Computational Intelligence Approach to Capturing the Implied Volatility

Fahed Mostafa[1], Tharam Dillon[1(✉)], and Elizabeth Chang[2]

[1] Deptartment of Computer Science and IT, La Trobe University, Melbourne, Australia
fahedm@yahoo.com, Tharam.dillon7@gmail.com
[2] School of Business, ADFA and UNSW at Canberra, Canberra, Australia
E.Chang@adfa.edu.au

Abstract. In this paper, a Computational Intelligence Approach and more particularly a neural network is used to learn from data on the Black-Scholes implied volatility. The implied volatility forecasts, generated from the Neural Net, are converted to option price using the Black-Scholes formula. The neural network option pricing capabilities are shown to be superior to the Black-Scholes and the GARCH option-pricing model. The neural network has also shown that it is able to reproduce the implied volatility well into the future whereas the GARCH option-pricing model shows deterioration in the implied volatility with time.

Keywords: Computational intelligence · Neural networks · Option pricing · Implied volatility · GARCH option pricing model

1 Introduction

The Black-Scholes Option Pricing Model (BSOPM) published in 1973 (Black and Scholes 1973) remains one of the most cited scientific papers in this field. The model made a key contribution to option trading, allowing investors to calculate a fair value for an option contract. This model has its limitations, which stem from the unrealistic assumptions used in the derivation of the BSOPM. The wide success of this model was due to its ability to produce an option price instantaneously. The formulation of BSOPM was based on major assumptions which allowed for the derivation of this model. These assumptions significantly impacted the model pricing capabilities which are visible in the models pricing biases. The constant volatility assumption has captured the attention of researchers, as it contradicts the volatility behaviour observed in the underlying asset. Also, it was realised when the volatility was backed out from the BSOPM for different strike prices, the implied volatility displays a U-shape pattern (volatility smile) when plotted against the options moneyness. Moneyness is defined as the relative position of the current or future price of the underlying asset with respect to the option strike price. Nevertheless, implied volatility is derived directly from the options data, which contains vital information as it provides a forward looking view on how the market is anticipating movements in the underlying asset. This is a different view to the conditional volatility calculated from the stock returns which is

a retrospective view. The other interesting feature of the implied volatility is the deterioration of the volatility surface through time caused by the changes in the implied volatility levels. Many models have emerged to overcome the limitations of the BSOPM. In the present paper, we select the GARCH option pricing model (GOPM hereafter) as it allows for the volatility to follow a GARCH specification. This specification overcomes the constant volatility assumptions and at the same time caters for key stylised facts in the underlying returns series. This model has displayed superior improvements over other pricing models. The major shortcoming for this model is the absence of a closed form solution. This means the model parameters have to be optimised using a Monte Carlo simulation, which is a computationally intensive process that limits its ability to be applied to real-life problems.

Neural networks have been applied previously to option pricing problems (Karaali, Edelberg et al. 1997), (Meissner and Kawano 2001), (Amilon 2003), (Yao, Li et al. 2000), (Tino, Schittenkopf et al. 2001) and (Bennell and Sutcliffe 2004). The majority of research papers rely on neural networks to provide a mapping from inputs similar to the BSOPM to the option price. Training the neural network on the option price is a cumbersome task. The option itself can suffer from many issues such as stale prices due to infrequent trading (Rubinstein 1985). Therefore, using the option price to guide the training of the neural network can be misleading as the neural network will attempt to map the inputs to an output that has poor correlation with the input. To overcome this limitation, many variations of this method have been developed to improve the pricing accuracy. For instance, substantial pricing performance gains are achieved when the training data set is partitioned by moneyness and different neural networks are used for each of the training sets. Yet there has been no attempt to evaluate the neural networks' performance against more advanced models. The reason for this is that the pricing capabilities of the neural networks are inferior to those of more advanced pricing models such as the GOPM. In (Hanke 1997), a neural network was constructed to give an approximate price to GOPM. This was achieved by training the network on different input combinations for the GOPM allowing the neural networks to approximate the GOPM, thereby overcoming the numerical simulation issues of the GOPM. In (Mostafa and Dillon 2008), a new option pricing method was proposed by allowing the neural network to learn the implied volatility behaviour of the option through time.

