Int. J. of Engineering Intelligent Systems, Vol. 9, No. 4, December 2001, pp. 243-249.

Delta –Neutral Volatility Trading using Neural Networks *◆

Luiz Otavio M. Calôba[1], Luiz Pereira Calôba[2], Cláudio R. Contador[3]
[1]Banco BBM S.A., Rio de Janeiro, Brazil
[2] COPPE / Escola de Engenharia – UFRJ, Rio de Janeiro, Brazil
[3] COPPEAD– UFRJ, Rio de Janeiro, Brazil

E-mails: caloba@bbmbank.com.br, caloba@lps.ufrj.br, contador@coppead.ufrj.br

Abstract

In this paper we propose a technique to forecast the daily changes of the market volatility (ISD) using neural networks. We define the input variables based on specific literature but using lag and horizon of the training set calculated so that we guarantee the significancy of its correlation with the output variable through time. We use this technique on data from Telebrás PN Stock options from August 1994 through November 1996. Then, based on out-of-sample projections of the model we simulate a trading volatility strategy, creating Delta-hedged portfolios that result in abnormal returns in the market.

1. Introduction

Since 1973, when Fisher Black and Myron Scholes (1973) derived a closed formula solution to calculate the price of an option over an underlying asset, the options market suffered an enormous transformation increasing largely through time.

According to Black & Scholes formula the expected value of an European call option (price - c) is affected by five factors: the current asset price (S), the strike price(X), the time to expiration (t), the risk free interest rate (r) and the volatility of the asset price (σ).

$$c = c(S, X, t, \sigma, r) \qquad (1)$$

The sole parameter in Black & Scholes pricing formula that cannot be observed directly on the market is the volatility of the asset price. Nevertheless it is possible to calculate which volatility value makes the Black & Scholes formula option's price equal to the market price. This value is called implied volatility, and it is also known as market volatility.

Operations involving projections about the future behavior of the implied volatility (usually knew as volatility trading) represent an important part of the activities of trading desks of investment banks and hedge funds today. This behavior varies quickly through time and is extremely non-linear, being very difficult to predict using common techniques. This explains why neural networks can be more effective for this prediction.

The main objective of this paper is to try to forecast the behavior of the implied volatility of Telebrás PN options (traded at BOVESPA stock exchange) using a neural network technique. We also try to show that using this prediction it is possible to obtain abnormal returns in the market. That way we will be approaching this problem is on the same line as Harvey and Whaley (1992) and Malliaris and Salchenberger (1994) that use an auto-regressive model (GARCH) and a *feed-forward* neural-network with the same purpose.

This paper is divided into five main sections being this introduction the first of then. On the second section we present the variables used on the forecast of the behavior of the option's implied volatility. The variables were selected through two different stages: first we made a qualitative selection and then we tested these variables and selected which ones were to be used on the analysis and also the lag of each variable and the number of series we used to train the network. We also explain the origin of the data we used on this paper.

On the third section of this paper we present the two-layer feed-forward neural network we use. We explain which training method we applied, what kind of neurons we use on the output and intermediate layer and how we determined the number of neurons on the intermediate layer.

On the fourth section we present the numerical results obtained, explaining at first how the neural network selection is made. Following this, the volatility prediction results are compared with the real ones, discussing the forecasting quality. At last, we present how the trading volatility operations are made and test the capacity of applying the forecasting in order to obtaining abnormal returns in the market, trading costs included.

In the final section, we summarize the work as a whole and present some interesting suggestions and conclusions for further studies.

* This work is based on a Mastering thesis Volatility Trading using Neural Networks (Calôba, L. O. (2000)).
◆ An earlier version of this paper was presented at the V Brazilian Conference on Neural Networks, PUC, Rio de Janeiro, april 2001. We thank conference participants for their comments. We are grateful to Franklin de O. Gonçalves for comments and suggestions. Remaining errors are our own.

2. Input variables selection

In this section, we present the variables that will be used in order to forecast the daily changes in the option's implied volatility. At first, the data set used in this work will be presented. After that, we are going to explain how the variables were selected in a two-stages process. The first stage is a qualitative one where we define which variables are going to be studied. Finally, in a quantitative stage we define the number of series (time horizon) of the training set and the lag for each variable.

2.1 Data set

The raw data used in our analysis are the daily closing prices of the Telebrás options and stock, in a period between August, 1994 and October, 1996. In this period, Telebrás PN was the most liquid stock in the market, being solely responsible for 35% of the trading volume of the Brazilian stock market.

