Pricing American/Bermudan-style Options under Stochastic Volatility

Adam Jankelow

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

July 6, 2020

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy to the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

July 6, 2020

Abstract

A method to price American options under a stochastic volatility framework is introduced which is based on Rambharat and Brockwell (2010). We price American options under the Heston and Bates stochastic volatility models where volatility is assumed to be a latent process. The pricing algorithm is based on the least-squares Monte Carlo approach made popular by Longstaff and Schwartz (2001). Information about the volatility of the underlying asset is used to assist in solving the pricing problem. Since volatility is assumed to be a latent, a particle filter is used to estimate the filtering distribution of volatility. A summary vector is constructed which captures the essential features of the filtering distribution. At each time step before maturity, the elements of the summary vector and the current share price are used as explanatory variables in a regression function which estimates the continuation value of the option. Estimating the continuation value assists in finding the optimal time to exercise the option. This pricing approach is benchmarked against a method which assumes volatility is observable. Furthermore, our pricing approach is compared to simpler methods which do not use particle filtering. Results from our numerical experiments suggest the proposed approach produces accurate option prices.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

American option pricing has been heavily researched in quantitative finance. However, little attention has been given to pricing American options under a latent stochastic volatility framework. The solution to this problem has many practical applications since a large number of exchange-traded options are American. An American option can be exercised at any point in time from inception to expiry. In contrast, a European option can only be exercised at its expiry. The possibility of an early exercise makes pricing an American option more difficult than European options.

The price of an option can be calculated by evaluating the expectation of the option's discounted payoff under a risk-neutral measure (Karatzas *et al.*, 1998). Since the holder of an American option can choose to exercise the option at any time till maturity, the pricing of an American option involves an optimal stopping problem. The value of an American option is the supremum over a range of possible stopping times of the risk-neutral expectation of the discounted payoff of the option (Karatzas *et al.*, 1998). Therefore, valuing an American option involves finding the optimal time to exercise the option. The least-squares Monte Carlo (LSM) algorithm which was made popular by Longstaff and Schwartz (2001) is a method for solving this optimal stopping problem. The basic version of this algorithm assumes that all sources of randomness that affect the price process of the underlying stock are fully observable such as the volatility process. The LSM method requires that the stock price process is Markov. Under a stochastic volatility framework, the stock price process is not Markov. However, the stock and volatility processes are jointly Markov. An LSM method can be used to value an American option if the current stock price and volatility are known.

In this dissertation, it is assumed that the volatility process is latent. The option holder, not knowing the volatility, will exercise sub-optimally relative to the full optimal exercise strategy when volatility is observable. The volatility process of the underlying stock must first be estimated before an LSM type method is used to value an American option. This will ensure the option holder will exercise optimally. One way of estimating the volatility process is by using a particle filter. A particle filter is an algorithm which can be used for estimating the posterior distribution of latent states. It is one of the most popular filtering algorithms for non-linear non-Gaussian models (Douecet *et al.*, 2001).

This dissertation aims to price American options by combining an LSM algorithm with a particle filter for the simultaneous estimation of the value function of an American option and the latent volatility process of the underlying asset. In chapter 2, current research on American option pricing and particle filtering are discussed. In chapter 3, the stochastic volatility models used are presented and pricing methodology is described. More specifically, we show how we can combine the principles of dynamic programming and particle filtering to price American options. In chapter 4, numerical experiments are presented that test the accuracy of our pricing methodology and particle filtering estimation. In Chapter 5, a brief conclusion is provided.

Chapter 2

Literature Review

2.1 Pricing American Options in Stochastic Volatility Models

An option is a financial contract which gives the buyer the right but not the obligation to buy or sell an underlying asset at a specified price. An American options gives the buyer the right to exercise the option at any time up until expiry. Pricing an American option involves finding the maximum of the risk-neutral expectation of the discounted payoff of the option over a range of stopping times (Karatzas *et al.*, 1998). The price of an American option is given by

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}}[e^{-r\tau} g(S_\tau)],$$

where $\mathbb{Q}$ is a risk neutral probability measure and $g(s)$ represents the payoff function of the option. Many effective algorithms for pricing American options assume that the underlying asset process has a constant or directly observable volatility process such as the ones developed by Brennan and Schwartz (1977) and Carr *et al.* (1992).

Derman *et al.* (1996) and many others have recognised that the assumption of constant volatility is inconsistent with empirical findings. The smile is one example which shows evidence against models that assume volatility is constant. The smile effect indicates that options with differing strike prices where all other features of the option being the same, result in differing implied volatility.

These empirical findings motivate the use of stochastic volatility models to model stock price movements. A stochastic volatility model assumes that the volatility process of the underlying asset is stochastic. This class of models add greater flexibility compared to other modelling structures when pricing options.

Pricing American options under stochastic volatility poses both computational and theoretical problems. Rambharat (2012) recognises that it can be difficult to choose a risk-neutral pricing measure and accurate simulation methodology. The latency of volatility is an additional issue in stochastic volatility models that makes pricing American options difficult.

Several authors have investigated the problem of pricing American options under stochastic volatility, such as Clarke and Parrott (1999) and Zhang and Lim (2006). Most research under a stochastic volatility framework assumes that the volatility process is directly observable, including the approaches mentioned above. Clarke and Parrott (1999) describes an implicit finite difference method where the stock price and volatility are variables in a parabolic partial differential equation (PDE). Zhang and Lim (2006) uses a non-lattice pricing model which depends on volatility being observable.

Despite there being research on the problem of latent process estimation in an optimal stopping framework, very little research has focused on estimating latent volatility when pricing American options. Pham *et al.* (2005) demonstrates a quantization algorithm which solves optimal stopping problems when there is only partial information. Additionally, Ludkovski (2009) uses a regression/particle filter approach for optimal stopping problems which involve latent states. The majority of the approaches previously mentioned rely on Monte Carlo methods to price American options.

2.2 Monte Carlo Methods for Option Pricing

2.2.1 Least Squares Monte Carlo Algorithm

The least-squares Monte Carlo (LSM) algorithm is a method for pricing American options. The algorithm uses Monte Carlo simulations and the principles of dynamic programming. A finite set of exercise times $\{0 = t_0 < t_1 < t_2 \ldots < t_N = T\}$ is used to make the algorithm tractable. Since the set of exercise times is discrete, the algorithm is pricing a Bermudan option. The price will converge to that of an American option if the set of exercise times becomes large (Glasserman, 2013).

The LSM algorithm described below is under a stochastic volatility framework where we assume that the variance process V_t and stock price process S_t are jointly Markov and observable. Under this approach, M independent stock price paths are simulated. The discounted payoff of the option, $u_i^{(j)}$, for each path is calculated recursively backwards in time as shown below.

$$u_N^{(j)}(s) = g(s^{(j)}),$$
$$u_i^{(j)}(s, v) = \max(g(s^{(j)}), \mathbb{E}^{\mathbb{Q}}[e^{-r\Delta t_i} u_{i+1}^{(j)}(S_{i+1}^{(j)}, V_{i+1}^{(j)}) | S_i^{(j)} = s^{(j)}, V_i^{(j)} = v^{(j)}]),$$

$$u_0 = \frac{\sum_{j=1}^{M} u_0^{(j)}}{M}, \tag{2.1}$$

where $\Delta t_i = t_i - t_{i-1}$, $i = 1, ..., N$ and $j = 1, ..., M$.

$g(s)$ is the exercise value of the option and $\mathbb{Q}$ represents the risk neutral measure. The value of the option at each time step is the maximum of the exercise value and the continuation value. The continuation value is the expected value of the option conditional on the knowledge of the current stock price and that the holder of the option will only exercise the option in the future. The price of the option at t_0 is calculated by taking the average of the values u_0 over all sample paths which is shown in (2.1).