The main variable that can explain the option behaviour through time is the implied volatility, since the implied volatility is considered to be a forward looking view of the market anticipation. By allowing the neural network to capture the implied volatility surface through time, the neural network will be able to forecast future movements of the implied volatility across different strikes which can then be converted to an option price. This method was demonstrated to have superior option pricing over the Black-Scholes and GOPM. The option pricing of the neural network and GOPM was discussed and analysed in (Mostafa and Dillon 2008), in this paper we demonstrate the capabilities of neural networks in learning the implied volatility surface through time and its impact on option pricing. We also validate this method by studying how well the neural network can reproduce implied volatility into the future relative to the GOPM. In section 2, the option pricing models are reviewed; in section 3 the data used in this paper is explained. In section 4 the performance measuring criteria are explained, in section 5 the results from the experiments are discussed. In

section 6 we discuss and compare the implied volatility forecast behaviour of the neural network and GOPM. We then conclude this paper in section 7.

2   Option Pricing Models

In this section we review the option pricing models studied in this paper.

2.1   GARCH Option Pricing Model (GOPM)

It has been demonstrated by (Duan 1995) that under the Local Risk Neutral Valuation Relationship (LRNVR) the conditional variance does not change. However, under measure Q , the conditional expectation of rt is the risk free rate r_f.

$$E^Q\left[\exp(r_t)\mid\Omega_{t-1}\right]=\exp(r_f)$$
(1)

To derive the GOPM, the risk neutral valuation relationship has to be generalized to the LRNVR:

$$r_t = r_f - \frac{1}{2}\sigma_t^{*2} + \sigma_t^*\varepsilon_t^*, where \ \ \varepsilon_t^* \sim N(0,1)$$
(2)

$$\sigma_t^{*2} = \beta_0 + \beta_1(\varepsilon_{t-1}^* - \tilde{\gamma})^2\sigma_{t-1}^{*2} + \beta_2\sigma_{t-1}^{*2}$$
(3)

By having $\tilde{\gamma} = \lambda + \gamma$, the risk-neutral pricing measure is determined by four parameters: $\beta_0, \beta_1, \beta_2$ and $\tilde{\gamma}$. Using the above formulation, the asset terminal is then calculated at time T

$$S_T = S_t \exp\left(r_f(T-t) - \frac{1}{2}\sum_{i=t+1}^{T}\sigma_i^{*2} + \sum_{i=t+1}^{T}\sigma_i^*\varepsilon_i^*\right)$$
(4)

The terminal asset price is then calculated using Monte Carlo simulation. A set of N random path of residuals $(\varepsilon_{t+1,j}^*,...,\varepsilon_{T,j}^*)$ is generated with J = 1 to *N*. The residuals are used to calculate the asset prices $S_{T,j}$. The final option price is then approximated as follows:

$$C_{GARCH} = \exp(-r_f^*(T-t))\frac{1}{N}\sum_{j=1}^{M}\max(S_{T,j} - K, 0)$$
(5)

Empirical Martingale Simulation (EMS) of (Duan and Simonato 1998) has been used which has been shown to accelerate the convergence of the Monte Carlo prices estimates. The GOPM parameters were estimated for each day of the data sample. This was achieved by minimizing the average sum squared error eqn. (21) (Aboura 2005) (Lehnert) over all options on day t, with parameters from day t-1 as the initial values (Lehar, Scheicher et al. 2002).

$$\text{Min } SSE(\Theta) = \sum_{t=1}^{T} \sum_{n=1}^{N} \left(\frac{\left(\hat{C}_{t,n} - C_{t,n} \right)}{C_{t,n}} \right) \tag{6}$$

where $\Theta = \{\beta_0, \beta_1, \beta_2, \lambda\}$, $\hat{C}$ and C are the model theoretical price and the actual call price respectively. The number of Monte Carlo simulations was initially set to 15,000. This was sufficient to generate stable estimates. In some cases, the number of simulations had to be increased to 30,000.

