DO-AHEAD REPLACES RUN-TIME: A NEURAL NETWORK FORECASTS OPTIONS VOLATILITY

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: In this paper, we compare existing methods of estimating the volatility of daily S&P 100 Index for options. The implied volatility, calculated via the Black-Scholes model, is currently the most popular method of estimating volatility and is used by traders in the pricing of options. Historical volatility has been used to predict the implied volatility, but the estimates are poor predictors. A neural network for predicting volatility is shown to be far superior to the historical method while overcoming the run-time limitation of the Black-Scholes approach.

AI topics: neural networks, financial AI applications, forecasting

1. 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 may be made (Merville and Pieptea, 1989; Choi and Shastri, 1989; Haugen et.al. 1991; Lockwood and Linn, 1990; and Dubofsky 1991). The ability to accurately forecast volatility gives the trader a significant advantage in determining options premiums.

Both reseachers 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, but no significantly better forecasting model of volatility has emerged. The purpose of this research is to compare these two existing methods of predicting volatility for S&P 100 options with a new approach which uses neural networks. Neural networks, which have been shown to effectively model nonlinear relationships, prove to be a superior approach to predicting options volatility in all cases tested and can be used to develop monthly forecasts.

2. CALCULATING HISTORICAL AND IMPLIED VOLATILITIES

In their seminal work on pricing options, Black and Scholes (1973) assumed that the price of the underlying asset follows an Itô process

$$dS/S=\mu dt+\sigma dZ \tag{1}$$

where dS/S denotes the rate of return, μ is the instantaneous expected rate of return, σ is the expected instantaneous volatility and Z is a standardized Wiener process ,or dZ is a continuous-time random walk. To simplify their analysis, Black and Scholes assumed that both μ and σ were constants and by using an elegant arbitrage argument, they derived their call option pricing model. Their formula expresses the call price C, as a function of five inputs

$$C=C(S,X,T,\sigma,r) \tag{2}$$

where S is the current price of the underlying asset, X is the exercise or strike price, T is the time from now to expiration of the option, σ is the expected instantaneous volatility and r is the riskless short term rate of interest.

Observe that the μ of equation (1) does not appear in (2). The mathematical derivation of the call option pricing formula as shown in Lee, Finnerty and Wort (1990) or Malliaris (1982) shows that arbitrage requires that the per unit of risk excess returns between two appropriately designed portfolios must be equal. Making the necessary substitutions in this arbitrage relationship, the term containing μ drops out. With μ now out of the picture and with four of the five remaining variables directly observable, an estimate of the asset's volatility σ in (2) becomes the focal point of attention for both theorists and traders.

There are two main approaches to estimating and predicting the nonconstant σ: the historical approach and the implied volatility approach. The historical approach is the simplest because tomorrow's volatility σ_{t+1} is an estimate obtained from a sample, of a given size, of past prices of the underlying asset. Suppose that the sample size is n and let

$$S_{t-n+1}, \ . \ . \ . \ , S_{t-1}, \ S_t$$

denote daily historical prices for the underlying asset. To get an estimate for σ_{t+1}, first compute daily returns, r_{t-i}, $i=0,\ldots,n-2$, where

$$r_{t-i} = \ln(S_{t-i}) - \ln(S_{t-i-1}).$$

For a sample of n historical prices, we obtain (n-1) rates of daily return. The annualized standard deviation of these rates of return is defined as the historical volatility and can be used as an estimate of σ_{t+1}. The nearby historical volatility uses 30 days of data, the middle historical volatility uses 45, and the distant historical volatility has 60 daily prices.

An obvious problem with the historical approach is that it assumes that future volatility will not change and that history will exactly repeat itself. Markets, however, are forward looking and numerous illustrations can be presented to show that historical volatility does not always anticipate future volatility and a better estimate comes from the Black-Scholes option pricing model itself (Choi and Wohar, 1992).

Simply stated, supporters of implied volatility claim that tomorrow's volatility σ_{t+1} can only be estimated during trading tomorrow, i. e., in real time. As option prices are being formed by supply and demand considerations, each trader assesses the asset's volatility prior to making his or her bid or ask prices and, accepting the consensus price of a call as a true market price reflecting the corporate opinions of the trading participants, one solves the Black-Scholes model for the volatility that yields the observed call price. When volatility is calculated in this way, it is called the "implied volatility", with the adjective "implied" referring to the volatility estimate obtained from the Black-Scholes pricing formula. Unlike historical volatility, which is backwards looking to past returns, the implied volatility is forward looking to the stock's future returns from now to the time of the expiration of the option. This implied volatility technique has become the standard method of estimating volatility at the moment of trading.