One of the major difficulties of valuing an American option is calculating the continuation value. We let C_t represent the continuation value of the option at time t. The LSM algorithm uses a least-squares regression function to approximate C_t where the explanatory variables are the current stock price S_t and variance of the underlying asset V_t. It is assumed that the continuation value can be estimated by

$$C_t(s, v) = \mathbb{E}^{\mathbb{Q}}[e^{-r\Delta t_i} u_{i+1}(S_{i+1}, V_{i+1}) | S_i = s, V_i = v] = \sum_{p=0}^{P} \beta_{ip} L_p(s, v), \tag{2.2}$$

for basis functions L_p and constants β_{ip}, $p \in [0, 1, ..., P]$. (2.2) can be equivalently written as

$$C_t(s, v) = \beta_i^T L(s, v),$$

with $\beta_i^T = [\beta_{i0}, ..., \beta_{iP}]$ and $L(s, v) = [L_0(s, v), ..., L_P(s, v)]$. Vector β_i is given by

$$\beta_i = (\mathbb{E}^{\mathbb{Q}}[L(S_i, V_i) L(S_i, V_i)^T])^{-1} \mathbb{E}^{\mathbb{Q}}[L(S_i, V_i) e^{-r\Delta t_i} u_{i+1}(S_{i+1}, V_{i+1})] = B_L^{-1} B_{Lu},$$

where B_L is a $P \times P$ matrix and B_{Lu} is a vector of length P. Observations of S_i, V_i and $u_{i+1}(S_{i+1}, V_{i+1})$ are used to approximate the regression coefficients β_{ip}. The least-squares approximation of the vector β_i is

$$\bar{\beta}_i = \bar{B_L}^{-1} \bar{B}_{Lu},$$

where $\bar{B}_L$ and $\bar{B}_{Lu}$ are sample estimates of B_L and B_{Lu}. Finally, the continuation value can be estimated by

$$\bar{C}_t(s, v) = \bar{\beta}_i^T L(s, v). \tag{2.3}$$

The LSM approach was originally introduced by Tsitsiklis and Van Roy (2001). Various improvements have been made to the algorithm for pricing American options. Longstaff and Schwartz (2001) made some computational improvements. One improvement they suggested was that nodes $S_i^{(j)}$, where $g(S_i^{(j)}) = 0$, should not be included in estimating β_i. Fabozzi *et al.* (2017) presents an improvement to the LSM algorithm which applies a correction for heteroscedasticity. Furthermore, Clément *et al.* (2002) prove convergence of the LSM algorithm as M tends to infinity. Rambharat and Brockwell (2010) suggest that the LSM algorithm is very efficient and tractable. They note that more advanced regression methods such as weighted-least-squares or generalised method of moments can be used in more complex cases.

It is more common that the LSM method is applied when volatility is assumed to constant. Under constant volatility models, the basis functions used for estimating the continuation value are only functions of the current stock price. Figure 2.1 below gives a visual representation of the estimation of the continuation values at the time step before expiry under when volatility is assumed to be constant.

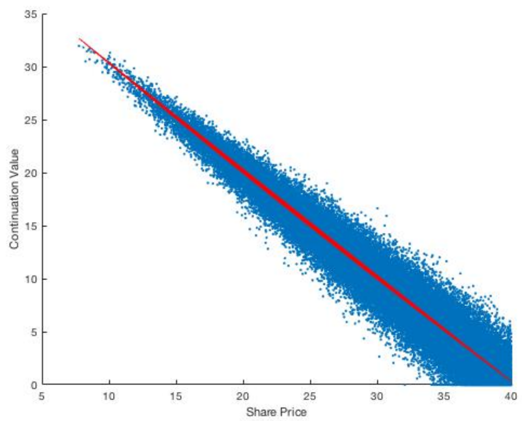

Fig. 2.1: Continuation value estimation using least-square regression.

A put option with 50 possible exercise times is used in this example. The scatter plot represents the realised continuation values and the red line represents the estimated expected continuation values as a function of the current share price. The estimated continuation value can be seen as a cross-sectional average of the realised continuation values. A more detailed description of the LSM method is available in Glasserman (2013).

2.2.2 Alternative Monte Carlo Methods

Rogers (2002) uses an alternative Monte Carlo approach to price American options. The method involves choosing a Lagrangian martingale and simulating paths of the options payoff function. An upper bound for the options price can be found by determining the path-wise supremum of the payoff less the martingale. Ibanez and Zapatero (2004) use Monte Carlo simulations to determine the optimal exercise frontier for multidimensional American options. The method uses a recursive algorithm to compute the optimal exercise frontier backwards in time. Glasserman (2013) presents an alternative method for pricing American options which simplifies the dynamic programming approach. The method is based upon using a parametric class of stopping rules. This approach is only suitable for options comprising of one underlying stock. Glasserman (2013) also suggests that random tree methods can be used to estimate the price of an American option with more than one underlying stock. This approach produces two consistent estimates with high and low bias respectively. One major challenge of the random tree method is that the computational cost of the algorithm increases exponentially as the number of exercise dates increases. The methods mentioned above either assume volatility is observable or they do not make use of the underlying volatility of the option to provide insight into the problem of pricing an American option. A particle filtering algorithm which is explained in section 2.4 could be used to estimate the underlying volatility before pricing an American option.

2.3 Pricing American Options when Volatility is Latent

Particle filtering may assist in the pricing of American options when volatility is not directly observable. There has been some research focusing on this problem such as the work done by Rambharat and Brockwell (2010) and Rambharat (2012). Both of these papers combine the use of the LSM algorithm and a particle filtering approach to estimate the latent volatility process and price an American option. The pricing methodology in this dissertation will be based on the research done by Rambharat and Brockwell (2010) and Rambharat (2012) and will be presented in detail in the chapter 3. In addition, Rambharat and Brockwell (2010) presents a "brute force" gridding approach as an alternative to the LSM algorithm. One of the major drawbacks of this method is that it is computationally costly. Rambharat and Brockwell (2010) uses an Ornstein-Uhlenbeck (OU) process for the log–volatility process in their approach for pricing American options. Rambharat and Brockwell (2010) notes that their approach would be able to apply to a range of stochastic volatility models. Rambharat (2012) uses a square root mean-reverting model for the volatility process in their approach for pricing American options. The model was similarly used by Heston (1993) to evaluate European options.

There has been some research which focuses on optimal stopping problems besides American option pricing under a limited information framework. Bezerra *et al.* (2017) presents an LSM algorithm to solve non-Markovian optimal stopping problems. They propose a Monte Carlo method which can solve optimal stopping problems which apply to several different frameworks such as stochastic volatility and non-Markovian systems. Their method is based upon the work done by Leão *et al.* (2017) who developed a discretisation approach for solving non-Markovian optimal stopping problems.

2.4 Particle Filtering

2.4.1 Overview of Particle Filtering

The estimation of a parameter of a model by using an observed process which depends on this parameter is a common problem. Some practical applications of this problem include using radar measurement to track aeroplanes and estimating the volatility of a financial instrument using share price data (Douecet *et al.*, 2001). The Kalman filter was the first algorithm which was used for the estimation of the posterior distribution of a latent state given an observed state (Kalman *et al.*, 1960). This type of filter was suitable for linear models which had Gaussian noise. An increased focus on solving this type of problem lead to extensions to the algorithm which incorporate non-linear models with non-Gaussian noise. The extended Kalman filter and the Gaussian Sum algorithm were some of the early extensions to the Kalman filter.

Particle filtering is a method that utilizes sequential Monte Carlo simulations for estimating latent states in non-linear models with non-Gaussian noise. Gordon *et al.* (1993) was the first to introduce latent state estimation through sequential Monte Carlo methods. Liu and Chen (1998) improved this approach and their research became known as particle filtering. Some attractive features of this filtering technique is that it is generally tractable and easy to implement compared to other algorithms. Furthermore, particle filtering can be applied to complex models to estimate the distribution of latent states. In general, particle filtering comprises of four main steps propagation, measurement, forecasting and update. A formulation of the problem of estimating latent spaces by using particle filtering is shown below:

Suppose we are working in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where $\{Y_t, t \geq 1\}$ and $\{X_t, t \geq 0\}$. are stochastic processes. We assume Y_t is an observable process which is dependant on the latent process X_t. The posterior distribution of X_t given Y_t needs to be determined in order to estimate the latent process $\{X_t, t \geq 0\}$ conditional on observations of $\{Y_t, t \geq 1\}$. The posterior distribution is denoted by $p(X_{0:t}|Y_{1:t})$ where $X_{0:t} = \{X_0, ..., X_t\}$ and $Y_{1:t} = \{Y_1, ..., Y_t\}$.