2.2   The Black-Scholes Option Pricing Model (BSOPM)

The BSOPM theorem was first published in 1971 (Black and Scholes 1973). It is the most widely used pricing model. The interest rate and volatility are constant; the call option on the asset, expiring at time T and with strike price K will have the following value at time t:

$$C(t,T) = S(t)N(d_1) - Ke^{-r(T-t)}N(d_2) \tag{7}$$

where, $d_1 = \dfrac{\ln\left(\dfrac{S(t)}{K}\right) + \left(r + \dfrac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}$ and $d_2 = \dfrac{\ln\left(\dfrac{S(t)}{K}\right) + \left(r - \dfrac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}$

$N(x)$ is the cumulative probability distribution for a standard normally distributed variable. The BSOPM delta is given by $N(d_1)$, which also measures the sensitivity to the underlying asset.

2.2.1 Implied Volatility

The implied volatility is a volatility parameter, which equates the market price with the price given by the BSOPM formula. The implied volatility, $\sigma_t^{BS}(K,T)$ is a function of K (Strike) and T (time to maturity) (Hull 2003). The two most interesting features of the volatility surface, which have been studied and analysed by researchers, are the volatility smile (skew) and the term structure or the level of implied volatility changes with time. The volatility smile is a key indicator of an unrealistic assumption of constant volatility. Whereas, the changes in the implied volatility level with time is seen by the deformation of the volatility surface with time. Therefore, the ability to capture the implied volatility term structure will lead to accurate option pricing. (Cont and da Fonseca 2001) have expressed mathematically the implied volatility surface as follows:

$$C_{BS}\left(S_t, K, \tau, \sigma_t^{BS}(K,T)\right) = C_t^*(K,T), \text{ where } \sigma_t^{BS}(K,T) > 0 \tag{8}$$

The value of the call option as a function of the volatility is a monotonic mapping from $[0,+\infty[$ to $]0, S_t\text{-}Ke^{-rt}[$. The implied volatility $\sigma_t^{BS}(K,T)$ of a call option with strike K and price maturity of T is dependent on K and T. If K and T are fixed, $\sigma_t^{BS}(K,T)$ can

be generalised to follow a stochastic process. For a fixed t, the value will depend on the options characteristics such as maturity and strike level *K*. Equation 9 represents the volatility surface at time t.

$$\sigma_t^{BS} : (K,T) \longrightarrow \sigma_t^{BS}(K,T) \tag{9}$$

$$I_t(m,\tau) = \sigma_t^{BS}(mS(t),t+\tau) \tag{10}$$

Where *m* is the moneyness, $I_t(m,\tau)$ is the implied volatility function.

The two most important features of this surface are the volatility smile and the changes in the volatility levels with time. Thus, the evolution in time of this surface will reflect the evolution of market option prices. Many previous studies have attempted to explain the information contained in the implied volatility (Day and Lewis 1992; Canina and Figlewski 1993; Christensen and Prabhala 1998; Blair, Poon et al. 2001; Jiang and Tian 2005) where the majority of studies have shown that the implied volatilities contain relevant measurement errors with quantifiable magnitude. Earlier research shows the implied volatility to be an inefficient forecast of volatility (Canina and Figlewski 1993). However, more recent studies such as (Jiang and Tian 2005) show that the implied volatility subsumes all information contained in the historical volatility and is capable of providing an accurate forecast of future volatility. This information can be utilised in many aspects since the option prices reflect the participant's expectations of the market's future movements. So, if the option market is efficient such that all information is observed and the correct option pricing model is specified, the implied volatility should subsume the information contained in other variables in explaining future volatility. That is, the implied volatility should be an efficient forecast of future trends over the life of the option.