As for stocks, options on Telebrás PN were traded at BOVESPA (São Paulo's Stock market) and had as expiration date the third Monday of February, April, June, August, October and December. Usually only the options of the first expiration date were traded in the market, and more than that, only the call options with 6 or 7 strike prices for each expiration date. Within these, only three had liquidity.

To avoid overlapping of the series and distortions caused from liquidity problems on periods extremely far or close to the expiration date, we normalized the size of the series, analyzing the options only between 34 and 5 working days to the expiration date. This procedure also makes it easier to handle the data.

We also decided to forecast only the behavior of the implied volatility of the most liquid option (the one at-the-money). This choice is also reinforced by the fact that we will verify the possibility of obtaining abnormal returns in the market trading volatility with the forecast results. These options are the most recommended for this kind of operations because it presents the larger Vega (option's sensibility to changes in volatility).

During the analyzed period, there were 13 expiration dates which gives us a total of 390 days to be used in the analysis (30 points for each series). These series are presented in figure 1, below.

2.2 Input variables selection – qualitative stage

In this work we are not trying to make an extension of previously researches in the artificial neural network prediction for options [Hutchinson, Lo and Poggio, 1994; Anders, Korn and Schmitt, 1998; Geigle and Aranson, 1999] or trying to predict the volatility of the underlying asset in as Lachtermacher e Gaspar (1995)[1].

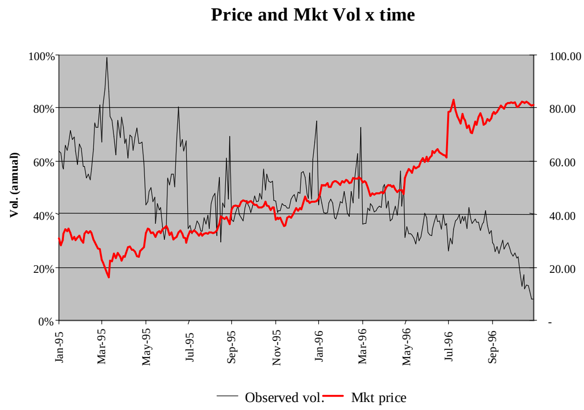

Figure 1: Price and implied volatility x time

We will try to focus this research on the forecasting of the implied volatility behavior, approaching the problem in a similar way as Harvey and Whaley (1992)[2] and Malliaris and Salchenberger (1994)[3]. This way, instead of using only the input values of Black & Scholes (1973) formula as other works, we will test the possibility of using other variables. In order to achieve this, we analyze the literature and define the variables to be analyzed.

At first, the variables involved in Black & Scholes (1973) formula were pre-selected, specifically: the asset value, strike price, time to expiration, risk free rate, implied volatility and the option price.

Besides that, studying other works [Mayhew, 1995, Harvey and Whaley 1992; Malliaris and Salchenberger, 1994] we decided to test other variables: monday (dummy), friday (dummy), stock market's trading volume, changes in the asset price, changes in the implied volatility, Vega and Moneyness function (strike price over asset's future price).

On the following section, we test these variables, verifying it's efficiency in the forecasting of the implied volatility.

[1] This was the only reference found using neural networks in the Brazilian derivatives market. In this work, the authors successfully developed a neural network with the capacity to calculate the market value of Telebrás options using Black & Scholes (1973) inputs.

[2] In this work, the authors approach the problem using na autoregressive model (GARCH) rejecting the hypothesis of non-predictability of the market volatility of S&P 100. However, they also reject the possibility of obtaining abnormal returns with these predictions.

[3] In this work, the authors developed a neural network with the intent to predict the value of implied volatility on the following day of the at-the-money option. They used trial and error process to select input variables, arriving in 13 final variables. The correlation coefficient between the real and predicted volatility (R2) was of 0,85. The average of direction changes correctly forecasted was 79,4%.

2.3 Input variables selection – quantitative stage

As the implied volatility behavior is a very noisy and time changing process, we verify that one of the largest problems to face was to define the optimum number of training sets (time horizon) where this process could be treated as a pseudo-stable process[4].

Being more specific there were 13 series that could be used on the neural network training. We use the N first series to train the neural network and then we use this network in the forecast of the volatility behavior of the next series (N+1). After this we use the second N (from 2 to N+1) series to train another network that we use on the following series (N+2) and so on. We need to define the training horizon (N) in a way to maximize the efficiency of the network prediction.

