

PRICING AMERICAN OPTIONS UNDER PARTIAL OBSERVATION OF STOCHASTIC VOLATILITY

Fan Ye
Enlu Zhou

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

ABSTRACT

Stochastic volatility models capture the impact of time-varying volatility on the financial markets, and hence are heavily used in financial engineering. However, stochastic volatility is not directly observable in reality, but is only “partially” observable through the inference from the observed asset price. Most of the past research studied American option pricing in stochastic volatility models under the assumption that the volatility is fully observable, which often leads to overpricing of the option. In this paper, we treat the problem under the more realistic assumption of partially observable stochastic volatility, and propose a numerical solution method by extending the regression method and the martingale duality approach to the partially observable case. More specifically, we develop a filtering-based martingale duality approach that complements a lower bound on the option price with an approximate upper bound. Numerical experiments show that our method reduces overpricing of the option with a moderate computational cost.

1 INTRODUCTION

American options give the holder the right to sell or to buy an underlying asset at a pre-determined price before the expiration date. Most of the research on pricing American options assumes that the volatility is a constant or follows some stochastic process that is fully observable. Even with these assumptions, the problem usually does not have a closed-form solution, so various numerical methods have been developed, among which those based on Monte Carlo simulation enjoy the advantage of robustness and the hope to overcome the curse of dimensionality. Some important work along this line includes (Tilley 1993) (Carriere 1996) (Raymar and Zwecher 1997) (Broadie and Glasserman 1997) (Tsitsiklis and Roy 2001) (Longstaff and Schwartz 2001) (Garcia 2003) (Ibanez and Zapatero 2004) (Broadie and Glasserman 2004). In particular, (Longstaff and Schwartz 2001) and (Tsitsiklis and Roy 2001) proposed approximate dynamic programming techniques based on the regression idea, to obtain a suboptimal exercise policy with a low-biased estimator for the option price. (Rogers 2002) and (Haugh and Kogan 2004) proposed a martingale duality approach to complement a lower bound with an upper bound. (Andersen and Broadie 2004) later developed an alternative duality-based method, which generates the suboptimal martingale based on the approximate polices rather than the approximate value functions as in (Haugh and Kogan 2004). (Jamshidian 2007) studied multiplicative duality, which was comprehensively compared with the former three additive duality methods by (Chen and Glasserman 2007). One computational disadvantage of these duality approaches is the expensive nested simulation (sometimes also called “simulation within simulation”), so (Glasserman and Yu 2002) and (Belomestny, Bender, and Schoenmakers 2009) suggested non-nested simulation methods: the former uses a choice of basis functions that preserves the martingale property, while the latter utilizes the martingale representation theorem. A different type of approaches based on Jensen’s inequality, including the random tree method by (Broadie and Glasserman 1997) and the stochastic mesh method by (Broadie and Glasserman 2004), generates both low-biased and high-biased estimators for the option price with

a proven asymptotic convergence. On a higher level, the two different types of approaches based on martingale duality and Jensen's inequality can be both viewed as special cases of information relaxation (Brown, Smith, and Sun 2010), utilizing future information to gain a high-biased estimator of the value function. A nice summary of Monte Carlo methods for American option pricing can be found in the book (Glasserman 2004).

In reality, stochastic volatility is not directly observable, but is only “partially” observable in the sense that it can be inferred from the observation of the asset price. Under the assumption that the volatility is fully observable, pricing an American option is equivalent to solving an optimal stopping problem, and the resultant optimal exercise policy is a function of the state variables, i.e., the asset price and the volatility. This exercise policy cannot be replicated in reality due to the unavailability of the volatility, and hence it often leads to overpricing of the option. On the contrary, under the more realistic assumption of partially observable stochastic volatility, it is a partially observable optimal stopping problem, for which the resultant optimal exercise policy only depends on all the available observations of the asset price, and hence it is replicable in reality. To the best of our knowledge, only (Pham, Runggaldier, and Sellami 2005), (Florescu and Viens 2008), (Ludkovski 2009), and (Rambharat and Brockwell 2010) have studied American option pricing under the assumption of partially observable stochastic volatility. (Florescu and Viens 2008) proposed a multinomial tree method that combines with particle filtering; (Ludkovski 2009) and (Rambharat and Brockwell 2010) also utilize the particle filtering technique, and incorporate it into the regression-based approximate dynamic programming approach; whereas (Pham, Runggaldier, and Sellami 2005) used a grid-based method to approximate the filtering distribution.

All of the above four approaches provide approximate solutions, and some are proven to converge asymptotically to the true option price in the limit. However, in practice with a finite computational power, the difference between their approximate solutions and the true value is not known. So it motivates us to develop a lower-and-upper-bound approach such that the gap between the two bounds gives an indication of the quality of the approximate solution. To guarantee a high-quality approximate solution, we can increase the computation until the gap between the two bounds decreases to a desirable tolerance level. To this end, we first use a regression-based method following the idea of (Longstaff and Schwartz 2001) to generate a suboptimal exercise policy and the associated lower bound, and then complement this lower bound with an approximate upper bound by a filtering-based martingale duality approach, which we propose for the partially observable optimal stopping problem. Our lower-bound algorithm is similar to the methods proposed in (Ludkovski 2009) and (Rambharat and Brockwell 2010), but the main difference is that our method guarantees a true lower bound so that it can be used in pair with an upper bound.

