Monte Carlo Pricing of Derivative Securities and Uncertainty in Volatility Estimation

Monte Carlo Pricing of Derivative Securities and Uncertainty in Volatility Estimation

George A. Joplin, V

Rhodes College
Memphis, Tennessee

Submitted in partial fulfillment
of the requirements for the Bachelor of Arts
degree with Honors in Mathematics and Economics

May 20, 2011

This Honors paper by George A. Joplin, V has been read and approved for Honors in Mathematics and Economics.

Dr. Jeff Hamrick
Project Advisor

____

Dr. Marshall Gramm
Chair, Economics and Business Department

____

Dr. Michael Sheard
Chair, Mathematics and Computer Science Department

____

Contents

1 Introduction **1**

2 Review of Relevant Literature **7**
2.1 Pricing Options with Monte Carlo Simulation 7
2.2 Nonparametric Estimation of the Diffusion Function via the Maximum Penalized Quasi-Likelihood Method 10
2.3 Nonparametric Estimation of the Diffusion Function via Kernel-Based Methods 11

3 Methodology **15**
3.1 Two Methods of Monte Carlo Simulation 15
3.2 Types of Options Studied in the Research 16
3.3 Local Mean Estimation of the Diffusion Function Using Kernel-Based Techniques 18
3.4 Bandwidth Selection in the Local Mean Estimation 19
3.5 Tests on Variance: Ratio of Variances 19

4 Results and Implications **21**
4.1 Results - IBM 22
4.2 Results - Exxon Mobil (XOM) 24
4.3 Results - Microsoft (MSFT) 27
4.4 Results - GE 29
4.5 Implications 33

A Appendix **35**

List of Figures

1 IBM European . . . 22
2 IBM Asian . . . 23
3 IBM Barrier . . . 23
4 IBM Binary . . . 24
5 XOM European . . . 24
6 XOM Asian . . . 25
7 XOM Barrier . . . 25
8 XOM Binary . . . 26
9 MSFT European . . . 27
10 MSFT Asian . . . 28
11 MSFT Barrier . . . 28
12 MSFT Binary . . . 29
13 GE European . . . 29
14 GE Asian . . . 30
15 GE Barrier . . . 30
16 GE Binary . . . 31
17 IBM price time series . . . 35
18 Exxon-Mobil price time series . . . 35
19 Microsoft price time series . . . 35
20 General Electric price time series . . . 35

List of Tables

1 List of All Results 32

Abstract: In the standard time-inhomogeneous diffusion model, estimation of the volatility function is far more important for Monte Carlo pricing than estimation of the drift function (due to a standard application of Girsanov's Theorem). As such, we study the distribution of option prices under the uncertainty of volatility function estimation. First, we run Monte Carlo simulations to price a variety of options using a fixed estimate of the volatility function. Then, we run Monte Carlo simulations to price a variety of options using a bootstrapped re-estimation of volatility function in each Monte Carlo trial. The differences in the resulting distributions of option prices may have implications for thinking about the bid-ask spread on an option price, and can be compared to historical data to gain a more complete perspective on the acceptability of various American-style option prices.

1 Introduction

One of the most central and widely-acknowledged assumptions of option pricing theory is that it is impossible to predict the values of tomorrow's asset prices with absolute certainty. That is, for an asset price at time t denoted by X_t, it is impossible to predict the price at X_{t+1}. Through statistical analysis and under the assumption that the stochastic process is time-inhomogeneous however, the past history of any asset can be studied to determine mean, variance, skewness, and distribution (among other measures) and *estimate* the fair price of a given asset. It is also assumed that asset prices move randomly due to the *efficient market hypothesis*, which states that (1) the past history of an asset is fully reflected in the present price, but does not contain any further information, and (2) markets respond immediately to new information about an asset [25]. Each relative change in asset price is measured by a change denoted by dX_t. The most common model of these changes has two parts: a deterministic, anticipated contribution $\mu(X_t)dt$, where μ is a measure of the average rate of growth, or drift, of an asset price when the asset is at the level X_t, and a stochastic contribution $\sigma(X_t)dW_t$ that models random change, where σ measures the average rate of growth, or volatility at the level X_t, of an asset price, and dW_t is approximately a normally distributed random variable. Thus, the return on an asset price is given by the time-inhomogeneous stochastic differential equation

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t. \tag{1}$$

Let Equation (1), also known as the *random walk* of an asset price, have a probability transition density $p(X, t, X', t')$. Then it satisfies the partial differential equation

$$\frac{\partial p}{\partial t} + \frac{1}{2}\sigma(X_t)^2\frac{\partial^2 p}{\partial X^2} + \mu(X_t)\frac{\partial p}{\partial X} = 0 \tag{2}$$

which is called the backward Kolmogorov equation. Suppose Equation (1) represents the value of an asset at expiry time T. It is natural, then, to find the present value of the expected amount received. Multiply $p(X,t)$ by the discount factor to give

$$V(x,t) := e^{-r(T-t)}p(X,t) \tag{3}$$

where r is a constant deterministic interest rate. Then Equation (3) satisfies

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma(X_t)^2\frac{\partial^2 V}{\partial x^2} + \mu(X_t)\frac{\partial V}{\partial x} - rV = 0 \tag{4}$$

which is almost exactly like the Black-Scholes partial differential equation for pricing an option. The only difference is that there is no μ in the Black-Scholes equation, but instead the Black-Scholes equation has an r in place of μ. Replace μ with r and consider a new random walk given by

$$dX_t = r(X_t)dt + \sigma(X_t)dW_t \tag{5}$$

Equation (5) gives a *risk-neutral random walk* for X [24]. Unlike many differential equations, this stochastic differential equation cannot be solved to generate a deterministic path for asset price X, but it can give interesting information about the behavior of X in a probabilistic sense [25].

One of the most useful applications of Equation (1) is its use in estimating the fair value of a call option. Strictly defined, a call option is a derivative that gives the holder the *right*, but not the obligation, to buy or sell an asset at a specified price until a specified expiration date [5]. Its payoff function is therefore nonlinear. Thus, the buyer of a single call option with a strike price of \$50 and an expiration date of May 1 has the right to purchase 100 shares of a given stock for \$50 on or before May 1. This right will be exercised only if it is profitable to do so. The fair value of an option in the Black-Scholes model is the present value of the expected payoff at the expiration date under the risk-neutral random walk for the underlying [24]. The fair value of the option at time t is given by

$$V(t) = e^{-r(T-t)} E_*[f(S_T)] \tag{6}$$

where r is the interest rate, T is the expiration date, and $E_*[f(S_T)]$ is the expectation of the risk-neutral random walk with respect to the derivative's payoff function. As previously discussed, the expectation of the risk-neutral random walk is subject to a degree of uncertainty, namely the component dW_t in its stochastic term. Pricing the call option, then, requires many simulations of Equation (6) to account for the randomness of the term. This is done effectively using Monte Carlo simulation, which involves repeated random sampling to estimate a mean—in our case, the average discounted risk-neutral payoff of some option. Financial researcher Paul Wilmott gives an algorithm in [24] to execute Monte Carlo simulation to price an option, which is as follows:

1. Simulate the risk-neutral random walk from Equation (5), starting at today's value of the asset X_t over the required time horizon (until the option expires). This is one realization of the underlying price path.

2. Calculate the option payoff using Equation (6) for this realization.

3. Repeat steps 1 and 2 n times, where n is an arbitrarily large integer.

4. Find the average payoff over all realizations.

5. Take the present value of this average. This is the option value.

Monte Carlo simulation is a viable technique for estimating option prices because correlations between assets can be easily modeled, accuracy is directly related to the number of simulations run, and it is relatively easy to code and develop. In Monte

Carlo, the accuracy of the result is directly related to the number of simulations run. By the Law of Large Numbers, it must be true that the average of results obtained from a number of trials should converge to the expected value as the number of trials tends to infinity.

Of crucial importance in approximating the price of the option is accurately estimating the parameters of interest. In this case, the principal parameter of interest is the volatility function σ. Methods by which σ is estimated are varied across the financial literature. Volatility can be estimated nonparametrically using kernel-smoothing estimation of σ by employing kernel estimation of a local mean function. The methodology for kernel smoothing is the following.

