Technology: Journal of The Franklin Institute, Vol. 335A, pp. 57-63, 1998

1072-9240/98 $5.00 +.00

THE PRICING OF PATH-DEPENDENT EUROPEAN OPTIONS VIA MONTE CARLO SIMULATION

Michael A. Pizzi and George L. Montgomery
Montgomery Investment Technology, Inc., 2 Radnor Corporate Center, Suite 121, Radnor, PA 19087

TF 9606-526 M (Received 14 June 1996; accepted 16 March 1998)

The traditional valuation methodologies established in the early 1970s for standard traded options are generally inadequate to handle the valuation of today's more complex exotic options. In order to address the valuation of sophisticated new options, a Monte Carlo technique was developed. Based on efficient market assumptions, the behavior of a stock price over time was assumed to be a generalized Wiener process. Several hypothetical European-style options were valued and compared to close form solutions. Some possible solutions to the problem of the relatively high computational time required for precise estimation are suggested.

INTRODUCTION

Derivative securities have witnessed incredible innovation over the past years. The traditional valuation methodologies established in the early 1970s for standard traded options are generally inadequate to handle the valuation of today's more complex exotic options. This paper presents an approach, based on Monte Carlo simulation, to value European-style path-dependent exotic options. While the accuracy of the Monte Carlo method of option valuation has been known since Boyle's (1977) seminal paper on the subject, it has been generally assumed that the method is too computationally intensive to be of practical use. With recent hardware and software developments and further advances in stratified sampling, this problem can, in part, be solved. Furthermore, many of today's path-dependent options, especially employee stock options, are so complex that the Monte Carlo method must be employed. This paper is divided into five main sections. The first section defines what path-dependent options are and shows how they are used. The next section reviews the Monte Carlo method. The third section develops an equation for the simulation of the path of a stock price based on conditions given by the efficient market hypothesis (EMH), while the fourth section applies the Monte Carlo technique for the valuation of standard options and several types of path-dependent options. The last section is reserved for conclusions.

PATH-DEPENDENT OPTIONS

Options are contracts that give their holder the right, but not obligation, to buy or sell some underlying asset at a fixed price for a specified period of time. A call option is an option to buy a security or commodity at some time in the future. If the time in the future is only the day of expiration of the option, this type of option is referred to as a European-style option. If the option can be exercised at any time over the option life, it is referred to as an American-style option. Put options, on the other hand, are options to sell a security or com-

NOTATIONS

C	Present value of a European-call option
E	Expected value
r	Risk-free interest rate
s	An estimate of the stock volatility
S	Current stock price
t	Current time value
T	Expiration date
X	Exercise price
ε	Random number normally distributed with average = 0 and standard deviation = 1
μ	Stock return rate
σ	Stock volatility

modity and can be either American- or European-style. The payoff of a standard European-call option is the maximum between zero and the difference between the stock price at expiration and the exercise price. Similarly, the payoff of a put option can be expressed as the maximum between zero and the difference between the exercise price and the price at expiration (Table 1). The present value of an option is the discounted value of the expected payoff. For example, for a European-call option, this is expressed as:

$$C = e^{-r(T-t)} E[\max(S - X,0)]. \quad (1)$$

Notice that C is independent of the path that the security follows over the life of the option. If the price at expiration is \$100 and the exercise price is \$90, the payoff if \$10. It makes no difference if the price is \$100 over the entire life of the option or if the price climbed from \$25. The payoff is path-independent, therefore, the discounted expected value of the payoff is also path-independent. A path-dependent option makes S, X, or both statistics based on the path the security follows over the life of the option. The best way to understand path-dependent options is to consider examples. The three most common types of path-dependent options will be considered.

