Electronic Journal of Applied Statistical Analysis
EJASA, Electron. J. App. Stat. Anal.
http://siba-ese.unisalento.it/index.php/ejasa/index
e-ISSN: 2070-5948
DOI: 10.1285/i20705948v10n1p14

Short and long-term forecasting using artificial neural networks for stock prices in Palestine: a comparative study

By Safi, White

Published: 26 April 2017

Electronic Journal of Applied Statistical Analysis
Vol. 10, Issue 01, April 2017, 14-28
DOI: 10.1285/i20705948v10n1p14

Short and long-term forecasting using artificial neural networks for stock prices in Palestine: a comparative study

Samir Safi*[a] and Alexander White[b]

[a] *Department of Economics and Applied Statistics, The Islamic University of Gaza, Palestine*
[b] *Department of Mathematics, Texas State University, San Marcos TX, USA*

Published: 26 April 2017

To compare the forecast accuracy, Artificial Neural Networks and Autoregressive Integrated Moving Average models were fit with training data sets and then used to forecast prices in a test set. Three different measures of accuracy were computed: Root Mean Square Error, Mean Absolute Error and Mean Absolute Percentage Error. To determine how the accuracy depends on sample size, models were compared between short and long-term time series of stock closing prices from Palestine.

Keywords: Artificial Neural Network; Time Series, Forecasts; ARIMA; Stock Prices.

1 Introduction

Economic indicators, for example, stock prices, poverty rate, unemployment rate, etc, are vital measures of economic health in any country. Forecasting of economic indicators plays an important role in setting policy. Predicting the future values of economic indicators helps the decision makers take necessary steps and apply the required resources to avoid troubles and problems in a given sector.

Over the past three decades, there has been growing literature on applications of artificial neural networks (ANNs) to business and financial domains. In fact, a great deal of attention has been placed in the area of stock return forecasting. Prediction of stock price index movement is regarded as a challenging task of financial time series

*Corresponding author: samirsafi@gmail.com

ISSN: 2070-5948
http://siba-ese.unisalento.it/index.php/ejasa/index

prediction. However, once artificial neural network applications are successful, monetary rewards will be substantial (Kara et al., 2011). Many studies have reported promising results in successfully applying various types of ANN architectures for predicting stock returns. The results show that ANNs are an emerging and promising computational technology that will continue to be a challenging tool for future research (Thawornwong and Enke, 2003).

In recent years, many studies have come to a conclusion that the relationship between the financial and economic variables and the stock returns is nonlinear, and that ANNs can be accurately used to model problems involving nonlinearities. ANNs do not require a pre-specification during the modeling process because they independently learn the relationships inherent in the variables. Thus, ANNs are capable of performing nonlinear modeling without a priori knowledge about the relationship between input and output variables (Abhyankar et al., 1997).

To forecast a short-term movement of stock returns, daily data is likely to be selected. Researchers focused on a longer-term horizon are likely to use weekly or monthly data as inputs to the ANNs. Studies relying on economic variables would likely be restricted to monthly or quarterly data. This is due to the lag associated with the publication of economic indicators. A review of the literature prior to 2003 indicated that 26 studies modeled the ANNs using daily data, whereas monthly data was used in 10 studies. In addition, two studies examined the predictability of ANNs on data of different time frequencies, while quarterly and weekly data were each selected by two studies (Thawornwong and Enke, 2003).

The main purpose of this paper is to find a more accurate and reliable forecasting model for the stock prices. We use ANN and ARIMA models to forecast stock prices of the Bank of Palestine. This paper is structured as follows. The next section presents literature review; the third section describes the methodology including description of the data and measures of forecasting accuracy; in the fourth section, we present the empirical results on in-sample training for the two forecasting cases fitting ARIMA and ANN models for stock closing prices data; and the last section concludes some important results of this paper and offers future research.

2 Literature Review

ANNs are one of the most powerful tools for pattern classification due to their nonlinear and non-parametric adaptive-learning properties. Since they were popularized by Rumelhart and colleagues in 1986 , ANNs have garnered considerable attention of research workers, as they can handle the complex non-linearity problems better than the conventional statistical techniques. Shrivastava et al. (2012) identify two main drawbacks of the conventional numerical and statistical models: firstly, the statistical models are not useful to study the highly nonlinear relationships between response variable and its predictors; secondly, there is no ultimate end in finding the best predictors. ANNs are able to get rid of these two drawbacks. Many studies compare ANNs with other traditional techniques, see for example White and Safi (2016), Valipour et al. (2012), Safi

(2013), Aksoy and Dahamsheh (2009), Zhang and Kline (2007), Lee and Chen (2005), among others.

