Optimal Stopping of a Partially Observed Long-Memory Process with Applications in Finance

Rendani Emmanuel Tshisikule (tshisikule@aims.ac.za)

African Institute for Mathematical Sciences (AIMS)

Supervised by: Dr. Dennis Ikpe
University of South Africa (UNISA), South Africa

Co-supervised by: Mr Aidoo Edward
Barclays Africa Group, South Africa

26 October 2017

Submitted in partial fulfillment of a structured masters degree at AIMS South Africa

Abstract

In this project, we use a combination of observable real market data and simulation approach to solve an optimal stopping problem under partial observation of a stochastic volatility model. In particular, we consider a long-memory process driven by a fractional Brownian motion with Hurst parameter $H \in (0.5, 1)$. Our result provides the shortest distance between the lower and upper bounds on our value function. We demonstrate an application in mathematical finance by pricing an American style option on an underlying with both stochastic volatility and convenience yield through a simulation method. Our numerical results show that the upper bound on the option price increases with the number of sample paths in our pricing algorithm and decreases as the option maturity increases.

Keywords: optimal stopping, partial observation, particle filtering, long-memory, martingale duality, Fractional Brownian motion.

Declaration

I, the undersigned, hereby declare that the work contained in this research project is my original work, and that any work done by others or by myself previously has been acknowledged and referenced accordingly.

Rendani Emmanuel Tshisikule, 26 October 2017

Contents

1. Introduction

In this project, we solve the optimal stopping problem of a partially observed fractional Brownian motion and demonstrate applications in mathematical finance. Optimal stopping is a class of optimal control problems in stochastic analysis and often used in decision making on when is the best moment to enter or exit a given process. Optimal stopping problems under partial observation of stochastic processes are of interest to many researchers in mathematics with application ranging from optimal investment under partial information (Décamps et al., 2005), selling of financial derivatives (Vannestål, 2011), quality control, reliability (Jensen and Hsu, 1993) to pricing of options and optimal stock selling (Rishel and Helmes, 2006). The theory of optimal stopping under partial observation has been in existence since 1970's when Myron Scholes and Fischers Black introduced the formula for estimating stock option prices (Hill, 2009). Since then, the study of optimal stopping problem has gained wide popularity in the mathematical finance community.

Solving optimal stopping problems under partial observation is typically very challenging since the inference of the unobserved state and the stopping decision should be done at the same time. This difficulty is further compounded in the non-markovian setting as the interaction between learning and optimization is more complicated in processes with memory (Leão et al., 2017; Chronopoulou and Spiliopoulos, 2017). Several attempts to addressing these difficulties can be found in the literature. In the Markovian setting, Mazziotto (1986) used the assumed density filter in the filtering step and for the optimisation step, used PDE solver. Pham et al. (2005) proposed to use optimal quantization in filtering step and dynamic programming principle in obtaining optimal stopping time. In Olusaya and Ikpe (2017), the Kalman-Bucy Filter is used for filtering step and then dynamic programming for maximization step. Ludkovski (2009) considered interacting particle system to estimate the distribution of the unobservable process given the filtration of the observable and then simulation-based method which depends on Snell-envelope technique to solve the optimization problem. In both the Markovian and non-Markovian cases, a solution to an optimal stopping problem under partial observation involves a transformation to an equivalent fully observable optimal stopping problem by introducing a new state variable (filtering distribution). However, this transformation does not reduce the complexity of the problem, since the state variable is usually infinite dimensional (Ye and Zhou, 2013). In addition, the "curse of dimensionality" problem affects the accurate implementation of the dynamic programming principle in the optimization step.

The main purpose of this study is to determine the optimal stopping policies for partially observable stochastic volatility models in mathematical finance. In particular, we applied our method in pricing an American option when the underlying asset is described by a stochastic volatility model with long memory. Using real market stock price time series, we adapt an interacting particle filtering algorithm (Genetic-type algorithm) in Chronopoulou and Viens (2012a), to estimate the empirical distribution of the unobservable stock volatility while optimising using a filtering-based martingale duality approach. The filtering-based martingale duality approach is used to handle the computational complexity due to the infinite dimensionality of our filtering distribution. Pricing American option is one of the most difficult computational problems in derivatives' market, particularly when there are many factors affecting the value of the option (Longstaff and Schwartz, 2001). Our method has many advantages as a framework for valuing American options under stochastic volatility. Precisely, it allows for non-Markovian stochastic processes. Also, our simulation method is compatible with parallel computing, which allows for a significant gain in computational speed and efficiency. Our focus is on developing lower and upper bounds on the price (i.e., the optimal value of the stopping problem) with reasonable computational cost.

The rest of this project is structured as follows. In Chapter 2, we describe the optimal stopping of stochastic processes, the transformation from partial observation to an equivalent fully observable optimal stopping problem. We also discuss the main properties of our model. Chapter 3 is devoted to the description of the filtering-based duality approach, a combination of genetic-type particle filtering and approximate martingale duality method. We also present error analysis and convergence of our numerical method. Finally we demonstrate an application and conclude in Chapter 4.

2. Problem Formulation

2.1 Stochastic Volatility Models with Long-Memory

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space satisfying the usual conditions, hosting the processes $\{X_t\}$ and $\{Y_t\}$ in continuous time with the filtration $\mathbb{F} := \{\mathcal{F}_t\}_{t\geq 0}$. The two processes satisfy the following dynamics:

$$dX_t = (r - \frac{\sigma^2(Y_t)}{2})dt + \sigma(Y_t)dW_t, \tag{2.1.1}$$

$$dY_t = \alpha Y_t dt + \beta dB_t^H, \tag{2.1.2}$$

where σ is a deterministic function, α is the rate of mean reversion , β is the volatility constant, r is the short-term risk free rate of interest, $\{B_t^H; t \geq 0\}$ is a standard fractional Brownian motion (fBm) with $H \in (0.5, 1)$ and $\{W_t; t \geq 0\}$ is a standard Weiner process (Brownian motion). We call X the observable process in a continuous space $\mathcal{X} \subseteq \mathbb{R}^{d_x}$ and Y the unobservable process in a continuous state space $\mathcal{Y} \subseteq \mathbb{R}^{d_y}$. We assume that $\mathbb{P}$ is a martingale measure and that B_t^H and W_t are independent of each other. Depending on the function $\sigma(\cdot)$, we have a variety of stochastic volatility models. We shall denote the filtration generated by X_t up to time t by $\{\mathcal{F}_t^X := \sigma(X_s), s \leq t\}$ and the filtration generated by Y as $\mathcal{F}_t^Y$ respectively. In a similar manner, we take $\mathcal{T}^X$ to be the set of $\mathcal{F}_T^X$ stopping times.

Unlike other stochastic volatility models in the literature, the model we use is more general in that it extends to the non-Markovian setting as it makes provision for long-range dependency in the unobservable volatility process, Y. We also assume that the initial unobservable process Y_0 follows a distribution say π_0 which is known and formulated from past data including observable process X. With the above setup, we then consider the finite-horizon partially observable optimal stopping problem given by:

$$S(X_0, Y_0) = \sup_{\tau \in \mathcal{T}^X} \mathbb{E}[h(\tau, X_\tau, Y_\tau) \mid X_0 = x, Y_0 \sim \pi_0], \tag{2.1.3}$$

where h is the pay-off function depending on the stopping time τ, the unobserved process Y and observed process X. The right hand side of equation (2.1.3) is usually difficult to compute because we can't observe Y. This suggests that the optimal stopping time τ^* is not adapted to $\mathcal{T}^X$. Many researchers have shown that we can overcome this difficulty by utilizing the availability of some noisy observation $\{X_t\}_{0\leq t\leq T}$ that contains some partial information about the unobserved process Y such that the optimal stopping time τ^* is adapted to the filtration $\{\mathcal{F}_t^X\}$. Thus, we have an optimal stopping problem under partial information. Optimal stopping problems of partially observable processes occur very often in financial mathematics and applied probability, where the decision maker is not fully aware of the situation or process at hand. The most interesting features of this type of problems is the interaction between learning and optimization, in which the observation process X plays an important role as a source of information needed about the underlying unobservable (system state) Y. In the sequel, the decision maker has to consider the trade-off between further learning of X in order to obtain a more satisfying inference of Y, concerning stopping early if the market is not favourable. In partial observation setting, the decision making has to wait due to the demand for learning as opposed to full observation, where no inference is required.

Continuous-state optimal stopping problems under partial observation are generally hard to solve analytically and also pose a great challenge to numerical solutions. Generally we can resolve this difficulty by utilizing a two step filtering and optimization approach. From equation (2.1.3), suppose that τ is also adapted to $\mathcal{F}_t^X$, then we would have the problem of optimal stopping under full information which can then be solved directly using dynamic programming principle or other methods but this is not the case. Thus, the main focus in this project is to achieve the following:

- Estimate the empirical distribution of the unobservable volatility process Y using interacting particle stochastic filtering algorithm (Chronopoulou and Viens, 2012a).
- Use filtering-based duality approach to approximate the optimal stopping time.
- Perform error analysis and convergence on the algorithms that we use.
- Consider an application in the valuation of an American option.

2.2 Fractional Brownian Motion

Before discussing further properties of our model, we present a formal definition of a fractional Brownian motion, which is the driving noise of this process, (for a more detailed exposition see Chronopoulou and Viens (2012a), Chronopoulou and Viens (2012b), Chronopoulou and Spiliopoulos (2017) and Pipiras and Taqqu (2017)).

2.2.1 Definition. The fractional Brownian motion (fBm) with Hurst parameter $H \in (0,1)$ is a centred Gaussian process $\{B_t^H; t \in \mathbb{R} \mid t \geq 0\}$ whose paths are continuous with probability 1 and whose distribution is defined by its covariance as follows:

$$cov(B_t^H, B_s^H) = \frac{1}{2}(\mid t \mid^{2H} + \mid s \mid^{2H} - \mid t-s \mid^{2H}), t, s \geq 0$$

The covariance of fBm implies that it has H-self-similar increments; for every $b > 0$ the process $\{B_{bt}^H; t \in \mathbb{R} \mid t \geq 0\}$ and $\{b^H B_t^H; t \in \mathbb{R} \mid t \geq 0\}$ have the same distribution. The mean square of the increments of fBm is computed as follows:

$$E(\mid B_t^H - B_s^H \mid^2) = \mid t-s \mid^{2H},$$

which clearly indicates stationary increments. When $H = 0.5$, the process is standard Brownian motion and Markovian. However, in contrast to standard Brownian motion, fBm is not semi-martingale nor a Markov process when $H \neq 0.5$.

