VALUATION OF MULTIPLE EXERCISE OPTIONS

(Spine title: Valuation of Multiple Exercise Options)

(Thesis format: Monograph)

by

Thomas James <u>Marshall</u>

Graduate Program in Applied Mathematics with Scientific Computing

A thesis submitted in partial fulfillment
of the requirements for the degree of
Doctorate of Philosophy

The School of Graduate and Postdoctoral Studies
Western University of Canada
London, Ontario, Canada

© Thomas James Marshall 2012

WESTERN UNIVERSITY OF CANADA

School of Graduate and Postdoctoral Studies

CERTIFICATE OF EXAMINATION

Examiners:

Supervisor:

....
Dr. Matt Davison

....
Dr. Mark Reesor

Supervisory Committee:

....
Dr. Colin Denniston

....
Dr. Matt Davison

....
Dr. Hubert Pun

....
Dr. Allan MacIsaac

....
Dr. Antony Ware

The thesis by

Thomas James Marshall

entitled:

Valuation of Multiple Exercise Options

is accepted in partial fulfillment of the
requirements for the degree of
Doctorate of Philosophy

....
Date

....
Chair of the Thesis Examination Board

Co-Authorship Statement

The following thesis contains material co-authored by Dr. R. Mark Reesor, Mathew Cox, and Jordan Dickson. The authors retain the right to re-use all or part of their article(s) in subsequent publications provided the source is acknowledge and no financial gain is realized.

Acknowlegements

First and foremost I would like to thank my advisor Dr. Mark Reesor for his guidance over the past 3 years. His insights and advice have been invaluable at each step of way and will continue to be in the future.

I'd also like to thank Dr. Matt Davison and Dr. Adam Metzler who have on many occasions shared there wisdom with me. I owe all the faculty in the Financial Mathematics group a debt for the support they have given me.

Finally, I'd like to thank my fellow grad students especially Piers, Walid, Melissa, Adrian, Alex, Francis, Matt, Bernard and all the others for their always interesting discussion both during and after the work day. I wish all of you success in your future careers and I hope you will look back on these times as fondly as I do.

London, Canada   T. James Marshall

March 2012

Abstract

Multiple exercise options may be considered as generalizations of American-style options as they provide the holder more than one exercise right. Examples of financial derivatives and real options with these properties have become more prevalent over the past decade and appear in sectors ranging from insurance to energy industries. Throughout the thesis particular attention is paid to swing options although the methods described are equally applicable to other types of multiple exercise options. This thesis presents two novel methods for pricing multiple exercise option by simulation; the forest of stochastic trees and the forest of stochastic meshes. The proposed methods are of particular use in cases where there is potentially a large number (≥ 3) of assets underlying the contract and/or if a number of risk factors are desirable for modelling the underlying price process.

These valuation methods result in positively- and negatively-biased estimators for the true option value. We prove the sign of the estimator bias and show that these estimators are consistent for the true option value. A confidence interval for the true option value is easily constructed. Examples confirm that the implementation of these methods is correct and consistent with the theoretical properties of the estimators.

This thesis also explores in detail a number of methods meant to enhance the effectiveness of the proposed simulation methods. These include using high performance computing techniques which include both parallel computing techniques on CPU-clusters and General purpose Graphics Processing Units (GPGPU) that take advantage of relatively inexpensive processors. Additionally we explore bias-corrected estimators for the option values which attempt to estimate the bias introduced at each time step by the estimator and then subtract this result. These improvements are desirable due to the computationally intensive nature of both methods.

Keywords: Monte Carlo, Multiple Exercise Options, Dynamic Programming, Stochastic Optimal Control, High Performance Computing

Contents

List of Figures

List of Tables

List of Appendices

Chapter 1

Pricing American-Style Options

An American-style option gives the holder the right but not the obligation to buy or sell the underlying asset(s) at a prespecified price at any point up to and including the expiry of the option. As discussed later in this thesis multiple exercise options are a generalization of this type of financial derivative so this chapter includes an introduction to the pricing of American-style options with the binomial tree and Monte Carlo methods. These pricing techniques serve as a basis for pricing algorithms for multiple exercise options, the main focus of this thesis.

1.1   American-Style Options

An American-style option is a type of financial derivative that gives the holder the right but not the obligation to buy or sell the underlying asset(s) at a prespecified price at any point up to and including the expiry of the option. American-style options comprise most of the exchange traded options in the United States.

The pricing of American options poses a challenging problem due to the fact that the option may be exercised at any time up to and including its expiry. This means that the pricing problem is path dependent. Indeed finding the value of an American option involves finding the optimal exercise rule and using this to compute the expected discounted payoff of the option. As such the valuation of these options is an example of an optimal stopping time problem.

The broader class of optimal stopping time or early exercise problems includes many important problems in Management Sciences and Operations Research such as development of natural resources, project initiation or abandonment, maintenance scheduling, land use decisions, as well as many others. Very few of these types of problems have closed form analytical solutions, for example the Black-Scholes-Merton formula for pricing European-style options does not have an American-style option analog. Other approaches such as binomial lattice methods, PDE methods, variational inequalities and integral equations have been adopted for pricing these types of derivatives. However all of these methods mentioned are limited in the number of sources of uncertainty and the dimensionality of the underlying asset that can be practically incorporated.

Of particular interest to this thesis are the binomial method of Cox et. al. [1] and its extension to more general diffusion processes [2]. Both have computational costs that grow exponentially in the number of state variables. Another drawback is that convergence proofs for lattice based methods are notoriously difficult, in general [3]. A detailed description of the binomial tree method is given later in this chapter in Section 1.2.

Simulation based methods for asset pricing were introduced in [4] for European-style options. The generality of these methods allow for a large variety of assets to be handled. Unlike lattice methods the convergence rates here are independent of the number of state variables. However, with these methods the speed of computation can be a significant problem. In the seminal paper of Tilley [5] it was shown that simulation based approaches were suitable for optimal stopping problems. Prior to this work the major issue with valuing early exercise options was that they are generally solved via backward in time algorithms due to the optimal exercise policy being easily determined at expiry. Simulation methods are inherently forward based approaches where the paths of the state variable are simulated forward in time and then a pre-specified exercise policy is applied to the state variable trajectory to determine the path

price. Tilley's method used a single state variable and at each time period simulated paths are ordered by asset price and bundled into groups and an optimal exercise decision is estimated for each group. Drawbacks of this approach are outlined in [5, 6].

Two other simulation methods of particular importance to this thesis are the Stochastic Tree and Stochastic Mesh methods. In [6, 7] the authors develop and provide theoretical support for each method. In both cases positively and negatively biased estimators are used to form upper and lower bounds for the value of an American option with all estimators converging asymptotically to the true price. The main drawback of the stochastic tree method [6] is that its computational effort grows exponentially with the number of exercise opportunities. This results in a method that is only practical with a relatively modest number of exercise opportunities. In contrast the stochastic mesh method allows for a larger, but still finite number of exercise dates to be computationally feasible. This method is linear in the number of exercise opportunities and quadratic in the size of the mesh. Detailed descriptions of both the Stochastic Tree and Mesh methods are provided in Section 1.3.

Other simulation methods for pricing American Options include the regression based methods of [8, 9]. These methods all combine simulation with regression on a set of basis functions to develop low dimensional approximations to high dimensional dynamic problems. Although the performance of these methods does not decrease with dimensionality, these methods suffer from other drawbacks. In Least-squares Monte Carlo methods one must select a set of basis functions on which to regress to estimate continuation values. In general only a complete (infinite) set of basis functions results in continuation value estimators that are consistent for the true option value. In practice, of course, a finite set of basis functions is used and introduces an approximation error. Continuation value estimators are consistent for the true approximation value and not the true option value [10, 11].

Duality based methods as in [12] pair an upper bound to a lower bound generated by a simulation based method and in doing so create an interval estimate for the option value. This method then inherits the benefits and drawbacks of the simulation method used to generate its lower bound. An alternative to solving a backward dynamic program are policy iterations for dynamic programs. Where backward schemes require the calculation of nested conditional expectations prior to the time 0 value being approximated, iteration methods such as [13] yield approximations of the time 0 value at each iteration of the dynamic program.

1.1.1 Problem Formulation

Define a process $U(t)$ on $0 \leq t \leq T$, where T is the expiry time of the option, to represent the discounted payoff from exercising an American option at time t. If we are to represent the

set of all possible stopping times in $[0, T]$ by $\mathcal{T}$ the problem is to find the optimal expected discounted payoff,

$$\sup_{\tau \in \mathcal{T}} \mathrm{E}\left[U(\tau)\right], \tag{1.1}$$

which may be considered to be the option price under appropriate conditions [14].

We may specialize this to the case of an American put option with strike price K and a single underlying asset with the time t price, $S(t)$, and constant risk free rate r. Here the time 0 option value is

$$\sup_{\tau \in \mathcal{T}} \mathrm{E}\left[e^{-r\tau} \left(K - S(\tau)\right)^+\right], \tag{1.2}$$

where, $\mathrm{E}[\cdot]$ denotes a risk-neutral expectation and

$$(x)^+ = \begin{cases} x & \text{if } x \geq 0 \\ 0 & \text{if } x < 0. \end{cases}$$

To obtain numerical approximations for the American option price we typically restrict the set $\mathcal{T}$ to be composed of m fixed exercise opportunities $t_0 < t_1 < \ldots < t_m$. We assume that the path of the underlying asset is a Markov chain, $\{S_i\}$, $i = 0, \ldots, m$, that can be generated given the interval $t_{i+1} - t_i$; which is to say that simulated values are not subject to discretization error.

1.1.2   Dynamic Programming Formulation

The option value can now be characterised via dynamic programming. Let the payoff function for a given timestep t_i be denoted by h_i and the option value at the same time step be denoted $B_i(s)$ given that $S_i = s$. The initial option value $B_0(S_0)$ recursively computed through

$$\begin{aligned} B_m(s) &= h_m(s) \\ B_i(s) &= \max\left\{h_i(s), \mathrm{E}\left[D_{i+1,i}(S_i)B_{i+1}(S_{i+1})|S_i = s\right]\right\}, \\ i &= m - 1, \ldots, 0. \end{aligned} \tag{1.3}$$

Here the function $D_{i+1,i}(S_i)$ is the discount factor from time t_{i+1} to t_i. More forms of discounting, such as stochastic interest rates, can easily be incorporated into our proposed valuation methods.

This general form in Equation 1.3 is the basis for most dynamic programming schemes. This process of computing or estimating the conditional expectations in Equation 1.3 is the main difficulty in pricing American options with a numerical scheme. This conditional expectation is referred to as the continuation or hold value. In this discrete time setting this is the

value in holding on to the option's right to exercise until at least the next exercise opportunity.

An approximation to the continuation value, $\hat{H}_i(S_i)$, at time-t_i determines a stopping rule for the option

$$\hat{\tau} = \min \left\{ t_i \in \mathcal{T} : h_i(S_i) \geq \hat{H}_i(S_i) \right\}. \tag{1.4}$$

With these formulations of the American option pricing problem and dynamic programing we dedicate the remainder of this chapter to describing in detail the specific valuation methods most relevant to this thesis.

1.2 Tree Based Methods

In the case of [1] a binomial tree structure as shown in Figure 1.1 was suggested to describe the price paths of the underlying asset. Here given a value at the beginning of a time period the asset may take just one of two values at the end of that period. That is to say that if at some initial time the asset has a price of S_0 then at the end of the first time period it has a value of uS_0 or dS_0 where $u > 1 + r > d$, with r being the risk free rate. If this inequality does not hold then there would be an arbitrage opportunity involving a riskless bond and the asset. This process continues to generate all possible paths through the tree.
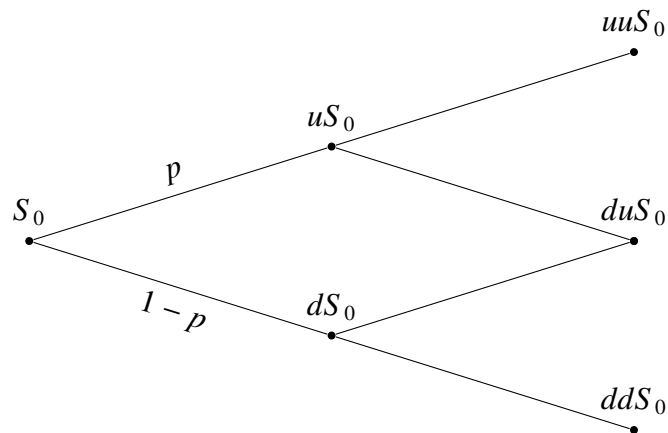
Figure 1.1: Binomial tree with 2 periods

Consider now the initial option value B_0 and the option values B_1^u and B_1^d after up and down stock price moves, respectively. A portfolio containing a long position in Δ shares and a short position in a single option then we may calculate the value of Δ which makes this portfolio riskless. The value after an up move is $\Delta uS_0 - B_1^u$ and after a down move is $\Delta dS_0 - B_1^d$.

Equating these and solving for Δ gives

$$\Delta = \frac{B_1^u - B_1^d}{S_0(u - d)}. \tag{1.5}$$

For this value of Δ the portfolio is riskless and therefore must grow at the riskfree rate, r regardless of the change in the asset value. It then follows that $S_0\Delta - B_0 = (uS_0\Delta - B_1^u)e^{-r\Delta t}$, where Δt is the length of a period. Rearranging gives

$$B_0 = e^{-r\Delta t}\left(pB_1^u + (1 - p)B_1^d\right), \tag{1.6}$$

where,

$$p = \frac{e^{r\Delta t} - d}{u - d}. \tag{1.7}$$

Here p is interpreted as the risk-neutral probability of an upward move in the asset price.

This method is easily generalized to any number of periods. In Figure 1.1 the option values at expiry, B_2^{uu}, B_2^{ud} and B_2^{dd}, are determined from the payoff function and the remaining option values at prior time steps are given by

$$B_1^u = e^{-r\Delta t}\left(pB_2^{uu} + (1 - p)B_2^{ud}\right), \tag{1.8}$$

$$B_1^d = e^{-r\Delta t}\left(pB_2^{ud} + (1 - p)B_2^{dd}\right), \tag{1.9}$$

$$B_0 = e^{-r\Delta t}\left(pB_1^u + (1 - p)B_1^d\right).$$

The choice of u and d are in practice determined from the stock price volatility, σ. They are given by

$$u = A - \sqrt{A^2 - 1}, \tag{1.10}$$

$$d = A + \sqrt{A^2 - 1}, \tag{1.11}$$

where,

$$A = \frac{1}{2}\left(e^{-r\Delta t} + e^{(r+\sigma^2)\Delta t}\right) \tag{1.12}$$

To return briefly to our discussion in the previous section regarding Equation 1.3 the time t_i, the conditional expectation seen there is what is being calculated in Equations 1.6, 1.8, and 1.9. In this case the hold value is being calculated exactly from the underlying binomial model.

As mentioned previously lattice methods like the binomial tree are computationally efficient compared to simulation based methods for cases of small number of underlyings and one or two sources of uncertainty. However beyond this these methods become unfeasible. In the case

of high dimensional underlyings the computational cost grows exponentially with the number of state variables and in cases with more risk factors or more complicated payoff structures these methods also become intractable. In these cases simulation based methods such as the ones described in the following two section become preferable.

1.3 Monte Carlo Methods

1.3.1 Basics of Monte Carlo

Before describing the two simulation based methods for pricing American-style options we first review some fundamental concepts of Monte Carlo simulation. In their simplest form Monte Carlo (MC) methods consist of randomly sampling from the universe of possible outcomes and taking the fraction of these which lie in a given set as an estimate of that set's volume. Given a sample size of n, an estimator $\tilde{\beta}_n$ is defined as a statistic (a function of the data) that is used to infer the value of an unknown parameter, β.

The statement of the consistency of an estimator to the unknown parameter is

$$\tilde{\beta}_n \xrightarrow{p} \beta \quad \text{as } n \to \infty, \tag{1.13}$$

which is to say that for any $\epsilon > 0$,

$$\lim_{n \to \infty} P\left(\left|\tilde{\beta}_n - \beta\right| > \epsilon\right) = 0. \tag{1.14}$$

The bias of an estimator is defined as

$$\text{bias}\left(\tilde{\beta}_n, \beta\right) \equiv \text{E}\left[\tilde{\beta}_n\right] - \beta. \tag{1.15}$$

An estimator is said to be unbiased if bias $\left(\tilde{\beta}_n, \beta\right) = 0$, biased high (positive) if, bias $\left(\tilde{\beta}_n, \beta\right) > 0$ and biased low (negative) if, bias $\left(\tilde{\beta}_n, \beta\right) < 0$. The bias and its sign play important roles in this thesis. The combination of high and low biased estimators allow us to construct confidence intervals for the value of multiple exercise options.

Since this is so crucial to this work we now show that the combination of high and low estimators can be nearly as effective as a single unbiased estimator. Take $\hat{\beta}_n^+(\alpha)$ and $\hat{\beta}_n^-(\alpha)$ to be sample means of n independent replications, for every value of a simulated parameter α. Suppose that it can be shown that

$$\text{E}\left[\hat{\beta}_n^+(\alpha)\right] \geq \beta \geq \text{E}\left[\hat{\beta}_n^-(\alpha)\right], \tag{1.16}$$

which is to say that $\hat{\beta}_n^+$ and $\hat{\beta}_n^-$ are high and low biased estimators of β. Also suppose that $H_n(\alpha)$ is the two standard deviation bound of $\hat{\beta}_n^+(\alpha)$ and similarly $L_n(\alpha)$ is the two standard deviation bound of $\hat{\beta}_n^-(\alpha)$, where both bounds go to zero as $n \to \infty$. Then

$$\hat{\beta}_n^+(\alpha) \pm H_n(\alpha) \tag{1.17}$$

$$\hat{\beta}_n^-(\alpha) \pm L_n(\alpha), \tag{1.18}$$

define the 95% confidence intervals for $\mathrm{E}[\hat{\beta}_n^+(\alpha)]$ and $\mathrm{E}[\hat{\beta}_n^-(\alpha)]$ respectively.

As shown in [6] by taking the upper limit of the high estimator's confidence interval and the lower limit of the low estimator's confidence interval, we get the interval

$$\left(\hat{\beta}_n^+(\alpha) + H_n(\alpha), \hat{\beta}_n^-(\alpha) - L_n(\alpha)\right). \tag{1.19}$$

which is a conservative 95% confidence interval for the value of β. Hence we are able to construct a conservative confidence interval by combining the confidence intervals for two oppositely biased estimators.

In the case of an unbiased estimator the central limit theorem gives us information about the likely error in the estimate. For a random variable X, an estimator $\tilde{\beta}_n$ is the mean of n independent and identically distributed samples. This is to say that

$$\tilde{\beta}_n = \frac{1}{n} \sum_{i=1}^n X_i \tag{1.20}$$

with $\mathrm{E}[X_i] = \beta$ and $\mathrm{Var}[X_i] = \sigma^2 < \infty$. The central limit theorem asserts that the estimator $(\tilde{\beta} - \beta)/(\sigma_X/\sqrt{n})$ converges in distribution to a standard normal random variable. Formally this means that

$$\lim_{n \to \infty} \mathrm{P}\left(\frac{\tilde{\beta}_n - \beta}{s_X/\sqrt{n}} \le x\right) = \Phi(x) \tag{1.21}$$

with $\mathrm{P}(\cdot)$ the probability of some event occurring, Φ the cumulative standard normal distribution function and where the standard deviation, σ_X, is replaced by the sample standard deviation, s_X, since the former is rarely known in practice.

Another motivating feature of MC methods that we will mention repeatedly during this thesis that the convergence of a MC estimator is independent of the dimension of the problem. A simple illustration of this comes from numerical integration. The integral

$$\beta = \int_a^b f(x)dx = (b - a) \int_a^b \frac{f(x)}{(b - a)}dx = (b - a)\mathrm{E}\left[f(U)\right] \tag{1.22}$$

may be represented as the expectation, $E[f(U)]$, with U being a uniform random variable (RV) on $[a, b]$.

If we draw points $U_1, U_2, \ldots, U_n$ independently and evaluate f at these n points then averaging the results and accounting for the volume produces the estimator,

$$\tilde{\beta}_n = \frac{b-a}{n} \sum_{i=1}^{n} f(U_i). \tag{1.23}$$

The estimator, $\tilde{\beta}_n$, is unbiased and consistent to the true value of the integral β. If f is square integrable and we set

$$\sigma_f^2 = \int_a^b (f(x) - \beta)^2 \, dx, \tag{1.24}$$

then the error in the estimate, $\tilde{\beta}_n - \beta$, by using (1.21) is approximately $N\left(0, \frac{\sigma_f^2}{n}\right)$, where $N(\mu, \sigma^2$ describes a random variable that is distributed normally with mean μ and variance σ^2. This implies a square-root convergence rate of the estimate.

Typically σ_f would be unknown but could be estimated via the sample standard deviation

$$s_f = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} \left((b-a)f(U_i) - \tilde{\beta}_n\right)^2} \tag{1.25}$$

In comparision even the simple trapezoidal rule of numerical quadrature

$$\beta = \frac{b-a}{n} \left(\frac{f(a) + f(b)}{2} + \sum_{i=1}^{n-1} f\left(a + \frac{(b-a)i}{n}\right) \right), \tag{1.26}$$

has error $O(n^{-2})$. So in the case of one-dimensional integrals MC is not competitive with traditional methods.

However in the case of estimating an integral over $[a, b]^d$ the f and σ_f change but the error in the MC estimator is still of the form $\frac{\sigma_f}{\sqrt{n}}$. Hence the $O(n^{-1/2})$ convergence rate holds for all d whereas the error in the trapezoidal rule goes as $O(n^{-1/d})$ giving a clear advantage to the MC approach for $d \geq 3$.

A key implication of asset pricing theory is that the value of a derivative security can be represented as an expected value. Often the dimension of these expectation integrals can be very large or even infinite.

1.3.2 Generating Random Numbers

The heart of all MC simulations is a sequence of seemingly random numbers (RN) that are used in the simulation. Since the computers used to generate these random numbers are inherently deterministic the sequence of random numbers will not be truly random. However there is potential for the sequence to mimic true randomness sufficiently for our needs.

A so called uniform pseudo random number generator (PRNG) is used as the basis of MC simulations to generate a finite sequence of numbers $u_1, u_2, \ldots, u_n$ which have the following properties,

(i) u_i is $U[0, 1]$,

(ii) u_i's are mutually independent.

The general considerations that should be taken into account are as follows:

Period Length MC simulations generate large volumes of RNs. All PRNGs will eventually repeat themselves so in order to ensure a quality sequence of numbers it is preferred to only use a small fraction of the generator's period. Hence a longer period is desirable.

Speed Again due to the large volume of RNs required the speed at which the sequence can be generated is paramount to the method being practically useful.

Quality A sequence of RNs is of no use if it does not successfully mimic true randomness. Therefore the sequence must perform satisfactorily under various statistical tests.

Reproducibility It is often important to rerun a simulation while varying inputs. In this case it is desirable to be able to use the same sequence of random numbers.

Skipping Ahead Given the great advantages of implementing parallel computing techniques in MC simulation it is important to be able to skip ahead to different portions of the sequence while still maintaining independent sequences of RNs.

Portability An algorithm should be able to generate RNs on all computing platforms. For instance some generators rely on how overflow is handled on particular computers which limits their portability.

In this thesis two generators in particular are used. The first is the Mersenne-Twister of [15] which is a twisted generalised feedback shift register generator. The particular implementation we use is MT19937-64 which has a period of $2^{19937} - 1$. This is used in the serial and MPI versions of our code in Chapters 3 and 4. In Section 4.5 where we consider GPGPU techniques

we use the method of [16]. This method is shown to be more efficient on GPUs then other traditional methods.

RNGs generate uniform random numbers however most MC simulations entail sampling from non-uniform distributions. Of particular interest in financial mathematics is the standard normal distribution, $N(0, 1)$. The density function of a standard normal is given by

$$\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}, \qquad \infty < x < \infty$$

and the cumulative distribution function is

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-u^2/2} du.$$

If we wish to generate non-standard normal distributions we may use the fact that if $Z \sim N(0, 1)$ then $\mu + \sigma Z \sim N(\mu, \sigma^2)$.

When generating non-uniform RVs two common methods are inverse transform and acceptance-rejection. In the case of Normal RVs the lack of a closed form solution for the inverse of the normal CDF and the relatively slow speed of acceptance-rejection methods have lead to alternative methods. Although there are many approximation methods, [17], to generate normal RVs, the method used for this thesis is the Box-Muller algorithm [18]. It generates two independent bivariate standard normals from two $U[0, 1]$ RVs.

The method is based on the following two properties of independent normal variates; $R = Z_1^2 + Z_2^2$ is exponentially distributed with mean 2; the ordered pair (Z_1, Z_2) is uniformly distributed on a circle of radius $\sqrt{R}$ centered at the origin. The algorithm begins by generating two uniform RVs, $U_1, U_2 \sim U[0, 1]$. U_1 is used to generate an exponential RV R and U_2 is transformed to become $V = 2\pi U_2 \sim U[0, 2\pi]$. The components of the point ($\sqrt{R} \sin V$, $\sqrt{R} \cos V$) are then independent normally distributed RVs. The algorithm is as follows,

- Generate two uniform [0,1] RVs, U_1, U_2

- Calculate $R = -2 \log U_1$

- Calculate $V = 2\pi U_2$

- Calculate $Z_1 = \sqrt{R} \sin V$ and $Z_2 = \sqrt{R} \cos V$

1.3.3 Sample Path Generation

Being able to generate sample paths of the underlying asset(s) is fundamental to MC methods. An important stochastic process used in quantitative finance is geometric Brownian motion

(GBM). GBM is the result of exponentiating Brownian motion (BM) and therefore the methods for simulating BM are also methods for simulating GBM. Note that the valuation algorithms in Chapters 3 and 4 accommodate general stochastic processes. Here we describe sample path generation for GBM.

To review, BM on the interval $[0, T]$ is a stochastic process $\{W(t), 0 \leq t \leq T\}$ such that,

(i) $W(0) = 0$.

(ii) the mapping $t \mapsto W(t)$ is, with probability one, continuous for all $[0, T]$.

(iii) the increments $W(t_m) - W(t_{m-1}), \ldots, W(t_1) - W(t_0)$ are independent for all $0 \leq t_0 \leq \cdots \leq t_m \leq T$.

(iv) $W(t_{i+k}) - W(t_i) \sim N(0, t_{i+k} - t_i)$ for all $0 \leq t_i \leq t_{i+k} \leq T$.

From $W(t)$ and constants μ and σ we may construct a process by setting $X(t) = \mu t + \sigma W(t)$. This process has the dynamics given by

$$dX(t) = \mu dt + \sigma dW(t). \tag{1.27}$$

Due to the independent increments, simulating the points $W(t_i)$ or $X(t_i)$ is straightforward. For independent standard normal random variables $Z_1, \ldots, Z_k$ and beginning from the initial value, subsequent path values may be generated as follows

$$W(t_{i+1}) = W(t_i) + \sqrt{t_{i+1} - t_i} Z_{i+1}, \tag{1.28}$$

$$X(t_{i+1}) = X(t_i) + \mu(t_{i+1} - t_i) + \sqrt{t_{i+1} - t_i} Z_{i+1}, \tag{1.29}$$

for $i = 1, \ldots, m$.

In comparison with a true BM process the values of Equations 1.28–1.29 are exact in that their joint distributions are exact at the time points $t_1, \ldots, t_m$.

As a model for prices, $S(t_i)$, GBM is more desirable than BM because it does not allow for negative values and also it is the percentage changes, $\frac{S(t_{i+1}) - S(t_i)}{S(t_i)}$, that are independent in GBM rather then the absolute changes as in BM.

The dynamics of GBM are typically written as

$$\frac{dS(t)}{S(t)} = \mu dt + \sigma dW(t), \tag{1.30}$$

where the parameter μ is referred to as the drift parameter and σ is referred to as the volatility parameter. Using Itô's lemma it can be shown that the above stochastic differential equation

has solution

$$S(t) = S(0) \exp\left((\mu - \frac{1}{2}\sigma^2)t + \sigma W(t)\right), \tag{1.31}$$

where $S(0)$ is the initial value of the process. A simple recursive relation may be used to generate successive values of this process for $0 \le t_1 \cdots \le t_m \le T$;

$$S(t_{i+1}) = S(t_i) \exp\left((\mu - \frac{1}{2}\sigma^2)(t_{i+1} - t_i) + \sigma \sqrt{t_{i+1} - t_i}Z_{i+1}\right), \tag{1.32}$$

for $i = 1, \ldots, m-1$ and where $Z_1, \ldots, Z_k$ are independent standard normal random variables. This method again introduces no discretization error at times $t_1, \ldots, t_m$.

1.3.4 Stochastic Tree Method

We now define the vector valued Markov process $\{\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_m\}$ on $\mathbb{R}^d$ as the vector of underlying asset prices, where d is the number of underlyings. The value of the American option written on the underlyings at time t_i is

$$B_i = \max_{t_i \le \tau \le t_m} \mathrm{E}\left[h_\tau(\mathbf{S}_\tau)\right], \tag{1.33}$$

where $h_\tau(\mathbf{S}_\tau)$ is the discounted payoff of the option at time τ. Using dynamic programming this can be solved by the recursive relation

$$B_m = h_m(\mathbf{S}_m), \tag{1.34}$$

$$H_i = \mathrm{E}\left[D_{i+1,i}B_{i+1}|\mathcal{Z}_i\right], \tag{1.35}$$

$$B_i = \max\left(h_i(\mathbf{S}_i), H_i\right), \tag{1.36}$$

$$i = m - 1, \ldots, 0, \tag{1.37}$$

where $\mathcal{Z}_i$ is the filtration of the system at time-t_i and $D_{i+1,i}$ is the discount factor which without loss of generality we disregard for the remainder of this section. Here H_i is the time-t_i hold or continuation value representing the expected value to the holder if they decided not to exercise at t_i.

The stochastic tree method produces estimates for the value of an American option. Two main estimators are outlined below, both of which converge to the true option value and may be used to generate confidence intervals. The requirements for this method are little more than the ability to generate Markov chains to represent the underlying asset.

The tree itself is constructed beginning at an initial state $\mathbf{S}_0^1 = \mathbf{S}_0$ from which b independent successor node states ($b \ge 2$) are randomly simulated. Then for each time-t_1 node, $\mathbf{S}_1^{j_1}$ again

randomly simulate b successor nodes, continuing this process until the m^{th} time step and in doing so create a tree that recombines with probability zero. A generic node is then represented by $\mathbf{S}_i^{j_1 j_2 \cdots j_i}$ which for convenience we represent as $\mathbf{S}_i^{\mathbf{j}}$ to make the notation more compact. The basic structure of a stochastic tree with 2 timesteps and a branching factor of 3 is shown in Figure 1.2. As can be seen from this diagram one of the main drawbacks of this method is that the computational effort is exponential in the number of time steps although it is not in the number of dimensions.
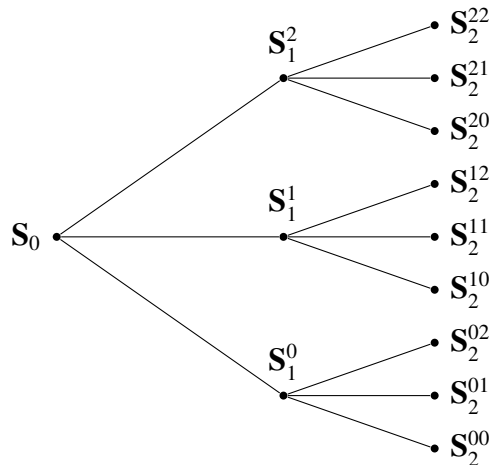
Figure 1.2: Stochastic Tree at timestep 2, $b = 3$

The first estimator we discuss is the high-biased estimator, denoted $\hat{V}$. Working backward, the above recursion relation takes the form

$$\hat{V}_m^{\mathbf{j}} = h_m(\mathbf{S}_m^{\mathbf{j}}) \tag{1.38}$$

$$\hat{H}_i^{\mathbf{j}+} = \frac{1}{b} \sum_{k=1}^{b} \hat{V}_{i+1}^{\mathbf{k}} \tag{1.39}$$

$$\hat{V}_i^{\mathbf{j}} = \max\left\{ h_i\left(\mathbf{S}_i^{\mathbf{j}}\right), \hat{H}_i^{\mathbf{j}+} \right\}, \tag{1.40}$$

where $\mathbf{k} = \{\mathbf{j}, k\}$ and the $+$ signifies the positive bias of the hold estimator.

Intuitively the high bias of the estimator can be reasoned by the fact that it unfairly looks into the future when making the decision. This conditioning on future information allows for better decision than otherwise possible and hence this estimator should give an estimate greater than the true value. A rigorous proof of the high bias property of this estimator is found in Theorem 2 of [6] and is restated here due to its importance. In what follows we expand the

notation to reduce the number of indices that are used and include the branching factor, b, as an argument for the estimator.

Theorem 1 *(Stochastic tree high estimator bias) The stochastic tree high estimator $\hat{V}_0(b, \mathbf{S}_0)$ is biased high. That is,*

$$\mathrm{E}\left[\hat{V}_0(b, \mathbf{S}_0)\right] \geq B_0(\mathbf{S}_0), \tag{1.41}$$

where $B_0(\mathbf{S}_0)$ is the true value.

The consistency of this estimator is described, and subsequently proven, in Theorem 1 of [6] and again we restate it here.

Theorem 2 *(Stochastic tree high estimator consistency) The Stochastic tree high estimator $\hat{V}_0(b, \mathbf{S}_0)$ converges to the true value $B_0(\mathbf{S}_0)$ in the p-norm as $b \to \infty$. That is,*

$$\| \mathrm{E}\left[\hat{V}_0(b, \mathbf{S}_0)\right] - B_0(\mathbf{S}_0) \|_p \to 0. \tag{1.42}$$

An alternative to the high estimator is the low estimator which we denote $\hat{v}$. As its name suggest this estimator is biased low. Consider an arbitrary node in the tree with $i < m$. Then divide its successor nodes into two disjoint sets $\mathcal{X}_1$ and $\mathcal{X}_2$ with sample means $\bar{\mathcal{X}}_1$ and $\bar{\mathcal{X}}_2$. We then define the exercise decision as

$$\hat{v} = \begin{cases} h & \text{if } \bar{\mathcal{X}}_1 \leq h \\ \bar{\mathcal{X}}_2 & \text{otherwise.} \end{cases} \tag{1.43}$$

This leads to many possible choices for the exact definition of the low estimator. Here we have chosen the set $\mathcal{X}_2$ to be a single branch; then we average over all b possibilities. This leads to the iteration scheme

$$\hat{v}_m^{\mathbf{j}} = h_m(\mathbf{S}_m^{\mathbf{j}}) \tag{1.44}$$

$$\hat{H}_{il}^{\mathbf{j}-} = \frac{1}{b-1} \sum_{\substack{k=1 \\ k \neq l}}^{b} \hat{v}_{i+1}^{\mathbf{k}} \tag{1.45}$$

$$\hat{v}_{il}^{\mathbf{j}} = \begin{cases} h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) & \hat{H}_{il}^{\mathbf{j}-} \leq h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) \\ \hat{v}_{i+1}^{\mathbf{l}} & \hat{H}_{il}^{\mathbf{j}-} \geq h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) \end{cases} \tag{1.46}$$

$$\hat{v}_i^{\mathbf{j}} = \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il}^{\mathbf{j}}, \tag{1.47}$$

where again $\mathbf{k} = \{\mathbf{j}, k\}$ and $\mathbf{l} = \{\mathbf{j}, l\}$ and the $-$ represents the negative bias of the hold value estimator.