The application of ANNs to short-term load forecasting has gained a lot of attention in the last two decades, see for example Potočnik et al. (2015), Grant (2014), Rodrigues et al. (2014), Taylor (2012), Beccali et al. (2008), Kandil et al. (2006), Hippert et al. (2001), Toth et al. (2000), Szkuta et al. (1999), El-Sharkawi and Niebur (1996), among others.

Since the 1990's, ANN models have been used to model financial data. Kohzadi et al. (1996) compared neural network and ARIMA models to forecast US monthly live cattle and wheat cash prices from 1950 to 1990. They showed that the neural network forecasts were considerably more accurate than those of the traditional ARIMA models, which were used as a benchmark. The mean squared error, absolute mean error, and mean absolute percent error were all lower on average for the neural network forecast than for the ARIMA model. Kohzadi et al. (1996) conjectured that the neural network model performed better than ARIMA because the data contained non-linear or chaotic behavior, which could not be fully captured by the linear ARIMA model. Desai and Bharati (1998) compared linear regression and neural network methods for predicting excess returns for the S&P 500 index. They showed that one cannot say that the linear regression forecasts are conditionally efficient with respect to the neural networks forecasts with any degree of confidence, however, the neural networks forecasts are conditionally efficient with respect to the linear regression forecasts with some confidence.

More recently, Dhamija and Bhalla (2011) showed that ANNs can be effectively used in forecasting exchange rates and hence in designing trading strategies. They showed that neural networks can simultaneously and effectively extract the non-linear functional form as well as model parameters. In addition, neural networks provide quantitative finance with strong support in problems related to non-parametric regression. Li et al. (2004), reported on an application of recurrent neural networks (RNNs) to model and forecast short-term exchange rate movements. They showed that that a discrete-time RNN performs better than the traditional methods in forecasting short-term foreign exchange rates. While, Masoud (2014) used ANN models to predict the direction of movement for the Libyan stock market from January 2, 2007 to March 28, 2013 and found that ANNs are a better alternative technique for forecasting the daily stock market prices. In particular, he showed that the ANN model accurately predicted the direction of movement with the average prediction rate 91%.

Grant (2014) tested the ANN model against other forecasting methods including simple moving average (SMA), linear regression, and multivariate adaptive regression splines (MARSplines) and showed that ANN was effective at forecasting peak building electrical demand in a large government building sixty minutes into the future. The ANN model outperformed the other forecasting methods tested with a mean absolute percentage error (MAPE) of 3.9% as compared to the SMA, linear regression, and MARSplines MAPEs of 7.7%, 17.3%, and 7.0% respectively. Additionally, the ANN model realized an absolute maximum error (AME) of 18.2% as compared to the SMA, linear regression, and MARSplines AMEs of 26.2%, 45.1%, and 22.5% respectively

Rodrigues et al. (2014) have used ANN for Short Term Load Forecasting (STLF).

The ANNs are recognized to be a potential methodology for modeling hourly and daily energy consumption and load forecasting. They concluded that a feed-forward ANN using the Levenberg-Marquardt algorithm performed well providing a reliable model for forecasting household electric energy consumption for 93 real households, in Lisbon, Portugal, between February 2000 and July 2001.

Valipour et al. (2012) used monthly discharges data from 1960 to 2007 in Dez reservoir inflow at the Taleh Zang station. Using root mean square error (RMSE) and mean bias error (MBE) they compared various forecasting methods. The results indicated that the ARIMA model performed better than ARMA model due to the non-stationarity of the time series in both training and forecasting phases. But the dynamic autoregressive artificial neural network was superior to static autoregressive artificial neural network, due to the output delay effect as input to network and increase in the power of network training compared to autoregressive static neural network and in general compared to the ARMA and ARIMA models in both training and forecasting stages. This comparison is done by.

Kandil et al. (2006) suggested a simple multi-layered feedforward ANN and their results showed that the ANN is able to interpolate among the load and weather variables pattern data of training sets to provide the future load pattern. However, these are preliminary results. The possibility for better results exists and can be achieved by using: (1) more advanced types of ANN, (2) better selection of input variables, (3) better ANN architecture and (4) better selection of the training set.