Long-range dependence (long-memory) is one very important property of fBm for some values of parameter H. When $H \neq 0.5$, the increments of fBm over disjoint intervals are not independent; their correlation function is

$$\rho_H(m) = \frac{1}{2}\Big((m+1)^{2H} + (m-1)^{2H} - 2m^{2H}\Big), m \in \mathbb{N}$$

It is observed that when $H < 0.5$ then $\rho_H(m) < 0$ and the increments over disjoint intervals are negatively correlated. When $H > 0.5$ then $\rho_H(m) > 0$ and the increments over disjoint intervals are positively correlated. More specifically, when $H > 0.5$, the X_m exhibits long-rage dependence in the sense that $\sum_{m=1}^{\infty} \rho_H(m) = \infty$, which follows from the asymptotic

$$\rho_H(m) = H(2H-1)m^{2H-2} + O(m^{2H-2})$$

When $H < 0.5$, then $\sum_{m=1}^{\infty} \mid \rho_H(m) \mid < \infty$, this simply means that the process has short-memory. However, in this setting our focus is on long-range dependence (long-memory);

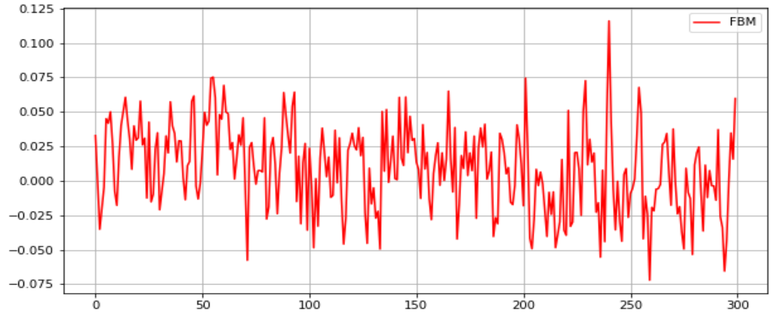

Figure 2.1: Fractional Brownian Motion with Hurst parameter $H = 0.7$.

2.2.1 Long-Memory Ornstein-Uhlenbeck Process. The fractional Ornstein-Uhlenbeck process is a continuous-time first-order autoregressive process $Y = \{Y_t; t \geq 0\}$ which is the solution of an one-dimensional homogeneous linear stochastic differential equation driven by a fractional Brownian motion $\{B_t^H; t \geq 0\}$ with Hurst parameter $H \in (\frac{1}{2}, 1)$. Specifically it is the unique Gaussian process satisfying the following linear stochastic integral equation;

$$Y_t = \alpha \int_0^t Y_s ds + \beta B_t^H, \quad t \geq 0, \tag{2.2.1}$$

where α and β are constant drift and variance parameters, respectively. The solution to equation (2.2.1) is stationary, almost surely continuous, and H-self-similar. The decay of the auto-covariance function of $\{Y_t; t \geq 0\}$ is similar to that of the increments of fractional Brownian motion, and thus its exhibits long-range dependence (Chronopoulou and Viens, 2012a). Below is an example of the graphical representation of fractional Ornstein–Uhlenbeck process.

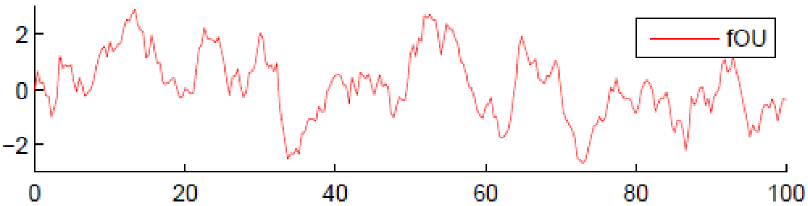

Figure 2.2: Fractional Ornstein–Uhlenbeck Process with Hurst parameter $H = 0.7$ from Kaarakka and Salminen (2007).

2.3 Equivalent Transform and Adapted Stopping Times

We can transform the optimal stopping problem in (2.1.3) with partial observation to an equivalent full observation counterpart. This can be achieved by introducing an additional state variable Π_t known as the filtering distribution. It is the conditional distribution of Y_t given the observations $\{X_0, \ldots, X_t\}$. In particular, given a set A in the Borel σ-algebra over $\mathcal{Y}$, we define $\Pi_t(A) := p(Y_t \in A \mid X_0, \dots, X_t), \quad t = 0, \dots, T$. Given a realisation of the observations $\{x_0, \dots, x_t\}$, the probability density of Π_t is correspondingly denoted as π_t and evolves as follows:

$$\pi_t(y_t) = \frac{\int_{\mathcal{Y}} p(x_t, y_t \mid x_{0:t-1}, y_{0:t-1}) \pi_{0:t-1}(y_{0:t-1}) d(y_{0:t-1})}{\int_{\mathcal{Y}} p(x_t \mid x_{0:t-1}, y_{0:t-1}) \pi_{0:t-1}(y_{0:t-1}) d(y_{0:t-1})}, \quad t \in (0, T], \tag{2.3.1}$$

where the conditional densities $p(x_t \mid x_{0:t-1}, y_{0:t-1})$ and $p(x_t, y_t \mid x_{0:t-1}, y_{0:t-1})$ are influenced by the model in (2.1.1) , (2.1.2) and the distributions of B_t^H and W_t. One can see that π_t in (2.3.1) depends on the entire history of the observations, $x_{0:t}$ and the past conditional densities before the current time $\pi_{0:t-1}$, and replacing realisation $\{x_0, \dots, x_t\}$, by the sequence of random variables $\{X_0, \dots, X_t\}$, then filtering recursion (2.3.1) can be abstractly rewritten as

$$\Pi_t = \vartheta(\Pi_{0:t-1}, X_{0:t-1}, X_t), \quad t \in (0, T].$$

Then our optimal stopping problem in (2.1.3) can be transformed to an equivalent problem with fully observable state (Π_t, X_t) and can be defined as

$$S(x_0, \pi_0) = \sup_{\tilde{\tau} \in \mathcal{T}^X} \mathbb{E}[\tilde{h}(\tilde{\tau}, X_{\tilde{\tau}}, \Pi_{\tilde{\tau}}) \mid X_0 = x_0, Y_0 \sim \pi_0], \tag{2.3.2}$$

where

$$\tilde{h}(t, X_t, \Pi_t) := \mathbb{E}[h(t, X_t, Y_t) \mid \mathcal{F}_t^X] = \int h(t, X_t, y_t) \Pi_t(y_t) dy_t.$$

In an economic view point, the value function, S represents the possible optimal pay-off that we can get under the given time horizon with the initial conditions $X_0 = x_0$ and $Y_0 \sim \pi_0$.

Theoretically, we can compute the optimal stopping rule based on a backward dynamic programming principle recursively:

$$S(t, X_t, \Pi_t) = \max(\tilde{h}(t, X_t, \Pi_t), C(t, X_t, \Pi_t)), \quad t = T, \dots, 0, \tag{2.3.3}$$

where $C(t, X_t, \Pi_t)$ is the continuation value at time t defined recursively as

$$\begin{aligned} C(T, X_T, \Pi_T) &:= \tilde{h}(T, X_T, \Pi_T) \ ; \\ C(t, X_t, \Pi_t) &:= \mathbb{E}[S(t+1, X_{t+1}, \Pi_{t+1}) \mid X_t, \Pi_t]. \end{aligned}$$

In this case $\mathbb{E}[\cdot \mid X_t, \Pi_t] = \mathbb{E}[\cdot \mid X_t, Y_t \sim \Pi_t]$. $S(0) = C(0)$ and the optimal stopping time is defined as

$$\tilde{\tau}^* := \min\{t \in \mathcal{K} \mid \tilde{h}(t, X_t, \Pi_t) \geq C(t, X_t, \Pi_t)\}$$

for $\tilde{\tau}^*$, we define its associated t-indexed $\tilde{\tau}_t^*$ for each t taking values in $\mathcal{K} = [0, \dots, T]$ as

$$\tilde{\tau}_t^* := \min\{i \in \mathcal{K}_t \mid \tilde{h}(i, X_i, \Pi_i) \geq C(i, X_i, \Pi_i)\} \tag{2.3.4}$$

with $\mathcal{K}_t := \{t, t+1, \dots, T\}$. The recursion presented above tells us that the state (X_t, Π_t) is adequate statistic that define the optimal stopping time. The Snell envelope process defined in (2.3.3) of the process $\tilde{h}(t, X_t, \Pi_t)$ is the smallest $\mathcal{F}^X$-supermartingale that dominates $\tilde{h}$ in the sense that

$S(t, X_t, \Pi_t) \geq \tilde{h}(t, X_t, \Pi_t)$. By moving the time index in the original optimal stopping problem, we can define the value function S as

$$S(x_0, \pi_0) = \sup_{\tau \in \mathcal{T}^X, t \leq \tau \leq T} \mathbb{E}[h(\tau, X_\tau, Y_\tau) \mid X_0 = x_0, Y_0 \sim \pi_0] \tag{2.3.5}$$

$$= \mathbb{E}[h(\tau_t^*, X_{\tau_t^*}, Y_{\tau_t^*}) \mid X_t = x_t, Y_t \sim \pi_t] \text{ for } t \in (0, T] \tag{2.3.6}$$

However, it is often very difficult if not impossible to compute the optimal stopping time, based on the recursive dynamic programming principle (2.3.3). This difficulty is two folds. The first is that in general, the filtering distribution Π_t is infinite dimensional hence the filtering recursion (2.3.1) can only be approximated. The other difficulty lies in the accurate estimation of the continuation value $C(t, X_t, \Pi_t)$. In this project, we propose to use a new approach in addressing these challenges. In particular, we adapt a genetic type algorithm with a martingale duality framework to approximate the upper bound on our value function.