3. NEURAL NETWORKS FOR PREDICTION

Neural networks are an information processing technology which model mathematical relationships between inputs and outputs. Based on the architecture of the human brain, a set of processing elements or neurons (nodes) are interconnected and organized in layers. These layers of nodes can be structured hierarchically, consisting of an input layer, an output layer, and middle (hidden) layers. Each connection between neurons has a numerical weight associated with it which models the influence of an input cell on an output cell. Positive weights indicate reinforcement; negative weights are associated with inhibition. Connection weights are "learned" by the network through a training process, as examples from a training set are presented repeatedly to the network. Each processing element has an activation level, specified by continuous or discrete values. If the neuron is in the input layer, its activation level is determined in response to input signals it receives from the environment. For cells in the middle or output layers, the activation level is computed as a function of the activation levels on the cells connected to it and the associated connection weights. This function is called the transfer function or activation function and may be a linear discriminant function, i.e., a positive signal is output if the value of this function exceeds a threshold level, and 0 otherwise. It may also be a continuous, nondecreasing function. Feedforward networks map inputs into outputs with signals flowing in one direction only, from the input layer to the output layer.

While there are dozens of network paradigms, the backpropagation 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 institute failure (Salchenberger, Cinar, and Lash, 1992), and estimating option prices (Malliaris and Salchenberger, 1993). The term backpropagation technically refers to the method used to train the network, although it is commonly used to characterize the network architecture. In this learning algorithm, mean squared error and gradient descent are employed to determine a set of weights for the trained network. At each iteration, current weights are updated by minimizing the mean squared differences between the actual response of the system to a given example and the desired response. The nonlinear response functions generate gradients of the error function with respect to the weights and the chain rule is used to determine the appropriate weight changes which propagate back through the layers of the network. For more details of this method, see Rumelhart and McClelland (1986). Currently, a number of variations on this method exist which overcome some of its limitations.

Nonlinear, multilayer, feedforward networks differ from traditional modelling techniques in several ways. Relationships between inputs and outputs are learned during a training process in which the network is repeatedly presented with historical examples. Neural networks possess the ability to approximate arbitrary mappings with no apriori assumptions about the nature of the underlying model required. Also, no assumptions about the distributions of the variables are required and the variables may be highly correlated.

There are several limitations to the use of neural networks for classification and prediction. First, there are no formal methods for determining the optimal network topology for a given task. Although it has been shown that for certain types of networks, one middle layer is sufficient, the number of middle layer nodes can affect network performance and must be determined largely through experimentation. While the number of middle layer nodes can be arbitrarily large, too many middle layer nodes can result in the problem of overfitting and result in a network which lacks the ability to generalize. Other model-building decisions like the choice of transfer function and learning parameters also must be determined experimentally. Unlike other analytical models, neural networks are viewed as black boxes and it is difficult to interpret the significance of individual input variables.

4. 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-money, S&P 100 Index prices, call volume, put volume, call open interest and put open interest were collected from the *Wall Street Journal* for the 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 the predictin periods, with each trading cycle ending the third Friday of the month.

5. A COMPARISON OF HISTORICAL AND IMPLIED VOLATILITY ESTIMATES

The historical and implied volatility for the nearby contract are graphed together in Figure 1 for June 22 through December 30, 1992. As can be observed, the historical estimate significantly underestimates the volatility used by most traders, i.e., 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 which 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.

The average MAD and MSE for the entire forecasting period, from June 22 through Dec. 30 were 0.0331 and 0.0016. The proportion of times which the historical volatility correctly predicted that the implied volatility would increase or decrease are shown in the last column of the table. An overall average of the number of times a change was correctly indicated is .4439, i.e., a little less than half of the time.

Table 1. A Comparision of Historical and Implied Volatilities

Dates of Forecast	MAD	MSE	Proportion of Correct Directions
Jun 22--Jul 19	.0318	.0012	8/19 = .421
Jul 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

6. DEVELOPMENT OF THE NEURAL NETWORKS

To develop a neural network which is capable of generalizing a relationship between inputs and outputs, the training set selected must contain a sufficient number of examples which are representative of the process which is being modelled. 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 from 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 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 new variables which 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 3 to 7 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 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, better predictions were achieved than with past attempts to predict volatility.

