Monte Carlo Simulation for Advanced Option Pricing: A Simplifying Tool

Jason Fink
James Madison University
College of Business

Phone (voice): (540)568-8107
(fax): (540)568-3017

finkjd@jmu.edu

Kristin E. Fink
James Madison University
College of Business

Phone (voice): (540)568-8740
(fax): (540)568-3017

finkke@jmu.edu

Monte Carlo Simulation for Advanced Option Pricing: A Simplifying Tool

Because of the proliferation of quantitative leaning undergraduate finance programs, an increasing number of finance students at the undergraduate and MBA levels are in possession of intermediate to advanced levels of mathematical knowledge. This increase in mathematical skill has opened the door for the Black Scholes model to be presented to advanced undergraduate and MBA students. However, although these students can grasp the weaknesses of the Black Scholes model, they are often not mathematically advanced enough to handle more realistic option pricing models. We demonstrate how Monte Carlo simulation may be employed to open the field of advanced option pricing to students without requiring any more mathematical knowledge than basic calculus and intermediate statistics. As an example, we demonstrate how to simulate option values when the underlying process follows Heston's stochastic volatility process, and motivate the example by demonstrating the significant improvement of a properly specified stochastic volatility model over the Black Scholes model.

1. INTRODUCTION

Since Black and Scholes' path-breaking work on the pricing of options, there has been an exponential rise in the importance of option-related products in the global economy [Black and Scholes, 1973]. Not coincidentally, there have been legions of articles documenting the various deviations of option prices from those given by the Black-Scholes formula (for some recent examples, see Alexander [2004], Bakshi, Cao and Chen

[1997], Heston [1993], and Heston and Nandi [2000]). The typical approach in explaining these deviations is to point out the inaccuracy of the lognormal distribution implied by the geometric Brownian motion (GBM) assumption, and then to posit some more complicated set of dynamics in order to improve this accuracy. The more complicated set of dynamics is then solved, often by employing some Fourier transform, given boundary conditions as dictated by the nature of the derivative.

This set of facts presents difficulties to the professor of advanced undergraduates or first year MBA students. It is often the case that these groups of students are well-versed enough in financial theory to recognize the weaknesses of the baseline models, but not sufficiently well-versed in advanced mathematics to follow the derivation of an option pricing model based on multiple, possibly non-Markovian stochastic processes. Because of the proliferation of quantitative leaning undergraduate programs, the number of students that fall into this category seems to be increasing (see, for example Albert, Buetow, Francfort and Hobson [1999]). The professor is then left to either provide pricing formulas without their derivation, or to leave the students with a set of pricing tools that do not appropriately address reality.

This paper aims to demonstrate how Monte Carlo simulation may be employed to bridge this gap. Arnold and Henry [2003] demonstrate that Monte Carlo simulation can be used pedagogically to visualize the stochastic process underlying the Black-Scholes pricing model. We show that with an intermediate level of statistical training and comfort with a common computer analysis tool (such as Excel), students may be presented with a means to price options on assets whose dynamics are substantially more complicated than the students' mathematical training would otherwise permit.

Section 2 reviews briefly the well established shortcomings of the Black-Scholes model, and introduces the Heston stochastic volatility model as one of many possibilities that can help rectify some of these shortcomings. Section 3 then introduces a framework for Monte Carlo simulation that is grounded in intermediate statistics, and easily extendable to more advanced stochastic processes. Section 4 concludes.

2. THE SHORTCOMINGS OF THE BLACK-SCHOLES MODEL

Inaccuracies of option prices that result from assuming geometric Brownian motion dynamics, as is done in the Black-Scholes model, have been well-documented. Under this assumption, the risk-neutral measure stock prices evolve according to

1) $$dS = (r - \delta)Sdt + S\sigma dW \ ,$$

where S is the current price of the asset underlying the option, δ (possibly zero) is a continuous dividend paid by the underlying asset, r is the risk free rate of interest, and σ is the constant volatility of the option. dt is an infinitesimally short period of time, and dW is a Wiener process – a randomly determined value with mean zero and a variance of dt. This stochastic process, along with a specification of the type of European option, gives rise to the Black-Scholes formula. This formula may be thought of as a function of volatility, the only unknown parameter in the model.