Lookback options

If the strike price, X, is replaced with the minimum value of the security or commodity over the life of the option, a Lookback call option is created. The Lookback period can be over the entire option's life or some sub-period of the option's life. To value a Lookback put, the maximum price reached during the option's life is needed. Lookback options are always at or in the money. The zero in the corresponding formulas from Table 1 can be eliminated without altering the option's payoff. A Lookback can be interpreted as an option to buy at the low or sell at the high. For this reason, Lookback options are always more expensive than standard options. Lookback options began trading in 1982, when the Macotta Metals Corporation began trading Lookbacks on gold, silver, and platinum (Hunter and Stowe 1992a). To consider how a Lookback option may be used, consider a corporation that requires a certain amount of a commodity in one year. If the company wanted to guarantee that it could purchase the commodity at the lowest price over the year, a one-year Lookback call option would be the proper option to purchase.

Table 1. Path-dependent options.

European-call payoff	= max $(S - X,0)$
European-put payoff	= max $(X - S,0)$
Lookback-call payoff	= max $(S_T - S_{min},0)$
Lookback-put payoff	= max $(S_{max} - S_T,0)$
Average-rate-call payoff	= max $(S_{ave} - X,0)$
Average-rate-put payoff	= max $(X - S_{ave},0)$
Average-strike-call payoff	= max $(S - S_{ave},0)$
Average-strike-put payoff	= max $(S_{ave} - S,0)$

Average or Asian options

Average or Asian options are options in which either S or X is replaced with some form of average. If X is replaced, the option is known as an average-strike option. If S is replaced, the option is known as an average-rate option. Options that use averaging terms can become considerably complex. The average can be either arithmetic or geometric, and can possibly be a moving average. The average can be taken over a subperiod of the option's life. Periods can even be given special weighting. Furthermore, the frequency of the averaging must also be taken into consideration. Average-strike options do not have a fixed strike price, but rather the strike is based on some average of the price path. Options involving averaging terms have become extremely popular in recent years. Hunter and Stowe (1992a) stated that, "Many multinational corporations use average rate put options on foreign currencies to hedge their estimated monthly foreign exchange income in an effort to achieve some budgeted average exchange rate for the year." Hunter and Stowe (1992a) then stated that commercial banks offer average rate currency options because they usually have "an average exchange rate on the books" that can be offset by selling average rate currency options.

Barrier options

Barrier options are options in which the payoff is dependent upon the price path reaching a certain barrier. There are four basic types of barrier options: down-and-out, down-and-in, up-and-out, and up-and-in. A down-and-out option goes out of existence if the price path of the underlier falls to a certain barrier. An up-and-out option goes out of existence if the price path of the underlier rises to a certain barrier. A down-and-in option comes into existence if the price path of the underlier falls to a certain barrier. Likewise, an up-and-in option comes into existence if the price path of the underlier rises to a certain barrier. Options ending with '-out' are sometimes referred to as knock-out, disappearance, or extinguishing options, whereas options ending in '-in' are referred to as knock-in or appearance options (Ong 1996). The payoff for barrier options is exactly as for European options except only the barrier condition is taken into consideration. In effect, barrier options are decomposing the payoff of a standard option. Purchasing a knock-out option and a knock-in option is equivalent to purchasing a standard option. From this, it is known that the premium on a knock-in or knock-out option should be less than the premium on a standard option (Ong 1996). The smaller premium for barrier options is the prime reason they are traded (Hunter and Stowe 1992a). It is possible to hedge against large movements in prices without having to pay the full premium for a standard option. Barrier options can also become extremely complex. Multiple barriers and curvilinear barrier options are two examples. Multiple barriers have two or more barrier levels. The simplest type of multiple barrier, the double barrier, will be valued later. Double barriers knock-in or knock-out if the price path reaches either an upper or lower barrier. In a curvilinear barrier, the barrier is not a constant, but rather a function. The most common type of curvilinear barrier is an exponential barrier (Ong 1996).

