adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$ with deterministic initial value $M_0 = 0$, then there exist a predictable process $\phi : \Omega \times [0, T] \to \mathbb{R}^{n_w}$ and a predictable random function $\psi : \Omega \times [0, T] \times \mathbb{R}^d \to \mathbb{R}$ such that

$$M_{T_j} = \int_0^{T_j} \phi_s dW(s) + \int_0^{T_j} \int_{\mathbb{R}^d} \psi(s, y) \tilde{\mathscr{P}}(ds, dy), \qquad j = 0, ..., \mathscr{J}, \tag{8}$$

where $\tilde{\mathscr{P}}$ is the compensated Poisson random measure induced by $\mathscr{P}$.

Remark 1 The proofs of all theorems, corollaries and propositions in this paper are provided in Zhu et al. (2013).

Inspired by Theorem 1 and following Belomestny et al. (2009)'s work, if one tries to approximate the martingale M_{T_j}, a natural idea is to first estimate the integrands ϕ_t and $\psi(t, y)$ in

$$M_{T_j} = \int_0^{T_j} \phi_t dW(t) + \int_0^{T_j} \int_{\mathbb{R}^d} \psi(t, y) \tilde{\mathscr{P}}(dt, dy), \quad j = 0, ..., \mathscr{J}, \tag{9}$$

at a finite number of time and space points. Then an approximation of M_{T_j} will be represented via ϕ_t and $\psi(t, y)$ using the Ito sum (similar to the Riemann sum).

We introduce a partition $\pi = \{t_l : l = 0, 1, ..., \mathscr{L}\}$ on $[0, T]$ such that $t_0 = 0$, $t_{\mathscr{L}} = T$ and $\pi \supset \Xi$, and a partition $\mathscr{A} = \{A_k : k = 0, 1, ..., \mathscr{K}\}$ on $\mathbb{R}^d$ such that $\{[t_l, t_{l+1}] \times A_k\}$ are μ-measurable disjoint subsets and $\bigcup_{k=1}^{\mathscr{K}} A_k = \mathbb{R}^d$. Then $P([t_l, t_{l+1}] \times A_k) = \int_{t_l}^{t_{l+1}} \int_{A_k} \mathscr{P}(ds, dy)$ is a Poisson random variable (regarded as Poisson increment), and $\tilde{P}([t_l, t_{l+1}] \times A_k) = \int_{t_l}^{t_{l+1}} \int_{A_k} \tilde{\mathscr{P}}(ds, dy)$ is the corresponding compensated Poisson random variable (regarded as compensated Poisson increment). We denote the magnitude of partitions π and $\mathscr{A}$ as $|\pi|$ and $|\mathscr{A}|$ respectively, i.e., $|\pi| = \max_{0 < l \le \mathscr{L}} (t_l - t_{l-1})$ and $|\mathscr{A}| = \max_{1 \le k \le \mathscr{K}} \int_{A_k} f(y) dy$.

From (6), we have

$$V_{T_{j+1}} - V_{T_j} = (M_{T_{j+1}} - M_{T_j}) + (U_{T_{j+1}} - U_{T_j}), \quad j = 0, ..., \mathscr{J}. \tag{10}$$

Combining with the Ito sum of $M_{T_{j+1}}$ in (9), we have

$$V_{T_{j+1}} - V_{T_j} \approx \sum_{T_j \le t_l < T_{j+1}} \phi_{t_l}(W(t_{l+1}) - W(t_l)) + \sum_{T_j \le t_l < T_{j+1}} \sum_{k=1}^{\mathscr{K}} \psi(t_l, y_k) \tilde{P}([t_l, t_{l+1}] \times A_k) + U_{T_{j+1}} - U_{T_j}, \tag{11}$$

where $y_k \in A_k$ is a representative value, and we will keep using this notation thereafter. Multiplying both sides of (11) by the Wiener process increment $(W(t_{l+1}) - W(t_l))$ and taking conditional expectations with respect to the filtration $\mathscr{F}_{t_l}$, we obtain

$$\phi_{t_l} \approx \frac{1}{t_{l+1} - t_l} E_{t_l} [(W(t_{l+1}) - W(t_l)) V_{T_{j+1}}], \quad T_j \le t_l < T_{j+1}, \tag{12}$$

where we use the $\mathscr{F}$-predictability of U, the independent increment property of $W(t)$ and the independence between W and $\mathscr{P}$.

Similarly, if we multiply both sides of (11) by the compensated Poisson random variable $\tilde{P}([t_l, t_{l+1}] \times A_k)$ and take the conditional expectations with respect to the filtration $\mathscr{F}_{t_l}$, we can obtain

$$\psi(t_l, y_k) \approx \frac{1}{\mu([t_l, t_{l+1}] \times A_k)} E_{t_l} [\tilde{P}([t_l, t_{l+1}] \times A_k) V_{T_{j+1}}], T_j \le t_l < T_{j+1}, 1 \le k \le \mathscr{K}. \tag{13}$$