Recall the continuous-time representation of the stochastic differential equation of return is given by Equation (1). For purpose of practicality, we will represent the equation in discrete time by

$$\Delta X_t = \mu(X_t)\Delta t + \sigma(X_t)\epsilon_{\Delta t} \tag{7}$$

where $\epsilon_{\Delta t}$ are independent normally distributed random variables with mean 0 and variance Δt. Then, define Y_t as follows:

$$Y_t := \frac{\Delta X_t}{\sqrt{\Delta t}} = \mu(X_t)\sqrt{\Delta t} + \sigma(X_t)\epsilon_t \tag{8}$$

where ϵ_t are standard normal random variables. For practical purposes, consider $\Delta t \ll 1$. This is reasonable to assume since, for daily returns, $\Delta t \approx \frac{1}{251}$. Since Δt is small, $\sqrt{\Delta t}$ is small, and we can reasonably ignore the first term of Equation (8) for practical purposes. Thus,

$$Y_t \approx \sigma(X_t)\epsilon_t \tag{9}$$

Equation (9) allows for heteroscedasticity in the model. Thus, one can use techniques similar to the theory of generalized least squares regression to reasonably estimate σ from Equation (9). This method is discussed and developed in [18].

No matter the method by which σ is estimated, the estimate $\widehat{\sigma}$ is subject to a degree of estimation error. Thus, the estimate has a confidence "tube" with which it is associated, where the true value of the the volatility lies with high probability. As such, a proper and accurate estimation of $\widehat{\sigma}$ is crucial for correctly pricing the option via Monte Carlo simulation. Typically, one estimates $\widehat{\sigma}$ via a method above, assigns the value $\widehat{\sigma}$ to the approximation, and completes the algorithm above. Thus, the steps for pricing the option are as follows:

1. Estimate $\widehat{\sigma}$ via a method of one's choosing.
2. Set $\widehat{\sigma} = \sigma$.
3. Complete Steps 1-5 of Monte Carlo Algorithm as stated in [24]. This is the option value.

However, one can undertake a different method for pricing the option via Monte Carlo simulation. Since $\widehat{\sigma}$ is subject to estimation error, it can vary from estimation to estimation. Thus, one can bootstrap the estimation of σ and run Monte Carlo simulation based on a re-sampled value of $\widehat{\sigma}$ each time. The steps for pricing the options via this method are as follows:

1. Estimate $\widehat{\sigma}$ via a method of one's choosing.
2. Set $\widehat{\sigma} = \sigma_1$. Complete Steps 1-5 of Monte Carlo Algorithm.
3. Repeat step 1 and set $\widehat{\sigma} = \sigma_2$. Complete Steps 1-5 of Monte Carlo Algorithm.
4. Repeat step 3 n times and average all results. This is the option value.

Theoretically, simulation with a fixed σ playing the role of σ could differ drastically from a simulation with bootstrapped value of σ. The fundamental questions at hand are as follows: What is the distribution of option prices under the uncertainty of volatility estimation? What is the shape of the distribution of option prices with a fixed estimate of σ as opposed to many bootstrapped re-estimates of σ?

2 Review of Relevant Literature

2.1 Pricing Options with Monte Carlo Simulation

Monte Carlo methods are based upon the relationship between probability and volume [14]. The notion of probability associates an event with a set of outcomes and defines the probability of that event occuring to be its volume relative to the universe of possible outcomes. Monte Carlo methods use this relationship in reverse by calculating the volume of a set and interpreting that volume as a probability. The simplest form of Monte Carlo simulation is implemented by taking a large number of random samples within a universe of possible values and finding the number of samples that fall within a subset of interest. The law of large numbers guarantees that the estimate approaches its true value as the number of random samples approaches infinity [14].

The seminal work on Monte Carlo simulation's applications to pricing derivative securities is Phelim C. Boyle's 1976 article "Options: A Monte Carlo Approach." Findings from Cox and Ross [8] show that if the distribution of the final stock value of an asset can be found, then the value of the option can be determined through integration. The integrals will not necessarily have analytic solutions, but can nonetheless be computed through numerical methods. Chen [6] and Parkinson [20] used two different numerical integration schemes to compute these integrals. Boyle's paper [2] shows that Monte Carlo simulation is a third method for obtaining numerical solutions to option valuation problems. The Monte Carlo numerical scheme uses the fact that the distribution of stock prices at the time of expiration is determined by the process generating future stock movements. This process of producing asset "trajectories" can be facilitated using computer simulation, which determines potential terminal stock values that can be used to obtain an estimate of the option value through the payoff functions.

According to Boyle, it is best to discuss Monte Carlo simulation in terms of evaluating a definite integral. Consider

$$\int_A g(y)f(y)dy = \bar{g} \tag{10}$$

where $g(y)$ is an arbitrary function, $f(y)$ is a probability density function, and A is some range of integration. To obtain an estimate for $\bar{g}$, a number n of samples y_i is selected from $f(y)$. Then the estimate of $\bar{g}$ is given by

$$\widehat{g} = \frac{1}{n}\sum_{i=1}^{n} g(y_i) \tag{11}$$

with sample variance

$$\widehat{s}^2 = \frac{1}{n-1}\sum_{i=1}^{n}(g(y_i) - \widehat{g})^2. \tag{12}$$

The distribution

$$\frac{\widehat{g} - \bar{g}}{\sqrt{\widehat{s}^2/n}} \tag{13}$$

tends to a standardized normal distribution as n approaches infinity. Boyle includes two variance reduction techniques in order to improve the accuracy of Monte Carlo Simulation: the control variate method and the method of antithetic variates.

In [8], Cox and Ross have shown that the assumptions of risk-neutrality can be used to give expected rates of return on stocks, given by

$$\mathbb{E}_*[S_T/S_t] = e^{r(T-t)} \tag{14}$$

where S_T is the stock price at expiration, S_t is the stock price at time t, and r is the risk-free rate. Furthermore, the ratio of successive stock values (S_t to S_{t+1}) follow a lognormal distribution, where

$$S_{t+1} = S_t e^{r - \frac{\sigma^2}{2} + \sigma x} \tag{15}$$

where σ is the standard deviation (per quarter) of the underlying asset, and x is a normally distributed random variable with zero mean and unit variance. With this notation in mind, Boyle describes the process of Monte Carlo simulation using a dividend-paying stock, with dividends paid at time t denoted by D_t. A value of S_{t+1} is generated. If this value is greater than D_{t+1}, then $S_{t+1} - D_{t+1}$ is used as the initial value at the start of the second period and the procedure continues until S_T is obtained. If at some point $S_{t+m}, m = 1, 2, ..., T - t - 1$ is less than or equal to D_{t+m}, the process stops and another simulation trial starts from time t. The simulation is carried out n times, and the expected value of $\max[S_T - K, 0]$, where K is the strike price, is found. This is then discounted at the risk-free rate and used as an estimate of the option value $\widehat{v}$. Ninety-five percent confidence intervals are given by $\widehat{v} \pm 2\widehat{s}/\sqrt{n}$. Boyle uses the simulation process on three assets with ten different periods to maturity, using 5000 trials for each simulation. All thirty Monte Carlo estimates were within two standard deviations of accurate option values founds using numerical integration.

In recent years, Monte Carlo simulation has been used by both financial practitioners and academic financial mathematicians to accurately price financial derivatives such as call options. In [21], Monte Carlo simulation is used to confirm that option prices are more informative about estimation parameters than asset prices in continuous-time stochastic volatility models. In [9], Duffie, Darrell and Glynn discuss an efficient Monte Carlo algorithm for estimating continuous-time security prices, focusing particularly on the tradeoff between number of time increments and number of simulations. In [4], two direct Monte Carlo methods, a pathwise method and a likelihood ratio method method, are uesd for estimating derivatives of security prices using simulation. Furthermore, [7] discusses the ways in which Monte Carlo methods have become of renewed interest in recent years because of their flexibility and computational speed.