If a long horizon is used, we won't be able to train, for the system time variability is very large so the behavior of the last series will be much different of the first one. On the other hand, if the horizon is short, we may not be able to map the problem correctly (the network will not be trained correctly due to lack of data).

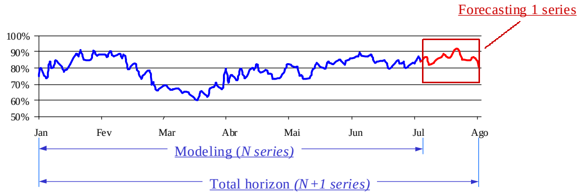

Figure 2: training horizon

Another important issue is the lag to be used on each input variable. To answer these questions, we elaborated a method to verify the significance of the correlations between input and output variables, as a function of each variable lag and the series horizon applied.

The first step of this method was to calculate the lagged correlations between all input variables and the return of implied volatility (we also calculated the lagged autocorrelation of this variable), as we show in figure 3. As a result of this analysis, we obtain 13 matrixes of correlations (one for each series), with 17 columns (number of variables) and 30 lines (number of possible lags).

After that, we selected the matrixes of correlation for the N first series, calculating the average and standard deviation for each cell of the matrix. Following this, we tested the significancy of this average correlation (cell) through a hypothesis test, with 95% probability (1.65 standard deviations) that the correlation won't change its signal. In case the test doesn't fail, we attribute the value of 1 to the cell; in case it fails, we attribute 0 to the cell value. Doing this to all cells we have as a result a matrix of equivalent dimensions as the correlation matrix, but only with binary values (0-1). This matrix will be denominated as significancy matrix.

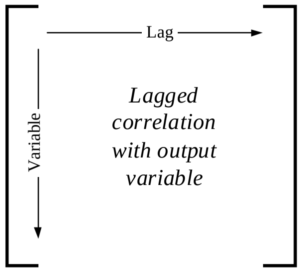

Figure 3: Series correlation matrix.

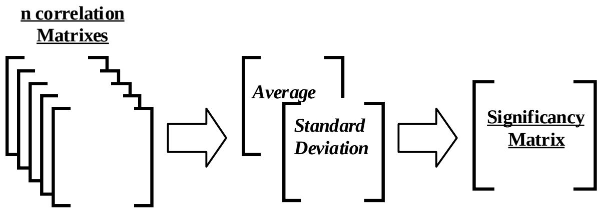

Figure 4: Significancy Matrix.

The same process will be repeated to the correaltion matrixes between the series 2 and N+1 and so on and so forth, until the series 13-N and N. As a result, we will have 13-N resultant matrixes.

Then, cell by cell, the average is calculated and multiplicated by the value of the significancy probability used (95%). Doing this, we generate an average significancy matrix for a *total horizon* N+1.

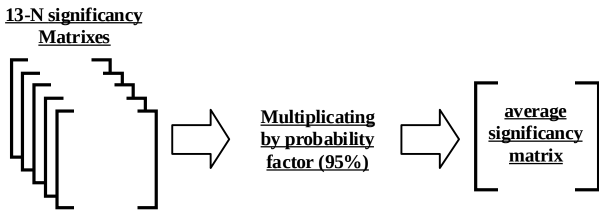

Figure 5: Average Significancy Matrix.

This process is repeated for all possible horizons (from 2 to 13), having as a final product 12 average significancy matrixes, one for each total horizon.

Then we separate the lines related to the i*th* variable in each average significancy matrix, generating a single average significancy matrix for the variable i, composed by 12 lines (number of total horizons) and 30 columns (number of possible lags). In that way, each point will show the average medium significancy of the correlation between the lagged variable and the output variable (implied volatility) (with a j lag), for a total horizon N+1.

[4] A pseudo-stable process is a process where the variables statistic parameters and its inter relations changes through time but are locally constants.