The particle filtering algorithm relies on certain information being known at the outset. In, particular it is assumed that we know the

- initial distribution of X_0,
- transitional distribution $p(X_t|X_{t-1}), \quad t \geq 1,$
- conditional distribution $p(Y_t|X_t), \quad t \geq 1.$

Our goal is to recursively estimate the joint posterior distribution $p(X_{0:t}|Y_{1:t})$ over time. This will allow us to determine its marginal distribution $p(X_t|Y_{1:t})$ which is known as the filtering distribution. Bayes Theorem can be used to determine the posterior distribution shown below as

$$p(X_{0:t}|Y_{1:t}) = \frac{p(Y_t|X_t)p(X_{0:t})}{\int p(Y_t|x_t)p(x_{0:t})dx_{0:t}}.$$

This formulation leads to a recursive method for estimating the posterior and marginal distributions through time where new information about the observed process is taken into account as it becomes available. The recursive method for the joint distribution and marginal distribution are

$$p(X_{0:t}|Y_{1:t}) = p(X_{0:t-1}|Y_{1:t-1})\frac{p(Y_t|X_t)p(X_t|X_{t-1})}{p(Y_t|Y_{1:t-1})},$$

$$p(X_t|Y_{1:t}) = \frac{p(Y_t|X_t)p(X_t|Y_{1:t-1})}{\int p(Y_t|x_t)p(x_t|Y_{1:t-1})dx_t}.$$

A sequential Monte Carlo procedure is used to give an empirical estimation of the filtering distribution over time. This estimation converges to the true filtering distribution as the number of particles estimated tends to infinity (Douecet *et al.*, 2001). Particle filtering has many applications in option pricing.

Aihara *et al.* (2009) uses a particle filtering algorithm to estimate stochastic volatility using stock price data under the Heston model. Jasra and Del Moral (2011) demonstrates how sequential Monte Carlo methods can be utilised to estimate derivatives of expectations. These estimates can help calculate Greeks of options under the likelihood ratio method. Moreover, Johannes *et al.* (2009) uses a particle filtering methodology in jump-diffusion models. The particle filter is used to disentangle stochastic volatility from jumps, filter option prices, and returns. The reader can refer to Douecet *et al.* (2001) for a more comprehensive breakdown of sequential Monte Carlo methods.

2.4.2 Sequential Importance Sampling

The Sequential Importance Sampling (SIS) algorithm is a particle filtering method for estimating latent states. More specifically, it produces particles from the filtering distribution $p(X_t|Y_{1:t})$.Time is discretised and the time horizon $[0, T]$ is broken up into the partition $\{0 < t_1 < t_2.. < t_N = T\}$ where each interval is of length Δt. Particles are generated in a step-wise manner over time. Importance weights are attached to each particle and these weights assist in estimating the posterior distribution. The weights are updated at each time step. Shown below is a basic formulation of the algorithm which is based off the method provided by Douecet *et al.* (2001). Since a closed-form solution may not be available for the filtering distribution $p(X_t|Y_{1:t})$, an importance function $\pi(X_t|Y_{1:t})$ must be chosen which is an approximation to the filtering distribution.

- A sample of M particles are generated from the initial distribution of X_0.
- At time-step n:
  - Let $\{x_{n-1}^{(i)}, \tilde{w}_{n-1}^{(i)}, i = 1, 2, ..., M\}$ represent M particles and their weights respectively at time-step $n-1$.
  - Simulate M independent and identically distributed (iid) particles $x_n^{(i)}$ $\sim \pi(X_n|Y_{1:n})$.
  - Use the formula below to update the weights:

$$w_n^{(i)} = \tilde{w}_{n-1}^{(i)} \frac{p(Y_n|X_n^{(i)})p(X_n|X_{n-1})}{\pi(X_t|Y_{1:t})}.$$

  - Normalise these updated weights to obtain $\{x_n^{(i)}, \tilde{w}_n^{(i)}, i = 1, 2, ..., M\}$ which is used in the next time step.
- The filtering distribution is estimated by

$$\bar{p}(x|Y_{1:n}) = \sum_{i=1}^{N} \tilde{w}_n^{(i)} \delta_{X_n^i}(x),$$

  where $\delta_{X_n^i}(x)$ is the Dirac delta function.

The SIS algorithm is attractive as it is easy to implement. It does have drawbacks such as the degeneracy of particles. Over time, their is a tendency for the importance weights attached to some particles to become very small which results in these particles not contributing significantly to the approximation of the filtering distribution. This can cause the filtering estimates to be inaccurate. The choice of the importance function may affect the level of degeneracy of the algorithm. Zaritskii *et al.* (1975) proposes that the transition equation $p(X_n|X_{n-1})$ should be used as the importance function as this distribution minimises the variance of the importance weights. This simplifies the calculation of the weights such that

$$w_n \propto w_{n-1} p(Y_n|X_n^{(i)}).$$

SIS Filtering Example

We now present an example of the SIS algorithm in order to estimate stock prices when we are able to observe option prices of this share. Let Y_t represent the price observed at time t of an European put option with strike price K and time to maturity T. X_t is the underlying stock price at time t which is assumed to be latent. The option price is calculated under the Black-Scholes model where the dynamics of the underlying stock price follow geometric Brownian motion. The transition equation of the latent stock price, $p(X_i|X_{i-1})$, under this modeling framework is

$$X_i = X_{i-1} e^{(r-0.5\sigma^2)\Delta t + \sigma\sqrt{\Delta t} W_i},$$

where X_0 is the observable share price at $t = 0$ and $W_i \sim \mathcal{N}(0, 1)$. We assume that $p(Y_i|X_i^j) = \phi(Y_i, Y_i^j, 1)$, where ϕ is the Gaussian density function with mean Y_i^j and variance 1. We assume that Y_i^j equals the Black-Scholes put option price $(\mathcal{P}(X_i^j; K, T))$ evaluated at the simulated share price (X_i^j). Let M represent the number of particles generated in each time step and N represent the number of time steps. We choose the importance function to be the transition equation $p(X_i|X_{i-1})$ which implies that $w_i^j \propto \phi(Y_i, Y_i^j, 1)$.

The algorithm for estimating the stock price is presented below. The parameters used in this example were $r = 0.1, \sigma = 0.2; S_0 = 50; K = 50; \Delta t = \frac{1}{200}; T = 1; M = 100$ and $N = 400$.

1. Initialise: $X_0 = 50$ and $Y_0 = P(X_0; K, T)$.
2. For $i = 1 : N$ and $j = 1 : M$
   - $X_i^j = X_{i-1} e^{(r-0.5\sigma^2)\Delta t + \sigma\sqrt{\Delta t} W_i}$,
   - $Y_i^j = \mathcal{P}(X_i^j; K, T)$,
   - $w_i^j = \frac{\phi(Y_i, Y_i^j, 1)}{\sum_1^M \phi(Y_i, Y_i^j, 1)}$.
3. $\bar{X}_i = \sum_{j=1}^M w_i^j X_i^j$ represents the estimate of the stock price at time i which is the expectation of the particles according to their empirical distribution.

Figure 2.2 below shows the observable and estimated stock prices as a function of time. The solid blue line represents the true value of the stock price and the dashed red line represents the SIS estimate of the stock price.