The rest of the paper is organized as follows. In Section 2 we describe the problem formulation of pricing American options under partial observation of the stochastic volatility. In Section 3 we briefly introduce the theory of dynamic programming and filtering that is underpinning of our method. In Section 4, we describe the regression-based method for finding a suboptimal exercise policy and the corresponding lower bound on the option price. In Section 5, we develop the filtering-based martingale duality approach to generate an approximate upper bound on the option price. We present some numerical examples in Section 6 and finally conclude in Section 7.

2 PROBLEM FORMULATION

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. The asset price S_t satisfies the following stochastic differential equation (SDE)

$$dS_t = S_t(rdt + \sigma(X_t)dW_t^1), \quad (1)$$

where r is the constant interest rate, $\{W_t^1\}$ is a Wiener process, and the volatility $\sigma(X_t)$ is the deterministic function $\sigma(\cdot)$ of the stochastic process X_t that evolves as

$$dX_t = \alpha(X_t)dt + \beta(X_t)d\tilde{W}_t^2, \quad (2)$$

where $\{\tilde{W}_t^2\}$ is another Wiener process. The two processes $\{W_t^1\}$ and $\{\tilde{W}_t^2\}$ are correlated such that $dW_t^1 d\tilde{W}_t^2 = \rho dt$ with $\rho \in [-1, 1]$. Hence, $\{\tilde{W}_t^2\}$ can be written as a linear combination of two independent Wiener processes $\{W_t^1\}$ and $\{W_t^2\}$ in the form of $\tilde{W}_t^2 = \rho W_t^1 + \sqrt{1 - \rho^2} W_t^2$.

To fix ideas in the following, we consider a more specific form of the above model: the function $\sigma(X_t) = \exp(X_t)$ in (1), and (2) is an Ornstein-Uhlenbeck process

$$dX_t = \lambda(\theta - X_t)dt + \gamma d\tilde{W}_t^2, \quad (3)$$

where the constant λ is the rate of reversion, the positive constant θ is the mean reversion value, and γ is another constant.

Assume that the asset price is observed at equally-spaced time points t_0, t_1, \dots, t_T , simply denoted by $0, 1, \dots, T$, with time period $\delta = t_{i+1} - t_i$. Following the analytical solutions to (1) and (3), they can be discretized as follows:

$$X_{t+1} = X_t e^{-\lambda\delta} + \theta(1 - e^{-\lambda\delta}) + \gamma \sqrt{\frac{1 - e^{-2\lambda\delta}}{2\lambda}} \left(\rho Z_{t+1}^1 + \sqrt{1 - \rho^2} Z_{t+1}^2 \right), \quad t = 0, \dots, T-1, \quad (4)$$

$$S_{t+1} = S_t \exp \left\{ \left(r - \frac{\sigma(X_{t+1})^2}{2} \right) \Delta + \sigma(X_{t+1}) \sqrt{\delta} Z_{t+1}^1 \right\}, \quad t = 0, \dots, T-1, \quad (5)$$

where $\{Z_t^1\}$ and $\{Z_t^2\}$ are two independent sequences of independent and identically distributed (i.i.d.) random variables with standard normal distribution. Then the transition probability $P(X_t | X_{t-1})$ and the observation probability $P(S_t | X_t, S_{t-1})$ are determined. We assume that the initial price S_0 is a known constant, and the initial volatility X_0 follows a known distribution π_0 , which in practice can be chosen as the conditional distribution of X_0 given the historical data.

Let $\mathcal{J} = \{1, 2, \dots, T\}$. Denote by $\mathcal{F}_t^S = \sigma\{S_1, \dots, S_t\}$ the filtration generated by the process (5). In the setting that we cannot observe the volatility directly, the price of the American option with payoff function $g(t, S_t)$ is

$$V_0(s_0, \pi_0) = \sup_{\tau \in \mathcal{J}, \{\mathcal{F}_t^S\}\text{-adpted}} \mathbb{E}[g(\tau, S_\tau) | S_0 = s_0, X_0 \sim \pi_0]. \quad (6)$$

Please note that τ is a $\{\mathcal{F}_t^S\}$ -**adapted stopping time**, i.e., the exercise decision at time t only depends on the observations of the asset price up to time t . For an American put option with strike price K , the payoff function $g(t, S_t) = \max(e^{-r\delta t}(K - S_t), 0)$. For convenience, in the following we use $g(S_t)$ and $g(S_\tau)$ to be short for $g(t, S_t)$ and $g(\tau, S_\tau)$. In the rest of the paper, We will not differentiate ‘‘stopping time’’ and ‘‘exercise policy’’ in the rest of the paper.

3 TRANSFORMATION

The above partially observable problem can be transformed to an equivalent fully observable one by introducing a new state variable, often referred to as the *filtering distribution*, which is the conditional distribution of X_t given the observations of the asset price $\{S_0, S_1, \dots, S_t\}$. Given a set A in the Borel σ -algebra on \mathbb{R} , define

$$\Pi_t(A) \triangleq P(X_t \in A | S_0, \dots, S_t, X_0) = \int_A \pi_t(x_t) dx_t, \quad t = 1, \dots, T.$$