The source of its low bias comes from a separation of the exercise decision from the value received upon continuation which leads to suboptimal decisions. A statement and rigorous proof of this is shown in Theorem 4 of [6] and restated here. Again we expand our notation for convenience and include the branching factor, b, as an argument for the estimator.

Theorem 3 *(Stochastic tree low estimator bias) The stochastic tree low estimator $\hat{v}_0(b, \mathbf{S}_0)$ is biased low. That is,*

$$\mathrm{E}\left[\hat{v}_0(b, \mathbf{S}_0)\right] \leq B_0(\mathbf{S}_0), \tag{1.48}$$

where $B_0(\mathbf{S}_0)$ is the true value.

As with the high biased estimator the low estimator is also consistent. The following theorem is stated and proven in Theorem 3 of [6],

Theorem 4 *(Stochastic tree low estimator consistency) The stochastic tree low estimator $\hat{v}_0(b, \mathbf{S}_0)$ converges to the true value $B_0(\mathbf{S}_0)$ in the p-norm as $b \rightarrow \infty$. That is,*

$$\| E\left[\hat{v}_0(b, \mathbf{S}_0)\right] - B_0(\mathbf{S}_0) \|_p \rightarrow 0. \tag{1.49}$$

A third estimator, the interleaving estimator, may also be considered. It is called the interleaving estimator because it is a blend of the high and low estimators. The implementation involes applying a backward induction estimate identical to that of the high estimator and then applying a sub-optimal stopping rule. However in this thesis we do not consider results generated by this estimator. For details on the implementation of the stochastic tree method refer to [6] for pseudo code describing a *Depth-First Processing* approach to the problem.

1.3.5   Bias Reduced Stochastic Tree

As an extension of the stochastic tree described in the previous section the authors of [19] propose a statistical method for reducing price estimator bias. An approximation for estimator bias is derived using large sample theory which is easily evaluated in a simulation. Subtracting the bias approximation results in significantly improved bias-corrected estimators at each exercise opportunity. This unique approach to bias correction is quite useful in that it; retains the favourable properties of the stochastic tree method; it is relatively simple to implement by modifying a few lines of code in existing algorithms; it does not increase the computational time; and it does not increase estimator variance.

The numerical results displayed in Section 3 of [19] show that the corrected estimators are always better than the uncorrected versions across all combinations of option moneyness (strike to spot price ratio), number of exercise opportunities and sample size. In addition it is shown that, for a desired level of accuracy, the corrected estimators can be computed much faster than the uncorrected estimators, and for a given amount of computing time, the corrected estimators are much more accurate than the uncorrected estimators. Also, by allowing for a better tradeoff between decreasing branching factor in exchange for an increased number of runs, the technique permits increased computational efficiencies over trivial parallel implementations of existing algorithms.

In Section 2 of [19] the authors derive expressions for the time-t_i bias of the high- and low-biased estimators. Their expressions are in the form of $\mathcal{F}_i$-conditional expectations. By letting $\bar{H}_i^{j\pm} = \mathrm{E}[\hat{H}_i^{j\pm}|\mathcal{F}_i]$, where the $\pm$ represents the postive and negative bias estimators respectively, they define the time-t_i high- and low-bias as $\bar{H}_i^{j+} - H_i^j = \mathrm{E}[\hat{V}_{i+1}^j - B_{i+1}^j|\mathcal{F}_i]$ and $\bar{H}_i^{j-} - H_i^j = \mathrm{E}[\hat{v}_{i+1}^j - B_{i+1}^j|\mathcal{F}_i]$, respectively. Here we have suppressed the l index for the low estimator because $\hat{H}_{il}^{j-}$, $l = 1, \ldots, b$ are iid given $\mathcal{F}_i$ and so these statements hold regardless of the choice of l. In what follows we will continue this practice when it is unambiguous to do so.

Consider the high-biased estimator, expanding the inner terms and then adding and subtracting $\mathrm{E}[\max(\bar{H}_{i+1}^{j+}, h_{i+1}^j)|\mathcal{F}_i]$ inside the time-t_i bias definition splits this expression into what the authors refer to as a local (Equation 1.50) and a global (Equation 1.51) component, namely,

$$\mathrm{E}\left[\max\left(\hat{H}_{i+1}^{j+}, h_{i+1}^j\right) - \max\left(\bar{H}_{i+1}^{j+}, h_{i+1}^j\right)|\mathcal{F}_i\right] \tag{1.50}$$

$$+ \mathrm{E}\left[\max\left(\bar{H}_{i+1}^{j+}, h_{i+1}^j\right) - \max\left(H_{i+1}^j, h_{j+1}^j\right)|\mathcal{F}_i\right]. \tag{1.51}$$

Focusing on the local component the authors show that the main source of bias is due to exercising incorrectly. This implies that significant contributions are limited to the region about the exercise boundary as even poor estimators are unlikely to result in incorrect exercise away from the boundary. Similar analysis in [19] shows that the low-biased estimator may also be split into a local and global component and that the local component for the low-estimator arises from exercising incorrectly. They summarize these contributions in tables which we reproduce here as Tables 1.1 and 1.2, respectively. Here, $\hat{Y}_i^{j\pm} = \hat{H}_i^{j\pm} - h_i^j$ and $\bar{Y}_i^{j\pm} = \bar{H}_i^{j\pm} - h_i^j$. From these results it is also evident why the estimators have their respective biases.

Equations involving $\bar{Y}_i^{j\pm}$ are not of much use from a numerical perspective as they are not directly observable and replacing them with the estimators $\hat{Y}_j^{j\pm}$ immediately collapses the bias expressions to zero. It is therefore necessary to incorporate additional knowledge of the distribution. The time-t_{i+1} hold value estimators are averages of the time-t_{i+2} option value estimators which, for the stochastic tree, are independent and identically distributed conditional

	Held: $\hat{Y}^{\mathbf{j}+}_{i+1} > 0$	Exercised: $\hat{Y}^{\mathbf{j}+}_{i+1} \leq 0$
Should Hold: $\bar{Y}^{\mathbf{j}+}_{i+1} > 0$	0	$-\hat{Y}^{\mathbf{j}+}_{i+1}$
Should Exercise: $\bar{Y}^{\mathbf{j}+}_{i+1} \leq 0$	$\hat{Y}^{\mathbf{j}+}_{i+1}$	0

Table 1.1: The local error in the time-t_{i+1} high-biased hold value estimator. Note that this error is always non-negative.

	Held: $\hat{Y}^{\mathbf{j}-}_{i+1} > 0$	Exercised: $\hat{Y}^{\mathbf{j}-}_{i+1} \leq 0$
Should Hold: $\bar{Y}^{\mathbf{j}-}_{i+1} > 0$	0	$-\bar{Y}^{\mathbf{j}-}_{i+1}$
Should Exercise: $\bar{Y}^{\mathbf{j}-}_{i+1} \leq 0$	$\bar{Y}^{\mathbf{j}-}_{i+1}$	0

Table 1.2: The local error in the time-t_{i+1} low-biased hold value estimator. Note that this error is always non-positive.

on $\mathcal{F}_{i+1}$.

The Central Limit Theorem is used to approximate the distribution of the estimators in deriving these bias approximations. Assuming that $\hat{Y}^{\mathbf{j}\pm}_{i}$ can be replaced with $\hat{Y}^{\mathbf{j}\pm*}_{i}$ which is normally distributed with mean $\bar{Y}^{\mathbf{j}\pm}_{i}$ and variances given by the conditional variances of the time-t_{i+2} option value estimators the approximate bias for the high and low biased estimators are

$$\left|\hat{Y}^{\mathbf{j}+*}_{i+1}\right| \Phi\left(\frac{-\left|\hat{Y}^{\mathbf{j}+*}_{i+1}\right|}{\sqrt{\hat{W}^{\mathbf{j}}_{i+1}/b}}\right) \tag{1.52}$$

and

$$\left|\hat{Y}^{\mathbf{j}-*}_{i+1}\right| \Phi\left(\frac{-\left|\hat{Y}^{\mathbf{j}-*}_{i+1}\right|}{\sqrt{\hat{w}^{\mathbf{j}}_{i+1}/b}}\right) - \sqrt{\frac{\hat{w}^{\mathbf{j}}_{i+1}}{b}} \,\phi\left(\frac{\left|\hat{Y}^{\mathbf{j}-*}_{i+1}\right|}{\sqrt{\hat{w}^{\mathbf{j}}_{i+1}/b}}\right), \tag{1.53}$$

respectively, where $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal cdf and pdf respectively and $\hat{W}_i^{\mathbf{j}}$ and $\hat{w}_i^{\mathbf{j}}$ are the sample variances given by

$$\hat{W}_i^{\mathbf{j}} = \frac{1}{b-1} \sum_{k=1}^{b} \left(\hat{V}_{i+1}^{\mathbf{k}} - \frac{1}{b} \sum_{p=1}^{b} \hat{V}_{i+1}^{\mathbf{p}} \right)^2 \tag{1.54}$$

and

$$\hat{w}_{il}^{\mathbf{j}} = \frac{1}{b-1} \sum_{\substack{k=1 \\ k \neq l}}^{b} \left(\hat{v}_{i+1}^{\mathbf{k}} - \frac{1}{b} \sum_{\substack{p=1 \\ p \neq l}}^{b} \hat{v}_{i+1}^{\mathbf{p}} \right)^2 , \tag{1.55}$$

with $\mathbf{k} = \{\mathbf{j}, k\}$ and $\mathbf{p} = \{\mathbf{j}, p\}$.

Then replacing the quantities $\bar{Y}_i^{\mathbf{j}\pm*}$ with their sample estimates $\hat{Y}_i^{\mathbf{j}\pm*}$ Equations 1.52 and 1.53 are subtracted from their corresponding uncorrected estimators yielding

$$\hat{V}_i^{\mathbf{j}} = \max\left\{ \hat{H}_i^{\mathbf{j}+}, h_i^{\mathbf{j}} \right\} - \left| \hat{Y}_i^{\mathbf{j}+} \right| \Phi\left(\frac{-\left| \hat{Y}_i^{\mathbf{j}+} \right|}{\sqrt{\hat{W}_i^{\mathbf{j}}/b}} \right) \tag{1.56}$$

and

$$\hat{v}_i^{\mathbf{j}} = \frac{1}{b} \sum_{l=1}^{b} \left(\begin{cases} h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) & \hat{H}_{il}^{\mathbf{j}-} \leq h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) \\ \hat{v}_{i+1}^{\mathbf{l}} & \hat{H}_{il}^{\mathbf{j}-} \geq h_i\left(\mathbf{S}_i^{\mathbf{j}}\right) \end{cases} \right.$$
$$\left. - \left| \hat{Y}_{i,l}^{\mathbf{l}-} \right| \Phi\left(\frac{-\left| \hat{Y}_{i,l}^{\mathbf{l}-} \right|}{\sqrt{\hat{w}_{i,l}^{\mathbf{l}}/b}} \right) - \sqrt{\frac{\hat{w}_{i,l}^{\mathbf{l}}}{b}} \, \phi\left(\frac{\left| \hat{Y}_{i,l}^{\mathbf{l}-} \right|}{\sqrt{\hat{w}_{i,l}^{\mathbf{l}-}/b}} \right) \right), \tag{1.57}$$

respectively, where $\mathbf{l} = \{\mathbf{j}, l\}$.

Using these bias-corrected estimators, the time-t_i bias becomes

$$\mathrm{E}\left[\max\left(\hat{H}_{j+1}^{\mathbf{j}+}, h_{i+1}^{\mathbf{j}} \right) - \max\left(\bar{H}_{i+1}^{\mathbf{j}+}, h_{i+1}^{\mathbf{j}} \right) \,\middle|\, \mathcal{F}_i \right] \tag{1.58}$$

$$+ \mathrm{E}\left[\max\left(\bar{H}_{i+1}^{\mathbf{j}+}, h_{i+1}^{\mathbf{j}} \right) - \max\left(H_{i+1}^{\mathbf{j}}, h_{i+1}^{\mathbf{j}} \right) \,\middle|\, \mathcal{F}_i \right] \tag{1.59}$$

$$- \mathrm{E}\left[\left| \hat{Y}_{i+1}^{\mathbf{j}+} \right| \Phi\left(\frac{-\left| \hat{Y}_{i+1}^{\mathbf{j}+} \right|}{\sqrt{\hat{W}_{i+1}^{\mathbf{j}}/b}} \right) \,\middle|\, \mathcal{F}_i \right] \tag{1.60}$$

for the high-biased version and

$$\mathrm{E}\left[I_{\hat{H}^{\mathbf{j}-}_{i+1} > h^{\mathbf{j}}_{i+1}} \tilde{H}^{\mathbf{j}-}_{i+1} + I_{\hat{H}^{\mathbf{j}-}_{i+1} \le h^{\mathbf{j}}_{i+1}} h^{\mathbf{j}}_{i+1} - \max(\bar{H}^{\mathbf{j}-}_{i+1}, h^{\mathbf{j}}_{i+1}) \,\Big|\, \mathcal{F}_i \right] \tag{1.61}$$

$$+ \mathrm{E}\left[\max\left(\bar{H}^{\mathbf{j}-}_{i+1}, h^{\mathbf{j}}_{i+1} \right) - \max\left(H^{\mathbf{j}}_{i+1}, h^{\mathbf{j}}_{i+1} \right) \,\Big|\, \mathcal{F}_i \right] \tag{1.62}$$

$$- \mathrm{E}\left[\left| \hat{Y}^{\mathbf{j}-}_{i+1} \right| \Phi\left(\frac{-\left| \hat{Y}^{\mathbf{j}-}_{i+1} \right|}{\sqrt{\hat{w}^{\mathbf{j}}_{i+1}/b}} \right) + \sqrt{\hat{w}^{\mathbf{j}}_{i+1}/b}\, \phi\left(\frac{\left| \hat{Y}^{\mathbf{j}}_{i+1} \right|}{\sqrt{\hat{w}^{\mathbf{j}-}_{i+1}/b}} \right) \,\Big|\, \mathcal{F}_i \right] \tag{1.63}$$

for the low-biased version, where Equations 1.58 and 1.61, 1.59 and 1.62, and 1.60 and 1.63 are the local bias, global bias and corrections components, respectively.

The local and correction components asymptotically cancel as the sample size gets large, leaving just the global components. Applying Jensen's inequality to move the absolute value inside the expectation and applying the inequality $|\max(x, y) - \max(u, v)| \le |x - u| + |y - v|$ to the absolute value of the global component gives

$$\left| \mathrm{E}\left[\max\left(\bar{H}^{\mathbf{j}\pm}_{i+1}, h^{\mathbf{j}}_{i+1} \right) - \max\left(H^{\mathbf{j}}_{i+1}, h^{\mathbf{j}}_{i+1} \right) |\mathcal{F}_i \right] \right|$$

$$\le \mathrm{E}\left[\left| \bar{H}^{\mathbf{j}\pm}_{i+1} - H^{\mathbf{j}}_{i+1} \right| |\mathcal{F}_i \right] = \mathrm{E}\left[\left| \mathrm{E}\left[\hat{H}^{\mathbf{j}\pm}_{i+1} |\mathcal{F}_{i+1} \right] - H^{\mathbf{j}}_{i+1} \right| |\mathcal{F}_i \right],$$

which shows it to be bound by the time-t_{i+1} bias. Similarly the time-t_{i+1} bias is bound by the time-t_{i+2} bias. Continue in this fashion through to the next-to-last exercise opportunity $m-1$ and note that the time-t_{m-1} hold-value estimator is unbiased. Thus, the global bias is also accounted for. Specifically, the propagation of bias across exercise opportunities is at most of the same order as the difference between the local bias and the correction component.

1.3.6   Stochastic Mesh Method

The stochastic mesh is a method for American option pricing and more general optimal stopping time problems. The mesh construction begins with the initial state vector $\mathbf{S}_0$ and then the state transition function is applied to produce state vectors $\mathbf{S}^j_i$ for $j = 1, \ldots, b$ and $i = 1, \ldots, m$, where b is the chosen mesh size. This can be seen illustratively on the left in Figure 1.3. The individual mesh points $\mathbf{S}^j_i$ and $\mathbf{S}^k_{i+1}$ are then tied together by a set of weights $\omega(i, \mathbf{S}^j_i, \mathbf{S}^k_{i+1}) = \omega^{jk}_i$, which are defined below. This is illustrated on the right of Figure 1.3. Following the description given in [22] Section 8.5.2 we now describe how the mesh is generated and the weights are determined.

Suppose our state space to be a Markov chain $(\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_m)$ in $\mathbb{R}^d$ with transition densities
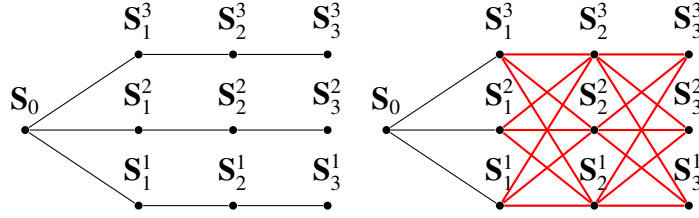
Figure 1.3: Stochastic Mesh with 3 timesteps and $b = 3$

$f_1, \ldots, f_m$. This is to say that for $\mathbf{s} \in \mathbb{R}^d$ and for $A \subseteq \mathbb{R}^d$, we have that

$$P\{\mathbf{S}_i \in A | \mathbf{S}_{i-1} = \mathbf{s}\} = \int_A f_i(\mathbf{s}, \mathbf{y}) d\mathbf{y}, \tag{1.64}$$

for $i = 1, \ldots, m$.

Now consider the marginal density of some $\mathbf{S}_i$ in our chain. For fixed $\mathbf{S}_0$ we have the marginal density for $\mathbf{S}_1$, $g_1(\cdot) = f_1(\mathbf{S}_0, \cdot)$. The marginal densities for other $\mathbf{S}_i$, $i = 2, \ldots, m$ are given by,

$$g_i(\mathbf{y}) = \int g_{i-1}(\mathbf{s}) f_i(\mathbf{s}, \mathbf{y}) d\mathbf{s}. \tag{1.65}$$

The true hold value, $H_i(\mathbf{s})$, for an option in state $\mathbf{s}$ at time t_i is

$$H_i(\mathbf{s}) = E\left[B_{i+1}(\mathbf{S}_{i+1}) | \mathbf{S}_i = \mathbf{s}\right] = \int B_{i+1}(\mathbf{y}) f_{i+1}(\mathbf{s}, \mathbf{y}) d\mathbf{y}, \tag{1.66}$$

and in what follows it is shown that the purpose of the weights as mentioned is to correct for the discrepancy in how the future time nodes were generated. There are several possible ways that one could imagine for how these nodes are generated and each way could possibly have a different weight function associated with it. We now explore some of these possibilities in order to justify the choice used in [7] and this thesis.

First, suppose that each $\mathbf{S}_{i+1}^k$ for $k = 1, \ldots, b$ is generated independently from all others, with density g. If we also suppose that we know all the true option values, $B_{i+1}(\mathbf{S}_{i+1}^k)$ and that g is the marginal density from our Markov chain then the average over the true option values as $b \to \infty$ is

$$\frac{1}{b} \sum_{k=1}^{b} B_{i+1}(\mathbf{S}_{i+1}^k) \to E\left[B_{i+1}(\mathbf{S}_{i+1}^k)\right]$$
$$= \int B_{i+1}(\mathbf{y}) g(\mathbf{y}) d\mathbf{y},$$

which is not in general equal to 1.66 and is in fact the unconditional expectation at t_{i+1} instead

of the desired conditional expectation.

Here we see the meaning of our previous statement that the weights are meant to correct for the fact that the mesh nodes were generated from density g instead of density f. If we take the weight function to be

$$\omega_i^{jk} = \frac{f_{i+1}(\mathbf{S}_i^j, \mathbf{S}_{i+1}^k)}{g(\mathbf{S}_{i+1}^k)}, \tag{1.67}$$

which is the likelihood ratio connecting the transition density to the mesh density, then we have that

$$\frac{1}{b} \sum_{k=1}^{b} \omega_i^{jk} B_{i+1}(\mathbf{S}_{i+1}^k) \rightarrow \mathrm{E}\left[\omega_i^{jk} B_{i+1}(\mathbf{S}_{i+1}^k)\right]$$
$$= \int \frac{f_{i+1}(\mathbf{S}_i^j, \mathbf{y})}{g(\mathbf{y})} V_{i+1}(\mathbf{y}) g(\mathbf{y}) d\mathbf{y}$$
$$= \int V_{i+1}(\mathbf{y}) f_{i+1}(\mathbf{S}_i^j, \mathbf{y}) d\mathbf{y}$$
$$= H_i(\mathbf{S}_i^j).$$

Above we had assumed that all of the nodes were generated independently, however in our case this is not true. Even so there is still some flexibility in the choice of weights and construction method.

First we begin by considering the independent path method shown in left panel of Figure 1.3. Here sequences of asset values $(\mathbf{S}_1^j, \dots, \mathbf{S}_m^j)$ for $j = 1, \dots, b$ are independent of one another. Hence for $k \neq j$, $\mathbf{S}_{i+1}^k$ is independent of $\mathbf{S}_i^j$ and therefore the conditional distribution of $\mathbf{S}_{i+1}^k$ given $\mathbf{S}_i^j$ is just the unconditional distribution which has density g_{i+1}. In the case of $k = j$ the distribution has the density of the transitional density f_{i+1} and so no weight is needed. This means that

$$\omega_i^{jk} = \begin{cases} \frac{f_{i+1}(\mathbf{S}_i^j, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^k)} & \text{if } k \neq j \\ 1 & \text{if } k = j \end{cases} \tag{1.68}$$

Another mesh construction that may be considered consists of generating nodes at time-t_{i+1} by randomly selecting nodes at time-t_i, with replacements, from which to produce the successor by sampling from the transition density $f_{i+1}(\mathbf{S}_i^j, \cdot)$. This procedure is repeated until all b nodes are filled. Then given, $\mathbf{S}_i$, the nodes at time-t_{i+1} are i.i.d. with density equal to the average of the transition densities

$$\frac{1}{b} \sum_{j=1}^{b} f_{i+1}(\mathbf{S}_i^j, \cdot), \tag{1.69}$$

since each time-t_i node is equally likely to be chosen. This gives a weight function of

$$\omega_i^{jk} = \frac{f_{i+1}(\mathbf{S}_i^j, \mathbf{S}_{i+1}^k)}{\frac{1}{b} \sum_{l=1}^b f_{i+1}(\mathbf{S}_i^l, \mathbf{S}_{i+1}^k)}. \tag{1.70}$$

Further consideration of this construction method reveals that if we make the modification to use stratified sampling over the index j of $\mathbf{S}_i^j$, we choose exactly b strata and we draw exactly one from each stratum then this is exactly the same as the independent path method. As such the weights in Equation 1.70 are equally applicable to the independent path method.

The authors of [7] make their choice of which mesh construction method and weight function based on considering the pricing of a European option with the stochastic mesh. It is important to note that there is no reason to use this method to price a European option, however, what follows is taken as evidence that their particular choices are beneficial when carried over to the case of pricing American option.

When considering the pricing of a European option with the mesh, the lack of early exercise rights leads to the exercise decision at the time prior to expiry collapsing to just the hold value at that node in the mesh. It can be easily shown that this leads to a option price given by

$$\hat{V}_0 = \frac{1}{b^m} \sum_{j_1,\dots,j_m} \prod_{i=2}^m \omega_{j_{i-1}j_i}^{i-1} h_m(\mathbf{S}_m^i), \tag{1.71}$$

where $j_i = 1,\dots,b$, and $i = 1,\dots,m$. This is simply the average over all b^m paths through the mesh of the payoffs at expiry multiplied by the weights of each of the arcs in those paths. Rearranging gives

$$\hat{V}_0 = \frac{1}{b} \sum_{j_m=1}^b h_m(\mathbf{S}_m^{j_m}) \left(\frac{1}{b^{m-1}} \sum_{j_1,\dots,j_{m-1}} \prod_{i=2}^m \omega_{j_{i-1}j_i}^{i-1} \right). \tag{1.72}$$

The term in parenthesis has an expected value of 1 when the weight functions are given by likelihood ratios. Rewriting that term gives

$$\frac{1}{b} \sum_{j_{m-1}}^b \omega_{m-1}^{j_{m-1},j_m} \cdots \frac{1}{b} \sum_{j_1}^b \omega_{m-1}^{j_1,j_2}. \tag{1.73}$$

Returning to Equation 1.70 it is easily seen that

$$\frac{1}{b} \sum_j^b \omega_i^{j,k} = 1. \tag{1.74}$$

Therefore choosing weights given by (1.70) leads automatically to the term in the parenthesis

of Equation 1.70 being identically 1 and a price for the European option being given by

$$\hat{V}_0 = \frac{1}{b} \sum_{j_m=1}^{b} h_m(\mathbf{S}_m^{j_m}),\tag{1.75}$$

as expected.

With other choices of weights, including those shown above, the term in question is not necessarily identically equal to 1 and therefore an exponentially growing variance build up may occur along a trajectory through the mesh as weights are multiplied together.

We may now define the mesh estimator, $\hat{V}_0$, as a positively biased estimate of the option determined by moving backward in time through the mesh beginning at expiry. Using dynamic programming the recursive relation for valuation is

$$\hat{V}_m^j = h_m(\mathbf{S}_m^j)\tag{1.76}$$

$$\hat{H}_i^j = \frac{1}{b} \sum_{k=1}^{b} \omega_i^{jk} \hat{V}_{i+1}^k\tag{1.77}$$

$$\hat{V}_i^j = \max\left(h_i(\mathbf{S}_i^j), \hat{H}_i^j\right),\tag{1.78}$$

$$i = m - 1, \ldots, 0,\tag{1.79}$$

where we have neglected a discounting factor without loss of generality.

With this we now have a complete description of the algorithm for the mesh estimator. In terms of computational intensity the effort involved in generating the state vectors is proportional to $b \times m$. The effort for evaluating the mesh via Equations (1.76)– (1.79) is $b^2 \times m$. Hence the algorithm is quadratic in the mesh parameter b and linear in the number of exercise opportunities.

There are two key properties of the mesh estimator which we now discuss. The first regards the bias of the estimator described in Equations (1.76)– (1.79). We restate here for convenience Theorem 1 from Broadie and Glasserman, 2004,

Theorem 5 *(Mesh estimator bias) The mesh estimator $\hat{V}_0(b, \mathbf{S}_0)$ is biased high. That is,*

$$E\left[\hat{V}_0(b, \mathbf{S}_0)\right] \geq B_0(\mathbf{S}_0)$$

for all b.

For a proof of this theorem refer to [7]. Intuitively the high bias of the estimator can be reasoned by the fact that it unfairly looks into the future when making the exercise decision.

This conditioning on future information allows for better decisions than otherwise possible and hence this estimator is greater then the true value on average.

The second property regards the convergence of the mesh estimator. For a rigorous proof and assumptions on the density $f_i(\mathbf{S}_i, \mathbf{S}_{i+1})$ we again refer the reader to Broadie and Glasserman, 2004. For convenience we restate Theorem 2 of that paper here.

Theorem 6 *(Mesh estimator convergence) Let $\tilde{p} > p > 1$. Under the assumptions for $f_i(\mathbf{S}_i, \mathbf{S}_{i+1})$,*

$$\| \hat{V}_i(b, \mathbf{s}) - B_i(\mathbf{s}) \| \to 0$$

as $b \to \infty$, for all $\mathbf{s}$ and i.

Here $\| X \| = \mathrm{E}[|X|^p]^{1/p}$ is the p-norm and $\tilde{p}$ defines moment conditions given in [7].

This convergence in the p-norm implies a convergence in probability of the mesh estimator to the true value and hence $\hat{V}_0(b, \mathbf{S}_0)$ is a consistent estimator of the option value.

The remaining piece to the stochastic mesh is a second estimator that is biased low, called the path estimator. The path estimator is constructed by simulating a path $\{\mathbf{X}_0, \mathbf{X}_1, \ldots, \mathbf{X}_m\}$, which is independent of the mesh points $\mathbf{S}_i^j$, according to the density function $f_i(\mathbf{x}, \cdot)$. The mesh is used to estimate the path estimator stopping rule. Along the path the optimal policy chooses to exercise the option at time $\hat{\tau} = \min\{i : h_i(\mathbf{X}_i) \geq \hat{H}_i(\mathbf{X}_i)\}$ giving a low estimate of the option value to be $\hat{v} = h_{\hat{\tau}}(\mathbf{X}_{\hat{\tau}})$. An average low estimator which is conditional on the mesh is calculated by simulating multiple paths independently each of which follows the stopping rule defined by the mesh.

To evaluate the path estimator we can use the following dynamic programming scheme

$$\hat{v}_m^j = h_m(\mathbf{X}_m^j) \tag{1.80}$$

$$\hat{H}_i^j = \frac{1}{b} \sum_{k=1}^{b} \omega_i^{jk} \hat{V}_{i+1}^k \tag{1.81}$$

$$\hat{v}_i^j = \begin{cases} \hat{v}_{i+1}^j & \text{if } h_i(\mathbf{X}_i^j) \leq \hat{H}_i^j \\ h_i(\mathbf{X}_i^j) & \text{if } h_i(\mathbf{X}_i^j) > \hat{H}_i^j \end{cases} \tag{1.82}$$

$$i = 0, \ldots, m - 1 \tag{1.83}$$

for $j = 1, \ldots, n_p$, where n_p is the number of paths, $\hat{V}_t^i$ is the mesh estimator described in Equations (1.76)– (1.79) and $\omega_i^{jk} = \omega_i(\mathbf{X}_i^j, \mathbf{S}_{i+1}^k)$ is the same weight function previously described with $\mathbf{S}_i$ being the set of time-t_i mesh points. Again we have suppressed the discounting factor without a loss of generality.

To see how to extend the weights, ω_i^{jk}, to this case consider the weight along a path arc from state $\mathbf{x}$ to mesh node $\mathbf{S}_{i+1}^k$ denoted as $\omega_i^k(\mathbf{x})$. Take $\mathbf{x} \in \mathbb{R}^d$ but $\mathbf{x} \notin A_i$ as in the description of the path estimator. Noting that in all choices of mesh construction and weight choices the current node $\mathbf{X}_i^j$ appears as an explicit argument of the transition density. Therefore we are free to extend the arguments to any arbitrary points in $\mathbb{R}^d$.

As in the previous section we now, for convenience, state two theorems from [7] regarding the path estimator; in that publication Theorem 3 refers to the bias of the path estimator and Theorem 4 refers to its consistency.

Theorem 7 *(Path estimator bias) The path estimator $\hat{v}_i(b, \mathbf{X}_0)$ is biased low. That is,*

$$E[\hat{v}_0(b, \mathbf{X}_0)] \le B_0(\mathbf{X}_0)$$

for all b.

Intuitively the source of the low bias in this estimator can be seen from equation (1.83). The information used in determining whether to exercise or hold the option is independent of that which is propagated in the event of a decision to exercise. Hence a suboptimal decision is expected and therefore a lower valuation for the option is expected because no other policy can be better then the optimal policy.

Theorem 8 *(Path estimator convergence) Suppose the same condition hold as in Theorem 6 and that $E[h_i(S_i)^{1+\epsilon}] < \infty$ for all $i = 0, \ldots, m$, for some $\epsilon > 0$. Suppose also that $P(h_i(S_i) = B_i(S_i)) = 0$ for all $i = 0, \ldots, m-1$. Then*

$$E[\hat{v}_0(b, \mathbf{X}_0)] \to B_0(\mathbf{X}_0)$$

as $b \to \infty$.

So the path estimator is asymptotically unbiased.

The computational effort for the path estimator in addition to the mesh estimator is $n_p \times m$ to generate the paths and $n_p \times b^2 \times m$ to evaluate the dynamic programing scheme in Equation (1.83).

In comparison with regression methods the mesh requires $O(b^2)$ whereas the regression based methods are $O(b)$. However the mesh will not suffer from approximation error from the choice of regression basis functions as previously mentioned. In addition the likelihood ratio method for calculating mesh weights is independent of the payoff of the option and therefore the weights may be saved and used to price many different options unlike the regression based methods which require a recalculation of the regression variables with each option being priced.

Above we have outlined the computational cost associated with both the mesh and path estimators. Another limitation in using the mesh is that, as described above, there is a need for a transition density when determining the weights due to the choice of using likelihood ratios. If the transition densities for the underlying Markov chain are unknown or are not easily computable then it may be necessary to use approximations to easily computable densities in order to generate the likelihood ratios.

Although cases where the transition densities are unknown or non-existent are not considered in this thesis we briefly overview how using a constrained optimization problem can be used to determine the weights as shown in [20].

Suppose that we have some function G on the state space of the underlying Markov chain. The conditional expectation

$$g(\mathbf{s}) = \mathrm{E}\left[G(\mathbf{S}_{i+1})|\mathbf{S}_i = \mathbf{s}\right] \tag{1.84}$$

is a known function of the state $\mathbf{s}$. If we fix a node $\mathbf{S}_i^j$ and consider weights ω_i^{jk}, $k = 1, \ldots, b$ satisfying

$$\frac{1}{b}\sum_{k=1}^{b}\omega_i^{jk}G\left(\mathbf{S}_{i+1}^k\right) = g\left(\mathbf{S}_i^j\right), \tag{1.85}$$

then we can see that these are the weights that fix the pricing of the payoff G, to be received at t_{i+1}, from the perspective of node $\mathbf{S}_i^j$. If the option estimates B_{i+1} are well approximated by a linear combination of components of G, then such weights should provide a reasonable approximation to the continuation value at $\mathbf{S}_i^j$.

From all feasible weights the authors of [20] select those that minimize an objective function of the form

$$\sum_{k=1}^{b}C\left(\omega_i^{jk}\right), \tag{1.86}$$

for any convex function C. A draw back of this method is that different choices of C will result in differing sets of weights and therefore differing option values. In [20] it is shown that either $C(x) = x^2/2$ or $C(x) = x\log x$ result in similar price estimates to those produced when using likelihood ratio weights. They also show that using the quadratic function does lead to a simpler optimization, however it could produce negative weights whereas the entropy objective will provide non-negative results if they exist.

Because this method leads to weights that are no longer independent of the payoff function for the option this method forces us to recalculate the weights for each new type of option we wish to price. It is shown in [21] that this approach for generating mesh weights through a constrained optimization problem is closely related to the regression based methods for pricing American options [8, 9]. The difference being that here the method produces weights that

depend on $\mathbf{S}_i^j$ and $\mathbf{S}_{i+1}^k$ whereas with the regression methods the weights depend on $\mathbf{S}_i^j$ and $\mathbf{S}_i^k$.

Chapter 2

Introduction to Multiple Exercise Options

This chapter begins our discussion of the focus of thesis, valuation of multiple exercise options. It includes a detailed description of multiple exercise options along with a summary of previous methods to value them. It finishes with a detailed look at one particular method, the Forest of Trees [23].

2.1   Examples of Multiple Exercise Options

Multiple exercise right options are generalizations of American-style options as they provide the holder more than one exercise right and sometimes control over one or more other variables, such as the amount exercised. Examples of financial derivatives and real options with these properties have become more prevalent over the past decade and appear in sectors ranging from insurance to energy industries. We describe some of these instruments here.

Tolling agreements are a real options example of multiple exercise options from the energy industry. Power markets have extremely high price volatilities making participants as well as potential lenders to producers wary of price risk. Other difficulties facing power producers include the facts that electricity is not a traded asset, generation costs depend on the price of the underlying physical asset(s), potential emissions costs, and power transmission to consumers. A tolling agreement is a structured transaction between a plant owner and typically a financial institution (FI). In the agreement, for an up front premium paid to the plant owner, the buyer is given the ability to determine when the plant is in operation and to take the electricity output at a predetermined price. Along with the operational constraints inherent to the plant the buyer is also often constrained by other contractual obligations. This agreement divides the risks associated with power generation into operational risks associated with the plant and market risk associated with the fluctuating prices of the plant's output. It then leaves the management of each risk to the party involved which is best suited to handle them. The plant operator is guaranteed a stable cash flow from a FI with a typically higher credit rating which in turn makes it easier for the plant to raise capital to fund expansions and improvements, for example. The buyer of the contract uses the agreement to financially replicate the operation of a power plant without having to take on the associated operational risk.