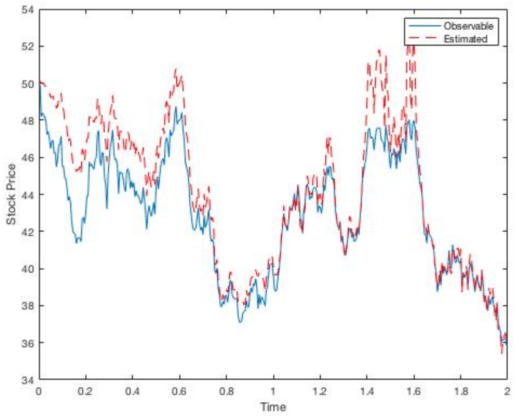

Fig. 2.2: SIS estimation of X_t.

2.4.3 Sequential Importance Resampling

The Sequential Importance Resampling (SIR) algorithm is an extension of the SIS algorithm. This algorithm incorporates an additional selection step into the SIS algorithm where the particles are resampled according to their importance weights. The resampling step aims to transform the M particles of different weights that represent a sample from the filtering distribution into a sample of M particles with equal weights from the filtering distribution (Gellert and Schlogl, 2018). An example of a method used for resampling particles is presented below.

Produce an ordered number u_k for $k = (1, ..., M)$ where

$$u_k = \frac{(k-1) + \tilde{u}}{M},$$
$$\tilde{u} \sim U[0, 1).$$

Particles X_t^j are chosen for the new sample only if

$$u_k \in [\sum_{s=1}^{j-1} w_t^s, \sum_{s=1}^{j} w_t^s).$$

When a particle has a large weight attached to it, it is likely that several ordered numbers (u_k, $k = (1, ..., M)$) will fall into the interval $[\sum_{s=1}^{j-1} w_t^s, \sum_{s=1}^{j-1} w_t^s)$. This will result in duplicates of this particle being present in the new sample after the resampling step. Douecet *et al.* (2001) suggests that this extra step assists in reducing the particle degeneracy effect. The SIR algorithm aims to increase the number of particles with high importance weights and to eliminate particles with low importance weights.

SIR Filtering Example

We now present a SIR algorithm for the problem of estimating the stock price when option prices are observable. This example was introduced in section 2.4.2.

1. Initialise: $X_0 = 50$ and $Y_0 = P(X_0; K, T)$.
2. For $i = 1 : N$ and $j = 1 : M$,
   - $X_i^j = X_{i-1}e^{(r-0.5\sigma^2)\Delta t+\sigma\sqrt{\Delta t}W_i}$,
   - $Y_i^j = \mathcal{P}(X_i^j; K, T)$,
   - $w_i^j = \frac{\phi(Y_i, Y_i^j, 1)}{\sum_1^M \phi(Y_i, Y_i^j, 1)}$.
3. Resample $\{X_i^1, ..., X_i^M\}$ according to $\{w_i^1, ..., w_i^M\}$.
4. $\bar{X}_i = \sum_{j=1}^M \frac{X_i^j}{M}$ represents the estimate of the stock price.

Figure 2.3 displays the observable and estimated stock prices as a function of time. The solid blue line represents the true value of the stock price and the dashed red line represents the SIR estimate. The SIR estimate seems to track the true value more accurately than the SIS estimate as it minimises the degeneracy effect.

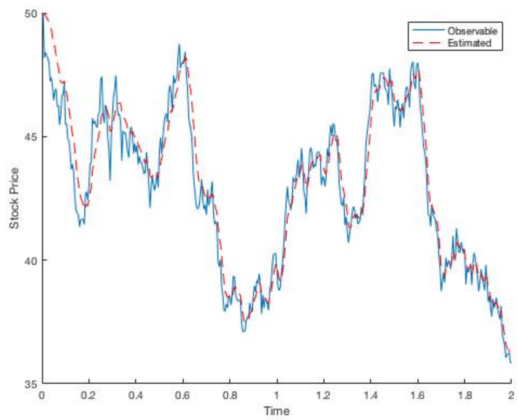

Fig. 2.3: SIR estimation of X_t.

Chapter 3

Option Pricing using Dynamic Programming

In this chapter, the stochastic volatility models and the pricing methodology used are presented. We show how the principles of dynamic programming and particle filtering can be utilised to price American options. This approach is based off the research of Rambharat and Brockwell (2010) and Rambharat (2012).

3.1 Stochastic Volatility Models

3.1.1 Heston Stochastic Volatility Model

Let $\{S_t, t \geq 0\}$ be a stochastic process that describes the price process of a stock over time where $(\Omega, \mathcal{F}, P)$ is the probability space that S_t is defined upon. Assume that S_t changes over time under a risk-neutral measure according to the Itô stochastic differential equations (SDEs)

$$dS_t = rS_t dt + \sqrt{V_t} S_t d\bar{W}_t^{(1)},$$

$$dV_t = \alpha(\beta - V_t)dt + \gamma\sqrt{V_t} dW_t^{(2)},$$

where $\sigma_t^2 = V_t$, r is the risk-free interest rate and V_t is the variance process of the underlying asset. α, β and γ are constants which represent the mean reversion rate, mean reversion level and volatility of V_t. $\bar{W}_t^{(1)}$ and $W_t^{(2)}$ are correlated Wiener processes with a correlation of p such that

$$d\bar{W}_t^{(1)} = \sqrt{1 - \rho^2} dW_t^{(1)} + \rho dW_t^{(2)},$$

where $W_t^{(1)}$ and $W_t^{(2)}$ are independent Wiener processes. The stochastic volatility model shown above was first proposed by Heston (1993). It is an example of a non-linear state-space model with non-Gaussian noise. Since our observations are made at discrete times, we will use the discrete solutions of the SDEs described above to simulate the processes S_t and V_t. It is assumed that stock prices are observed on the time horizon $[0, T]$ at the points, $\{0 < t_1 < t_2 ... < t_N = T\}$ where each interval is of length Δt. Parameters of the variance process are chosen that satisfy the Feller conditions. This ensures that the variance process is always positive (Rambharat, 2012). The exact solution and approximate Euler solution to the SDEs of S_t and V_t respectively are

$$S_{t+1} = S_t \cdot e^{[(r - \frac{V_{t+1}}{2})\Delta t + \sqrt{V_{t+1}\Delta t}(\sqrt{1-\rho^2} Z^{(1)}_{t+1} + \rho Z^{(2)}_{t+1})]}, \tag{3.1}$$

$$V_{t+1} = V_t + \alpha(\beta - V_t)\Delta t + \gamma\sqrt{V_t \Delta t} Z^{(2)}_{t+1}, \tag{3.2}$$

where $Z^{(1)}_t$ and $Z^{(2)}_t$ are independent standard normal random variables. The expression in (3.1) allows use to express the log returns of the stock price $R_{t+1} = \log(S_{t+1}/S_t)$ as

$$R_{t+1} = (r - \frac{V_{t+1}}{2})\Delta t + \sqrt{V_{t+1}\Delta t}(\sqrt{1-\rho^2} Z^{(1)}_{t+1} + \rho Z^{(2)}_{t+1}). \tag{3.3}$$

3.1.2 Bates Stochastic Volatility Model

The Bates stochastic volatility model is an extension of the Heston model. Specifically, a random jump component is added to the stock price process. The jumps are modeled according to a Poisson process. This model was first proposed by Bates (1996). S_t changes over time under a risk-neutral measure according to the Itô stochastic differential equations (SDEs)

$$dS_t = (r - \lambda m) S_t dt + \sqrt{V_t} S_t d\bar{W}^{(1)}_t + \bar{m} S_t dK_t,$$
$$dV_t = \alpha(\beta - V_t)dt + \gamma\sqrt{V_t} dW^{(2)}_t,$$

where

$$K_t = \sum_{i=0}^{N(t)} Y_i,$$
$$N(t) \sim \mathcal{P}(\lambda),$$
$$Y_i \sim \mathcal{N}(\mu_x, \sigma_x),$$
$$\bar{m} = e^{Y_i} - 1,$$
$$m = e^{u_x + \frac{1}{2}\sigma_x^2} - 1.$$

