Beating the Best: A Neural Network Challenges the Black-Scholes Formula

Mary Malliaris
Management Science Department
Loyola University Chicago
820 N. Michigan Ave.
Chicago, IL 60611

Linda Salchenberger
Management Science Department
Loyola University Chicago
820 N. Michigan Ave
Chicago, IL 60611

Abstract

A neural network model which processes financial input data is developed to estimate the market price of options. The network's ability to estimate option prices is compared to estimates generated by the Black-Scholes model, a traditional financial model. Comparisons reveal that the neural network outperforms the Black-Scholes model in about half of the cases examined.

AI Topic: Neural networks
Domain Area: Financial modelling
Language/Tool: NeuralWorks Professional II™
Status: Prototype has been successfully developed
Effort: Approximately 2 man-years of effort
Impact: Our results show that a neural network can successfully compete with a sophisticated, well-established mathematical model.

The Black-Scholes option pricing model

In 1973, Black and Scholes [1] proposed a model for computing the current market worth of an option. An option is an agreement giving the holder the right to purchase [a call] or sell [a put] some asset at an agreed upon future time, called the date of expiration. The price that will be paid at this future date is called the exercise price of the option. The market price of the option is the price you pay now for the privilege of buying or selling the underlying asset on or before the expiration date. The Black-Scholes model uses five input variables [exercise price of the option, volatility of the underlying asset, price of the underlying asset, number of days until the option expires, and interest rate] to estimate the price which should be charged for an option. The Black-Scholes option pricing formula for calculating the equilibrium price of call options is shown in (1)

$$C = S \cdot N(d_1) - Xe^{-rT} \cdot N(d_2) \qquad (1)$$

where C is the market price to be charged for the option, N is the cumulative normal distribution, T is the number of days remaining until expiration of the option expressed as a fraction of a year, S is the price of the underlying asset, r is the risk-free interest rate prevailing at period t, X is the exercise price of the option and d_1 and d_2 are given by (2) and (3)

$$d_1 = \frac{\ln\left(\frac{S}{X}\right) + \left(r + \frac{\sigma^2}{2}\right) \cdot T}{\sigma\sqrt{T}} \qquad (2)$$

$$d_2 = d_1 - \sigma\sqrt{T} \qquad (3)$$

where σ^2 is the variance rate of return for the underlying asset. For any time interval [0,t] of length t, the return on the underlying asset is normally distributed with variance $\sigma^2 t$ [4].

One of the critical assumptions underlying this model is that the distribution of prices is log-normal and the volatility is constant [1], [3]. For a rigorous presentation of the derivation of the Black-Scholes model, see [5]. The exercise price, number of days to expiration, and closing price are observable. The volatility cannot be directly observed so it is computed implicitly. Most observers use the Implied Standard Deviation of observed option prices as an estimate of volatility [2]. We used at-the-money call options for this estimate and then used that estimate of volatility to calculate all the call options for that day.

Options with an exercise price equal to the closing price of the index are said to be at-the-money. In the pricing of calls, exercise prices less than the closing price are in-the-money, and exercise prices greater than the closing price are out-of-the-money.

Since its introduction in 1973, the Black-Scholes

1043-0989/93 $03.00 © 1993 IEEE

connection weights were initially randomized, and were then determined during the training process.

The generalized Delta rule was used with the backpropagation of error to transfer values from internal nodes. (For a more detailed explanation of backpropagation learning and the generalized Delta rule, see [6].) The sigmoidal function is the activation function specified in this neural network and is used to adjust weights associated with each input node.

Supervised learning was conducted with training sets consisting of the seven predictor variables and the corresponding market price of the option for each exercise price, for each trading day. For the input nodes in which the data was not in ratio form, the values were scaled to be within a range of 0 to 1. This minimizes the effect of magnitude among the inputs and increases the effectiveness of the learning algorithm. The selection of the examples for the training set focused on quality and the degree to which the data set represented the population. The size of the training set is important since a larger training set may take longer to process computationally, but it may accelerate the rate of learning and reduce the number of iterations required for convergence.