Mohammadi et al. (2005) forecasted spring inflow to the Amir Kabir reservoir in the Karaj river watershed, located to the northwest of Tehran (Iran). They used three different methods, ANN, ARIMA time series and regression analysis. The results showed that ANN can be an effective tool for reservoir inflow forecasting in the Amir Kabir reservoir using snowmelt equivalent data.

Darbellay and Slama (2000) examined the dependence structure of the electric load time series of the Czech Republic and indicated that the autocorrelations in this time series are predominantly linear. For univariate modelling, we found that, indeed, the forecasting abilities of a linear model and a nonlinear model were not very different. These models were, respectively, an ARIMA model and a neural network.

3 Methodology

3.1 Data Description

We use a data set of stock closing prices from Palestine Exchange, http://www.pex.ps/. We considered two daily data sets for stock of Bank of Palestine and Stock of Jerusalem during the period of the 2006-2015. The results obtained for the two stocks were quite similar, so we present only the results for the Bank of Palestine. Results for the Stock of Jerusalem are available on request. The time-series plot of the Bank of Palestine data is presented in the top-left of Figure 1. From this plot, we can see that the prices are not linear over time and show large fluctuations. This indicates that one must be cautious using ARIMA models as they may not provide accurate forecasts.

In this study, 10% of the sample size is used as the testing sample. A training sample is used for the model building, and the testing sample is used for the model validation at the end of analysis. To see the impact of sample size on the comparison, series of different lengths were considered. First, we considered all 2449 observations for daily stock prices of the Bank of Palestine. This “long term” series was separated into two sub-samples- the training sample (2204 observations) and testing sample (245 observations). By examining the most recent observations, we also modeled moderate and short term series. The final 300 observations were chosen as a moderate size series, which was separated into two sub-samples- the training sample (270 observations) and testing sample (30 observations). Similarly the final 50 observations were used as short term time series and separated into two sub-samples- the training sample (45 observations) and testing sample (5 observations).

3.2 Forecast Models

For each of the data sets, models fit on the training sample were used to forecast the test sample. Two types of forecast models were used: ARIMA and ANN models.

3.2.1 ARIMA Models

The general ARIMA(p, d, q) model is given by Box et al. (2015)

$$\phi(B)\, \nabla^d Y_i \;=\; \theta(B)\, \varepsilon_i, \tag{1}$$

where $d \geq 1$ is the degree of differencing, $\nabla = 1 - B$ is the differencing operator, the *lag* operator B, is defined as $BY_t = Y_{t-1}$, the operator which gives the previous value of the series. $\phi(B)$ and $\theta(B)$ are polynomials of degree p and q in B,

$$\phi(B) = \; 1 - \phi_1 B - \phi_2 B^2 - \cdots - \phi_p B^p$$

and

$$\theta(B) = \; 1 - \theta_1 B - \theta_2 B^2 - . \cdots - \theta_q B^q$$

3.2.2 Artificial Neural Network

R-software (2015) was used for fitting ANN and ARIMA models for the stock closing price time series data. The *nnetar* function from the R-package *forecast* (2015) was used to fit neural networks. This function creates feed-forward neural networks with a single hidden layer using lagged inputs for forecasting univariate time series. The *nnetar* function fits an Neural Network Autoregressive models $NNAR(p, P, k)$ model. For non-seasonal time series, the default is the optimal number of lags (according to the AIC) for a linear $AR\,(p)$ model. For seasonal time series, the default values are $P = 1$ and p is chosen from the optimal linear model fitted to the seasonally adjusted data, $k = \frac{1}{2}\,(p + P + 1)$ (rounded to the nearest integer). By default, 25 networks with random starting values are trained and their predictions averaged (Hyndman, 2012).

3.3 Measures of Forecasting Accuracy

Accuracy is an important issue in forecasting. Researchers tend to add more and more variables in the proposed forecasting model. Does a more complex model necessarily do a better job than a simpler one? The conclusion reached when evaluating forecasts can vary for identical data when applying different measures of evaluation. Therefore, it is of interest to select several complementary measures that can expose differences in the forecasts (Lyhagen et al., 2015). For this reason, many measures of forecasting accuracy have been developed, and several authors have discussed the usage for these measurements and made comparisons among the accuracy of forecasting methods in univariate time series data, see for example Hyndman and Athanasopoulos (2014), Cryer and Chan (2008), Wei (2006), among others.