In interest rate markets a chooser flexible cap is a multi-exercise interest rate derivative. These are a variant of an over the counter interest rate cap agreement. An interest rate cap is a derivative consisting of a series of caplets in which the i^{th} caplet provides the holder at time t_i with a payment of the notional multiplied by the difference in the current interest rate and a fixed strike rate, if positive, over $\mathcal{N}$ time intervals. These derivatives protect the holder, who may for example be an issuer of floating rate debt, from rises in the short term interest rates in exchange for an up front premium. A chooser flexible cap is almost identical to a interest rate cap except it only allows for $\tilde{\mathcal{N}} < \mathcal{N}$ caplets to be exercised at the holders discretion. The holder must decide how to optimally allocate the $\tilde{\mathcal{N}}$ exercise rights across the $\mathcal{N}$ opportunities. This multiple exercise option was valued in [24].

A Segregated Fund is a type of investment fund administered by insurance companies in the form of variable life insurance contracts offered to individuals with certain guarantees such as

reimbursement of capital upon death. Like mutual funds, segregated funds consist of a pool of investments in securities such as bonds and stocks. The value of the segregated fund varies with the market value of the underlying securities. Segregated funds do not issue units or shares, instead, the investor is the holder of a segregated fund contract. The policy holder is given a number of funds (possibly 10 or more) from which to pick.

The maturity date for a segregated fund contract is the time horizon at which there is a guarantee to retain some predefined level (typically 75 percent) of initial capital. Holding periods to reach contract maturity are usually 10 or more years. A death benefit provides the guarantee if death occurs prior to the time horizon. Many contracts have a reset option allowing the contract holder to lock in investment gains if the market value of the underlying fund increases. This resets the contract's deposit value to the greater of the deposit value or current market value and extends the maturity date. Contract holders are limited to a certain number of resets during the lifetime of the contract and at the same time the holder may switch between underlying investment funds. As of 2009 Canadians owned about $140-billion worth of these funds.

Another multiple exercise option example is a swing option. Swing options have typically been used in energy markets to help producers manage the raw materials used in energy production in the face of uncertain demand. They allow the holder to have some control over both the timing and delivery amount of the underlying asset at predetermined prices. A swing option is typically noted as the swing portion of a base-loaded futures contract that gives a predetermined price for an amount of a commodity over a given period of time. The swing part of this overall contract allows for a variable delivery amount of the underlying above or below the amount determined by the base-loaded contract. However, the two pieces of the contract can be detached from one another and treated individually for valuation purposes. During the length of the contract the holder may exercise a given number swing rights. Typically these rights can only be exercised at a discrete set of times. Upon swinging, the holder can choose the additional volume bought/sold of the underlying asset. Refer to [26] for a more detailed discussion on swing options.

2.2 Multiple Exercise Options

As with the pricing of American-style options the valuation of multiple exercise options is a problem in stochastic optimal control. For American-style options the solution provides both a value and optimizing exercise rule, or stopping time. For multiple exercise options the solution also gives both a value and optimizing exercise policy. In the case in which the holder controls only the exercise times (e.g., chooser flexible cap) the exercise policy is a sequence of stopping

times. For options in which the holder controls the exercise times and amounts (e.g., swing options) the exercise policy is a paired sequence of stopping times and exercise amounts. The policy generalizes to other control variables.

Multiple exercise option-valuation algorithms are generalizations of those used for pricing American-style options. Most of the work in the literature has focused on swing options. As such, the remainder of this thesis will focus specifically on the valuation of swing options.

There are continuous-time solutions to both the American-style and swing option valuation problems; these are computed by solving a system of Hamilton-Jacobi-Bellman quasi-variational inequalities [27]. These methods give more accurate and stable price and sensitivity estimates than those computed using simpler tools (e.g., trees). However, these methods are quite complex mathematically and break down in higher dimensions financial examples of which we consider to be the inclusion of multiple underlying assets and associated risk factors.

Approaches that use dynamic programming to solve a time-discretized version of the optimal control problem are less complex, easier to implement and hence more prevalent than the continuous-time methods described above. For American-style options standard methods include tree-based algorithms and numerical solutions of the associated partial differential equation which require discretization of the state space (coarse in the case of tree based algorithms). The analogous method for swing option valuation is given by the forest of trees [23, 26]. These methods are extensions of the corresponding methods for pricing American-style options and have similar properties. One crucial property is that these methods fail as the dimensionality of the problem increases.

The simulation approach is the obvious tool to overcome the curse of dimensionality, as the rate of convergence of Monte Carlo estimators is independent of the dimension. As previously mentioned Tilley [5] was the first to show that the forward-in-time Monte Carlo approach could be used to solve the backward-in-time dynamic programming problem arising from valuation of an American-style option. Since this seminal paper, numerous other methods for the Monte Carlo valuation of American style options have appeared. These include methods that attempt to parameterize the exercise region [28] and those that discretize the state space [29]. Both of these approaches, however, also suffer from the curse of dimensionality and do not easily generalize to arbitrary payoffs and underlying price processes.

Monte Carlo methods that do not break down with the dimensionality and that accommodate general payoff and price processes include those that solve the optimal stopping-time problem through estimation of the hold or continuation value. These include the stochastic tree and mesh techniques of Broadie and Glasserman [6, 7] and the regression-based approach first appearing in [8] and the subsequently generalized in [9]. For each of these valuation techniques, high- and low-biased estimators are easily generated, along with a hybrid interleaving

estimator that has properties of both. Duality-based methods solve the optimal control problem in the dual space, by approximating an optimal martingale, typically by regression [12, 30].

Methods that parameterize the early-exercise region have been extended to value swing options by parameterizing the set of exercise level curves [31]. Similarly state space aggregation methods have been used for swing option valuation [29]. As in American option valuation, however, these methods deteriorate as the dimension increases and do not easily generalize to arbitrary payoffs and price processes.

The Least-squares Monte Carlo and duality methods have been modified for the pricing of swing options in [24, 32], respectively. Although increased dimensionality does not decrease the performance of these methods, they suffer from other drawbacks. In Least-squares Monte Carlo methods one must select a set of basis functions on which to run regressions to estimate continuation values. In general only a complete (infinite) set of basis functions results in continuation value estimators that are consistent for the true option value. In practice, of course, a finite set of basis functions is used and introduces an approximation error. Continuation value estimators are consistent for the true approximation value and not the true option value [10, 11]. This approximation error can propagate backwards through the exercise opportunities and produce high- and low-biased estimators that do not converge to the same value [33].

Duality-based methods typically use regression on a finite set of basis functions to approximate the optimal martingale, implying similar issues as Least-squares Monte Carlo. These issues persist in extensions of these algorithms to the pricing of swing options. Policy iteration methods such as [34], yield approximations of the time-0 value at each iteration of the dynamic program. As with the pricing of American-style options this method is advantageous because it removes the requirement to calculate nested conditional expectations prior to the time-0 value being approximated.

In this thesis we propose two new methods, the Forest of Stochastic Trees and the Forest of Stochastic Meshes, for valuing high-dimensional swing options. The trees used to model the underlying price process in the forest of trees algorithm [23], are replaced with stochastic trees [6] and meshes [7]. High and low-biased estimators are constructed which are consistent for the true option value. Furthermore the methods can accommodate general price processes and payoffs, multiple risk factors, and a confidence interval for the true value can easily be constructed. In the case of the forest of stochastic trees the method is limited to only a modest number of exercise opportunities due to the exponential scaling of the computational effort described in Chapter 3. The forest of stochastic meshes described in Chapter 4 of this thesis avoids this exponential scaling and allows for a large number of exercise opportunities.

2.2.1 Mathematical Description of Swing Options

The type of contract that will be the primary focus of this thesis is a swing option. A swing option is typically more correctly noted as a swing portion of a base-loaded futures contract that gives a predetermined price for an amount of a commodity over a given period of time. The swing part of this overall contract allows for a variable delivery amount of the underlying above or below the amount determined by the base-loaded contract. However the two pieces of the contract can be detached from one another and treated individually.

Take 0 to be the time the contract is signed and suppose the option is in effect for time $t_i \in [T_1, T_2]$, where $0 \le T_1 < T_2$. During the length of the contract the holder may exercise a given number of up and down swing rights denoted $\mathcal{N}_u$ and $\mathcal{N}_d$ respectively. Typically these rights can only be exercised at discrete set of times $\{t_1, \ldots, t_m\}$ with $T_1 \le t_1 < \ldots < t_m \le T_2$. Also the contract may stipulate a holding time of Δt between successive swings which if imposed would be greater than $t_{i+1} - t_i$ for all $1 \le i \le m - 1$. In subsequent sections we do not include this feature, but it can be incorporated into our new valuation algorithms in a straightforward manner.

In general the act of exercising a swing right can have two different effects on the overall contract which are broken down into the categories of Local and Global effects. Local effects change the delivery volume only on the date of the exercise and then for any subsequent deliveries the volume reverts back to the original base-loaded volume. Global effect change the delivery volume on the exercise date and all remaining deliveries unless another swing right is exercised. In this thesis we consider examples of swing options having only local effects, though we could also accommodate exercise rights having global effects.

When the holder chooses to swing up or down they may also have a choice of exercise volumes. These amounts may be continuous or discrete but in either case the volumes at a given opportunity at τ_i will take the form $[u_i^1, u_i^2] \cup [u_i^3, u_i^4]$ for $1 \le i \le m$ and $u_i^1 \le u_i^2 \le 0 \le u_i^3 \le u_i^4$.

Another feature that is included in some contracts are penalties which restrict the total volume which may be exercised during the contract. We define the usage level, U, as the net swing volume above or below the base-loaded volume stipulated by the forward part of the contract. A contract will then stipulate that U must be in a range $[U_{min}, U_{max}]$ for there to be no penalties incurred at the completion of the contract where U_{min} and U_{max} are agreed in the contract. The penalty function therefore takes the form

$$\tilde{\phi}(U) = \begin{cases} \mathcal{P}_1 & \text{if } U(T_2) < U_{min} \\ 0 & \text{if } U_{min} \le U(T_2) \le U_{max} \\ \mathcal{P}_2 & \text{if } U(T_2) > U_{max} \end{cases} \tag{2.1}$$

The final aspect of swing options is that of the strike price which specifies the price at which the underlying asset may be bought or sold. Swing options also include this, but because of the ability to both buy or sell, these options include both up and down strike prices which may be functions of time. In general for the up strike price, K_u, and the down strike price, K_d, we have $0 < K_d \leq K_u$.

Given the above conditions if we define the exercise and usage decision variables, $\sigma_i^\pm$ and $u_i^\pm$ as

$$\sigma_i^\pm = \begin{cases} 1 & \text{if swing up/down} \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

$$u_i^\pm = \begin{cases} \text{volume bought/sold} & \text{if swing up/down} \\ 0 & \text{otherwise,} \end{cases} \tag{2.3}$$

then we may list the following set of equations as the precise mathematical definition of a swing option for all $1 \leq j < i \leq m$,

$$0 \leq \sigma_i^+ + \sigma_i^- \leq 1 \tag{2.4}$$

$$\left(\sigma_j^+ + \sigma_j^-\right) + (\sigma_i^+ + \sigma_i^-) \leq 1 + \frac{\tau_i}{\tau_j + \Delta\tau} \tag{2.5}$$

$$0 \leq \sum_{i=1}^{m} \sigma_i^+ \leq \mathcal{N}_u \tag{2.6}$$

$$0 \leq \sum_{i=1}^{m} \sigma_i^- \leq \mathcal{N}_d \tag{2.7}$$

$$u_i^3 \sigma_i^+ \leq u_i^+ \leq u_i^4 \sigma_i^+ \tag{2.8}$$

$$u_i^1 \sigma_i^- \leq u_i^- \leq u_i^2 \sigma_i^- \tag{2.9}$$

2.3 Forest of Trees

The Forest of Trees method for multiple exercise option valuation [23] is a generalization of the pricing of American options via trees in regard to both extending the number of trees and the exercise decision. When pricing an American option using dynamic programming techniques one can use a single binomial (or trinomial) tree. The Forest of Trees method generalizes this for swing option valuation to construct a forest which contains a tree for every possible

combination of swing rights and usage levels.

In order to numerically implement this the dynamic programming algorithm must be modified so as to not only move backward in time through a tree but also to move throught the forest beginning on the tree which has no swing rights remaining. An additional restriction which we set is that the volume choices must be discretized giving the holder a finite list from which to choose. Even though we are not theoretically limited to the number of choices in the list we are computationally restricted to keep the list short.

The valuation of multiple exercise options is a stochastic optimal control problem with three relevant state variables; the underlying state variable (S) (here restricted to be $S \in \mathbb{R}$), number of exercise rights remaining $(\mathcal{N})$, and usage level (U) assuming some volume control. At each exercise opportunity and given $(S, \mathcal{N}, U)$, the current values of the state variables, the holder must choose between

- exercising u units plus continuing with an option having $\mathcal{N} - 1$ remaining exercise rights and usage level $U + u$; and

- continuing with an option having $\mathcal{N}$ exercise rights and usage level U (i.e., no exercise).

Note that with volume control the payoff from exercising u units changes with u (as does the continuation value of the option). Thus, the holder chooses the value-maximizing u when deciding to exercise.

Dynamic programming can be used for valuing multiple exercise options. In all variables, let the subscript i denote time-t_i and let $\mathcal{U}_i$ be the time-t_i set of admissible volume choices which includes the zero volume choice (ie hold). The recursive equations for the dynamic program are

$$H_i(S_i, \mathcal{N}_{i+1}, U_{i+1}) = \mathrm{E}[B_{i+1}(S_i, \mathcal{N}_{i+1}, U_{i+1})|\mathcal{Z}_i] \quad \text{and} \tag{2.10}$$

$$B_i(S_i, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[h_i(S_i, \mathcal{N}_i, U_i, u) + H_i(S_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right], \tag{2.11}$$

with the terminal conditions

$$H_m(S_m, \mathcal{N}_m, U_m) = \tilde{\phi}(U_m) \quad \text{and} \tag{2.12}$$

$$B_m(S_m, \mathcal{N}_m, U_m) = \max_{u \in \mathcal{U}_m} \left[h_m(S_m, \mathcal{N}_m, U_m, u) + H_m(S_m, \mathcal{N}_m - I_{\{u \neq 0\}}, U_m + u) \right], \tag{2.13}$$

where $H_i(S, \mathcal{N}, U)$ and $B_i(S, \mathcal{N}, U)$ are the time-t_i, state-$\mathcal{Z}_i$ continuation and option values, respectively, $h_i(S, \mathcal{N}, U, u)$ is the payoff from exercising u units where $h_i(S, \mathcal{N}, U, 0) = 0$, and $\mathcal{Z}_i$ is the time-t_i information set generated by the paths of $(S, \mathcal{N}, U)$. This dynamic program

can be modified in obvious ways to accommodate different multiple exercise option specifications. For example a swing option contract may specify a certain number of *up* and *down* swing rights, $\mathcal{N}_u$ and $\mathcal{N}_d$. Another variation is to allow for multiple rights to be exercised at each opportunity where each right corresponds to a fixed volume amount [24, 25].
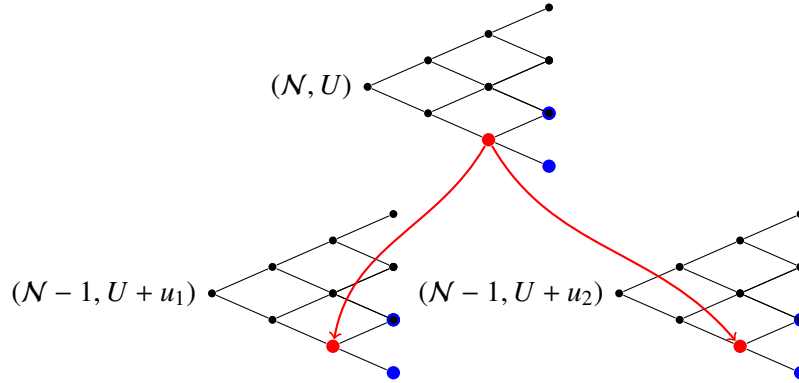
Figure 2.1: Section of a Forest of Trees, $\mathcal{N}$ = # of Swing rights remaining, U = usage level.

The Forest of Trees method generalizes the tree method for pricing American options. In this method, many replications of the tree describing the underlying state variable process are used, each corresponding to a different state (number of rights remaining, usage level). Only a discrete finite set of volume choices is allowed, otherwise there would be an infinite number of trees in the forest. The dynamic program in (2.10)-(2.13) is solved exactly using the trees to compute the continuation value for each possible state (2.10). Comparing this with the analogous equation for an American option (1.3) it is easy to see that here we have an extra term in (2.11) which is due to swing options having multiple exercise rights that may be used during the contract's lifetime.

Figure 2.1 is a diagram of a section in a forest with two volume choices, u_1 and u_2, and no distinction between up and down rights. It illustrates the nodes in the forest which need be considered when making an exercise decision. Clearly as the number of volume choices increases, so does the number of accessible nodes required to compute the option value.

Chapter 3

Forest of Stochastic Trees

This chapter describes the first of two methods developed in this thesis, the Forest of Stochastic Trees. This method, as described, is a generalization of the previously introduced Forest of Trees [23] method which replaces the binomial trees with stochastic trees [6]. Detailed descriptions of the dynamic program and the properties of the resulting estimators are given. The chapter ends with numerical results illustrating the method along with a detailed discussion and more results about a bias corrected version of the forest of stochastic trees.

3.1   Valuation Via Dynamic Programming

To value multiple exercise options using the forest of stochastic trees we use the framework
of [23, 26] and generalize it by replacing their binomial and trinomial trees with stochastic
trees. The stochastic trees are constructed identically as described in Section 1.3.4, however
now we have multiple replications of this stochastic tree, each representing a different state
(i.e., number of exercise rights remaining and usage level).

In the dynamic program for the forest of stochastic trees, the continuation values in Equa-
tions (2.10) and (2.12) are replaced by stochastic tree-type estimators. As with the original
stochastic tree technique, high- and low-biased option value estimators are constructed by us-
ing the analogous high- and low-biased continuation value estimators, respectively, on each
stochastic tree in the forest. The recursive equations for the high-biased estimator are

$$\hat{H}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}) = \frac{1}{b} \sum_{k=1}^{b} \hat{V}_{i+1}(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_{t+1}, U_{t+1}), \quad \text{and} \tag{3.1}$$

$$\hat{V}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[h_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) + \hat{H}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right], \tag{3.2}$$

with the terminal conditions

$$\hat{V}_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = \max_{u \in \mathcal{U}_m} \left[h_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m, u) + \tilde{\phi}(U_m + u) \right], \tag{3.3}$$

where $\hat{H}_i(\mathbf{S}, \mathcal{N}, U)$ and $\hat{V}_i(\mathbf{S}, \mathcal{N}, U)$ are the time-t_i, state-$\mathcal{Z}_i$ continuation and option values
estimators, respectively, $h_i(\mathbf{S}, \mathcal{N}, U, u)$ (with $h_i(\mathbf{S}, \mathcal{N}, U, 0) = 0$) is the time-$t_i$, state-$\mathcal{Z}_i$ payoff
from exercising u units, b is the branching factor, I is an indicator function and $\tilde{\phi}(U_m + u)$ is
a global penalty term. The superscript $\mathbf{j} = \{j_0, j_1, \ldots, j_i\}$ indicates the specific node within a
given stochastic tree and $\mathbf{k} = \{\mathbf{j}, k\}$.

Figure 3.1 is a diagram of a section in a forest of stochastic trees with two volume choices,
u_1 and u_2. It illustrates the nodes in the forest which need be considered when making an
exercise decision given state $(\mathcal{N}, U)$. The three choices are no exercise, exercise u_1 units, and
exercise u_2 units.

The low estimator is similarly defined using the low estimator on each stochastic tree via
the dynamic program,

$$\hat{g}_{il}(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) = h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + \frac{1}{b-1} \sum_{\substack{k=1 \\ k \neq l}}^{b} \hat{v}_{i+1}(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u), \tag{3.4}$$

$$\hat{H}_{il}(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[\hat{g}_{il}(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right], \tag{3.5}$$
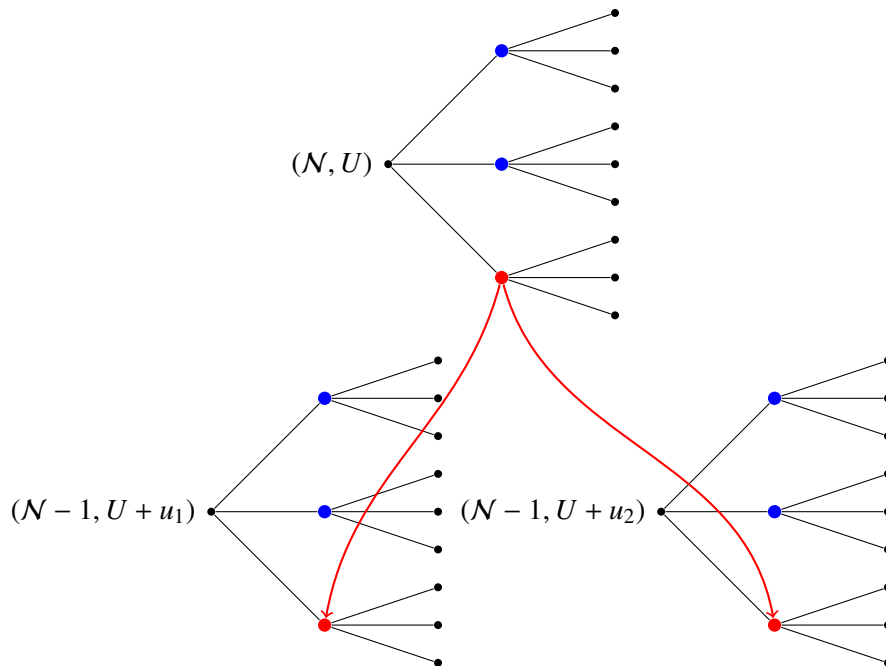
Figure 3.1: Section of a Forest of Trees, $\mathcal{N}$ = # of exercise rights remaining, U = usage level.

$$\hat{v}_{il}(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = h_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^*) + \hat{v}_{i+1}(\mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{u^* \neq 0\}}, U_i + \hat{u}^*), \quad \text{and} \quad (3.6)$$

$$\hat{v}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il}(\mathbf{S}_i^{\mathbf{j}}, N_i, U_i) \quad (3.7)$$

where $\hat{H}_{il}(\mathbf{X}_i^j, \mathcal{N}_i, U_i)$ is the l–th leave-one-out hold value estimator and $\hat{u}^*$ is the estimated optimal exercise amount which depends on i and l. The terminal conditions associated with this dynamic programming scheme are,

$$\hat{v}_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = \max_{u \in \mathcal{U}_m} \left[h_m(\mathbf{S}_m^j, \mathcal{N}_m, U_m, u) + \tilde{\phi}(U_m + u) \right], \quad (3.8)$$

where $\tilde{\phi}(U_m + u)$ is a global penalty term.

3.2   Bias of Estimators

In order to justify using the high- and low-biased estimators to define the upper and lower bounds respectively of our confidence intervals we must prove that indeed the high estimator is always positively biased and that the low estimator is always negatively biased. In addition we include a comparison of the estimators which orders their values on any realization of the simulated forest.

The theorems that follow are direct extensions of those in [6]. Below, the branching factor, b, appears as an argument in the estimators. For example, $\hat{V}_0(b, \mathbf{S}_0, \mathcal{N}_0, U_0)$ refers to the time-0, state-$\mathcal{Z}_0$ high-biased estimator with a stochastic tree branching factor of b. This argument has been suppressed to this point for convenience. We begin with the theorem regarding the bias of the high estimator.

Theorem 9 *(High estimator bias) The high estimator is biased high, i.e.,*

$$\mathrm{E}\left[\hat{V}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)\right] \geq B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right) \tag{3.9}$$

for all b.

Next the theorem stating the bias of the low estimator is as follows.

Theorem 10 *(Low estimator bias) The low estimator is biased low, i.e.,*

$$\mathrm{E}\left[\hat{v}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)\right] \leq B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right) \tag{3.10}$$

for all b.

Finally, a comparison of the high and low estimator bias is stated in the following theorem.

Theorem 11 *(Comparison of Estimators) On every realization of the forest the low estimator is less than or equal to the high estimator. That is,*

$$\hat{v}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \leq \hat{V}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \tag{3.11}$$

*with probability one for all **j**, i.*

Proofs for all of these theorems are found in Appendix A.

3.3 Convergence of Estimators

An advantage of the stochastic tree method over some other MC valuation methods is that its estimators are consistent to the true value of the option. This property continues to hold true for the forest of stochastic trees estimators. In this section we state two theorems; one for the consistency of the high estimator, and the other for the consistency of the low estimator. Here convergence is in probability to the true option value and as above the argument b that appears with the estimators refers to an arbitrary branching factor size of b with convergence being shown as $b \to \infty$.

Theorem 12 *(High estimator convergence) Suppose* $\mathrm{E}\left[|h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i\right)|^{p'}\right] < \infty$, *for all* t_i, *and some* $p' > 1$. *Then* $\hat{V}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)$ *converges to* $B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right)$ *in p-norm for any* $0 < p < p'$ *as* $b \to \infty$ *with the number of repeated valuations of the forest, R, arbitrary. In particular* $\hat{V}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)$ *converges to* $B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right)$ *in probability and is thus a consistent estimator of the option value, i.e..*

$$\mathrm{E}\left[\hat{V}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)\right] \to B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right) \tag{3.12}$$

as $b \to \infty$.

Theorem 13 *(Low estimator convergence) Suppose that,*

$$\mathrm{P}\left[h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u^1\right) + H_i\left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u^1 \neq 0\}}, U_i + u^1\right)\right.$$
$$\left. \neq h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u^2\right) + H_i\left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u^2 \neq 0\}}, U_i = u^2\right)\right] = 1,$$

for $u^1, u^2 \in \mathcal{U}_i$, $u^1 \neq u^2$ *and all i. Then Theorem 12 also holds for the low estimator.*

The additional condition imposed in Theorem 13 is analogous to that used in Theorem 3 of [6]. This condition says that, with probability one, the optimal exercise policy is never indifferent between the choices of volumes to exercise (including $u = 0$). As in [6] imposing this condition simplifies the analysis of the estimator.

Proofs for these theorems showing estimator convergence are found in Appendix B.

3.4   Numerical Results

Pricing swing contracts involves a large number of parameters and in this section we provide some results which illustrate the validity of our method across a variety of specifications. We assume that the underlying assets follow a risk neutral stochastic process, there are no transaction costs and other than penalties, there are no other constraints considered. We also assume a constant *risk free* rate of interest and the volatilities of all assets are known constant functions of time.

The option swing rights may be exercised at discrete times up to and including expiry and the volume choices given are in discrete amounts. This is to say that at any time that the holder chooses to exercise one of their rights, they must choose from a finite list of possible volume amounts. The rational behind allowing all time steps to be exercise opportunities is the exponential growth in computational time caused by adding intermediate non-exercise times.

However, the method can easily be modified to incorporate these extra time steps. As previously mentioned penalties can be implemented globally and are based on the net volume swung during the contract.

3.4.1 One-dimension

Beginning with the one dimensional case, we have based our simulations on an underlying asset with a risk neutralized price process that satisfies the following stochastic differential equation,

$$dS_i = S_i [(r - \delta) \, dt + \sigma dZ_i] . \tag{3.13}$$

In this equation, r is the riskless interest rate, Z_i is a standard Brownian motion process, σ is a constant volatility parameter and the underlying asset itself pays a continuous dividend yield δ.

All simulations in this subsection were completed on the SHARCNet cluster Whale. Whale is located at the University of Waterloo and consists of Opteron 2.2 GHz processors (4 per node) with a Gigabit Ethernet interconnect. Timing results listed below are given in total cpu time accumulated which is approximately equal to (program runtime) × (number of processors used). In this example the following set of parameters is held constant; expiry is 3.0 years, the risk free rate is $r = 0.05$, there is a continous dividend yield of $\delta = 0.1$, a base volume of the underlying, The up and down swing right strikes are $K_u = K_d = 40.0$, and the volatility of the asset is $\sigma = 0.2$. For comparison purposes the results in this subsection include a *binomial* value which is calculated using the forest of trees [23].

In cases where a list of multiple volume choices is given, the volume values are consecutive integer multiples of the base amount (20 units) and the up and down volumes have the same magnitude. In addition all tables list the number of repeated valuations, R, and the number of time steps (exercise opportunities), m.

In Figure 3.2 the swing option being priced consists of 1 up and 1 down swing rights that may be exercised at 3 opportunities to receive (up right) or sell (down right) 60 units of the underlying. The initial price is $40, there are no penalties. The results illustrate how the two estimators are affected by increasing the branching factor. Mainly it illustrates the properties in the previous section which states the high-biased estimator is positively biased, the low-biased estimator is negatively biased and both are consistent to the true option value. It shows a tightening in the spread between the high- and low-estimators which appears consistent with the result from the forest of trees algorithm.

The results in Table 3.1 show how the estimates vary with moneyness and the effects of penalties. In these simulations global cash penalties were accrued if the final net usage level

Single Assets
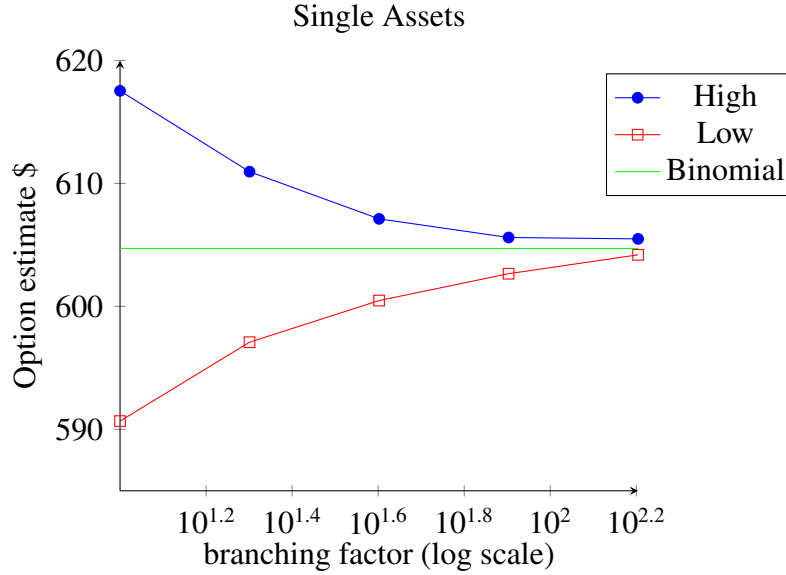
Figure 3.2: Option value estimate (\$) vs log branching factor. Here $R = 32000 \left(\frac{10}{b}\right)$, where b is the branching factor and standard error $\approx 0.07\%$ of estimator value

had a magnitude that was greater than 90 units. The option gave the holder the right to swing up or down 2 times each with a list containing 3 volume choices which, given the base volume of 20 units, the choices included 20, 40 and 60 units of the asset. The penalty multiplication factors when included are taken to be $\mathcal{P}_u = \mathcal{P}_d = 10$. Average time per row (not including the binomial forest valuation) for the case with penalties was 5.6 hours and for penalities off was 1.1 hours. The reduction in runtime for the case with no penality can be described as follows. If there are no constraints (penalities, storage, etc.) on the holder of the option then if a decision is made to swing it is always optimal for the holder to swing the maximum amount. Therefore with no penalities the above valuation is equivalent to that of a swing option with no volume choices and a base volume of 60 units but otherwise identical. The latter has fewer trees in its forest and is therefore quicker to evaluate. In Table 3.1 we have chosen to exploit this as a convient way to save computational time. For the binomial method run times were on the order of a few seconds.

Table 3.2 shows the increase in the estimate of the option value as the number of swing rights are increased. As can be noted the cases $\mathcal{N}_u = \mathcal{N}_d = 3$ and 5 are less than 3 and 5 times the estimate for the $\mathcal{N}_u = \mathcal{N}_d = 1$ case respectively. This result matches with the intuition that a swing with a given number of rights is less valuable than a basket of American Put and Call options with otherwise identical parameters and $K_d \leq K_u$. Here, for the case of $\mathcal{N}_u = \mathcal{N}_d = 5$ the high and low estimator are the same due to the fact that the number of up and down swing

S_0	Penalty	High	Error	Low	Error	Binomial
60	ON	2271.153	1.418	2240.319	1.378	2259.845
	OFF	2422.781	1.576	2392.872	1.523	2411.844
50	ON	1445.468	0.844	1408.843	0.904	1429.645
	OFF	1542.053	0.978	1503.963	0.980	1526.055
40	ON	1018.104	0.859	968.793	1.044	989.651
	OFF	1156.591	0.911	1134.093	0.903	1145.801
30	ON	1345.556	1.205	1309.214	1.280	1326.266
	OFF	1562.347	1.316	1532.854	1.343	1546.055
20	ON	2189.531	0.905	2147.623	1.018	2157.976
	OFF	2443.877	0.924	2402.192	1.034	2412.354

Table 3.1: Swing option values as a function of moneyness and penalties. Parameters: $\mathcal{N}_u = \mathcal{N}_d = 2$, $\mathcal{U}_i = \{20, 40, 60\}$, $b = 20$, $R = 4000$, $m = 5$, $U_{min} = -90$, $U_{max} = 90$

rights are equal to the number of time steps and the up and down strike prices are the same. Therefore the holder of the option at each time step would swing either up or down. For the low estimator if each of the repeated exercise decisions are the same (i.e. all hold or all exercise) then it becomes equal to the high estimator and that is what is illustrated here.

Figure 3.3 shows a comparison between a basket of American options and a swing option with a comparable number of exercise rights. For the case of 1 up and 1 down swing right the basket of American options would contain a single call and a single put option and so forth with equal strike prices for both types of rights. As expected the graph of the basket of American options is linear and the graph of the swing option value is less then the values of the American option when the number of rights is greater then one. This follows from the restriction that only one swing right may be exercised at an exercise opportunity whereas all American options of a particular type could be exercised at a given time. In the case of 1 up and 1 down right the the two are equal since it would never be optimal to exercise both the put and call style rights at the same time. The low-biased estimator is used for both the basket and swing option priced in Figure 3.3.

$\mathcal{N}_u = \mathcal{N}_d$	High	Error	Low	Error	Binomial
1	630.054	0.449	605.394	0.453	617.832
3	1573.237	1.449	1559.517	1.437	1567.344
5	1852.788	2.128	1852.788	2.128	1852.627

Table 3.2: Swing option values as a function of the number of exercise rights. Parameters: base volume = 60 units, $S_0 = 40$, $b = 20$, $R = 4000$, $m = 5$, no penalties

Option Estimate vs Exercise rights - Single Asset - Low Estimator

Figure 3.3: Comparing values of a basket of American calls and puts vs swing option value using the low-biased estimator as a function of the number of exercise rights

3.4.2   Calibrated Forward Curve

As a second example we use the trinomial-tree model given in Figure 5 of [26] from which price movements are simulated. This model is a 1-factor model with mean reversion that is seasonally adjusted and calibrated to the forward curve. The option we value is that of [26] Section 4.2 Example (a), which is a 2 up right swing option with each right allowing the holder to take delivery of either 1 or 2 MMBTus of natural gas. It is simplified to have 4 exercise opportunities and 4 months until expiry. Upon exercise the holder gets

$$\max \left(U_i \left(A_i S_i - K \right), 0 \right) \tag{3.14}$$

where U_i is the volume chosen, A_i is the seasonality factor and S_i is the deseasonalized spot price.