Other path-dependent options

While the above are the three major categories of path-dependent options, there are several other path-dependent options that are sometimes referred to in the literature; examples include Ladder, Ratchet, and Shout options. All of these options can be related in some way back to the Lookback option. There is a type of path-dependent option that is exchange traded. These are barrier price options referred to as CAP options traded on both the Chicago Board Options Exchange (CBOE) and the American Stock Exchange (AMEX). A CAP option is a type of barrier option. If the barrier is reached, the option pays, on that day, the difference between the underlier (a stock index) and the ceiling (cap). The ceiling on CBOE CAPs is \$30 (Hull 1993).

Any of the general categories of path-dependent options presented above can be combined into a new type of option. Examples can be constructed such as setting a barrier equal to an average (e.g., the price path must have a 30-d average above some constant). Also, path dependency can be incorporated with other non-path-dependent exotics. While these types of options may not be frequently traded, they arise in many employee stock options.

MONTE CARLO SIMULATION

Monte Carlo simulation is a powerful tool that was originally developed to solve problems in atomic physics. Monte Carlo methods essentially reduce complex problems to problems of expected value through the use of simulated random quantities. The description of the Monte Carlo method is usually given in the form of approximating a definite integral. Since such a description requires knowledge of probability theory and integration, Sobol's (1974) more intuitive approach will be used. Consider the problem of estimating the area within the closed curve within the unit square shown in Fig. 1.

One can estimate the area of the curve by sampling N random points in the unit square. The number of points that fall inside the closed curve N' divided by N is a Monte Carlo estimate of the area within the curve. As N approaches ∞, the estimate approaches the true value. An error associated with this estimate can also be calculated. The above is given only for an intuitive understanding of the Monte Carlo method. Sobol (1974) notes that one would not use the Monte Carlo method to estimate an area because superior methods exist.

SIMULATING A PRICE PATH

The weak form of the EMH states that the price in a market should reflect all historical price information. This is to say that the best predictor of the next period's price of a security is the current period price. The fol-

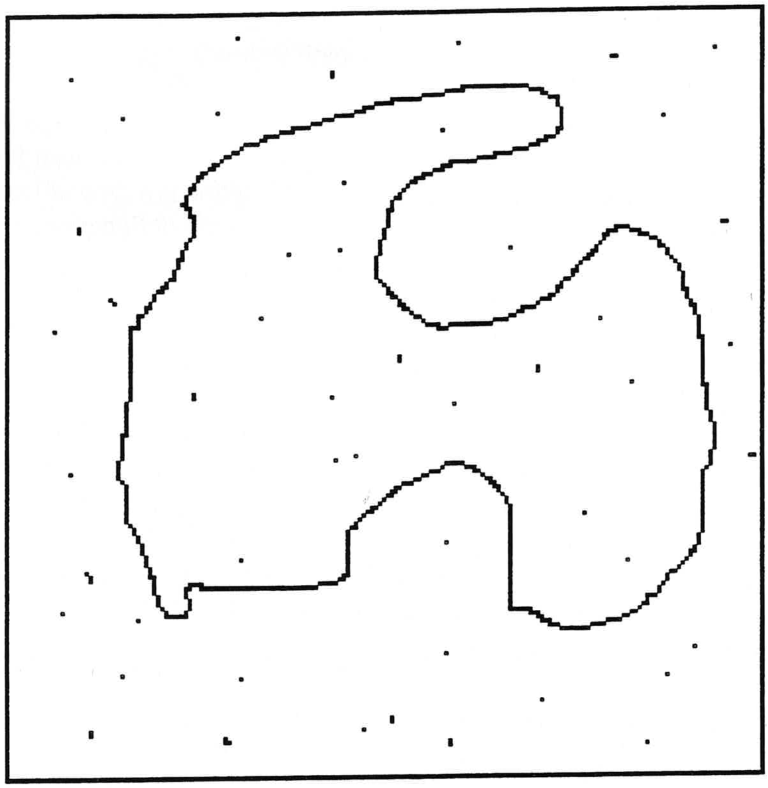

Fig. 1. Estimating the area within the closed curve.