When choosing models, it is common to use a portion of the available data for testing, and use the rest of the data for estimating (or "training") the model. Then the testing data can be used to measure how well the model is likely to forecast on new data. In other words, forecast accuracy should be computed by using test data that was not used when computing the forecasts. The size of the test data set is typically about 20% of the total sample, although this value depends on how long the sample is and how far ahead you want to forecast. The size of the test set should ideally be at least as large as the maximum forecast horizon required (Hyndman and Athanasopoulos, 2014).

Suppose $y_1,\, y_2, \ldots, y_n$ denotes the data set, and we split it into two parts: the training data $y_1,\, y_2, \ldots, y_t$ and the test data $y_{t+1},\, y_{t+2}, \ldots, y_n$. To check the accuracy of the forecasting method, we will estimate the model's parameters using the training data, and forecast the next $n - t$ observations. Then, we compare the test data with these forecasts.

Definition 1. The forecast errors are the difference between the actual values in the test set data and the forecasts produced using the data in the training set. Thus

$$e_i = y_i - \hat{y}_{i|t},\ \ i = t+1,\ \ldots,\ n \tag{2}$$

The selected best forecasting models will be compared using three different forecasting criteria: Mean Absolute Error (MAE), the Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE).

Definition 2. The Mean Absolute Error (MAE) is defined as

$$MAE = (n-t)^{-1} \sum_{i=t+1}^{n} \left| y_i - \hat{y}_{i|t} \right| \tag{3}$$

Definition 3. The Root Mean Squared Error (RMSE) is defined as

$$RMSE = \sqrt{(n-t)^{-1} \sum_{i=t+1}^{n} \left(y_i - \hat{y}_{i|t} \right)^2} \tag{4}$$

The RMSE has been popular, largely because of its theoretical relevance in statistical modeling (Hyndman and Koehler, 2006). However this measure is more sensitive to outliers than MAE which has led some authors, for example, see Armstrong (2001) to recommend using other forecast accuracy measures. MAE is preferable in case of the existence of outliers. It is recommended to use MAE or RMSE when comparing forecast methods on a single data set. This means, the MAE and RMSE are used if all forecasts are measured on the same scale.

Definition 4. The Mean Absolute Percentage Error (MAPE) is defined as

$$MAPE = (n-t)^{-1} \sum_{i=t+1}^{n} \left| \frac{y_i - \hat{y}_{i|t}}{y_i} \right| \times 100 \tag{5}$$

MAPE presents the forecast error in terms of percentage and hence it is scale invariant and unit free (Lyhagen et al., 2015). However, MAPE has the disadvantage of being infinite or undefined if $y_i = 0$ for any i in the period of interest, and having an extremely skewed distribution when any y_i is close to zero. It is recommended to use MAPE when comparing the accuracy of the same or different methods on different time series data with different scales, unless the data contain zeros or small values (Hyndman and Koehler, 2006).

The evaluation criterion for these measures of forecasting accuracy is that the smaller value obtained, the better is the forecasting ability of the model (McKenzie, 2011).

Definition 5. The efficiency of the proposed forecast method relative to that of benchmark method in terms of the RMSE, ρ, is given by

$$\rho = \frac{RMSE_p}{RMSE_b}, \tag{6}$$

where $RMSE_p$ and $RMSE_b$ are the RMSE from the proposed and benchmark methods, respectively. Usually the benchmark method is the naive method (Hyndman and Koehler, 2006). A ratio less than one indicates that the forecast performance of the proposed method is more efficient than benchmark method and if ρ is close to one, then the proposed forecast methods is nearly as efficient as the benchmark forecast. Otherwise, the proposed method performs poorly (White and Safi, 2016).

4 Empirical Results on In-Sample Training

This section presents the empirical results on In-Sample Training for fitting models for stock closing prices data for bank of Palestine by using two different approaches, ANN and ARIMA(p,d,q) models. The forecasting results are presented in the following subsections.

4.1 Fitting Models for Long Term Series of Daily Stock Closing Prices

In this section we fit daily stock closing prices using two forecasting methods, namely: ANNs (the proposed method) and ARIMA. Figure 1 contains a time series plot of the original stock price data, the fitted model from the training data set and the forecasts for the test data set. Table 1 shows the three measures of forecasting accuracy for the

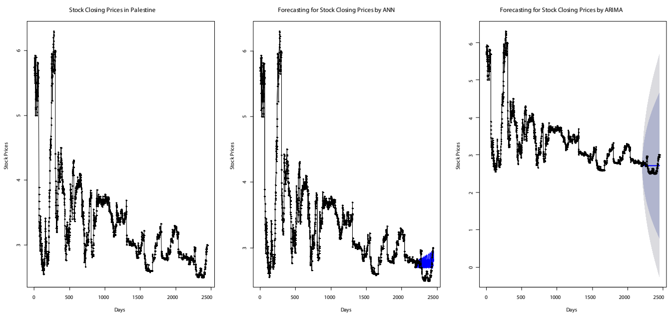

Figure 1: Predictions for Daily Data

two methods, and the relative efficiencies of the error of ANN model to the ARIMA models for daily stock prices.