In Figure 3.4 we see that, with a branching factor of only 8, the confidence intervals for the high- and low-biased estimators begin to overlap and quickly become almost indistinguishable for higher branching factors, numerical illustration of both estimators' consistency. The number of realizations used to generate the results shown were $\frac{160000}{b}$ and the serial computational times for branching factors of 8 and 32 were approximately 4.5 and 110 seconds respectively using a 2.1 Ghz Core 2 Duo processor. The results shown in Figure 3.4 are consistent with the results in [26] but we note that the valuation method in that publication breaks down in higher
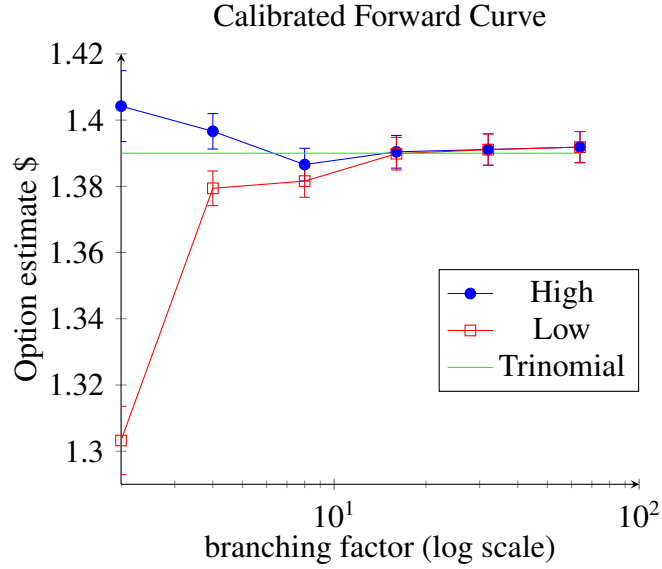
Figure 3.4: Option-value estimates versus (log) mesh size. Approximate pointwise 95% confidence intervals for each estimate are given by the vertical bars.

dimensions and in cases where the inclusion of more risk factors is desirable.

3.4.3 Five-dimension

Due to the computationally intensive nature of this method it only becomes truly useful in cases where PDE or tree based methods fail. In this subsection we provide numerical results for the simplest of these cases, GBM with d underlying assets ($d > 3$). The payoff function used,

$$u \times \left(\max_{k=1,\dots,d} S^k - K_u, K_d - \max_{k=1,\dots,d} S^k \right)^+.$$
(3.15)

in the case of d underlying assets and where u is the exercise volume. This is an extension of the example given in [6] and [7].

The asset prices are taken to follow a risk neutralized correlated GBM described by the stochastic differential equation,

$$dS_i^k = S_i^k \left[\left(r - \delta^k \right) dt + \sigma^k dZ_i^k \right],$$
(3.16)

where Z_i^k is a standard Brownian motion process where the correlation between Z^k and Z^s is ρ^{ks}. In all simulated results that follow it is assumed that $\delta^k = \delta$, $\sigma^k = \sigma$ for all k and that $\rho^{ks} = 0$ for all $k \neq s$. In addition to this we also take all assets to have the same initial value,

S_0.
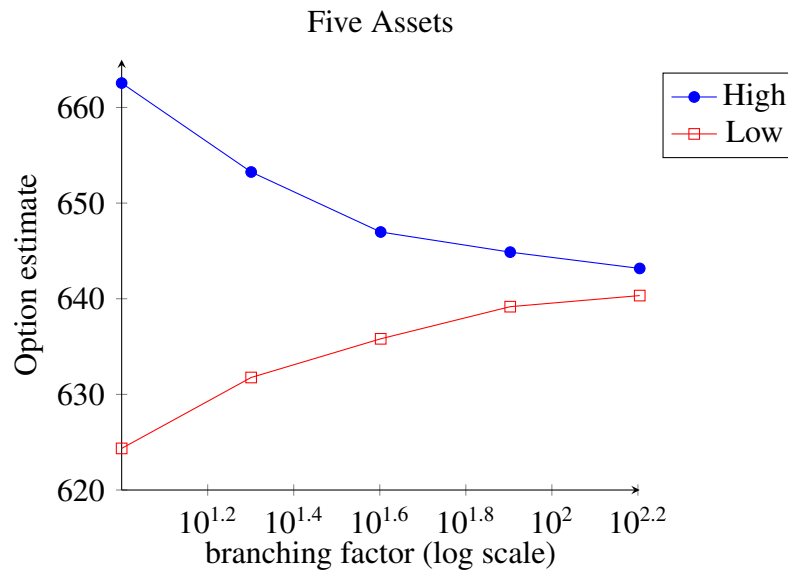
Figure 3.5: $\mathcal{N}_u = \mathcal{U}_d = 1$, $\mathcal{U}_i = \{60\}$, $S_0 = 40$, $R = 32000\left(\frac{10}{b}\right)$, $m = 3$, no penalties, std. err. $\approx$ 0.09%

Figure 3.5 shows that the high- and low-biased estimators appear to be converging to a common value for the option estimate as the branching factor increases. The swing option being priced here gives the holder 1 up right and 1 down right to exercise over 3 exercise opportunities.

The results in Table 3.3 show how the estimates vary with moneyness and the effects of penalties. In these simulations global cash penalties were accrued if the final net usage level had a magnitude that was greater than 90 units. The option gave the holder the right to swing up or down 2 times each with a list containing 3 volume choices which, given the base volume of 20 units, the choices included 20, 40 and 60 units of the asset. The penalty multiplication factors when included are taken to be $\mathcal{P}_u = \mathcal{P}_d = 10$. Average time per row for penalties *ON* was 5.9 hours and for penalities off was 1.5 hours.

Table 3.4 shows the increase in the estimate of the option as the number of swing rights are increased. The intuition here is the same as that describing Table 3.2

Figure 3.6 shows a comparison between a basket of American options and a swing option with a comparable number of exercise rights. The intuition here is the same as that describing the equivalent one-dimensional case.

S_0	Penalty	High	Error	Low	Error
60	ON	3577.280	2.864	3517.297	2.845
	OFF	3832.050	2.286	3772.123	2.856
50	ON	2246.657	2.280	2197.957	2.259
	OFF	2479.081	2.341	2431.065	2.318
40	ON	1221.847	1.595	1189.610	1.564
	OFF	1257.171	1.499	1226.370	1.467
30	ON	1105.831	0.453	1087.851	0.447
	OFF	1209.179	0.393	1196.255	0.391
20	ON	1937.615	0.445	1930.860	0.472
	OFF	2177.194	0.489	2177.031	0.513

Table 3.3: Swing option values as a function of moneyness and penalties. Parameters: $\mathcal{N}_u = \mathcal{N}_d = 2$, $\mathcal{U}_i = \{20, 40, 60\}$, $b = 20$, $R = 4000$, $m = 5$, $U_{min} = -90$, $U_{max} = 90$.

$\mathcal{N}_u = \mathcal{N}_d$	High	Error	Low	Error
1	683.144	0.741	652.481	0.721
3	1728.947	2.279	1709.497	2.248
5	2087.495	3.114	2087.495	3.114

Table 3.4: Swing option values as a function of the number of exercise rights. Parameters: $\mathcal{U}_i = \{60\}$, $S_0 = 40$, $b = 20$, $R = 4000$, $m = 5$, no penalties.

3.5 Algorithmic Enhancements

In this section we describe two methods that may be adopted to increase the efficiency of the forest of stochastic trees algorithm. The first method described is the implementation of parallel computing techniques to the existing algorithm to shorten overall run times without reducing the number of computations required to produce the estimators. The second method is an extension of the bias corrected estimator method for the stochastic tree method discussed in 1.3.5. This method successfully reduces the branching factor required to obtain the desired accuracy for the option value. These enhancements are not mutually exclusive and may both be applied to further improving the algorithm's performance. Indeed the results shown at the end of this section were simulated using a combination of the two as discussed below.

3.5.1 Parallel Processing

One method for enhancing the computational efficiency of this algorithm is by taking advantage of multi-processor computing techniques. The simplest and most obvious implementation would be to parallelize across repeated valuations of the forest resulting in effectively serial

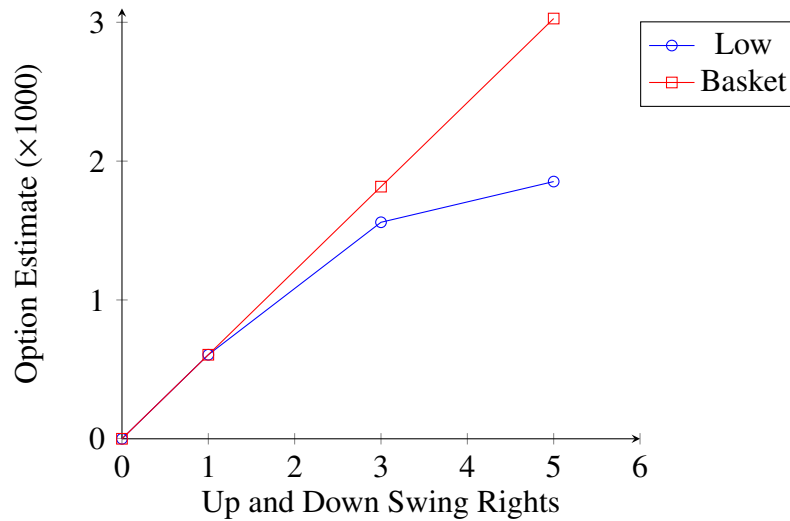Option Estimate vs Exercise rights - Five Assets - Low Estimator
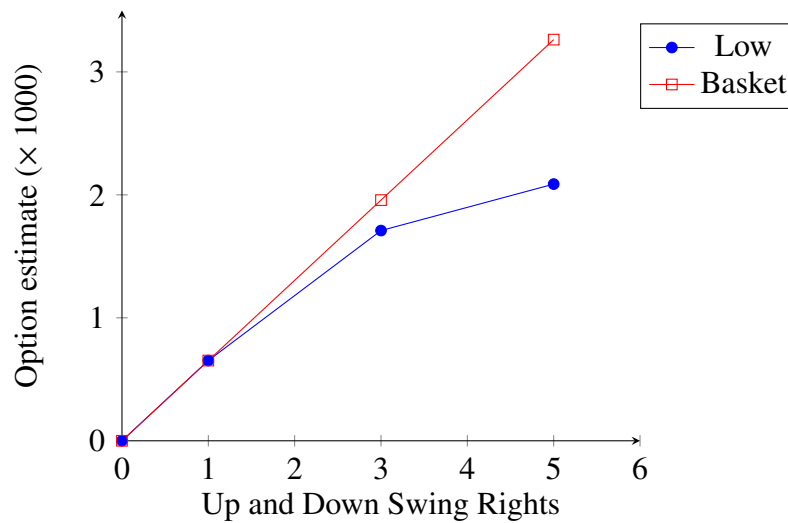
Figure 3.6: Comparing values of a basket of American calls and puts vs swing option values using the low-biased estimator as a function of the number of exercise rights.

farming of the repeated valuations. Since each repeated valuation results in an iid random value for the option estimate, the generation of all the results may be completed independently of one another, removing the need for communication between processors. This method is simple and effective. However we state here without numerical evidence that it results in a near perfect speed up without the need for expensive interconnections. With this method the minimum run time that can be produced is determined by the number of processors available, the number of repeated valuations necessary for the desired accuracy and the run time of a single forest.

A variation on the aforementioned parallel implementation is to parallelize the forest of stochastic tree computations internally within the forest. In the results shown in Figure 3.7 the forest of stochastic trees algorithm has been modified so that the computation of the individual trees within the forest is done using multiple processors. Here we have begun the parallelization after the first time step by dividing up the computation of the remaining subtrees across different processors. Upon completion the results are gathered and the option value at the initial time step is determined. In Figure 3.7 we see that this method results in a near perfect speed up due to the small ratio of communication time versus computational time. This implementation may be combined with serial farming resulting in further computational time efficiency.

In Figure 3.7 the swing option being priced is identical to the no volume choice swing option of Section 3.4.3 with a branching factor $b = 160$. The results found here were generated using the SHARCNET cluster Hound which comprises 2.2 GHz Opteron processors with 4 GB

per core and Infini-Band interconnections. Run times are normalized to the run time of a single processor.
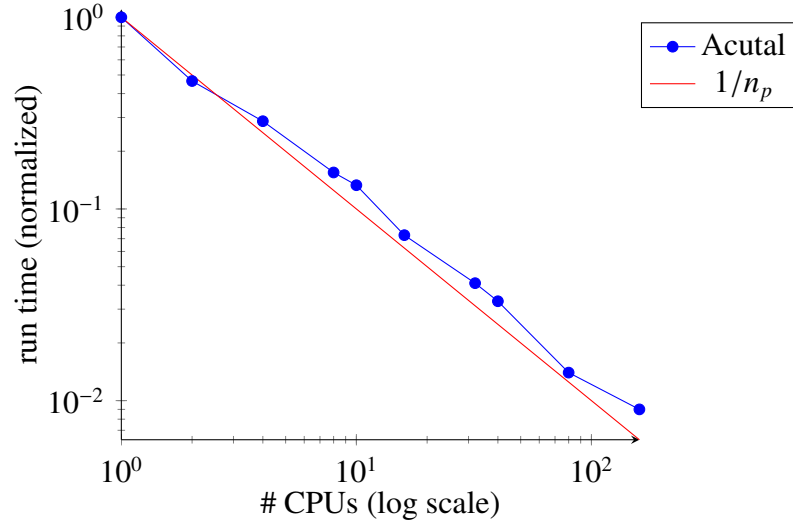
Figure 3.7: Normalized speedup graph using MPI. Forest of Stochastic Trees with branching factor $b = 160$. n_p appearing in the legend is the number of processors

3.5.2   Bias Reduction

In addition to parallel computing techniques the bias reduction techniques discussed in Section 1.3.5 are generalizable to the forest of stochastic tree estimators [35]. In doing so we must account for the possibility of differing volume choices that may be available to the holder of the multiple exercise option. If we define the time-t_i bias analogously to its definition in Section 1.3.5 then we have that the bias of the high and low estimators are given by

$$
\begin{aligned}
\bar{H}_i^+ &\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) - H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) \\
&= \mathrm{E}\left[\hat{H}_i^+\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) | \mathcal{Z}_i\right] - H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) \\
&= \mathrm{E}\left[\hat{V}_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) | \mathcal{Z}_i\right]
\end{aligned}
\tag{3.17}
$$

and

$$
\begin{aligned}
\bar{H}_i^- &\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) - H_i^{\mathbf{j}}\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) \\
&= \mathrm{E}\left[\hat{H}_i^-\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) | \mathcal{Z}_i\right] - H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) \\
&= \mathrm{E}\left[\hat{v}_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right) | \mathcal{Z}_i\right]
\end{aligned}
\tag{3.18}
$$

respectively, where the superscript ± denotes the corresponding estimator by whether it is positively or negatively biased.

As with the bias reduced stochastic tree these biases may be broken down into local and global components. To demonstrate this consider the high-biased estimator in a case where there are two available volume choices (i.e. $u \in \mathcal{U}_i = \{0, u_1, u_2\}$). Then upon expanding the terms of (3.17) inside the conditional expectation we obtain

$$\begin{aligned}
\mathrm{E}\Big[\max\Big\{&h(u_1) + \hat{H}^+(u_1), h(u_2) + \hat{H}^+(u_2), \hat{H}^+(0)\Big\} \\
&- \max\Big\{h(u_1) + H(u_1), h(u_2) + H(u_2), H(0)\Big\} \Big| \mathcal{Z}_i\Big],
\end{aligned} \tag{3.19}$$

where we have adopted the following condensed notation

$$\begin{aligned}
h(u) &= h_{i+1}(\mathbf{S}^{\mathbf{j}}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}, u), \\
\hat{H}^\pm(u) &= \hat{H}^\pm_{i+1}(\mathbf{S}^{\mathbf{j}}_{i+1}, \mathcal{N}_{i+1} - I_{\{u \neq 0\}}, U_{i+1} + u), \\
\bar{H}^\pm(u) &= \bar{H}^\pm_{i+1}(\mathbf{S}^{\mathbf{j}}_{i+1}, \mathcal{N}_{i+1} - I_{\{u \neq 0\}}, U_{i+1} + u).
\end{aligned}$$

Adding and subtracting the term

$$\mathrm{E}\Big[\max\Big\{h(u_1) + \bar{H}^+(u_1), h(u_2) + \bar{H}^+(u_2), \bar{H}^+(0)\Big\} \Big| \mathcal{Z}_i\Big] \tag{3.20}$$

splits the expression in (3.19) into a local and global component, given by

$$\begin{aligned}
\mathrm{E}\Big[\max\Big\{&h(u_1) + \hat{H}^+(u_1), h(u_2) + \hat{H}^+(u_2), \hat{H}^+(0)\Big\} \\
&- \max\Big\{h(u_1) + \bar{H}^+(u_1), h(u_2) + \bar{H}^+(u_2), \bar{H}^+(0)\Big\} \Big| \mathcal{Z}_i\Big]
\end{aligned} \tag{3.21}$$

and

$$\begin{aligned}
\mathrm{E}\Big[\max\Big\{&h(u_1) + \bar{H}^+(u_1), h(u_2) + \bar{H}^+(u_2), \bar{H}^+(0)\Big\} \\
&- \max\{h(u_1) + H(u_1), h(u_2) + H(u_2), H(0)\} \Big| \mathcal{Z}_i\Big],
\end{aligned} \tag{3.22}$$

respectively. Before proceeding, note that we hold provided

$$\hat{H}^\pm(0) \geq h(u) + \hat{H}^\pm(u), \forall u \in \mathcal{U}_i \tag{3.23}$$

and exercise u_1 units if

$$h(u_1) + \hat{H}^\pm(u_1) \geq \max\Big\{h(u_1) + \hat{H}^\pm(u_2), \hat{H}^\pm(0)\Big\}. \tag{3.24}$$

Likewise for the bias of the low estimator we can expand the inner terms of Equation 3.18 which results in

$$\mathrm{E}\left[I_{\{\hat{0}\geq\max\{\hat{u}_1,\hat{u}_2\}\}}\hat{H}^-(0) + I_{\{\hat{u}_1\geq\max\{\hat{0},\hat{u}_2\}\}}\left(h(u_1) + \hat{H}^-(u_1)\right)\right. \tag{3.25}$$

$$\left.+I_{\{\hat{u}_2\geq\max\{\hat{0},\hat{u}_1\}\}}\left(h(u_2) + \hat{H}^-(u_2)\right) - B_i\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right)|\mathcal{Z}_i\right], \tag{3.26}$$

where $\hat{0}, \hat{u}_1, \hat{u}_2$ are the potential estimated usage choice decisions. Adding and subtracting $\mathrm{E}[\max\{\bar{u}_1, \bar{u}_2, \bar{0}\}|\mathcal{Z}_i]$ splits the expression into a local (3.27) and global (3.28) component given by

$$\mathrm{E}\left[I_{\{\hat{0}\geq\max\{\hat{u}_1,\hat{u}_2\}\}}\hat{H}^-(0) + I_{\{\hat{u}_1\geq\max\{\hat{0},\hat{u}_2\}\}}\left(h(u_1) + \hat{H}^-(u_1)\right)\right.$$

$$\left.+I_{\{\hat{u}_2\geq\max\{\hat{0},\hat{u}_1\}\}}\left(h(u_2) + \hat{H}^-(u_2)\right) - \max\left\{\bar{u}_1, \bar{u}_2, \bar{0}\right\}|\mathcal{Z}_i\right] \tag{3.27}$$

$$+ \mathrm{E}\left[\max\left\{\bar{u}_1, \bar{u}_2, \bar{0}\right\} - B_i\left(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}\right)|\mathcal{Z}_i\right], \tag{3.28}$$

where $\bar{0}, \bar{u}_1, \bar{u}_2$ are the potential optimal usage choice decisions. Focusing on the local component of the bias it is possible to show that as in the case of an American-style option the bias is mainly caused by incorrect exercise decisions. For the case of two volume choices we summarize the contributions to the bias for the high- and low-biased forest of stochastic trees estimators in Tables 3.5 and 3.6 respectively where we define $\hat{Y}^\pm(u) = \hat{H}^\pm(0) - h(u) - \hat{H}^\pm(u)$.

	Held	Exercised u_1 units	Exercised u_2 units
Should hold	0	$-\hat{Y}^+(u_1)$	$-\hat{Y}^+(u_2)$
Should exercise u_1 units	$\hat{Y}^+(u_1)$	0	$\hat{Y}^+(u_1) - \hat{Y}^+(u_2)$
Should exercise u_2 units	$\hat{Y}^+(u_2)$	$\hat{Y}^+(u_2) - \hat{Y}^+(u_1)$	0

Table 3.5: The local error in the time-t_{i+1} high-biased hold value estimator. Note that this error is always non-negative.

In the general case of a z-volume choices swing option we have that $u \in \mathcal{U}_t = \{u_1, \ldots, u_z : z \in \mathbb{N}\}$. Equations (3.19)-(3.22) and Equations (3.25)-(3.28) are easily extended to the high-bias estimator for z-volume choices and as expected the local bias is mainly caused by incorrect

	Held	Exercised u_1 units	Exercised u_2 units
Should hold	0	$-\bar{Y}^-(u_1)$	$-\bar{Y}^-(u_2)$
Should exercise u_1 units	$\bar{Y}^-(u_1)$	0	$\bar{Y}^-(u_1) - \bar{Y}^-(u_2)$
Should exercise u_2 units	$\bar{Y}^-(u_2)$	$\bar{Y}^-(u_2) - \bar{Y}^-(u_1)$	0

Table 3.6: The local error in the time-t_{i+1} low-biased hold value estimator. Note that this error is always non-positive.

exercise decisions. Tables 3.7 and 3.8 list the possible contributions to the bias for the high- and low-biased estimators respectively.

	Held	Exercised u_1 units	$\cdots$	Exercised u_z units
Should hold	0	$-\hat{Y}^+(u_1)$	$\ldots$	$-\hat{Y}^+(u_z)$
Should exercise u_1 units	$\hat{Y}^+(u_1)$	0	$\ldots$	$\hat{Y}^+(u_1) - \hat{Y}^+(u_z)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Should exercise u_z units	$\hat{Y}^+(u_z)$	$\hat{Y}^+(u_z) - \hat{Y}^+(u_1)$	$\ldots$	0

Table 3.7: The local error in the time-t_{i+1} high-biased hold value estimator, z-volume case. Note that this error is always non-negative.

Equations built from of these bias contributions are not useful in their own right, as the $\bar{Y}^\pm(u)$ cannot be observed directly and upon replacing them with the corresponding $\hat{Y}^\pm(u)$ results in the expressions collapsing to zero. As in the American option case, it is necessary to incorporate additional distributional knowledge. The time-t_i hold value estimators are averages of the time-t_{i+1} option value estimators which given $\mathcal{Z}_i$, are independent and identically distributed. At this point we now employ the Multivariate Central Limit Theorem to govern the joint (limiting) distribution of $\hat{H}^\pm(u_1), \ldots, \hat{H}^\pm(u_z)$ and $\hat{H}^\pm(0)$. Replacing all $\hat{Y}^\pm(u)$ with $\hat{Y}^{\pm*}(u)$

	Held	Exercised u_1 units	$\cdots$	Exercised u_z units
Should hold	0	$-\bar{Y}^-(u_1)$	$\ldots$	$-\bar{Y}^-(u_z)$
Should exercise u_1 units	$\bar{Y}^-(u_1)$	0	$\ldots$	$\bar{Y}^-(u_1) - \bar{Y}^-(u_z)$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
Should exercise u_z units	$\bar{Y}^-(u_z)$	$\bar{Y}^-(u_z) - \bar{Y}^-(u_1)$	$\ldots$	0

Table 3.8: The local error in the time-t_{i+1} low-biased hold value estimator, z-volume case. Note that this error is always non-negative.

which is assumed to be normally distributed with $\mathcal{Z}_i$-conditional mean $\bar{Y}^\pm(u)$ and variance

$$\mathrm{Var}\left[\hat{Y}^{\pm*}(u)|\mathcal{Z}_i\right] = \mathrm{Var}\left[\hat{H}^{\pm*}(0) - h(u) - \hat{H}^{\pm*}(u)|\mathcal{Z}_i\right]$$
$$= \mathrm{Var}\left[\hat{H}^{\pm*}(0)|\mathcal{Z}_i\right] + \mathrm{Var}\left[\hat{H}^{\pm*}(u)|\mathcal{Z}_i\right] + \mathrm{Cov}\left[\hat{H}^{\pm*}(0), \hat{H}^{\pm*}(u)|\mathcal{Z}_i\right]$$

This substitution eventually leads to approximations for the z-volume choice local high-bias being given by

$$\mathrm{E}\left[\sum_{a\in\{u_1,\ldots,u_z\}} I_{[\hat{Y}^{+*}(a)\leq\min\{0,\hat{Y}^{+*}(c):c\neq a\}]}\left[-\hat{Y}^{+*}(a)\right] g\left[\hat{Y}^{+*}(u_1),\ldots,\hat{Y}^{+*}(u_z),\bar{\Sigma}_Y^+\right] \right.$$
$$+ \sum_{a\in\{u_1,\ldots,u_z\}} I_{[0\leq\min\{\hat{Y}^{+*}(v_1),\ldots,\hat{Y}^{+*}(v_r)\}]}\left[\hat{Y}^{+*}(a)\right] h\left[\hat{Y}^{+*}(c):c\neq a, \hat{Y}^{+*}(a),\bar{\Sigma}_Y^+\right]$$
$$\left. + \sum_{\substack{a,d\in\{u_1,\ldots,u_z\} \\ a\neq d}} I_{[\hat{Y}^{+*}(a)\leq\min\{0,\hat{Y}^{+*}(c):c\neq a\}]}\left[\hat{Y}^{+*}(d) - \hat{Y}^{+*}(a)\right] h\left[\hat{Y}^{+*}(e):e\neq d, \hat{Y}^{+*}(d),\bar{\Sigma}_Y^+\right] \,\right|\, \mathcal{Z}_i\right],$$

$$(3.29)$$

where

$$g\left[\hat{Y}^{+*}(u_1),\ldots,\hat{Y}^{+*}(u_z),\bar{\Sigma}_Y^+\right] = \int_0^\infty \cdots \int_0^\infty MVN_z\left[\hat{y}^{+*}|\bar{y}^+,\bar{\Sigma}_Y^+\right] d\bar{y}_{1,\ldots,z} \qquad (3.30)$$

$$h\left[\hat{Y}^{+*}(u_z),\ldots,\hat{Y}^{+*}(u_1),\bar{\Sigma}_Y^+\right] = \int_{-\infty}^0 \int_{\hat{y}^{+*}(u_1)}^\infty \cdots \int_{\hat{y}^{+*}(u_1)}^\infty MVN_z\left[\hat{y}^{+*}|\bar{y}^+,\bar{\Sigma}_Y^+\right] d\bar{y}_{z,\ldots,1} \qquad (3.31)$$

$$MVN_z\left[\hat{y}^{\pm*}|\bar{y}^{\pm},\bar{\Sigma}_Y^{\pm}\right] = \frac{1}{(2\pi)^{z/2}\left|\bar{\Sigma}_Y^{\pm}/b\right|^{1/2}}\exp\left\{-\frac{1}{2}\left(\hat{y}^{\pm*}-\bar{y}^{\pm}\right)'\left(\bar{\Sigma}_Y^{\pm}/b\right)^{-1}\left(\hat{y}^{\pm*}-\bar{y}^{\pm}\right)\right\}, \quad (3.32)$$

$$\hat{y}^{\pm*} = \begin{pmatrix}\hat{y}^{\pm*}(u_1)\\ \vdots \\ \hat{y}^{\pm*}(u_z)\end{pmatrix}, \quad \bar{y}^{\pm} = \begin{pmatrix}\bar{y}^{\pm}(u_1)\\ \vdots \\ \bar{y}^{\pm}(u_z)\end{pmatrix} \quad \text{and} \quad d\bar{y}^{\pm} = d\bar{y}^{\pm}(u_1)\cdots d\bar{y}^{\pm}(u_z). \quad (3.33)$$

The order of the $\hat{y}^{+*}(u)$ arguments of h determines the order of integration of the corresponding $\bar{y}^+(u)$, i.e. $\hat{Y}^{+*}(u_z)$ appearing as the first argument of h means $\bar{Y}^+(u_z)$ is to be integrated out first.

Similarly, the expression for the z-volume choice local low-bias is

$$\mathrm{E}\left[\sum_{a\in\{u_1,\dots,u_z\}} I_{[\hat{Y}^{-*}(a)\leq\min\{0,\hat{Y}^{-*}(c):c\neq a\}]}m\left[-\bar{Y}^-(a),\hat{Y}^{-*}(u_1),\dots,\hat{Y}^{-*}(u_z),\bar{\Sigma}_Y^-\right]\right.$$

$$+ \sum_{a\in\{u_1,\dots,u_z\}} I_{[0\leq\min\{\hat{Y}^{-*}(u_1),\dots,\hat{Y}^{-*}(u_z)\}]}n\left[\bar{Y}^-(a),\hat{Y}^{-*}(c):c\neq a,\hat{Y}^{-*}(a),\bar{\Sigma}_Y^-\right]$$

$$\left.+ \sum_{\substack{a,d\in\{u_1,\dots,u_z\}\\ a\neq d}} I_{[\hat{Y}^{-*}(a)\leq\min\{0,\hat{Y}^{-*}(c):c\neq a\}]}n\left[\bar{Y}^-(d)-\bar{Y}^-(a),\hat{Y}^{-*}(e):e\neq d,\hat{Y}^{-*}(d),\bar{\Sigma}_Y^-\right]\,\middle|\,\mathcal{Z}_i\right] \quad (3.34)$$

where

$$m\left[\bar{Y}^-(u_1),\hat{Y}^{-*}(u_1),\dots,\hat{Y}^{-*}(u_z),\bar{\Sigma}_Y^-\right] = \int_0^\infty\cdots\int_0^\infty[\bar{y}^-(u_1)]\,MVN_z\left[\hat{y}^{-*}|\bar{y}^-,\bar{\Sigma}_Y^-\right]d\bar{y}_{1,\dots,z} \quad (3.35)$$

$$n\left[\bar{Y}^-(u_1),\hat{Y}^{-*}(u_z),\dots,\hat{Y}^{-*}(u_1),\bar{\Sigma}_Y^-\right] = \int_{-\infty}^0\int_{\hat{y}^{-*}(u_1)}^\infty\cdots\int_{\hat{y}^{-*}(u_1)}^\infty[\bar{y}^-(u_1)]\,MVN_z\left[\hat{y}^{-*}|\bar{y}^-,\bar{\Sigma}_Y^-\right]d\bar{y}_{z,\dots,1}.$$

$$(3.36)$$

It remains to replace the theoretical quantities, $\hat{Y}^{\pm*}$, $\bar{\Sigma}_Y^{\pm}$ with their corresponding sample quantities, $\hat{Y}^{\pm}$, $\hat{\Sigma}_Y^{\pm}$ in Equations 3.29 and 3.34. To do that we now subtract the modified Equations 3.29 and 3.34 from the uncorrected time-t_i estimators in Equations 3.1 and 3.4. This yields the following expressions for the corrected high- and low-biased estimators for the forest of stochastic trees

$$\hat{V}_i^{\mathbf{j}} = \max_{u\in\mathcal{U}_i}\left[h_i(u)+\hat{H}_i^{\mathbf{j}}(u)\right] - \left[\sum_{a\in\{u_1,\dots,u_z\}} I_{[\hat{Y}^+(a)\leq\min\{0,\hat{Y}^+(c):c\neq a\}]}\left[-\hat{Y}^+(a)\right]g\left[\hat{Y}^+(u_1),\dots,\hat{Y}^+(u_z),\bar{\Sigma}_Y^+\right]\right.$$

$$+ \sum_{a\in\{u_1,\dots,u_z\}} I_{[0\leq\min\{\hat{Y}^+(u_1),\dots,\hat{Y}^+(u_z)\}]}\left[\hat{Y}^+(a)\right]h\left[\hat{Y}^+(c):c\neq a,\hat{Y}^+(a),\bar{\Sigma}_Y^+\right]$$

$$+ \sum_{\substack{a,d\in\{u_1,\ldots,u_z\} \\ a\neq d}} I_{[\hat{Y}^+(a)\leq\min\{0,\hat{Y}^+(c):c\neq a\}]} \left[\hat{Y}^+(d) - \hat{Y}^+(a)\right] h\left[\hat{Y}^+(e) : e \neq d, \hat{Y}^+(d), \bar{\Sigma}_Y^+\right] \Bigg| \mathcal{Z}_i \Bigg],$$

$$(3.37)$$

and

$$\hat{v}_{il}^{\mathbf{j}} = h_i(\hat{u}^*) + \hat{v}_{i+1}^{\mathbf{l}}\left(\mathcal{N}_i - I_{\{u^*\neq 0\}}, \hat{u}^*\right)$$

$$\hat{v}_i^{\mathbf{j}} = \frac{1}{b}\sum_{l=1}^{b}\Bigg[\hat{v}_{il} - \sum_{a\in\{u_1,\ldots,u_z\}} I_{[\hat{Y}_l^-(a)\leq\min\{0,\hat{Y}_l^-(c):c\neq a\}]} m\left[-\bar{Y}_l^-(a), \hat{Y}_l^-(u_1), \ldots, \hat{Y}_l^-(u_z), \bar{\Sigma}_Y^-\right]$$

$$+ \sum_{a\in\{u_1,\ldots,u_z\}} I_{[0\leq\min\{\hat{Y}_l^-(u_1),\ldots,\hat{Y}_l^-(u_l)\}]} n\left[\bar{Y}_l^-(a), \hat{Y}_l^-(c) : c \neq a, \hat{Y}_l^-(a), \bar{\Sigma}_Y^-\right]$$

$$+ \sum_{\substack{a,d\in\{u_1,\ldots,u_z\} \\ a\neq d}} I_{[\hat{Y}_l^-(a)\leq\min\{0,\hat{Y}_l^-(c):c\neq a\}]} n\left[\bar{Y}_l^-(d) - \bar{Y}_l^-(a), \hat{Y}_l^-(e) : e \neq d, \hat{Y}_l^-(d), \bar{\Sigma}_Y^-\right] \Bigg| \mathcal{Z}_i \Bigg].$$

$$(3.38)$$

Following a inductive argument analogous to that given in [19], it is straightforward to show that the global components in Equations (3.22) and (3.28) are small compared to the local component's bias and correction terms. Note that in these expressions we have assumed only one type of swing right and that only one right can be used at each exercise opportunity. Extensions to the case of different types of swing rights and/or to the use of multiple rights at each exercise opportunity is relatively straightforward.

3.5.3   Example - Bias Reduced

We now show the effectiveness of these corrected estimators in the following example. Beginning with the one dimensional case, we priced a swing option similar to the no volume choice swing option from Section 3.4.1. That is the underlying asset follows a risk neutralized price process that satisfies the following stochastic differential equation

$$dS_i = S_i\left[(r - \delta)\,dt + \sigma dZ_i\right], \qquad (3.39)$$

where r is the riskless interest rate, Z_i is a standard Brownian motion process, σ is a constant volatility parameter and the underlying asset itself pays a continuous dividend yield δ.

All simulations in this subsection were completed on the SHARCNet cluster Orca. Orca consists of AMD Opteron 6174 2.2 GHz processors (12 per socket) with 32 GB of memory (per

2 sockets) and a QDR InfiniBand interconnect. In this example the following set of parameters is held constant; expiry is 3.0 years; the risk free rate is $r = 0.05$; there is a continous dividend yield of $\delta = 0.1$; a base volume of 10 units of the underlying; the up and down swing right strikes are $K = 100.0$; and the volatility of the asset is $\sigma = 0.2$. In the one-dimensional example we include a binomial value, calculated using the forest of trees [23], for comparsion.