The backpropagation 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, and 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.

6. 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 Table 2, where comparisons are made between the volatility forecasted by the network and tomorrow's implied volatility. The overall MAD 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 1 and 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.

Table 2. Neural Network and Implied Volatilities

Dates of Forecast	MAD	MSE	Proportion of Correct Directions
Jun 22--Jul 19	.0148	.0003	16/19 = .842
Jul 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

7. 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 like the 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 to the traders which 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 which 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. Secondly, and more importantly, the network forecasts, in the cases we tested, were very accurate estimates of the volatility preferred by traders.

REFERENCES

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

Choi, J.Y. and K. Shastri, "Bid-Ask Spreads and Volatility Estimates: The Implications for Option Pricing", *Journal of Banking and Finance,* 13 (1989), 207-219, North-Holland.

Choi, S. and M.E. Wohar, "Implied Volatility in Options Markets and Conditional Heteroscedasticity in Stock Markets", The Financial Review, Vol. 27, No. 4, Nov. 1992, pp. 503-530.

Collins, E., S. Ghosh, and C. Scofield, "An Application of a Multiple Neural-Network Learning System to Emulation of Mortgage Underwriting Judgments," in *Proc. of the IEEE International Conference on Neural Networks,* 1988, pp. 459-466.

Dubofsky, D.A., "Volatility Increases Subsequent to NYSE and AMEX Stock Splits", The Journal of Finance, vol. XLVI, no. 1, March 1991, pp. 421-431.

Dutta, S. and S. Shekhar, "Bond-rating: A Non-Conservative Application of Neural Networks," in *Proc. of the IEEE International Conference on Neural Networks,* 1988, pp. 142-150.

Engle, R.F. and M. Rothschild, "Statistical Models for Financial Volatility", *Journal of Econometrics,* 52 (1992), 1-4, North-Holland.

Haugen, R.A., E. Talmor, and W. Torous, "The Effect of Volatility Changes on the Level of Stock Prices and Subsequent Expected Returns", *The Journal Of Finance,* vol. XLVI, No. 3, July 1991, pps. 985-1007.

Lee, C.F., J.E. Finnerty and D.H. Wort, Security Analysis and Portfolio Management, Scott,Foresman and Company, Glenview, Illinois.

Lockwood, L.J. and S.C. Linn, "An Examination of Stock Market Return Volatility During Overnight and Intraday Periods,1964-1989", *The Journal of Finance*, vol. xlv, No. 2, June 1990, pp. 591-601.

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

Malliaris, M.E. and L. Salchenberger, "Beating the Best: A Neural Network Challenges the Black-Scholes Formula", *Proceedings of the Ninth Conference on Artificial Intelligence for Aplications*, Orlando, Florida, March 1-5, 1993.

Merville, L.J. and D.R. Pieptea, "Stock-Price Volatility, Mean-Reverting Diffusion, and Noise", *Journal of Financial Economics*, 24 (1989), 193-214, North-Holland.

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

Salchenberger, L., E.M. Cinar, and N. Lash, "Neural Networks: A New Tool for Predicting Thrift Failures", *Decision Sciences*, 23(4), 1992, pp. 899-916.

Sheikh, A.M., "Stock Splits, Volatility Increases, and Implied Volatilities", *The Journal of Finance*, Vol XLIV, No. 5, Dec. 1989.