Motivated by expressions (12) and (13), we denote the approximation of ϕ_{t_l} and $\psi(t_l, y_k)$ by $\phi_{t_l}^{\pi,\mathscr{A}}$ and $\psi^{\pi,\mathscr{A}}(t_l, y_k)$ respectively, which are defined as follows:

$$\phi_{t_l}^{\pi,\mathscr{A}} := \frac{1}{\Delta_l^{\pi}} E_{t_l}\left[(\Delta^{\pi} W_l) V_{T_{j+1}}\right], \quad T_j \le t_l < T_{j+1}, \tag{14}$$

and

$$\psi^{\pi,\mathscr{A}}(t_l, y_k) := \frac{1}{\mu\left([t_l, t_{l+1}] \times A_k\right)} E_{t_l}\left[\tilde{P}\left([t_l, t_{l+1}] \times A_k\right) V_{T_{j+1}}\right], T_j \le t_l < T_{j+1}, 1 \le k \le \mathscr{K}, \tag{15}$$

where Δ_l^{π} and $\Delta^{\pi} W_l$ represent the increments of time t and the Wiener process $W(t)$ respectively, i.e. $\Delta_l^{\pi} = (t_{l+1} - t_l)$ and $\Delta^{\pi} W_l = (W_{l+1} - W_l)$. Therefore we can construct an approximation of M_{T_j}, denoted by $M_{T_j}^{\pi,\mathscr{A}}$, as

$$M_{T_j}^{\pi,\mathscr{A}} := \sum_{0 \le t_l < T_j} \phi_{t_l}^{\pi,\mathscr{A}} (\Delta^{\pi} W_l) + \sum_{0 \le t_l < T_j} \sum_{k=1}^{\mathscr{K}} \psi^{\pi,\mathscr{A}}(t_l, y_k) \tilde{P}\left([t_l, t_{l+1}] \times A_k\right). \tag{16}$$

The construction procedure of $M_{T_j}^{\pi,\mathscr{A}}$ can be summarized in the following Algorithm 1.

Algorithm 1   *Construction of the Martingale Approximation $M^{\pi,\mathscr{A}}$*

Step 1: Express M_{T_j} as an integral of $\phi(t)$ and $\psi(t, y)$ via (9).

Step 2: Approximate ϕ_{t_l} by $\phi_{t_l}^{\pi,\mathscr{A}}$ via (14) and $\psi(t_l, y_k)$ by $\psi^{\pi,\mathscr{A}}(t_l, y_k)$ via (15) respectively.

Step 3: Construct the approximation of M_{T_j}, denoted by $M_{T_j}^{\pi,\mathscr{A}}$, via (16).

Notice that $M^{\pi,\mathscr{A}} = \{M_{T_j}^{\pi,\mathscr{A}} : j = 0, ..., \mathscr{J}\}$ remains to be a martingale adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$, based on its structure. We formally state this result in the following theorem.

Theorem 2 (*True Martingale*) If an approximation of M, denoted by $M^{\pi,\mathscr{A}}$, is constructed using Algorithm 1, then $M^{\pi,\mathscr{A}}$ is still a martingale adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$.

According to Theorem 2, if we plug $M^{\pi,\mathscr{A}}$ in (5), it is easy to see that $V_0^{up}(M^{\pi,\mathscr{A}})$ is a true upper bound on the Bermudan option price V_0^* in the sense that $V_0^{up}(M^{\pi,\mathscr{A}})$ is an unbiased expectation for a valid upper bound. Moreover, if we adopt the L-S algorithm to solve the primal problem (3), we will obtain a suboptimal exercise strategy $\tilde{\tau}$. Exercising $\tilde{\tau}$ along a certain number of sample paths yields an approximation V_{T_j} of the Bermudan option price at time T_j via $V_{T_j} = E_{T_j}\left[H_{\tilde{\tau}_j}\right]$, where $\tilde{\tau}_j$ means the stopping time $\tilde{\tau}$ takes value greater than or equal to j. Due to the tower property of conditional expectations, we can rewrite (14) and (15) as

$$\phi_{t_l}^{\pi,\mathscr{A}} := \frac{1}{\Delta_l^{\pi}} E_{t_l}\left[(\Delta^{\pi} W_l) H_{\tilde{\tau}_{j+1}}\right], \quad T_j \le t_l < T_{j+1}, \tag{17}$$

and

$$\psi^{\pi,\mathscr{A}}(t_l, y_k) := \frac{1}{\mu\left([t_l, t_{l+1}] \times A_k\right)} E_{t_l}\left[\tilde{P}\left([t_l, t_{l+1}] \times A_k\right) H_{\tilde{\tau}_{j+1}}\right], T_j \le t_l < T_{j+1}, 1 \le k \le \mathscr{K}. \tag{18}$$

Through this we avoid the computation of conditional expectations in (14) and (15), which would incur nested simulation in implementation. Therefore we can estimate $M_{T_j}^{\pi,\mathscr{A}}$ in (16) via non-nested simulation, and hence significantly improve the computational efficiency.

The following theorem provides the convergence analysis of the above true martingale approximation.

Theorem 3 Let M_{T_j} be the martingale component of $V_{T_j} = v(T_j, X_{T_j})$ and $M_{T_j}^{\pi,\mathscr{A}}$ be its approximation obtained via Algorithm 1, where $v(T_j, \cdot)$ are Lipschitz continuous functions. Then there exists a constant $C > 0$ such that

$$E\left[\max_{0 \le j \le \mathscr{J}} |M_{T_j}^{\pi,\mathscr{A}} - M_{T_j}|^2\right] \le C|\pi|.$$

According to the relationship between M and $V_0^{up}(M)$ in (5), we can immediately obtain the following corollary on the quality of uppers bounds.