2.2 Nonparametric Estimation of the Diffusion Function via the Maximum Penalized Quasi-Likelihood Method

One possibility for nonparametric estimation of σ in the diffusion function is accomplished via the maximum penalized quasi-likelihood method developed in [16]. Consider the discrete anlalog of the stochastic differential equation in Equation (1)

$$\Delta X_t = X_{t+\Delta t} = \mu(X_t)dt + \sigma(X_t)\epsilon_{\Delta t}. \tag{16}$$

where $\epsilon_{\Delta t} \sim$ N(0,1). Assume that $\Delta t \ll 1$. This gives

$$\Delta X_{t+\Delta t} - X_t = \sigma(X_t)\sqrt{\Delta t}\epsilon. \tag{17}$$

Then define the quasi-likelihood function of X_t by

$$\mathrm{ql}(\sigma | X_{t_0}, ..., X_{t_n}) = \prod_{i=1}^{n} \frac{1}{\sigma(X_{t_i})\sqrt{2\pi}} e^{\frac{(X_{t_{i+1}} - X_{t_i})^2}{2\sigma(X_t)^2}} \tag{18}$$

and the subsequent quasi-log-likelihood function is given by

$$\begin{aligned} \mathrm{qll}(\sigma | X_{t_0}, ..., X_{t_n}) &= \log(\prod_{i=1}^{n} \frac{1}{\sigma X_{t_i}\sqrt{2\pi}} e^{\frac{(X_{t_{i+1}} - X_{t_i})^2}{2\sigma(X_t)^2}}) \\ &= \sum_{i=1}^{n} \left\{ \log\sigma(X_{t_i}) - \log\sqrt{2\pi} - \frac{(X_{t_{i+1}} - X_{t_i})^2}{2\sigma(X_t)^2} \right\} \end{aligned} \tag{19}$$

which is more mathematically tractable and maintains all relevant information from the likelihood function since the log function is monotonic and therefore one-to-one [17]. Then, define the weighted quasi log-likelihood function by

$$\text{wqll}(\sigma|X_{t_0},...,X_{t_n}) = \frac{1}{n}\sum_{i=1}^{n}\Big\{\log\sigma(X_{t_i}) - \log\sqrt{2\pi} - \frac{(X_{t_{i+1}} - X_{t_i})^2}{2\sigma(X_t)^2}\Big\} \tag{20}$$

Equation (20) is maximized by any points that pass through the collection of points $\{(Y_{t_i}, |R_{t_i}|)\}_{i=1}^{n}$. However, this approach is mathematically untractable since the estimators obtained often oscillate wildly. Therefore, we define a function that penalizes estimators that have infinite or nearly-infinite derivatives at various points. Define a function the *penalized weighted quasi log-likelihood function* by

$$\text{pwqll}(\sigma|X_{t_0},...,X_{t_n}) = \frac{1}{n}\sum_{i=1}^{n}[\log\sigma(X_{t_i}) - \log\sqrt{2\pi} - \frac{(X_{t_{i+1}} - X_{t_i})^2}{2\sigma(X_t)^2}] - \frac{\lambda}{2}\int_{m_T}^{M_T}|\sigma^{(m)}(\upsilon)|d\upsilon \tag{21}$$

where m_T and M_T are the lowest and highest prices of the path of the stock price over $[0,T]$. Then the following must be true:

1. There exists a $\sigma^* = \text{argmax}_{\sigma\in W^{(2,m)}}\, \text{pwqll}(\sigma|X_{t_0},...,X_{t_n})$
2. σ^* is unique.
3. σ^* is a natural spline of order $2m-1$.
4. There exists a numerical scheme that finds σ^*.

The theory and methodology for this estimation scheme is developed in [16].

2.3 Nonparametric Estimation of the Diffusion Function via Kernel-Based Methods

In recent years, interest in nonlinear regression, and particular interest in estimation of the diffusion function, has provided an extensive amount of rich and relevant literature on the topic. One significant paper on the subject was Jianqing Fan and Qiwei Yao's "Efficient Estimation of Conditional Variance Functions in Stochastic Regression" [13]. In the paper, the authors propose a residual-based method for estimating the conditional variance in a set of statistical data. Let $\{(T_i, X_i)\}$ be a two-dimensional strictly stationary process. Define

$$m(x) = E(Y|X = x) \tag{22}$$

and

$$\sigma^2(x) = \text{var}(Y|X = x) > 0 \tag{23}$$

Then we can define a regression model of Y on X by

$$Y_i = m(X_i) + \sigma(X_i)\epsilon_i \tag{24}$$

In order to effectively create a model, the authors apply a local linear regression to the squared residuals. Previous models suffered from issues related to high bias of the estimator. The idea of using a residual-based estimator of conditional variance had been used in previous research, but the particular estimator described in Fan and Yao's research is regression-adaptive, which means that with or without knowing m, it is possible to conditionally estimate σ^2. Furthermore, in using this approach, it is not necessary to undersmooth the regression function m in order to obtain an estimator for σ^2. Instead, allows the user to use a data-driven bandwidth selector in estimating m, and use the same selector for the squared residuals to estimate σ^2.

Since $m(.)$ is typically unknown in practice, it is naturally replaced by a nonparametric regression estimator [13]. Then $\widehat{m}(x) = \widehat{a}$ is the local linear estimator that solves the weighted least-squares problem:

$$(\widehat{a}, \widehat{b}) = \underset{a,b}{\operatorname{argmin}} \sum_{i=1}^{n} \{Y_i - a - b(X_i - x)\}^2 K\Big(\frac{X_i - x}{h_2}\Big) \tag{25}$$

where K is a density function and $h_2 > 0$ is a bandwidth. Define the squared residuals $\widehat{r}_i = \{Y_i - \widehat{m}(X_i)\}^2$. Then $\widehat{\sigma}^2(x) = \widehat{\alpha}$ with kernel W and bandwidth h_1 is the locally linear residual-based estimator which solves

$$(\widehat{\alpha}, \widehat{\beta}) = \underset{\alpha,\beta}{\operatorname{argmin}} \sum_{i=1}^{n} \{\widehat{r}_i - \alpha - \beta(X_i - x)\}^2 W\Big(\frac{X_i - x}{h_1}\Big). \tag{26}$$

Fan and Yao find that the variance estimator $\widehat{\sigma_2}$ is asymptotically normal and is an efficient estimator. Furthermore, the authors use a bandwidth selection rule for a local linear regression subject to data of the form $(X_1, Y_1), \ldots, (X_n, Y_n)$. Much research has been done on proper bandwidth selection for estimation of the diffusion function, among other kernel-based estimation methods. As is known in statistics and computing, the bandwidth is a smoothing parameter that is often calculated from the relevant data. Its proper selection is important for a proper estimate of, for example, the volatility function.

In [12], Jianqing Fan and Irene Gijbels discuss the methodology of properly selecting the bandwidth parameter. The performance of local polynomial regression estimators are dependent primarily on two choices: the order of the local polynomial fit and the bandwidth. The bandwidth selection can be constant or variable; constant bandwidth is typically sufficient when the unknown curve is smooth, while variable bandwidth is preferred when the curve is not smooth. Local polynomial fitting captures spatially inhomogeneous curves by using the order of the fit (for fixed bandwidth) or choosing an appropriate bandwidth at each point when using a fixed order of fit [12]. Choosing these two parameters is often difficult to do, and Fan and Gijbels propose a method to choose bandwidth that enhances spatial adaptation of local polynomial fitting. The variable bandwidth selection, which is most often used for diffusion function estimation, involves initially partitioning the interval of estimation $[a, b]$ into subintervals I_k and using a bandwidth selector in each interval. Let $\text{RSC}(y; h) = \widehat{\sigma}^2(y)\{1+(p+1)V\}$ be the residual squares criterion which estimates the local mean-squared errors, and let $\widehat{\beta}_p$ be estimated regression coefficients from fitting a p-order polynomial selected by RSC. Then, for each interval I_k, fit a $(p+2)$-order polynomial to select the proper bandwidth for estimating β_{p+1} by minimizing the integrated residual squares criterion, given by

$$\text{IRSC}(h) = \int_{I_k} RSC(y; h)dy. \tag{27}$$

Then, smooth the resulting bandwidth step function by averaging locally on each subinterval. Use this bandwidth function to fit a $(p+2)$-order polynomial and obtain $\widehat{\beta}_{p+1}(x_0), \widehat{\beta}_{p+2}(x_0)$, and $\widehat{\sigma}^2(x_0)$. Do this for each subinterval I_k, and fit a p-ordered polynomial. This approach gives the final estimated curve.

In [13], Fan and Yao use a bandwidth selection scheme similar to that presented above in their local linear fit. Their method is as follows: Let $\widehat{h}(X_1, \ldots, X_n; Y_1, \ldots, Y_n)$ be a data-driven bandwidth selection based on data $(X_1, Y_1), \ldots, (X_n, Y_n)$. Use bandwidth $h_2 = \widehat{h}(X_1, \ldots, X_n; Y_1, \ldots, Y_n)$ to estimate $\widehat{m}(X_i)$. Compute the squred residuals, as given above. Then apply the same bandwidth $h_1 = \widehat{h}(X_1, \ldots, X_n; Y_1, \ldots, Y_n)$ to obtain $\widehat{\sigma}^2$.