$\bar{W}^{(1)}_t$ and $W^{(2)}_t$ are correlated Wiener processes and $N(t)$ follows a Poisson distribution with intensity λ. We use Euler approximations to simulate the stock price and variance processes. The discrete solutions of S_t, V_t and R_t are

$$S_{t+1} = S_t \cdot e^{[(r - \lambda m - \frac{1}{2}V_{t+1})\Delta t + \sqrt{V_{t+1}\Delta t}(\sqrt{1-\rho^2} Z^{(1)}_{t+1} + \rho Z^{(2)}_{t+1})] + J_{t+1} Y_{t+1}}, \tag{3.4}$$

$$V_{t+1} = V_t + \alpha(\beta - V_t)\Delta t + \gamma\sqrt{V_t \Delta t} Z^{(2)}_{t+1}, \tag{3.5}$$

$$R_{t+1} = (r - \lambda m - \frac{1}{2}V_{t+1})\Delta t + \sqrt{V_{t+1}\Delta t}(\sqrt{1-\rho^2} Z^{(1)}_{t+1} + \rho Z^{(2)}_{t+1}) + J_{t+1}Y_{t+1}, \tag{3.6}$$

where $Z^{(1)t+1}$ and $Z^{(2)}_{t+1}$ are independent standard normal random variables. It is assumed that only a single jump can occur in the interval Δt. This implies that the probability of a jump occurring in the interval Δt is $\lambda \Delta t$. In the expression above, J_i follows a Bernoulli distribution with

$$\mathbb{P}(J_i = 1) = \lambda \Delta t,$$

$$\mathbb{P}(J_i = 0) = 1 - \lambda \Delta t.$$

3.2 Statement of the Problem

The arbitrage-free price of an American option is

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}}[e^{-r\tau} g(S_\tau)],$$

where $\mathbb{Q}$ is a risk neutral probability measure and $g(s)$ represents the payoff function of the option. τ is a random stopping time representing the possible times the option is exercised and $\mathcal{T}$ is the set of all possible stopping times with respect to the filtration $\{\mathcal{F}_t = \sigma(S_0, ..., S_t); t = 0, 1, ...\}$. The pricing methodology presented will be applied to put options as it is never optimal to exercise an American call option on a non-dividend paying asset before maturity.

In this context, $g(s) = \max(K - s, 0)$. The problem is to find the optimal expected discounted payoff of the option. This is equivalent to finding the stopping time τ that results in the supremum of the expectation of the discounted payoff of the option.

3.3 General Method

We utilise the principles of dynamic programming to price American options. A finite set of discrete exercise times $\{0 < t_1 < t_2 .. < t_N = T\}$ is used to make the algorithm tractable. The dynamic programming algorithm calculates the price of the option backwards in time starting from the terminal decision point T.

We let τ be a stopping time which represents the exercise time of the option. The stopping time τ must be found which results in the supremum of the risk-neutral expectation of the discounted payoff of the option. To price an American option, we must find the optimal stopping rule which takes into account the unobserved stochastic volatility process. Rambharat and Brockwell (2010) suggests that the stopping rule can be represented by an optimal decision function.

Let $d_t \in \{H, E\}$ represent the decision of either to hold (H) or to exercise (E) the option at time t after the share price S_t has been observed. The optimal decision would be a function of the share price and volatility if the volatility process was observable. Since volatility is not observable, the holder of an American option only has information about S_t and must estimate the volatility σ_t conditional on the observed share price process. Since $\sqrt{V_t} = \sigma_t$, estimating σ_t is equivalent to estimating the variance process V_t. We denote the optimal decision function by

$$d_t(s_t, \pi_t(v_t)) \in \{H, E\},$$

where

$$\pi_t(v_t) = p(V_t | S_0 = s_0, ..., S_t = s_t), \quad t = 0...T.$$

π_t represents the posterior filtering distribution of the variance process V_t conditional on the stock price observations up until time t. The usual convention is used where upper case letter represent random variables and lowercase letters represent particular observations of the respective random variable. In this particular case, S_j represents the random variable for the share price at time j and s_j represents an observed share price.

Now, let

$$u_T(s_T, \pi_T, d_T) = \begin{cases} g(s_T) & d_T = E, \\ 0 & d_T = H, \end{cases}$$

$$u_t(s_t, \pi_t, d_t) = \begin{cases} g(s_t) & d_t = E, \\ \mathbb{E}^{\mathbb{Q}}[e^{-r\Delta t_i} u_{t+1}(S_{t+1}, \pi_t) | S_t = s_t, \pi_t] & d_t = H. \end{cases}$$

u_t represents the discounted expected payoff of the option at time t where we assume that the optimal decisions at times $t_1, ..., t_N$ are made. $g(s_t)$ represents the payoff of the option at time t. At the time T, the option is exercised if the payoff of the option is positive. At any time t before maturity, the value of u_t is the maximum of the exercise value and continuation value of the option.

Rambharat and Brockwell (2010) shows that the information present in the share price history is equivalent knowing the current share price and filtering distribution. This explains why the conditional expectation in the equation above is conditional on knowing the current share price and filtering distribution. The optimal decision function is calculated by

$$d_t(s_t, \pi_t) = \underset{d_t \in \{E,H\}}{\text{argmax}} (u_t(s_t, \pi_t, dt).$$

The decision function will equal E if the exercise value is greater than the continuation value at time t. The decision function will equal H if the converse is true. It is assumed that the option will be exercised if the exercise value and continuation value are equal. The stopping time τ for each price path is simply

$$\tau = \min(\{t \in \{0, ..., T\} | d_t = E\} \cup \{\infty\}). \tag{3.7}$$

We use the convention that $\tau = \infty$ if the option is never exercised. The risk-neutral option price at inception is

$$\sup_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}}[e^{-r\tau} g(S_\tau)]. \tag{3.8}$$

3.4 Particle Filtering and Summary Vectors

We would need to determine the filtering distributions π_t for each time step to implement the dynamic programming algorithm described above. It would be very complex to use the filtering distributions directly since these distributions have infinite dimensions. The algorithm is simplified by calculating a k-dimensional summary vector Q_t which captures key features of the filtering distribution. The summary vector is

$$Q_T = [f_1(\pi_t), ..., f_k(\pi_t)],$$

where $f_1(\pi_t), ..., f_k(\pi_t)$ are functions of the filtering distribution. Rambharat and Brockwell (2010) suggests that the moments of the filtering distribution could be chosen for the elements of Q_t. As the size of the vector Q_t increases, it is likely that Q_t will more accurately summarise the filtering distribution. However, it will also increase the computational cost of the algorithm.

Now, u_t can be approximated by

$$u_T(s_T, Q_T, d_T) = \begin{cases} g(s_T) & d_T = E, \\ 0 & d_T = H, \end{cases}$$

$$u_t(\mathrm{s}_t, Q_t, d_t) = \begin{cases} g(s_t) & d_t = E, \\ \mathbb{E}^{\mathbb{Q}}[e^{-r\Delta t_i} u_{t+1}(S_{t+1}, Q_{t+1}) | S_t = s_t, Q_t = q_t] & d_t = H. \end{cases}$$

where the optimal decision function is

$$d_t(s_t, Q_t) = \underset{d_t \in \{E,H\}}{\operatorname{argmax}}\, u_t(s_t, Q_t, dt).$$

Once the optimal decision functions are calculated for each price path, (3.7) and (3.8) can be used to find the optimal stopping time and the price of the option at inception.

One of the challenges of the dynamic programming algorithm is calculating the continuation value at each time step, denoted by C_t. Rambharat and Brockwell (2010) approximate the conditional expectation C_t by an LSM regression function. Both the share price S_t and the elements of the summary vector Q_t are used as explanatory variables when computing C_t. Therefore, Q_t must be estimated before the LSM regression algorithm can be used. We used a particle filtering approach in order to estimate the filtering distributions and construct Q_t. A Sequential Importance Sampling (SIR) method is used and the algorithm is presented below. The reader can refer to Douecet *et al.* (2001) for more details on particle filtering methods.