There have been several theoretical explanations for the volatility smile such as the distribution assumption and stochastic volatility (Poon 2005). There have been other explanations proposed based on market microstructure, measurement errors and investor risk preference. However, the BSOPM assumes the asset price to have a lognormal distribution or the (logarithmic) returns to follow a normal distribution. Given a high strike price the call option is deep out-of-the-money and the option has a very low probability that it will be exercised. The leptokurtic right tail will give the option a higher probability than the normal distribution for the asset price to exceed the strike price, which means the call will be in-the-money at expiry. If we consider a low strike price, the same argument applies as above but an out-of-the-money put option is used instead. Due to the thicker left tail of the leptokurtic distribution, the probability of the out-of-the-money put option to finish in-the-money will be higher than that of the normal distribution. Therefore, the put option price should be greater than the price given by the BSOPM. If the BSOPM is used to back out the implied volatility, the BSOPM implied volatility will be higher than the actual volatility. This causes the volatility smile where the implied volatility is much higher at very low strike prices and low at high strike prices. The implied volatility for different strikes and maturities do not evolve independently and they are a highly correlated multivariate system (Cont and Fonseca 2002). To construct a model in terms of the implied volatility rather than volatility of the underlying asset volatility may complicate

the modelling procedure. However, there are advantages of modelling the implied volatility directly. For instance, the implied volatility is observable and is derived from the market data, whereas, the asset volatility is not directly observable. The implied volatility gives an insight to the option markets which can be analysed by the practitioners. The shifts in levels of implied volatility are highly correlated across strikes and maturities, which allows for the modelling of the joint dynamics. The implied volatility is now widely used by practitioners and especially with the emergence of implied volatility indexes and derivative instruments (Cont and Fonseca 2002).

2.3   Artificial Neural Network (ANN)

The ANN consists mainly of an input layer, one more hidden layer and an output layer. The layers are connected via a set of weights. The hidden layer and the output layer consist of individual neurons. The inputs are multiplied by the weights and a bias term is added which then constitutes the input to the activation function. This then serves as the input to the following layer. The activation of the output layer is given by:

$$F(x) = G(\sum_J w_{ij} H(\sum_K w_{jk} x_k + B_k) + B_J)$$ (11)

The activation functions of the neurons could be chosen to be linear or non-linear functions. A sum of error-squared function is normally used as the objective function for the training of the MLP. So the MLP is trained to minimize this function with respect to the in-sample data.

$$E = \frac{1}{N} \sum_1^N (t_i - F_i(x))^2$$ (12)

The MLP performance is dependent on the initial values of the weights. To overcome this issue, the neural network is trained 50 times using different initial values for the weights. The weight set that introduces the least error is then adopted.

Given the limited success in applying neural networks to option pricing, we adopt the same method used in (Mostafa and Dillon 2008) for training the neural network with the intention of capturing the underlying asset dynamics of the instrument which then can be translated to an option price using known market variables. This new methodology allows the neural network to capture the dynamics of the underlying asset and the changes in the volatility levels through time. Using the option price to train the neural network has been proven to be ineffective; therefore, instead of training the neural network to price the option directly, we allow the neural network to learn the implied volatility. So the neural network will be optimized on the implied volatility, thereby allowing the neural network to forecast the implied volatility for different maturities and moneyness. Many studies have demonstrated that vital information is contained in the implied volatility especially since the implied volatility contains the market expectation of the underlying asset future movements. To allow the neural network to capture the entire volatility surface, the data set was not parti-

tioned by moneyness, thereby allowing the neural network full visibility of the volatility surface. This approach is a mapping exercise described in equation 14. The neural networks must be tested and verified on the training data set to overcome key issues such as over-fitting (Mostafa and Dillon 2008). The neural network is trained using the BSOPM implied volatility as the target output. This method allows the neural network to predict the future implied volatility, which also means that the neural network can produce the implied volatility surface over any trading horizon. The implied volatility can be converted to an option price by plugging it into the BSOPM. The neural networks were optimized by varying the number of inputs, length of in-sample and validation sets, and the number of hidden units. The combination that gives the lowest error is considered to be the optimal model and is chosen accordingly. The initial weight values were initialized randomly, and to eliminate the dependency on the initial weight values, this process was repeated 50 times.