3. Main Results

3.1 Filtering-Based Martingale Duality

In this chapter, we formulate a dual problem to our optimal stopping problem, following the standard approach in Rogers (2002) and Haugh and Kogan (2004). We then use a numerical method that gives an asymptotic upper bound on the value function S. The following definition is useful in our construction.

3.1.1 Definition. A random process $X = \{X_t\}$ is called **martingale** relative to a filtration $\{\mathcal{F}_t : t \geq 0\}$, if the following are satisfied;

(i) X_t is adapted to $\mathcal{F}_t$

(ii) $\mathbb{E}(| X_t |) < \infty \ \forall \ t$, and

(iii) $\mathbb{E}(X_{t+1} \mid \mathcal{F}_t) = X_t$, almost surely (LaLonde, 2013) .

Next, we present our duality result, an important theorem in the dual formulation and numerical implementation of the primal optimal stopping problem stated in Lamberton et al. (2009).

3.1.1 Theorem. *Let $\mathfrak{M}$ represent the space of $\mathcal{F}_t^X$-adapted martingale $\{M_t\}$ with $M_0 = 0$ and $\sup_{t\in\mathcal{K}} \mathbb{E}[| M_t |] < \infty$. Then*

$$S(x_0, \pi_0) = \min_{M\in\mathfrak{M}} \{\mathbb{E}[\max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi_t) - M_t\} \mid X_0 = x_0, Y_0 \sim \pi_0]\} \tag{3.1.1}$$

The optimal martingale is presented by $\{M_t^*\}$ which achieves the minimum on the right hand side of (3.1.1) is of the form;

$$M_t^* = \amalg_1^* + \cdots + \amalg_t^* \ \ t \in \mathcal{K}, \tag{3.1.2}$$

where $\amalg_t^*$ is the martingale difference sequence given by:

$$\amalg_t^* := \mathbb{E}[S(t) \mid \mathcal{F}_t^X] - \mathbb{E}[S(t) \mid \mathcal{F}_{t-1}^X], \ \ t \in \mathcal{K}. \tag{3.1.3}$$

Finally the following equality hold path-wisely:

$$S(x_0, \pi_0) = \max(\tilde{h}(t, X_t, \Pi_t) - M_t^*) \ \text{ almost surely.}$$

Proof. We prove the above duality result by following the approach in Lamberton et al. (2009) while extending the arguments to a non-Markovian setting. Given $M \in \mathfrak{M}$, we have for any stopping time $\tilde{\tau} \in \mathcal{T}^X$,

$$\begin{aligned} \mathbb{E}[\tilde{h}(\tilde{\tau}, X_{\tilde{\tau}}, \Pi_{\tilde{\tau}})] =& \mathbb{E}[\tilde{h}(\tilde{\tau}, X_{\tilde{\tau}}, \Pi_{\tilde{\tau}}) - M_{\tilde{\tau}}] \\ \leq& \mathbb{E}[\max_{0\leq s\leq T} \{\tilde{h}(s, X_s, \Pi_s) - M_s\}], \end{aligned}$$

therefore

$$\sup_{\tilde{\tau}\in\mathcal{T}^X} \mathbb{E}[\tilde{h}(\tilde{\tau}, X_{\tilde{\tau}}, \Pi_{\tilde{\tau}})] \leq \inf_{M\in\mathfrak{M}} \mathbb{E}[\max_{0\leq s\leq T} \{\tilde{h}(s, X_s, \Pi_s) - M_s\}]. \tag{3.1.4}$$

On the other hand, the Snell envelope (With finite horizon T) of $\tilde{h}$ admits a Doob-Meyer decomposition

$$S^T(t, x_t, \pi_t) = M_t - \mathcal{A}_t,$$

where M_t is a martingale and $\mathcal{A}_t$ is the increasing process. The process $M_t^* \in \mathfrak{M}$ is defined as $M_t^* = M_t - M_0$ and we have

$$S^T(t, x_t, \pi_t) = S^T(0, x_0, \pi_0) + M_t^* - \mathcal{A}_t,$$

so that

$$S^T(0, x_0, \pi_0) = \mathcal{A}_t - M_t^* + S^T(t, x_t, \pi_t) > \tilde{h}(t, X_t, \Pi_t) - M_t^*.$$

Hence

$$\begin{aligned}\sup_{\tilde{\tau} \in \mathcal{T}^X} \mathbb{E}[\tilde{h}(\tilde{\tau}, X_{\tilde{\tau}}, \Pi_{\tilde{\tau}})] =& \mathbb{E}[S^T(0, x_0, \pi_0)] \\ &\geq \sup_{0 \leq t \leq T} \mathbb{E}[\max_{0 \leq s \leq T} \{\tilde{h}(s, X_s, \Pi_s) - M_s\}],\end{aligned}$$

which proves that we have equality in (3.1.4) and the infimum is achieved by taking $M_t = M_t^*$ □

The theorem above distinguishes a strong duality relation between our original optimal stopping problem and its corresponding dual problem in equation (3.1.1). This duality propose that any $\mathcal{F}_t^X$-adapted martingale M_t can lead to an upper bound on $S(0, x_0, \pi_0)$ and that the martingale $\{M_t^*\}$ is derived from the Doob-Meyer decomposition of the supermartingale $S(t)$. To be specific we can rewrite (3.1.3) as follows:

$$\begin{aligned}\Pi_t^* =& \mathbb{E}[S(t) \mid X_t, \Pi_t] - \mathbb{E}[S(t) \mid X_{t-1}, \Pi_{t-1}], && (3.1.5)\\ =& \mathbb{E}[h(\tilde{\tau}_t^*, X_{\tilde{\tau}_t^*}, Y_{\tilde{\tau}_t^*}) \mid X_t, \Pi_t] - \mathbb{E}[h(\tilde{\tau}_t^*, X_{\tilde{\tau}_t^*}, Y_{\tilde{\tau}_t^*}) \mid X_{t-1}, \Pi_{t-1}] \;\; t \in \mathcal{K} && (3.1.6)\end{aligned}$$

However it is often impossible to compute the optimal martingale $\{M_t^*\}$, due to the fact that the martingale difference term (3.1.6) involves the optimal stopping time $\tilde{\tau}_t^*$ and filtering distribution Π_t. To address this problem, we introduce approximation scheme in the coming section. We approximate the intractable filtering distribution Π_t by a discrete distribution using interactive particle stochastic filtering algorithm. Equation (3.1.6) propose that we estimate Π_t^* using suboptimal $\mathcal{F}_t^X$-adapted stopping times that estimate $\tilde{\tau}_t^*$. Adding to this, some other heuristic construction can be taken into account for instance, we can define

$$\Pi_t = \mathbb{E}[V(t, X_t, Y_t) \mid \mathcal{F}_t^X] - \mathbb{E}[V(t, X_t, Y_t) \mid \mathcal{F}_{t-1}^X],$$

where $V(t, X_t, Y_t)$ is the value function, which correspond to the optimal stopping problem with fully observable state (X_t, Y_t) defined by,

$$V(t, x_t, \pi_t) = \sup_{\hat{\tau} \in \mathcal{T}_t} \mathbb{E}[h(\hat{\tau}, X_{\hat{\tau}}, Y_{\hat{\tau}}) \mid X_t = x_t, Y_t = y_t], \quad (3.1.7)$$

where $\mathcal{T}_t$ is the set of $\mathcal{F}_t$-stopping times $\hat{\tau}$ that takes values in $\mathcal{K}_t$. On the other hand we can also take

$$\Pi_t = \mathbb{E}[h(\hat{\tau}_t^*, X_{\hat{\tau}_t^*}, Y_{\hat{\tau}_t^*}) \mid X_t, \Pi_t] - \mathbb{E}[h(\hat{\tau}_t^*, X_{\hat{\tau}_t^*}, Y_{\hat{\tau}_t^*}) \mid X_{t-1}, \Pi_{t-1}]$$

where $\hat{\tau}_t^*$ is the optimal $\mathcal{F}_t$-stopping time to problem (3.1.7). The approximation of $\hat{\tau}_t^*$ and $V(t)$ can be used in $\amalg_t$ and still preserve martingale difference property, even though their explicit forms are unknown. In addition, the simple structure of $\hat{\tau}_t^*$ or $V(t)$ as functions of only (X_t, Y_t) gives us the advantage of approximation and τ_t^* or $S(t)$ is a function of $(X_0, \ldots, X_t)$. Then it is simple to generate martingale difference terms using $\hat{\tau}_t^*$ or value function $V(t)$. Note that this may produce less optimal values.

Throughout this chapter we focus on approximating $\amalg_t^*$ in equation (3.1.6) by $\amalg_t^n$ using fixed stopping time τ which is $\mathcal{F}_t$ or $\mathcal{F}_t^X$-adapted:

$$\amalg_t^n := \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_t, \Pi_t^n] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_{t-1}, \Pi_{t-1}^n]$$

where τ_t is the t-indexed stopping time corresponding with $\mathcal{F}_t^X$-adapted τ and Π_t^n (with the superscript n denoting the number of particles) is the approximate filtering distribution at time t, obtained from particle filtering. The notation Λ_t^n represent approximate filtering distribution on the given historical stock prices. Then we define M_t^n as

$$M_0^n = 0; \quad M_t^n = \sum_{i=1}^{t} \amalg_i^n \tag{3.1.8}$$

The next goal is to explain how to generate an estimate filtering distribution using interacting particle stochastic filtering algorithm (genetic-type mutation-selection algorithm) and how to compute the approximate difference using nested simulation algorithm.

3.2 The Filtering Algorithm

In this section we shall consider the discrete-time observation of the historical stock price $\{X_{t_1}, X_{t_2}, \ldots X_{t_k}\}$, even though the model we use is in continuous-time, since computations are usually done in discrete-time. It is difficult to obtain the observation of arbitrarily high volatile of X, particularly when the process Y_t is unobservable. In order to perform optimal stopping under this situation, we begin by estimating the filtered stochastic volatility distribution based on a given observable historical stock price X. We then employ filtering-based duality approach to solve the optimal stopping problem.