The swing option being priced consists of 2 swing rights that may be exercised at 3 opportunities (not including the initial time) and 10 units of the underlying. No option was given to the holder as to the volume received upon exercise. The payoff upon exercise is $10 \times \max(S_i - K, 0)$. The initial price is \$100 and there are no penalties. The results in Figure 3.8 illustrate how the bias reduced estimators significantly improve upon the uncorrected estimators. The dip in the corrected estimator as seen here in Figure 3.8 is reflective of the fact that although the corrected estimators are guaranteed to be consistent estimators to the true option value the sense of there bias is not necessarily maintained.
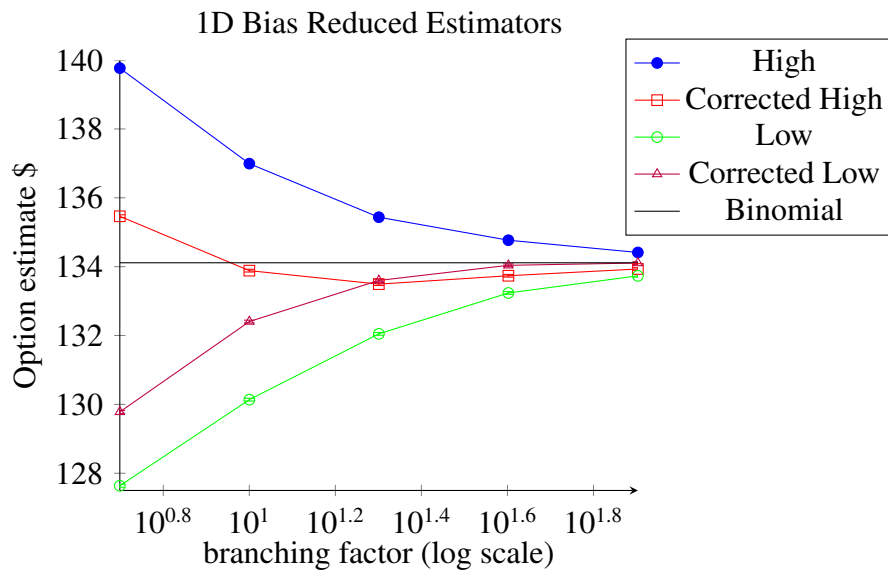
Figure 3.8: Option value estimate (\$) vs log branching factor. Here $R = 4194304 \left(\frac{5}{b} \right)$, run time for $b = 10$ is 326s and $b = 80$ is 24846s and Std. Err.$\approx$\$0.04 for all estimators and mesh sizes.

Next we give an example using 5 underlying assets while maintaining all other parameters as they were in Figure 3.8. The payoff function used is

$$10 \times \left(\max_{k=1,\dots,d} S^k - K_u, K_d - \max_{k=1,\dots,d} S^k \right)^+ \tag{3.40}$$

in the case of d underlying assets and where 10 units is the exercise volume. This is an extended example to that given in [6] and [7].

The asset prices are taken to follow a risk neutralized correlated GBM described by the stochastic differential equation

$$dS_i^k = S_i^k \left[\left(r - \delta^k \right) dt + \sigma^k dZ_i^k \right], \tag{3.41}$$

where Z_i^k is a standard Brownian motion process and where the correlation between Z^k and Z^s is ρ^{ks}. In all simulated results that follow it is assumed that $\delta^k = \delta$, $\sigma^k = \sigma$ for all k and that $\rho^{ks} = 0$ for all $k \neq s$. In addition to this we also take all assets to have the same initial value, S_0. Figure 3.9 plots the corrected and uncorrected estimators as a function of (log) sample size. As in the 1-dimensional case, the corrected estimators are significantly better than their uncorrected counterparts.
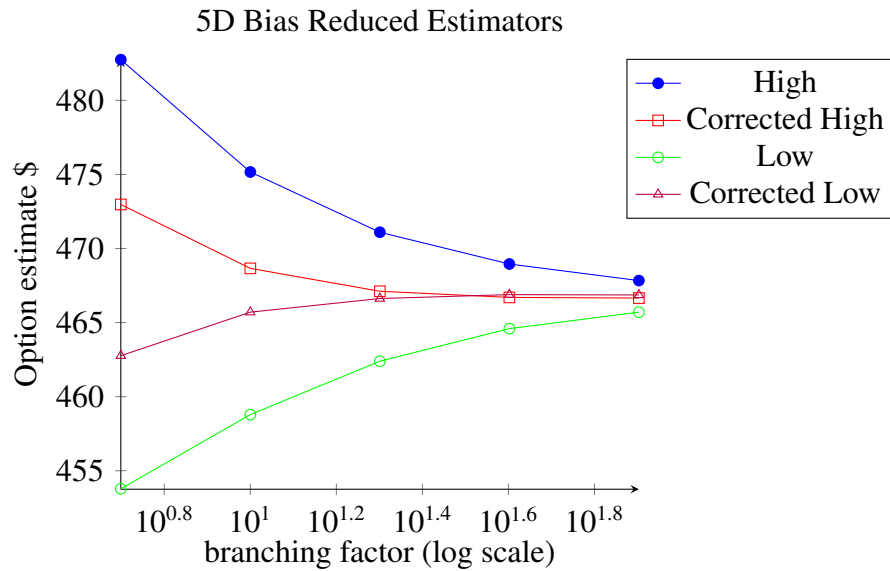
Figure 3.9: Option value estimate (\$) vs log branching factor. Here $R = 4194304 \left(\frac{5}{b} \right)$, run time for $b = 10$ is 447s and $b = 80$ is 22649s and Std. Err.$\approx$\$0.05 for all estimators and mesh sizes.

Chapter 4

Forest of Stochastic Meshes

This chapter describes the second of two methods developed in this thesis, the Forest of Stochastic Meshes [37]. This method, as described, is a variation of the previously introduced Forest of Trees [23] method which replaces the binomial trees with stochastic meshes [7]. Detailed descriptions of the dynamic program and the properties of the resulting estimators are given and the implementation of the estimators is justified by numerical results. The chapter ends by exploring enhancements to the proposed algorithms using HPC techniques [38].

4.1 Dynamic Programming

Similarly to the valuation of multiple exercise options using the forest of stochastic trees given in Chapter 3 the valuation using the forest of stochastic meshes uses the framework of [23, 26] and generalizes it by replacing their binomial and trinomial trees with stochastic meshes. The meshes are constructed identically as described in Section 1.3.6, however now we have multiple meshes each representing a different state (ie number of swing rights remaining and usage level). In addition to the high-biased mesh estimator we also construct a low-biased path estimator again in a way analogous to the one previously described.

In the Forest of Stochastic Meshes, the continuation values in (2.10) and (2.12) are replaced by stochastic mesh-type estimators. As with the original stochastic mesh technique, high- and low-biased option value estimators are constructed by using the analogous mesh and path continuation value estimators, respectively, on each mesh in the forest. The recursive equations for the high-biased mesh estimator are

$$\hat{H}_i^+(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}) = \frac{1}{b} \sum_{j=1}^{b} \omega_i^{jk} \hat{V}_{i+1}(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_{i+1}, U_{i+1}), \quad \text{and} \tag{4.1}$$

$$\hat{V}_i(\mathbf{S}_i^{j}, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[h_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) + \hat{H}_i^+(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u^* \neq 0\}}, U_i + u) \right], \tag{4.2}$$

with the terminal conditions

$$\hat{V}_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = \max_{u \in \mathcal{U}_m} \left[h_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m, u) + \tilde{\phi}(U_m + u) \right], \tag{4.3}$$

where $\hat{H}_i^+(\mathbf{S}, \mathcal{N}, U)$ and $\hat{V}_i(\mathbf{S}, \mathcal{N}, U)$ are the time-t_i, state-$\mathcal{Z}_i$ continuation and option values estimators, respectively and $h_i(\mathbf{S}, \mathcal{N}, U, u)$ is the time-t_i, state-$\mathcal{Z}_i$ payoff, with $h_i(\mathbf{S}, \mathcal{N}, U, 0) = 0$ from exercising u units, b is the mesh size and ω_i^{jk} are the mesh weights identical to those described in Section 1.3.6 and $\tilde{\phi}(U_m + u)$ is a global penalty term.

Figure 4.1 is a diagram of a section in a forest of meshes with two volume choices, u_1 and u_2. It illustrates the nodes in the forest which need to be considered when making an exercise decision given state $(\mathcal{N}, U)$. The three choices are no exercise, exercise u_1 units, and exercise u_2 units. The red nodes are the equivalent nodes on the different meshes and the blue nodes are those which are used to calculate the continuation value estimators for the different exercise choices.

The path estimator is similarly defined using the path estimator on each mesh. Namely,

$$\hat{g}_i(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}) = h_i(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) + \frac{1}{b} \sum_{k=1}^{b} \omega_i^{jk} \hat{V}_{i+1}(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_{i+1}, U_{i+1}), \tag{4.4}$$
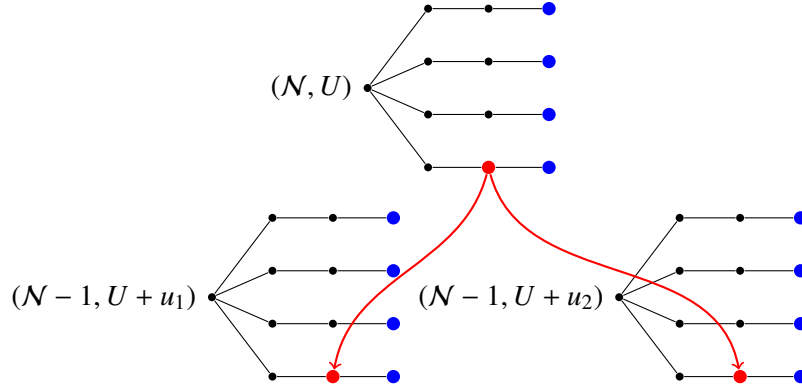
Figure 4.1: Section of a Forest of Meshes, $\mathcal{N}$ = # of exercise rights remaining, U = usage level.

$$\hat{H}_i^-(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[\hat{g}_i(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right] \quad \text{and} \tag{4.5}$$

$$\hat{v}_i(\mathbf{X}_i^j, \mathcal{N}_i, U_i) = h_i(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^*) + \hat{v}_{i+1}(\mathbf{X}_{i+1}^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + \hat{u}^*), \tag{4.6}$$

with the terminal conditions,

$$\hat{v}_m(\mathbf{X}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = \max_{u \in \mathcal{U}_m} \left[h_m(\mathbf{X}_m^{\mathbf{j}}, \mathcal{N}_m, U_m, u) + \tilde{\phi}(U_m + u) \right], \tag{4.7}$$

where $\hat{H}_i^-(\mathbf{X}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ is the continuation value, $\hat{u}^*$ is the estimated optimal exercise amount, I is an indicator function and where $\tilde{\phi}(U_m + u)$ is a global penalty factor.

4.2 Bias of Estimators

As with the forest of stochastic trees in order to justify using the mesh and path estimators to define the upper and lower bounds respectively of our confidence intervals we must prove that indeed the mesh estimator is positively biased and that the path estimator is negatively biased. In the theorems that follow the argument b that appears with the estimators refers to an arbitrary mesh and path size of b. We begin with the theorem regarding the bias of the mesh estimator.

Theorem 14 *(Mesh estimator bias) The mesh estimator $\hat{V}_0(b, \mathbf{S}_0, \mathcal{N}_0, U_0)$ is biased high. That is,*

$$E\left[\hat{V}_0(b, \mathbf{S}_0, \mathcal{N}_0, U_0) \right] \geq B_0(\mathbf{S}_0, \mathcal{N}_0, U_0) \tag{4.8}$$

for all b.

A proof of this theorem is found in Appendix C. The next the theorem states that the bias of the path estimator is non-positive.

Theorem 15 *(Path estimator bias) The path estimator* $\hat{v}_0(b, \mathbf{X}_0, \mathcal{N}_0, U_0)$ *is biased low. That is,*

$$E\left[\hat{v}_0\left(b, \mathbf{X}_0, \mathcal{N}_0, U_0\right)\right] \leq B_0\left(\mathbf{X}_0, \mathcal{N}_0, U_0\right) \tag{4.9}$$

for all b.

The proof is similar to the proof of the low bias of the path estimator in [7]. Since the exercise policy that follows from the formulation of the path estimator is not necessarily optimal and no policy can be better than the optimal policy then the path estimator defined in Equation 4.4 must be a lower bound for the option value. Intuitively, the information used for exercise decisions is independent of the values propagated and hence the estimated exercise policy is not peering into the future along the independant path. Therefore, this estimated optimal exercise policy is suboptimal along that path.

4.3 Convergence of Estimators

Similarly to the stochastic tree an advantage that the stochastic mesh method for valuing American-style options has over some other MC methods is that its estimators are consistent to the true value of the option. This property continues to hold true for the generalization to the forest of stochastic meshes. In this section we state two theorems; one for the convergence of the mesh estimator, and the other for the convergence of the path estimator. As above the argument b that appears with the estimators refers to an arbitrary mesh or path size of b with convergence being shown as $b \to \infty$.

In order to state the convergence results for the mesh estimator we require three additional moment assumptions and some additional notation. First, for $i = 1, \ldots, m$ and $k = 0, \ldots, m - i$ define

$$R(t_i, t_{i+k}) = \left(\prod_{j=0}^{k-1} \frac{f_{i+j}\left(\mathbf{S}_{i+j}^1, \mathbf{S}_{i+j+1}^1\right)}{g_{i+j+1}\left(\mathbf{S}_{i+j+1}^1\right)}\right) h_{i+k}\left(\mathbf{S}_{i+k}, \mathcal{N}_{i+k}, U_{i+k}, u\right). \tag{4.10}$$

where $\prod_{j=0}^{-1} \equiv 1$. For $\tilde{p} > p > 1$ we make the following assumptions:

Assumption 4.3.1

$$E\left[\left(\frac{g\left(t_1, \mathbf{S}_1\right)}{f\left(t_1, \mathbf{S}_1\right)}\right)|h_2\left(\mathbf{S}_2, \mathcal{N}_2, U_2, u\right)|^{\tilde{p}}\right] < \infty$$

for all $t_2 = t_1, \ldots, t_m$.

Assumption 4.3.2

$$E\left[|R(t_1, t_2)|^{\tilde{p}}\right] < \infty$$

for all $t_2 = t_1, \ldots, t_m$.

Assumption 4.3.3

$$E\left[\left(\frac{f\left(t_i, \mathbf{s}, \mathbf{S}^1_{i+1}\right)}{g_{i+1}\left(\mathbf{S}^1_{i+1}\right)}\right)^q\right] < \infty$$

for all $\mathbf{s}$ *and* $i = 0, 1, \ldots, m - 1$, *for all* $q \geq 1$.

We now proceed to the statement of the convergence theorems.

Theorem 16 *(Mesh estimator convergence) Let* $\tilde{p} > p > 1$. *Under Assumptions 4.3.1–4.3.3*

$$\mathrm{E}\left\|\hat{V}_i\left(b, \mathbf{S}_i, \mathcal{N}_i, U_i\right) - B_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i\right)\right\| \to 0,$$

as $b \to \infty$, *for all* $(\mathbf{S}_i, \mathcal{N}_i, U_i)$ *and* i. *Since convergence in the p-norm implies convergence in probability, we have that* $\hat{V}_i(b, \mathbf{S}_i, \mathcal{N}_i, U_i)$ *is consistent for* $B_i(\mathbf{S}_i, \mathcal{N}_i, U_i)$ *for all* i, $(\mathbf{S}_i, \mathcal{N}_i, U_i)$. *Taking* $i = 0$ *we have that,*

$$\mathrm{E}\left[\hat{V}_0\left(b, \mathbf{S}_0, \mathcal{N}_0, U_0\right)\right] \to B_0\left(\mathbf{S}_0, \mathcal{N}_0, U_0\right),$$

implying that the mesh estimator is asymptotically unbiased.

Theorem 17 *(Path estimator convergence) Suppose again that Assumptions 4.3.1–4.3.3 hold and that* $\mathrm{E}[|h_i(\mathbf{X}_i, \mathcal{N}_i, U_i, u)|^{1+\epsilon}] < \infty$ *for* $u \in \mathcal{U}_i$ *and all* $i = 1, \ldots, m$, *for some* $\epsilon > 0$. *Suppose also that* $\mathrm{P}[h_i(\mathbf{X}_i, \mathcal{N}_i, U_i, u^1) + [H_i(\mathbf{X}_i, \mathcal{N}_i - I_{\{u^1 \neq 0\}}, U_i + u^1) = h_i(\mathbf{X}_i, \mathcal{N}_i, U_i, u^2) + [H_i(\mathbf{X}_i, \mathcal{N}_i - I_{\{u^2 \neq 0\}}, U_i + u^2)] = 0$ *for all* $u^1, u^2 \in \mathcal{U}_i$, $u^1 \neq u^2$, $(\mathbf{S}_i, \mathcal{N}_i, U_i)$ *and for all* i. *Then*

$$E\left[\hat{v}_0\left(b, \mathbf{X}_0, \mathcal{N}_0, U_0\right)\right] \to B_0\left(\mathbf{X}_0, \mathcal{N}_0, U_0\right)$$

as $b \to \infty$, *and hence the path estimator is asymptotically unbiased.*

Proofs of Theorems 16 and 17 are found in Appendix D. Note that one of the conditions in Theorem 17 states that the optimal exercise policy is almost surely never indifferent between the exercise volume choices. This is also assumed in Theorem 13.

4.4 Numerical Results

4.4.1 One-dimension

Beginning with the one dimensional case, we have based our simulations on an underlying asset with a risk neutralized price process that satisfies the following stochastic differential equation

$$dS_i = S_i \left[(r - \delta)\, dt + \sigma dZ_i \right]. \tag{4.11}$$

In this equation, r is the riskless interest rate, Z_i is a standard Brownian motion process, σ is a constant volatility parameter and the underlying asset itself pays a continuous dividend yield δ.

These simulations were completed on the SHARCNet cluster Whale. Whale is located at the University of Waterloo and consists of Opteron 2.2 GHz processors (4 per node) with a Gigabit Ethernet interconnect. In this example the following set of parameters is held constant; expiry is 1.0 years; the risk free rate is $r = 0.05$; there is a continous dividend yield of $\delta = 0.1$; a base volume of the underlying; the up and down swing right strikes are $K_u = K_d = 40.0$; and the volatility of the asset is $\sigma = 0.2$. For comparison purposes the results in this subsection include a binomial value which is calculated using the forest of trees [23].

In Figure 4.2 the swing option being priced consists of 2 up and 2 down swing rights that may be exercised at 6 opportunities to receive (up right) or sell (down right) 1 unit of the underlying. The initial price is $40, there are no penalties. The results illustrate how the two estimators are affected by increasing the branching factor. Mainly it illustrates the properties in the previous section which states the high-biased estimator is positively biased, the low-biased estimator is negatively biased and both are consistent to the true option value. It shows a tightening in the spread between the high- and low-estimators which appears consistent with the result from the forest of trees algorithm.

4.4.2 Calibrated Forward Curve

Here we repeat the same example found in Section 3.4.2 using the forest of stochastic meshes. This model is a 1-factor model with mean reversion that is seasonally adjusted and calibrated to the forward curve. The option we value is the first one (a) in Section 4.2 of [26], which is a 2 up right swing option with each right allowing the holder to take delivery of either 1 or 2 MMBTus of natural gas. It is simplified to have 4 exercise opportunities and 4 months until

Single Assets
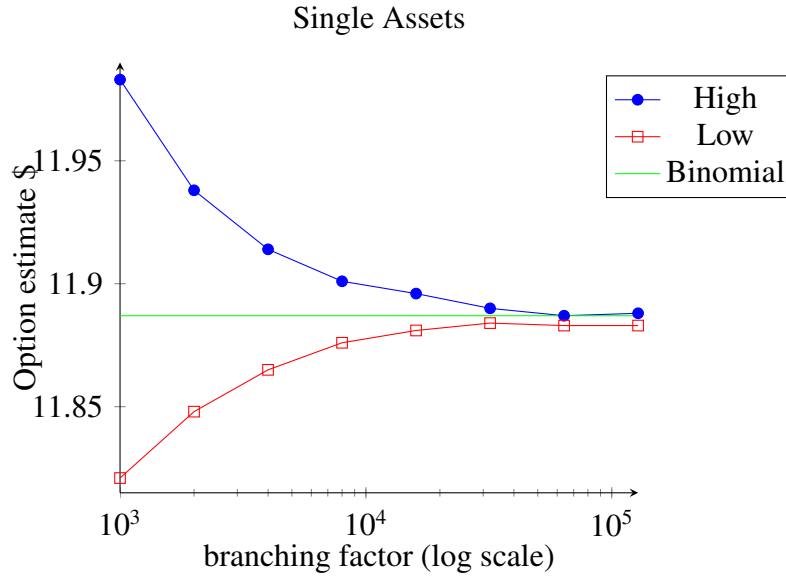
Figure 4.2: Option value estimate (\$) vs log branching factor. Here $R = 16384\left(\frac{1000}{b}\right)$ and standard error $\approx 0.01\%$ of estimator value

expiry. Upon exercise the holder gets

$$\max\left(V\left(A_i S_i - K\right), 0\right) \tag{4.12}$$

where V is the volume chosen, A_i is the seasonality factor and S_i is the deseasonalized spot price.

In Figure 4.3 we see that with a branching factor of 160 the confidence intervals for the high- and low-biased estimators begin to overlap and quickly become almost indestinguishable for higher branching factors. We see that the mesh and path estimators are high- and low-biased, respectively. Furthermore, we see evidence of estimator convergence as the mesh size increases. The number of realizations used in order to generate the results show were $\frac{160000}{b}$ and the serial computational times for mesh sizes of 160 and 2560 were approximately 45 seconds and 12 minutes respectively using a 2.1 Ghz Core 2 Duo processor. The results shown in Figure 4.3 are consistent with the results in [26] however in cases where the inclusion of more risk factors is desirable we note that the valuation method in that publication breaks down in higher dimensions.
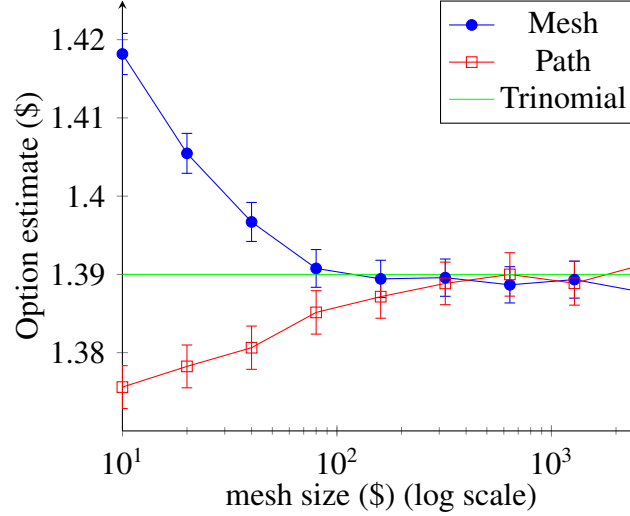
Figure 4.3: Mesh and path option-value estimates versus (log) mesh size. Approximate point-wise 95% confidence intervals for each estimate are given by the vertical bars.

4.4.3 Five-Dimensional Underlying

Here we focus on a high-dimensional example by pricing a swing option whose value depends on the price evolution of five stocks. This example is an extended version of an example appearing in [6] and [7] illustrating the stochastic tree and mesh techniques for valuing American-style options. Each asset (stock) price, S^k, follows a risk–neutralized geometric Brownian motion (GBM) described by

$$dS_t^k = S_t^k \left[\left(r - \delta^k \right) dt + \sigma^k dZ_t^k \right], \tag{4.13}$$

where r is the risk-free rate of interest, δ^k is the continuously-paid dividend rate, σ^k is the volatility and Z^k is a standard Brownian motion. Furthermore, we assume that stock prices evolve independently and that r, δ^k and σ^k are constants. Finally, we assume that all stocks have the same initial value, S_0.

Upon exercising u units, the payoff of the option is

$$u \times \left(\max(S^1, S^2, \ldots, S^5) - K_u, K_d - \max(S^1, S^2, \ldots, S^5) \right)^+, \tag{4.14}$$

where K_u and K_d are the up and down strike prices respectively. We distinguish between up and down swing rights with $\mathcal{N}_u$ and $\mathcal{N}_d$ denoting the number of such rights, respectively.

To completely specify the parameters of this example, we take an option expiry of $T = 1$ year, $r = 0.05$, $\delta^k = 0.1$ and $\sigma^k = 0.2$ for $k = 1, \ldots, 5$. In addition $K_u = K_d = 40$ there are no penalties and the volume exercised is held constant at 1 (i.e., take $u = 1$ in equation (4.14)
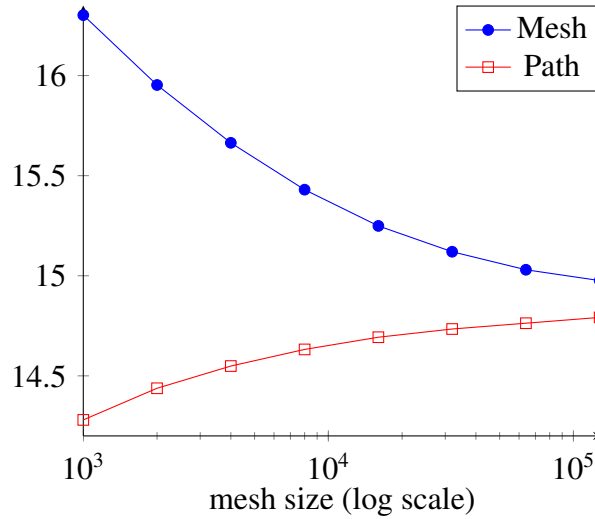
Figure 4.4: Mesh and path option-value estimates versus (log) mesh size.

and the holder has no control over the amount exercised). The number of up and down swing rights are $\mathcal{N}_u = \mathcal{N}_d = 2$ and the number of exercise opportunities is 6.

Using the above parameters the option is valued for different mesh size and number of repeated valuation pairs, (b, R), with changes to R satisfying $R = 16384(\frac{1000}{b})$. These are given in Table 4.1 along with the approximate CPU time. It is worth noting that the number of repeated valuations is chosen to produce estimator standard errors of approximately 0.01%. Furthermore, increasing R has no effect on estimator bias. Only by increasing the mesh size can estimator bias be reduced with this algorithm.

Figure 4.4 plots the mesh and path option value estimates as a function of the mesh size. The high-biased mesh estimator decreases with mesh size while the low-biased path estimator increases with mesh size. Both estimators appear to be converging towards the same value, numerical evidence of the consistency of the estimators. Due to the scale, confidence intervals for each estimator are not noticeable in Figure 4.4. With a branching factor of 128,000 the confidence interval for the true price has a width of approximately $0.18 which represents about 1.3% of the option value (approximately 14.88). Improving this interval estimate of option value requires increasing the mesh size. Addition numerical results are given in Appendix E.

Simulations for the high dimensional example were completed on the SHARCNET cluster Hound. Hound consists of Opteron 2.2 GHz processors with 128 GB of memory. Run time is determined by 3 factors, the calculation time for a single forest, the number of repeated valuations desired and the number of processors available. The timing results given in the Serial time column of Table 4.1 are given in total CPU time across all valuations. Serial refers to serial computing, MPI external refers to a naive parallel implementation (64 CPUs), MPI

internal refers to an internally parallelized implementation (64 CPUs) with serial processing of the repeated valuations and GPGPU refers to a graphics processing unit implementation with serial processing of the repeated valuations. Note that simultaneously doubling b and halving R results in a doubling of the computational time due to the quadratic scaling of computational effort with mesh size, implying a tradeoff between computational time (mesh size) and accuracy (estimator bias).

Table 4.1: Approximate total run time for various combinations of mesh size and number of repeated valuations (b, R).

(b, R)	Serial	MPI external	MPI internal	GPGPU
(4032,1)	105 sec	105 sec	1.9 sec	2.7 sec
(8000,1)	464 sec	464 sec	8.1 sec	10.7 sec
(16000,1)	40.78 min	40.78 min	0.64 min	0.70 min
(32000,1)	174.38 min	174.38 min	2.58 min	2.87 min
(4032,4096)	5 days	113 mins	132 mins	187 mins
(8000,2048)	11 days	248 mins	278 mins	366 mins
(16000,1024)	29 days	11 hrs	11 hrs	12.0 hrs
(32000,512)	62 days	23 hrs	22 hrs	24.5 hrs

4.5   Algorithmic Enhancements

To reduce the run time for the Forest of Stochastic meshes we investigate HPC methods involving both CPUs and GPGPUs. Our first method of using HPC techniques to decrease the computational time involved a naive parallel implementation of the repeated valuations across many CPUs (i.e., essentially serial farming of the independent repeated valuations). No communication between CPUs is required to evaluate a given forest. Increasing the number of CPUs by a factor of n decreases the run time by a factor of $\frac{1}{n}$. Approximate timing results from this implementation can be found in the MPI external column of Table 4.1 and is computed by dividing the serial timing results by 64, the number of processors used. Note that for a single valuation ($b = 1$), the run time is the same as that for serial processing and if the number of repeated valuations is not a multiple of the number of processors then some CPUs will be idle (see Figure 4.6).

The second method involves only CPUs and we perform an internal parallelization of the meshes, which requires communication between CPUs, and perform the repeated valuations by serial processing. There are three major computational components of the simulation i) generation of the state; ii) computation of the weight denominators; and iii) calculations of

weights, hold value, and exercise decision at each node in the mesh. In our scheme, CPUs are assigned tasks in the following manner,

(i) The underlying state vector mesh is generated completely and stored. To generate this mesh, the sample paths are generated in parallel with a single CPU generating an entire path from 0 to T. Upon path completion a CPU leapfrogs to the next sample path to be generated. For example if 256 CPUs are used the sample paths are generated in batches of 256, one for each available CPU.

(ii) As can be seen from Equation (1.74) all weights going into a given node share the same denominator. There are $b(m-2)$ of these denominators (where m is the number of time steps) as compared to b^2m weights. As such the denominators are calculated and stored prior to the evaluation of the meshes. Here a single processor is given the entire calculation of an individual denominator. Once this is complete the processor leapfrogs to the next available denominator calculation. Calculations performed in this way effectively reduce the computing work for the weights by a factor of $\frac{1}{b}$.

(iii) Beginning at expiry the meshes are evaluated. At a given time-step all work to be done for the ith node on all meshes in the forest is given to a single processor and then that processor leapfrogs to the next available set of nodes. That is the loop that is parallelized is the loop over sample paths for the high-biased mesh estimator. The loop over the forest index (indicating which mesh in the forest) is internal to the loop over sample paths for the mesh estimator. The work at each node entails: i) computing the weights needed for the hold value calculation using the stored denominators (not done at expiry); ii) the hold value calculation (not done at expiry); and iii) the exercise decision.

(iv) The results for a given time step are then collected on all processors using an MPI AllReduce before moving back to the prior time step and repeating.

Figure 4.5 plots the run time (normalized to the run time of a single CPU) against the number of CPUs for the internal parallelization scheme described above and with $(b, R) = (2000, 1024)$. There is a near perfect tradeoff between run time and number of processors. That is, increasing the number of processors by a factor of n reduces the run time by a factor of $\frac{1}{n}$. This is similar to the efficiency gains realized by the naive parallel implementation. The reason for this is that the communication time is negligible compared to the overall run time (see Table 4.2). The communication time given in Table 4.2 is the total (across all CPUs) average waiting time until other CPUs that are working on shared information are finished. This is a blocking communication protocol. Note that although the run time monotonically

decreases with the number of processors, the communication time does not. This is due to delays from system synchronization. The run time for this implementation is given in the MPI internal column of Table 4.1. For a single valuation there is a significant reduction in the run time over serial computing and the naive parallel implementation since all processors are utilized. For example, the internal parallel scheme can perform a single valuation with $b = 32000$ in approximately the same time as the naive parallel implementation takes for $b = 4032$, producing less biased estimators. Figure 4.6 plots the run time versus number of repeated valuations for $b = 16000$ and for 64 CPUs showing that the internal parallelization is better than the naive implementation over that range of R (though as R increases its advantage over the naive parallel implementation diminishes due to the communication time). However due to the algorithm being quadratic in the mesh size versus linear in the number of repeated valuations the internal parallelization scheme is preferable to the external scheme given that more is gained by increasing the mesh size as opposed to increasing the number of repeated valuations.

Table 4.2: Total run time and communication time versus number of processors for an internal CPU parallelization of the meshes with $(b, R) = (2000, 1024)$.

# procs	run time (sec)	avg comtime (sec)
2	13875.6	51.7
4	6962.7	58.0
8	3608.0	72.4
16	1858.9	91.0
20	1504.3	128.1
40	798.4	84.9

In our third method, a FirePro V9800 graphics accelerator with 1600 processing elements is used for valuation, with repeated valuations done in serial. The shared memory on the GPU allows for communication between processing units, effectively results in an internal parallelization of the meshes and is implemented by assigning parallel tasks using the OpenCL framework (an open specification application programming interface) from a single CPU as described below. Note that many modifications to the code (not detailed here due to space constraints) have resulted in a highly efficient GPGPU implementation.

(i) The state vector mesh is generated by a single OpenCL kernel call. The approach in [16] is used to generate the random numbers used for the paths, with care taken to draw unique numbers across repeated valuations. On GPUs, this method is much more efficient than other traditional methods (e.g., Mersenne Twister).

(ii) Prior to the evaluation of the meshes the mesh weight denominators (Equation (1.70)) are
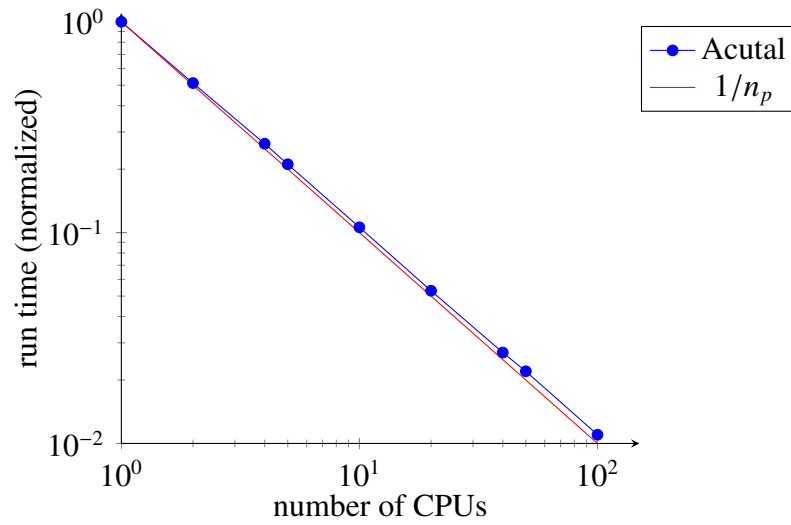
Figure 4.5: Run time (normalized to the run time of a single CPU) versus number of CPUs for an internal CPU parallelization of the meshes with $(b, R) = (2000, 2014)$. n_p appearing in the legend is the number of processors
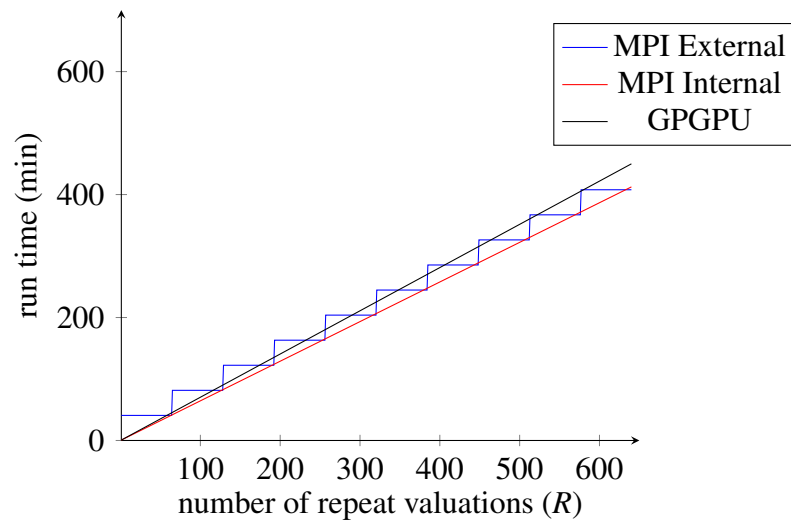
Figure 4.6: Run time versus the number of repeated valuations (R) for mesh size $b = 16000$. CPU implementations use 64 processors.

all calculated using an OpenCL kernel. Exploiting the memory hierarchy of the GPU is crucial to the performance of this phase of computation, particularly by holding batches of global values. Substantial effort has gone into optimizing how this kernel assign tasks to compute units and exploiting the particular characteristics of the memory hierarchy, particularly by performing the memory-intensive summation using fast local memory.