Algorithm 1: Particle Filtering Estimation

At $t = 0$, it is assumed that the current volatility $\sigma_0 = \sqrt{V_0}$ of the underlying asset is observable. This is not strictly necessary, as the algorithm will typically converge on V_t even if the incorrect value of V_0 is used. We simulate M particles at each time step.

For $t = 1, ..., T$, implement the following steps:

1. For $i = 1, 2, ..., M$, simulate $v_t^{(i)}$ from the transition equation $p(v_t|V_{t-1})$. (The particles are easily obtained from (3.2) and (3.5) for the Heston and Bates model respectively).

2. Calculate the weights,

$$w_t^{(i)} = p(r_t|V_t = v_t^{(i)}).$$

   The weights are equal to the density of log-returns R_t conditional on $V_t = v_t^{(i)}$. (The weights are easily obtained from (3.3) and (3.6) for the Heston and Bates model respectively).

3. Re-sample with replacement from $\{v_t^{(1)}, ..., v_t^{(m)}\}$ with probabilities proportional to $w_t^{(1)}, ..., w_t^{(m)}$, to obtain a new sample $\{\tilde{v}_t{}^{(1)}, ..., \tilde{v}_t{}^{(m)}\}$.

Algorithm 1 produces $T+1$ approximate samples of particles $\{\tilde{v_t}^{(1)}, ..., \tilde{v_t}^{(m)}\}$ from the filtering distribution π_t for each time step $t = 0, ..., T$. As m increases, the empirical distributions of the particles converge to the distribution π_t. However, increasing the size of m results in the algorithm becoming more computationally costly. Algorithm 1 can be improved by reducing the effects of degeneration of particles. For more details and improvements to the algorithm, the reader can refer to Liu and Chen (1998).

3.5 Pricing Algorithm

In our approach, the algorithm used for pricing American options is based on the least-squares Monte Carlo (LSM) approach. This method uses the principles of dynamic programming described in section 3.4 to price an option recursively backward in time. Furthermore, a regression function is used to estimate the conditional expectation representing the continuation value of the option at each time step. The conditional expectation is a function of the filtering distribution π_t. This suggests that the estimation of the continuation value could be improved by including summary statistics of the filtering distribution as additional explanatory variables in the regression function. The current share price S_t and the elements of Q_t are the explanatory variable used in the regression function in our approach. The key measures statistics included in the summary vector $Q_t = (u_t, sd_t, \psi_t)$ are the sample mean, standard deviation and skewness of π_t. Algorithm 2 describes the simulation component necessary for our pricing approach.

Algorithm 2: Preliminary step

Let N represent the number of simulated stock price paths. For $n = 1, ..., N$ implement the following steps:

1. Simulate a stock price path $S_n = \{S_0, ..., S_T\}$ under the risk - neutral stochastic volatility model. (This is readily achieved by using equations (3.1) and (3.2) or (3.2) and (3.4)).
2. Apply Algorithm 1 (particle filtering algorithm) in order to provide estimations to the filtering distributions $\{\pi_0, ..., \pi_T\}$.
3. Calculate the key measure statistics which comprise the summary vector Q_n by using the estimations of the filtering distributions.
4. Store the vector $Z_n = (S_n, Q_n)$.

Once the simulated stock prices and summary vectors are obtained, algorithm 3 can be used to price the option. This algorithm describes the LSM regression step which makes use of the summary vector Q_t.

Algorithm 3: Least-Squares Monte Carlo

Simulate N independent paths, where each path consists of realizations of the share price S_t and the summary vector Q_t for time steps $t = 1, 2, ..., T$.

Compute the option price at $t = T$ for each of the N paths by calculating the payoff function $g(S_T)$. The option's payoffs computed are represented by $\{u_T^{(1)}, ..., u_T^{(M)}\}$.

For $t = T-1, T-2, ..., 1$, implement the following steps:

1. Compute the exercise value $g(S_t^{(i)})$ for $i = 1, ..., N$.
2. Evaluate the basis functions $\{\phi_1^{(i)}, ..., \phi_p^{(i)}\}$, of $S_t^{(i)}$ and $Q_t^{(i)}$ for $i = 1, ..., N$. Details of the basis functions chosen are given at the end of the algorithm.
3. Estimate the continuation value of the option at time t by:

   $$E^Q[e^{-r\Delta t_i} u_{t+1}(S_{t+1}, Q_{t+1}) | S_t = s_t, Q_t = q_t] \approx \sum_{k=1}^{p} \beta_{tk}\phi_k(S_t, Q_t),$$

   Note: β_{tk} represent the coefficients of the regression function with p explanatory variables.
4. Compute the value of the option at time t by choosing the maximum of the exercise value and continuation value.

Calculate a Monte Carlo estimate for the price of the option at $t = 0$ by taking the average of the option values at this time over the N paths.

Laguerre functions of S_t and Q_t are chosen as basis functions in step 3 of Algorithm 3. The elements of the summary vector $Q_t = (\mu_t, sd_t, \psi_t)$ are the sample mean, standard deviation and skewness of the filtering distribution. The first two Laguerre functions of S_t, μ_t, sd_t and ψ_t and a few cross terms of these functions are used to make up the design matrix which is used in the regression estimation at each time step. The first two Laguerre functions are $L_0(x) = e^{\frac{-x}{2}}$ and $L_1(x) = e^{\frac{-x}{2}}(1-x)$. Specifically, the basis functions at time step n are

$$L_0(S_n),\quad L_1(S_n),\quad L_0(\mu_n),\quad L_1(\mu_n),\quad L_0(sd_n),\quad L_1(sd_n),\quad L_0(\psi_n),\quad L_1(\psi_n),$$

$$L_0(S_n) \times L_0(\mu_n),\quad L_0(S_n) \times L_0(sd_n),\quad L_0(sd_n) \times L_0(\mu_n),$$

$$L_0(S_n) \times L_0(\psi_n),\quad L_0(\psi_n) \times L_0(sd_n),\quad L_0(\psi_n) \times L_0(\mu_n),$$

$$L_1(S_n) \times L_1(\mu_n),\ L_1(S_n) \times L_1(sd_n),\quad L_1(sd_n) \times L_1(\mu_n),$$
$$L_1(S_n) \times L_1(\psi_n),\quad L_1(\psi_n) \times L_1(sd_n),\quad L_1(\psi_n) \times L_1(\mu_n).$$

Chapter 4

Numerical Experiments

4.1 Description of Numerical Experiments

Several numerical experiments were used to evaluate the performance of the pricing algorithm outlined in the previous chapter. The price of American options was computed under four different methods. All methods use an LSM type algorithm where the current stock price S_t and a measure of volatility are used as explanatory variables in the regression function at each time step. The difference in the methods is how volatility is measured.

1. Method A (Benchmark): Under this approach, the simulated asset price and volatility processes are assumed to be observable. The observed volatility was used as a measure of volatility. This is the benchmark case that the other methods are tested against. The filtration under this method is larger than the filtration under the three other methods as volatility is assumed to be observable under this method. The filtration under this method includes information about the share price and volatility history whereas the filtration under the other methods contain information about the share price history only. A larger filtration could lead to more stopping times which would tend to increase the price of an American option.

2. Method B (Particle Filtering (PF) approach): The asset price process is simulated assuming that volatility is a latent process. This approach uses particle filtering to construct a summary vector Q_t that encapsulates key features of the filtering distribution π_t. This method utilises the approach outlined in the previous chapter.

3. Method C (Basic LSM): The asset price process is simulated assuming that volatility is a latent process. The realization of the asset price process at the previous time step (S_{t-1}) is used as a measure of volatility. The method is the most similar to the original LSM approach.