Given a realization $\{s_0, s_1, \dots, s_t\}$, the filtering distribution evolves as follows:

$$\begin{aligned} \pi_t(x_t) &= p(X_t = x_t | S_0 = s_0, \dots, S_t = s_t, X_0 \sim \pi_0) \\ &= \frac{p(X_t = x_t, S_t = s_t | S_0 = s_0, \dots, S_{t-1} = s_{t-1}, X_0 \sim \pi_0)}{p(S_t = s_t | S_0 = s_0, \dots, S_{t-1} = s_{t-1}, X_0 \sim \pi_0)} \\ &= \frac{\int p(X_t = x_t, S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1}) \pi_{t-1}(x_{t-1}) dx_{t-1}}{\int p(S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1}) \pi_{t-1}(x_{t-1}) dx_{t-1}}, \quad t = 1, \dots, T, \end{aligned}$$

where the second equality follows Bayes' rule and the third equality follows from the fact $\{(X_t, S_t)\}$ is a bivariate Markov process. The conditional probability density functions $p(X_t = x_t, S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1})$ and $p(S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1})$ are induced by (4), (5), and the distributions of Z_t^1 and Z_t^2 . Noticing that π_t only depends on π_{t-1} , s_{t-1} , and s_t , and letting the realization $\{s_0, \dots, s_t\}$ be replaced by the random variables $\{S_0, \dots, S_t\}$, we can abstractly rewrite the filtering recursion as

$$\Pi_t = \Phi(\Pi_{t-1}, S_{t-1}, S_t), \quad t = 1, \dots, T.$$

Therefore, the partially observable process (X_t, S_t) is transformed to an equivalent fully observable process with state (Π_t, S_t) , which turns out to be Markovian with respect to \mathcal{F}_t^S . Theoretically, the option price V_0 can be solved following the dynamic programming recursion (with details in Appendix 8.1):

$$V_t(s_t, \pi_t) = \max(g(s_t), C_t(s_t, \pi_t)), \quad t = T, \dots, 1, \quad (7)$$

where

$$\begin{aligned} C_T(s_T, \pi_T) &\triangleq g(s_T), \\ C_t(s_t, \pi_t) &\triangleq \mathbb{E}[V_{t+1}(S_{t+1}, \Pi_{t+1}) | S_t = s_t, X_t \sim \pi_t], \quad t = T-1, \dots, 1, \end{aligned}$$

denote the *continuation value* at time t , which is the expectation of the option value conditional on the information up to time t .

Hence, the optimal stopping time is

$$\tau^* = \min\{t \in \mathcal{J} \mid g(S_t) \geq C_t(S_t, \Pi_t)\},$$

which can be derived from the following recursion:

$$\begin{aligned} \tau_T^* &= T, \\ \tau_t^* &= \tau_{t+1}^* \cdot \mathbf{1}_{\{C_t(S_t, \Pi_t) > g(S_t)\}} + t \cdot \mathbf{1}_{\{C_t(S_t, \Pi_t) \leq g(S_t)\}}, \quad t = T-1, \dots, 1, \end{aligned} \quad (8)$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. Hence, $\tau^* = \tau_1^*$. The above recursion also shows that (S_t, Π_t) are the sufficient statistics that determine the optimal exercise policy. With τ_t^* , it can be shown that

$$C_t(s_t, \pi_t) = \mathbb{E}[g(S_{\tau_{t+1}^*}) | S_t = s_t, X_t \sim \pi_t], \quad t = T-1, \dots, 1.$$

However, it is often impossible to solve the problem exactly following (7) or (8) due to two main difficulties. One is that in general the filtering distribution π_t is infinite dimensional and the filtering recursion cannot be computed exactly, and the other difficulty lies in the accurate estimation of the continuation value $C_t(s_t, \pi_t)$. So we introduce an approximate method in the next two sections.

4 LOWER BOUND: REGRESSION-BASED METHOD

In this section, we follow the idea in (Longstaff and Schwartz 2001) to present a regression-based algorithm for obtaining a suboptimal policy that is expected to yield a tight lower bound on the true option price. The main idea is to iteratively estimate the continuation value using regression and update the optimal stopping time on a number of independent sample paths.

Let $S = \{S_1, \dots, S_T\}$. Given a sequence of approximate continuation functions $\{\tilde{C}_t(S_t), t = 1, \dots, T\}$ with $\tilde{C}_T(S_T) \equiv 0$, a suboptimal exercise policy τ is determined by

$$\tau(S) = \min\{t \in \mathcal{J} \mid g(S_t) \geq \tilde{C}_t(S_t)\}, \quad (9)$$

which is indeed an \mathcal{F}_t^S -adapted stopping time. It can be computed recursively according to

$$\begin{aligned}\tau_T(S) &= T, \\ \tau_t(S) &= \tau_{t+1} \cdot \mathbf{1}_{\{\tilde{C}_t(S_t) > g(S_t)\}} + t \cdot \mathbf{1}_{\{\tilde{C}_t(S_t) \leq g(S_t)\}}, \quad t = T-1, \dots, 1.\end{aligned}\tag{10}$$

Applying the policy τ on a sample path of the asset price $s \triangleq \{s_1, \dots, s_T\}$, the resultant payoff is $g(s_{\tau(s)})$. Suppose $s^{(i)} = \{s_1^{(i)}, \dots, s_T^{(i)}\}$, $i = 1, \dots, N$ are independent sample paths generated from the dynamics (4) and (5) with initial condition $S_0^{(i)} = s_0$ and $X_0^{(i)} \sim \pi_0$. Then we can take the average of the payoffs over all the sample paths to obtain an estimator for the option price.

$$L_N^\tau = \frac{1}{N} \sum_{i=1}^N g(s_{\tau}^{(i)}),$$

where $g(s_{\tau}^{(i)})$ is short for $g(s_{\tau(s^{(i)})}^{(i)})$. The following theorem shows that L_N^τ is an asymptotically (as $N \rightarrow \infty$) lower bound on the option price.