Corollary 4 Under the assumptions of Theorem 3, we have

$$|V_0^{up}(M^{\pi,\mathscr{A}}) - V_0^{up}(M)|^2 \leq C|\pi|.$$

4 TRUE MARTINGALE ALGORITHM

We will formally describe the T-M algorithm based on the construction of the martingale approximation $M^{\pi,\mathscr{A}}$ in section 3. The outline of the T-M algorithm consists of four steps in order: generating a suboptimal exercise strategy $\tilde{\tau}$, approximating the integrands $\phi^{\pi,\mathscr{A}}$ and $\psi^{\pi,\mathscr{A}}$, constructing the martingale approximation $M^{\pi,\mathscr{A}}$, and generating true upper bounds $V_0^{up}(\hat{M}^{\pi,\mathscr{A}})$ on the option price.

First, let's start with generating the suboptimal exercise strategy $\tilde{\tau}$. It not only provides the lower bound, but also plays an important role in approximating the integrands $\phi^{\pi,\mathscr{A}}$ and $\psi^{\pi,\mathscr{A}}$. We adopt the L-S algorithm to generate the suboptimal exercise strategy $\tilde{\tau}$ and the corresponding approximation of option price process $\bar{V}_{T_j}$ at time T_j, of the form $\bar{V}_{T_j} = v\left(T_j, X_{T_j}^{\bar{\pi},\bar{\mathscr{A}}}\right)$, where $\bar{\pi} \supset \pi$, $\bar{\mathscr{A}} \supset \mathscr{A}$ are employed to simulate the discretized asset price process $\{X^{\bar{\pi},\bar{\mathscr{A}}}\}$.

Second, let us approximate the integrands $\phi^{\pi,\mathscr{A}}$ and $\psi^{\pi,\mathscr{A}}$. To avoid confusion, we denote

$$\begin{cases} \bar{\phi}_{t_l}^{\pi,\mathscr{A}} = \frac{1}{\Delta_l^{\pi}} E_{t_l}\left[(\Delta^{\pi}W_l) v\left(T_{j+1}, X_{T_{j+1}}^{\bar{\pi},\bar{\mathscr{A}}}\right)\right], \quad T_j \leq t_l < T_{j+1} \\ \bar{\psi}^{\pi,\mathscr{A}}(t_l, y_k) = \frac{1}{\mu([t_l,t_{l+1}]\times A_k)} E_{t_l}\left[\tilde{P}([t_l,t_{l+1}]\times A_k) v\left(T_{j+1}, X_{T_{j+1}}^{\bar{\pi},\bar{\mathscr{A}}}\right)\right], T_j \leq t_l < T_{j+1}, 1 \leq k \leq \mathscr{K} \end{cases}$$

as the counterparts of $\phi^{\pi,\mathscr{A}}_{t_l}$ and $\psi^{\pi,\mathscr{A}}(t_l, y_k)$ respectively, under the discretized asset price $X_{T_{j+1}}^{\bar{\pi},\bar{\mathscr{A}}}$. Inspired by Longstaff and Schwartz (2001)'s least-squares regression approach to approximating the continuation values, we apply a similar regression technique to approximate $\bar{\phi}^{\pi}$ and $\bar{\psi}^{\pi,\mathscr{A}}$. Specifically, the function bases chosen to regress $\bar{\phi}_{t_l}^{\pi,\mathscr{A}}$ and $\bar{\psi}^{\pi,\mathscr{A}}(t_l, y_k)$ are row function vectors $\rho^W\left(t_l, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}}\right) = \left(\rho_i^W(t_l, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}})\right)_{i=1,\dots,I_1}$ and $\rho^{\mathscr{P}}\left(t_l, y_k, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}}\right) = \left(\rho_i^{\mathscr{P}}(t_l, y_k, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}})\right)_{i=1,\dots,I_2}$ respectively, where I_1 and I_2 are the dimensions of the function bases. If we simulate N independent samples of Wiener increments $\Delta^{\pi}W_l$, denoted by $\{\Delta_n^{\pi}W_l: l = 1, \dots \mathscr{L}, n = 1, \dots, N\}$, and N independent samples of Poisson increments $P([t_l, t_{l+1}]\times A_k)$, denoted by $\{P_n([t_l,t_{l+1}]\times A_k): l = 1, \dots, \mathscr{L}, k = 1, \dots, \mathscr{K}, n = 1, \dots, N\}$, and based on which we construct the sample paths of the asset price $\{X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\}_{l=0,\dots,\mathscr{L},n=1,\dots,N}$, then we can obtain the regressed coefficients $\hat{\alpha}_{t_l}$ and $\hat{\beta}_{t_l,k}$, for $T_j \leq t_l < T_{j+1}$ and $1 \leq k \leq \mathscr{K}$, via