3.2.1 Empirical Distribution of the Volatility Process. Under this subsection, we make use of the observable historical stock prices X, this implies that the probability measure $\mathbb{P}$ correspond to all past observation. Let the conditional probability distribution defined by;

$$\rho_j(dy) = \mathbb{P}[Y_{t_j} \in dy \mid X_{t_1}, X_{t_2}, \ldots X_{t_j}], \tag{3.2.1}$$

for each $j = 1 \ldots k$, where k is the number of observable historical data. Since $\rho_j(dy)$ depend on random variables $\{X_{t_1}, X_{t_2}, \ldots X_{t_j}\}$, then it is an ideal to conclude that it is also random and this is also supported by the stochastic filter of unobservable Y given discrete-values of observable X. It is usually a good idea to consider that ρ_j is not random at time j, since it entirely depend on j values that are known at that particular time.

Now we point out how to set up n times-varying particles $\{Y_{t_j}^i : j = 1 \ldots k, i = 1 \ldots n\}$ along with their corresponding probabilities $\{\rho_j^i : j = 1 \ldots k, i = 1 \ldots n\}$, such that the empirical distribution of the particles and the given observations $\{X_{t_1}, X_{t_2}, \dots X_{t_j}\}$ converges to the probability measure $\rho_j(dy)$ as n appraoches infinity. We then apply two steps genetic-type algorithm: selection and mutation steps. The algorithm can be described as follows:

Genetic-type Mutation-Selection Algorithm

Step 1:

Firstly assume that we have $k+1$ past observations at time $t_0 < t_1 < \cdots < t_k$, where t_k is the time of the most recent observation and these times have constant intervals: $t_{j+1} - t_j = \Delta t$. We choose an integrable function Ψ that is bounded in $\mathbb{R}_+$ such that

$$\int \Psi(x)dx = 1 \text{ and } \int \mid x \mid \Psi(x)dx < +\infty$$

where $\Psi(x) = e^{-2|x|}$ and the contraction $\Psi(x) = n^{\frac{1}{3}}\Psi(xn^{\frac{1}{3}})$ for $n > 0$. We begin by dividing the interval $[t_0, t_1]$ into N sub-intervals like $t_0 = s_0 < s_1 < \cdots < s_N = t_1$. As a results, N will represent the number of Euler steps in each time interval $[t_j, t_{j+1}]$.

Iteration 1:

Here we initialise out iteration scheme for the particles that represent our observable and unobservable process as follows:

$$\begin{cases} X_{t_0}, \text{which can be described as the observed value of } X \text{ from the past data at } t_0. \\ Y_{t_0}, \text{ which can be described as an arbitrary value or an estimate of } Y \text{ based on past data.} \end{cases}$$

3.2.2 Mutation Step:.

(a) Simulate N values of fBm, described by

$$(B^{H,i}_{t_0}(s_0), B^{H,i}_{t_0}(s_1), \dots, B^{H,i}_{t_0}(s_N))$$

(b) Recursively simulate the equation (2.1.1) and (2.1.2) in the model using an Euler scheme for $k = 0, \dots, N-1$.

$$\begin{aligned} Y^i_{t_0}(s_{k+1}) =& Y^i_{t_0}(s_k) + \alpha Y^i_{t_0}(s_{k+1})(s_{k+1} - s_k) + \beta(B^{H,i}_{t_0}(s_{k+1}) - B^{H,i}_{t_0}(s_k)) \\ X^i_{t_0}(s_{k+1}) =& X^i_{t_0}(s_k) + \left(r - \frac{\sigma^2(Y^i_{t_0}(s_k))}{2}\right))(s_{k+1} - s_k) + \sigma(Y^i_{t_0}(s_{k+1}))W_k\sqrt{s_{k+1} - s_k} \end{aligned}$$

where W_k are independent standard normal random variables.

(c) In this step, only final values from previous recursive step is kept.

$$Y^i_{t_1} := Y^i_{t_0}(s_N)$$

$$X^i_{t_1} := X^i_{t_0}(s_N)$$

Sub-steps (a) to (c) above are independently repeated n times for each i, in order to obtain the pairs $\{X^i_{t_1}, Y^i_{t_1}\}$ $i = 1, \dots, n$ at the end of mutation step.

3.2.3 Selection Step:. In this step we introduce the discrete probability measure build up from the mutation step pairs $\{X^i_{t_1}, Y^i_{t_1}\}\;\; i = 1, \ldots, n$ as follows:

$$\Lambda^n_1(dy) = \begin{cases} \frac{1}{D}\sum_{i=1}^n \Psi_n(X^i_{t_1} - x_1)\delta_{Y^i_{t_1}}, & \text{if } D > 0 \\ \\ \delta_0, & \text{otherwise} \end{cases}$$

Simulated particles that are close to the observed return value of $x_1 = \log X_{t_1}$ will have a higher weight, because Ψ_n act as an approximation of the Dirac-delta function. The constant D is chosen such that $\Lambda^n_1(dy)$ is a probability measure. $D := \sum_{i=1}^n \Psi_n(X^i_{t_1} - x_1)$.

Iteration From 2 to K:

Apply the mutation step above for each iteration $j = 2, 3, \ldots, k$ with initial values for the Euler scheme, known observation $X_{t_{j-1}}$, at time t_{j-1} and sample volatility value $Y_{t_{j-1}}$ from the distribution $\Lambda^n_{t_{j-1}}$, in which each sample is performed independently from each i. However, the sampling distribution depend entirely on the particles system. In this way we obtain n pairs $\{X^i_{t_j}, Y^i_{t_j}\}\;\; i = 1, \ldots, n$ at the end of mutation step. After obtaining n pairs from mutation step, we then apply selection step and get the following probability measure:

$$\Lambda^n_j(dy) = \begin{cases} \frac{1}{D}\sum_{i=1}^n \Psi_n(X^i_{t_j} - x_j)\delta_{Y^i_{t_j}}, & \text{if } D > 0 \\ \\ \delta_0, & \text{otherwise} \end{cases}$$

Output:

At each step j, the discrete-value distribution Λ^n_j is an estimate of the probability measure $\rho_j(dy)$, this simply means the distribution of Y_{t_j} at time t_j, given the observables $x_1, x_2, x_3, \ldots, x_j$.

Martingale Difference Approximation

In constructing our filtering-based duality approach we show how to approximate the martingale difference. In the rest of the following section, we assume that the suboptimal stopping time τ is defined as

$$\tau = \min\{t \in [1, \ldots, T] \mid h(t, X_t, Y_t) \geq \hat{C}(t, X_t, Y_t)\} \tag{3.2.2}$$

where $\hat{C}(t, X_t, Y_t)$ is a sequence of approximate continuation functions of the value function $V(t)$.Even thought the martingale difference can be computed using $\mathcal{F}_t$ or $\mathcal{F}^X_t$-stopping times, here we choose $\mathcal{F}_t$-stopping times τ defined above to make it easy for presentation.

Now we utilize nested simulation to obtain an approximate of $\amalg^n_t$ from equation (3.1.8). Given discrete-value random distribution Λ^n_j from the particle filtering algorithm above, then we have:

$$\begin{aligned}\amalg_t^n =& \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_t = x_t, Y_t = y_t^i] \\ &- \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_{t-1} = x_{t-1}, Y_{t-1} = y_{t-1}^i]\end{aligned} \tag{3.2.3}$$

where τ_t is defined as

$$\tau_t = \min\{i \in \mathcal{K}_t \mid h(i, X_i, Y_i) \geq \hat{C}(i, X_i, Y_i)\}$$

We now generate m sub-paths that are stopped according to τ_t, to estimate

$$\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, y_t^i] \text{ and } \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, y_{t-1}^i]$$

given initial conditions $X_t = x_t, Y_t = y_t^i, X_{t-1} = x_{t-1}$ and $Y_{t-1}^i = y_{t-1}$ for each t and i. Then we average $h(\tau_t, X_{\tau_t}, Y_{\tau_t})$ over m sub-paths. Finally there are total number of $n \times m$ sub-paths generated to approximate each expectation term which is the main result of nested simulation in the following algorithm. Below we describe Nested Simulation Algorithm.

Approximate Martingale Difference Algorithm

Step 1: In this step we simulate the following:

$$\{(x_t^{ij}, y_t^{ij}), \ldots, (x_t^{ij}, y_T^{ij})\}, \text{ for } j = 1, \ldots, m,$$

from our model (2.1.1) and (2.1.2) with the initial conditions: $X_{t-1} = x_{t-1}$ and $Y_{t-1} = y_{t-1}^i$. Then we compute

$$t_{ij} = \min\{\hat{\tau} \in \mathcal{K}_t : h(\hat{\tau}, x_{\hat{\tau}}^{ij}, y_{\hat{\tau}}^{ij}) \geq \hat{C}(\hat{\tau}, x_{\hat{\tau}}^{ij}, y_{\hat{\tau}}^{ij})\},$$

in order to use τ_t on these sample paths. We then set

$$b_i = \frac{1}{m}\sum_{j=1}^{m} h(t_{ij}, x_{t_{ij}}^{ij}, y_{t_{ij}}^{ij}).$$

For each iteration $i = 1, 2, \ldots, n$ we apply Step 1 above.

Step 2: In this step we set $Z_{t-1,t}^{n,m} := \frac{1}{n}\sum_{i=1}^{n} b_i$ which can be described as an unbiased estimator of

$$\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \Lambda_{t-1}^n].$$

Step 3: Here we evaluate the inequality, for $i = 1, \ldots, n$ as follows:

we check if $h(t, x_t, y^i) \geq \hat{C}(t, x_t, y^i)$, then (t, x_t, y^i) is in the stopping region and set $\tilde{b}_i = h(t, x_t, y^i)$. Otherwise we repeat Step 1 with the initial condition $Y_t = y_t^i$ and $X_t = x_t$ to get $\tilde{b}_i$.