(iii) To evaluate the forest at a particular time step, a hold value kernel is enqueued with three indices, one each for sample path, estimator type (path/mesh) and mesh. Once the hold values are computed, an exercise decision kernel compares the hold value on the current mesh with option value on all other accessible meshes. These two kernels use an in-place update of a global buffer, requiring concurrent storage only of two time steps of option values, thus optimizing memory usage. Another kernel is used to process the expiry time step where all hold values are 0.

The run time for this implementation is given in the GPGPU column of Table 4.1, showing that in terms of computing time the GPGPU implementation is similar to the internal parallel method. Hence the above comments about the internal parallel method regarding the bias of an estimator in a fixed amount of time (e.g., a single valuation) and the right panel of Figure 4.5 also apply. Additionally note that

(i) We have shown that this valuation problem is well-suited for a GPU implementation as the computing time scales approximately with that of a naive parallel method using no communication between processors.

(ii) A single GPU costs a few hundred dollars, orders of magnitude less than the cost of CPU clusters, giving similar computational improvements for a much lower dollar cost.

(iii) Multiple GPUs can be used to perform the repeated valuations in parallel, reducing the computing time by the reciprocal of the number of GPUs.

This shows that HPC techniques can significantly reduce the run times of the computationally intensive Forest of Stochastic Meshes. The naive parallel, internal parallel and GPGPU implementations are quite effective and all give approximately the same run time when performing repeated valuations that are a multiple of the number of CPUs used. For a single valuation, the internal parallel and GPGPU methods give similar computing times and are much faster than the naive parallel and serial methods. The low cost and ability to use multiple GPUs suggests this as the method of choice. Note that additional computational efficiencies can be realized through variance and bias-reduction techniques. The methods presented here can be combined with these to yield extremely efficient pricing algorithms for multiple exercise options.

Chapter 5

Conclusions

In this thesis we have introduced two new methods for pricing multiple exercise options. The thesis concludes with general comments on the effectiveness of these methods, other research not included in this thesis, and a summary of results.

In this thesis we have introduce two new methods for pricing multiple exercise options; the forest of stochastic trees and the forest of stochastic meshes. Multiple exercise options may be considered as generalizations of American-style options as they provide the holder more than one exercise right and as such the methods we proposed, and all other currently available methods, are built off of two well known methods for pricing American-style options. Examples of multiple exercise options have become more prevalent over the past decade and appear in sectors ranging from insurance to energy industries. In Chapter 2 we motivated our discussion by briefly describing some of these specific types of contracts before focusing the remainder of our discussions on swing options which are a type of multiple exercise option prevalent in energy markets.

The methods proposed in this thesis are of particular use in cases where there are potentially a large number (≥ 3) of assets underlying the contract, if a number of risk factors are desirable for modelling the price process of these assets or if the option has a general pay-off function. This is a failing point for many of the other proposed methods for pricing these types of options. Simulation based methods are typically a way to overcome these hurdles. In these cases the Least-squares Monte Carlo and duality methods have also been modified for the pricing of swing options in [24, 32], respectively. Although the performance of these methods does not decrease with dimensionality, these methods suffer from other drawbacks. In Least-squares Monte Carlo methods one must select a set of basis functions on which to run regressions to estimate continuation values. In general only a complete (infinite) set of basis functions results in continuation value estimators that are consistent for the true option value. In practice, of course, a finite set of basis functions is used and introduces an approximation error. Continuation value estimators are consistent for the true approximation value and not the true option value [10, 11]. This approximation error can propagate backwards through the exercise opportunities and produce high- and low-biased estimators that do not converge to the same value [33].

Duality-based methods typically use regression on a finite set of basis functions to approximate the optimal martingale, implying similar issues as Least-squares Monte Carlo. These issues persist in extensions of these algorithms to the pricing of swing options.

In Chapters 3 and 4 we describe in detail the dynamic programming schemes used for the forest of stochastic trees and forest of stochasic meshes respectively. Both methods proposed result in the generation of confidence intervals for the true value of the option by utilizing both positively- and negatively-biased estimators to the true value. Detailed proofs are given in Appendices A-D to confirm both the bias and consistency of all the estimators proposed. For the forest of stochastic trees the high- and low-biased estimators are described in Equations 3.1 and 3.4 respectively and for the stochastic forest of meshes the mesh- and path-estimators are

shown in Equations 4.1 and 4.4 respectively.

Throughout the thesis effective use of examples show confirmation that our numerical implementation of these methods is correct and consistent with theoretical properties of the estimators. In both Chapters 3 and 4 we show results for the pricing of one- and five-dimensional assets which follow Geometric Brownian motion which underly a swing option with two up and two down swing rights. The results for the forest of stochastic trees are shown in figures 3.2, 3.4 and 3.5 and the results for the forest of stochastic meshes are shown in figures 4.3 and 4.4. In the case of the one-dimensional options where we take the true value to be the one generated from the forest of trees algorithm of [23] we see that all our proposed estimators have the expected biases and appear to converge to the true value. In the five-dimensional case where no true value is available we see again that the estimators appear to have appropriate biases and again seem to be converging to a common value which our analytical arguments would suggest is the true option value.

A drawback of the methods proposed in this thesis is that they are computationally intensive. For this reason a significant portion of this thesis is devoted to describing algorithmic enhancements for these methods. These include using high performance computing techniques which includes both parallel computing techniques on CPU-clusters as well as GPGPU computing techniques that take advantage of relatively inexpensive yet highly parallel graphics processor available. In Section 3.5 we briefly describe attempts to improve the efficiency of the forest of stochastic trees algorithm through the use of parallel CPU implementations. The effectiveness of these methods is obvious from Figure 3.7. Along the same lines in Section 4.5 we describe in detail two implementations of the forest of stochastic meshes, one which uses traditional parallel CPU techniques like those implemented for the forest of stochastic trees as well as another method which uses GPGPU computing to take advantage of the highly parallelizable internal structure of the stochastic mesh. Table 4.1 displays the significant time savings that can be found when using these methods. Especially interesting here is the performance of the GPGPU implementation for which one graphics processing unit costing on the order of $1000 is able to give relatively the same performance boost as a 64-processor cluster which would cost on the order of $30 000.

Additionally we also explore bias-corrected estimators for the option values which attempt to estimate the bias introduced at each time step by the estimator and then subtract this result at each exercise opportunity. These result in bias-reduce estimators for the forest of stochastic trees given by Equations 3.37 and 3.38. Numerical results shown in Figure 3.8 and 3.9 show that by implementing these estimators that we can significantly reduce the branching factor necessary to obtain the desired accuracy for the option value without significantly increasing the computational burden. The combination of the bias reduced estimators along with the HPC

improvements lead to efficient algorithms for valuing high-dimensional swing options.

Future work on the topics covered in this thesis could include a bias reduced version of the forest of stochastic meshes algorithm which would further enhance the algorithm's efficiency. Another avenue for realizing computational efficiencies is through the use of variance reduction techniques and these will be explored. There are a few possible GPU implementations for the forest of stochastic trees, a more detailed study could be conducted to determine the optimal approach for future implemenation. In terms of risk analysis, these algorithms may be modified to determine various option hedging parameters.

Appendix A

Proof of Estimator Bias - Forest of Stochastic Trees

In this appendix we show detailed proofs for the theorems stated in Section 3.2. These theorems refer to the bias of the estimators for the forest of stochastic trees. In the proofs of Theorems 9–11 we introduce the following notation. If X is a random variable, we write $\|X\|$ for the p-norm $(E|X|^p)^{1/p}$ of X. The conditional p-norm of X on $\mathcal{Z}_i$, $(E|X|^p|\mathcal{Z}_i)^{1/p}$, is denoted $\|X\|_{\mathcal{Z}_i}$. Here we include a summary of notation used for the proofs contained in this appendix as well as all subsequent appendices.

- Time is indexed by i for t_i, $i = 0, 1, \ldots, m$.

- R is the number of repeated valuations of the forest.

- b is the branching factor (forest of stochastic trees) or mesh size (forest of stochastic meshes).

- $\mathbf{S}_i^{\mathbf{j}}$ is the spot price vector at time t_i for branch $\mathbf{j} = \{j_0, j_1, \ldots, j_i\}$. For convenience we may suppress the superscript if there is no ambiguity in doing so, in these cases $\mathbf{S}_{i+1}^{j}$ refers to the time t_{i+1} along the branch path $\mathbf{j} = \{j_0, j_1, \ldots, j_i, j\}$.

- $\mathcal{Z}_i$ represents the time-t_i history of the set of state variables $(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$, where we suppress the branching history index.

- $\hat{V}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$ is the time-t_i, state-$\mathcal{Z}_i$ high estimator.

- $\hat{v}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ is the time-t_i, state-$\mathcal{Z}_i$ *leave one out* low biased estimator which does not include node l at time-t_{i+1}.

78

- $\hat{H}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)$ is the time-t_i, state-$\mathcal{Z}_i$ *leave one out* hold value estimator for exercising u units which does not include node l at time-t_{i+1},

$$\hat{H}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) = \frac{1}{b-1} \sum_{\substack{k=1 \\ k \neq l}}^{b} D_{i+1} \hat{v}_{i+1}^{\mathbf{k}}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)$$

$$\mathbf{k} = \{\mathbf{j}, k\}$$

- $\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) = h_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) + \hat{H}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)$

- $\hat{v}_i(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ is the time-t_i, state,$\mathcal{Z}_i$ low estimator,

$$\hat{v}_i = \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$$

- $\mathcal{N}_i$ is the time-t_i number of exercise rights remaining.

- U_i is the time-t_i cumulative volume.

- $\mathcal{U}_i$ is the time-t_i discretized set of available volume choices available.

$$\mathcal{U}_i = \{u_0, u_1, u_2, \ldots, u_z : z \in \mathbb{N}\}$$

  where $u_0 = 0$.

- u is the time-t_i volume exercised. Here $u \in \mathcal{U}_i$.

- D_{i+1} is the discount factor from t_{i+1} to t_i.

- $h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right)$ is the time-t_i, state-$\mathcal{Z}_i$ payoff from exercising u units with $h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, 0\right) = 0$.

- $H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$ is the time-t_i, state-$\mathcal{Z}_i$ true hold value,

$$H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) = \mathrm{E}\left[D_{i+1} B_{i+1}(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_{i+1}, U_{i+1})|\mathcal{Z}_i\right]$$

- $B_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ is the time-t_i, state-$\mathcal{Z}_i$ true option value,

$$B_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) = \max_{u \in \mathcal{U}_i} \left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right]$$

  where $I_{\{A\}}$ is the indicator function for set A.

We now proceed with the proof of Theorem 9 regarding the bias of the high bias estimator.

Proof (Proof of Theorem 9) Here we prove the more general statement that $\mathrm{E}\,[\hat{V}_i(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)|\mathcal{Z}_i] \geq$ $B_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ for $i = 0, 1, \dots, m$. The proof proceeds by backward induction. At expiry the inequality holds trivially since $\hat{V}_m(b, \mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m)$ so that $\mathrm{E}[\hat{V}_m(b, \mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m)|\mathcal{Z}_m] \geq$ $B_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m)$. We now assume the inductive hypothesis, $\mathrm{E}[\hat{V}_{i+1}(b, \mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1})|\mathcal{Z}_{i+1}] \geq$ $B_{i+1}(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1})$ and proceed to the time-t_i case. We have,

$$
\begin{aligned}
\mathrm{E} &\left[\hat{V}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)|\mathcal{Z}_i \right] \\
&= \mathrm{E}\left[\max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \right] \Bigg| \mathcal{Z}_i \right] \\
&\geq \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + \mathrm{E}\left[D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \Bigg| \mathcal{Z}_i \right] \right] \\
&= \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + \mathrm{E}\left[D_{i+1}\mathrm{E}\left[\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \Bigg| \mathcal{Z}_{i+1} \right] \Bigg| \mathcal{Z}_i \right] \right] \\
&\geq \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + \mathrm{E}\left[D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_{i+1}, U_{i+1}\right) \Bigg| \mathcal{Z}_i \right] \right] \\
&= \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \right] \\
&= B_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right),
\end{aligned}
$$

where the first equality comes from the definition of the high estimator, the first inequality comes from the conditional Jensen's inequality and note that $\mathcal{N}_{i+1} = \mathcal{N}_i - I_{\{u^* \neq 0\}}$ and $U_{i+1} = u_i + u^*$ where u^* is the value-maximizing volume choice, the second equality uses the tower law and the fact that D_{i+1} is $\mathcal{Z}_i$-measurable, and the second inequality invokes the inductive hypothesis. ∎

Next we prove Theorem 10 regarding the bias of the low estimator.

Proof (Proof of Theorem 10) As with the proof of the bias of the high estimator we prove the more general statement that $\mathrm{E}\,[\hat{v}_i(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)|\mathcal{Z}_i] \leq B_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$ for $i = 0, 1, \dots, m$ by backward induction. Again at expiry the inequality holds trivially since $\hat{v}_m(b, \mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m)$. We now assume the inductive hypothesis, $\mathrm{E}[\hat{v}_{i+1}(b, \mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1})|\mathcal{Z}_{i+1}] \leq B_{i+1}(\mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1})$. We also note that since the $\hat{v}_{il}$'s are iid we have that, $\mathrm{E}\,[\hat{v}_i|\mathcal{Z}_i] = \mathrm{E}\,[\hat{v}_{il}|\mathcal{Z}_i]$. In what follows we define $\hat{u}_l^* \in \mathcal{U}_i$ to be the volume choice which maximizes a particular $\hat{v}_{il}$. That is,

$$\hat{u}_l^* = \arg\max_{u \in \mathcal{U}_i}\left[\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) \right]. \tag{A.1}$$

Note that $\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)$ is conditionally independent of $\hat{v}_{i+1,l}(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_{i+1}, U_{i+1}, u)$ given $\mathcal{Z}_i$ and subsequently $\hat{u}_l^*$ is also independent of $\hat{v}_{i+1,l}$ given $\mathcal{Z}_i$ since it is a function of $\hat{g}_{il}$.

Now,

$$
\mathrm{E}\left[\hat{v}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right] = \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i, U_i\right) I_{\{\hat{u}_l^*=0\}} \,\Big|\, \mathcal{Z}_i\right]
$$

$$
+ \mathrm{E}\left[\left(h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) + D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_1\right)\right) I_{\{\hat{u}_l^*=u_1\}} \,\Big|\, \mathcal{Z}_i\right]
$$

$$
+ \ldots + \mathrm{E}\left[\left(h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) + D_{i+1}\hat{v}_{i+1,l}\left(i+1, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_z\right)\right) I_{\{\hat{u}_l^*=u_z\}} \,\Big|\, \mathcal{Z}_i\right]
$$

$$
= \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right] \mathrm{P}\left(\hat{u}_l^* = 0 \,\Big|\, \mathcal{Z}_i\right) + h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) \mathrm{P}\left(\hat{u}_l^* = u_1 \,\Big|\, \mathcal{Z}_i\right)
$$

$$
+ \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_1\right) \,\Big|\, \mathcal{Z}_i\right] \mathrm{P}\left(\hat{u}_l^* = u_1 \,\Big|\, \mathcal{Z}_i\right)
$$

$$
+ \ldots + h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) \mathrm{P}\left(\hat{u}_l^* = u_z \,\Big|\, \mathcal{Z}_i\right)
$$

$$
+ \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_z\right) \,\Big|\, \mathcal{Z}_i\right] \mathrm{P}\left(\hat{u}_l^* = u_z \,\Big|\, \mathcal{Z}_i\right)
$$

$$
= \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right] p_0
$$

$$
+ h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) p_1 + \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_1\right) \,\Big|\, \mathcal{Z}_i\right] p_1
$$

$$
+ \ldots + h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) p_z + \mathrm{E}\left[D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_z\right) \,\Big|\, \mathcal{Z}_i\right] p_z
$$

where in the second equality we have used the conditional independence of $\hat{g}_{il}$ and $\hat{v}_{i+1,l}$. Here $p_0 = \mathrm{P}(\hat{u}_l^* = 0|\mathcal{Z}_i)$ and $p_j = \mathrm{P}(\hat{u}_l^* = u_j|\mathcal{Z}_i)$ for $1 \le j \le z$ and $p_0 + \ldots + p_z = 1$. Thus, using the tower law, we have,

$$
\mathrm{E}\left[\hat{v}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right] = \mathrm{E}\left[\hat{v}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right]
$$

$$
= \mathrm{E}\left[D_{i+1}\,\mathrm{E}\left[\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_{i+1}\right] \,\Big|\, \mathcal{Z}_i\right] p_0
$$

$$
+ h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) p_1 + \mathrm{E}\left[D_{i+1}\,\mathrm{E}\left[\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_1\right) \,\Big|\, \mathcal{Z}_{i+1}\right] \,\Big|\, \mathcal{Z}_i\right] p_1
$$

$$
+ \ldots + h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) p_z + \mathrm{E}\left[D_{i+1}\,\mathrm{E}\left[\hat{v}_{i+1,l}\left(b, \mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_z\right) \,\Big|\, \mathcal{Z}_{i+1}\right] \,\Big|\, \mathcal{Z}_i\right] p_z
$$

$$
\le \mathrm{E}\left[D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i, U_i\right) \,\Big|\, \mathcal{Z}_i\right] p_0
$$

$$+ h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) p_1 + \mathrm{E}\left[D_{i+1} B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_1\right) \Big| \mathcal{Z}_i\right] p_1$$

$$+ \ldots + h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) p_z + \mathrm{E}\left[D_{i+1} B_{i+1}\left(\mathbf{S}_{i+1}^{\mathbf{k}}, \mathcal{N}_i - 1, U_i + u_z\right) \Big| \mathcal{Z}_i\right] p_z$$

$$= H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) p_0 + \left(h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_1\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - 1, U_i + u_1\right)\right) p_1$$

$$+ \ldots + \left(h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u_z\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - 1, U_i + u_z\right)\right) p_z$$

$$\leq \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right]$$

$$= B_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i)$$

Where the first inequality follows from the inductive hypothesis and the remaining steps follow from the definitions for B_i and H_i. ∎

Finally we proceed to the proof of Theorem 11. This theorem justifies that the high estimator is always greater than the low estimator for any realization and any b.

Proof At expiry we have that $\hat{v}_m(b, \mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = \hat{V}_m(b, \mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m^{\mathbf{j}}, \mathcal{N}_m, U_m)$ so the relation holds trivially. We now take the inductive hypothesis to be $\hat{v}_{i+1}(b, \mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1}) \leq \hat{V}_{i+1}(b, \mathbf{S}_{i+1}^{\mathbf{j}}, \mathcal{N}_{i+1}, U_{i+1})$ for $j_{i+1} = 1, \ldots, b$. Using the $\hat{g}_{il}$ as defined above we first consider the case where for a given tree,

$$\hat{u}_l^* = \arg \max_{u \in \mathcal{U}_i}\left[\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)\right], \tag{A.2}$$

is the same for all l (i.e., $\hat{u}_l^* = \hat{u}^*$, for all l).
    Then,

$$\hat{v}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) = \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$$

$$= \frac{1}{b} \sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^*\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^* \neq 0\}}, U_i + \hat{u}^*\right)\right]$$

$$\leq \frac{1}{b} \sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^*\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^* \neq 0\}}, U_i + \hat{u}^*\right)\right]$$

$$= h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^*\right) + \frac{1}{b} \sum_{l=1}^{b}\left[D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^* \neq 0\}}, U_i + \hat{u}^*\right)\right]$$

$$\leq \max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) + \frac{1}{b} \sum_{l=1}^{b}\left[D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right]\right]$$

$$= \hat{V}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$$

where the first inequality comes from the inductive hypothesis and the remaining relations come from the parameter definitions.

Next consider the case where the low estimator gives two different estimated optimal exercise amounts, $\hat{u}^1$, $\hat{u}^2$, across all l branches where $\hat{u}^1 \neq \hat{u}^2$. That is $\hat{u}_l^* = \hat{u}_1$ or $\hat{u}_l^* = \hat{u}^2$ for all $l = 1, \ldots, b$. As above we take $\hat{u}_l^*$ to be the optimal exercise amount determined by the l-th leave one out estimator, then,

$$\hat{v}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) = \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right)$$

$$= \frac{1}{b} \sum_{l=1}^{b} \left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}_l^*\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}_l^* \neq 0\}}, U_i + \hat{u}_l^*\right)\right]$$

$$= \frac{1}{b} \sum_{l=1}^{b} \left\{\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right] I_{\{\hat{u}_l^* = \hat{u}^1\}}\right.$$

$$\left. + \left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right] I_{\{\hat{u}_l^* = \hat{u}^2\}}\right\}$$

$$= \frac{\left(\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^1\}}\right) \times \left(\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right] I_{\{\hat{u}_l^* = \hat{u}^1\}}\right)}{\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^1\}}}$$

$$+ \frac{\left(\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^2\}}\right) \times \left(\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right] I_{\{\hat{u}_l^* = \hat{u}^2\}}\right)}{\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^2\}}}$$

$$= p \times \frac{\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right] I_{\{\hat{u}_l^* = \hat{u}^1\}}}{\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^1\}}}$$

$$+ (1 - p) \times \frac{\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right] I_{\{\hat{u}_l^* = \hat{u}^2\}}}{\frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^2\}}},$$

where $p = \frac{1}{b}\sum_{l=1}^{b} I_{\{\hat{u}_l^* = \hat{u}^1\}}$.

Without loss of generality, suppose that $\hat{u}_l^* = \hat{u}^1$ for $l = 1, \ldots, k$ and $\hat{u}_l^* = \hat{u}^2$ for $l = k+1, \ldots, b$. Then the above ratios become

$$\frac{\sum_{l=1}^{k}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right]}{k} \tag{A.3}$$

and

$$\frac{\sum_{l=k+1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right]}{b - k}, \tag{A.4}$$

respectively. Now for any $i^* \leq k < j^* \leq b$ we have

$$\hat{g}_{ii^*}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) > \hat{g}_{ij^*}\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right)$$

which from the definition of $\hat{g}_{il}$, implies that

$$D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{i}^*}, \mathcal{N}_i, U_i + \hat{u}^1\right) \leq D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{j}^*}, \mathcal{N}_i, U_i + \hat{u}^1\right).$$

Therefore,

$$\max_{1 \leq a \leq k}\left[D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{a}}, \mathcal{N}_i, U_i + \hat{u}^1\right)\right] \leq \min_{k+1 \leq a \leq b}\left[D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{a}}, \mathcal{N}_i, U_i + \hat{u}^1\right)\right].$$

This implies that Equation A.3

$$\frac{1}{k}\sum_{l=1}^{k}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right]$$

$$\leq \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right],$$

and similarly for Equation A.4

$$\frac{1}{b-k}\sum_{l=k+1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right]$$

$$\leq \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right].$$

Therefore

$$\hat{v}_i\left(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i\right) \leq p \times \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right]$$

$$+ (1-p) \times \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right]$$

$$\leq p \times \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right]$$

$$+ (1-p) \times \frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right]$$

$$\leq \max\left[\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^1\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^1 \neq 0\}}, U_i + \hat{u}^1\right)\right],\right.$$

$$\left.\frac{1}{b}\sum_{l=1}^{b}\left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, \hat{u}^2\right) + D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{\hat{u}^2 \neq 0\}}, U_i + \hat{u}^2\right)\right]\right]$$

$$\leq \max_{u \in \mathcal{U}_i} \left[\frac{1}{b} \sum_{l=1}^{b} \left[h_i \left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u \right) + D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u \right) \right] \right]$$

$$= \hat{V}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i),$$

where the second inequality comes from the inductive hypothesis, the third inequality is an application of Jensen's inequality, the fourth inequality comes from maximizing over a larger set, and the final equality is the definition of the high-biased estimator.

For the cases where the low estimator gives z^* distinct estimated optimal exercise amounts, $\hat{u}^1, \ldots, \hat{u}^{z^*}$, across all z branches, $z^* = 3, \ldots, z$, arguments similar to those given above (for 2 distinct estimated optimal exercise amounts) show that,

$$\hat{v}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i) \leq \hat{V}_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i).$$

Since we restrict the number of volume choices to be finite, the theorem is proven. ∎

This concludes the proofs of Theorems 9–11.

Appendix B

Proof of Estimator Convergence - Forest of Stochastic Trees

In this appendix we prove Theorems 12 and 13 found in Section 3.3. These proofs verify that both the high and low bias estimators converge to the true option value as the branching factor goes to infinity. Refer to Appendix A for notation used here. Prior to proving Theorems 12 and 13 we first state and prove the following two preliminary results.

Lemma B.0.1 *If* $\|h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u)\| < \infty$ *for all* t_i, *for some* $p \geq 1$, *then the following are true for all* $0 \leq t_i \leq t_k \leq t_m$:

$$\|B_k(\mathbf{S}_k, \mathcal{N}_k, U_k)\|_{\mathcal{Z}_i} < \infty \tag{B.1}$$

$$\sup_b \left\|\hat{V}_k(b, \mathbf{S}_k, \mathcal{N}_k, U_k)\right\|_{\mathcal{Z}_i} < \infty \tag{B.2}$$

$$\sup_b \|\hat{v}_k(b, \mathbf{S}_k, \mathcal{N}_k, U_k)\|_{\mathcal{Z}_i} < \infty \tag{B.3}$$

Proof (Proof of Lemma B.0.1) If every $h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u)$ has finite p-th moment, then each $\|h_i(\mathbf{S}_k, \mathcal{N}_k, U_k, u)\|_{\mathcal{Z}_i}$ is finite. Since the max, discounting, and conditional expectation operators preserve finiteness of moments then it follows that $\|B_k(\mathbf{S}_k, \mathcal{N}_k, U_k)\|_{\mathcal{Z}_i}$ and also $\|H_k(\mathbf{S}_k, \mathcal{N}_k, U_k)\|_{\mathcal{Z}_i}$ must also be finite.

Proceeding to B.2, fix t_i and proceed by backward induction on t_k from t_m to t_i. At expiry B.2 follows from B.1. Then for $t_k < t_m$,

$$\sup_b \left\|\hat{V}_k(b, \mathbf{S}_k, \mathcal{N}_k, U_k)\right\|_{\mathcal{Z}_i} = \sup_b \left\|\max_{u \in \mathcal{U}_k} \left[h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u) + \hat{H}_k(b, \mathbf{S}_k, \mathcal{N}_k - I_{\{u \neq 0\}}, U_k + u)\right]\right\|_{\mathcal{Z}_i}$$

$$= \sup_b \left\|\max_{u \in \mathcal{U}_k, u} \left[h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u) + \frac{1}{b} \sum_{j=1}^b D_{k+1} \hat{V}_{k+1}(b, \mathbf{S}_{k+1}^j, \mathcal{N}_k - I_{\{u \neq 0\}}, U_k + u)\right]\right\|_{\mathcal{Z}_i}$$

86

$$\leq \|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, 0)\|_{\mathcal{Z}_i} + \sup_b \left\| \frac{1}{b} \sum_{j=1}^b D_{k+1} \hat{V}_{k+1}(b, \mathbf{S}_{k+1}^j, \mathcal{N}_k, U_k) \right\|_{\mathcal{Z}_i}$$

$$+ \|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u_1)\|_{\mathcal{Z}_i} + \sup_b \left\| \frac{1}{b} \sum_{j=1}^b D_{k+1} \hat{V}_{k+1}(b, \mathbf{S}_{k+1}^j, \mathcal{N}_k - 1, U_k + u_1) \right\|_{\mathcal{Z}_i}$$

$$+ \dots + \|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u_r)\|_{\mathcal{Z}_i} + \sup_b \left\| \frac{1}{b} \sum_{j=1}^b D_{k+1} \hat{V}_{k+1}(b, \mathbf{S}_{k+1}^j, \mathcal{N}_k - 1, U_k + u_r) \right\|_{\mathcal{Z}_i}$$

$$\leq \sup_b \left\| \hat{V}_{k+1}(b, \mathbf{S}_{k+1}, \mathcal{N}_k, U_k) \right\|_{\mathcal{Z}_i} + \|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u_1)\|_{\mathcal{Z}_i} + \sup_b \left\| \hat{V}_{k+1}(b, \mathbf{S}_{k+1}, \mathcal{N}_k - 1, U_k + u_1) \right\|_{\mathcal{Z}_i}$$

$$+ \dots + \|h_k(\mathbf{S}_{k+1}, \mathcal{N}_k, U_k, u_r)\|_{\mathcal{Z}_i} + \sup_b \left\| \hat{V}_{k+1}(b, \mathbf{S}_{k+1}, \mathcal{N}_k - 1, U_k + u_r) \right\|_{\mathcal{Z}_i},$$

where $h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, 0) = 0$. This is the sum of a finite number of terms, each of which is finite.

For (iii) the proof is similar to that of (ii). ∎

The second preliminary result that we prove is as follows:

Lemma B.0.2 *Let* $a_1, \dots, a_n, b_1, \dots, b_n, c_1, \dots, c_n$ *be real numbers. Then,*

$$A_n \equiv |\max(a_1 + b_1, \dots, a_n + b_n) - \max(a_1 + c_1, \dots, a_n + c_n)| \leq 2 \sum_{i=1}^{n+1} |b_i - c_i| \equiv B_n. \qquad \text{(B.4)}$$

Proof (Proof of Lemma B.0.2) In order to prove Lemma B.0.2 we proceed by induction by considering the cases for $n = 1$ and $n = 2$. For $n = 1$,

$$A_1 = |\max(a_1 + b_1) - \max(a_1 + c_1)| = |b_1 - c_1| < B_1$$

therefore $A_1 \leq B_1$. Now, for $n = 2$

$$A_1 = |\max(a_1 + b_1, a_2 + b_2) - \max(a_1 + c_1, a_2 + c_2)|$$

Consider the following,

(i) $a_1 + b_1 > a_2 + b_2$

    (a) $a_1 + c_1 > a_2 + c_2$
       Then

$$A_2 = |a_1 + b_1 - a_1 - c_1| = |b_1 - c_1| \leq B_2$$

    (b) $a_1 + c_1 < a_2 + c_2$

Note that conditions (i) and (b) imply that

$$b_2 - b_1 < a_1 - a_2 < c_2 - c_1 \tag{B.5}$$

and we have

$$
\begin{aligned}
A_2 &= |a_1 + b_1 - a_2 - c_2| \\
&= |a_1 + b_1 - c_1 + c_1 - a_2 - c_2| \\
&\leq |b_1 - c_1| + |(a_1 - a_2) - (c_2 - c_1)| \\
&\leq |b_1 - c_1| + |(b_2 - b_1) - (c_2 - c_1)| \\
&= |b_1 - c_1| + |b_2 - c_2 - (b_1 - c_1)| \\
&\leq 2|b_1 - c_1| + |b_2 - c_2| \\
&\leq 2|b_1 - c_1| + 2|b_2 - c_2| \\
&= B_2
\end{aligned}
$$

where the first inequality comes from the triangle inequality, the second comes from Inequality B.5 and the third inequality comes from another application of the triangle inequality.

(ii) $a_2 + b_2 > a_1 + b_1$

    (a) $a_2 + c_2 > a_1 + c_1$

    Then
$$A_2 = |a_2 + b_2 - a_2 - c_2| = |b_2 - c_2| \leq B_2$$

    (b) $a_2 + c_2 < a_1 + c_1$

    Note that conditions (ii) and (b) imply that

$$b_1 - b_2 < a_2 - a_1 < c_1 - c_2 \tag{B.6}$$

    and we have

$$
\begin{aligned}
A_2 &= |a_2 + b_2 - a_1 - c_1| \\
&= |a_2 + b_2 - c_2 + c_2 - a_1 - c_1| \\
&\leq |b_2 - c_2| + |(a_2 - a_1) - (c_1 - c_2)| \\
&\leq |b_2 - c_2| + |(b_1 - b_2) - (c_1 - c_2)| \\
&= |b_2 - c_2| + |b_1 - c_1 - (b_2 - c_2)|
\end{aligned}
$$

$$\leq 2|b_2 - c_2| + |b_1 - c_1|$$

$$\leq 2|b_2 - c_2| + 2|b_1 - c_1|$$

$$= B_2$$

where again the first inequality comes from the triangle inequality, the second comes from Inequality B.6 and the third inequality comes from another application of the triangle inequality.

Therefore $A_2 \leq B_2$.

Now assume that the inductive hypothesis $A_n \leq B_n$ is true. We need to show that $A_{n+1} \leq B_{n+1}$. First define i_n and j_n such that

$$a_{i_n} + b_{i_n} = \max(a_1 + b_1, \ldots, a_n + b_n)$$

and

$$a_{i_n} + c_{i_n} = \max(a_1 + c_1, \ldots, a_n + c_n)$$

respectively. Now,

$$A_{n+1} = |\max(a_1 + b_1, \ldots, a_n + b_n, a_{n+1} + b_{n+1}) - \max(a_1 + c_1, \ldots, a_n + c_n, a_{n+1} + c_{n+1})|$$

$$= \left|\max(a_{i_n} + b_{i_n}, a_{n+1} + b_{n+1}) - \max(a_{j_n} + c_{j_n}, a_{n+1} + c_{n+1})\right|$$

Consider the following,

(i) $a_{i_n} + b_{i_n} > a_{n+1} + b_{n+1}$

    (a) $a_{j_n} + c_{j_n} > a_{n+1} + c_{n+1}$

$$A_{n+1} = |a_{i_n} + b_{i_n} - a_{j_n} - c_{j_n}|$$

$$\leq 2 \sum_{i=1}^{n} |b_i - c_i|$$

$$\leq 2 \sum_{i=1}^{n+1} |b_i - c_i|$$

$$= B_{n+1}$$

where the first inequality comes from the inductive hypothesis.

    (b) $a_{j_n} + c_{j_n} < a_{n+1} + c_{n+1}$

By the definitions of i_n and j_n and (b) we have

$$a_{i_n} + c_{i_n} \leq a_{j_n} + c_{j_n} < a_{n+1} + c_{n+1}.$$

This combined with (i) gives

$$b_{n+1} - b_{i_n} < a_{i_n} - a_{n+1} < c_{n+1} - c_{i_n} \tag{B.7}$$

Then

$$
\begin{aligned}
A_{n+1} &= |a_{i_n} + b_{i_n} - a_{n+1} - c_{n+1}| \\
&= |a_{i_n} + b_{i_n} - c_{i_n} + c_{i_n} - a_{n+1} - c_{n+1}| \\
&\leq |b_{i_n} - c_{i_n}| + |(a_{i_n} - a_{n+1}) - (c_{n+1} - c_{i_n})| \\
&\leq |b_{i_n} - c_{i_n}| + |(b_{n+1} - b_{i_n}) - (c_{n+1} - c_{i_n})| \\
&= |b_{i_n} - c_{i_n}| + |b_{n+1} - c_{n+1} - (b_{i_n} - c_{i_n})| \\
&\leq 2|b_{i_n} - c_{i_n}| + |b_{n+1} - c_{n+1}| \\
&\leq B_{n+1}
\end{aligned}
$$

where again the first inequality comes from the triangle inequality, the second comes from Inequality B.7 and the third inequality comes from another application of the triangle inequality.