The learning rate and momentum were set initially at 0.9 and 0.6, respectively and the learning rate was adjusted downward and the momentum was adjusted upward to improve performance. The training examples were presented to the network in random order to maximize performance and to minimize the introduction of bias. Training was halted after a minimum of 40,000 iterations. The network was implemented using the software package Neuralworks Explorer running on a 386-based microcomputer with a math co-processor.

Results

To compare the estimations made by each model, we compute and report the mean absolute deviation (MAD), mean absolute percent error (MAPE), and mean squared error (MSE) for each of the 5 two-week periods for both in-the-money and out-of-the-money prices. Option prices were estimated from the Black-Scholes model using a computer program based on equations (1)-(3). Neural network estimations were developed by inputting the estimation sets into a trained network.

The initial results showed that, compared to the actual prices, the neural network estimations had a lower MAPE than Black-Scholes for 4 of the 5 two-week periods for the out-of-the-money case, but Black-Scholes was superior for 4 of 5 two-week periods for in-the-money trades. These results are reported in Tables 2 and 3.

Paired sample comparisons tests were run on the Black-Scholes estimates and actual market prices and on the neural network estimates and actual market prices. These results show that the Black-Scholes consistently overprices the options, while the neural network underprices them. We also observe that the standard deviation of the differences is smaller in the neural network prices.

Results of the paired sample comparisons test for the in-the-money cases show that there is a statistically significant difference between the means of the sample of neural network predictions and the sample of actual market prices. The Black-Scholes however, did not show a significant difference from zero, hence it provides a better model for in-the-money, for this data set.

A few observations about the results can be made. First, although we have only presented summary statistics, one can observe similarities between the individual price estimates made by the two models. Each model has difficulty computing prices when the trades are deep in-the-money. This is expected for the neural network because the majority of trades are close to at-the-money and thus, there are insufficient examples to present to the network for these cases. Secondly, we would not expect to achieve results with the neural network which are significantly different than those of Black-Scholes if many traders are using the Black-Scholes model and the market prices reflect their strategies. The neural network is only capable of learning the relationships which are imbedded in the observations. The neural network exhibited a bias of underpricing the options and in fact, may be best utilized as input into another pricing mechanism.

Summary and conclusions

This empirical examination of the Black-Scholes option valuation model and the neural network option pricing model leads to some interesting conclusions. First, while the two modelling approaches differ fundamentally in their methodology to determining option prices, some common results emerge. While the neural network performs better than Black-Scholes on prices out-of-the money, estimations near the expiration date are accurate for both. The neural network may play a valuable role in some type of preliminary data analysis for in-the-money, rather than directly computing prices.

Second, are several limitations which may restrict the use of neural network models for estimation. There is no formal theory for determining optimal network topology and therefore, decisions like the appropriate number of layers and middle layer nodes must be determined using experimentation. The development and interpretation of neural network models requires more expertise from the user than traditional analytical models. Training a neural network can be computationally intensive and the results

closing price are out-of-the-money.

Since its introduction in 1973, the Black-Scholes options pricing model has performed better overall than any model. Empirical tests show that Black-Scholes remains superior among option pricing equilibrium models, with the possible exception of cases in which trades are made deep-in and deep-out-of-the-money. The volume of research which continues to proliferate related to the Black-Scholes model, even 20 years after its introduction, indicates there is considerable interest and value in developing a model which is more robust than Black-Scholes. In addition, there is some reason to believe that the trading process itself may reveal underlying strategies as well as analytical models and there is information to be gained from historical pricing data. Neural networks have been shown to be useful in modelling nonstationary processes and nonlinear dependencies and thus, may represent a channel of investigation in the search for another type of option pricing model.

Methodology

The data set used for this research was developed using option price transactions data published in the Wall Street Journal during the period from January 1, 1990 to June 30, 1990. The exercise price, market price of the option, and closing price of the S&P 100 index are reported for each trading day. The interest rate came from the results of the 3-month US Treasury Bill Monday auction, as reported each Tuesday in the Wall Street Journal. The data set selected for testing includes pricing data from April 23 to June 29, 1990 and includes in-the-money options and out-of-the-money options with time to expiration between 30 and 60 days. Typically, 6 different call prices per day are quoted.