Step 4: As in Step 2, we set $Z_{t,t}^{n,m} := \frac{1}{n}\sum_{i=1}^{n} \tilde{b}_i$, which can also described as an unbiased estimator of

$$\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \Lambda_t^n].$$

Step 5: Finally we set $\tilde{\amalg}_t^n = Z_{t,t}^{n,m} - Z_{t-1,t}^{n,m}$.

3.3 The Pricing Algorithm

In this section, we integrate all the information we have gathered in this chapter, particularly information from the particle filtering algorithm and martingale difference estimator algorithm. We then present an algorithm that yield an asymptotic upper bound on value function $S(0)$. The algorithm is described as follows:

Step 1:

In this step generate a path of observations $\{x_1^k, \ldots, x_T^k\}$ from our model (2.1.1) and (2.1.2) with the initial conditions $Y_0 \sim \pi_0$ and $X_0 = x_0$. Then use particle filtering algorithm to estimate filtering distribution $\{\Lambda_1^{n(k)}, \ldots, \Lambda_T^{n(k)}\}$. Now we compute $\tilde{\amalg}_t^{n(k)}$ from martingale difference algorithm for $t = \{1, \ldots, T\}$, which is an estimation of $\amalg_t^{n(k)}$ defined as

$$\amalg_t^{n(k)} = \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t^k, \Lambda_t^{n(k)}] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}^k, \Lambda_{t-1}^{n(k)}].$$

Summing martingale difference estimate we get

$$\tilde{M}_t^{n(k)} = \sum_{i=1}^{t} \tilde{\amalg}_i^{n(k)}.$$

Evaluate

$$S^k = \max_{t \in \mathcal{K}} \left(\tilde{h}(t, x_t^k, \Lambda_t^{n(k)}) - \tilde{M}_t^{n(k)} \right).$$

For each iteration $k = 1, 2, \ldots, N$, we apply Step 1 above.

Step 2:

Finally, we compute an asymptotic upper bound on the value function $S(0, x_0, \pi_0)$,defined as

$$S_N^\tau = \frac{1}{N} \sum_{k=1}^{N} S^k.$$

3.4 Error Analysis and Convergence

In this section, we analyse the errors accumulated in computing the asymptotic upper bound in our algorithm. In particular, we analyse the error bound and asymptotic convergence. The following assumption will be used throughout this section.

3.4.1 Assumption.

1. $\|h\|_\infty := \max_{t \in \mathcal{K}} \|h(t, \cdot, \cdot)\|_\infty < \infty$
2. $\sup_{y_t \in \mathcal{Y}} p(x_t \mid x_{t-1}, y_t) < \infty$, for all $t \in \mathcal{K}$ and $\{x_0, \ldots, x_T\}$.

For convenience, we will use $\mathbb{E}^{\oplus}[\cdot]$ in short for $\mathbb{E}[\cdot \mid X_0 = x_0, Y_0 \sim \pi_0]$ throughout this section. Now we present martingale difference sequence $\amalg_t^\tau$ and martingale M_t^τ which is an $\mathcal{F}_t^X$-adapted and is influenced by an $\mathcal{F}_t^X$(or$\mathcal{F}_t$)-stopping time τ respectively. This can be formulated as follows:

$$\amalg_t^\tau = \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_t, \Pi_t] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_{t-1}, \Pi_{t-1}], \quad \text{and}$$

$$M_0^\tau = 0; M_t^\tau := \sum_{i=1}^{t} \Pi_t^\tau, \text{ for } t \in \mathcal{K}$$

By Theorem (3.1.1), $\mathbb{E}^{\oplus}[\max_{t\in\mathcal{K}}(\hat{h}(t, X_t, \Pi_t) - M_t^\tau)]$ is an upper bound on the value function $S(0, x_0, \pi_0)$, since M_t^τ is an $\mathcal{F}_t^X$-adapted martingale. Also recall that the approximate martingale difference Π_t^n, following realization of observables $\{x_0, \ldots, x_t\}$ is defined as

$$\Pi_t^n = \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \Lambda_t^n] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \Lambda_{t-1}^n]$$

The empirical estimates of $\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \Lambda_t^n]$ is represented by $Z_{t,t}^{n,m}$ and $\mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \Lambda_{t-1}^n]$ is represented by $Z_{t-1,t}^{n,m}$ as shown from the approximate martingale difference algorithm. This implies that we can use

$$\tilde{\Pi}_t^n = Z_{t,t}^{n,m} - Z_{t-1,t}^{n,m} \text{ and } \tilde{M}_t^n = \tilde{\Pi}_1^n + \cdots + \tilde{\Pi}_t^n,$$

for $t \in \mathcal{K}$ to estimate M_t^n and Π_t^n.

The goal is to evaluate the value $\max_{t\in\mathcal{K}}\{\tilde{h}(t, \Lambda_t^n, x_t) - M_t^n\}$ directly along each path of the observables $\{x_0, \ldots, x_T\}$. However, it was shown (Ye and Zhou, 2013) that this evaluation is infeasible. Therefore Ye and Zhou (2013) suggested an approximation of $\max_{t\in\mathcal{K}}\{\tilde{h}(t, \Lambda_t^n, x_t) - M_t^n\}$ by $\max_{t\in\mathcal{K}}\{\tilde{h}(t, \Lambda_t^n, x_t) - \tilde{M}_t^n\}$. From the fixed observation, there is a transformation from constant term to a random term, due to sampling. The difference are as a result of two sources of noise: the first difference is between π_t and Λ_t^n, the deterministic density and the random measure respectively. This difference converges to zero, if we increase the number of particle n under Assumption (3.4.1). The second difference is as a results of the variability of the nested simulation and can be eliminated by adding the number of sample paths $n \cdot m$. In the following, we introduce a useful proposition to prove Theorem (3.4.3). This theorem shows the convergence of the approximation stated above to the exact solution. In addition, unlike (Ye and Zhou, 2013) which applies the theorem in a Markovian setting, we apply it in a non-Markovian setting.

3.4.2 Proposition. Let $\{\Lambda_0^n, \ldots, \Lambda_T^n\}$ be the random measure generated by particle filtering algorithm for the observation sequence $\{x_0, \ldots, x_T\}$. Assume that the following assumption holds:

$$\|f\|_\infty < \infty \text{ and } \sup_{y_t\in\mathcal{Y}} p(x_t \mid x_{0:t-1}, y_t) < \infty, \quad t \in \mathcal{K}.$$

Then

$$\mathbb{E}\left[\left(\int_{\mathcal{Y}} f(y_t)\pi_t(x_t)dy_t - \int_{\mathcal{Y}} f(y_t)\Lambda_t^n(x_t)dy_t\right)^2\right] \le \frac{\hat{\tau}_t^2 \|f\|_\infty^2}{n}, \quad t = [0, \ldots, T]$$

where the constant $\hat{\tau}_t$ depend on t and $\{x_0, \ldots, x_t\}$, but not on n. More specifically $\hat{\tau}_0 = 1$.

3.4.3 Theorem. *Assume τ is an $\mathcal{F}_t^X$(or $\mathcal{F}_t$)-stopping time. Then*

$$\lim_{n\to\infty} \mathbb{E}^{\oplus}[\max_{t\in\mathcal{K}} \tilde{h}(t, X_t, \Pi_t^n) - \tilde{M}_t^n] = \mathbb{E}^{\oplus}[\max_{t\in\mathcal{K}} \tilde{h}(t, X_t, \Pi_t) - M_t^\tau] \tag{3.4.1}$$

In addition, we have the following inequalities:

$$\mathbb{E}^{\oplus}[\max_{t\in\mathcal{K}} \tilde{h}(t, X_t, \Pi_t) - M_t^{\tau}] - S(0, x_0, \pi_0) \tag{3.4.2}$$

$$\leq 2\sqrt{\sum_{t=1}^{T} \mathbb{E}^{\oplus}[(\Pi_t^* - \Pi_t^{\tau})^2]} \tag{3.4.3}$$

$$\leq 2\sqrt{\sum_{t=1}^{T} \mathbb{E}^{\oplus}[(\mathbb{E}[h(\tau_t^*, X_{\tau_t^*}, Y_{\tau_t^*}) \mid X_t, \Pi_t] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid X_t, \Pi_t])^2]} \tag{3.4.4}$$

Proof. We first prove (3.4.1). Given a sample path of the observation $\{x_0, \ldots, x_T\}$, the difference of $\tilde{h}(t, x_t, \pi_t)$ and $\tilde{h}(t, x_t, \Lambda_t^n)$ is

$$\varphi_t^n := \int_{\mathcal{Y}} h(t, x_t, x_t)\pi_t(y_t)dy_t - \int_{\mathcal{Y}} h(t, x_t, x_t)\Lambda_t^n(y_t)dy_t.$$

From Proposition (3.4.2), we have that

$$\mathbb{E}[| \varphi_t^n |] \leq \sqrt{\mathbb{E}[(\varphi_t^n)^2]} < \frac{\hat{\tau}_t \, \|h\|_\infty^2}{n}$$

for some constant $\hat{\tau}_t$. The difference between $\tilde{M}_t^n$ and M_t^{τ} is described as the sum of the difference between $\tilde{\Pi}_t^n$ and Π_t^{τ}:

$$\Pi_t^{\tau} - \tilde{\Pi}_t^n = \chi_{t,t}^n - \chi_{t-1,t}^n + \xi_{t,t}^{n,m} - \xi_{t-1,t}^{n,m}. \tag{3.4.5}$$

where

$$\begin{cases} \chi_{t,t}^n := \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \pi_t] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \Lambda_t^n] & (3.4.6) \\ \chi_{t-1,t}^n := \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \pi_{t-1}] - \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \Lambda_{t-1}^n] & (3.4.7) \\ \xi_{t-1,t}^{n,m} := \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_{t-1}, \Lambda_{t-1}^n] - Z_{t-1,t}^{n,m} & (3.4.8) \\ \xi_{t,t}^{n,m} := \mathbb{E}[h(\tau_t, X_{\tau_t}, Y_{\tau_t}) \mid x_t, \Lambda_t^n] - Z_{t,t}^{n,m} & (3.4.9) \end{cases}$$