$$\begin{cases} \hat{\alpha}_{t_l} = \arg\min_{\alpha \in \mathbb{R}^{I_1}} \left\{\sum_{n=1}^N \left|\frac{\Delta_n^{\pi}W_l}{\Delta_l^{\pi}} H_{\tilde{\tau}_{j+1}}\left(X_{T_{j+1},n}^{\bar{\pi},\bar{\mathscr{A}}}\right) - \rho^W\left(t_l, X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\right)\alpha\right|^2\right\} \\ \hat{\beta}_{t_l,k} = \arg\min_{\beta \in \mathbb{R}^{I_2}} \left\{\sum_{n=1}^N \left|\frac{\tilde{P}_n([t_l,t_{l+1}]\times A_k)}{\mu([t_l,t_{l+1}]\times A_k)} H_{\tilde{\tau}_{j+1}}\left(X_{T_{j+1},n}^{\bar{\pi},\bar{\mathscr{A}}}\right) - \rho^{\mathscr{P}}\left(t_l, y_k, X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\right)\beta\right|^2\right\} \end{cases}, \quad (19)$$

where we employ the tower property to avoid nested simulation, as described in (17) and (18). Therefore we can compute the approximations of the integrands $\bar{\phi}_{t_l}^{\pi,\mathscr{A}}$ and $\bar{\psi}^{\pi,\mathscr{A}}(t_l, y_k)$, denoted by $\hat{\phi}^{\pi,\mathscr{A}}(t_l, x)$ and $\hat{\psi}^{\pi,\mathscr{A}}(t_l, y_k, x)$ respectively, via

$$\hat{\phi}^{\pi,\mathscr{A}}(t_l, x) = \rho^W(t_l, x)\hat{\alpha}_{t_l} \quad \text{and} \quad \hat{\psi}^{\pi,\mathscr{A}}(t_l, y_k, x) = \rho^{\mathscr{P}}(t_l, y_k, x)\hat{\beta}_{t_l,k}. \quad (20)$$

Next, with fixed $\hat{\alpha}$ and $\hat{\beta}$, we construct an approximation of $M^{\pi,\mathscr{A}}$, denoted by $\hat{M}^{\pi,\mathscr{A}}$, by combining the approximation $\hat{\phi}^{\pi,\mathscr{A}}$ and $\hat{\psi}^{\pi,\mathscr{A}}$ of the integrands with the Euler scheme of system (1). Precisely, we have

$$\hat{M}_{T_j}^{\pi,\mathscr{A}} := \sum_{0 \leq t_l < T_j} \hat{\phi}^{\pi,\mathscr{A}}\left(t_l, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}}\right)(\Delta^{\pi}W_l) + \sum_{0 \leq t_l < T_j}\sum_{k=1}^{\mathscr{K}} \hat{\psi}^{\pi,\mathscr{A}}\left(t_l, y_k, X_{t_l}^{\bar{\pi},\bar{\mathscr{A}}}\right)\tilde{P}([t_l,t_{l+1}]\times A_k). \quad (21)$$

Obviously, $\hat{M}^{\pi,\mathscr{A}}$ remains to be a martingale adapted to the filtration $\{\mathscr{F}_{T_j} : j = 0, ..., \mathscr{J}\}$. Consequently, $V_0^{up}\left(\hat{M}^{\pi,\mathscr{A}}\right)$ is a true upper bound on the Bermudan option price V_0^*.

Finally, let's estimate $V_0^{up}\left(\hat{M}^{\pi,\mathscr{A}}\right)$ via (5) by simulating a new set of $\bar{N}$ independent sample paths $\{X_n^{\bar{\pi},\bar{\mathscr{A}}} : n = 1, ..., \bar{N}\}$. Precisely, an unbiased estimator for $V_0^{up}\left(\hat{M}^{\pi,\mathscr{A}}\right)$ is given as follows:

$$\hat{V}_0^{up}\left(\hat{M}^{\pi,\mathscr{A}}\right) = \frac{1}{\bar{N}} \sum_{n=1}^{\bar{N}} \max_{0 \le j \le \mathscr{J}} \left[h\left(T_j, X_{T_j,n}^{\bar{\pi},\bar{\mathscr{A}}}\right) - \hat{M}_{T_j,n}^{\pi,\mathscr{A}} \right], \tag{22}$$

where $\hat{M}_{T_j,n}^{\pi,\mathscr{A}}$ represents the realization of $\hat{M}_{T_j}^{\pi,\mathscr{A}}$ along the sample path $X_{T_j,n}^{\bar{\pi},\bar{\mathscr{A}}}$. We can formally summarize these steps in the following Algorithm 2.

Algorithm 2   *True Martingale Algorithm*

Step 1: Apply the L-S algorithm to generate a suboptimal exercise strategy $\tilde{\tau}$.

Step 2: Simulate N independent samples of Wiener increments $\Delta^\pi W_l$ and N independent samples of Poisson increments $P\left([t_l, t_{l+1}] \times A_k\right)$, for $l = 0, ..., \mathscr{L} - 1$ and $k = 1, ..., \mathscr{K}$; construct the sample paths of the asset price $\{X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\}_{l=0,...,\mathscr{L},n=1,...,N}$.