3 Methodology

3.1 Two Methods of Monte Carlo Simulation

No matter the method by which σ is estimated, the estimate $\widehat{\sigma}$ is subject to a degree of estimation error. Thus, the estimate has a confidence "tube" with which it is associated, where the true value of the the volatility lies with high probability. As such, a proper and accurate estimation of $\widehat{\sigma}$ is crucial for correctly pricing the option via Monte Carlo simulation. Typically, one estimates $\widehat{\sigma}$ via a method above, assigns the value $\widehat{\sigma}$ to the approximation, and completes the algorithm above. As previously mentioned, the steps for pricing the option are as follows:

1. Estimate $\widehat{\sigma}$ via a method of one's choosing.
2. Set $\widehat{\sigma} = \sigma$.
3. Complete Steps 1-5 of Monte Carlo Algorithm as stated in [24]. This is the option value.

Since $\widehat{\sigma}$ is subject to estimation error, it can vary from estimation to estimation. Thus, one can bootstrap the estimation of σ and run Monte Carlo simulation based on a re-sampled value of $\widehat{\sigma}$ each time. As previously mentioned, the steps for pricing the options via this method are as follows:

1. Estimate $\widehat{\sigma}$ via a method of one's choosing.
2. Set $\widehat{\sigma} = \sigma_1$. Complete Steps 1-5 of Monte Carlo Algorithm.
3. Repeat step 1 and set $\widehat{\sigma} = \sigma_2$. Complete Steps 1-5 of Monte Carlo Algorithm.
4. Repeat step 3 n times and average all results. This is the option value.

Theoretically, simulation with a fixed σ playing the role of σ could differ drastically from a simulation with bootstrapped value of σ. To reiterate, the fundamental questions at hand are as follows: What is the distribution of option prices under the uncertainty of volatility estimation? What is the shape of the distribution of option prices with a fixed estimate of σ as opposed to many bootstrapped re-estimates of σ?

3.2 Types of Options Studied in the Research

As previously defined, a call option is a derivative with a nonlinear payoff function that gives the holder the *right*, but not the obligation, to buy or sell an asset at a specified price until a specified expiration date [5]. Thus, the buyer of a single call option with a strike price of $50 and an expiration date of May 1 has the right to purchase 100 shares of a given stock for $50 on or before May 1. This right will be exercised only if it returns a profit. The fair value of an option in the Black-Scholes Model is the present value of the expected payoff at the expiration date under the risk-neutral random walk for the underlying [24]. The value of the option at time t is given by

$$V(t) = e^{-r(T-t)} E_*[f(S_T)] \tag{28}$$

where r is the interest rate, T is the expiration date, and $E_*[f(S_T)]$ is the expectation of the risk-neutral random walk with respect to the derivative's payoff function.

There are hundreds of different options in existence, each with different characteristics that make them unique. Some options are *path independent*—thus their value depends only on the final value of the underlying on the date of expiration. Others are *path dependent*, and their value depends on the entire path that the security has taken until the date of expiration. In the research, we consider four different types of options: European call, Asian call, barrier call, and binary call.

European call options are often considered the most "basic" type of option. European options are path independent, and thus depend only on the final value of the security at expiry. A European call option is a contract in which the holder has the right to buy a particular asset for a particular price, the strike price K, at some prescribed time in the future, T [25]. If the asset price is below K at time T, then the holder will not exercise the call option because he or she can buy the asset on the open market at lower cost. If the asset price is above K at time T, the holder will purchase the asset at price K. Thus the payoff of a European call option is given by $\max\{S_T - K, 0\}$.

Asian call options give the holder the right to buy a security for a prescribed strike price, or the average price of the security during an agreed upon period. For this reason, an Asian option is path dependent, because its value depends on the path that the underlying takes until the expiration date. The price of the Asian option is much less than other vanilla options because the payoff depends on the average of the path through its life [24]. The payoff of the Asian option is given by $\max\{A-K, 0\}$, where K is the strike price and A is the average value of the security during the agreed upon period. The type of averaging done is important when considering the Asian option. In this project, we studied the geometric Asian call option, whose payoff is the exponential of the sums of all the logarithms of prices, equally weighted, and divided by the number of prices used. Furthermore, the geometric average was calculated using discrete sampling. Only closing prices from trading days were considered, not prices of the security throughout the day, which could affect the payoff.

Barrier call options are another type of path dependent option whose payoff structure relies on the path the option takes during the agreed upon period. Certain properties of the contract come into play when the security price reaches a certain level. In this project, I consider a discretely measured up-and-out barrier call option. The up-and-out call option has the payoff of $max\{S - K, 0\}$, where K is the strike price, unless the security's closing price exceeds some agreed upon price S_u during the agreed upon period, in which case the contract is terminated. If the security price never exceeds S_u, then the payoff is the difference $S_T - K$. Barrier call options are popular for traders with very particular views about the market, particularly traders who believe that movement of the security will be limited prior to expiration time T and who do not want to pay for the upside potential of a European call option.

Binary call options are path independent, but have a payoff at expiration that is discontinuous in the underlying security price [24]. A binary call option pays a particular agreed upon price—denote by J, if the security price is greater than the strike price K at expiration time T. If the security price is less than the strike price K, it pays \$0. As opposed to a European call, whose payoff grows linearly beyond the strike price, a binary call can pay J or \$0. A binary call is useful to hold when you believe that the security price will be greater than K, but not much greater—otherwise you would choose to hold a European call.

3.3 Local Mean Estimation of the Diffusion Function Using Kernel-Based Techniques

In our research, we estimate the diffusion function nonparametrically using a kernel-weighted least squares approach. First, combine a proper set of closing price data $\{Y_t\}_{t=1}^n$ with a set of time data $\{X_t\}_{t=0}^n$ in n steps to form a set of n tuples $\{X_t, Y_t\}$. We then make the transformation $\{Y_t, \frac{Y_{t+1}-Y_t}{\sqrt{X_{t+1}-X_t}}\}$. Then, at each point Y_t, we solve a kernel-weighted least-squares problem to fit a polynomial of order p at each point. The data are weighted by the Epanechnikov kernel given by $K_h(u) = \frac{1}{h}\frac{3}{4}(1-u^2)^+$, where $u = \frac{x-x_0}{h}$, x_0 is a design point in the data, and h is the kernel bandwidth, the smoothing parameter for the estimate. In [22], it has been shown that the Epanechnikov kernel is the only admissible kernel in nonparametric kernel density estimation.

By creating a least-squares estimate at each point in a mesh of points, it is possible obtain an estimate of the volatility function and its first and second derivatives at each point. [15] This gives us information to fit a spline through the entire mesh of points. We then apply an interpolating function to the spline to give a computationally tractable volatility function that fits the data.

3.4 Bandwidth Selection in the Local Mean Estimation

As mentioned in the previous section, the selection of the bandwidth parameter h is very important in estimating the volatility function. At each design point x_0 in the data, h controls the number of and extent to which nearby points affect the estimate of the function σ at x_0. [15] Lower values of h will allow too much influence of other data points, thus making the estimate overly oscillatory while an overestimate of h might oversmooth the data. In [3], an algorithm for selecting an appropriate data-driven bandwidth level is given and was modified and used in our project.

3.5 Tests on Variance: Ratio of Variances

In the project, we are interested in the distribution of option prices under the uncertainty of volatility estimation. We use real-world stock data to generate Monte Carlo-simulated option prices under a "no bootstrap" and "with bootstrap" condition. As result, we have several thousand estimated option prices for different types of stocks. We are interested in whether there is a statistically significant difference in the variance of prices, holding all other parameters equal. Thus, we perform statistical tests on variance on the data. Given a sample size of n_1 and n_2 from the two sets of data, we are interested in testing the null hypothesis $\sigma_1^2 = \sigma_2^2$ against the alternative hypothesis $\sigma_1^2 \neq \sigma_2^2$. The appropriate critical region for this situation is given by

$$\frac{s_1^2}{s_2^2} \geq f_{\alpha/2,n_1-1,n_2-1} \text{ if } s_1^2 \geq s_2^2 \tag{29}$$

and

$$\frac{s_2^2}{s_1^2} \geq f_{\alpha/2,n_2-1,n_1-1} \text{ if } s_1^2 < s_2^2 \tag{30}$$

where s_1^2 and s_2^2 are sample variances and α is the probability of rejecting the null hypothesis when it is actually true [19].