lowing is a brief mathematical description for the generation of a price path. For a more complete description, see Hull (1993). Following Hull (1993), efficient markets can be described by a Markov process. A Markov process is a stochastic process in which only the present value of a variable is relevant to predicting the future. It is generally assumed that markets follow a more particular type of Markov process known as a Weiner process. According to Hull (1993), a variable Z follows a Weiner process if it meets the following two properties. The first property is

$$\Delta Z = \varepsilon\sqrt{\Delta t}, \tag{2}$$

where ε is a random drawing from an N (0,1) distribution and t is time. The second property is that the Δ's are independent. It follows then that Δ has mean 0 and variance equal to Δt. Rewriting Eq. 2 for continuous time yields:

$$dz = \varepsilon\sqrt{dt}. \tag{3}$$

To fully describe the price path of a security, it is necessary to generalize the Weiner process in Eq. 3 as follows:

$$dx = Adt + Bd, \tag{4}$$

where A and B are constants.

The first term on the right-hand side represents a constant drift per unit time. The second term can be though of as the addition of a random component. The randomness is a constant times the Weiner process. While the above model could be used as the basis for the generation of a price path, it fails to take into account two assumptions regarding stock price behavior (Hull 1993). The first assumption is that the required return of an investor is independent of a stock's price (e.g., if an investor requires a 10% return, it is required whether the stock price is $10 or $75). This will invalidate the constant drift term in Eq. 4. The term will be replaced by a drift rate expressed as a proportion of the stock price:

$$dS = \mu\, Sd, \tag{5}$$

where μ is the expected rate of return on the stock and S is the stock price.

The second assumption is similar except that it deals with the random portion in Eq. 4. An investor should be uncertain as to the return on a stock regardless of the stock's price. If σ^2 is defined as variance of a proportional change in the stock price, then $\sigma^2 S^2$ is the instantaneous variance rate of S. From this, the proper model for the basis of a stock price path is:

$$dS = \mu\ Sdt + \sigma Sd\,. \tag{6}$$

The above equation is a specific type of generalized Weiner process known as an Ito process. For a complete description, see Ito (1951) and Hull (1993). Returning to discrete time, Eq. 6 becomes:

$$\Delta S = \mu S \Delta t + \sigma \varepsilon S \sqrt{\Delta t}\,. \tag{7}$$

This equation can be used to simulate stock price paths, but there are several alterations that can be made to make Eq. 7 more useful. From Eq. 7, it results (Hull 1993):

$$S + \Delta S = S e^{\left[\left(\mu - \frac{\sigma^2}{2}\right)\Delta t + \sigma\varepsilon\sqrt{\Delta t}\right]} \tag{8}$$

Since σ^2 is generally unknown, some estimate of the volatility must be used. Furthermore, Cox and Ross (1976) showed, in a risk-neutral world, that the equilibrium rate of return on common stock is equal to the risk-free rate. Therefore, Eq. 8 becomes:

$$S_{t+1} = S_t e^{\left[\left(r - \frac{s^2}{2}\right)\Delta t + s\varepsilon\sqrt{\Delta t}\right]} \tag{9}$$

where s is a measure of volatility over the next period and r is the risk-free rate of interest. This is an equation by which the price path of a stock can be simulated.

PRICING OPTIONS BY MONTE CARLO SIMULATION

The model that was derived in the previous section can now be used to price options. To demonstrate the technique, various options will be priced via Monte Carlo simulation, and the results will be compared to closed form analytic solutions (Table 2). In general, the technique involves simulating the possible price paths based on Eq. 9. The value of the stock price in the next period depends on a random drawing from a N (0,1) distribution. Thus by using different random drawings, one can simulate possible stock prices in the next period. These prices compose possible values of S_{t+1}. These prices can then be placed in Eq. 9, replacing S_t to get possible prices for S_{t+2}. This process is continued until S_{t+n} is reached.