If we observe the prices of an option in the market, we can invert (numerically) the Black-Scholes formula to find the volatility parameter "implied" by the market price, since all other parameters are known. If we have more than one option on the same stock, we should recover the same volatility however many times we repeat this process – if the Black Scholes model assumptions are correct. When this is done with market data, however, we find significant differences in these implied volatilities of options written on the same underlying asset. Graphs of these implied volatilities are by now familiar to those acquainted with the derivatives markets. Options markets in foreign currencies exhibit volatility “smiles” – implied volatilities are higher for in- or out-of-the-money strike prices than for strike prices at-the-money. Equity options have since 1987 exhibited a “volatility skew” – implied volatilities for low-strike options are typically higher than implied volatilities for high-strike options.

Implied volatility smiles for foreign currency options have been documented, for example by Hull and White [1998]. Discussion of both implied volatility smiles of foreign currency options and implied volatility skews of equity options may be found in Hull [2003]. For concreteness, we will examine in greater detail the deviations found in equity options.

Figure 1 presents the implied volatility at the time of market close on Feb. 7, 2005 for a collection of liquid American Call options on Diamonds (DIA), an exchange-traded fund that aims to replicate the Dow-Jones industrial average.[1] Data were collected for February, March, and June expiring options – these groupings are presented in the three panels of Figure 1. Valueline was used to collect the data on the options, and the appropriate interest rates were taken from the Bloomberg website (www.bloomberg.com). Figures of this nature are often used to describe the shortcomings of the Black-Scholes assumptions.[2] Of course, if the GBM assumption is accurate, the skews presented in Figure 1 would be horizontal lines with the implied volatilities at the level of the true volatility of the stock. The observed differences in the implied volatility can represent significant deviations of the Black-Scholes model price from the observed price and confusion concerning the appropriate delta to employ when hedging the stock.

Ad-hoc modifications to the Black-Scholes method, such as using the implied volatility in constructing the delta of the option, have demonstrated hedging effectiveness. However, this effectiveness does not equal that of an appropriately specified, theoretically consistent model. For example, Bakshi, Cao, and Chen [1997] find that, for hedging purposes, a stochastic volatility model such as Heston [1993] provides better dynamic hedges than the ad-hoc Black-Scholes model. Further, many professors grimace at the notion of presenting a model that has no theoretical basis.

The Heston pricing model assumes that stock prices follow the process:

2)
$$\begin{aligned} dS &= (r-\delta)S_t dt + \sqrt{v_t}S_t dW_1 \\ dv &= [\alpha - \beta v_t]dt + \sigma_v \sqrt{v_t}dW_2 \end{aligned}$$

where v_t is the instantaneous variance of the underling asset, α/β is the long-run mean of the instantaneous variance, β is a parameter that indicates the speed of mean-reversion of the variance, and σ_v is the volatility of the variance process (sometimes referred to as the "volatility of volatility"). dW_1 and dW_2 are Wiener processes as defined above, but with correlation ρ. Thus, the Heston model allows for the volatility of the underlying asset to be randomly determined, and assumes that it follows a mean-reverting (Orenstein-Uhlenbeck) process. The solution for the value of European options under these dynamics, while well-known, is beyond the reach of most undergraduate and MBA students. However, the Heston model, along with a number of others, is much more accurate in describing observed option prices than the Black-Scholes model.

Figure 2 demonstrates the absolute value of the percentage pricing error for each option in the collection presented in Figure 1. This pricing error is shown both for the Black-Scholes constant volatility model, and for the Heston stochastic volatility model. For the Black-Scholes reported values, volatility is assumed to be constant at the implied volatility of the closest-to-expire at-the-money option, the February expiring call with a strike price of 107. For the Heston model, parameters are chosen by fitting the five unknown parameters, α, β, σ_v, ρ, and v_t, to the observed option prices.[3] Given the in-sample nature of this comparison between these two models, it must be the case that the Heston model, with a greater number of parameters, fits the data better. However, Bakshi, Cao and Chen [1997] demonstrate that this is also true out-of-sample, as well as in the context of both pricing and hedging. We provide Figure 2 as merely a demonstration of the magnitude of the improvement.

Examining the first panel of Figure 2, we see that the Heston model only outperforms the Black-Scholes model marginally for the soon to expire February options. For these six options, the sum of the absolute *dollar* pricing errors for the Black Scholes model is \$.25, where for the Heston model that value is \$.15. However, it is apparent from the next two panels that the Heston model dramatically outperforms Black-Scholes for the longer dated options. This is an expected outcome, as the effect of stochastic volatility on option prices for short dated options can be expected to be negligible. However as the time horizon lengthens from a short to an intermediate term, the effect of stochastic volatility will become more important. For the 17 longer dated options, in only two cases does the Black-Scholes model provide a better price estimate than the Heston model.