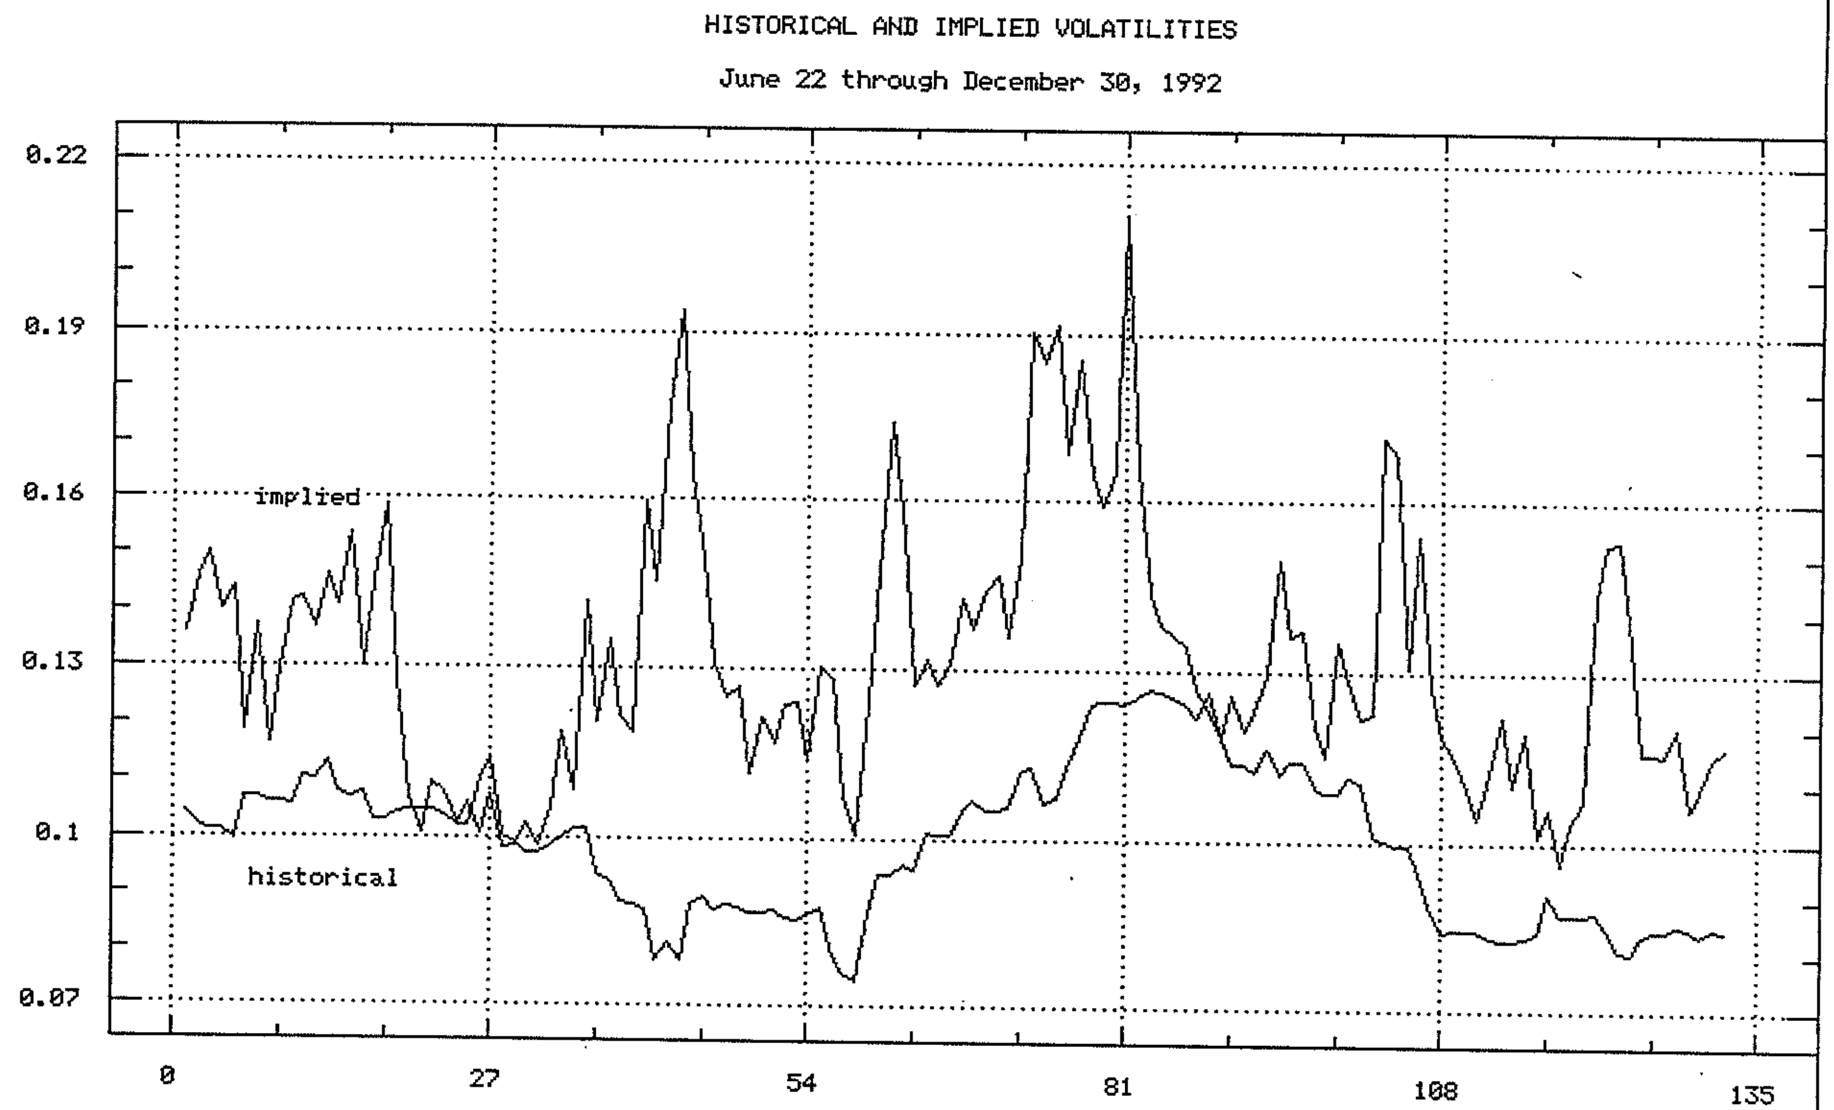
HISTORICAL AND IMPLIED VOLATILITIES
June 22 through December 30, 1992
0.22
0.19
0.16
0.13
0.1
0.07
implied
historical
0
27
54
81
108
135