Theorem 1 For any τ of the form (9), we have the following inequality almost surely,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N g(s_{\tau}^{(i)}) \leq V_0(s_0, \pi_0).$$

Proof. According to the strong law of large numbers, it holds almost surely

$$\begin{aligned}\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N g(s_{\tau}^{(i)}) &= E[g(S_\tau) | S_0 = s_0, X_0 \sim \pi_0] \\ &\leq E[g(S_{\tau^*}) | S_0 = s_0, X_0 \sim \pi_0] \\ &= V_0(s_0, \pi_0),\end{aligned}$$

where the inequality holds as τ is a suboptimal exercise policy. □

The lower bound gets closer to the true option price if \tilde{C}_t is a better approximation of the true continuation value C_t . We employ the regression idea in (Longstaff and Schwartz 2001) to approximate $C_t(s_t, \pi_t)$. More specifically, we expect a good approximation using some function \hat{C}_t in \mathcal{H} , where $\mathcal{H} = \text{span}\{h_j, j = 1, 2, \dots, p\}$ is a p -dimensional function space with the basis functions h_j 's. Given a suboptimal exercise policy τ_{t+1} at time $t+1$, define

$$\hat{C}_t(s_t) = \text{pr}_{\mathcal{H}} \circ \mathbb{E}[g(S_{\tau_{t+1}}) | S_t = s_t, X_t \sim \pi_t],$$

where $\text{pr}_{\mathcal{H}}$ is the projection operator onto \mathcal{H} . However, the right hand side often cannot be carried out exactly, so in practice it is approximately computed through empirical regression. The idea is to simulate N independent sample paths of the process S , denoted by $s^{(i)} = \{s_1^{(i)}, \dots, s_T^{(i)}\}$ for $i = 1, \dots, N$, and then approximate $\hat{C}_t(s_t)$ by the empirical least-square regression. To be precise,

$$\hat{C}_t(s_t) \approx \arg \min_{f \in \mathcal{H}} \frac{1}{N} \sum_{i=1}^N |f(s_t^{(i)}) - g(s_{\tau_{t+1}}^{(i)})|^2 = \sum_{j=1}^p \beta_{t,j} h_j(s_t) \triangleq \tilde{C}_t(s_t).\tag{11}$$

Hence, $\tilde{C}_t(s_t^{(i)}) = \sum_{j=1}^p \beta_{t,j} h_j(s_t^{(i)})$.

Therefore, the procedure of finding a suboptimal policy and its corresponding value is to carry out the updating of the stopping time and the estimation of the continuation value in an interleaving way.

Given τ_{t+1} , we compute the approximate continuation functions $\tilde{C}_t(S_t)$ according to (11), and evaluate the function values on each sample path to obtain the approximate continuation value $\tilde{C}_t(s_t^{(i)})$, based on which we carry out (10) to obtain τ_t . This is repeated for $t = T - 1, \dots, 1$, until we obtain τ_1 as a suboptimal exercise policy τ . The algorithm is stated below.

Algorithm 1. Lower Bound and Suboptimal Policy

- *Step 1: Simulate $s^{(i)} = \{s_1^{(i)}, \dots, s_T^{(i)}\}, i = 1, \dots, N$ independently from processes (4)-(5) with initial condition $X_0 \sim \pi_0$ and $S_0 = s_0$.*
- *Step 2: Set $\tau_T(s^{(i)}) = T, i = 1, \dots, N$.*
- *Step 3: For $t = T - 1, \dots, 1$, do*
 - Evaluate the basis function $h_j(s_t^{(i)})$ for each $i = 1, \dots, N$ and $j = 1, \dots, p$.*
 - Least square regression:*

$$\beta_t = (\beta_{t1}, \dots, \beta_{tp}) \triangleq \arg \min_{\beta_1, \dots, \beta_p} \frac{1}{N} \sum_{i=1}^N |g(s_{\tau_{t+1}}^{(i)}) - \sum_{j=1}^p \beta_j h_j(s_t^{(i)})|^2,$$

- *Evaluate the approximate continuation value for each path $i = 1, \dots, N$*

$$\tilde{C}_t(s_t^{(i)}) = \sum_{j=1}^p \beta_{tj} h_j(s_t^{(i)}).$$

- *Update the stopping time for each path $i = 1, \dots, N$*

$$\tau_t(s^{(i)}) = \begin{cases} t & \text{if } \tilde{C}_t(s_t^{(i)}) \leq g(s_t^{(i)}); \\ \tau_{t+1}(s^{(i)}) & \text{otherwise.} \end{cases}$$

end

- *Step 4: Set $L_N^\tau = \frac{1}{N} \sum_{i=1}^N g(s_\tau^{(i)})$, where $\tau = \tau_1(s^{(i)})$. L_N^τ is an asymptotically lower bound on $V_0(s_0, \pi_0)$.*

In our algorithm, the continuation value is estimated using the updated policy and the empirical regression. Since the decision in our context is either exercise or hold, the gap between the true continuation value C_t and the approximate continuation value \tilde{C}_t will not affect the decision unless it is large enough to flip the decision. Hence, we can expect our method to obtain a good policy if $\mathbb{P}(\{C_t(S_t, \Pi_t) > g(S_t)\} \cap \{\tilde{C}_t(S_t) \leq g(S_t)\})$ is small for each t .

Our method is similar to those proposed in (Rambharat and Brockwell 2010) and (Ludkovski 2009). The main difference is that they assume the basis function h_j used in regression is a function of both s_t and an approximate filtering distribution. Since the approximate filtering distribution has its own randomness and cannot be determined purely by s_t , it enlarges the filtration \mathcal{F}_t^S so that the resultant policy is not \mathcal{F}_t^S -adapted. As a result, their estimates of the option price are not guaranteed to be an asymptotically lower bound. However, we need a lower bound estimate to be used in pair with an upper bound estimate.