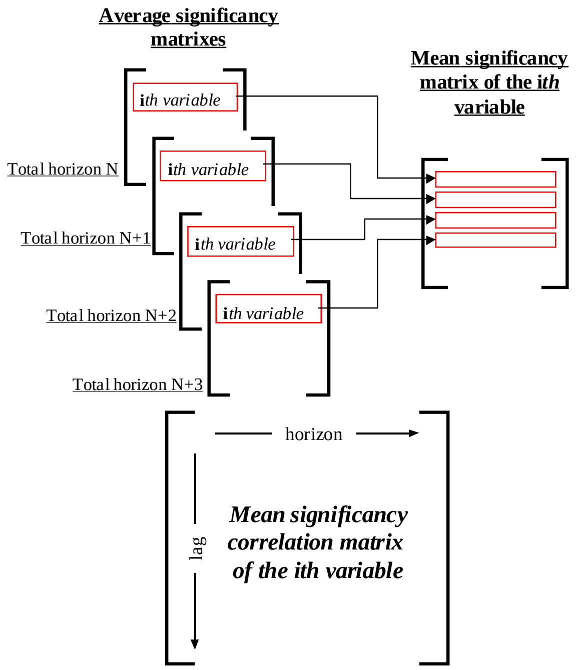

Figure 6: Mean significancy correlation matrix of the i*th* variable

Then, analyzing the results, we perceived that as the horizon is increased, the less significant the averages are (the correlations change signal much more times within the horizon). We also noticed that from a total horizon (N+1) equal to 3 the number of significant input variables is small and as the total reaches 4 there are barely no variables with stable correlations. As a conclusion of this result, we adopted a training horizon of 2 series in order to forecast a third one. So we'll have a total of 11 sets of 3 series each (2 for training/testing and 1 for the forecasting).

In order to choose the input variables we eliminated all those with less than 50% of medium significancy. Being clear, we cut off all variables where the chance of changing the correlation signal between series is larger than 50%, assuring a more stable training process.

Then we select the variables with larger significancy and lag smaller than 7 days (the larger the lag, the less data we'll have to train and trade). In the end, we get to the eight following variables:

Table 1: Lag and significancy of input variables.

Variable	Lag	Significancy
Implied vol. Daily changes	1	84%
Moneyness Function	5	76%
Asset value	2	76%
Strike price	0	67%
Implied volatility	2	67%
Option's price	1	59%
Time to expiration	0	59%
Monday *(Dummy)*	0	59%

We get then to the part of definition, training and forecasting through the neural network.

3. The Neural Network

Artificial neural networks are powerful tools essentially for multivariate non-linear process [e.g. Wasserman, 1989; Hecht-Nielsen, R., 1990; Cichocki and Unbehauen, 1993; Calôba, 1997; Haykin, 1999]. In this work, we used a feed-forward neural network with 2 layers of n neurons. This network has as an advantage its easy handling and the fact that it can approach, theoretically, any function [Hecht-Nielsen, R., 1990]. Similar networks were used by Hutchinson, Lo and Poggio, 1994, Anders, Korn and Schmitt,1998 and Geigle and Aranson, 1999.

The methodology of training used in this work was back-propagation, using an epoch training method, developed by Silva and Almeida (described by Cichocki and Unbehauen, 1993), modified to avoid overtraining problems. Specifically, after the net training, we select the synapse set that presented, throughout the training, the smaller error on the test set.

The network used on the forecasting has as inputs the eight variables described in the previous section and as output the change in the volatility. The neurons on the intermediary layer (hidden layer) have a sigmoid type activation function, while the output neuron has a linear activation function.

In order to define the number of neurons in the hidden layer, we studied the behavior of the validation error (out-of-sample) as a function of the number of neurons. This study was made using the mean square error and also as a function of the direction error (percentage of signal error on the change of volatility). To avoid problems in this analysis caused by a bad training (adverse starting points), we generated initially 10 networks for each number of neurons / series data set. We selected then, in this set of 10 networks, the one with smaller test error, using the hypothesis that this is the best available indicator for the series with smaller validation error.

As a continuation, we calculated the arithmetic average of the training, test and validation errors for the 11 sets of series (Square Average Error and Direction Error). The process is repeated for other amounts of neuron in the hidden layer.

As a result, we generated the two graphics below, which show the behavior curve of the average mean square error and the average direction error as a function of the number of neurons in the hidden layer:

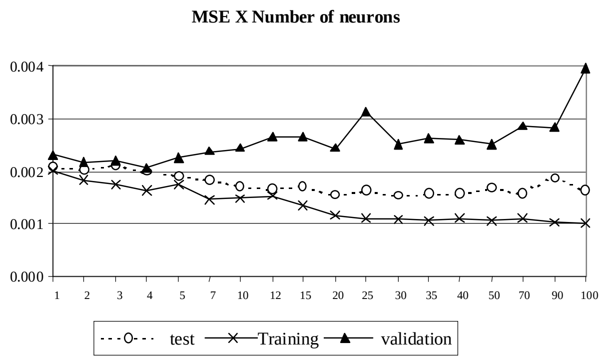

Figure 7: Average of the MSE observed as a function of the number of neurons in the hidden layer.