4. Method D (Realised volatility (RV)): The asset price process is simulated assuming that volatility is a latent process. A measure of realised volatility, RV, is calculated for each time step $(k = 1, ..., N)$ and for every path $(l = 1, ..., M)$.

$$RV_{p,l} = \frac{1}{p}\sum_{k=1}^{p} R_{k,l}^2$$

$R_{k,l}$ is the return of stock at time p along path l. The realised volatility is used as an estimate of underlying volatility in the LSM regression function.

The numerical experiments were applied over various option inputs and model parameters. Specifically, the accuracy of the methods were compared under high/low volatility of V_t, mean reversion levels, time to maturity and different degrees of moneyness.

We also assessed the accuracy of the particle filtering algorithm. The root mean square error (RMSE) is used as a test for the performance of estimating the underlying volatility. This measure indicates how close the estimated and observed volatility lie together. Smith and Hussain (2012) used the RMSE in their paper for measuring the performance of their particle filtering estimate of stochastic volatility. The formula for this measure is

$$\text{RMSE} = \sqrt{\sum_{i=1}^{N}(O_i - E_i)^2},$$

where O_i and E_i represent the expected and observed volatility at time step i. N represents the number of time steps. In addition, graphical analysis was performed to illustrate the accuracy of the particle filtering algorithm in estimating the underlying volatility process of the option. All numerical experiments were done using a Macintosh with 4 GB of RAM. The processor of the computer is a $1,3$ GHz Intel Core i5.

4.2 Results of Numerical Experiments

In figure 4.1 below, the price of American options is plotted for different initial volatility. The prices are computed under method A (blue) and method C (red) under the Heston (left) and Bates (right) model respectively. The prices computed for method A are always greater than the prices computed for method B under this set of parameters. These methods represent the extremes of minimum observation and full observation of volatility. Parameters are used in calculating the option prices where stochastic volatility is prevalent. The difference in the prices under the two methods demonstrates the importance of estimating the underlying stochastic volatility when pricing American options.

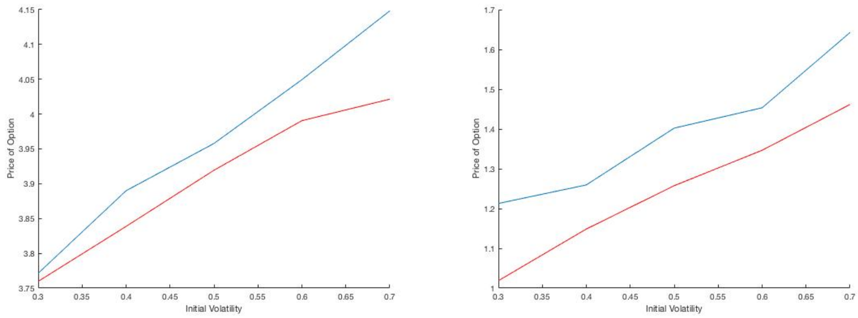

Fig. 4.1: Options prices under the Heston (right) and Bates (left) models as a function of initial volatility.

4.2.1 Particle Filtering Analysis

The following section illustrates the accuracy of the particle filtering estimate of volatility for the Heston and Bates model. Figures 4.2 and 4.3 below represent paths of the volatility of the underlying stock under the parameters values of experiment 1. The solid blue line represents observable volatility and the dashed red line represents the mean of the particle filtering estimates of volatility at each time step. It seems that the particle filtering seems to be estimating underlying quite closely under both stochastic volatility models. The graphs do illustrate that estimations contain some errors. The mean RMSE of all 10 000 volatility paths is below 0.03 under both models which indicates that the estimates of volatility lie close to the observed volatility value.

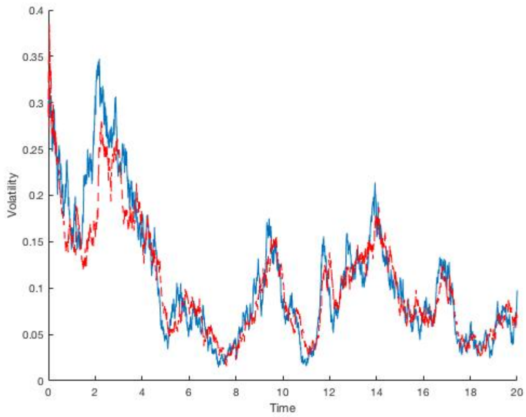

Fig. 4.2: Observable and estimated volatility under the Heston model.

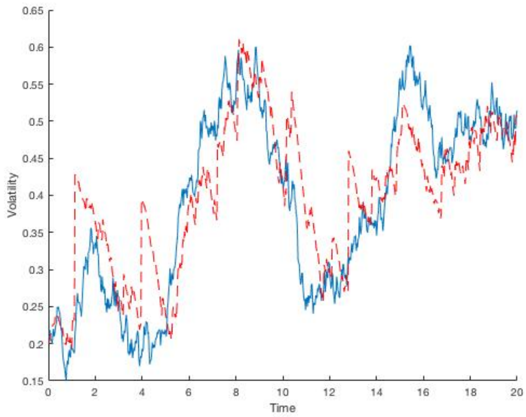

Fig. 4.3: Observable and estimated volatility under the Bates model.

4.2.2 Pricing Results

We experimented with a range of option and model parameters and found cases where simpler methods were just as accurate as the particle filtering approach. However, there were cases where the particle filtering approach produced more accurate results than simpler methods. The table below summarises the values used for the inputs of the American put option for the model parameters. Specifically, the inputs are the strike price, time to expiry in days, initial volatility, initial share price and interest rate which are denoted respectively as K, T, σ_0, S_0 and r.

Experiment no	Option inputs (K, T, σ_0, S_0, r)	Model parameters ($\alpha, \beta, \gamma, p, \lambda$)
1	$95, 40, 0.04, 85, 0.0325$	$3, 0.1, 0.4, -0.2, 40$
2	$95, 40, 0.04, 85, 0.0325$	$3, 0.1, 6.4, -0.2, 40$
3	$50, 20, 7.1, 48, 0.01$	$0.1, 7.1, 0.9 - 0.25, 10$
4	$50, 20, 0.5, 48, 0.01$	$0.2, 0.5, 1.1, -0.25, 10$
5	$60, 50, 0.3, 48, 0.01$	$0.9, 0.3, 0.1, -0.025, 5$
6	$60, 50, 0.3, 60, 0.01$	$0.9, 0.3, 0.1, -0.025, 5$

Tab. 4.1: Description of inputs and parameters values for numerical experiments.

Table 4.2 and 4.3 show the pricing results and standard errors for American options under the Heston and Bates model respectively. The prices are shown under the four proposed methods. 10000 stock price paths were simulated in the pricing exercise, 1000 particles were generated in the particle filtering step for estimating volatility and $\Delta t = \frac{1}{252}$. The mean and variance of the jump sizes under the bates model were -0.1 and 0.1 under all experiments. Experiments 1 and 2 illustrate the difference in pricing results under low and high volatility of variance (γ) respectively. Experiments 3 and 4 illustrate the difference in pricing results under low and high rate of mean reversion (β). Finally, experiments 5 and 6 illustrate the difference in pricing results when the option is initially in-the-money and out-the-money respectively. The experiments also display pricing results under differing lengths of maturity and levels of mean reversion.

Experiment no	A (Benchmark)	B (PF approach)	C (Basic LSM)	D (RV)
1	10.219 (0.0483)	10.203 (0.035)	10.171 (0.0371)	10.158 (0.0369)
2	17.237 (0.152)	17.444 (0.145)	16.257 (0.125)	17.0185 (0.148)
3	32.617 (0.189)	32.987 (0.186)	33.2259 (0.175)	32.640 (0.189)
4	4.802 (0.0475)	4.798 (0.0485)	4.767 (0.0472)	4.745 (0.0469)
5	12.370 (0.0480)	12.442 (0.0491)	12.391 (0.0531)	12.319 (0.0498)
6	1.0993 (0.0251)	1.177 (0.0267)	1.156 (0.0263)	1.1106 (0.0251)

Tab. 4.2: American option pricing results and standard errors under the Heston model.

