32

NEURAL NETWORKS FOR PREDICTING OPTIONS VOLATILITY

Mary Malliaris and Linda Salchenberger

INTRODUCTION

The desire to forecast volatility of financial markets has motivated a large body of research during the past decade (Engle and Rothschild 1992). Volatility is a measure of price movement often used to ascertain risk. Relationships between volatility and numerous other variables have been studied in an attempt to understand the underlying process so that accurate predictions can be made. The ability to accurately forecast volatility gives the trader a significant advantage in determining options premiums.

Both researchers and traders use two estimates of option volatility: the historical volatility and the implied volatility. It is almost routinely reported in various publications of exchanges that these two series differ, technique has become the standard method of estimating volatility at the moment of trading

NEURAL NETWORKS FOR PREDICTION

While there are dozens of network paradigms, the back-propagation network has frequently been applied to classification, prediction, and pattern-recognition problems. Financial applications of neural networks include underwriting (Collins, Ghosh, and Scofield 1988), bond rating (Dutta and Shekhar 1988), predicting thrift institution failure (Salchenberger, Cinar, and Lash 1992), and estimating option prices (Malliaris and Salchenberger 1993). The term *back-propagation* technically refers to the method used to train the network, although it is commonly used to characterize the network architecture. For details of this method, see Rumelhart and McClelland (1986). Currently, a number of variations on this method exist that overcome some of its limitations.

DATA AND METHODOLOGY

Data have been collected for the most successful options market: the S&P 100 (OEX), traded at the Chicago Board Options Exchange. Daily closing call and put prices and the associated exercise prices closest to at-the-month, S&P 100 Index prices, call volume, put volume, call open interest, and put open interest were collected from *The Wall Street Journal* for calendar year 1992.

Three estimates for the historical volatilities using Index price samples of sizes 30, 45, and 60 were computed for each trading day in 1992. We also used the Black–Scholes model to calculate implied volatilities for the closest at-the-money call for three contracts: those expiring in the current month, those expiring one month away, and those expiring two months away (nearby, middle, and distant, respectively). Thus, we have approximately 250 observations for six series of volatilities for use in our study.

Comparisons were made between the nearby historical, implied, and network volatility estimates. Because the neural network must have sufficient previous data in order to generalize, these estimates were developed using each method for June 22 through December 30, 1992. Trading cycles

were used as th
the third Friday

A COMPARIS
VOLATILITY I

The historical a
together in Figı
be observed, th
tility used by m
volatility is an
not surprising
for each day tł
fluctuations. Tl
information fro
Thus, the impl

Figure 32.1
Historical an
June 22 throu

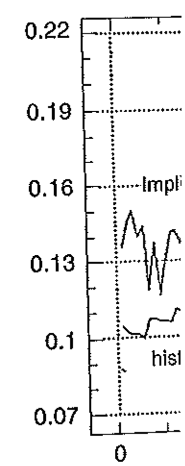

were used as the prediction periods, with each trading cycle ending on the third Friday of the month.

A COMPARISON OF HISTORICAL AND IMPLIED VOLATILITY ESTIMATES

The historical and implied volatility for the nearby contract are graphed together in Figure 32.1 for June 22 through December 30, 1992. As can be observed, the historical estimate significantly underestimates the volatility used by most traders, that is, the implied volatility. Since the historical volatility is an average based on returns from 30 preceding days, it is not surprising that the estimate smoothes out the peaks, giving a value for each day that is less variable and thus less sensitive to daily market fluctuations. The implied volatility for any given day uses only trading information from that day, not a previous time period, to generate a value. Thus, the implied volatility is more reflective of market changes.

Figure 32.1

Historical and Implied Volatilities

June 22 through December 30, 1992

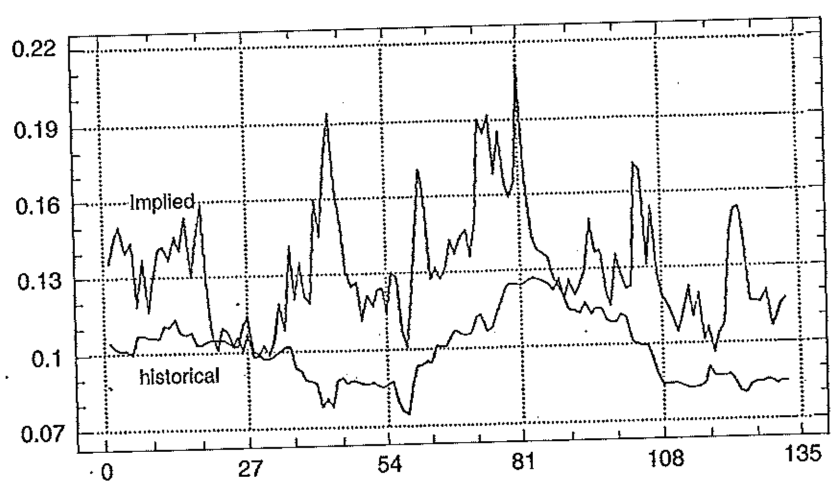

The average MAD (mean absolute deviation) and MSE (mean squared error) for the entire forecasting period, from June 22 through December 30 were 0.0331 and 0.0016. The proportion of times the historical volatility correctly predicted that the implied volatility would increase or decrease are shown in the last column of Table 32.1. An overall average of the number of times a change was correctly indicated is .4439, that is, a little less than half of the time.