4 Results and Implications

In our research, we ran Monte Carlo simulations under the "no bootstrap" and "with bootstrap" condition for real-world stock data. We used stock data from IBM, Exxon Mobil, Microsoft, and General Electric stock starting from January 1, 2000 until the present. For each sample path, we started at the closing price for the stock closest to January 1, 2000, with a strike price \$10 out. Furthermore, each Monte Carlo simulation was run with 1000 trials, so $n_1 = n_2 = 1000$ in each case. Each statistical test on variance was performed with $\alpha = 0.01$. Barrier options were modeled as "up-and-out" call options with \$1 given to the holder if the underlying's closing price did not break the barrier. Binary options were modeled as giving \$1 to the holder if the underlying's closing price was not above the barrier at the end of the option's life. For each of the following cases, we give s_n^2, the sample variance under the no bootstrap condition, and s_w^2, the sample variance under the with bootstrap condition.

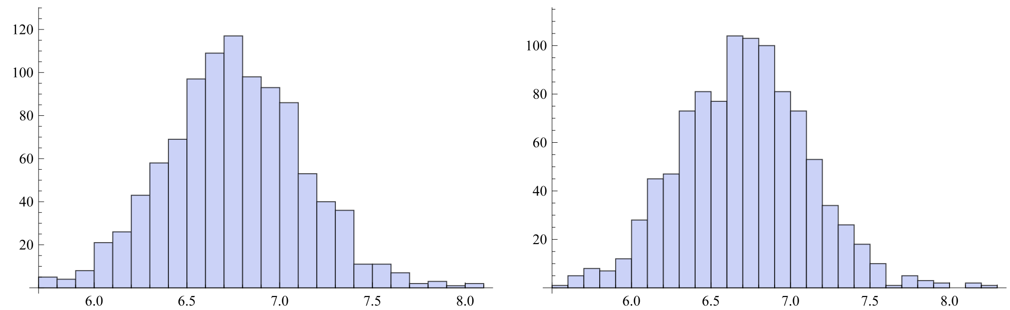

Figure 1: IBM European
Monte Carlo simulation for a European option with no bootstrapping (left) and with bootstrapping (right) for an IBM stock with starting price \$100 and strike price \$110.

4.1 Results - IBM

European: In the no bootstrapping case, $s_n^2 = 0.135$. In the with bootstrapping case, $s_w^2 = 0.163$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 1.203 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Asian: In the no bootstrapping case, $s_n^2 = 2.321 \times 10^{-5}$. In the with bootstrapping case, $s_w^2 = 2.305 \times 10^{-5}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 1.58 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Barrier: In the no bootstrapping case, $s_n^2 = 1.22 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 2.44 \times 10^{-4}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 2.003 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

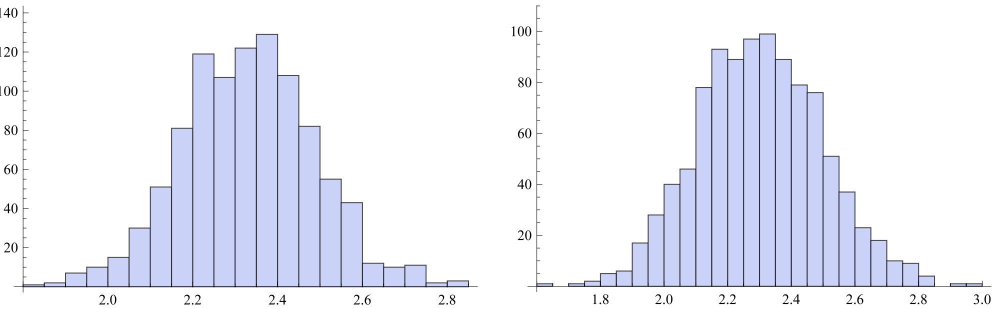

Figure 2: IBM Asian
Monte Carlo simulation for a Asian option with no bootstrapping (left) and with bootstrapping (right) for an IBM stock with with starting price \$100 and strike price \$110.

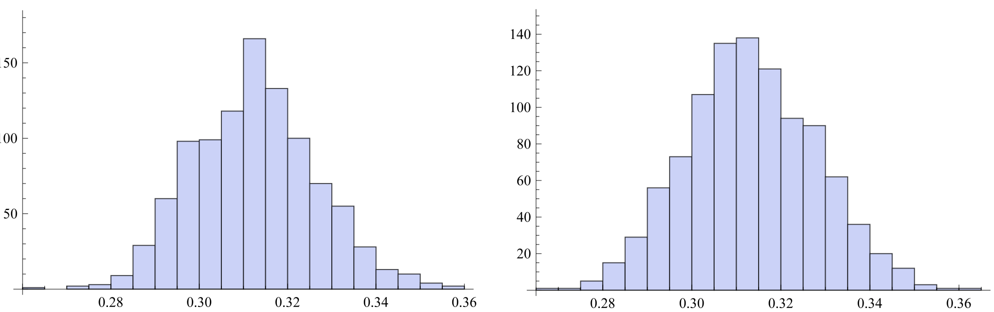

Figure 3: IBM Barrier
Monte Carlo simulation for a barrier option with no bootstrapping (left) and with bootstrapping (right) for an IBM stock with starting price \$100 and strike price \$110.

Binary: In the no bootstrapping case, $s_n^2 = 1.35 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 8.0 \times 10^{-5}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 1.68 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

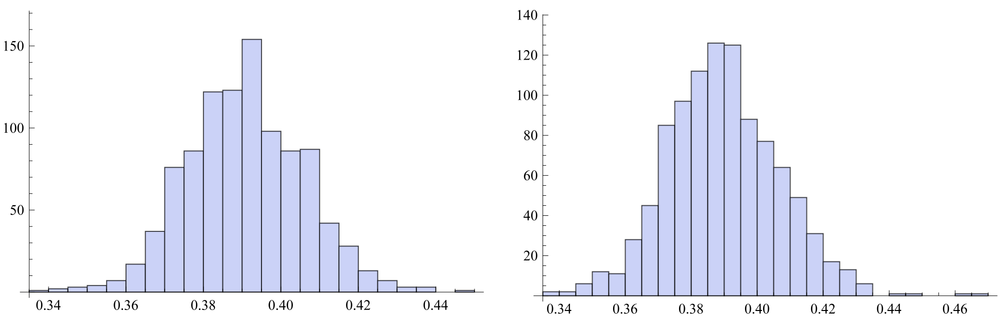

Figure 4: IBM Binary
Monte Carlo simulation for a binary option with no bootstrapping (left) and with bootstrapping (right) for an IBM stock with starting price \$100 and strike price \$110.

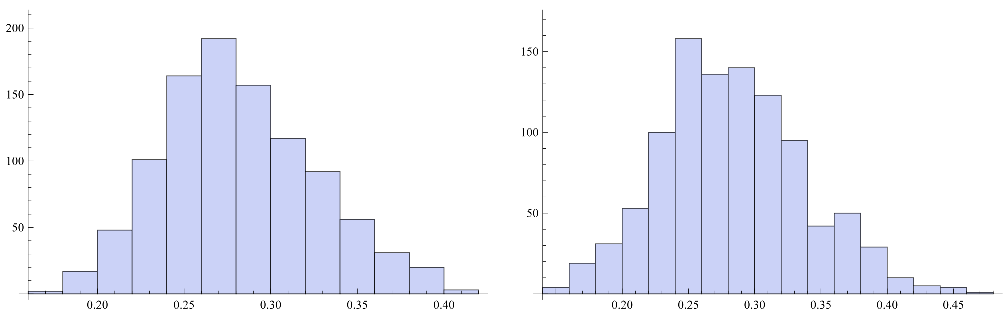

Figure 5: XOM European
Monte Carlo simulation for a European option with no bootstrapping (left) and with bootstrapping (right) for XOM stock with starting price \$30.34 and strike \$40.34.

4.2 Results - Exxon Mobil (XOM)

European: In the no bootstrapping case, $s_n^2 = 1.94 \times 10^{-3}$. In the with bootstrapping case, $s_w^2 = 2.94 \times 10^{-3}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 1.565 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Asian: In the no bootstrapping case, $s_n^2 = 2.321 \times 10^{-5}$. In the with bootstrapping case, $s_w^2 = 2.305 \times 10^{-5}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 1.58 > 1$, we **reject the null hypothesis** and conclude

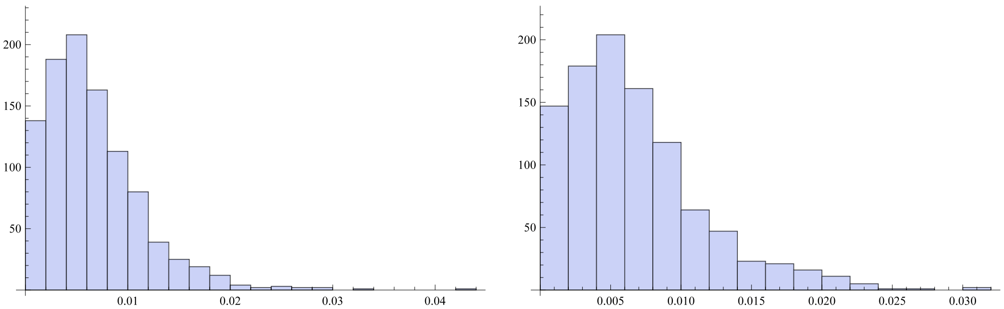

Figure 6: XOM Asian
Monte Carlo simulation for a Asian option with no bootstrapping (left) and with bootstrapping (right) for XOM stock with starting price \$30.34 and strike \$40.34.