5 APPROXIMATE UPPER BOUND: FILTERING-BASED DUALITY APPROACH

In this section, we construct the dual problem to the original optimal stopping problem, and combine with the particle filtering technique to find an approximate upper bound of (6). We first state Theorem 2 below, which is the theoretical basis for our algorithm. It was first used by (Rogers 2002), (Haugh and Kogan 2004) and (Andersen and Broadie 2004) to determine the American option price under constant volatility.

Theorem 2 (c.f. (5) in (Andersen and Broadie 2004)) Let \mathcal{M} represent the space of \mathcal{F}_t^S -adapted martingale M_t with $M_0 = 0$ and $\sup_{t \in \mathcal{J}} |M_t| < \infty$. Then

$$V_0(s_0, \pi_0) = \inf_{M \in \mathcal{M}} \left\{ \mathbb{E}[\max_{t \in \mathcal{J}} (g(S_t) - M_t) | S_0 = s_0, X_0 \sim \pi_0] \right\}.$$

This theorem suggests that any \mathcal{F}_t^S -adapted martingale M_t will lead to an upper bound on $V_0(s_0, \pi_0)$. Therefore, the natural question is how we can find a proper M_t to compute a tight upper bound in practice. We define such a martingale based on a suboptimal stopping time τ , which may be generated by Algorithm 1, i.e., $\tau = \min\{t \in \mathcal{J} \mid g(S_t) \geq \tilde{C}_t(S_t)\}$. Define the t-indexed stopping time τ_t for each $t \in \mathcal{J}$:

$$\tau_t = \min \{i \in \mathcal{J}_t : g(S_i) \geq \tilde{C}_i(S_i)\}$$

with $\mathcal{J}_t = \{t, t+1, \dots, T\}$. Note that τ_t here coincides with the one we derived from the backward recursion (10). Based on τ_t , we define the martingale difference

$$\begin{aligned} \Delta_t &= \mathbb{E}[g(S_{\tau_t}) | \mathcal{F}_t^S] - \mathbb{E}[g(S_{\tau_t}) | \mathcal{F}_{t-1}^S] \\ &= \mathbb{E}[g(S_{\tau_t}) | S_t, X_t \sim \Pi_t] - \mathbb{E}[g(S_{\tau_t}) | S_{t-1}, X_{t-1} \sim \Pi_{t-1}], \end{aligned} \quad (12)$$

where the second line follows from the fact that (S_t, Π_t) is Markovian with respect to \mathcal{F}_t^S . It can be seen that $\mathbb{E}[\Delta_t | \mathcal{F}_{t-1}^S] = 0$. Hence, the process $\{M_t\}$ defined as

$$M_0 = 0, \quad M_t = \Delta_1 + \dots + \Delta_t, \forall t \in \mathcal{J},$$

is an \mathcal{F}_t^S -adapted martingale.

5.1 Approximate Martingale Difference

The remaining issue is how to compute the martingale difference (12). First notice that given a sequence of observed asset price $\{s_1, \dots, s_T\}$, the first term on the righthand side of (12) can be rewritten as

$$\mathbb{E}[g(S_{\tau_t}) | S_t = s_t, X_t \sim \pi_t] = \begin{cases} g(s_t), & \text{if } g(s_t) \geq \tilde{C}_t(s_t); \\ \mathbb{E}[g(S_{\tau_{t+1}}) | S_t = s_t, X_t \sim \pi_t], & \text{if } g(s_t) < \tilde{C}_t(s_t). \end{cases} \quad (13)$$

Hence, the only term we need to estimate is $\mathbb{E}[g(S_{\tau_{t+1}}) | S_t = s_t, X_t \sim \pi_t]$ for $t = 0, \dots, T-1$. However, an unbiased estimator requires i.i.d. samples from the intractable filtering distribution π_t . One approach to partially solve this problem is to apply the particle filtering method, which approximates π_t by a finite number of samples $\{x_t^{(1)}, \dots, x_t^{(m)}\}$, i.e., a discrete distribution $\hat{\pi}_t$ written as follows

$$\hat{\pi}_t = \frac{1}{m} \sum_{i=1}^m \delta_{x_t^{(i)}}, \quad (14)$$

where δ is the Dirac measure. As the number of samples m goes to infinity, particle filtering ensures that $\hat{\pi}_t$ converges to π_t in certain sense. A good tutorial on particle filtering can be found in (Arulampalam, Maskell, Gordon, and Clapp 2002), (Doucet and Johansen 2009), and in the book (Doucet, de Freitas, and Gordon 2001). Here we adopt the basic version of particle filtering, which is stated in Algorithm 2 below.

Algorithm 2. Particle Filtering

Input: $X_0 \sim \pi_0$ and a sequence of observed asset price $\{s_1, \dots, s_T\}$.

Output: The approximate filtering distribution $\hat{\pi}_0, \dots, \hat{\pi}_T$.