After experimentation with the different inputs, the following combination was selected for optimal performance. The in-sample and validation set were 20 and 5 days respectively, and the output sample was 1 day. The inputs used were moneyness (index/strike), time to expiry and historical volatility. The target data set was the BSOPM implied volatility. Once the training of the ANN is completed, the implied volatility produced by the ANN would be plugged into the Black-Scholes formula to obtain an analytical price of the option. This model would be referred to as NNiv.

3   Data

The data used in this research consists of European call options on the FTSE 100 index traded at the London International Financial and Options Exchange (LIFFE). The data for this research was obtained from SIRCA (http://www.sirca.org.au/), which covered a two-year period starting from 2/1/2000 and ending on 31/12/2001. The sample was made up of 63,094 call options and the daily Index value (adjusted for dividends). The following constraints were applied to filter the data series: moneyness (index/strike) outside [-1.01,0.9], maturity greater 175 days and less than 7 days and close price less than five where removed. The data series was then reduced to 26533 options. GB Libor rates were used for the risk-free rate inputs to the models. The parameters for the GOPM and BSOPM were estimated on a daily basis. However, when optimising the ANN, the data was split into 3 sets: in-sample, validation and out-of-sample. The in-sample set is used as input to the ANN, thus impacting on the weights. The validation set is used to evaluate the error at each epoch. The training of the network is terminated when the error of the validation increases. The out-of-sample data set is used for evaluating the ANN model.

4   Performance Measure

To analyze the model's performance, Absolute Relative Pricing Error (ARPE) was calculated for each model. where $\hat{C}$ and C is the model theoretical call price and the

actual call price. The errors are reported against maturity and moneyness. The models pricing performance is further analyzed by running a regression analysis to study the pricing behavior based on the moneyness and TTE (Time to Expiry).

$$ARPE \quad = \frac{1}{n} \sum_{i=1}^{n} \left| \frac{\hat{C}_i - C_i}{C_i} \right| \tag{13}$$

$$ARPE \quad = aTTE \; + bMon \; + c + \varepsilon \;\text{, where } \varepsilon \sim N(0, \sigma^2) \tag{14}$$

5   Results

Below the estimated mean values of the GOPM parameters are given in Table 1.

Table 1. GOPM average parameters

β_0	β_1	β_2	γ	$\sqrt{\sigma}$
3.109E-05	1.848E-01	3.184E-01	6.745E-01	22.37%

The risk premium is extracted from the option prices rather than from the index time series. It was found to be a little weak which in turn will have a minor impact on the option price. The annual volatility was found to be approx 22.37% which is a close approximation to the historical vol atility when estimated on the index prices 20.7%.

The ARPE figures are displayed in Table 2. On average, the performances of GOPM and NNiv models increase with maturity. Again, the BSOPM performs the worst on all accounts. On average, the NNiv has a slightly better performance than the GOPM. The GOPM model performs slightly better for MT contracts and the NNiv has a better performance for ST and LT contracts. All models have performed worst for ST-DOTM contracts. In addition, the GOPM performs the best for LT-DITM, whereas the NNiv and BSOPM are at their best for ST-DITM contracts. The NNiv seems to perform best for the MT-ITM contracts. The ARPE is regressed on TTE (Time to Expiry) and Mon (Moneyness). The estimated coefficients are displayed in Table 3. The coefficients R^2 and *F-test* for NNiv and the GOPM are similar. This indicates that the models price the options in a similar manner. For all models, the ARPE is smaller for longer maturities and increased moneyness. Therefore, ITM options are priced more accurately than the OTM options and the models price LT contracts with better accuracy than ST contracts.