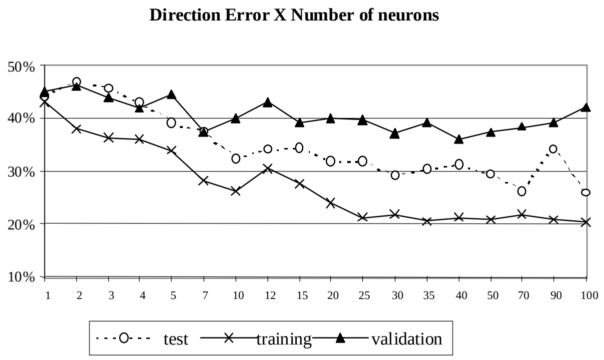

Figure 8: Average of the Direction Error observed as a function of the number of neurons in the hidden layer.

The network with better performance considering the MSE of the validation set, as seen in the graphics, is the one with 4 neurons in the hidden layer. On the other hand, the network with better performance considering the direction error of the validation set, as seen in the graphics, is the one with 40 neurons in the hidden layer. Since we want to trade volatility with this forecasting, the numerical precision of this forecasting is more important than the direction itself. So, we decided to use the network with 4 neurons in the hidden layer.

4. Numerical Results

Once the neural network was defined, we trained 50 networks for each one of the 11 series. We selected then, from the 50 networks, the one with the least accumulated test error. Following we have an example of the prediction (validation) accumulated error of the 50 series with the chosen one highlighted. As we can see, the chosen method is efficient (the testing error is a good indicator for the prediction error). Finally, through these nets we then made the forecasting of the change in volatility.

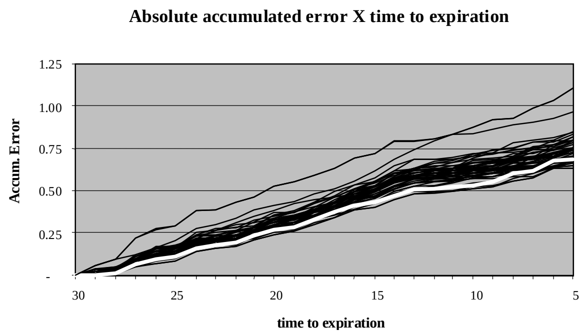

Figure 9: Absolute accumulated forecasting error of the selected training net (white) compared to the others.

4.1 Forecasting the change in volatility

After the selection of the best networks, we calculated from the observed variables on a trading day (D0) the value of change predicted for implied volatility between D0 and D+1 (the following day). Then, we calculated the expected implied volatility for D+1. After that, we compared this prediction with the real value. The graphic below shows an evolution of these values:

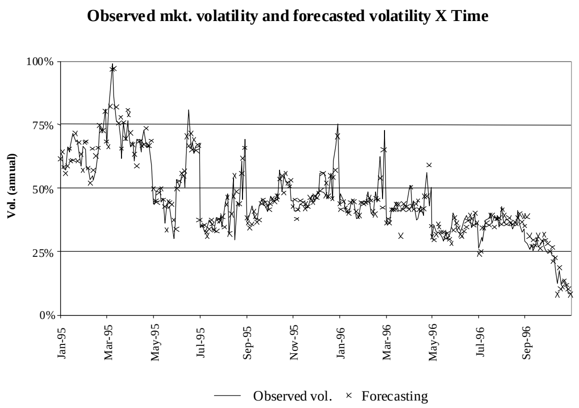

Figure 10: Forecasting of implied volatility against market volatility (out-of-sample).

As we can see, the forecasting, except for some points, is very accurate. Table 2 exhibits the comparison statistics between the forecasted and real values. We made two different analysis, one with all the series and the other excluding the period of strongest influence of Mexican Crisis.