- *Step 1. Initialization:* Set $t = 0$. Draw m i.i.d. sample $x_0^{(1)} \dots, x_0^{(m)}$ from the distribution π_0 . Set $\hat{\pi}_0 = \frac{1}{m} \sum_{i=1}^m \delta_{x_0^{(i)}}$.
- *Step 2. For* $t = 1, \dots, T$, **do**
 - *Prediction:* Sample $\tilde{x}_t^{(i)}$ from the state equation (4) by generating standard normal distributed random variable $z_t^{(i)}$ with $X_{t-1} = x_{t-1}^{(i)}$, $i = 1, \dots, m$.
 - *Bayes' Updating:* Compute $w_t^{(i)} = \frac{p(S_t = s_t | S_{t-1} = s_{t-1}, X_t = \tilde{x}_t^{(i)})}{\sum_{i=1}^m p(S_t = s_t | S_{t-1} = s_{t-1}, X_t = \tilde{x}_t^{(i)})}$ according to (5), $i = 1, \dots, m$.
 - *Resampling:* Draw i.i.d. samples $\{x_t^{(1)}, \dots, x_t^{(m)}\}$ from the discrete distribution $\text{Prob}(\tilde{x}_t^{(i)}) = w_t^{(i)}$, $i = 1, \dots, m$. Set $\hat{\pi}_t = \sum_{i=1}^m \delta_{x_t^{(i)}}$. **end**

So instead of estimating $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t \sim \pi_t]$ directly, we estimate $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t \sim \hat{\pi}_t]$ as its approximation. Since $\hat{\pi}_t$ is of the form (14),

$$\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t \sim \hat{\pi}_t] = \sum_{i=1}^m \frac{1}{m} \mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]. \quad (15)$$

For each i and t , $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]$ can be estimated with noise using Algorithm 3.

Algorithm 3. Estimation of $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]$

- *Step 1:* Simulate $\{s_{t+1}^{(i)}, \dots, s_T^{(i)}\}$ from the processes (4)-(5) starting at time t with initial condition $X_t = x_t^{(i)}$, $S_t = s_t$.
- *Step 2:* To apply τ_{t+1} on this sample path, find

$$t_i = \min \left\{ k \in \mathcal{J}_{t+1} : g(s_k^{(i)}) \geq \tilde{C}_k(s_k^{(i)}) \right\}.$$

Then $g(s_{t_i}^{(i)})$ is an unbiased estimator of $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]$.

Notice that in Algorithm 3 only one subpath is simulated to estimate $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]$. It is of no doubt that more subpaths can be generated to reduce the variance of the estimator. Since we can estimate $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t = x_t^{(i)}]$ for each i by Algorithm 3, (15) can be estimated with noise by taking the average of these terms, that is, $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t \sim \hat{\pi}_t] \approx \frac{1}{m} \sum_{i=1}^m g(s_{t_i}^{(i)})$. As m becomes greater, we have less bias in estimating $\mathbb{E}[g(S_{\tau_{t+1}})|S_t = s_t, X_t \sim \pi_t]$.

5.2 Approximate Upper Bound

So far we have derived a way to compute the approximate martingale difference $\hat{\Delta}_t$ for a given stopping time τ and a sample path $\{s_1, \dots, s_T\}$. Based on this, we propose the following algorithm that yields an approximate upper bound on the option price.

Algorithm 4. A Filtering-Based Duality Approach

- *Step 1.* Generate N_1 independent paths of the asset price $\{s_1^{(k)}, \dots, s_T^{(k)}\}$ according to the processes (4)-(5), and then follow Algorithm 2 to generate the approximate filtering distribution $\{\hat{\pi}_1^{(k)}, \dots, \hat{\pi}_T^{(k)}\}$ for $k = 1, \dots, N_1$.
- *Step 2. For* $k = 1, 2, \dots, N_1$, **do**
 - *For* $t = T, \dots, 1$ **compute**

$$\hat{\Delta}_t^{(k)} = \mathbb{E}[g(S_{\tau_t})|S_t = s_t^{(k)}, X_t \sim \hat{\pi}_t^{(k)}] - \mathbb{E}[g(S_{\tau_t})|S_t = s_{t-1}^{(k)}, X_t \sim \hat{\pi}_{t-1}^{(k)}]$$

according to (13) with π_t replaced by $\hat{\pi}_t$ and Algorithm 3.

- *Sum the approximate martingale differences to get*

$$\hat{M}_t^{(k)} = \hat{\Delta}_1^{(k)} + \dots + \hat{\Delta}_t^{(k)}, \quad t = 1, \dots, T.$$

- *Evaluate* $U^{(k)} = \max_{t \in \mathcal{J}} \left(g(S_t^{(k)}) - \hat{M}_t^{(k)} \right)$. **end**

- *Step 3.* Set $U_{N_1}^\tau = \frac{1}{N_1} \sum_{k=1}^{N_1} U^{(k)}$. $U_{N_1}^\tau$ is an approximate upper bound on the option price $V_0(s_0, \pi_0)$.

6 NUMERICAL EXAMPLES

In this section, we consider pricing an American put option with $g(S_t) = \max(e^{-r\Delta t}(K - S_t), 0)$. We use the parameters as follows:

1. Volatility parameter: $\lambda = 1$, $\theta = \log(0.2)$, $\gamma = 1$, $\rho = 0$;
2. Asset price parameter: $r = 0.05$, $K = 40$;
3. Time parameter: $\delta = 0.1$, and $T = 5, 10, 15$;
4. Initial condition: $S_0 = 36$, $x_0 = \log(0.2)$.

Besides the parameters for (4)-(5), we also need to assign the basis functions in the regression step of Algorithm 1. In our simulation, we use the following set of basis functions

$$H_1 = \{L_0(S_t), L_0^2(S_t), L_1(S_t), L_1^2(S_t), L_0(S_t)L_1(S_t), 1\},$$

where $L_0(x) = x$, $L_1(x) = \min(K - x, 0)$.