The errors (3.4.6) and (3.4.7) are filtering errors, since we re-define $\chi_{t,t}^n$ as

$$\chi_{t,t}^n = \mathbb{E}\Big[\sum_{j=t}^{T} h(j, X_j, Y_j)\mathbf{1}_{\{\tau_t=j\}} \mid x_t, \pi_t\Big] - \mathbb{E}\Big[\sum_{j=t}^{T} h(j, X_j, Y_j)\mathbf{1}_{\{\tau_t=j\}} \mid x_t, \Lambda_t^n\Big] \tag{3.4.10}$$

$$= \int_{\mathcal{Y}} \mathbf{I}(t, x_t, y_t)\pi_t(y_t)dy_t - \int_{\mathcal{Y}} \mathbf{I}(t, x_t, y_t)\Lambda_t^n(y_t)dy_t \tag{3.4.11}$$

where $\mathbf{I}(t, x_t, y_t)$ is defined as the integrand of $\mathbb{E}\Big[\sum_{j=t}^{T} h(j, X_j, Y_j)\mathbf{1}_{\{\tau_t=j\}} \mid x_t, \pi_t\Big]$, that is

$$\mathbf{I}(t, x_t, y_t) := h(t, x_t, y_t)\mathbf{1}_{\{\tau_t = t\}} + \sum_{j=t+1}^{T} \int h(t, x_t, y_t)\mathbf{1}_{\{\tau_t = j\}} p(dx_{t+1}dy_{t+1} \dots dx_t dy_t \mid x_t, y_t),$$

where $p(dx_{t+1}dy_{t+1} \dots dx_t dy_t \mid x_t, y_t)$ represent the joint probability distribution of $(x_{t+1}, y_{t+1}, \dots, x_j, y_j)$ which is conditional on (t, x_t, y_t). The disjoint set $\{\tau_t = j\}$ for each $t \leq j \leq T$ implies that $\|\mathbf{I}_t\|_\infty \leq \|h\|_\infty$. Following (3.4.11) and Proposition (3.4.2) with $f = \mathbf{I}_t$, it is clear that $\mathbb{E}[| \ \chi^n_{t,t} \ |] \leq \frac{\hat{\tau}'_t \|h\|_\infty}{\sqrt{n}}$ for some constant $\hat{\tau}'_t$. Also $\mathbb{E}[| \ \chi^n_{t-1,t} \ |] \leq \frac{\hat{b}_{t-1}\|h\|_\infty}{\sqrt{n}}$ for some constant $\hat{b}_{t-1}$. The last two errors (3.4.8) and (3.4.9) comes from the variability sampling from nested simulation. The error bound are ensured by Proposition (3.4.2) with $t = 0$. That is

$$\mathbb{E}[| \ \xi^{n,m}_{t-1,t} \ |] \leq \frac{\|h\|_\infty}{\sqrt{nm}} \quad \text{and} \quad \mathbb{E}[| \ \xi^{n,m}_{t,t} \ |] \leq \frac{\|h\|_\infty}{\sqrt{nm}}$$

Now we have

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{E}[| \ (\tilde{h}(t, x_t, \pi_t) - M^\tau_t) - (\tilde{h}(t, x_t, \Lambda^n_t) - \tilde{M}^n_t) \ |] \\
&= \lim_{n\to\infty} \mathbb{E}[| \ \varphi^n_t + \left(\sum_{i=1}^{t} (\tilde{\Pi}^n_i - \Pi^n_i) \right) |] \\
&= 0,
\end{aligned}$$

given a sample path of the sequence of observations $\{x_0, \dots, x_t\}$ for each $t \in \mathcal{K}$. Since

$$\begin{aligned}
&| \max_{t\in\mathcal{K}} \{\tilde{h}(t, x_t, \pi_t) - M^\tau_t\} - \max_{t\in\mathcal{K}} \{| \ \tilde{h}(t, x_t, \Lambda^n_t) - \tilde{M}^n_t\} \ | \\
&\leq \max_{t\in\mathcal{K}} \{| \ (\tilde{h}(t, x_t, \pi_t) - M^\tau_t) - (\tilde{h}(t, x_t, \Lambda^n_t) - \tilde{M}^n_t) \ |\} \\
&\leq \sum_{t=1}^{T} | \ (\tilde{h}(t, x_t, \pi_t) - M^\tau_t) - (\tilde{h}(t, x_t, \Lambda^n_t) - \tilde{M}^n_t) \ |,
\end{aligned}$$

Taking expectation and as n turns to infinity, we have

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{E}[| \ (\tilde{h}(t, x_t, \Lambda^n_t) - \tilde{M}^n_t) - \max_{t\in\mathcal{K}} \{(\tilde{h}(t, x_t, \pi_t) - M^\tau_t)\}] \\
&= 0.
\end{aligned}$$

It can be observed that $\tilde{\Pi}^n_t$ is bounded by $2\|h\|_\infty$ for each $t \in \mathcal{K}$. Therefore, $\{\tilde{h}(t, X_t, \Pi^n_t) - \tilde{M}^n_t\}$ is bounded by $(2t+1)\|h\|_\infty$ and $\max_{t\in\mathcal{K}}\{\tilde{h}(t, X_t, \Pi^n_t) - \tilde{M}^n_t\}$ is bounded by $(2T+1)\|h\|_\infty$. The same applies for Π^τ_t, $\{\tilde{h}(t, X_t, \Pi^n_t) - M^\tau_t\}$ and $\max_{t\in\mathcal{K}}\{\tilde{h}(t, X_t, \Pi^n_t) - M^\tau_t\}$. Then

$$\begin{aligned}
&\lim_{n\to\infty} \mathbb{E}^\oplus[\ | \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - \tilde{M}^n_t\} - \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - M^\tau_t\} \ | \ ] \\
&= \lim_{n\to\infty} \mathbb{E}^\oplus[\mathbb{E}[| \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - \tilde{M}^n_t\} - \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - M^\tau_t\} \ || \ \mathcal{F}^X_T]] \\
&= \mathbb{E}^\oplus[\lim_{n\to\infty} \mathbb{E}[| \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - \tilde{M}^n_t\} - \max_{t\in\mathcal{K}} \{\tilde{h}(t, X_t, \Pi^n_t) - M^\tau_t\} \ || \ \mathcal{F}^X_T]] \qquad (3.4.12) \\
&= 0
\end{aligned}$$

where (3.4.12) follows from the boundedness of the integrand and the dominated convergence theorem. Thus,

$$\lim_{n\to\infty} \mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t^n)-\tilde{M}_t^n\}\big] = \mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t^n)-M_t^{\tau}\}\big].$$

Now we prove the second part of our theorem. We have

$$\begin{aligned}
&\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{\tau}\}\big]-S(0,x_0,\pi_0)\\
&=\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{\tau}\}\big]-\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{*}\}\big]\\
&\le \mathbb{E}^{\oplus}\max_{t\in\mathcal{K}}[\{M_t^*-M_t^{\tau}\}],
\end{aligned}$$

based on the fact that

$$\begin{aligned}
&\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{\tau}\}]-\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{*}\}\\
&\le \max_{t\in\mathcal{K}}\{M_t^*-M_t^{\tau}\}.
\end{aligned}$$

Then

$$\mathbb{E}^{\oplus}\max_{t\in\mathcal{K}}[\{M_t^*-M_t^{\tau}\}\le 2\sqrt{\mathbb{E}^{\oplus}[(M_t^*-M_t^{\tau})^2]} \tag{3.4.13}$$

$$=2\sqrt{\sum_{t=1}^{T}\mathbb{E}^{\oplus}[((M_t^*-M_t^{\tau})-(M_{t-1}^*-M_{t-1}^{\tau}))^2]} \tag{3.4.14}$$

$$=2\sqrt{\sum_{t=1}^{T}\mathbb{E}^{\oplus}[(\Pi_t^*-\Pi_t^{\tau})^2]} \tag{3.4.15}$$

$$=2\sqrt{\sum_{t=1}^{T}\mathbb{E}^{\oplus}[(\mathbb{E}[h(\tau_t^*,X_{\tau_t^*},Y_{\tau_t^*})\mid X_t,\Pi_t]-\mathbb{E}[h(\tau_t,X_{\tau_t},Y_{\tau_t})\mid X_t,\Pi_t])^2]}; \tag{3.4.16}$$

where the inequality in (3.4.13) is induced by the fact that $\{M_t^*-M_t^{\tau}\}$ is a martingale and applying Doob's martingale inequality, (3.4.13) uses the orthogonality property of martingale difference. Lastly recall that

$$\begin{aligned}
\Pi_t^*-\Pi_t^{\tau}=&\Big(\mathbb{E}[h(\tau_t^*,X_{\tau_t^*},Y_{\tau_t^*})\mid\mathcal{F}_t^X]-\mathbb{E}[h(\tau_t,X_{\tau_t},Y_{\tau_t})\mid\mathcal{F}_t^X]\Big)\\
&-\Big(\mathbb{E}[h(\tau_t^*,X_{\tau_t^*},Y_{\tau_t^*})\mid\mathcal{F}_{t-1}^X]-\mathbb{E}[h(\tau_t,X_{\tau_t},Y_{\tau_t})\mid\mathcal{F}_{t-1}^X]\Big);
\end{aligned}$$

then the last equality can be shown by simple algebra and iterated expectation on $\mathcal{F}_{t-1}^X$.

□

From Theorem (3.4.3) above, it is observable that $\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t^n)-\tilde{M}_t^n\}\big]$ converges to $\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{\tau}\}\big]$ when the particle n turns to infinity. As a results, $\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t^n)-\tilde{M}_t^n\}\big]$ is an asymptotic upper bound on $S(0,x_0,\pi_0)$ as n turns to infinity. Furthermore, the difference between $\mathbb{E}^{\oplus}\big[\max_{t\in\mathcal{K}}\{\tilde{h}(t,X_t,\Pi_t)-M_t^{\tau}\}\big]$ and $S(0,x_0,\pi_0)$ is clearly due to the suboptimal stopping time τ.

4. Application and Numerical Results

In this chapter we shall demonstrate the effectiveness of our method in American style derivative pricing, precisely an American option.