(ii) $a_{i_n} + b_{i_n} < a_{n+1} + b_{n+1}$

(a) $a_{j_n} + c_{j_n} < a_{n+1} + c_{n+1}$

$$
\begin{aligned}
A_{n+1} &= |a_{n+1} + b_{n+1} - a_{n+1} - c_{n+1}| \\
&= |b_{n+1} - c_{n+1}| \\
&\leq 2 \sum_{i=1}^{n+1} |b_i - c_i| \\
&= B_{n+1}
\end{aligned}
$$

(b) $a_{j_n} + c_{j_n} > a_{n+1} + c_{n+1}$

By the definitions of i_n and j_n and (ii) we have

$$a_{j_n} + b_{j_n} \leq a_{i_n} + b_{i_n} < a_{n+1} + b_{n+1}.$$

This combined with (b) gives

$$b_{j_n} - b_{n+1} < a_{n+1} - a_{j_n} < c_{j_n} - c_{n+1} \tag{B.8}$$

Then

$$
\begin{aligned}
A_{n+1} &= |a_{n+1} + b_{n+1} - a_{j_n} - c_{j_n}| \\
&= |a_{n+1} + b_{n+1} - c_{n+1} + c_{n+1} - a_{j_n} - c_{j_n}| \\
&\leq |b_{n+1} - c_{n+1}| + |(a_{n+1} - a_{j_n}) - (c_{j_n} - c_{n+1})| \\
&\leq |b_{n+1} - c_{n+1}| + |(b_{j_n} - b_{n+1}) - (c_{j_n} - c_{n+1})| \\
&= |b_{n+1} - c_{n+1}| + |b_{j_n} - c_{j_n} - (b_{n+1} - c_{n+1})| \\
&\leq 2|b_{n+1} - c_{n+1}| + |b_{j_n} - c_{j_n}| \\
&\leq B_{n+1}
\end{aligned}
$$

where again the first inequality comes from the triangle inequality, the second comes from Inequality B.8 and the third inequality comes from another application of the triangle inequality.

Therefore $A_{n+1} \leq B_{n+1}$ and the Lemma is proven. ∎

We are now able to proceed beginning with the proof of Theorem 12 regarding the convergence of the high bias estimator.

Proof (Proof of Theorem 12) Here we take $R = 1$ and state that if the convergence holds for a single realization of the forest then it will hold for the mean of any number of realizations due to the independence of each repeated valuation. Here we prove by backward induction the more general statement $\|\hat{V}_i(b, \mathbf{S}_i, \mathcal{N}_i, U_i) - B_i(\mathbf{S}_i, \mathcal{N}_i, U_i)\|_{\mathcal{Z}_i} \to 0$ for any generic node in a given tree and for all $i = 0, \ldots, m$. At expiry the relation holds trivially since at $t_i = t_m$ we have that $\hat{V}_m(b, \mathbf{S}_m, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m, \mathcal{N}_m, U_m)$. The inductive hypothesis is taken to be $\|\hat{V}_{i+1}(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}) - B_{i+1}(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1})\|_{\mathcal{Z}_{i+1}} \to 0$.

Now,

$$
\begin{aligned}
&\left\| \hat{V}_i(b, \mathbf{S}_i, \mathcal{N}_i, U_i) - B_i(\mathbf{S}_i, \mathcal{N}_i, U_i) \right\|_{\mathcal{Z}_i} \\
&= \left\| \max_{u \in \mathcal{U}_i} \left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \frac{1}{b} \sum_{j=1}^{b} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^{j}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u \right) \right] \right. \\
&\quad \left. - \max_{u \in \mathcal{U}_i} \left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + H_i(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right] \right\|_{\mathcal{Z}_i}
\end{aligned}
$$

$$= \left\| \left| \max_{u \in \mathcal{U}_i} \left[h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u \right) + \frac{1}{b} \sum_{j=1}^{b} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u \right) \right] \right. \right.$$
$$\left. \left. - \max_{u \in \mathcal{U}_i} \left[h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u \right) + H_i \left(b, \mathbf{S}_i^j, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u \right) \right] \right| \right\|_{\mathcal{Z}_i}$$

$$\leq \left\| 2 \sum_{k=0}^{z} \frac{1}{b} \left| \sum_{j=1}^{b} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) \right| \right\|_{\mathcal{Z}_i}$$

$$\leq 2 \sum_{k=0}^{z} \left\| \frac{1}{b} \sum_{j=1}^{b} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) \right\|_{\mathcal{Z}_i}$$

$$\leq 2 \sum_{k=0}^{z} \left\| \frac{1}{b} \sum_{j=1}^{b} D_{i+1} \left[\hat{V}_{i+1} \left(i, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) - B_{i+1} \left(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) \right] \right\|_{\mathcal{Z}_i}$$

$$+ 2 \sum_{k=0}^{z} \left\| D_{i+1} B_{i+1} \left(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) \right\|_{\mathcal{Z}_i}$$

$$= 2 \sum_{k=0}^{z} \left(E_k + C_k \right)$$

where the first equality comes from the definitions of the estimator and the true value. The third step comes as a result of Lemma B.0.2, the fourth step comes from a generalization of the triangle inequality. In the final step we rewrite the expression for convience in what follows.

First we deal with the C_k's. Given $\mathcal{Z}_i$ we have that $D_{i+1} B_{i+1}(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k)$ for $j = 1, \ldots, b$ and $k = 0, \ldots, z$ are iid with means of $H_i(\mathbf{S}_i, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k)$ and finite p-norms. Then by Theorem I.4.1 of [36] we have that all C_k's in the above expression go to zero.

Next we consider the E_k's. Here we have, by the properties of p-norms and the fact that the terms being averaged are iid, that

$$E_k \leq \left\| \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) - B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u_k \neq 0\}}, U_i + u_k \right) \right\|_{\mathcal{Z}_i}$$

since E_k is bounded by the p-norm of any one of the terms being averaged. By the inductive hypothesis

$$\left\| \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1} \right) \right\|_{\mathcal{Z}_{i+1}} \to 0,$$

Where $\mathcal{N}_{i+1} = \mathcal{N}_i - I_{\{u_k \neq 0\}}$ and $U_{i+1} = U_i + u_k$.

Also by a standard condition for uniform integrability (see [36] p178) we have that

$$\left\| \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1} \right) \right\|_{\mathcal{Z}_i} \to 0, \tag{B.9}$$

provided

$$\sup_b \mathrm{E}\left[\left|\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{p+\epsilon} \middle| \mathcal{Z}_i\right] < \infty$$

for some ϵ. From Lemma B.0.1 we know that

$$\sup_b \mathrm{E}\left[\left|\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{p+\epsilon} \middle| \mathcal{Z}_i\right] < \infty$$

and that

$$\mathrm{E}\left[\left|B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{p+\epsilon} \middle| \mathcal{Z}_i\right] < \infty.$$

Thus B.9 holds for each $k = 0, \ldots, z$ and hence the result is proven. ■

The final proof of this appendix is that of Theorem 13 regarding the consistency of the low biased estimator.

Proof (Proof of Theorem 13) As with the proof of Theorem 12 we proceed by backward induction. Again at expiry the relation holds trivially since $\hat{v}_m(b, \mathbf{S}_m, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m, \mathcal{N}_i, U_i)$. The inductive hypothesis is taken to be $\|\hat{v}_{i+1}(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}) - B_{i+1}(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1})\|_{\mathcal{Z}_{i+1}} \to 0$.

Let $\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)$ be as defined at the start of Appendix A and note that, with probability one,

$$h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u^1\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u^1 \neq 0\}}, U_i + u^1\right)$$
$$\neq h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u^2\right) + H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u^2 \neq 0\}}, U_i + u^2\right),$$

for all $u^1, u^2 \in \mathcal{U}_i, u^1 \neq u^2$.

Before proceeding, we stop to make three claims:

(i) $\left\|\frac{1}{b}\sum_{l=1}^b D_{i+1}\hat{v}_{i+1}\left(b, \mathbf{S}_{i+1}^{\mathbf{l}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) - H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right\|_{\mathcal{Z}_i} \to 0$

(ii) $\left\|\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) - \left[h_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u\right) - H_i\left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right]\right\|_{\mathcal{Z}_i} \to 0$

(iii) $\left\|I_{\{\hat{u}_l^* = u\}} - I_{\{u^* = u\}}\right\|_{\mathcal{Z}_i} \to 0$

for all $u \in \mathcal{U}_i$ and where

$$\hat{u}_l^* = \arg\max_{u \in \mathcal{U}_i}\left[\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u)\right] \quad \text{and}$$

$$u^* = \arg \max_{u \in \mathcal{U}_i} \left[h_i \left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u \right) - H_i \left(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u \right) \right]$$

The proof of item (i) is the same as the proof of the corresponding step in Theorem 12. Since the estimators in (i) and (ii) differ only in the omission of one term in $\hat{g}_{il}$, similar arguments prove that (ii) also holds.

Now for (iii), if $\hat{u}_l^* = u^*$ then the result holds trivially. Now suppose that $\hat{u}_l^* = u \neq u^*$ for some $u \in \mathcal{U}_i$. Then,

$$
\begin{aligned}
\left\| I_{\{\hat{u}_l^* = u\}} - I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i} &= \left\| I_{\{\hat{u}_l^* = u\}} \right\|_{\mathcal{Z}_i} = \left[P \left(\hat{u}_l^* = u | \mathcal{Z}_i \right) \right]^{\frac{1}{p}} \\
&= \left[P \left(\hat{g}_{il}(b, \mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) \geq h_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i, U_i, u) + H_i(\mathbf{S}_i^{\mathbf{j}}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u) \right) \right]^{1/p} \\
&\to 0.
\end{aligned}
$$

Since (ii) holds and convergence in p-norm implies convergence in probability. Thus (iii) is proven.

Now proceeding from the definition of the low estimator and the true option value for all $u \in \mathcal{U}_i$,

$$
\begin{aligned}
& \left\| \hat{v}_i \left(b, \mathbf{S}_i, \mathcal{N}_i, U_i \right) - B_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i \right) \right\|_{\mathcal{Z}_i} \\
&= \left\| \frac{1}{b} \sum_{l=1}^{b} \hat{v}_{il} \left(b, \mathbf{S}_i, \mathcal{N}_i, U_i \right) - B_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i \right) \right\|_{\mathcal{Z}_i} \\
&= \left\| \frac{1}{b} \sum_{l=1}^{b} \left(D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i, U_i \right) I_{\{\hat{u}_l^* = 0\}} \right. \right. \\
& \quad + \left[h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_1 \right) + D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - 1, U_i + u_1 \right) \right] I_{\{\hat{u}_l^* = u_1\}} \\
& \quad + \ldots + \left. \left. \left[h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_z \right) + D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - 1, U_i + u_z \right) \right] I_{\{\hat{u}_l^* = u_z\}} \right) - B_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i \right) \right\|_{\mathcal{Z}_i} \\
&\leq \left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i, U_i \right) I_{\{\hat{u}_l^* = 0\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i \right) I_{\{u^* = 0\}} \right\|_{\mathcal{Z}_i} \\
& \quad + \left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - 1, U_i + u_1 \right) I_{\{\hat{u}_l^* = u_1\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i - 1, U_i + u_1 \right) I_{\{u^* = u_1\}} \right\|_{\mathcal{Z}_i} \\
& \quad + \ldots + \left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - 1, U_i + u_z \right) I_{\{\hat{u}_l^* = u_z\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i - 1, U_i + u_z \right) I_{\{u^* = u_z\}} \right\|_{\mathcal{Z}_i} \\
& \quad + \left\| h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_1 \right) I_{\{\hat{u}_l^* = u_1\}} - h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_1 \right) I_{\{u^* = u_z\}} \right\|_{\mathcal{Z}_i} \\
& \quad + \ldots + \left\| h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_z \right) I_{\{\hat{u}_l^* = u_z\}} - h_i \left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_z \right) I_{\{u^* = u_z\}} \right\|_{\mathcal{Z}_i} \quad \text{(B.10)}
\end{aligned}
$$

where the inequality in the third step is due to a generalization of the triangle inequality.

The immediate consequence of claim (iii) above is that all terms in Equation B.10 with the form

$$\left\| h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) I_{\{\hat{u}_l^* = u\}} - h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i} \to 0$$

for all $u \in \mathcal{U}_i$. Thus

$$\sum_{k=0}^{z} \left\| h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u_k) I_{\{\hat{u}_l^* = u_k\}} - h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u_k) I_{\{u^* = u_k\}} \right\|_{\mathcal{Z}_i} \to 0.$$

It remains to show that the remaining terms in B.10 converge in the p-norm. Taking one of these terms, that is, fix a $u \in \mathcal{U}_i$, we now show this converges in the p-norm to zero.

$$\left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{\hat{u}^* = u\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u_l^* = u\}} \right\|_{\mathcal{Z}_i}$$

$$= \left\| \frac{1}{b} \sum_{l=1}^{b} \left[D_{i+1} \hat{v}_{i+i,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{\hat{u}_l^* = u\}} - [D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} \right] \right.$$

$$\left. + \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i}$$

$$\leq \left\| \frac{1}{b} \sum_{l=1}^{b} \left[D_{i+1} \hat{v}_{i+i,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{\hat{u}_l^* = u\}} - D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} \right] \right\|$$

$$+ \left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i}$$

$$\leq \left\| D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \right\|_{\mathcal{Z}_i} \cdot \left\| I_{\{\hat{u}_l^* = u\}} - I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i}$$

$$+ \left\| I_{\{u^* = u\}} \right\|_{\mathcal{Z}_i} \cdot \left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \right\|_{\mathcal{Z}_i},$$

where the first step comes from adding and subtracting the same term, the second comes from applying the triangle inequality and the third step comes from factoring out common terms.

Now by (iii),

$$\left\| I_{\{\hat{u}_l^* = u\}} - I_{\{\hat{u}^* = u\}} \right\|_{\mathcal{Z}_i} \to 0,$$

by (i),

$$\left\| \frac{1}{b} \sum_{l=1}^{b} D_{i+1} \hat{v}_{i+1,l} \left(b, \mathbf{S}_{i+1}^l, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \right\|_{\mathcal{Z}_i} \to 0,$$

and we note that

$$\left\| I_{\{u^*=u\}} \right\|_{\mathcal{Z}_i} < \infty \quad \text{and}$$

$$\left\| D_{i+1}\hat{v}_{i+1,l}\left(b, \mathbf{S}^l_{i+1}, \mathcal{N}_i - I_{\{u\neq 0\}}, U_i + u\right) \right\|_{\mathcal{Z}_i} < \infty,$$

by B.3.

Hence we have proven the consistency of the low-biased estimator. ∎

Appendix C

Proof of Estimator Bias - Forest of Stochastic Meshes

In this appendix we show a detailed proof for Theorem 14 stated in Section 4.2. This theorem refers to the bias of the mesh estimators for the forest of stochastic meshes. The notation in this appendix follows from the definitions given in Appendix A. We begin the proof of Theorem 14.

Proof (Proof of Theorem 14) We note that at expiry the relation holds trivially since $\hat{V}_m(b, \mathbf{S}_m, \mathcal{N}_m, U_m) = B_m(\mathbf{S}_m, \mathcal{N}_m, U_m)$. We then proceed by backward induction at time $t_i < t_m$ with the inductive hypothesis taken to be $\mathrm{E}[\hat{V}_{i+1}(b, \mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1})|\mathcal{Z}_i] \geq B_{i+1}(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1})$. Now,

$$\mathrm{E}\left[\hat{V}_i(b, \mathbf{S}_i, \mathcal{N}_i, U_i) \,\Big|\, \mathcal{Z}_i\right]$$

$$= \mathrm{E}\left[\max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \frac{1}{b}\sum_{j=1}^{b}\frac{f_i\left(\mathbf{S}_i, \mathbf{S}_{i+1}^j\right)}{g_i\left(\mathbf{S}_{i+1}^j\right)}D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\right]\,\Big|\,\mathcal{Z}_i\right]$$

$$\geq \max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \mathrm{E}\left[\frac{1}{b}\sum_{j=1}^{b}\frac{f_i\left(\mathbf{S}_i, \mathbf{S}_{i+1}^j\right)}{g_i\left(\mathbf{S}_{i+1}^j\right)}D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\,\Big|\,\mathcal{Z}_i\right]\right]$$

$$= \max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \mathrm{E}\left[\frac{f_i\left(\mathbf{S}_i, \mathbf{S}_{i+1}^1\right)}{g_i\left(\mathbf{S}_{i+1}^1\right)}D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\,\Big|\,\mathcal{Z}_i\right]\right]$$

$$= \max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \mathrm{E}\left[\frac{f_i\left(\mathbf{S}_i, \mathbf{S}_{i+1}^1\right)}{g_i\left(\mathbf{S}_{i+1}^1\right)}\mathrm{E}\left[D_{i+1}\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\,\Big|\,\mathcal{Z}_{i+1}\right]\,\Big|\,\mathcal{Z}_i\right]\right]$$

$$\geq \max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \mathrm{E}\left[\frac{f_i\left(\mathbf{S}_i, \mathbf{S}_{i+1}^1\right)}{g_i\left(\mathbf{S}_{i+1}^1\right)}D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\,\Big|\,\mathcal{Z}_i\right]\right]$$

$$= \max_{u \in \mathcal{U}_i}\left[h_i(\mathbf{S}_i, \mathcal{N}_i, U_i, u) + D_{i+1}\mathrm{E}\left[B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right)\,\Big|\,\mathcal{Z}_i\right]\right]$$

$$= B_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i\right)$$

The first three steps come from the definition of the estimator, followed by the use of Jensen's inequality and then utilizing the fact that at each time step the mesh points are identically distributed. Next comes the law of iterative expectations which is followed by the inductive hypothesis. The sixth step comes from the identity,

$$\mathrm{E}\left[\frac{f_i\left(\mathbf{S}_i, \mathbf{S}^1_{i+1}\right)}{g_i\left(\mathbf{S}^1_{i+1}\right)} B_{i+1}\left(\mathbf{S}^1_{i+1}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \,\middle|\, \mathcal{Z}_i\right] = \int f_i\left(\mathbf{s}, \mathbf{y}\right) B_{i+1}\left(\mathbf{y}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) d\,\mathbf{y},$$

and finally the definition of the true option value. ∎

Appendix D

Proof of Estimator Convergence - Forest of Stochastic Meshes

In this appendix we give proofs of the theorems found in Section 4.3. These proofs verify that both the mesh and path estimators converge to the true option value in the limit that the mesh and path sizes go to infinity. For all proceeding results refer to Appendix A for notation. Prior to showing the proofs for Theorems 16 and 17 we first state and prove the following preliminary result:

Lemma D.0.3 *For any $p \geq 1$ if Assumption 4.3.1 holds then,*

$$\mathrm{E}\left[\left|B_i\left(\mathbf{S}_i^1, \mathcal{N}_i, U_i\right)\right|^p\right] < \infty \tag{D.1}$$

and if Assumption 4.3.2 also holds then,

$$\sup_{b \geq 1} \mathrm{E}\left[\left|\hat{V}_i\left(b, \mathbf{S}_i^1, \mathcal{N}_i, U_i\right)\right|^p\right] < \infty \tag{D.2}$$

Proof (Proof of D.1 and D.2) First consider,

$$
\begin{aligned}
|B_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i\right)| &= \left|\max_{u \in \mathcal{U}_i}\left[h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u\right) + \mathrm{E}\left[D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \,\Big|\, \mathcal{Z}_i\right]\right]\right| \\
&\leq \max_{u \in \mathcal{U}_i}\left[\left|h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u\right)\right| + \left|\mathrm{E}\left[D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u \neq 0\}}, U_i + u\right) \,\Big|\, \mathcal{Z}_i\right]\right|\right] \\
&\leq \sum_{u_i \in \mathcal{U}_i} \left|h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_i\right)\right| + \left|\mathrm{E}\left[D_{i+1}B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_i, U_i + u_i\right) \,\Big|\, \mathcal{Z}_i\right]\right|
\end{aligned} \tag{D.3}
$$

where u_i is the volume choice at time-t_i, which we have added for reasons that will become

apparent below, and $I_i = I_{\{u_i \neq 0\}}$. Now consider,

$$
\mathrm{E}\left[B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_i, U_i + u_i\right) \,\middle|\, \mathcal{Z}_i\right] = \mathrm{E}\left[\max_{u_{i+1} \in \mathcal{U}_{i+1}}\left[h_i\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_i, U_i + u_i, u_{i+1}\right)\right.\right.
$$

$$
\left.\left. + \mathrm{E}\left[D_{i+2}B_{i+2}\left(\mathbf{S}_{i+2}, \mathcal{N}_i - I_i - I_{i+1}, U_i + u_i + u_{i+1}\right)\right]\right]\,\middle|\, \mathcal{Z}_i\right]
$$

Then returning to D.3 we have the following.

$$
|B_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i\right)| \leq \sum_{u_i \in \mathcal{U}_i}\left[|h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_i\right)| + \sum_{u_{i+1} \in \mathcal{U}_{i+1}}\left[\mathrm{E}\left[|h_i\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_i, U_i + u_i, u_{i+1}\right)|\,\middle|\, \mathcal{Z}_i\right]\right.\right.
$$

$$
\left.\left. + \left|D_{i+2}\mathrm{E}\left[B_{i+2}\left(\mathbf{S}_{i+2}, \mathcal{N}_i - I_i - I_{i+1}, U_i + u_i + u_{i+1}\right)\,\middle|\, \mathcal{Z}_i\right]\right|\right]\right]
$$

$$
\leq \sum_{u_i \in \mathcal{U}_i}\left[|h_i\left(\mathbf{S}_i, \mathcal{N}_i, U_i, u_i\right)| + \sum_{u_{i+1} \in \mathcal{U}_{i+1}}\left[\mathrm{E}\left[|h_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_i, U_i + u_i, u_{i+1}\right)|\,\middle|\, \mathcal{Z}_i\right]\right.\right.
$$

$$
\left.\left. + \ldots + \sum_{u_m \in \mathcal{U}_m}\mathrm{E}\left[\left|h_m\left(\mathbf{S}_m, \mathcal{N}_i - \sum_{k=i}^{m}I_k, U_i + \sum_{k=i}^{m-1}u_k, u_m\right)\right|\,\middle|\, \mathcal{Z}_i\right]\cdots\right]\right]
$$

So $B_i(\mathbf{S}_i, \mathcal{N}_i, U_i)$ has finite p-th absolute moment if each of the terms above do, i.e. if,

$$
\int \mathrm{E}\left[|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u)|\,\middle|\, \mathcal{Z}_i\right]^p g_i(\mathbf{s})d\mathbf{s} < \infty
$$

for all $k \geq i$. By Jensen's inequality we have that,

$$
\int \mathrm{E}\left[|h_k(\mathbf{S}_k, \mathcal{N}_k, U_k, u)|\,\middle|\, \mathcal{Z}_i\right]^p g_i(\mathbf{s})d\mathbf{s}
$$

$$
\leq \int \mathrm{E}\left[|h_k\left(\mathbf{S}_k, \mathcal{N}_k, U_k, u\right)|^p\,\middle|\, \mathcal{Z}_i\right] g_i(\mathbf{s})d\mathbf{s}
$$

$$
= \int \mathrm{E}\left[|h_k\left(\mathbf{S}_k, \mathcal{N}_k, U_k, u\right)|^p\,\middle|\, \mathcal{Z}_i\right]\frac{g_i(\mathbf{s})}{f_i(\mathbf{s}, \mathbf{S}_k)}f_i(\mathbf{s}, \mathbf{S}_k)d\mathbf{s}
$$

$$
= \mathrm{E}\left[|h_k\left(\mathbf{S}_k, \mathcal{N}_k, U_k, u\right)|^p\left(\frac{g_i(\mathbf{S}_i)}{f_i(\mathbf{S}_i, \mathbf{S}_k)}\right)\right]
$$

$$
< \infty
$$

by Assumption 4.3.1.

Similarly,

$$
\left| \hat{V}_{m-k} \left(b, \mathbf{S}_{m-k}^1, \mathcal{N}_{m-k}, U_{m-k} \right) \right| \le \sum_{u_{m-k} \in \mathcal{U}_{m-k}} \Bigg[\left| h_{m-k} \left(\mathbf{S}_{m-k}, \mathcal{N}_{m-k}, U_{m-k}, u_{m-k} \right) \right|
$$

$$
+ \sum_{u_{m-k+1} \in \mathcal{U}_{m-k+1}} \left[\frac{1}{b} \sum_{j_1=1}^b \left| \omega \left(m - k, \mathbf{S}_{m-k}^1, \mathbf{S}_{m-k+1}^{j_1} \right) h_{m-k} \left(\mathbf{S}_{m-k+1}^{j_1}, \mathcal{N}_{m-k} - I_{m-k}, U_{m-k} + u_{m-k}, u_{m-k+1} \right) \right|
$$

$$
+ \ldots + \frac{1}{b^k} \sum_{u_m \in \mathcal{U}_m} \sum_{j_1=1}^b \cdots \sum_{j_k=1}^b \left| \frac{f(m - k, \mathbf{S}_{m-k}^1, \mathbf{S}_{m-k+1}^{j_1})}{g_{m-k+1}(\mathbf{S}_{m-k+1}^{j_k})} \right.
$$

$$
\times \cdots \times \frac{f_{m-1}(\mathbf{S}_{m-1}^{j_{k-1}}, \mathbf{S}_m^{j_k})}{g_m(\mathbf{S}_m^{j_k})} h_m \left(\mathbf{S}_m^{j_k}, \mathcal{N}_m - \sum_{i=m-k}^m I_i, U_i + \sum_{i=m-k}^{m-1} u_i, u_m \right) \left. \Bigg] \Bigg] \right]
$$

The p-norm of any average with identically distributed terms is bounded by any one of the terms in the average. Thus,

$$
\mathrm{E} \left[\left| \hat{V}_{m-k} \left(b, X_{m-k}^1, \mathcal{N}_{m-k}, U_{m-k} \right) \right|^p \right]^{1/p}
$$

$$
\le \mathrm{E} \left[|R(m - k, m - k)|^p \right]^{1/p} + \ldots + \mathrm{E} \left[|R(m - k, m)|^p \right]^{1/p}
$$

The right hand side is independent of the mesh size b and its finiteness is given by Assumption 4.3.2. ∎

After the proof of Lemma D.0.3 we are in a position to prove Theorem 16 regarding the convergence of the mesh estimator to the true option value.

Proof (Proof of Theorem 16) We proceed by backward induction beginning with the case at expiry where the relation holds trivially. We then state the inductive hypothesis to be

$$
\left\| \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1} \right) \right\|_{p''} \to 0
$$

for all $(\mathbf{S}_{i+1}, \mathcal{Z}_{i+1}, U_{i+1})$ and for some $p'' > p$. Now, take any $p' \in (p, p'')$, then,

$$
\left\| \hat{V}_i (b, \mathbf{S}_i, \mathcal{N}_i, U_i) - B_i (\mathbf{S}_i, \mathcal{N}_i, U_i) \right\|_{p'}
$$

$$
= \left\| \max_{u \in \mathcal{U}_i} \left[h_i (\mathbf{S}_i, \mathcal{N}_i, U_i, u) + \frac{1}{b} \sum_{j=1}^b \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u \ne 0\}}, U_i + u \right) \right] \right.
$$

$$
\left. - \max_{u \in \mathcal{U}_i} \left[h_i (\mathbf{S}_i, \mathcal{N}_i, U_i, u) + H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u \ne 0\}}, U_i + u \right) \right] \right\|_{p'}
$$

$$
\le 2 \sum_{u_i \in \mathcal{U}_i} \left\| \frac{1}{b} \sum_{j=1}^b \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} D_{i+1} \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \ne 0\}}, U_i + u_i \right) - H_i \left(\mathbf{S}_i, \mathcal{N}_i - I_{\{u_i \ne 0\}}, U_i + u_i \right) \right\|_{p'}
$$

$$\leq 2 \sum_{u_i \in \mathcal{U}_i} \left\| \frac{1}{b} \sum_{j=1}^{b} \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} D_{i+1} \left[\hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) - B_{i+1} \left(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) \right] \right\|_{p'}$$

$$+ 2 \sum_{u_i \in \mathcal{U}_i} \left\| \frac{1}{b} \sum_{j=1}^{b} \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) \right.$$

$$\left. - \mathrm{E} \left[B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) \,\middle|\, \mathcal{Z}_i \right] \right\|_{p'}$$

$$= 2 \sum_{u_i \in \mathcal{U}_i} \left(\Delta_b^{u_i} + \Delta^{u_i} \right),$$

where,

$$\Delta_b^{u_i} = \left\| \frac{1}{b} \sum_{j=1}^{b} \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} D_{i+1} \left[\hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) - B_{i+1} \left(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) \right] \right\|_{p'}$$

and,

$$\Delta^{u_i} = \left\| \frac{1}{b} \sum_{j=1}^{b} \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} B_{i+1} \left(\mathbf{S}_{i+1}^j, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) - \mathrm{E} \left[B_{i+1} \left(\mathbf{S}_{i+1}, \mathcal{N}_i - I_{\{u_i \neq 0\}}, U_i + u_i \right) \,\middle|\, \mathcal{Z}_i \right] \right\|_{p'}$$

The first inequality above is a result of applying the generalized contraction property Lemma B.0.2. Since the terms in the summands of $\Delta_b^{u_i}$ are identically distributed we have, for any $u \in \mathcal{U}_i$, $\mathcal{N}_{i+1} = \mathcal{N}_i - I_{\{u \neq 0\}}$, $U_{i+1} = U_i + u$, and

$$\Delta_b^u \leq \left\| \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^1)}{g_{i+1}(\mathbf{S}_{i+1}^1)} D_{i+1} \left[\hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) \right] \right\|_{p'}$$

$$\leq \left\| \frac{f_i(\mathbf{S}_t, \mathbf{S}_{i+1}^1)}{g_i(\mathbf{S}_{i+1}^1)} \right\|_{\frac{qp'}{q-1}} \left\| D_{i+1} \left[\hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) \right] \right\|_{qp'}$$

for any $q > 1$, where the second inequality comes from applying Hölder's inequality. Now by Assumption 4.3.3 the first factor is finite. Also, we can choose a $q > 1$ such that $qp' < p''$ since $p' < p''$. From the inductive hypothesis, we have,

$$\mathrm{E} \left[\left| \hat{V}_{i+1} \left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) - B_{i+1} \left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1} \right) \right|^{qp'} \,\middle|\, \mathcal{Z}_i \right] \to 0, \tag{D.4}$$

almost surely. Furthermore, for $\epsilon > 0$ and small enough that $p''(1 + \epsilon) \leq \tilde{p}$ (with $\tilde{p}$ as in the

statement of the theorem),

$$
\begin{aligned}
\sup_{b \geq 1} & \mathrm{E}\left[\left(\mathrm{E}\left[\left|\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{qp'} \Big| \mathcal{Z}_i\right]\right)^{1+\epsilon}\right] \\
& \leq \sup_{b \geq 1} \mathrm{E}\left[\left|\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{qp'(1+\epsilon)}\right] \\
& \leq \sup_{b \geq 1} \mathrm{E}\left[\left(\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right)\right)^{\tilde{p}}\right] + \mathrm{E}\left[\left(B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right)\right)^{\tilde{p}}\right] \\
& < \infty,
\end{aligned}
$$

where the first inequality comes from Jensen's and the properties of conditional expectations, the second inequality is from the definition of $\tilde{p}$ and the properties of p-norms and finiteness is given by Lemma D.0.3. Thus the sequence in Equation D.4 is uniformly integrable and hence,

$$
\mathrm{E}\left[\left|\hat{V}_{i+1}\left(b, \mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right) - B_{i+1}\left(\mathbf{S}_{i+1}^1, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{qp'}\right] \to 0
$$

as $b \to \infty$ and therefore $\Delta_b^u \to 0$ as $b \to \infty$ for all $u \in \mathcal{U}_i$.

We now show that $\Delta^u \to 0$ for all $u \in \mathcal{U}_i$. Note that each of the b terms in the sum are independent and identically distributed with mean,

$$
\begin{aligned}
\mathrm{E}&\left[\frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} B_{i+1}\left(\mathbf{S}_{i+1}^j, \mathcal{N}_{i+1}, U_{i+1}\right)\right] \\
&= \int \frac{f_i(\mathbf{S}_i, \mathbf{y})}{g_{i+1}(\mathbf{y})} B_{i+1}(\mathbf{y}, \mathcal{N}_{i+1}, U_{i+1}) g_{i+1}(\mathbf{y}) d\mathbf{y} \\
&= \int f_i(\mathbf{S}_i, \mathbf{y}) B_{i+1}(\mathbf{y}, \mathcal{N}_{i+1}, U_{i+1}) d\mathbf{y} \\
&= \mathrm{E}\left[B_{i+1}\left(\mathbf{S}_{i+1}, \mathcal{N}_{i+1}, U_{i+1}\right) \Big| \mathcal{Z}_i\right]
\end{aligned}
$$

So we obtain the needed result provided that,

$$
\frac{1}{b} \sum_{j=1}^b \frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} B_{i+1}\left(\mathbf{S}_{i+1}^j, \mathcal{N}_{i+1}, U_{i+1}\right)
$$

converges to its expectation in the p'-norm as $b \to \infty$. From [36] this holds provided that each of the terms in the sum have finite p'-norm. By Hölder's Theorem,

$$
\mathrm{E}\left[\left|\frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)} B_{i+1}\left(\mathbf{S}_{i+1}^j, \mathcal{N}_{i+1}, U_{i+1}\right)\right|^{p'}\right]
$$

$$\leq \mathrm{E}\left[\left(\frac{f_i(\mathbf{S}_i, \mathbf{S}_{i+1}^j)}{g_{i+1}(\mathbf{S}_{i+1}^j)}\right)^{\frac{\bar{p}p'}{\bar{p}-p'}}\right]^{\frac{\bar{p}-p'}{\bar{p}}} \mathrm{E}\left[B_{i+1}\left(\mathbf{S}_{i+1}^j, \mathcal{N}_{i+1}, U_{i+1}\right)\right]^{\frac{p'}{\bar{p}}}. \tag{D.5}$$

Here the first factor in Equation D.5 is finite by Assumption 4.3.3 and the finiteness of the second term follows from Equation D.1 in Lemma D.0.3. Thus we have shown that for all $u \in \mathcal{U}_i$,

$$\sum_{u_i \in \mathcal{U}_i}\left(\Delta_b^{u_i} + \Delta^{u_i}\right) \to 0 \quad \text{as} \quad b \to \infty. \quad \blacksquare$$

The final proof of this appendix is the proof of the convergence of the path estimator as stated in Theorem 17.

Proof (Proof of Theorem 17) Let $(\hat{\tau}_b^1, \hat{u}_b^1), \ldots, (\hat{\tau}_b^N, \hat{u}_b^N)$ be the estimated mesh exercise policy with mesh size b and $(\tau^1, u^1), \ldots, (\tau^N, u^N)$ be the optimal exercise policy. Then the estimated exercise policy differs from the optimal one if and only if $(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)$ for at least one $n = 1, \ldots, N$.