Consider a standard European-call option on a stock with the characteristics listed in Table 2 for Option 1. Since the price path is not required to price a standard European option, only the possible prices at expiration must be simulated. Notice that, because the path is not necessary, Δt is set to (T-t) in Eq. 9. Each random drawing is then associated with a possible stock price at expiration. The number of random drawings is referred to as the number of iterations. Running 500 iterations generates 500 possible values of the stock price at expiration. The payoff of the option is then calculated for each iteration according to:

$$\max\ (S_{t+1} - 20{,}0) \tag{10}$$

and these possible payoffs are then averaged. The average is then discounted to the present and is the Monte Carlo option price. The above example yields a Monte Carlo value of $5.9837. For just 500 iterations, this result compares well with the Black-Scholes analytic pricing for standard European options which is $5.9842.

To price a path-dependent option, a high number of possible paths must be generated. To demonstrate it, consider Option 2 from Table 2. To generate possible price paths, the above information is placed into Eq. 9. Notice that Δt equals 0.00274 which is approximately one day (1/365). Making a single random drawing will generate a possible value for one day into the future. This possible value is placed into Eq. 9 as the current period price, and this is continued until S_{t+30} is reached. This generates a possible price path. This process is repeated numerous times to generate a high number of price paths. The geometric average, S_{ave}, of each price path is then taken. The payoff for each path is then calculated according to:

$$\max\ (S_{ave} - 95{,}0) \tag{11}$$

The payoffs are then averaged and discounted. This is the Monte Carlo price estimate. The Monte Carlo estimate for the above example is $5.772 compared to a closed form value of $5.730.

Table 2. Option valuation.

	Option 1	Option 2	Option 3	Option 4
S	\$25	\$100	\$100	\$100
X	\$20	\$95	\$100	\$95
σ	35%	35%	35%	35%
r	5.5%	7.5%	7.5%	7.5%
T-t	0.5 y	30 d	30 d	30 d
Lower barrier	-	-	\$90	-
Upper barrier	-	-	\$115	-
Iterations	500	2000	2000	2000
	Option value			
Monte Carlo	\$5.9837	\$5.772	\$1.9619	\$6.72
Closed solution	\$5.9842	\$5.730	-	-

A Lookback option can be priced in a similar way. The path is generated as in the above example. Then, instead of taking the average of each path, the minimum value for each path is calculated. The payoff for each path is then calculated according to:

$$\max(S_r - S_{min}, 0) \quad (12)$$

where S_r is the terminal price (in the example above, it would be S_{30}) and S_{min} is the minimum for each path. The average payoff is then calculated and discounted and is the Monte Carlo price estimate. Average-rate and Lookback puts can be valued simply by changing to the appropriate payoff formula described in the first section.

The process to value a barrier option is similar to the above. To demonstrate the valuation of a double barrier, consider the call option listed Option 3 from Table 2. The 30-d price path is simulated as described above. This option will knock-out if the price goes above 115 or falls below 90. In order to value this option, each sample path must be considered. If a sample path reaches or breaks either barrier, its payoff is set to zero. If the sample path did not reach either barrier, then the payoff is calculated as if it were a standard option. The payoffs, including the zeros, are then averaged and discounted. This is the Monte Carlo price estimate. To see if the value in the above example is correct, one can estimate the value of a knock-in option based on the same sample paths. This is done by setting all payoffs to zero except those that break either of the barriers. Averaging the value of all payoffs and discounting will yield the Monte Carlo price. In the above example, the value is \$2.3850. Adding the value of the knock-in and the knock-out should result in a value that is close to the value of a standard option. Recall that buying a knock-out and a knock-in option is equivalent to purchasing a standard option. In this case, the sum is \$4.3469. The Black-Scholes price for a standard European option is \$4.3039.