4.1 American Options Valuation

In this section we consider an application of our method in pricing an American option on a long-memory stochastic volatility underlying. This model is widely used in asset pricing. In particular, it is used in financial mathematics to represent stock dynamics. Here, we let X_t represent the log-price of a stock at time t under the given (pricing) measure $\mathbb{P}$. We also let Y_t represent instantaneous volatility of X at time t. Recall that even though our model is in continuous time, the stock price X is only observed in a discrete time. We consider stochastic differential equation presented in Rambharat et al. (2010), formulated in discretised form as follows:

$$X_{t+1}^i = X_t^i \exp\{(r - \frac{(\sigma_{t+1}^i)^2}{2})\delta + \sigma_{t+1}^i \sqrt{\delta} Z_{t+1}^{i,1}\}, \quad i = 1, \dots, dx, \tag{4.1.1}$$

$$\sigma_{t+1} = \exp(Y_{t+1}), \tag{4.1.2}$$

$$Y_{t+1}^i = Y_t^i e^{-\lambda\delta} + \theta(1 - e^{-\lambda\delta}) + \gamma\sqrt{\frac{1 - e^{-2\lambda\delta}}{2\lambda}}(B_{t+1}^{H,i} - B_t^{H,i}), \quad i = 1, \dots, dy, \tag{4.1.3}$$

where $\{Z_{t+1}^{i,1}, t = [1, \dots, T]\}, i = 1, \dots, dx$ are independent and identical distribution sequences of Gaussian random variables with standard normal distribution $\mathcal{N}(0, 1)$, $\sigma_{t+1} = \exp(Y_{t+1})$ is referred to as the volatility which is a deterministic function of a dx-dimensional process $\{Y_t, t = [0, \dots, T]\}$, r represent the risk free interest rate, δ is a constant time difference, $B_t^{H,i}$ is a standard fractional Brownian motion with long-memory parameter $H = 0.7$, $\gamma > 0$ is a measure of the process volatility (volatility of volatility), $\theta > 0$ is the mean reversion value and δ is the difference in time points. We assume that $\{Z_t^{i,1}\}$ and $B_t^{H,i}$ are two independent distribution sequences. We also assume that the asset price process $\{X_t, t = [0, \dots, T]\}$ is observable and process $\{Y_t, t = [0, \dots, T]\}$ is unobservable. We consider the American option on asset X_t^1 with exercising opportunities taking values in $\mathcal{K}$. In this case we define our pay-off (earning) function as follows:

$$h(T, X_T, Y_T) = \max((X_T(c_1 + Y_T) - c_2), 0) \tag{4.1.4}$$

and explicitly define $\tilde{h}$ as

$$\tilde{h}(T, X_T, \Lambda_T^n) = \max((X_T(c_1 + \Lambda_T^n) - c_2), 0), \tag{4.1.5}$$

where c_1 and c_2 are constants. Since cannot the compute continuation value directly, we approximate it using Laguerre polynomials basis function (Longstaff and Schwartz, 2001). The approximation continuation function $\hat{C}_t$ is defined as follows:

$$\hat{C}_t = \sum_{j=0}^{1} a_j L_j(X_t) \quad t = [1, \dots, T],$$

where a_j are constant coefficients, $L_0(X_t) = \exp(\frac{-X_t}{2})$ and $L_1(X_t) = \exp(\frac{-X_t}{2})(1 - X_t)$. We note that the basis functions above only depend on the asset price X_t, but not on the volatility $\exp(Y_t)$. This implies that suboptimal stopping time is $\mathcal{F}_t^X$-adapted and the lower bound for the partially observable model is guaranteed.

4.1.1 Remark. In this section, we define conditional probability density function as follows:

$$\begin{aligned} p(X_t \mid Y_t, X_{1:t-1}) &= \prod_{i=1}^{dy} p(X_t^i \mid Y_t^i, X_{1:t-1}^i) \\ &= \prod_{i=1}^{dy} \frac{\exp\{-\frac{(A)^2}{B}\}}{X_t^i \sqrt{B}} \end{aligned}$$

where $A = \ln(\frac{X_t^i}{X_{1:t-1}^i}) - (r - \frac{\exp^2(Y_t^i)}{2})\delta$ and $B = 2\exp^2(Y_t^i)\delta\mu^2$.

The model (4.1.1)-(4.1.3) above represents a practical formulation of an optimal stopping problem under partial observations since the stochastic volatility is not observable in reality, but can only be partially estimated through inference from the observable asset price process $\{X_t, t = [0, \ldots, T]\}$.

4.2 Numerical Results

We compute the price of the above American option on a single asset using the methods developed in the previous chapters. In what follows, we demonstrate our algorithm through a number of numerical experiments with one underlying asset. We are interested in finding the upper bounds on our value function. The lower bounds are trivially determined since we used them as initial point on the filtering-based duality algorithm. In our numerical experiment, we use parameter values presented in Table (4.1) below adapted from Zhou (2013) and Ye and Zhou (2013) and historical stock market prices for DJI (Dow Jones Industrial Average) from Yahoo finance (https://finance.yahoo.com) presented in Figure (4.1) in which Figure (4.1a) is an Adjusted closing prices plot and Figure (4.1b) is a Log-returns of the Adjusted closing prices plot.

Parameter	Value
H	0.7
δ	0.1
λ	2
γ	1.0
r	0.05
θ	$\log(0.2)$
c_1	1
c_2	2

Table 4.1: Parameters Sets.

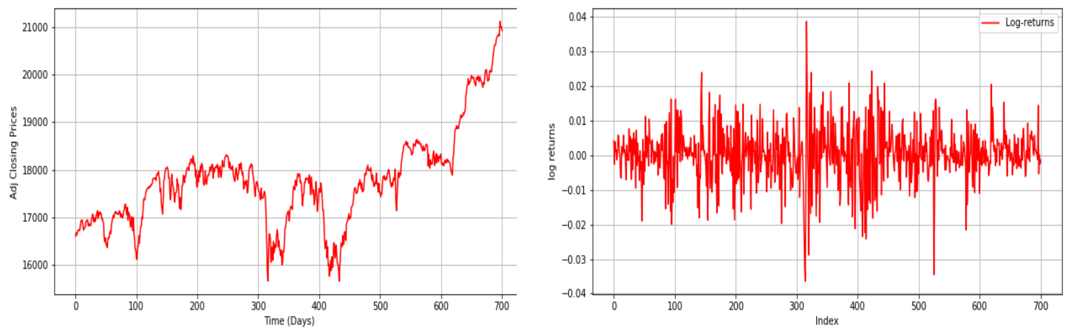

(a) Adjusted Closing Prices (b) Log Returns

Figure 4.1: Daily Stock prices for DJI (Dow Jones Industrial Average) from Yahoo finance.

We used DJI (Dow Jones Industrial Average) daily data from 2014-05-24 to 2017-03-08. Figure (4.1a) represent the evolution of adjusted closing prices over time. Figure (4.1b) represent the log of ration of each consecutive pair of adjusted closing prices. From Figure (4.1a), we observe a fall in adjusted closing price around day 320. This is confirmed in Figure (4.1b) by the highest variance around the same day. This data is used as the historical stock price observations $\{x_1, \ldots, x_j\}$ in the particle filtering algorithm.

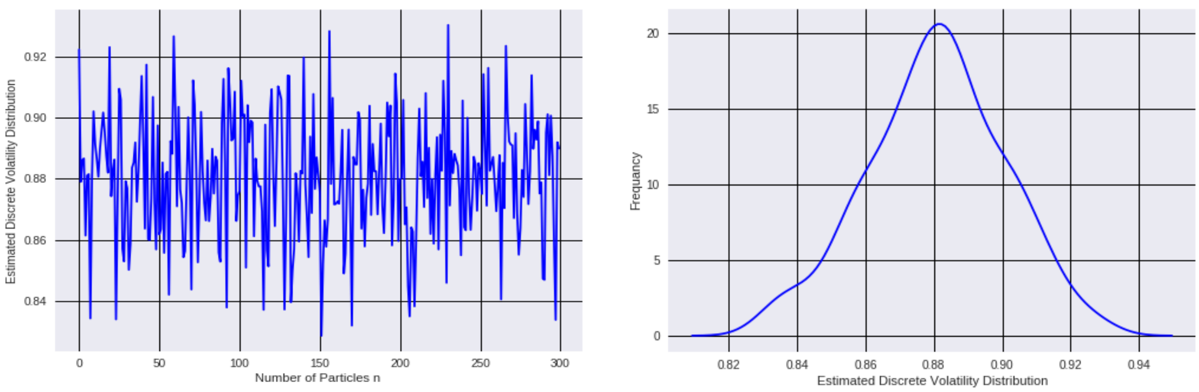

(a) (b)

Figure 4.2: Estimated Discrete Volatility Distribution of the Simulated Model for H=0.7

We constructed the empirical discrete volatility distribution of Y using the particle filter algorithm described in Section (3.2). To achieve that we used the following parameters: $N = 300$ Euler steps and $n = 300$ number of particles. Figure (4.2b) represents the smoothed histogram of an estimated discrete stochastic distribution and Figure (4.2a) represents the particle samples of the discrete random values in sequential order.

Below we present our numerical results with independent paths on the same parameter values in Table (4.1). The upper bounds (UB) are obtained from filtering-based duality method presented earlier for the partially observable model (4.1.1)-(4.1.3) above.

Run	Sample paths (N)	UB ($T = 5$)	UB ($T = 10$)	UB ($T = 15$)
1	50	0.757342576487	0.707996032443	0.700633746124
2	100	0.87369085178	0.8451834498	0.82594290678
3	150	0.923662417212	0.918598176493	0.897516495752
4	200	0.96412501393	0.9466672100	0.92413509776
5	250	0.98935907165	0.98211951171	0.973788366623
6	300	1.00933035366	0.9978427053	0.97744990168
7	350	1.02425361833	1.011990161	1.01077569481
8	400	1.04429408656	1.0220224487	1.02394164159
9	450	1.04585437942	1.03020705775	1.03220021621
10	500	1.04688708566	1.0452517458	1.04034072807
11	550	1.05393297088	1.05210302212	1.04822795933
12	600	1.05459397867	1.05432479663	1.05010101222

Table 4.2: American option prices on one asset.