For our complete collection of 23 options, the Black Scholes option price falls outside of the bid-ask spread six times. By comparison the Heston option price falls outside of the bid-ask spread only three times. The sum of the absolute dollar pricing errors for the Black Scholes model over all options is $1.68, for the Heston model it is $.44. Both theoretically and empirically, the Heston model outperforms the Black-Scholes model. Further, it is only one of many such models which do so. However, the Black-Scholes model is overwhelmingly the concluding model taught to even advanced undergraduate and MBA students. The only explanation for this, other than well-deserved homage to Fisher Black and Myron Scholes, is the simplicity of the model. For students capable of learning the better models, a more accurate description of the market is possible. In the next section, we describe how Monte Carlo techniques may be employed to obtain the results of these more advanced models.

3. MONTE CARLO SIMULATION FOR ADVANCED STOCHASTIC PROCESSES

The introduction of Monte Carlo simulation into the option pricing literature is due to Boyle [1977]. Numerous improvements to the basic approach have been developed over the years, but the basic pricing algorithm still proves to be an effective tool. Further, the exponential increase in processing speed of computers has helped substantially in mitigating the run-time obstacles associated with the technique. The basic approach can be presented to students with little more than knowledge of integration and some mathematical statistics. These certainly aren't skills acquired by all finance undergraduates. Nonetheless, with the proliferation of mathematical finance in the undergraduate curriculum, including the frequent cross-listings of these courses between mathematics and finance departments, it is often the case that there are many undergraduates with this level of mathematical knowledge and yet an inability to handle the difficult Fourier inversion problems associated with more realistic option pricing models.

Following Boyle [1977], students at this level may be reasonably asked to consider the integral

3) $$\int_A g(y)dF(y)$$

where g(y) is an arbitrary function, and dF(y) is a probability density function of y with support A and

4) $$\int_A dF(y) = 1.$$

The integral in (3) is:

5) $$E[g(y)]$$

where E[.] is the expectations operator. An application of the law of large numbers gives the result

6) $$\hat{g}(y) \equiv \frac{1}{N}\sum_{i=1}^{N} g(y_i) \overset{a.s}{\rightarrow} E[g(y)]$$

where $g(y_i)$ is a function of a series of iid draws of y_i from probability density function dF(y). The $\overset{a.s}{\rightarrow}$ notation indicates almost sure convergence. The appeal of this result from the perspective of the financial educator is that it is taught in most intermediate statistics courses. Further, if the variances of the $g(y_i)$'s are finite (and greater than zero), then we can appeal to the central limit theorem, and the distribution of the draws of $g(y_i)$ will be normal. This result, also from most standard intermediate statistics courses reads:

Let $g(y_1)$, $g(y_2)\ldots g(y_n)$ be iid with mean μ_g and variance σ_g^2. Then,

7) $$\frac{\hat{g}(y)-\mu_g}{\sigma_g / \sqrt{N}} \sim \Phi(0,1),$$

where $\Phi(0,1)$ represents the normal distribution with a mean of zero and a standard deviation of one. From this, one can easily demonstrate

8) $$\hat{g}(y) \sim \Phi\left(\bar{g}, \frac{\hat{s}}{\sqrt{N}}\right)$$

where $\hat{s}$ is an estimate of σ_g, and $\bar{g}$ is an estimate of μ_g. From this expression, students can form confidence bounds around the simulated estimator.

Of course, the relevance for option pricing comes when g(y) is the present value of the relevant derivative. To use the example of a European put option in an environment with a constant rate of interest

9) $$g(y) \equiv e^{-rT} \max(X - S_T, 0),$$

where S_T represents the asset price at time T, which is taken here to be the expiration of the European option. For a given draw, S_T corresponds to y_i above. The appropriate next step is to determine the driving process for S_T. The natural starting point for this discussion is geometric Brownian motion:

10) $$\frac{dS}{S} = \mu dt + \sigma dW_t$$

where μ is the expected return to the stock, σ is the volatility of the stock, and dWt is a Wiener process. This is an excellent starting point, particularly if the students are advanced enough to have worked through a derivation of the Black-Scholes formulas. Under the risk neutral measure, the expected return (in the absence of dividends) will be the constant interest rate r. If one appeals to Ito's Lemma, it is quickly discovered that

11) $$d \ln S = (\mu - \frac{\sigma^2}{2})dt + \sigma dW_t .$$

One may then design a simulation based on this stochastic differential equation. However, our goal here is to avoid the necessity of appealing to mathematical results beyond the experience of intermediate undergraduate mathematics students. Ito's lemma likely falls into this category. A more straightforward way to proceed would be to appeal straightforwardly to Euler's discretization, and present the differential equation in equation 10 as

12) $$\frac{\Delta S}{S} = \mu \Delta t + \sigma \sqrt{\Delta t} \varepsilon .$$

Here the discretized Wiener process has been expressed as $\sqrt{\Delta t}\varepsilon$, where ε is a standard normal random variable. This is a natural expression for most students who would be involved in a course at this level, particularly if they have seen Hull's discussion of Black-Scholes. This discretization then provides the foundation for the simulation.