The last option that is priced, Option 4 from Table 2, demonstrates the flexibility of the Monte Carlo method. Consider an option that knocks-in only if the arithmetic average value of the price of the underlier over the last 15 d of the option's life is above some barrier. It would be difficult, if not impossible, to arrive at a closed form expression for the value of such an option. In order to value the option via Monte Carlo, numerous sample price paths were created as above. The arithmetic average of the last 15 d of each price path were taken, and all payoffs were set to zero except those that meet or exceed the barrier. All payoffs were averaged and discounted in order to arrive at the Monte Carlo price.

The number of iterations used in the above examples is low. The standard errors associated with all of the above examples are all fairly large in relation to the options price, despite the quality of the point estimate. Several solutions to this problem are possible. The first is to perform a higher number of iterations. While this should be done, it should be noted that the standard error is inversely proportional to the square of the number of iterations. Thus, to decrease the error by a factor of 10 requires increasing the number of iterations by a factor of 100. The above examples were calculated using @RISK, a Monte Carlo add-in program for MS-Excel. Two thousand iterations of a 60-d price path take approximately 1 min on a 100 Mhz Pentium-based computer. Two hundred thousand iterations need to be performed in order to decrease the error by a factor of 10. Therefore, even with higher speed computers, Monte Carlo still presents a practical problem in regards to computation speed when an extremely precise value is necessary. There are two approaches to this problem. The classical approach is to use a method of variance reduction such as the control variate technique or the antithetic variable technique. Both techniques are briefly described by Hull (1993). Another possible solution is to use low discrepancy sequences which are sometimes referred to as quasi-random sequences or stratified sampling techniques. Instead of drawing random numbers from a distribution as in Monte Carlo, low discrepancy sequences sample evenly from the distribution. This allows for a superior estimate in far fewer iterations. Boyle (1996) showed that estimates based on low discrepancy sequences converged after 38 400 iterations, whereas Monte Carlo required 614 000 iterations in pricing an Asian oil option. The @RISK program has the ability to perform a type low discrepancy sequence sampling based on a Latin Hypercube. Using this technique usually resulted in a more accurate price than Monte Carlo.

CONCLUSION

In this paper, various types of path-dependent options were reviewed. The behavior of a stock price over time was shown to be a generalized Wiener process based on efficient market assumptions. From this process, a Monte Carlo technique of option valuation was determined. Several hypothetical options were valued and compared to close form solutions. Several possible solutions to the problem of the high computational time required for precise estimation were suggested. Further research is necessary on the use of low discrepancy sequences in option pricing. Also, there are several possible methods for valuing American-style path-dependent options via Monte Carlo. Further research is needed in this area. It may be possible to incorporate information as to the distribution of a particular underlier and incorporate this information into the path-generating formula. This is especially true if markets are not entirely efficient and the assumptions underlying the path-generating formula are invalid.

REFERENCES

Boyle, P.P. Options: A Monte Carlo approach. J. Finan. Econ. 4: 323-38; 1977.

Boyle, P.P. Valuation of exotic options using the Monte Carlo method. In: Nelken, I., ed. The handbook of exotic options. Chicago: Irwin; 1996.

Cox, J.C.; Ross, S.A. The valuation of options for alternative stochastic processes. J. Finan. Econ. 3: 145-166; 1976.

Hull, J. Options, futures, and other derivative securities. Englewood Cliffs: Prentice-Hall; 1993.

Hunter, C.; Stowe, D.W. Path-dependent options. Federal Reserve Bank of Atlanta. Econ. Rev. 77: 29-34; 1992a.

Hunter, C.; Stowe, D.W. Path-dependent options: Valuation and applications. Federal Reserve Bank of Atlanta. Econ. Rev. 77: 30-43; 1992b.

Ito, K. On stochastic differential equations. Mem. Am. Math. Soc. 4: 1-51; 1951.

Ong, M. Exotic options: The market and their taxonomy. In: Nelken, I., ed. The handbook of exotic options. Chicago, IL: Irwin; 1996.

Sobol, I.M. The Monte Carlo method. Chicago, IL: University of Chicago Press; 1974.