We present our numerical results in Table 1, where each entry shows the sample average and standard error (in parentheses) of 20 independent runs. The lower bound L on the option price is the average of 20 independent outputs, each of which is obtained by implementing Algorithm 1 with the number of sample paths $N = 40000$ to obtain a suboptimal policy τ , and then applying this policy on another independent set of 40000 paths. The dual upper bound U is the average of 20 independent outputs, each of which is obtained by implementing Algorithm 4 using one suboptimal policy τ with the number of sample paths $N_1 = 500$ and the number of particles $m = 1000$. To compare with the option price under full observation of stochastic volatility, we also implement Algorithm 1 in a similar way to obtain a lower bound $Full.LB$, with 40000 sample paths and a larger set of basis functions:

$$H_2 = H_1 \cup \{L_0(e^{X_t}), L_0(e^{X_t})L_1(S_t)\}.$$

Please note that we have incorporated the information of X_t in the new set H_2 . Then we implement the primal-dual method of (Andersen and Broadie 2004) with 500 sample paths and 1000 subpaths to obtain the dual upper bound $Full.UB$ on the option price of the fully observable model.

Table 1: American Put Option Values

T	L	U	\hat{V}	$Full.LB.$	$Full.UB.$	$Full.Obs.$	$Overprice$
5	4.386(0.002)	4.435(0.002)	4.411	4.470(0.002)	4.481(0.002)	4.475	0.064
10	4.897(0.003)	5.040(0.002)	4.969	5.101(0.002)	5.122(0.002)	5.112	0.143
15	5.342(0.003)	5.557(0.003)	5.450	5.612(0.004)	5.649(0.003)	5.631	0.181

We observe that in our numerical results the differences of the approximate upper bound U and the lower bound L are small, so they can be both regarded as approximate option prices. One possible choice of a point estimator for the true option price V_0 is the average of the two

$$\hat{V} = \frac{L+U}{2}.$$

The column $Overprice$ in Table 1 shows the difference between \hat{V} and $Full.Obs.$, which is the average of $Full.LB$ and $Full.UB$. The optimal solution to a fully observable problem is an upper bound on the optimal solution to the corresponding partially observable problem. However, the gap between the two solutions is usually unknown. So another advantage of introducing the dual upper bound U is that it provides a criterion to evaluate the performance of L : the smaller the gap between U and L , the better the bounds.

7 CONCLUSION

In this paper, we investigate the American option pricing problem under partial observation of stochastic volatility, for which we develop an efficient numerical method to compute tight lower bound and approximate upper bound on the true option price. The lower bound is obtained by a dynamic programming approach with regression on the continuation values, and the dual upper bound is based on a martingale duality

formulation on the joint process of the asset price and the filtering distribution. We apply this approach to price an American put option, and the numerical results show that our lower and upper bounds are close enough to be regarded as good approximation of the true option price. The numerical results also show that our method reduces overpricing of the option by treating the volatility more realistically as partially observable.

8 APPENDIX

8.1 Equivalent Formulation of the DP Problem

We first show that (S_t, Π_t) is Markovian with respect to \mathcal{F}_t^S . Given a realization of asset price $\{s_0, \dots, s_{t-1}\}$ and the initial distribution of X_0 ,

$$p(S_t = s_t | S_0 = s_0, \dots, S_{t-1} = s_{t-1}, X_0 \sim \pi_0) = \int p(S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1}) \pi_{t-1}(x_{t-1}) dx_{t-1},$$

where $p(S_t = s_t | X_{t-1} = x_{t-1}, S_{t-1} = s_{t-1})$ can be obtained from (4) and (5). Therefore, the probability distribution of S_t conditional on \mathcal{F}_{t-1}^S is determined by S_{t-1} and Π_{t-1} . Moreover, we have shown that $\Pi_t = \Phi(\Pi_{t-1}, S_{t-1}, S_t)$. Hence, conditional on \mathcal{F}_{t-1}^S , the distribution of (S_t, Π_t) is determined by (S_{t-1}, Π_{t-1}) , implying that (S_t, Π_t) is a \mathcal{F}_t^S -adapted Markov process.

The original dynamic programming problem is

$$\begin{aligned} \tilde{V}_T(s_0, \dots, s_T) &= g(s_T), \\ \tilde{V}_t(s_0, \dots, s_t) &= \max(g(s_t), \mathbb{E}[\tilde{V}_{t+1}(S_0, \dots, S_{t+1}) | X_0 \sim \pi_0, S_0 = s_0, \dots, S_t = s_t]), \end{aligned}$$

for $t = T-1, \dots, 1$. Next we show that the value function at time t , \tilde{V}_t , can be represented as a function of only π_t and s_t . To show this, first define $V_T(s_T, \pi_T) \triangleq \tilde{V}_T(s_0, \dots, s_T) = g(s_T)$. By induction, for $t = T-1, \dots, 1$

$$\begin{aligned} &\tilde{V}_t(s_0, \dots, s_t) \\ &= \max(g(s_t), \mathbb{E}[\tilde{V}_{t+1}(S_0, \dots, S_{t+1}) | X_0 \sim \pi_0, S_0 = s_0, \dots, S_t = s_t]) \\ &= \max(g(s_t), \mathbb{E}[V_{t+1}(S_{t+1}, \Pi_{t+1}) | X_0 \sim \pi_0, S_0 = s_0, \dots, S_t = s_t]) \\ &= \max\left(g(s_t), \int \mathbb{E}[V_{t+1}(S_{t+1}, \Pi_{t+1}) | X_t \sim \pi_t, S_t = s_t]\right), \end{aligned}$$

where the second equality holds due to the induction and the third equality holds as (S_t, Π_t) is a bivariate Markov process with respect to \mathcal{F}_t^S . Hence, we can define

$$V_t(s_t, \pi_t) \triangleq \max(g(s_t), \mathbb{E}[V_{t+1}(S_{t+1}, \Pi_{t+1}) | S_t = s_t, X_t \sim \pi_t]).$$

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under Grants ECCS-0901543 and CMMI-1130273.