Simulation of the discretized process is then straightforward, and found in many of the standard texts.[4] Continuity of thought merits a brief discussion of the simulation process here. One first discretizes time over the life of the option into M time steps. If the option expires at time T, then $\Delta t = T/M$. For each time step, a standard normal random variable, ε, is generated, which results in a stock price for time $S_{t+\Delta t}$. Since $\Delta S \equiv S_{t+\Delta t} - S_{t,}$

13) $$S_{t+\Delta t} = S_t + \mu S_t \Delta t + \sigma S_t \sqrt{\Delta t} \varepsilon .$$

At the conclusion of these M draws of ε, one potential future stock path has been generated, resulting in a potential S_T. An example of such a path is illustrated in Figure 3. Continuing with our example of a European put option, this possible S_T results in a possible value of the put option, given by max $(X-S_T,0)$. This value is then discounted to the present and recorded. This process of generating future option values is repeated N times. The resulting estimator is distributed as given in equation 8. Arnold and Henry [forthcoming] provide an instructive discussion concerning the use of Excel in constructing such a simulation in a geometric (and arithmetic) Brownian motion context that may be readily extended.

The simulation procedure for more elaborate stochastic processes is hardly more difficult. In the case of stochastic volatility, for example, the Euler discretization is very similar to that above. The chief difference is that in the case of stochastic volatility, we must simulate two paths, one for the stock price and one for the volatility. Here are the discretized processes:

$$\begin{aligned} \Delta S &= (r-\delta)S_t\Delta t + S_t\sqrt{v_t\Delta t}\varepsilon_1 \\ \Delta v &= [\alpha - \beta v_t]\Delta t + \sigma_v\sqrt{v_t\Delta t}\varepsilon_2 \,. \end{aligned} \tag{14}$$

Analogous with our discretization of GBM, each of the two Wiener processes have been replaced by $\sqrt{\Delta t}\varepsilon_i$ i=1,2. Since these two processes are correlated in this model, ε_1 and ε_2 have correlation ρ.

Just as the parameter σ must be known in order to simulate the Black-Scholes price of an option, so too must all the parameters for more advanced processes be known in order to simulate the resultant option prices. In the case of the Heston stochastic volatility model we have been using as an example, α, β, σ_v, ρ and v_0 must be known, where the latter is the instantaneous volatility. Further, the current stock price and the dividend yield, δ must also be known.

Let's assume we know these parameters – in practice they would need to be estimated. For this example, we use the parameters used in the Heston model in Section 2 to price the DIA call options. The procedure to simulate the Heston option prices for a European call option is as follows.

For step one, we need to simulate a stock price path and a variance path over the life of the options. First, discretize time into M pieces, where again $\Delta t = T/M$. Next, employing our Euler discretization, generate $S_{t+\Delta t} = S_t + (r-\delta)S_t\Delta t + S_t\sqrt{v(t)\Delta t}\varepsilon_1$. Then, generate $v_{t+\Delta t} = v_t + [\alpha - \beta v(t)]\Delta t + \sigma_v\sqrt{v(t)\Delta t}\varepsilon_2$. In both of these cases, everything on the right hand side of the equation is known except the ε_i terms. These are standard normal random variables, although appropriately simulating their correlation requires some care. Most standard analysis packages, such as Excel, provide a simple way to generate independent standard normal variables. It may be easily demonstrated that if x_1 and x_2 are two such standard normal variables, then $\varepsilon_1 = x_1$ and $\varepsilon_2 = \rho x_1 + (1-\rho^2)^{.5}x_2$ are also standard normal variables, but with correlation ρ. Using these simple transformations, $S_{t+\Delta t}$ and $v_{t+\Delta t}$ are easily simulated.

To complete step 1, repeat the one step ahead simulation M times, until a single complete path is generated for both the stock and the variance.

Step two is to then determine the payoff of the European call option at the last simulated stock price, max $(S_T\text{-}X,0)$, which will be at expiration date T. This payoff is then discounted at the risk free rate and recorded, as we did in the GBM case.