Comparing the results with the ones obtained by Malliaris and Salchenberger (1994), we obtained a larger correlation coefficient between the true and the predicted volatility (R2) (0,90 against 0,85). On the other hand, the average direction error was smaller (58% against 79%). These comparisons are of small significance for the difference of forecasting environments. Being more specific, in our case we are trying to predict the behavior for a derivative on an emergent market (Brazil), while Malliaris and Salchenberger (1994) deal with a developed economy.

Table 2: Forecasting data statistical analysis.

	11 Series	10 Series[5]
R2	0.8996	0.9035
Direction Error	41.82%	42.00%
Absolute Mean Error	3.42%	3.34%
Square Mean Error	0.0023	0.0023

We also calculated the behavior of the mean absolute error as a function of time to expiration. It was verified that the error increases as we approach to the expiration date, which indicates the possibility to use networks for smaller periods (for example, to train the network at 12 days or less, instead of 25) to improve results.

4.2 Trading volatility

Once we select the neural net we started using it to forecast the implied volatility of the next trading day. With these results we decided which operation we would do (buying or selling volatility).

To do so, we take the market value (c) of the option we will trade and calculate its implied volatility and delta. Them we make a *delta-neutral* portfolio that is compounded by option, stock (on a quantity that makes the portfolio delta equals to zero) and cash ($) that comes from the buying of the option and selling of the stocks:

$$portfolio_0 = Option(S, K, t, r, \sigma_{implied}) - \Delta \times S + \$(\Delta \times S - c) \quad (2)$$

This portfolio has a null value when we make it. Then we calculate the expected portfolio value on the following day (D+1), taking into account the forecasted volatility, the expected stock value (today's value plus risk free interest) and the cost of the position (interest rates).

$$portfolio_1 = Option(S, K, t-1, r, \sigma_{forecasted}) \\ - \Delta \times S + \$(\Delta \times S - C) \times (interest\ rate) \quad (3)$$

If the expected portfolio value of the next day is positive we buy the portfolio as it is described above (buying volatility). On the other hand, if the value is negative we do the opposite (selling volatility).

In the following graph we show the daily trading profit evolution of the 11 series on the analyzed period for a portfolio settled using the technique described above. The total exposure of the portfolio (cash) is 10 millions US Dollars (positive or negative). This graph does not take into account the trading costs.

Trading series profits x time to expiration

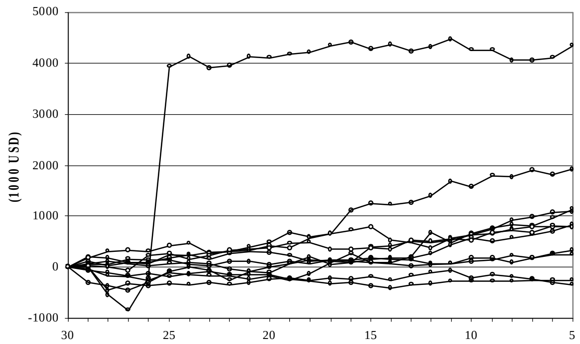

Figure 11: Trading series profit evolution as function of time to expiration. The detached series is located on the middle of the Mexican Crisis.

On the following table we show the total trading profit for each one of the series generated buying and selling volatility according with the network's forecast. On italic we show the Mexican Crisis period.

Table 3: Begin and end dates of the trading series and its profits with and without trading costs.

Begin Date	End Date	Trad. profits without costs (1000 USD)	Trad. profits with costs[6] (1000 USD)
Jan-1995	Feb-1995	1092	1030
Mar-1995	*Apr-1995*	*4336*	*4291*
May-1995	Jun-1995	1929	1868
Jul-1995	Aug-1995	1138	1109
Sep-1995	Oct-1995	828	800
Nov-1995	Dec-1995	781	714
Jan-1996	Feb-1996	-345	-374
Mar-1996	Apr-1996	799	749
May-1996	Jun-1996	322	266
Jul-1996	Aug-1996	246	185
Sep-1996	Oct-1996	-268	-302

We calculated the statistical summary of these results and showed then in the following table. We show that the medium daily profit is significantly positive, even if we take into account the operating costs. If we exclude the Mexican crisis period (series between March-1995 and April-1995) during which we had the largest profit we still have a positive result and

[5] Excluding the series between march-1995 and april-1995 where the Mexican crisis effects were concentrated.

[6] The trading costs at BOVESPA are of 1% of the trading volume. On the specific case of a financial institution trading with its own broker firm there are cost devolution between 99% and 95% of this value. In this paper we use 97% of devolution.