Step 3: Obtain the parameters $\hat{\alpha} = \{\hat{\alpha}_{t_l}\}_{l=0,...,\mathscr{L}}$ and $\hat{\beta} = \{\hat{\beta}(t_l,k)\}_{l=0,...,\mathscr{L},k=0,...,\mathscr{K}}$ via least-squares regression (19) by exercising $\tilde{\tau}$ along the sample paths $\{X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\}_{l=0,...,\mathscr{L},n=1,...,N}$.

Step 4: Simulate a new set of $\bar{N}$ independent sample paths $\{X_{t_l,n}^{\bar{\pi},\bar{\mathscr{A}}}\}_{l=0,...,\mathscr{L},n=1,...,\bar{N}}$; compute $\hat{\phi}^{\pi,\mathscr{A}}$ and $\hat{\psi}^{\pi,\mathscr{A}}$ via (20); construct the martingale approximation $\hat{M}^{\pi,\mathscr{A}}$ via (21); obtain an unbiased estimator $\hat{V}_0^{up}(\hat{M}^{\pi,\mathscr{A}})$ for the true upper bound on the Bermudan option price V_0^* via (22).

---

5   NUMERICAL EXPERIMENTS

In this section, we will conduct numerical experiments to illustrate the computational efficiency of our proposed T-M algorithm on a Bermudan option pricing problem under a jump-diffusion model. The exact model we consider here falls into the class of jump-diffusion models (see Merton (1976) and Kou (2008)) reviewed in section 1. Specifically, the asset prices evolve as follows:

$$\frac{dX(t)}{X(t^-)} = (r - \delta)\,dt + \sigma\,dW(t) + d\left(\sum_{i=1}^{P(t)} (V_i - 1)\right), \tag{23}$$

where r is the constant discount factor, δ is the constant dividend, σ is the constant volatility, $X(t) = [X_1(t), ..., X_n(t)]$ represents the asset price with a given initial price X_0, $W(t) = [W_1(t), ..., W_n(t)]$ is a Wiener process, $P(t)$ is a Poisson process with intensity λ, and $\{V_i\}$ is a sequence of independent identically distributed (i.i.d.) nonnegative random variables such that $J = \log(V)$ is the jump amplitude with p.d.f. $f(y)$. Here J can follow a normal distribution (see (Merton 1976)) or an exponential distribution (see Kou (2008)) or various other reasonable distributions. For simplicity, we assume J follows a one-dimensional ($d = 1$) normal distribution $N(m, \theta^2)$. We also assume $W(t)$, $P(t)$ and J are mutually independent.

To connect dynamics (23) with the jump-diffusion model (1) we have mainly focused on, we should construct a Poisson random measure $\mathscr{P}$ such that dynamics (23) can be easily transformed to an equivalent dynamics jointly driven by the Wiener measure and the Poisson random measure. The following proposition provides an intuitive criterion in selecting such a Poisson random measure $\mathscr{P}$ by explicitly defining the intensity function $\mu(dt \times dy)$ for the unique $\mathscr{P}$ induced by a compound Poisson process.

Proposition 5 (*Proposition 3.5 in Cont and Tankov (2003)*) Let $S(t)_{t>0}$ be a compound Poisson process with intensity λ and jump size distribution f. Then the Poisson random measure $\mathscr{P}_S$ induced by $S(t)_{t>0}$ on $[0, \infty] \times \mathbb{R}^d$ has intensity measure $\mu(dt \times dy) = \lambda f(y)\,dy\,dt$.

According to Proposition 5, for a compound Poisson process $S(t)$, the compensated Poisson random measure $\tilde{\mathscr{P}}_S$ induced by $S(t)$ can be simulated by $\tilde{\mathscr{P}}_S = \mathscr{P}_S - \lambda f(y)\,dy\,dt$. Although $X(t)$ satisfying (23) is not a compound Poisson process, $S(t) = \log(X(t))$ is usually a compound Poisson process, and thus its Poisson random measure $\mathscr{P}_S(t,y)$ can be easily simulated according to Proposition 5. Now if we incorporate $\mathscr{P}_S$ into the asset-price dynamics (23), we can obtain an equivalent dynamics as

$$\frac{dX(t)}{X(t^-)} = (r - \delta)\,dt + \sigma dW(t) + \int_{\mathbb{R}^d} y \mathscr{P}_S(t,y). \tag{24}$$

Unfortunately, the solution to asset dynamics (23) or (24) is not uniquely determined in the risk-neutral sense, caused by the incompleteness of the market under the jump-diffusion setting. However, we can construct pricing measures $\mathbb{Q}'s \sim \mathbb{P}$ such that the discounted price $\hat{X}(t)$ is a martingale under $\mathbb{Q}'s$ (c.f. Chapter 10 in Cont and Tankov (2003)). Here we will adopt the construction method proposed by Merton (1976). That is, changing the drift of the Wiener process but leaving other components of (23) unchanged to offset the jump results in a risk-neutral measure $\mathbb{Q}_M$, which is a generalization of the unique risk-neutral measure under the Black-Scholes model. Therefore, the solution under $\mathbb{Q}_M$ can be easily derived and efficiently simulated. Precisely, the solution to the asset-price dynamics (23) is given by:

$$X(t) = X_0 \exp\left[\mu^M t + \sigma W^M(t) + \sum_{i=1}^{P(t)} J_i\right], \quad t > 0, \tag{25}$$

where $\mu^M = r - \delta - \frac{1}{2}\sigma^2 - E\left[e^{J_i} - 1\right]$ is the new drift, $W^M(t)$ is a standard vector Wiener process and $J_i's$ are the i.i.d. random variables according to J.