We then repeat steps one and two N times. Recording each of the outcomes then allows us to average them to get an estimate for the call price, as well as to estimate the variance of the discounted payoffs in order to generate the distribution given in equation 8.

Illustrations of 50 sample paths for the stock process and the variance process are given in Figure 4. For the generation of these paths, the parameters were set equal to the ones from our DIA call option example. We can see in that figure that deviations from the current stock price become increasingly likely as we move forward in time, a characteristic this process shares with GBM. Examination of the variance process is also interesting. The mean reverting quality of the volatility process is evident, as the paths seem to rarely get very far from the mean before being drawn back to the long run average variance of α/β, which in this parameterization is approximately .0074 (implying a long run standard deviation of a very low .085, which is quite close to the Black-Scholes implied volatility from our Section 2 example).

A reasonable specification for N and M in these simulations depends on the degree of accuracy needed for the particular application. However, as the benefit of the more complicated dynamics is increased precision, N and M are likely to be chosen reasonably but not prohibitively high. For example, in a simulation conducted for this paper, we set N = 1,000 and M = 500. These settings were used to price the March expiring DIA call option with a strike price of 104, one of the options presented in Section 2. Accordingly, we continue to use our base set of parameters. For our simulation, this yields a simulated option price of \$3.551, with a standard error of .0475. Such values lead to a reasonably wide 95% confidence interval of \$3.458 – 3.644. The analytic Heston value for the option using these parameters is \$3.594, whereas the Black Scholes price using our estimated volatility in Section 2 gave a value of $3.495. The midpoint for the bid-ask spread for this particular option in the market was $3.65.

We have presented this discussion in terms of European options, as would be the case of most professors approaching this subject. However, advances in option pricing via simulation have given rise to the possibility of introducing American options in this topic as well. One technique with particular promise in this regard is that of Longstaff and Schwartz [2001], which allows for the simulation of American options prices while adding only the complication of linear regression to the process.

4. CONCLUSION

Finance students at the undergraduate and MBA levels are increasingly in possession of significant mathematical skills. Nonetheless, professors of classes with even advanced undergraduate and MBA students may find that these students, while knowledgeable enough to grasp the weaknesses of the Black Scholes model, are often not mathematically advanced enough to handle the difficult Fourier inversion problems associated with more realistic option pricing models. However, these students are usually well-versed in intermediate statistics. We show that simulation of more realistic processes requires no more than knowledge of intermediate statistics, opening the field of advanced option pricing to students earlier than has previously been the case.

References

Albert, J, G. Buetow, A. Francfort and H. Hobson "A Model for An Undergraduate Program in Quantitative Finance: The James Madison University Experience," Journal of Financial Education, 25, (Spring, 1999) 1-17.

Alexander, Carol "Normal Mixture Diffusion with Uncertain Volatility: Modelling Short- and Long- Term Smile Effects," *Journal of Banking and Finance*, 28, (December, 2004)

Arnold, T. and S. Henry "An Excel Application for Valuing European Options with Monte Carlo Analysis," *Journal of Financial Education*, *forthcoming*.

Arnold, T. and S. Henry "Visualizing the Stochastic Calculus of Option Pricing with Excel and VBA," *Journal of Applied Finance*, 13, (Spring/Summer, 2003) 56-65.

Bakshi, G., Cao, C. and Z. Chen. "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, (December, 1997) 2003- 2049.

Black, F. and M. Scholes. "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, (May-June, 1973) 637-654.

Boyle, P. "Options: A Monte Carlo Approach," *Journal of Financial Economics*, 4, (May, 1977) 323-338.

Heston, Steven. "A Closed-Form Solution for Options with Stochastic Volatility, with Applications to Bonds and Currency Options," *Review of Financial Studies*, 6, (1993) 327-343.

Heston, Steven and Saikat Nandi. "A Closed-Form GARCH Option Valuation Model," Review of Financial Studies, 13, (2000) 585-625

Hull, J.C. *Options, Futures and Other Derivatives* (Upper Saddle River, NJ), 5th Edition, Prentice Hall.

Hull, J.C. and A. White. "Value at Risk When Daily Changes in Market Variables Are Not Normally Distributed," *Journal of Derivatives*, 5, (Spring, 1998), 9 – 19.

Longstaff, F. and E. Schwartz. "Valuing American Options by Simulation: A Simple Least-Squares Approach," Review of Financial Studies, 14 (Spring, 2001), 113 – 147.