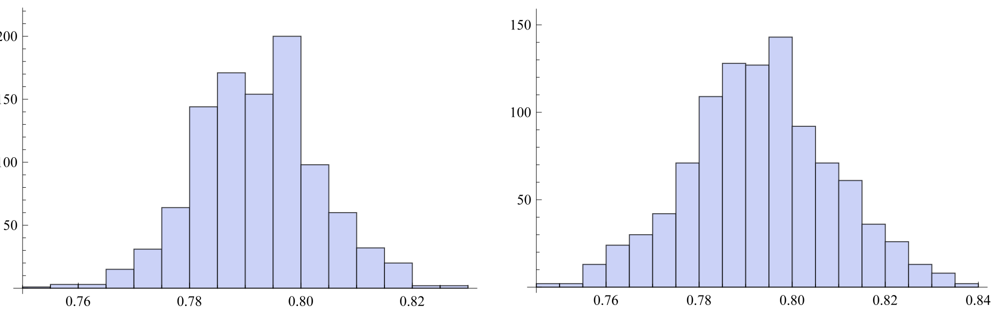

Figure 7: XOM Barrier
Monte Carlo simulation for a barrier option with no bootstrapping (left) and with bootstrapping (right) for XOM stock with starting price \$30.34 and strike \$40.34.

that the variances are significantly different.

Barrier: In the no bootstrapping case, $s_n^2 = 1.22 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 2.44 \times 10^{-4}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 2.003 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Binary: In the no bootstrapping case, $s_n^2 = 1.35 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 8.0 \times 10^{-5}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} =$

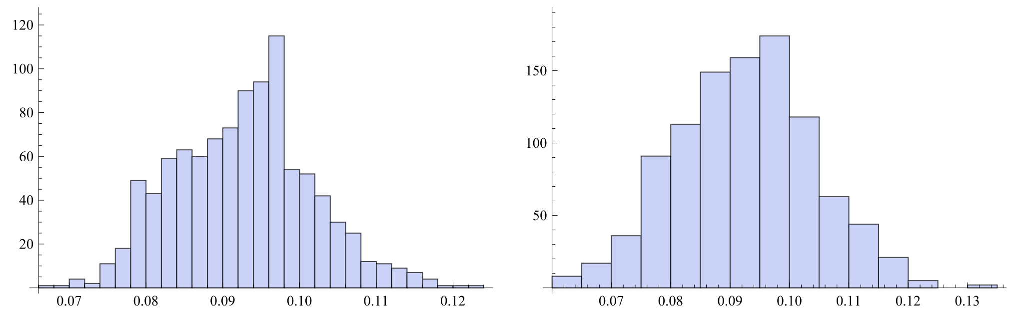

Figure 8: XOM Binary
Monte Carlo simulation for a binary option with no bootstrapping (left) and with bootstrapping (right) for XOM stock with starting price \$30.34 and strike \$40.34.

1. Since $\frac{s_w^2}{s_n^2} = 1.68 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

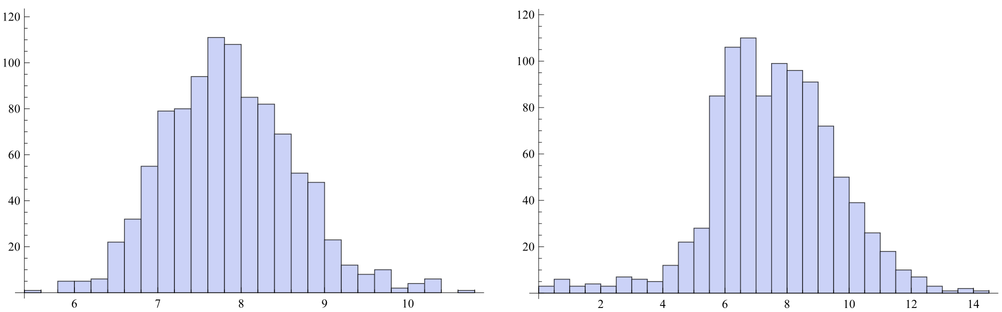

Figure 9: MSFT European
Monte Carlo simulation for a European option with no bootstrapping (left) and with bootstrapping (right) for MSFT stock with starting price \$46.09 and strike \$56.09.

4.3 Results - Microsoft (MSFT)

European: In the no bootstrapping case, $s_n^2 = 0.597547$. In the with bootstrapping case, $s_w^2 = 4.10555$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 6.87067 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Asian: In the no bootstrapping case, $s_n^2 = 868446 \times 10^{-2}$. In the with bootstrapping case, $s_w^2 = 9.76861 \times 10^{-1}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 11.2484 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Barrier: In the no bootstrapping case, $s_n^2 = 2.24773 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 1.32833 \times 10^{-3}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 5.90963 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

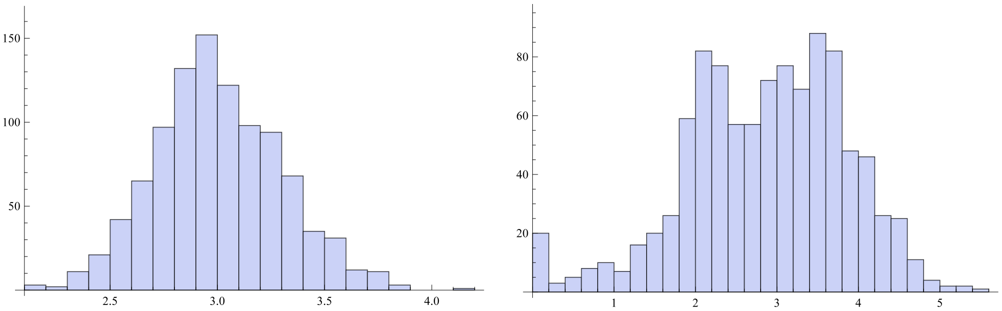

Figure 10: MSFT Asian
Monte Carlo simulation for a Asian option with no bootstrapping (left) and with bootstrapping (right) for MSFT stock with starting price \$46.09 and strike \$56.09.

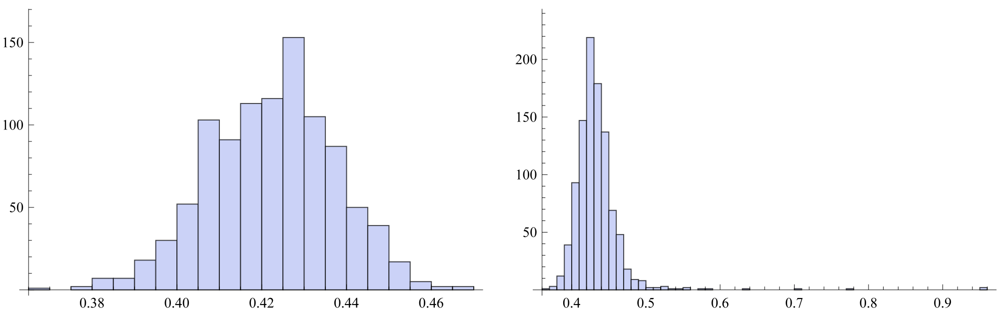

Figure 11: MSFT Barrier
Monte Carlo simulation for a barrier option with no bootstrapping (left) and with bootstrapping (right) for MSFT stock with starting price \$46.09 and strike \$56.09.