Experiment no	A (Benchmark)	B (PF approach)	C (Basic LSM)	D (RV)
1	32.950 (0.0646)	32.950 (0.0645)	32.947 (0.06451)	32.948 (0.0369)
2	36.0562 (0.139)	36.150 (0.138)	36.260 (0.141)	36.0561 (0..139)
3	34.328 (0.173)	34.341 (0.172)	34.4074 (0.167)	34.291 (0.172)
4	6.503 (0.0493)	6.532 (0.0499)	6.629 0.0522	6.488 (0.0493)
5	16.366 (0.0636)	16.456 (0.0644)	16.425 (0.0644)	16.389 0.0636
6	3.916 (0.0472)	3.954 (0.0481)	3.860 (0.0465)	3.927 (0.0474)

Tab. 4.3: American option pricing results and standard errors under the Bates model.

The results displayed in table 4.2 and 4.3 show some important observations. The pricing results under Method B are within one standard error of the pricing results under the benchmark method in the majority of the experiments under both the Heston and Bates model. This suggests that the pricing results under Method B seemed to be accurate. Method B's accuracy does not depend on specific option contract features (in/out-the- money, long/short maturity) or special cases (high/low mean reversion rate/level or volatility of the variance process). Experiments 3 and 4 represent cases when the effects of volatility are prevalent. It seems that the pricing accuracy of the particle filtering approach is much higher than the basic LSM method for these experiments. Furthermore, it seems on average that Method D produces more accurate results than Method C under the Heston and Bates models. This may suggest that using realised volatility as an estimate of volatility in the LSM regression improves the accuracy compared to the basic LSM method.

A major drawback of the particle filtering approach is its computational demand. The computational times for method B ranged from 20 to 350 seconds for the experiments above under both stochastic volatility models. The computational times for the three other methods ranged between 1 to 10 seconds. The table below shows the computational times (in seconds) for the first three numerical experiments under the Heston model. The computational times for Method B are significantly greater than the other methods.

Experiment no	A (Benchmark)	B (PF approach)	C (Basic LSM)	D (RV)
1	9.13	191.74	5.05	5.16
2	0.92	86.35	1.01	0.74
3	1.57	64.27	1.68	1.60

Tab. 4.4: Computational times (secs) for experiments 1–3.

Rambharat and Brockwell (2010) used similar numerical experiments to test the pricing accuracy of their particle filtering approach. Their results were similar to the results in this dissertation. In their paper, the particle filtering approach produced results which were within one standard error of the benchmark method in all cases. They noted that the particle filtering approach seemed to perform better than the simpler methods when stochastic volatility was prevalent.

Chapter 5

Conclusion

We introduced a pricing methodology for American options under a stochastic volatility framework where volatility is assumed to be latent. American options were valued under both the Heston and Bates stochastic volatility models. These models take into consideration the codependency between volatility and share prices. In this approach, particle filtering is used to produce estimates of the conditional filtering distribution π_t. It is assumed that π_t can be approximated by a summary vector Q_t which is constructed using the particle filtering estimates. A least-squares Monte Carlo approach is used to price American Options once Q_t is calculated for each time step.

Our numerical experiments show that in the majority of cases, the price of the options using the particle filtering method comes within one standard error of the price using the benchmark method for both stochastic volatility models. This demonstrates the accuracy of the proposed approach and that the adjustment for latency of the volatility process is not major. There are cases where simpler methods produce pricing results which are just as accurate as the particle filtering approach. Despite this, the particle filtering approach produces more accurate prices on average than these simpler methods. One major challenge of our approach, is its computational cost compared to simpler methods. The particle filtering pricing approach may be useful for financial decisions during a period of very volatile markets where simpler pricing methods may not suffice.

Bibliography

Aihara, S. I., Bagchi, A. and Saha, S. (2009). On parameter estimation of stochastic volatility models from stock data using particle filter-application to AEX index, *International journal of innovative computing, information and control* **5**(1): 17–27.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options, *The Review of Financial Studies* **9**(1): 69–107.

Bezerra, S. C., Ohashi, A., Russo, F. and de Souza, F. (2017). Discrete-type approximations for non-Markovian optimal stopping problems: Part ii, *arXiv preprint arXiv:1707.05250* .

Brennan, M. J. and Schwartz, E. S. (1977). The valuation of American put options, *The Journal of Finance* **32**(2): 449–462.

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

Clarke, N. and Parrott, K. (1999). Multigrid for American option pricing with stochastic volatility, *Applied Mathematical Finance* **6**(3): 177–195.

Clément, E., Lamberton, D. and Protter, P. (2002). An analysis of a least squares regression method for American option pricing, *Finance and Stochastics* **6**(4): 449–471.

Derman, E., Kani, I. and Goldman, N. C. (1996). Implied trinomial trees of the volatility smile, *Journal of Derivatives*, Citeseer.

Douecet, A., De Freitas, N. and Gordon, N. (2001). *Sequential Monte Carlo methods in practice*, Springer Science & Business Media.

Fabozzi, F. J., Paletta, T. and Tunaru, R. (2017). An improved least squares monte carlo valuation method based on heteroscedasticity, *European Journal of Operational Research* **263**(2): 698–706.

Gellert, K. and Schlogl, E. (2018). Parameter learning and change detection using a particle filter with accelerated adaptation.

Glasserman, P. (2013). *Monte Carlo methods in financial engineering*, Vol. 53, Springer Science & Business Media.

Gordon, N. J., Salmond, D. J. and Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian bayesian state estimation, *IEE proceedings F (radar and signal processing)*, Vol. 140, IET, pp. 107–113.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The Review of Financial Studies* **6**(2): 327–343.

Ibanez, A. and Zapatero, F. (2004). Monte Carlo valuation of American options through computation of the optimal exercise frontier, *Journal of Financial and Quantitative Analysis* **39**(2): 253–275.

Jasra, A. and Del Moral, P. (2011). Sequential Monte Carlo methods for option pricing, *Stochastic Analysis and Applications* **29**(2): 292–316.

Johannes, M. S., Polson, N. G. and Stroud, J. R. (2009). Optimal filtering of jump diffusions: Extracting latent states from asset prices, *The Review of Financial Studies* **22**(7): 2759–2799.

Kalman, R. E. *et al. (1960).* A new approach to linear filtering and prediction problems, *Journal of Basic Engineering* **82**(1): 35–45.

Karatzas, I., Shreve, S. E., Karatzas, I. and Shreve, S. E. (1998). *Methods of mathematical finance*, Vol. 39, Springer.

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

Liu, J. S. and Chen, R. (1998). Sequential Monte Carlo methods for dynamic systems, *Journal of the American statistical association* **93**(443): 1032–1044.

Longstaff, F. A. and Schwartz, E. S. (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 their Applications* **119**(12): 4061–4087.

Pham, H., Runggaldier, W. and Sellami, A. (2005). 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.

Rambharat, B. R. (2012). American option valuation with particle filters, *Numerical Methods in Finance*, Springer, pp. 51–82.

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

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

Smith, R. and Hussain, M. S. (2012). Genetic algorithm sequential Monte Carlo methods for stochastic volatility and parameter estimation, *Proceedings of the World Congress on Engineering*, Vol. 1, pp. 4–6.

Tsitsiklis, J. N. and Van Roy, B. (2001). Regression methods for pricing complex American-style options, *IEEE Transactions on Neural Networks* **12**(4): 694–703.

Zaritskii, V., Svetnik, V. and Šimelevič, L. (1975). Monte-Carlo technique in problems of optimal information processing, *Avtomatika i telemekhanika* (12): 95–103.

Zhang, Z. and Lim, K.-G. (2006). A non-lattice pricing model of American options under stochastic volatility, *Journal of Futures Markets: Futures, Options, and Other Derivative Products* **26**(5): 417–448.