Given the equivalence of (23) and (24), we can perform the Euler scheme on an equidistant partition $\bar{\pi}$ with $|\bar{\pi}| = 0.01$ and a continuously equi-probabilistic partition $\bar{\mathscr{A}}$ on $\mathbb{R}^d$ with $|\bar{\mathscr{A}}| = 0.1$ to simulate the Wiener increments $\{W_{t_l}\}$, the Poisson random measure increments $P([t_l, t_{l+1}] \times A_k)$, and the resulting sample paths of $X(t) = \exp(S(t))$ according to (25).

We consider a Bermudan Min-Puts on *n* assets, whose evolution is given by (25). In particular, at any time $t \in \Xi = \{T_0, T_1, ..., T_{\mathscr{J}}\}$, the option holder has the right to exercise his option to receive the payoff $h(X(t)) = (SK - \min(X_1(t), ..., X_n(t)))^+$. The maturity time of the option is $T = 1$ and can be exercised at 11 equally-spaced time points, i.e., $T_j = j \times T/10$, $j = 0, ..., 10$. Our objective is to solve the Bermudan option pricing problem by providing a lower bound and an upper bound on the exact option price.

5.1 Suboptimal Exercise Strategies and Lower Bounds

First, let's adopt the L-S algorithm to generate a suboptimal exercise strategy $\tilde{\tau}$ by regressing the continuation values on certain function bases, and compute the corresponding benchmark lower bound. Anderson and Broadie (2004) propose a function basis consisting of all monomials of underlying asset prices with degrees less than or equal to three, the European min-put option with maturity T, its square and its cube, since the European option under the pure-diffusion model has a closed-form that can be fast numerically evaluated (see Zhu et al. (2013) for explicit formula).

For the Bermudan option pricing problem under the jump-diffusion model (23), the corresponding European option still has a closed-form expression (see Zhu et al. (2013) for explicit formula). However, it is extremely difficult to be evaluated because of the infinite sum and the integral in the closed-form. Naturally, we try to approximate it directly by an European option under a closely-related pure-diffusion model. Surprisingly, the most intuitive one, i.e., the European option under the pure-diffusion model derived simply by removing the jump part of (23) works extremely well in our numerical experiments. To avoid confusion, we refer to it as "non-jump European option".

Now the function basis we choose includes all monomials of underlying asset prices with degrees less than or equal to three, the non-jump European option with maturity T, its square and its cube. With this basis, we implement the L-S algorithm, and obtain suboptimal exercise strategies $\tilde{\tau}'s$ and the corresponding lower bounds, as shown in Table 1.

5.2 Upper Bound by True Martingale Approach

Now let's implement our proposed T-M algorithm (Algorithm 2) described in section 4. Notice that we have addressed almost all the implementation details except the choices of partitions π and $\mathscr{A}$, and the bases ρ^W and $\rho^{\mathscr{P}}$.

First of all, the choice of partition π is essential to balance the tradeoff between the quality of the true martingale approximation and the computational efficiency. It has to be sufficiently small to reduce the overall mean square error between the true martingale approximation and the objective martingale, but not too small so that the computational effort for obtaining martingale approximation $M^{\pi,\mathscr{A}}$ is much less than the computational effort for obtaining the inner sample paths in A-B algorithm. In fact, a good way to achieve this tradeoff is to perform the regression on a rough partition in the beginning, and then interpolate them piece-wisely constant to a finer partition. To maximize the effect of this two-layer regression-interpolation technique, we choose to perform the regression procedure on the original exercisable dates $\Xi = \{T_0, T_1, ..., T_{\mathscr{J}} - 1\}$ and interpolate the regression coefficients piece-wisely constant to the partition $\bar{\pi}$ of the Euler scheme. Secondly, the choice of the partition $\mathscr{A}$ is less restrictive than the choice of π since $|\pi|$ will dominate the error between the martingale approximation and the objective martingale (see Theorem(3)) regardless of the choice of $\mathscr{A}$. For the sake of convenience, we let $\mathscr{A} = \bar{\mathscr{A}}$. Therefore the compensated Poisson increments $\{\tilde{P}([t_l, t_{l+1}] \times A_k)\}$ in (19) are obtained immediately from the simulation of $X^{\bar{\pi},\mathscr{A}}$, and $\mu([t_l, t_{l+1}] \times A_k)$ in (19) equals $\lambda \times 0.01 \times 0.1$ (see Proposition 5). Specifically, we obtain $\{\hat{\alpha}_{T_j}, j = 0, ..., \mathscr{J} - 1\}$ and $\{\hat{\beta}_{T_j,k}, j = 0, ..., \mathscr{J} - 1, k = 1, ..., \mathscr{K}\}$ via the regression (19), and set $\hat{\alpha}_{t_l} = \hat{\alpha}_{T_j}$ for $t_l \in [T_j, T_{j+1})$ and $\hat{\beta}_{t_l,k} = \hat{\beta}_{T_j,k}$ for $t_l \in [T_j, T_{j+1}), k = 1, ..., \mathscr{K}$.