(this is more interesting) the confidence coefficient enlarges.

Table 4: Statistical analysis of trading profits.

	11 Series		10 Series[7]	
	without costs	with costs	Without costs	with costs
Total profit (1000 USD)	10858	10335	6522	6034
Daily profit Standart deviation	275	275	104	104
Med. daily profit (1000 USD)	39	38	26	24
Med. daily profit Standart deviation	17	17	7	7
t test	2.38	2.27	3.98	3.68

5. Conclusions

This paper was the first trial of applying neural networks on implied volatility prediction at an emerging market (Brazil). We face some difficulties comparing with similar works at stock markets at developed countries.

At Brazil the liquidity of the options is extremely concentrated on the first expiration date and only at the calls. Therefore only a small number of options have enough liquidity to be traded (3 or 4 per expiration). Even with this obstacle we believe we reached some promising results.

The forecasted volatility presents a larger correlation with the market volatility but a worst direction error than the work of Malliaris and Salchenberger (1994). We believe this larger error comes from the complexity and the instability of the system we are trying to predict (the derivative market of an emerging country). Another evidence of that is the fact that we could use only two trading series on the training of the network.

We perceived that as closer we get to the option's expiration date, the forecasting error increases. These results show the possibility of trying to improve the forecast quality reducing the training set intervals, making one net for each week or day of prevision, for instance.

We also show the possibility of obtaining abnormal returns trading volatility during the analyzed period based on the neural networks forecasting results, even excluding the Mexican crisis period (series between March-1995 and April-1995) where we obtain the largest profit.

[7] Excluding the series between march-1995 and april-1995 (Mexican crisis period) where we found an abnormal profit.

As an extension to our work, beyond reducing the training interval data set of the networks, we believe it would be interesting trying to forecast the behavior of all the options with liquidity so we could have a forecast for the smile behavior which opens possibility of making *delta-gamma-neutral* trading. Another possibility is trying to use intraday data and making intraday predictions that, we believe, would take us to better results and perhaps to largest profits.

References

[1] Anders, U, Korn, O., Schmitt, C. Improving the Pricing of Options: A Neural Network Approach. *Journal of Forecasting*, v. 17, p. 369-388, 1998.

[2] Black, F., Scholes, M. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, v. 81, n. 3, p. 637-654, May/June 1973.

[3] Calôba, L. O., Previsão de volatilidade implícita utilizando Redes Neurais, Mastering thesis, Rio de Janeiro: UFRJ/COPPEAD, 2000, Original text.

[4] Cichocki, A.; Unbehauen, R. Neural Networks for Optimization and Signal Processing. Stuttgart: J. Willey, 1993.

[5] Geigle, D. S., Aronson, J. E. An Artificial Neural Network Approach to the Valuation of Options and Forecasting of Volatility. *The Journal of Computacional Inteligence in Finance*, v. 7, n. 6, p. 19-25, 1999.

[6] Harvey, C. R., Whaley, R. E. Market Volatility Prediction and the Efficiency of the S&P 100 Index Option Market. *Journal of Financial Economics*, v. 31, n. 1, p. 43-73, 1992.

[7] Haykin, S. Neural networks: a Comprehensive Foundation. 2.ed. S.l.: Prentice Hall, 1999.

[8] Hecht-Nielsen, R. Neurocomputing. Massachusetts: Addinson-Wesley, 1990.

[9]Hutchinson, J. M., Lo, A.W., Poggio, T. A Nonparametric Approach to pricing and Hedging Derivative Securities via Learning Networks. *The Journal of Finance*, v. 49, n. 3, p. 851-889, July 1994.

[10]Lachtermacher, G., Gaspar, L. A. R. Neural Networks in Derivatives Securities Pricing Forecasting in Brazilian Capital Markets. In: *Proceedings in the Third International Conference on Neural Networks in the Capital Market*. New Jersey: World Scientific, 1995. p. 92-97.

[11]Malliaris, M., Salchenberger, L. Neural Networks for Predicting Options Volatility In: Proceedings of the 1994 WCNN. San Diego: s.n., 1994. p. II-290-295;

[12]Maythew, S. Implied Volatility. *Financial Analysts Journal*; p. 8-20; July/Aug. 1995.

[13]Wasserman, P. D. Neural Computing, Theory and Pratice. New York: Van Nostrand Reinhold, 1989