Figure 1: Implied Volatilities of DIA options Feb. 7, 2005

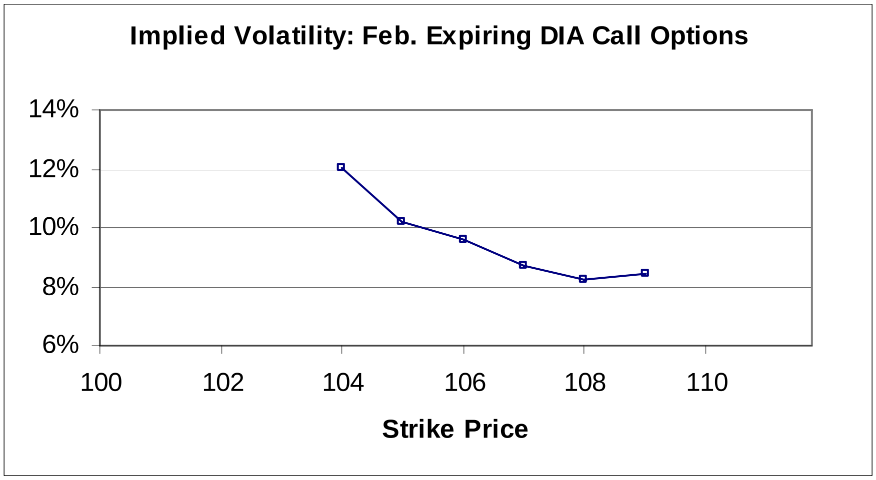

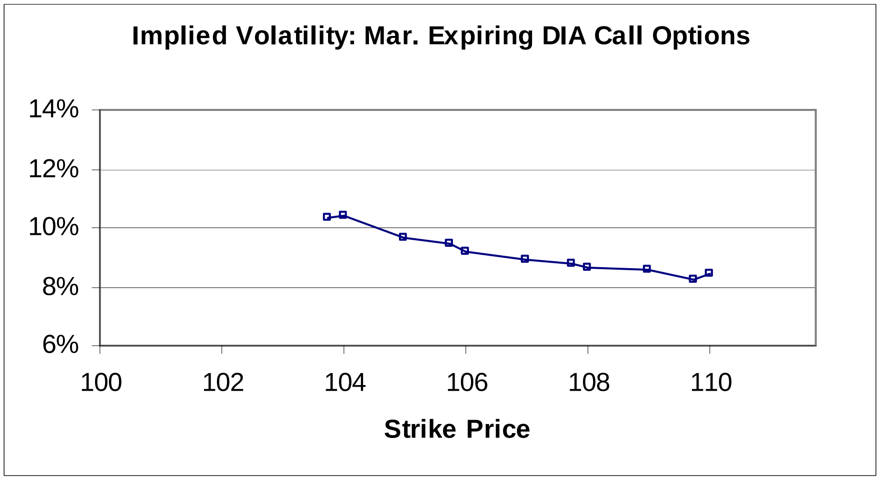

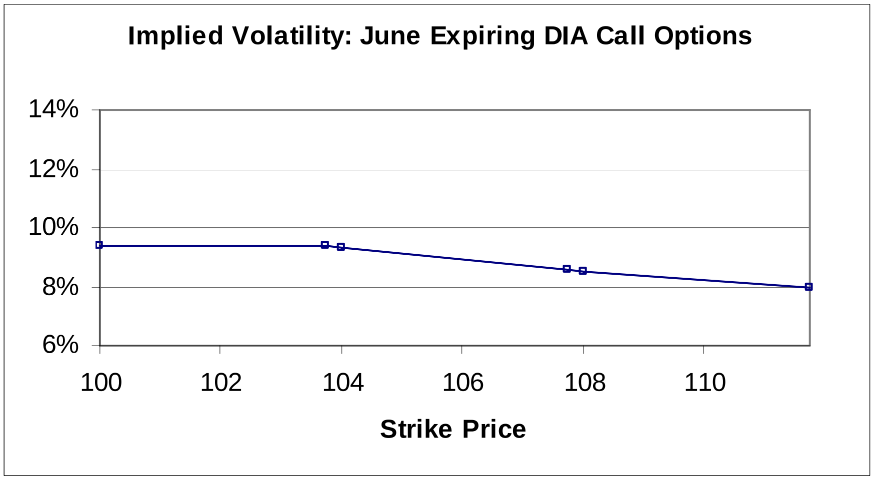

Figure 2: Absolute Percentage Errors of DIA Call Options, Feb. 7, 2005

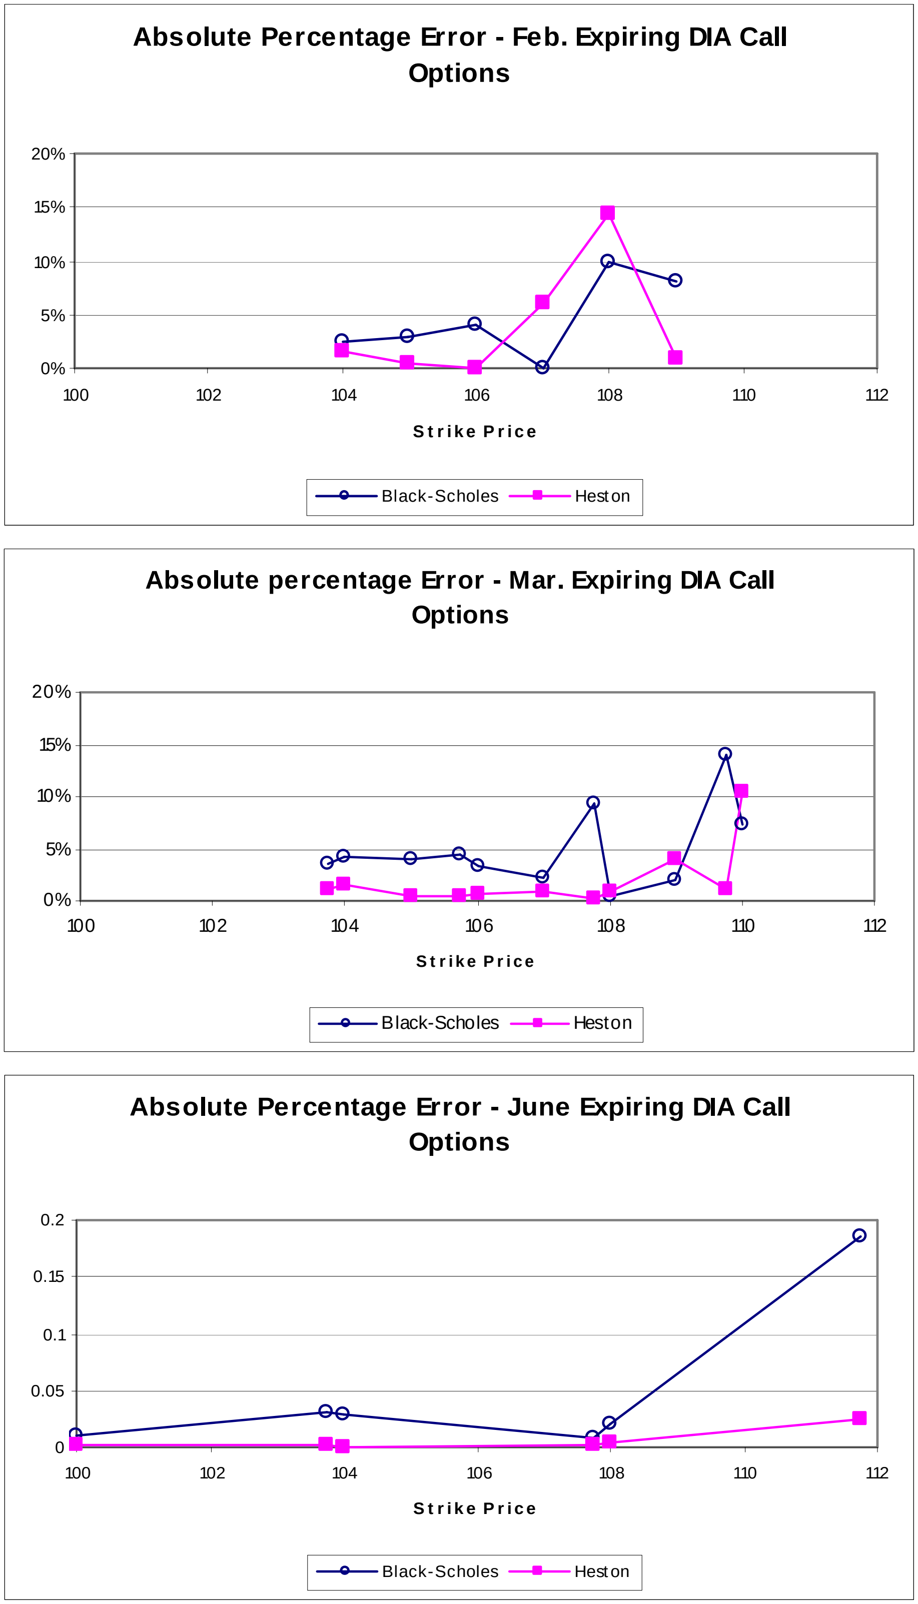

Figure 3: Sample Simulated Stock Price Path