DEVELOPMENT OF THE NEURAL NETWORKS

To develop a neural network that is capable of generalizing a relationship between inputs and outputs, the training set selected must contain a sufficient number of examples that are representative of the process being modeled. Therefore, the neural network models developed to predict volatility were trained with data sets from historical data from January 1 through July 18 and used to make predictions for six trading cycles beginning with the period July 20 through August 21 and ending with the period November 23 through December 31. All prior historical data was used when predicting the volatility for the next trading period. Predicting the volatility for the next cycle is a rather rigorous test of the forecasting capabilities of the network since we are asking it to predict volatility for up to 30 days in the future.

There is no well-defined theory to assist with the selection of input variables, and, generally, one of two heuristic methods is employed. One

Table 32.1

A Comparison of Historical and Implied Volatilities

Dates of Forecast	MAD	MSE	Correct Directions
June 22–Jul 19	.0318	.0012	8/19 = .421
July 20–Aug 21	.0292	.0019	11/25 = .440
Aug 24–Sep 18	.0406	.0018	12/18 = .667
Sep 21–Oct 16	.0479	.0027	7/20 = .350
Oct 19–Nov 20	.0213	.0008	14/25 = .560
Nov 23–Dec 18	.0334	.0014	8/18 = .444
Dec 21–Dec 30	.0294	.0009	2/6 = .333

approach is to include all the variables in the network and perform an analysis of the connection weights or a sensitivity analysis to determine which may be eliminated without reducing predictive accuracy. An alternative is to begin with a small number of variables and add to new variables that improve network performance. In this research, the latter was used, and variables were selected using existing financial theory, sensitivity analysis, and correlation analysis. Thus, a number of preliminary models were developed to determine which input variables of the group available in the data set would best predict volatility.

The first models were developed with variables representing volatility lagged from three to seven periods to determine an appropriate set of lag variables. Next, other networks were developed and trained to determine which variables were the best predictors of volatility. The final models include the following 13 variables: change in closing price, days to expiration, change in open put volume, the sum of the at-the-money strike price and market price of the option for both calls and puts for the current trading period and the next trading period, daily closing volatility for the current period, daily closing volatility for next trading period, and four lagged volatility variables. By including both the time-dependent path of volatility and related contemporaneous variables in our model, we obtained better predictions.

The back-propagation network developed to predict volatility has 13 input nodes representing the independent variables used for prediction, one middle layer consisting of 9 middle nodes, and an output node representing the volatility. The cumulative delta rule for training was selected, with an epoch size of 16, a decreasing learning rate initially set at 0.9, and an increasing momentum, initially set at 0.2. The networks were trained using Neuralworks Professional II software from Neuralware.

A COMPARISON OF THE NEURAL NETWORK AND IMPLIED VOLATILITY ESTIMATES

Using historical volatility as a benchmark, we evaluated the performance of the neural network by measuring mean absolute deviation, mean squared error, and the number of times the direction of the volatility (up or down) was corrected predicted. These results are shown in Figure 32.2 and Table 32.2, where comparisons are made between the volatility forecasted by the network and tomorrow's implied volatility. The overall MAD

Figure 32.2

Network and Implied Volatilities

June 22 through December 30, 1992

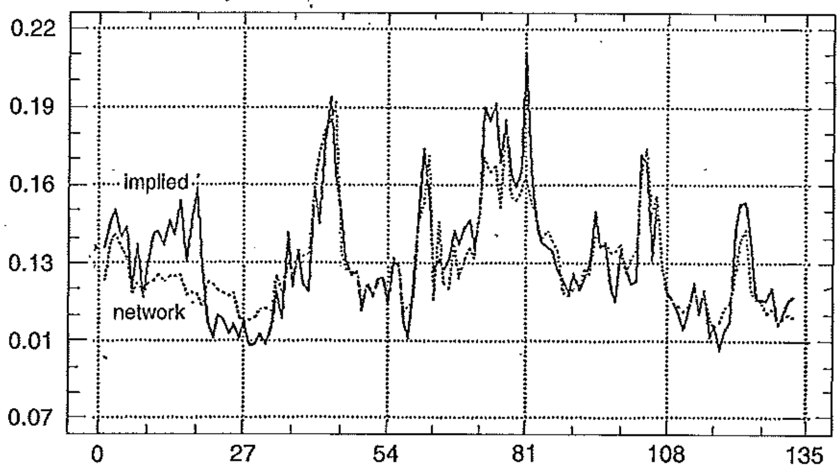

Table 32.2

Neural Network and Implied Volatilities

Dates of Forecast	MAD	MSE	Correct Directions
Jun 22–Jul 19	.0148	.0003	16/19 = .842
July 20–Aug 21	.0107	.0002	16/25 = .640
Aug 24–Sep 18	.0056	.0001	13/18 = .722
Sep 21–Oct 16	.0127	.0003	19/20 = .950
Oct 19–Nov 20	.0059	.0001	20/25 = .800
Nov 23–Dec 18	.0068	.0001	15/18 = .833
Dec 21–Dec 30	.0039	.0000	5/6 = .833