Finally, the choice of the bases ρ^W and $\rho^{\mathscr{P}}$ affects the accuracy of the martingale approximation $\hat{M}^{\pi,\mathscr{A}}$. Notice that European option price is a good basis function for the L-S algorithm. Inspired by Theorem 1, if we apply Ito's lemma on the European option price, the resulting integrands should be good candidates for basis functions. Precisely, we have

$$C^M(t, X_t; T_j) = h(X_{T_j}) - \int_t^{T_j} \frac{\partial C^M(u, X_{u^-}; T_j)}{\partial X} X_{u^-} \sigma dW_u^M - \int_t^{T_j} \int_{\mathbb{R}^d} \left[C^M(u, X_{u^-} \cdot e^y; T_j) - C^M(u, X_{u^-}; T_j) \right] \tilde{\mathscr{P}}_S(du, dy), \quad (26)$$

where $C^M(t, X_t; T_j)$ is the European option price with maturity T_j under pricing measure $\mathbb{Q}_M$. After simple numerical tests, we find out, for $t \in [T_j, T_{j+1})$ and $1 \le k \le \mathscr{K}$, $\rho^W(t, X_{t^-})$ consisting of 1, $\frac{\partial C^{BS}(t, X_{t^-}; T_{j+1})}{\partial X} X_{t^-}$ and $\frac{\partial C^{BS}(t, X_{t^-}; T)}{\partial X} X_{t^-}$, $\rho^{\mathscr{P}}(t, y_k, X_{t^-})$ consisting of $1, C^{BS}(t, X_{t^-} \cdot e^{y_k}; T_{j+1}) - C^{BS}(t, X_{t^-}; T_{j+1})$ and $C^{BS}(t, X_{t^-} \cdot e^{y_k}; T) - C^{BS}(t, X_{t^-}; T)$ yield the tightest upper bounds, where $C^{BS}(t, X_t; T_j)$ is the European option price under the corresponding pure-diffusion model and $y_k \in A_k$ is a representative value.

We report the numerical results on the lower bounds by the L-S algorithm, the benchmark upper bounds by the A-B algorithm and the true upper bounds by the T-M algorithm in Table 1. The small gaps between the lower bounds and the true upper bounds indicate that the T-M algorithm is quite effective in terms of generating tight true upper bounds. The small length of the confidence intervals of the true upper bounds indicates that T-M algorithm generates good approximations of the optimal dual martingales. The CPU time ratios indicate that T-M algorithm achieves a much higher numerical efficiency.

It is instructive to theoretically compare the computational complexity of the T-M algorithm and the A-B algorithm, since the CPU time ratios in Table 1 are quite different for 1-dimensional problems and 2-dimensional problems. We know that the total CPU time is mainly consumed by simulating sample paths and evaluating the basis functions. When $n = 1$, the time for simulating sample paths will significantly dominate the time for evaluating the basis functions because the basis functions, which are European options and their derivatives, reduce to the c.d.f.'s of a standard normal distribution, and hence can be evaluated extremely fast. Therefore, the CPU time ratio will be in the order of the ratio between the numbers of sample paths simulated in both algorithms, which is consistent with the result ($\approx$ 1:400). However, when $n \ge 2$, the basis functions reduce to infinite integrals of the c.d.f.'s of a standard normal distribution, which are relatively time-consuming to evaluate. Therefore, the CPU time ratio should be the ratio between the

total evaluation times of the basis functions in both algorithms. For the A-B algorithm, the total evaluation times is in the order of $(N_2 \times N_3 \times \mathscr{J} \times \mathscr{J})$; for the T-M algorithm, the total evaluation times is in the order of $(\bar{N} \times \mathscr{L} \times \mathscr{K})$. The ratio of the latter versus the former is around 1:20, which is in the approximately same order of the result ($\approx$ 1:9). We can expect the CPU time ratios (T-M algorithm versus A-B algorithm) to remain stable if the dimension of the problem increases, and even further decrease when the number of exercisable periods increases. See Zhu et al. (2013) for more detailed discussion.

Table 1: Bounds (with 95% confidence intervals) for Bermudan Min-put options. The payoff of the min-put option is: $(SK - min(X_1(t), ..., X_n(t)))^+$. The parameters are: $SK = 40, r = 4\%, \delta = 0, \sigma = 20\%, m = 6\%, \theta = 20\%, T = 1, \mathscr{J} = 10$. The jump intensity λ is 1 or 3 and the initial price is $X_0 = (X, ..., X)$ with $X = 36$, as shown in the table. We use $N = 5 \times 10^4$ sample paths to estimate the regression coefficients to determine the suboptimal exercise strategy, and we use $N = 5 \times 10^4$ sample paths to estimate the coefficients $\hat{\alpha}$ and $\hat{\beta}$. We use $N_1 = 10^5$ sample paths to determine the lower bounds. For the implementation of the A-B algorithm, we use $N_2 = 10^3$ outer sample paths and $N_3 = 5 \times 10^2$ inner sample paths to determine the benchmark upper bounds and the confidence intervals of appropriate length. For the implementation of the T-M algorithm, we use $\bar{N} = 2.5 \times 10^3$ sample paths to determine the true upper bounds and the confidence intervals of appropriate length.