Thus, we show that the probability that the estimated exercise policy differs from the optimal one goes to zero as the mesh size goes to infinity. That is,

$$\mathrm{P}\left[\bigcup_{n=1}^{N}\left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)\right]\right] \to 0 \quad \text{as} \quad b \to \infty. \tag{D.6}$$

Note that,

$$\mathrm{P}\left[\bigcup_{n=1}^{N}\left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau_b^n, u_b^n)\right]\right] \leq \sum_{n=1}^{N}\mathrm{P}\left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)\right].$$

So Equation D.6 holds if $\mathrm{P}[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)] \to 0$ for all $n = 1, \ldots, N$.

Now,

$$\mathrm{P}[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau_b^n, u_b^n)] = \mathrm{P}[\hat{\tau}_b^n \neq \tau^n] + \mathrm{P}[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u_b^n]$$

and we now show separately that,

$$\mathrm{P}\left[\hat{\tau}_b^n \neq \tau^n\right] \to 0 \quad \text{and,} \tag{D.7}$$

$$\mathrm{P}\left[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u_b^n\right] \to 0, \tag{D.8}$$

as $b \to \infty$ for all $n = 1, \ldots, N$.

Note that

$$\mathrm{P}\left[\hat{\tau}_b^n \neq \tau^n\right]$$

$$
= \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) < \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) > \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u\right)\right]\right], \text{ for some } i = n, \ldots, m\right\}
$$

$$
+ \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) > \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(b, \mathbf{S}_i, n, U_i\right) < \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u\right)\right]\right], \text{ for some } i = n, \ldots, m\right\}
$$

$$
\leq \sum_{i=n}^m \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) < \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) > \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right\}
$$

$$
+ \sum_{i=n}^m \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) > \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) < \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right\}.
$$

For convience we define

$$
\hat{u}_i = \arg \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i\right)\right] \quad \text{and} \tag{D.9}
$$

$$
u_i^* = \arg \max_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(b, \mathbf{S}_i, n-1, U_i\right)\right]. \tag{D.10}
$$

Then

$$
\mathrm{P}\left[\hat{\tau}_b^n \neq \tau^n\right]
$$

$$
\leq \sum_{i=n}^m \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) < \left[h_i\left(\mathbf{S}_i, n, U_i, \hat{u}_i\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + \hat{u}_i\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) > \left[h_i\left(\mathbf{S}_i, n, U_i, \hat{u}_i\right) + H_i\left(\mathbf{S}_i, n-1, U_i + \hat{u}_i\right)\right]\right]\right\}
$$

$$
+ \sum_{i=n}^m \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) > \left[h_i\left(\mathbf{S}_i, n, U_i, u_i^*\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u_i^*\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) < \left[h_i\left(\mathbf{S}_i, n, U_i, u_i^*\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u_i^*\right)\right]\right]\right\}
$$

$$
\leq \sum_{i=n}^m \sum_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) < \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right.
$$

$$
\left. \cap \left[H_i\left(\mathbf{S}_i, n, U_i\right) > \left[h_i\left(\mathbf{S}_i, n, U_i, u\right) + H_i\left(\mathbf{S}_i, n-1, U_i + u\right)\right]\right]\right\}
$$

$$
+ \sum_{i=n}^m \sum_{u \in \mathcal{U}_i \setminus \{\emptyset\}} \mathrm{P}\left\{\left[\hat{H}_i\left(b, \mathbf{S}_i, n, U_i\right) > h_i\left(\mathbf{S}_i, n, U_i, u\right) + \hat{H}_i\left(b, \mathbf{S}_i, n-1, U_i + u\right)\right]\right.
$$

$$\cap \, [H_i \, (\mathbf{S}_i, n, U_i) < [h_i \, (\mathbf{S}_i, n, U_i, u) + H_i \, (\mathbf{S}_i, n - 1, U_i + u)]]\}$$

$$= \sum_{i=n}^{m} \sum_{u \in \mathcal{U}_i \backslash \{\emptyset\}} \mathrm{P} \left\{ \left(\left[\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) < \left[h_i \, (\mathbf{S}_i, n, U_i, u) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u) \right] \right] \right. \right.$$

$$\cap \, [H_i \, (\mathbf{S}_i, n, U_i) > [h_i \, (\mathbf{S}_i, n, U_i, u) + H_i \, (\mathbf{S}_i, n - 1, U_i + u)]])$$

$$\cup \left(\left[\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) > \left[h_i \, (\mathbf{S}_i, n, U_i, u) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u) \right] \right] \right.$$

$$\cap \, [H_i \, (\mathbf{S}_i, n, U_i) < [h_i \, (\mathbf{S}_i, n, U_i, u) + H_i \, (\mathbf{S}_i, n - 1, U_i + u)]]) \} \, .$$

Under the hypothesis of the theorem, boundaries are hit with probability zero. Therefore almost surely there exists $\epsilon > 0$ for which

$$|H_i \, (\mathbf{S}_i, n, U_i) - [h_i \, (\mathbf{S}_i, n, U_i, u) + H_i \, (\mathbf{S}_i, n - 1, U_i + u)]| > \epsilon,$$

for $i = n, \ldots, m, u \in \mathcal{U}_i \backslash \{\emptyset\}$. If

$$H_i(\mathbf{S}_i, n, U_i) > h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n - 1, U_i + u)$$

then

$$H_i(\mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n - 1, U_i + u) - \epsilon > h_i(\mathbf{S}_i, n, U_i, u)$$

which implies

$$\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) < h_i \, (\mathbf{S}_i, n, U_i, u) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u)$$

$$\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) < H_i \, (\mathbf{S}_i, n, U_i) - H_i \, (\mathbf{S}_i, n - 1, U_i + u) - \epsilon + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u)$$

$$\Longleftrightarrow \epsilon < H_i \, (\mathbf{S}_i, n, U_i) - \hat{H}_i \, (b, \mathbf{S}_i, n, U_i) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u) - H_i \, (\mathbf{S}_i, n - 1, U_i + u) \, .$$

Similarly if

$$H_i(\mathbf{S}_i, n, U_i) < h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n - 1, U_i + u)$$

then

$$h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n - 1, U_i + u) - H_i(\mathbf{S}_i, n, U_i) > \epsilon$$

implying that

$$\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) > h_i \, (\mathbf{S}_i, n, U_i, u) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u)$$

$$\hat{H}_i \, (b, \mathbf{S}_i, n, U_i) > H_i \, (\mathbf{S}_i, n, U_i) - H_i \, (\mathbf{S}_i, n - 1, U_i + u) + \epsilon + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u)$$

$$\Longleftrightarrow \epsilon < \hat{H}_i \, (b, \mathbf{S}_i, n, U_i) - H_i \, (\mathbf{S}_i, n, U_i) + H_i \, (\mathbf{S}_i, n - 1, U_i + u) + \hat{H}_i \, (b, \mathbf{S}_i, n - 1, U_i + u) \, .$$

Therefore by combining these results we have,

$$\left|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i) + H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\right| > \epsilon \quad \text{(D.11)}$$

for $i = n, \ldots, m$, $u \in \mathcal{U}_i \backslash \{\emptyset\}$.

Therefore,

$$
\begin{aligned}
&\mathrm{P}\left[\hat{\tau}_b^n \neq \tau^n\right] \\
&\leq \sum_{i=n}^m \sum_{u \in \mathcal{U}_i \backslash \{\emptyset\}} \mathrm{P}\Big\{ \big|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i) \\
&\qquad\qquad + H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\big| > \epsilon \Big\} \\
&\leq \sum_{i=n}^m \sum_{u \in \mathcal{U}_i \backslash \{\emptyset\}} \mathrm{P}\Big\{ \big|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i)\big| \\
&\qquad\qquad + \big|H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\big| > \epsilon \Big\} \\
&\leq \sum_{i=n}^m \sum_{u \in \mathcal{U}_i \backslash \{\emptyset\}} \mathrm{P}\Big\{ \big|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i)\big| > \epsilon \Big\} \\
&\qquad\qquad + \mathrm{P}\Big\{ \big|H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\big| > \epsilon \Big\}. \quad \text{(D.12)}
\end{aligned}
$$

Note that ϵ depends on the path $(\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_m)$ and (n, U_i). Call this information $\bar{\mathcal{Z}}_i$ and note that ϵ is independent of the information used to generate the mesh. Since convergence in p-norm implies convergence in probability, we have from Theorem 16 that,

$$
\begin{aligned}
&\mathrm{P}\left\{ \left|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i)\right| > \epsilon \,\Big|\, \bar{\mathcal{Z}}_i \right\} \to 0 \quad \text{and} \\
&\mathrm{P}\left\{ \left|H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\right| > \epsilon \,\Big|\, \bar{\mathcal{Z}}_i \right\} \to 0
\end{aligned}
$$

almost surely for $i = n, \ldots, m$, $u \in \mathcal{U}_i \backslash \{\emptyset\}$.

The Dominated Convergence Theorem then implies that,

$$
\begin{aligned}
&\mathrm{P}\left\{ \left|\hat{H}_i(b, \mathbf{S}_i, n, U_i) - H_i(\mathbf{S}_i, n, U_i)\right| > \epsilon \right\} \to 0 \quad \text{and} \\
&\mathrm{P}\left\{ \left|H_i(\mathbf{S}_i, n-1, U_i+u) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i+u)\right| > \epsilon \right\} \to 0
\end{aligned}
$$

for all $i = n, \ldots, m$, $u \in \mathcal{U}_i \backslash \{\emptyset\}$.

This together with Equation D.12 proves that

$$\mathrm{P}\left[\hat{\tau}_b^n \neq \tau^n\right] \to 0 \quad \text{as} \quad b \to \infty.$$

for all $n = 1, \ldots, \mathcal{N}$.

Now we show that

$$P\left[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u^n\right] \to 0.$$

To begin

$$P\left[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u^n\right]$$

$$= P\left\{\left(\left[\hat{H}_i(b, \mathbf{S}_i, n, U_i) < \max_{u \in \mathcal{U}_i \setminus \{0\}}\left[h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u)\right]\right]\right.\right.$$

$$\cap\left[H_i(\mathbf{S}_i, n, U_i) < \max_{u \in \mathcal{U}_i \setminus \{0\}}[h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u)]\right] \text{ for } i = \tau^n\right)$$

$$\left.\cap\, [\hat{u}_b^n \neq u^n]\,\right\}$$

$$\leq P\left\{\left(\left[\hat{H}_i(b, \mathbf{S}_i, n, U_i) < \max_{u \in \mathcal{U}_i \setminus \{0\}}\left[h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u)\right]\right]\right.\right.$$

$$\cap\left[H_i(\mathbf{S}_i, n, U_i) < \max_{u \in \mathcal{U}_i \setminus \{0\}}[h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u)]\right] \text{ for some } i = n, \ldots, m\right)$$

$$\left.\cap\, [\hat{u}_b^n \neq u^n]\right\}$$

$$\leq \sum_{i=n}^{m} P\left\{\left[\hat{H}_i(b, \mathbf{S}_i, n, U_i) < h_i(\mathbf{S}_i, n, U_i, \hat{u}_i) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + \hat{u}_i)\right]\right.$$

$$\cap\, [H_i(\mathbf{S}_i, n, U_i) < h_i(\mathbf{S}_i, n, U_i, u_i^*) + H_i(\mathbf{S}_i, n-1, U_i + u_i^*)]$$

$$\left.\cap\, [\hat{u}_i \neq u_i^*]\right\}$$

$$\leq \sum_{i=n}^{m} \sum_{\substack{u,v \in \mathcal{U}_i \setminus \{0\} \\ u \neq v}} P\left\{\left[\hat{H}_i(b, \mathbf{S}_i, n, U_i) < h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u)\right]\right.$$

$$\cap\, [H_i(\mathbf{S}_i, n, U_i) < h_i(\mathbf{S}_i, n, U_i, v) + H_i(\mathbf{S}_i, n-1, U_i + v)]$$

$$\cap\left[h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) \geq h_i(\mathbf{S}_i, n, U_i, v) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v)\right]$$

$$\cap\, [h_i(\mathbf{S}_i, n, U_i, v) + H_i(\mathbf{S}_i, n-1, U_i + v) \geq h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u)]\}$$

$$\leq \sum_{i=n}^{m} \sum_{\substack{u,v \in \mathcal{U}_i \setminus \{0\} \\ u \neq v}} P\left\{\left[h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u)\right.\right.$$

$$\left.\geq h_i(\mathbf{S}_i, n, U_i, v) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v)\right]$$

$$\cap\, [h_i(\mathbf{S}_i, n, U_i, v) + H_i(\mathbf{S}_i, n-1, U_i + v) \geq h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u)]\}. \quad \text{(D.13)}$$

Again under the hypothesis of the theorem the boundaries are hit with probability zero. So

there exists some $\epsilon > 0$ almost surely such that

$$|h_i(\mathbf{S}_i, n, U_i, v) + H_i(\mathbf{S}_i, n-1, U_i + v) - [h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u)]| > \epsilon,$$

for all $i = n, \ldots, m$ and $u, v \in \mathcal{U}_i \backslash \{\emptyset\}$, $u \neq v$.

Now if

$$h_i(\mathbf{S}_i, n, U_i, v) + H_i(\mathbf{S}_i, n-1, U_i + v) \geq h_i(\mathbf{S}_i, n, U_i, u) + H_i(\mathbf{S}_i, n-1, U_i + u),$$

then,

$$h_i(\mathbf{S}_i, n, U_i, v) - h_i(\mathbf{S}_i, n, U_i, u) > \epsilon + H_i(\mathbf{S}_i, n-1, U_i + u) - H_i(\mathbf{S}_i, n-1, U_i + v)$$
$$\iff h_i(\mathbf{S}_i, n, U_i, u) - h_i(\mathbf{S}_i, n, U_i, v) < H_i(\mathbf{S}_i, n-1, U_i + v) - H_i(\mathbf{S}_i, n-1, U_i + u) - \epsilon$$

and

$$\left[h_i(\mathbf{S}_i, n, U_i, u) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) \geq h_i(\mathbf{S}_i, n, U_i, v) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v) \right]$$
$$= \left[h_i(\mathbf{S}_i, n, U_i, u) - h_i(\mathbf{S}_i, n, U_i, v) \geq \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) \right]$$
$$\subset \left[H_i(\mathbf{S}_i, n-1, U_i + v) - H_i(\mathbf{S}_i, n-1, U_i + u) - \epsilon > \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) \right]$$
$$= \left[H_i(\mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v) + \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) - H_i(\mathbf{S}_i, n-1, U_i + u) > \epsilon \right]$$
$$\subset \left(\left[|H_i(\mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v)| > \epsilon \right] \right.$$
$$\left. \cup \left[|\hat{H}_i(b, \mathbf{S}_i, n-1, U_i + u) - H_i(\mathbf{S}_i, n-1, U_i + u)| > \epsilon \right] \right).$$

Thus

$$\mathrm{P}\left[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u^n \right] \leq \sum_{i=n}^{m} \sum_{\substack{u,v \in \mathcal{U}_i \backslash \{\emptyset\} \\ u \neq v}} \mathrm{P}\left\{ |H_i(\mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v)| > \epsilon \right\}$$
$$+ \mathrm{P}\left\{ |\hat{H}_i(b, \mathbf{S}_i, n-1, U_i, u) - H_i(\mathbf{S}_i, n-1, U_i + u)| > \epsilon \right\}. \tag{D.14}$$

As above, note that ϵ depends on the path history $(\mathbf{S}_0, \mathbf{S}_1, \ldots, \mathbf{S}_m)$ and (n, U_i), which we have previously defined to be $\bar{\mathcal{Z}}_i$, but is independent of the mesh. Since convergence in p-norm implies convergence in probability, we have from Theorem 16 that

$$\mathrm{P}\left\{ |H_i(\mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v)| > \epsilon \,\middle|\, \bar{\mathcal{Z}}_i \right\} \to 0 \quad \text{and}$$

$$\mathrm{P}\left\{ |\hat{H}_i(b, \mathbf{S}_i, n-1, U_i, u) - H_i(\mathbf{S}_i, n-1, U_i + u)| > \epsilon \,\middle|\, \bar{\mathcal{Z}}_i \right\} \to 0,$$

for all $i = n, \ldots, m$, $u, v \in \mathcal{U}_i \setminus \{\emptyset\}$, $u \neq v$. The Dominated Convergence Theorem then implies that,

$$P\left\{ \left| H_i(\mathbf{S}_i, n-1, U_i + v) - \hat{H}_i(b, \mathbf{S}_i, n-1, U_i + v) \right| > \epsilon \right\} \to 0 \quad \text{and}$$

$$P\left\{ \left| \hat{H}_i(b, \mathbf{S}_i, n-1, U_i, u) - H_i(\mathbf{S}_i, n-1, U_i + u) \right| > \epsilon \right\} \to 0,$$

for all $i = n, \ldots, m$, $u, v \in \mathcal{U}_i \setminus \{\emptyset\}$, $u \neq v$. This together with Equation D.14 gives

$$P\left[\hat{\tau}_b^n = \tau^n, \hat{u}_b^n \neq u^n \right] \to 0 \quad \text{as} \quad b \to \infty,$$

for all $n = 1, \ldots, \mathcal{N}$. Recall that

$$P\left[\bigcup_{n=1}^{\mathcal{N}} \left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n) \right] \right] \leq \sum_{n=1}^{\mathcal{N}} P\left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n) \right]$$

and therefore

$$P\left[\bigcup_{n=1}^{\mathcal{N}} \left[(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n) \right] \right] \to 0 \quad \text{as} \quad b \to \infty.$$

It is left to show the asymptotic unbiasedness of the path estimator. We make the additional assumption that the payoff functions at the estimated and true optimal exercise times $h_{\hat{\tau}_b^n}$ and h_{τ^n}, respectively, are non-negative. Following from the low biasedness of the path estimator we have,

$$0 \leq B_0(\mathbf{S}_0, \mathcal{N}_0, U_0) - \mathrm{E}\left[\hat{v}_0(b, \mathbf{S}_0, \mathcal{N}_0, U_0) \right]$$

$$= \mathrm{E}\left[\sum_{n=1}^{\mathcal{N}_0} h_{\tau^n}(\mathbf{S}_{\tau^n}, \mathcal{N}_{\tau^n}, U_{\tau^n}, u^n) - \sum_{n=1}^{\mathcal{N}_0} h_{\hat{\tau}_b^n}(\mathbf{S}_{\hat{\tau}_b^n}, \mathcal{N}_{\hat{\tau}_b^n}, U_{\hat{\tau}_b^n}, \hat{u}_b^n) \right]$$

$$= \mathrm{E}\left[\left(\sum_{n=1}^{\mathcal{N}_0} h_{\tau^n}(\mathbf{S}_{\tau^n}, \mathcal{N}_{\tau^n}, U_{\tau^n}, u^n) - \sum_{n=1}^{\mathcal{N}_0} h_{\hat{\tau}_b^n}(\mathbf{S}_{\hat{\tau}_b^n}, \mathcal{N}_{\hat{\tau}_b^n}, U_{\hat{\tau}_b^n}, \hat{u}_b^n) \right) I_{\{\bigcup_{n=1}^{\mathcal{N}_0} [(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)]\}} \right]$$

$$\leq \mathrm{E}\left[\left(\sum_{n=1}^{\mathcal{N}_0} h_{\tau^n}(\mathbf{S}_{\tau^n}, \mathcal{N}_{\tau^n}, U_{\tau^n}, u^n) \right) I_{\{\bigcup_{n=1}^{\mathcal{N}_0} [(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)]\}} \right]$$

$$\leq \mathrm{E}\left[\left| \sum_{n=1}^{\mathcal{N}_0} h_{\tau^n}(\mathbf{S}_{\tau^n}, \mathcal{N}_{\tau^n}, U_{\tau^n}, u^n) \right|^{1+\epsilon} \right]^{\frac{1}{1+\epsilon}} \times P\left[\bigcup_{n=1}^{\mathcal{N}_0} [(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)] \right]^{\frac{\epsilon}{1+\epsilon}}$$

$$\leq \left(\sum_{n=1}^{\mathcal{N}_0} \mathrm{E}\left[\left| \sum_{n=1}^{\mathcal{N}_0} h_{\tau^n}(\mathbf{S}_{\tau^n}, \mathcal{N}_{\tau^n}, U_{\tau^n}, u^n) \right|^{1+\epsilon} \right] \right)^{\frac{1}{1+\epsilon}} \times P\left[\bigcup_{n=1}^{\mathcal{N}_0} [(\hat{\tau}_b^n, \hat{u}_b^n) \neq (\tau^n, u^n)] \right]^{\frac{\epsilon}{1+\epsilon}}$$

$$\to 0$$

where the first equality comes from the definition of v_0 and optimal and estimated exercise strategies and the second comes from multiplying by the indicator function. The first inequality comes from the sum of the optimal exercises being greater then the sum of the suboptimal exercises, the second comes from Hölder's inequality. On the line each of the terms in the sum is finite by the assumption in the Theorem and the second was shown to go to zero above.

Therefore we have that $E[\hat{v}_0(b, \mathbf{S}_0, \mathcal{N}_0, U_0)] \to B_0(\mathbf{S}_0, \mathcal{N}_0, U_0)$. ∎

Appendix E

Tables of Results - Forest of Stochastic Meshes

This Appendix contains data results from a SHARCNet Large Dedicated Resources project. The swing option priced had the same parameters as the one in Section 4.4.3 unless otherwise indicated that they have been varied. Run times listed are approximate aggregate time for all repeated valuations, all results listed were generated using serial farming of repeated valuations of a serial version of the forest of stochastic meshes on the SharcNet cluster Whale which consists of 2.2 GHz Opteron processor with 2Gb of memory per processor. Where listed the memory usage is the approximate memory required during a single valuation of the forest. The errors listed in the tables that follow are the standard error for each estimator representing one standard deviation.

We note here that relative to the results of [7] the spread in the high- and low-biased confidence intervals found here are large. This is due to not having implemented similar control variate techniques used in [7] or any other variance reduction technique. This is not to say that these enhancements are not applicable to the methods used here, indeed similar techniques could be incorported into our algorithm to improve its efficiency.

In Figure E.1 it is observed that the path estimator value actually decreases with the number of exercise opportunities. This may be explained by considering that the more exercise opportunities the option has the more chances there is for the path estimator to make suboptimal decisions. Similarly with the mesh estimator the more exercise opportunities leads to more chances for the mesh estimator to make better then optimal decisions which leads to the value increasing. However, it is hard to distinguish this effect for the mesh estimator since the true price also increases.

Exercise Opps	Mesh	Error	Path	Error	Point	Time (days)
12	35.147	0.005	26.845	0.004	30.996	104
24	42.857	0.006	27.754	0.003	35.306	211
36	51.411	0.007	26.415	0.003	38.913	300
50	59.005	0.008	24.794	0.003	41.900	422

Table E.1: $\mathcal{N}_u = \mathcal{N}_d = 4$, $b = 8000$, $R = 2048$
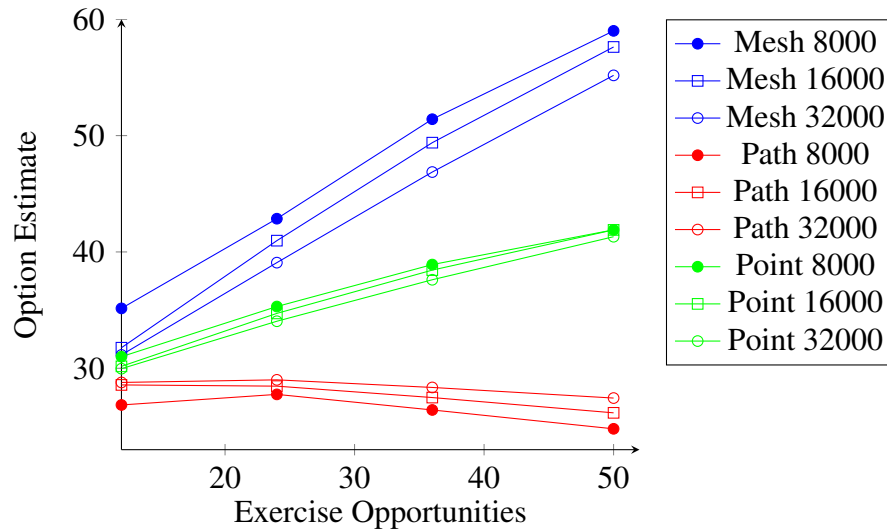
Figure E.1: Mesh and Path Estimators vs Exercise Opportunities: $\mathcal{N}_u = \mathcal{N}_d = 4$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
12	31.765	0.003	28.567	0.003	30.166	1.0
24	40.970	0.006	28.462	0.004	34.716	2.3
36	49.393	0.007	27.482	0.003	38.438	3.3
50	57.620	0.008	26.168	0.003	41.894	4.8

Table E.2: $\mathcal{N}_u = \mathcal{N}_d = 4$, $b = 16000$, $R = 2048$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
12	31.144	0.003	28.781	0.003	29.963	1.0
24	39.080	0.006	29.004	0.004	34.042	2.2
36	46.868	0.007	28.350	0.003	37.609	3.4
50	55.179	0.009	27.438	0.003	41.309	4.7

Table E.3: $\mathcal{N}_u = \mathcal{N}_d = 4$, $b = 32000$, $R = 512$, 0.8 Gb of memory for $m = 50$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	70.520	0.009	53.888	0.007	62.204	0.9
36	84.71	0.01	52.403	0.006	68.557	1.1
50	97.78	0.01	49.908	0.005	73.84	1.7

Table E.4: $\mathcal{N}_u = \mathcal{N}_d = 8$, $b = 8000$, $R = 2048$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	68.584	0.009	54.948	0.007	61.766	1.7
36	82.12	0.01	54.948	0.007	68.534	2.7
50	95.71	0.01	52.197	0.006	73.954	3.5

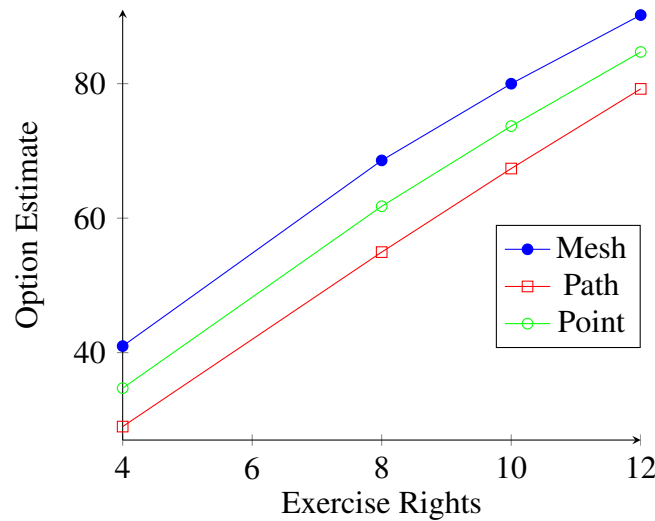Table E.5: $\mathcal{N}_u = \mathcal{N}_d = 4$, $b = 16000$, $R = 1024$, 1.6 Gb of memory for $m = 50$

Figure E.2: Mesh and Path Estimators vs Exercise Rights: $b = 16000$, $m = 24$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	66.629	0.008	55.794	0.006	61.212	2.1

Table E.6: $\mathcal{N}_u = \mathcal{N}_d = 8$, $b = 32000$, $R = 512$, 1.3 Gb of memory

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	81.79	0.01	66.252	0.008	74.021	1.1
36	98.39	0.01	65.128	0.007	81.759	1.3
50	113.94	0.01	62.444	0.006	88.192	1.8

Table E.7: $\mathcal{N}_u = \mathcal{N}_d = 10$, $b = 8000$, $R = 2048$, 1.3 Gb of memory for $m = 50$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	79.99	0.01	67.376	0.009	73.683	2.7
36	95.76	0.01	67.054	0.009	81.407	3.1
50	111.60	0.02	65.053	0.007	88.327	4.3

Table E.8: $\mathcal{N}_u = \mathcal{N}_d = 10$, $b = 16000$, $R = 1024$, 1.7 Gb of memory for $m = 50$

Exercise Opps	Mesh	Error	Path	Error	Point	Time (yrs)
24	90.20	0.01	79.22	0.01	84.71	2.8
36	108.08	0.02	79.73	0.01	93.91	4.4

Table E.9: $\mathcal{N}_u = \mathcal{N}_d = 12$, $b = 16000$, $R = 1024$, 1.8 Gb of memory for $m = 36$

Bibliography

[1] J. C. Cox, S. A. Ross, M. Rubinstein. Option pricing: A simplified approach. Journal of Financial Economics 7, (1979) 229-263.

[2] D. Nelson, K. Ramaswamy. Simple binomial processes as diffusion approximation in financial models. Review of Financial Studies 3, (1990) 393-430.

[3] K. I. Amin, A. Khanna. Convergence of American option values from discrete- to continuous-time financial models. Mathematical Finance 4, (1994) 289-304.

[4] P. P. Boyle. Options: A Monte Carlo approach. Journal of Financial Economics 4, (1977) 323-338.

[5] J. A. Tilley. Valuing American Options in a Path Simulation Model, Transactions of the Society of Actuaries 45, (1993) 83-104.

[6] M. Broadie, P. Glasserman. Pricing American-style securities using simulation. Journal of Economic Dynamics and Control 21, (1997) 1323-1352.

[7] M. Broadie, P. Glasserman. A stochastic mesh method for pricing high-dimensional American options. The Journal of Computational Finance 7(4), (2004) 35-72.

[8] J. Carriere. Valuation of Early-Excercise Price of Options Using Simulations and Non-parametric Regression. Insurance: Mathematics and Economics 19, (1996) 19-30.

[9] F. Longstaff, E. S. Schwartz. Valuing American Options by Simulation: A Simple Least-squares Approach. The Review of Financial Studies 14(1), (2001) 113-147.

[10] E. Clément, D. Lamberton, P. Protter. An analysis of a least-squares regression algorithm for American option pricing. Finance Stochastics 6, (2001) 449-471.

[11] L. Stentoft. Assessing the Least-Squares Monte-Carlo Approach to American Option Valuation. Review of Derivatives Research 7, (2004) 129-168.

[12] L. Andersen, M. Broadie. Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options. Management Science 50(9), (2004) 1222-1234.

[13] A. Kolodko, J. Schoenmakers. Iterative Construction of the Optimal Bermudan Stopping Time. Finance and Stochastics 10, (2004) 27-49.

[14] D. Duffie. Dynamic Asset Pricing Theory, Third Edition, Princeton University Press, Princeton, New Jersey (2001).

[15] M. Matsumoto, T. Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. ACM Transactions on Modeling and Computer Simulation 8, (1998) 3-30.

[16] C. L. Philips, J. A. Anderson, S. C. Glotzer. Pseudo-random number generation for Brownian Dynamics and Dissipative Particle Dynamics simulations on GPU devices. Journal of Computational Physics 230, (2011) 7191-7201.

[17] B. Moro. The full Monte. Risk 8, (1995) 57-58.

[18] G. E. P. Box, M. E. Muller. A note on the generation of random normal variates. Annals of Mathematical Statistics 29, (1958) 610-611.

[19] T. Whitehead, R. M. Reesor, and M. Davison. A bias-reduction technique of monte carlo pricing of early-exercise options. Journal of Computational Finance 15, (2012) 33-69.

[20] M. Broadie, P. Glasserman, and Ha. Z. Pricing American options by simulation using a stochastic mesh with optimized weights. Probabilistic Constrained Optimization: Methodology and Applications. Kluwer Publishers, Norwell, Mass. (2000).

[21] P. Glasserman, B. Yu. Pricing American options by simulation: regression now or regression later? Monte Carlo and Quasi-Monte Carlo Methods. Springer-Verlag, Berlin, 2003.

[22] P. Glasserman. Monte Carlo Methods in Financial Engineering, Springer, New York, 2004.

[23] A. Lari-Lavassani, M. Simchi, A. Ware. A Discrete Valuation of Swing Options. Canadian Applied Mathematics Quarterly 9, (2001) 35-73.

[24] N. Meinshausen, B. M. Hambly. Monte Carlo Methods For the Valuation of Multiple-Exercise Options. Mathematical Finance 14, (2004) 557-583.

[25] C. Bender. Dual pricing of multi-exercise options under volume constraints. Finance and Stochastics 15, (2011) 1-26.

[26] P. Jaillet, E. I. Ronn, S. Tompaidis. Valuation of Commodity-Based Swing Options. Management Science 50, (2004) 909-921.

[27] M. Dahlgren, R. Korn. The Swing Option on the Stock Market. The International Journal of Theoretical and Applied Finance 8(1), (2005) 123-129.

[28] J. Barraquand, D. Martineau. Numerical Valuation of High Dimensional Multivariate American Securities. Journal of Financial and Quantitative Analysis 30, (1995) 383-405.

[29] V. Bally, G. Pagés, J. Printems. A Quantization Tree Method for Pricing and Hedging Multidimensional American Options. Mathematical Finance 15, (2005) 119-168.

[30] M. B. Haugh, L. Kogan. Pricing American Options: A Duality Approach. Operations Research 52(2), (2004) 258-270.

[31] A. Ibáñez. Valuation by Simulation of Contingent Claims with Multiple Early Exercise Opportunities. Mathematical Finance 19, (1996) 19-30.

[32] C. Barrera-Esteve, F. Bergeret, C. Dossal, E. Gobet, A. Meziou, R. Munos, D. Reboul-Salze. Numerical Methods for the Pricing of Swing Options: A Stochastic Control Approach. Methodology and Computing in Applied Probability 8, (2006) 517-540.

[33] K. H. Kan, R. M. Reesor. Bias reduction of least-squares Monte Carlo estimators of American option value. Pre-print (2009).

[34] C. Bender, J. Schoenmakers. An Iterative Algorithm for Multiple Stopping: Convergence and Stability. Adv. in Appl. Probab. 38, (2006), 729-749.

[35] J. Dickson, T. J. Marshall, R. M. Reesor. A Bias-Reduction Technique for Monte Carlo Pricing of Multiple-Exercise Options. Working Paper, Western University of Canada (2012).

[36] A. Gut. Stopped Random Walks. Springer-Verlag, New York (1988).

[37] T. J. Marshall, R. M. Reesor. Forest of Stochastic Meshes: A Method for Valuing High Dimensional Swing Options. Operations Research Letters 39, (2011) 17-21.

[38] T. J. Marshall, R. M. Reesor, M. Cox. Simulation Valuation of Multiple Exercise Options. Proceedings of the 2011 Winter Simulation Conference.

Curriculum Vitae

Name:	Thomas James Marshall

Post-Secondary Education and Degrees:

Western University of Canada
London, ON, Canada
2009 - 2012 Ph.D.

Western University of Canada,
London, ON, Canada
2005 - 2007 M.Sc.

University of New Brunswick,
Fredericton, NB, Canada
2001 - 2005 B.Sc.

Honours and Awards:

2011 BMO Capital Markets Advanced Research Scholarship (Co-Winner)
2009–2011 NSERC PGS-D
2005–2007 NSERC PGS-M

Related Work Experience:

Associate, Financial Products Desk
BMO Capital Markets
May 2011 – August 2011

Instructor
Department of Physics, University of New Brunswick
2008

Teaching Assistant
The University of Western Ontario
2005 - 2007 & 2009 - 2011

Publications:

(i) T. J. Marshall, R. Mark Reesor, M. Cox, "Simulation Valuation of Multiple Exercise Options", Proceedings of the 2011 Winter Simulation Conference.

(ii) T. J. Marshall, R. Mark Reesor, "Forest of Stochastic Meshes: A Method for Valuing High Dimensional Swing Options", Operations Research Letters **39**, 17-21 (2011).

(iii) T. J. Marshall, D. G. C. McKeon, "Radiative Properties of the Stueckelberg Mechanism", International Journal of Modern Physics A **23**, 741 (2008)