for the entire period was .0116, and the MSE was .0001 as compared to 0.0331 and 0.0016, when the historical was compared to the implied volatility. Furthermore, for each forecasting period, the MAD and MSE were considerably lower (see Tables 32.1 and 32.2). In each of the time periods, the proportion of correct predictions of direction made by the neural network was greater than that of historical volatility. The overall proportion of correct direction predictions was 0.794, as compared to .4439 for the historical volatility estimate. This is not surprising since historical volatility smoothes out the estimate because it is an average of 30 values. The correlation between the implied volatility and the volatility predicted by the network is 0.85, as compared with 0.31 for the historical volatility, at the 5 percent level of significance.

DISCUSSION

The results of this comparative study of neural networks and conventional methods for forecasting volatility are encouraging. Because historical estimates are traditionally poor predictors, traders have been forced to rely on formulas such as Black–Scholes, which can be solved implicitly for the real-time volatility. But these models are difficult to use and limited since they can only provide estimates that are valid at that current time. Furthermore, they fail to incorporate knowledge of the history of volatility. The neural network model, on the other hand, employs both short-term historical data and contemporaneous variables to forecast future implied volatility.

The neural network approach has two advantages that make it more useable as a forecasting tool. First, predictions can be made for a full trading cycle, thus avoiding the problems associated with the need for real-time calculations. Second, and more importantly, the network forecasts, in the cases we tested, were very accurate estimates of the volatility preferred by traders.

The limitations of neural networks as financial modeling tools are well documented. Unlike the more familiar analytical models, a trained neural network does not provide information about the underlying model structure. It is often viewed as a black box since there are no theory-based methods available to interpret and analyze network parameters. Neural networks lack systematic procedures for developing network architecture, selecting training and testing sets, and setting network parameters and thus are difficult to develop. Explicit knowledge of the phenomenon being predicted is required to assist in variable selection.

There are several ways to extend this research. While the performance of these networks in forecasting volatility is superior to the use of historical volatility, improvement may be possible through experimentation with other variables and network architectures. In this chapter, we report results for predicting nearby volatility. However, networks for predicting middle and distant volatility have been developed, using different variables and different network architectures.

REFERENCES

Black, F., and M. Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81(1973) pp. 637–54.

Choi, J.S.,,and M. E. Wohar. Implied Volatility in Options Markets and Conditional Heteroscedasticity in Stock Markets. *The Financial Review* 27, no. 4 (1992) pp. 503–30. Collins, E., S. Ghosh, and C. Scofield. An Application of a Multiple Neural-Network Learning System to Emulation of Mortgage Underwriting Judgments. *Proc. of the IEEE International Conference on Neural Networks*, 1988, pp. 459–66.

Dutta, JS., and S. Shekhar. Bond-rating: A Non-Conservative Application of Neural Networks. *Proc. of the IEEE International Conference on Neural Networks*, 1988, pp. 142–50.

Engle, R.F., and M. Rothschild. Statistical Models for Financial Volatility. *Journal of Econometrics* 52 (1992) pp. 1–4.

Malliaris, A.G. *Stochastic Methods in Economics and Finance*. Amsterdam: North Holland Publishing Company, 1982.

Malliaris, M.E., and L. Salchenberger. Beating the Best: A Neural Network Challenges the Black–Scholes Formula. *Proc. of the Ninth Conference on Artificial Intelligence for Applications*. Los Alamitos, CA: IEEE Computer Society Press, 1993, pp. 445–49.

Rumelhart, D.E., and J. L. McClelland. *Parallel Distributed Processing*. Cambridge, MA: MIT Press, 1986.

Salchenberger, L., E. M. Cinar, and N. Lash. Neural Networks: A New Tool for Predicting Thrift Failures. *Decision Sciences* 23, no. 4 (1992), pp. 899–916.

A NONPA
APPROAC
HEDGING
SECURITI
NETWORI

James

This article reports research
Intelligence Laboratory and
Engineering. We thank Har
and Petr Adamek, Federico
seminar participants at the D
Finance Association for help
acknowledge the support of
Research contract N00014-92
Research under contract NOC
under contract ASC-9217041
HPCC program), by the Labc
of this research was conduct