Figure 1

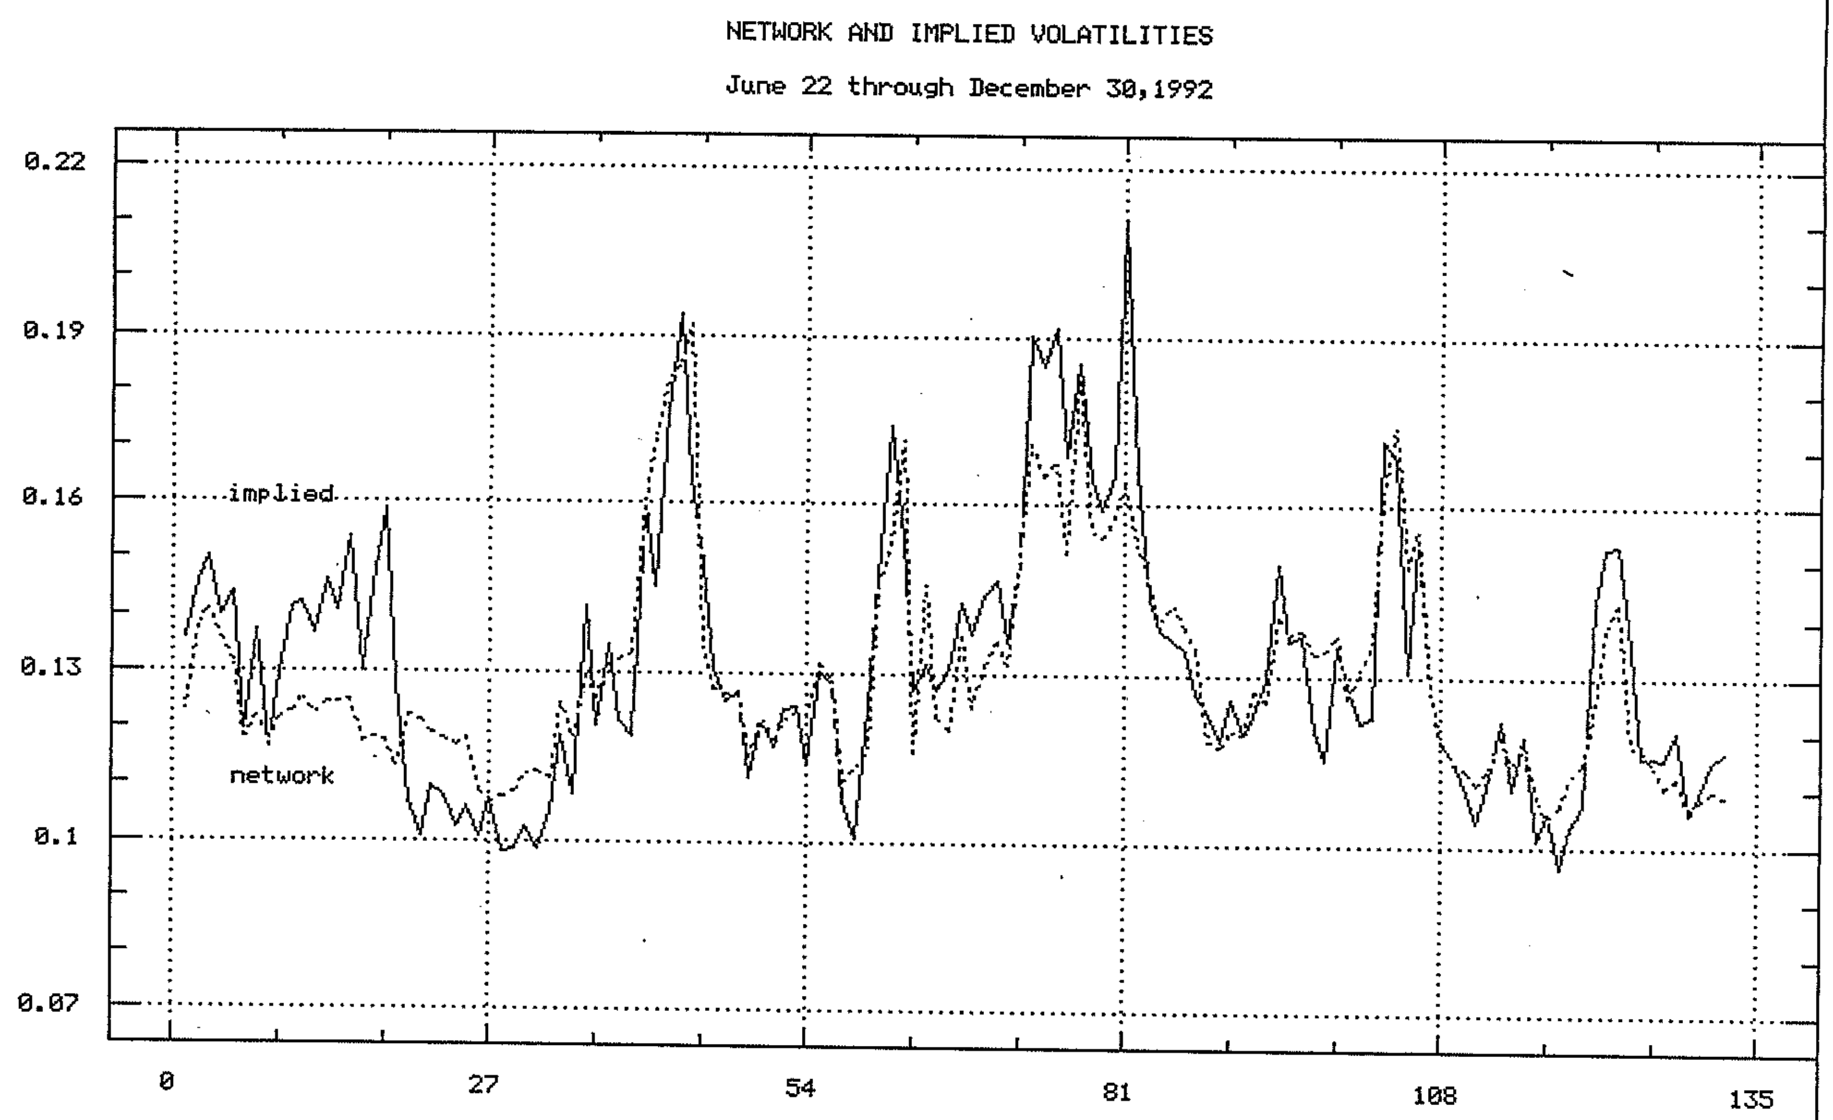
NETWORK AND IMPLIED VOLATILITIES
June 22 through December 30,1992
0.22
0.19
0.16
0.13
0.1
0.07
implied
network
0
27
54
81
108
135

Figure 2