Binary: In the no bootstrapping case, $s_n^2 = 1.54314 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 2.05195 \times 10^{-3}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 13.2973 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

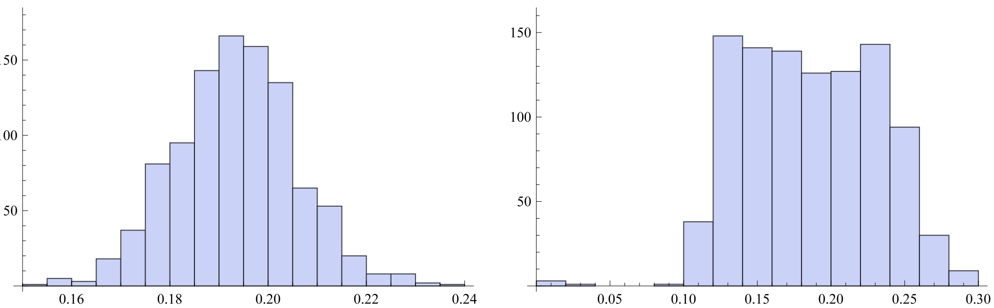

Figure 12: MSFT Binary
Monte Carlo simulation for a binary option with no bootstrapping (left) and with bootstrapping (right) for MSFT stock with starting price \$46.09 and strike \$56.09.

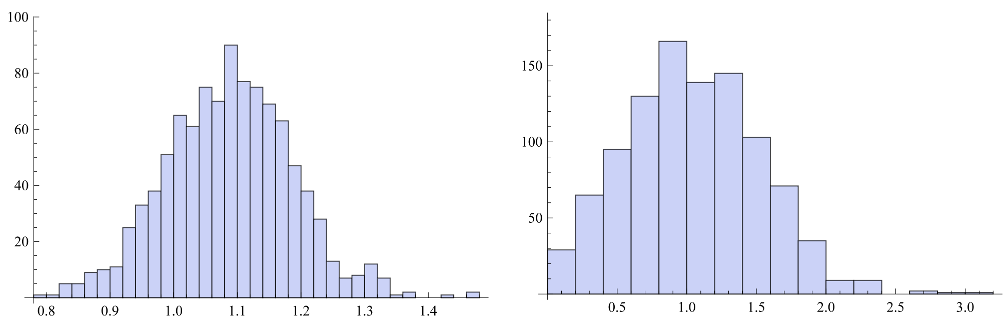

Figure 13: GE European
Monte Carlo simulation for a European option with no bootstrapping (left) and with bootstrapping (right) for GE stock with starting price \$36.09 and strike \$46.09.

4.4 Results - GE

European: In the no bootstrapping case, $s_n^2 = 9.97703 \times 10^{-3}$. In the with bootstrapping case, $s_w^2 = 2.32046 \times 10^{-1}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 23.258 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Asian: In the no bootstrapping case, $s_n^2 = 5.37762 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 9.44296 \times 10^{-3}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 17.5597 > 1$, we **reject the null hypothesis** and

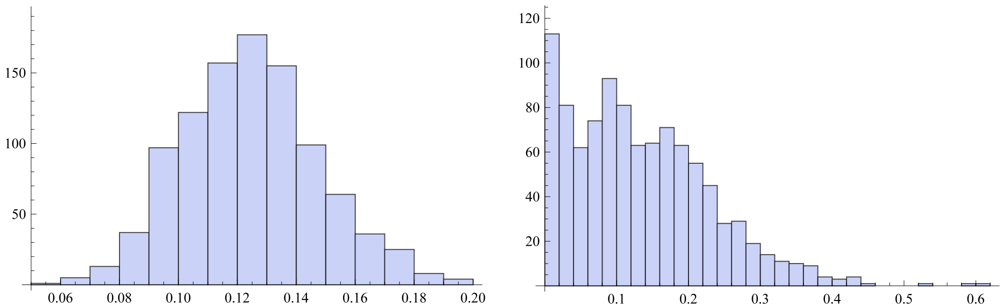

Figure 14: GE Asian
Monte Carlo simulation for a Asian option with no bootstrapping (left) and with bootstrapping (right) for GE stock with starting price \$36.09 and strike \$46.09.

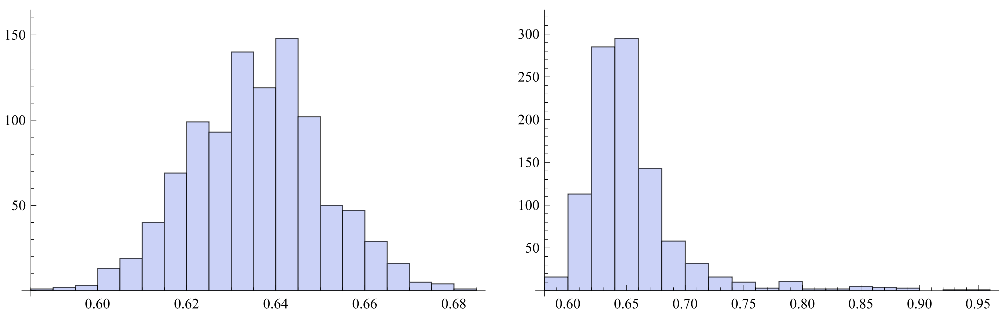

Figure 15: GE Barrier
Monte Carlo simulation for a barrier option with no bootstrapping (left) and with bootstrapping (right) for GE stock with starting price \$36.09 and strike \$46.09.

conclude that the variances are significantly different.

Barrier: In the no bootstrapping case, $s_n^2 = 2.17204 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 1.88664 \times 10^{-3}$. Thus $s_w^2 > s_n^2$, so the critical value becomes $\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 8.68599 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

Binary: In the no bootstrapping case, $s_n^2 = 1.36832 \times 10^{-4}$. In the with bootstrapping case, $s_w^2 = 5.09473 \times 10^{-4}$. Thus $s_w^2 > s_n^2$, so the critical value becomes

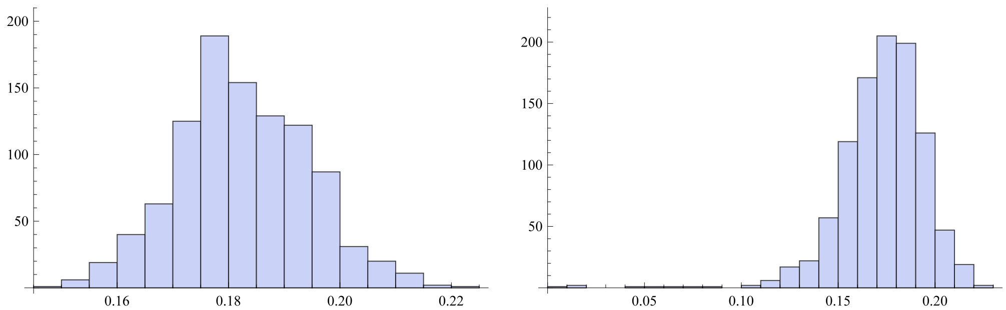

Figure 16: GE Binary
Monte Carlo simulation for a binary option with no bootstrapping (left) and with bootstrapping (right) for GE stock with starting price \$36.09 and strike \$46.09.

$\frac{s_w^2}{s_n^2} \geq f_{0.005,998,998} = 1$. Since $\frac{s_w^2}{s_n^2} = 3.72336 > 1$, we **reject the null hypothesis** and conclude that the variances are significantly different.

All of the previous results are presented in the following table that gives each stock, type of option, the start and strike prices used in the simulation, the sample variances of the no bootstrapping and with bootstrapping conditions, as well as the ratio of the variances.

Table 1: List of All Results

Stock	**Option**	**Start**	**Strike**	s_n^2	s_w^2	**Ratio**
IBM	European	100	110	1.35×10^{-1}	1.63×10^{-1}	1.203
	Asian			2.45×10^{-2}	3.87×10^{-2}	1.58
	Barrier			1.95×10^{-4}	2.15×10^{-4}	1.102
	Binary			2.9×10^{-4}	2.92×10^{-5}	1.275
XOM	European	30.34	40.34	1.94×10^{-3}	2.94×10^{-3}	1.565
	Asian			2.321×10^{-5}	2.305×10^{-5}	1.005
	Barrier			1.22×10^{-4}	2.44×10^{-4}	2.003
	Binary			1.35×10^{-4}	8.0×10^{-5}	1.68
MSFT	European	46.09	56.09	5.98×10^{-1}	4.105	6.871
	Asian			8.68×10^{-2}	9.77×10^{-1}	11.248
	Barrier			2.24×10^{-4}	1.32×10^{-3}	5.9096
	Binary			1.54×10^{-4}	2.05×10^{-3}	13.297
GE	European	36.09	46.09	9.97×10^{-3}	2.32×10^{-1}	23.258
	Asian			5.38×10^{-4}	9.44×10^{-3}	17.560
	Barrier			2.17×10^{-4}	1.88×10^{-3}	8.686
	Binary			1.37×10^{-4}	5.09×10^{-4}	3.723

As one can see, the the estimates under the with bootstrapping condition have higher measures of volatility in each case for the four stocks considered. This particular result has interesting implications for how financial practitioners buy and sell call options.