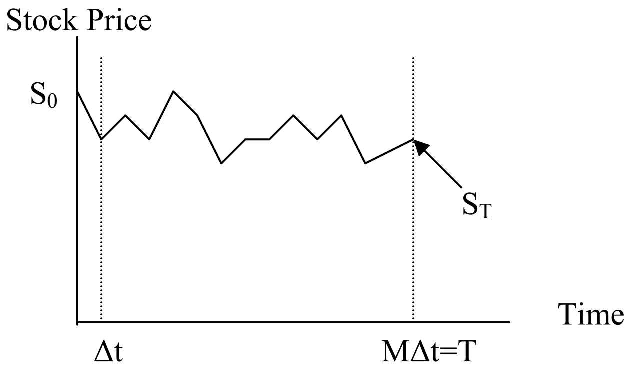

Figure 4: Simulated Paths: Stochastic Volatility

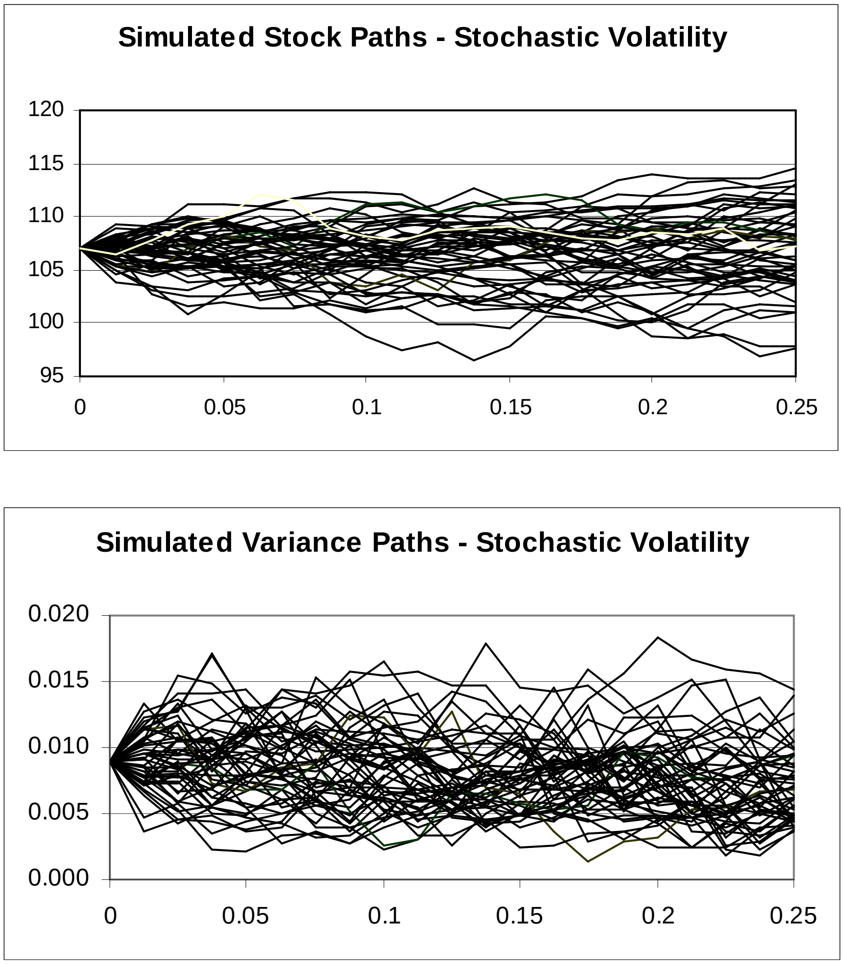

Simulated Stock Paths - Stochastic Volatility

Simulated Variance Paths - Stochastic Volatility

End Notes

---

[1] Since the underlying DIA index is assumed to pay a continuous proportional dividend, we may value American call options as if they were European options, as the two prices are approximately equal.

[2] See, for example, p. 334-336 of Hull [2003].

[3] The Black-Scholes $\sigma = .0868$ for all options. The Heston parameters for all options are: $\alpha = .094$, $\beta = 12.861$, $\sigma_v = .189$, $\rho = -.891$, and $v_t = .009$.

[4] The discussion in Hull [2003] may be found on pages 410-414.