6   The Empirical Dynamics of the Volatility Smile

In this section, we study the behavior of the implied volatility of the GOPM and the NNiv over time. The implied volatility of NNiv is given directly as the output of the model, whereas the GOPM produces a theoretical option price. To extract the implied

volatility from the GOPM, the option prices are equated to the BSOPM formula and the implied volatility is backed out. To study the volatility smile, a theoretical series is generated by an out-of-sample fit of the GOPM parameters calculated on the previous day. The implied volatility is then backed out for each option price using the BSOPM formula. Figure 1 shows the first day out-of-sample forecast for the GOPM. It can be seen that the volatilities are higher for the shorter maturities. As the option moves from ITM to ATM, they converge and the slopes are reversed.

Table 2. ARPE for all models against the options time to maturity.

Moneyness	Model	Time to Maturity			
		LT (Long TTE)	MT (Medium TTE)	ST (Short TTE)	TOTAL
DOTM (Deep Out of the money)	GOPM	0.179	0.246	0.384	0.267
	BSOPM	0.993	1.815	1.289	1.584
	NNiv	0.167	0.269	0.358	0.273
OTM (Out of The Money)	GOPM	0.105	0.157	0.272	0.202
	BSOPM	0.502	0.713	0.984	0.809
	NNiv	0.102	0.154	0.262	0.195
ATM (At the money)	GOPM	0.078	0.1	0.139	0.116
	BSOPM	0.274	0.262	0.277	0.27
	NNiv	0.086	0.094	0.135	0.112
ITM (In The Money)	GOPM	0.082	0.074	0.079	0.077
	BSOPM	0.169	0.12	0.077	0.104
	NNiv	0.081	0.071	0.075	0.074
DITM (Deep In The Money)	GOPM	0.039	0.061	0.05	0.053
	BSOPM	0.099	0.074	0.051	0.069
	NNiv	0.098	0.066	0.058	0.068
All	GOPM	*0.121*	*0.167*	*0.233*	*0.188*
	BSOPM	*0.589*	*0.924*	*0.742*	*0.818*
	NNiv	*0.119*	*0.172*	*0.222*	*0.185*

This has been highlighted by (Aboura 2005), where this point is of importance because the model is able to classify the volatilities according to their moneyness. Figure 2 displays the implied volatility for the end of the month using the same parameters as those for the first day of the month. The deformation of the smiles is obvious. The short maturity has a higher implied volatility; however, the skew of the longer maturities becomes more linear.

The theoretical volatilities calculated for NNiv at the first day of the month are displayed in Figure 3. The volatility smile is apparent in the graph. The graph also shows a similar behavior to the skews generated by the GOPM. For all maturities, the slope of the skews changes as it moves from ITM to OTM. So, as the option approaches ATM, the volatilities across all maturities become similar. For the ITM - LT option, the skew tends to be linear; then as it moves past the ATM point, the NN gives higher volatilities. For ITM-ST options, the NN generates much larger volatilities than the ITM-LT options. As ST option passes the ATM range, the slope changes direction forming the smile. However, the volatilities are smaller than the LT for OTM options.

Table 3. Regression results

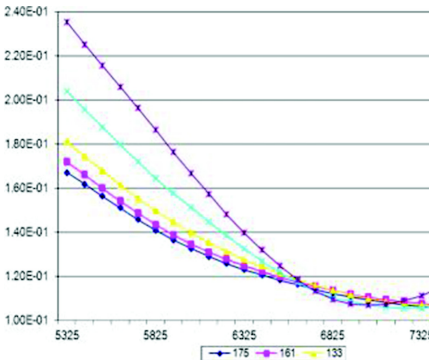|	NNiv	GOPM	BSOPM
C	2.097	2.052	13.931
	(-46.116)	(-45.072)	(69.214)
A	-0.327	-0.330	-0.769
	(-26.065)	(-26.306)	(-13.869)
B	-1.900	-1.851	-13.372
	(-40.983)	(-39.865)	(-65.148)
R^2	14.66%	13.71%	25.05%
F-test	1047.439	1014.297	2133.957