			Lower Bound	Upper Bound	Benchmark U-B	CPU Time Ratio
n	λ	X_0	(L-S algorithm)	(T-M algorithm)	(A-B algorithm)	(T-M vs A-B)
1	1	36	5.842 ± 0.031	5.970 ± 0.031	5.899 ± 0.038	$\approx$ 1:400
1	3	36	7.702 ± 0.043	7.899 ± 0.030	7.810 ± 0.053	$\approx$ 1:400
2	1	36	8.133 ± 0.033	8.308 ± 0.045	8.243 ± 0.040	$\approx$ 1:9
2	3	36	9.786 ± 0.045	10.038 ± 0.061	9.989 ± 0.057	$\approx$ 1:9

6 CONCLUSION AND FUTURE DIRECTIONS

In this paper, we propose a true martingale algorithm (T-M algorithm) to fast compute the upper bounds on the Bermudan option prices under the jump-diffusion models, as an alternative approach for the classic A-B algorithm proposed by Anderson and Broadie (2004), especially when the computational budget is limited. The theoretical analysis of our algorithm proves and the numerical results verify that our algorithm generates stable and tight upper bounds with significant reduction of computational effort. Moreover, we explore the structure of the optimal dual martingale for the dual problem and provide an intuitive understanding towards the construction of good approximations of the optimal dual martingale over the space of all adapted martingales.

Furthermore, from the information relaxation point of view (see Brown et al. (2010)), we can gain an intuitive understanding towards the structure of the optimal penalty function. It inspires us to construct good penalty functions over the space of "feasible penalty functions" for general dynamic programming problems, which is still an open area to explore (see Ye and Zhou (2012) for some initial exploration).

REFERENCES

Anderson, L., and M. Broadie. 2004. "Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options". *Management Science* 50 (9): 1222 – 1234.

Belomestny, D., C. Bender, and J. Schoenmakers. 2009. "True Upper Bounds for Bermudan Products via Non-Nested Monte Carlo". *Mathematical Finance* 19:53 – 71.

Brown, D., J. Smith, and P. Sun. 2010. "Information Relaxations and Duality in Stochastic Dynamic Programs". *Operations Research* 58 (4): 758 – 801.

Cont, R., and P. Tankov. 2003. *Financial Modelling with Jump Processes*, Volume 2. Chapman & Hall/CRC.

Desai, V., V. Farias, and C. Moallemi. 2012. "Pathwise Optimization for Optimal Stopping Problems". *Management Science*.

Feng, L., and V. Linetsky. 2008. "Pricing Options in Jump-diffusion Models: An Extrapolation Approach". *Operations Research* 56 (2): 304–325.

Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer.

Haugh, M., and L. Kogan. 2004. "Pricing American Options: A Duality Approach". *Operations Research* 52 (2): 258 – 270.

Kou, G. 2008. "Jump-diffusion Models for Asset Pricing in Financial Engineering". *Chapter 2 of Handbooks in Operations Research and Management Science*.

Kou, G., and H. Wang. 2004. "Option Pricing Under a Double Exponential Jump-diffusion Model". *Management Science*:1178–1192.

Longstaff, F., and E. Schwartz. 2001. "Valuing American Options by Simulation: A Simple Least-Squares Approach". *The Review of Financial Studies* 14 (1): 113 – 147.

Merton, R. 1976. "Option Pricing When Underlying Stock Returns are Discontinuous". *Journal of Financial Economics* 3 (1): 125–144.

Rogers, L. 2002. "Monte Carlo Valuation of American Options". *Mathematical Finance* 12 (3): 271 – 286.

Rogers, L. 2007. "Pathwise Stochastic Optimal Control". *SIAM J.Control Optimization* 46 (3): 1116 – 1132.

Tsitsiklis, J., and B. van Roy. 2001. "Regression Methods for Pricing Complex American-Style Options". *IEEE Transactions on Neural Networks* 12 (4): 694 – 703.

Ye, F., and E. Zhou. 2012. "Parameterized Penalties in the Dual Representation of Markov Decision Processes". In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 870–876. IEEE.

Ye, F., and E. Zhou. 2013a. "Information Relaxation and Duality in Controlled Markov Diffusions". Working Paper.

Ye, F., and E. Zhou. 2013b. "Optimal Stopping of Partially Observable Markov Processes: A Filtering-based Duality Approach". *IEEE Transactions on Automatic Control.. To Appear*.

Zhu, H., F. Ye, and E. Zhou. 2013. "Fast Estimation of True Bounds on Bermudan Option Prices Under Jump-diffusion Processes". Working Paper.

AUTHOR BIOGRAPHIES

HELIN ZHU is a second-year Ph.D student in the Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign. His research interest is on stochastic control. His email address is hzhu11@illinois.edu.

FAN YE is a third-year Ph.D student in the Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign. His research interest is on stochastic control. His email address is fanye2@illinois.edu.

ENLU ZHOU is an Assistant Professor in the Department of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana-Champaign. She received the B.S. degree with highest honors in electrical engineering from Zhejiang University, China, in 2004, and received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, in 2009. Her research interests include stochastic control and simulation optimization, with applications towards financial engineering. Her email address is enluzhou@illinois.edu and her web page is http://publish.illinois.edu/enluzhou/.