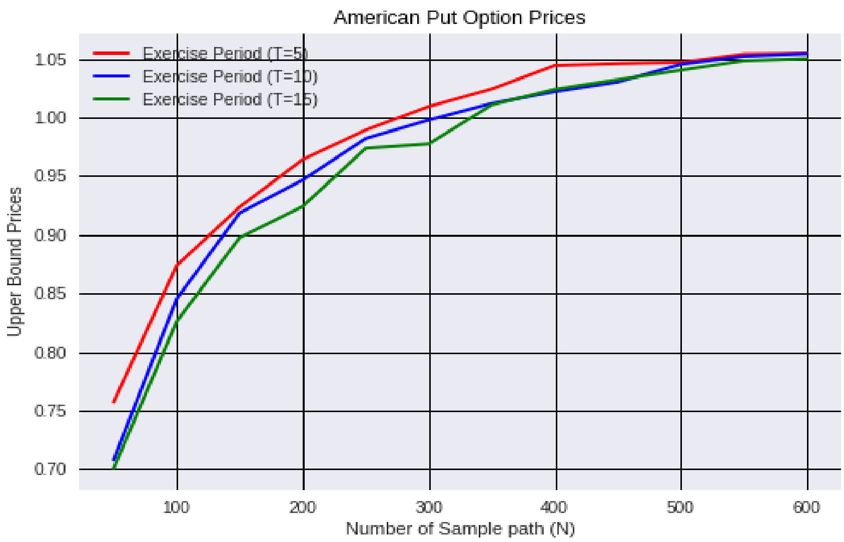

Figure 4.3: Graphical representation of American option prices on a single asset.

For the simulation purpose we used Python to implement our method. The results are shown in Figure (4.2), Table (4.2) and Figure (4.3). Each entry in Table (4.2) displays the average sample of option prices of the numerical results of 12 independent rums (simulations). To obtain upper bounds, we implemented filtering -based martingale duality approach algorithm using the suboptimal stopping time τ with the number of sample paths $N \in [50 : 50 : 600]$, number of particles $n = 50$, number of sub-paths $m = 10$ and the number of exercising opportunities $T \in [5, 10, 15]$.

For a given sample path N, Figure (4.3) shows that the upper bound value decreases as the option maturity (T) increases. This is also shown by each row of Table (4.2). Moreover, for a given option maturity T, the upper bound value is an increasing function of the sample paths N. From Figure (4.3), one can observe that the upper bound is converging for all option maturities. Another can also observe that the upper bound converges faster for longer option maturity and slower for shorter option maturity. Ye and Zhou (2013) claimed that the method we used achieves the smallest gap between the lower and the upper bounds.

4.3 Conclusion

In this project we have used simulation-based approach to compute the value function of an optimal stopping problem under partial observation. Using interactive particle filtering technique, we transformed the problem to a full observation optimal stopping problem by introducing the filtering distribution Π_t. We then used the filtering-based martingale duality approach to approximate the asymptotic upper bound on the value function. We considered an application of our method to a stochastic volatility model in valuing American option. A real world situation where the volatility is not directly observed but can be inferred from the stock prices was considered. From the numerical results obtained, one can observe that this approach is computationally efficient in approximating the value function when the pay-off is nonlinear and the underlying processes are not Makovian.

Acknowledgements

First of all, my appreciation goes to almighty God who kept and strengthen me throughout this project. I would like to take this opportunity to appreciate my supervisors Dr. Dennis Ikpe and Mr. Edward Aidoo who helped me tirelessly throughout this project. Thanks to my tutor Jordan Masakuna for his input and time. I would also like to thank AIMS staff for giving me such a wonderful opportunity to study. My stay at AIMS was indeed fruitful and exceeded my expectation.

Finally, special thank to my friends and family who constantly pray, support and encourage me throughout my studies. Particularly, my parents (Mrs Litshani Tshisikule and Mr Makonde Tshisikule), sister (Ms Tsireledzo Tshisikule), friend (Mphephu Ndivhuwo).etc

References

Applebaum, D. and Blackwood, S. The Kalman-Bucy filter for integrable Lévy processes with infinite second moment. *Journal of Applied Probability*, 52(3):636–648, 2015.

Carmona, R. and Ludkovski, M. Convenience yield model with partial observations and exponential utility. *Preprint. Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ*, 2004.

Chronopoulou, A. and Spiliopoulos, K. Sequential Monte Carlo for fractional stochastic volatility models. *Quantitative Finance*, pages 1–11, 2017.

Chronopoulou, A. and Viens, F. G. Estimation and pricing under long-memory stochastic volatility. *Annals of Finance*, 8(2):379–403, 2012a.

Chronopoulou, A. and Viens, F. G. Stochastic volatility and option pricing with long-memory in discrete and continuous time. *Quantitative Finance*, 12(4):635–649, 2012b.

Crisan, D. Particle approximations for a class of stochastic partial differential equations. *Applied Mathematics & Optimization*, 54(3):293–314, 2006.

Crisan, D. and Lyons, T. A particle approximation of the solution of the Kushner-Stratonovitch equation. *Probability Theory and Related Fields*, 115(4):549–578, 1999.

Crisan, D., Gaines, J., and Lyons, T. Convergence of a branching particle method to the solution of the Zakai equation. *SIAM Journal on Applied Mathematics*, 58(5):1568–1590, 1998.

Décamps, J.-P., Mariotti, T., and Villeneuve, S. Investment timing under incomplete information. *Mathematics of Operations Research*, 30(2):472–500, 2005.

Del Moral, P., Jacod, J., and Protter, P. The Monte-Carlo method for filtering with discrete-time observations. *Probability Theory and Related Fields*, 120(3):346–368, 2001.

Dieker, T. Simulation of fractional Brownian motion. *MSc theses, University of Twente, Amsterdam, The Netherlands*, 2004.

Dixit, A. Entry and exit decisions under uncertainty. *Journal of political Economy*, 97(3):620–638, 1989.

Dixit, A. K. and Pindyck, R. S. *Investment under uncertainty*. Princeton university press, 1994.

Gatarek, D., Świech, A., et al. Optimal stopping in Hilbert spaces and pricing of American options. *Mathematical methods of operations research*, 50(1):135–147, 1999.

Gozzi, F., Rouy, E., and Świech, A. Second order Hamilton-Jacobi equations in Hilbert spaces and stochastic boundary control. *SIAM Journal on Control and Optimization*, 38(2):400–430, 2000.

Haugh, M. B. and Kogan, L. Pricing American options: a duality approach. *Operations Research*, 52 (2):258–270, 2004.

Hill, T. P. Knowing when to stop: how to gamble if you must:the mathematics of optimal stopping. *American Scientist*, 97(2):126–133, 2009.

Jensen, U. and Hsu, G.-H. Optimal stopping by means of point process observations with applications in reliability. *Mathematics of Operations Research*, 18(3):645–657, 1993.

Kaarakka, T. and Salminen, P. On fractional Ornstein-Uhlenbeck processes. *arXiv preprint arXiv:0710.5024*, 2007.

Kushner, H. J. and Budhiraja, A. S. A nonlinear filtering algorithm based on an approximation of the conditional distribution. *IEEE Transactions on Automatic Control*, 45(3):580–585, 2000.

LaLonde, S. M. The martingale stopping theorem. 2013.

Lamberton, D. et al. Optimal stopping and American options. *Ljubljana Summer School on Financial Mathematics*, 2009.

Leão, D., Ohashi, A., and Russo, F. Discrete-type approximations for non-Markovian optimal stopping problems: Part i. *arXiv preprint arXiv:1707.05234*, 2017.

Longstaff, F. A. and Schwartz, E. S. Valuing American options by simulation: a simple least-squares approach. *The review of financial studies*, 14(1):113–147, 2001.

Ludkovski, M. A simulation approach to optimal stopping under partial information. *Stochastic Processes and their Applications*, 119(12):4061–4087, 2009.

Mazziotto, G. Approximations of the optimal stopping problem in partial observation. *Journal of applied probability*, 23(2):341–354, 1986.

McDonald, R. and Siegel, D. The value of waiting to invest. *The Quarterly Journal of Economics*, 101 (4):707–727, 1986.

Miao, J. and Wang, N. Experimentation under uninsurable idiosyncratic risk: An application to entrepreneurial survival. *Unpublished.[224, 225] Miller, RA (1984),"Job matching and occupational choice." Journal of Political Economy*, 92(6):1086–1120, 2007.

Olusaya, O. and Ikpe, D. Optimal stopping problem under partial information with application in finance. 2017.

Pham, H. *Continuous-time stochastic control and optimization with financial applications*, volume 61. Springer Science & Business Media, 2009.

Pham, H., Runggaldier, W., and Sellami, A. Approximation by quantization of the filter process and applications to optimal stopping problems under partial observation. *Monte Carlo Methods and Applications mcma*, 11(1):57–81, 2005.

Pipiras, V. and Taqqu, M. S. *Long-range dependence and self-similarity*, volume 45. Cambridge University Press, 2017.

Rambharat, B. R., Brockwell, A. E., et al. Sequential Monte Carlo pricing of American-style options under stochastic volatility models. *The Annals of Applied Statistics*, 4(1):222–265, 2010.

Rishel, R. and Helmes, K. A variational inequality sufficient condition for optimal stopping with application to an optimal stock selling problem. *SIAM journal on control and optimization*, 45(2):580–598, 2006.

Rogers, L. C. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.

Schwartz, E. S. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.

Trigeorgis, L. *Real options: Managerial flexibility and strategy in resource allocation*. MIT press, 1996.

Vannestål, M. Optimal stopping and the American put under incomplete information, 2011.

Ye, F. and Zhou, E. Optimal stopping of partially observable Markov processes: a filtering-based duality approach. *IEEE Transactions on Automatic Control*, 58(10):2698–2704, 2013.

Zhou, E. Optimal stopping under partial observation: Near-value iteration. *IEEE Transactions on Automatic Control*, 58(2):500–506, 2013.