Fig. 1. GOPM out-of-sample skews for 4/9/2000
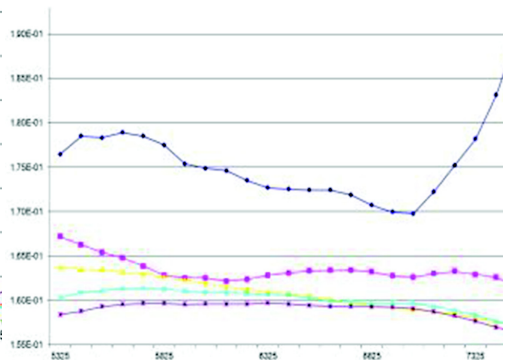

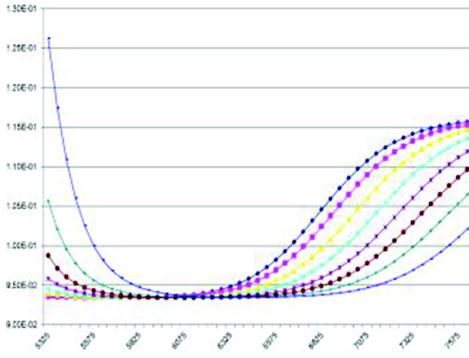**Fig. 2.** GOPM out-of-sample skews for 28/9/2000

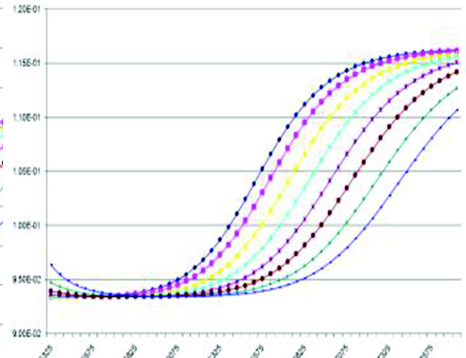**Fig. 3.** NNiv out of sample skews for 4/9/2000

Fig. 4. NNiv out of sample skews for 28/9/2000

In Figure 4, there does not seem to be any deformation of the skews. The skews also follow the same behavior as those in Figure 3. The deformation of the GOPM skews through time displays that the model is sensitive to changes in the underlying asset. The GOPM parameters are calibrated under LRNVR conditions by simulating the underlying asset process and pricing the options at expiry accordingly. However, the NNiv is designed and trained differently. The ability of the NNiv to price options out-of-sample is a key indicator of the ability of the NN to capture the dynamics of the options. The NNiv trained on mapping the input space (TTE, Moneyness, Historical Volatility) to Black-Scholes implied volatility. Due to the high number of variables in the NN, 1 day of historical data was not sufficient to produce a stable result. Hence, the NN was trained on 20 days of data, which contains substantial information. Therefore, when the smiles are generated out-of-sample towards the expiry of the options, the neural network is able to display the actual skew at that point in time. This is a good indication that the NN is a suitable tool for long-term forecasting.

7   Conclusion

The neural network was trained using the implied volatility as the target output of the neural network rather than the options price. This method has shown to have superior out-of-sample pricing performance over the BSOPM and GOPM. The neural network's performance was also compared to those of the GOPM and the BSOPM. The regression analysis shows the neural network behaves in a similar manner to the GOPM. This demonstrates the neural networks are capable of learning and reproducing option prices accurately by learning the implied volatility. The neural network was also shown to reproduce the implied volatility well into the future where as the GOPM start to show deformation. This behavior of the neural network demonstrates its capabilities to capture the underlying asset volatility structure and volatility levels through time. This can contain vital information on future market behaviors. In future research, this method will be further expanded and applied to Value-at-Risk models.