4.5 Implications

As previously mentioned, the Strong Law of Large Numbers and Central Limit Theorem show that Monte Carlo simulations converge to the true estimate of an option price with high probability because the simulations rely on random trials. However, using a fixed volatility estimate versus many bootstrapped volatility estimates has given estimates with lower volatility in all cases. Thus, accounting for the standard error in our estimate of the volatility function, the distribution of resulting option prices has generally higher variance. Therefore, we conclude that practitioners and researchers interested in pricing financial derivatives should have much less certainty in their price estimates when fixing a single volatility estimate.

One implication of this research is related to the *bid-ask spread*. Bid-ask spread is the difference between the prices quoted for immediate sale (ask) and immediate purchase (bid) of a derivative. Someone who believes the distribution of call option prices to have low variance, might think a bid-ask spread is too high and thus assumes that the cost of an instantaneous transaction is high. Someone using the “with bootstrapping” technique might believe distribution of option prices to have high variance and believe the opposite.

Financial practitioners use bid-ask spread instead of just a price in their decision to buy or sell call options. They might look at the upper and lower percentiles of the distribution of option prices to determine what is reasonable. Using Monte Carlo simulation with bootstrapped volatility estimation, practitioners might consider more trades to be reasonable. Thus, a practitioner who sees a particular bid-ask spread which deviates from his or her simulation might reconsider the way in which their simulation is done, particularly with respect to whether or not the volatility estimate is fixed or bootstrapped in their simulation.

Furthermore, there is potential for further research in the subject. There is potential to run Monte Carlo simulations under the two bootstrapping conditions on different assets, including commodities and other stocks, as well as with different strike prices. Furthermore, many of the options considered have different subcategories. For example, I considered only an "up-and-out" barrier call option, where the contract is terminated if the asset's price goes above a certain level. I could consider a "down-and-in" barrier call option, in which the contract is activated only when the asset's price goes below a certain level. In any case, the research has to potential to give interesting results with important implications.

A Appendix

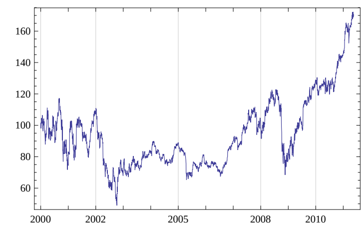

Figure 17: IBM price time series
IBM price fluctuations since January 1, 2000.

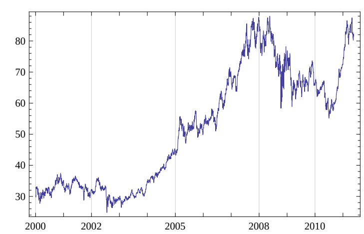

Figure 18: Exxon-Mobil price time series
Exxon-Mobil price fluctuations since January 1, 2000.

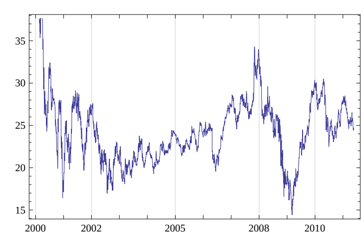

Figure 19: Microsoft price time series
Microsoft price fluctuations since January 1, 2000.

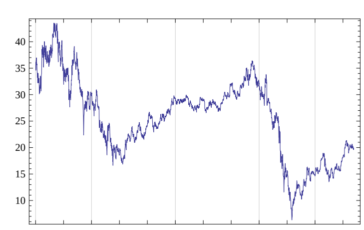

Figure 20: General Electric price time series
General Electric price fluctuations since January 1, 2000.

References

[1] Åkesson, Fredrik and John P. Lehoczky. "Path Generation for Quasi-Monte Carlo Simulation of Mortgage-Backed Securities." *Management Science*, vol. 46 No. 9, Stochastic Models and Simulation (Sep., 2000), pp. 1171-1187.

[2] Boyle, Phelim P. "Options: A Monte Carlo Approach." *Journal of Financial Economics*. Vol. 4 (1977): 323-338.

[3] Bradley, Brendan O. and Murad S. Taqqu. "Empirical Evidence on Spatial Contagion Between Financial Markets." *Finance Letters*, vol. 3, No. 1 (2005), pp. 77-86.

[4] Broadie, Mark and Paul Glasserman. "Estimating Security Price Derivatives Using Simulation." *Management Science*, vol. 42, No. 2 (Feb., 1996). pg 269-285.

[5] Capinski, Marek and Tomasz Zastawniak. *Mathematics for Finance: An Introduction to Financial Engineering*. London: Springer, 2003.

[6] Chen, H. "A Dynamic Programming Approach to the Valuation of Warrants." Ph.D thesis. University of California, Berkeley, 1969.

[7] Corwin Joy, Phelim P. Boyle, Ken Seng Tan. "Quasi-Monte Carlo Methods in Numerical Finance." *Management Science*, vol. 42, No. 6 (Jun., 1996), pp. 926-938.

[8] Cos, J.C. and S.A. Ross. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, vol. 3 (1976): 145-166.

[9] Duffie, Darrell and Peter Glynn "*Efficient Monte Carlo Simulation of Security Prices". The Annals of Applied Probability*, vol. 5, No. 4 (Nov., 1995), pp. 897-905.

[10] Dwight Grant, Gautam Vora, David Weeks. “Path-Dependent Options: Extending the Monte Carlo Simulation Approach.” *Management Science*, vol. 43, No. 11 (Nov., 1997), pp. 1589-1602.

[11] Egloff, Daniel. “Monte Carlo Algorithms for Optimal Stopping and Statistical Learning.” *The Annals of Applied Probability*, vol. 15, No. 2 (May, 2005), pp. 1396-1432.

[12] Fan, Jianqing and Irene Gijbels. “Data-driven Bandwidth Selection in Local Polynomial Fitting: Variable Bandwidth and Spatial Adaptation.” Journal of the Royal Statistical Society. Series B (Methodological). Vol. 57, No. 2 (1995). pg. 371-394.

[13] Fan, Jianqing and Qiwei Yao. “Efficient Estimation of Conditional Variance Functions in Stochastic Regression.” *Biometrika*, vol. 85 No. 3 (Sep., 1998) pg 645-660.

[14] Glasserman, Paul. *Monte Carlo Methods in Financial Engineering*. New York: Springer, 2004.

[15] Hamrick, Jeff. “Estimating the Local Mean Function” from the *Wolfram Demonstrations Project*. http://demonstrations.wolfram.com/EstimatingTheLocalMeanFunction/

[16] Hamrick, Jeff, Yifei Huang, Constantinos Kardaras, and Murad S. Taqqu. “Maximum Penalized Quasi-Likelihood Estimation of the Diffusion Function”. Forthcoming. August 14, 2010.

[17] Hogg, Robert V., Joseph W. McKean, Allen T. Craig. *Introduction to Mathematical Statistics, Sixth Edition*. Pearson: Upper Saddle River, 2005.

[18] Jiang, George J. and John L. Knight. "A Nonparametric Approach to the Estimation of Diffusion Processes, with an Application to a Short-Term Interest Rate Model." *Econometric Theory*, vol. 13, No. 5. (Oct., 1997), pp. 615-645.

[19] Miller, Iriwn and Maryless Miller. *John E. Freund's Mathematical Statistics, Sixth Edition*. Upper Saddle River: Prentice Hall, 1999.

[20] Parkinson, M. "Option Pricing: The American Put." *Journal of Business*, 1976.

[21] Pastorello, Sergio, Eric Renault, and Nizar Touzi "Statistical Inference for Random-Variance Option Pricing" *Journal of Business and Economic Statistics*, vol. 18, No. 3 (Jul., 2000), pp. 358-367.

[22] Samiuddin, M. and G. M. El-Sayyad. "On Nonparametric Kernel Density Estimates." *Biometrika*, vol. 77, No. 4 (Dec., 1990), pp. 865-874.

[23] Wilmott, Paul. *Paul Wilmott on Quantitative Finance, Vol. 1.* West Sussex: John Wiley & Sons, 2000.

[24] Wilmott, Paul. *Paul Wilmott on Quantitative Finance, Vol. 2.* West Sussex: John Wiley & Sons, 2000.

[25] Wilmott, Paul, Sam Howison and Jeff Dewynne. *The Mathematics of Financial Derivatives: A Student Introduction*. New York: Cambridge UP, 1995.