The five variables selected to estimate the market price of the option are those used in the Black-Scholes model; exercise price, time to expiration, closing price, volatility, and interest rate. For the neural network, we added two lagged variables: yesterday's closing price, LAG CLOSE PRICE, and yesterday's market price of the option, LAG MARKET PRICE.

Preliminary data analysis revealed dependencies and relationships between the variables which were used to partition the data sets for the neural network. Experimentation with different training sets showed that better results could be obtained in the neural networks when the data was separated into in-the-money and out-of-the-money groups. Prices in-the-money vary from $60.00 to $0.75; prices out-of-the-money vary from $15.50 to $0.0625. A larger proportion of observations exist for out-of-the-money prices than for in-the-money prices. Correlations were also found between time to expiration and market price of the option, and between the closing price and the market price of the option.

Under supervised learning, the feedforward, backpropagation neural network learns relationships between input and output variables during a training process, as data are presented to the network. One approach to testing the performance of the network is to check its accuracy in estimating values for a holdout sample generated from the training set. For evaluating the performance of the option price neural network, we selected a more realistic and more difficult performance measure. The network was trained using historical data and option price estimations for a future period were developed with the trained network and compared to actual prices.

To capture the volatile nature of the options market, a relatively short time frame was used for the training sets and testing sets. The testing sets were developed using a two-week time frame; this was a convenient choice because interest rate and volatility changed weekly and were relatively stable over a two-week period. Five two-week periods were selected for price estimation; the weeks beginning April 23, May 7, May 21, June 4, and June 18. To provide the neural network models with a variety of examples, each training set included as many observations as necessary to provide at least one full cycle (30 days prior to the estimation period) of pricing data.

The neural network model

Since feedforward, single hidden layer neural networks have been successfully used for classification and prediction, we selected this network model for our initial experiments and used the backpropagation training algorithm. A neural network consisting of 7 input nodes, 4 middle layer nodes, and 1 output node was developed. The input nodes represent the five financial variables used in the Black-Scholes model (EXER, DAYS, CLOSE PRICE, VOL, and INT) and two lag variables (LAG CLOSE PRICE and LAG MARKET PRICE), and the output node (MARKET PRICE) represents the market price of the option.

The network is fully connected, with a direct connection from exercise price (EXER) to the output node (MARKET PRICE). Better results were achieved with this additional connection because of the linear dependence between EXER and MARKET PRICE observed in the data set and verified with a series of regression models. All the

are sensitive to the selection of learning parameters, activation function, topology of the network, and the composition of the data set.

Thirdly, the paper illustrates that the neural network methodology offers a valuable alternative to estimating option prices to the traditional Black-Scholes model. The evidence reported here is encouraging, particularly in view of the essentially undisputed superiority of the Black-Scholes model. Analytically, it is remarkable, that the well-developed methodology of Black-Scholes, with its explicit formula for pricing options, derived using sophisticated financial arbitrage arguments and advanced stochastic calculus techniques, can actually be approximated by neural networks.

References

1. Black F and Scholes M (1973) The pricing of options and corporate liabilities. J. Pol. Econ. **81**, 637-654.

2. Chesney M and Scott L (1989) Pricing European currency options: a comparison of the modified Black-Scholes model and a random variance model. J. Fin. Quant. Anal. **24**, 267-284.

3.Chiras D and Manaster S (1978) The information content of option prices and a test of market efficiency. J. of Finl. Econ. **6**, 213-234.

4. Macbeth J and Merville L (1979) An empirical examination of the Black-Scholes call option pricing model. J. Fin. 34(5), 1173-1186.

5. Malliaris A and Brock W (1982) Stochastic Methods in Economics and Finance, North-Holland, Amsterdam.

6.Rumelhart D E and McClelland J L (1986) Parallel Distributed Processing. MIT Press, Cambridge, MA.