References

Aboura, S.: GARCH Option Pricing Under Skew. Journal: The ICFAI Journal of Applied Economics **4**(6), 78–86 (2005)

Amilon, H.: A neural network versus Black-Scholes: a comparison of pricing and hedging performances. Journal of Forecasting **22**(4), 317–335 (2003)

Bennell, J., Sutcliffe, C.: Black-Scholes versus artificial neural networks in pricing FTSE 100 options. Intelligent Systems in Accounting, Finance & Management **12**(4), 243–260 (2004)

Black, F., Scholes, M.: The Pricing of Options and Corporate Liabilities. The Journal of Political Economy **81**(3), 637–654 (1973)

Blair, B.J., Poon, S.H., et al.: Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. Journal of Econometrics **105**(1), 5–26 (2001)

Canina, L., Figlewski, S.: The informational content of implied volatility. Review of Financial Studies **6**(3), 659 (1993)

Christensen, B.J., Prabhala, N.R.: The relation between implied and realized volatility1. Journal of Financial Economics **50**(2), 125–150 (1998)

Cont, R., da Fonseca, J.: Deformation of implied volatility surfaces: an empirical analysis. In: Empirical Approaches to Financial Fluctuations (2001)

Cont, R., Fonseca, J.: Dynamics of implied volatility surfaces. Quantitative Finance **2**(1), 45–60 (2002)

Day, T.E., Lewis, C.M.: Stock Market Volatility and the Informational Content of Stock Index Options. Journal of Econometrics **52**(1) (1992)

Duan, J.C.: The GARCH Option Pricing Model. Mathematical Finance **5**(1), 13–32 (1995)

Duan, J.C., Simonato, J.G.: Empirical Martingale Simulation for Asset Prices. Management Science **44**(9), 1218–1233 (1998)

Dumas, B., Fleming, J., et al.: Implied volatility functions: Empirical tests. Journal of Finance **53**(6), 2059–2106 (1998)

Hanke, M.: Neural Network Approximation of Option Pricing Formulas for Analytically Intractable Option Pricing Models. Journal of Computational Intelligence in Finance **5**(5), 20–27 (1997)

Hull, J.: Options, futures, and other derivatives. Prentice Hall Upper Saddle River, NJ (2003)

Jiang, G.J., Tian, Y.S.: The model-free implied volatility and its information content. Review of Financial Studies **18**(4), 1305 (2005)

Karaali, O., Edelberg, W., et al.: Modelling volatility derivatives using neural networks. In: Proceedings of the IEEE/IAFE 1997 Computational Intelligence for Financial Engineering (CIFEr), pp. 280–286 (1997)

Lehar, A., Scheicher, M., et al.: GARCH vs. stochastic volatility: Option pricing and risk management. Journal of Banking and Finance **26**(2–3), 323–345 (2002)

Lehnert, T.: Mandelbrot and the Smile

Meissner, G., Kawano, N.: Capturing the Volatility Smile of Options on High-Tech Stocks-A Combined GARCH-Neural Network Approach. Journal of Economics and Finance **25**(3), 276–292 (2001)

Mostafa, F., Dillon, T.: Modelling volatility with mixture density networks. In: IEEE International Conference on Granular Computing, GrC 2008 (2008)

Mostafa, F., Dillon, T.: A neural network approach to option pricing. Computational Finance and its Applications III (2008)

Poon, S.H.: A practical Guide to Forecasting Financial Market Volatility. Wiley, New York (2005)

Rubinstein, M.: Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from August 23, 1976 through August 31, 1978. The Journal of Finance **40**(2), 455–480 (1985)

Tino, P., Schittenkopf, C., et al.: Financial volatility trading using recurrent neural networks. IEEE Transactions on Neural Networks **12**(4), 865–874 (2001)

Yao, J., Li, Y., et al.: Option price forecasting using neural networks. Omega **28**(4), 455–466 (2000)