REFERENCES

- Andersen, L., and M. Broadie. 2004. "Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options". *Management Science* 50 (9): 1222 – 1234.
- Arulampalam, S., S. Maskell, N. J. Gordon, and T. Clapp. 2002. "A Tutorial on Particle Filters for Online Non-linear/Non-Gaussian Bayesian Tracking". *IEEE Transactions on Signal Processing* 50 (2): 174–188.
- Belomestny, D., C. Bender, and J. Schoenmakers. 2009. "True Upper Bounds for Bermudan Products via Non-Nested Monte Carlo". *Mathematical Finance* 19:53 – 71.

- Broadie, M., and P. Glasserman. 1997. "Pricing American-Style Securities Using Simulation". *Journal of Economic Dynamics and Control* 21 (8-9): 1323 – 1352.
- Broadie, M., and P. Glasserman. 2004. "A Stochastic Mesh Method for Pricing High-Dimensional American Options". *Journal of Computational Finance* 7 (2): 59 – 79.
- Brown, D., J. Smith, and P. Sun. 2010. "Information Relaxations and Duality in Stochastic Dynamic Programs". *Operations Research* 58 (4): 758 – 801.
- Carriere, J. 1996. "Valuation of Early-Exercise Price of Options Using Simulations and Nonparametric Regression". *Insurance: Mathematics and Economics* 19 (1): 19 – 30.
- Chen, N., and P. Glasserman. 2007. "Additive and Multiplicative Duals for American Option Pricing". *Finance and Stochastics* 11:153 – 179.
- Doucet, A., J. F. G. de Freitas, and N. J. Gordon. (Eds.) 2001. *Sequential Monte Carlo Methods In Practice*. New York: Springer.
- Doucet, A., and A. M. Johansen. 2009. "A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later". In *Handbook of Nonlinear Filtering*, edited by D. Crisan and B. Rozovsky. Cambridge: Cambridge University Press.
- Florescu, I., and F. Viens. 2008. "Stochastic Volatility: Option Pricing Using a Multinomial Recombining Tree". *Applied Mathematical Finance* 15 (2): 151 – 181.
- Garcia, D. 2003. "Convergence and Biases of Monte Carlo Estimates of American Option Prices Using a Parametric Exercise Rule". *Journal of Economic Dynamics and Control* 27 (10): 1855 – 1879.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer.
- Glasserman, P., and B. Yu. 2002. "Pricing American Options by Simulation: Regression Now or Regression Later?". *Monte Carlo and Quasi-Monte Carlo Methods* (H. Niederreiter, ed.): 213 – 226.
- Haugh, M. B., and L. Kogan. 2004. "Pricing American Options: A Duality Approach". *Operations Research* 52 (2): 258 – 270.
- Ibanez, A., and F. Zapatero. 2004. "Monte Carlo Valuation of American Options Through Computation of the Optimal Exercise Frontier". *Journal of Financial and Quantitative Analysis* 39 (2): 253 – 275.
- Jamshidian, F. 2007. "The Duality of Optimal Exercise and Domineering Claims: A Doob-Meyer Decomposition Approach to the Snell Envelope". *Stochastics* 79:27 – 60.
- Longstaff, F. A., and E. S. Schwartz. 2001. "Valuing American Options by Simulation: A Simple Least-Squares Approach". *The Review of Financial Studies* 14 (1): 113 – 147.
- Ludkovski, M. 2009. "A Simulation Approach to Optimal Stopping under Partial Information". *Stochastic Processes and Applications* 119 (12): 2071 – 2087.
- Pham, H., W. Runggaldier, and A. Sellami. 2005. "Approximation by Quantization of the Filter Process and Applications to Optimal Stopping Problems under Partial Observation". *Monte Carlo Methods and Applications* 11 (1): 57 – 81.
- Rambharat, B. R., and A. E. Brockwell. 2010. "Sequential Monte Carlo Pricing of American-Style Options under Stochastic Volatility Models". *The Annals of Applied Statistics* 4, No. 1, 222-265 (1): 222 – 265.
- Raymar, S. B., and M. J. Zwecher. 1997. "A Monte Carlo Valuation of American Call Options On the Maximum of Several Stocks". *The Journal of Derivatives* 5 (1): 7 – 23.
- Rogers, L. C. G. 2002. "Monte Carlo Valuation of American Options". *Mathematical Finance* 12 (3): 271 – 286.
- Tilley, J. A. 1993. "Valuing American Options in a Path Simulation Model". *Transactions of the Society of Actuaries* 45:83 – 104.
- Tsitsiklis, J., and B. V. Roy. 2001. "Regression Methods for Pricing Complex American-Style Options". *IEEE Transactions on Neural Networks* 12 (4): 694 – 703.

AUTHOR BIOGRAPHIES

FAN YE is a Ph.D student in the Department of Industrial and Enterprise Systems Engineering, the University of Illinois at Urbana-Champaign. His research interest is 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 electri-

Ye and Zhou

cal 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.