The *nnetar* function was used to fit neural networks. Since the ANN model depends on a random starting values, the final model used was an average of 25 networks, each of which was a 22-12-1 network with 289 weights. Table 1 shows that for ANN method, the values for RMSE, MAE, and MAPE equal 0.0781, 0.0366, and 1.0416, respectively.

The *auto.arima* command in R was used to fit the ARIMA model. This algorithm uses maximum likelihood estimation (MLE) to fit an ARIMA(p, d, q) model for different choices of p, d and q, and then compares Akaike Information Criteria to determine the best model. Results for the final ARIMA(4,1,2) are shown in Table A1. We note that the P-value for each of the estimates of ARIMA(4,1,2) coefficients is significantly different from zero. In Table 1, we see that the RMSE, MAE, and MAPE values equal 0.0819, 0.0377, and 1.0624, respectively, for the ARIMA(4,1,2) model.

Comparing the measures of forecast accuracy in Table 1 between the models, we see that for RMSE, the relative efficiency of ANNs to ARIMA equals 0.9538. This result indicates that the RMSE for the ANN model equals 95.38% of ARIMA models. We obtained similar results for the MAE and MAPE. The relative efficiencies for MAE of ANNs to ARIMA equals 0.9701 and for MAPE is 0.9805. Focusing on the right of the graph in Figure 1, we see that the ARIMA forecasts are essentially constant over the test set. The ANN forecasts shown in the middle of the graph, on the other hand, capture some of the fluctuation shown in the data.

4.2 Fitting Models for Moderate Term Stock Closing Prices Data

In this section we fit moderate term series of stock prices consisting of the final 300 observations using the same two forecasting methods. Table 2 shows the three measures of forecasting accuracy for the two methods, and the relative efficiency of the error of ANN model to ARIMA model for moderate data for stock prices.

As before, the final ANN model is an average of 25 networks. In this case, due to the

Table 1: Accuracy Measures and Relative Efficiencies: Long Term Series

RMSE	MAE	MAPE	Method
0.0781	0.0366	1.0416	ANN
0.0819	0.0377	1.0624	ARIMA
0.9538	0.9701	0.9805	ANN/ARIMA

Table 2: Accuracy Measures and Relative Efficiencies: Moderate Term

RMSE	MAE	MAPE	Method
0.1507	0.0875	2.2129	ANN
0.1522	0.0875	2.2177	ARIMA
0.9902	1.0001	0.9978	ANN/ARIMA

smaller sample size of the training set, each network was a 1-1-1 network with 4 weights. Table 2 shows that for ANN method, the values for RMSE, MAE, and MAPE equal 0.1507, 0.0875, and 2.2129, respectively. Note these are 2-3 times larger than the values for the fits based on the entire daily price series.

The results from MLE of final ARIMA model is shown in Table A2. In this case, the best fit model is a much simpler ARIMA(1,2,1). As with the daily data, the P-value for the estimate of autoregressive and moving average coefficients in the ARIMA(1,2,1) model are significantly different from zero. In Table 2 the values for RMSE, MAE, and MAPE equal 0.1522, 0.0875, and 2.2177, respectively, for the ARIMA model which are 2-3 times larger than for the daily data.

Table 2 also shows that for RMSE, the relative efficiency of ANNs to ARIMA equal 0.9902. This result indicates that RMSEs for ANNs equal 99.02% of ARIMA model. Additionally, the results based on MAE and MAPE mimic the same as RMSE. As with the daily price data, the ANN forecasts are the most accurate.

4.3 Fitting Models for Short Term Stock Closing Prices Data

In this section we present the fitting short term stock prices using the same two forecasting methods. The size for this data is the final 50 observations. Table 3 shows the three measures of forecasting accuracy for the two methods, and the relative efficiency of the error of ANN model to ARIMA model for short term stock prices.

Once again, the final ANN model is an average of 25 networks, which in this case are 3-2-1 networks with 11 weights. Table 3 shows that for the ANN method, the values for

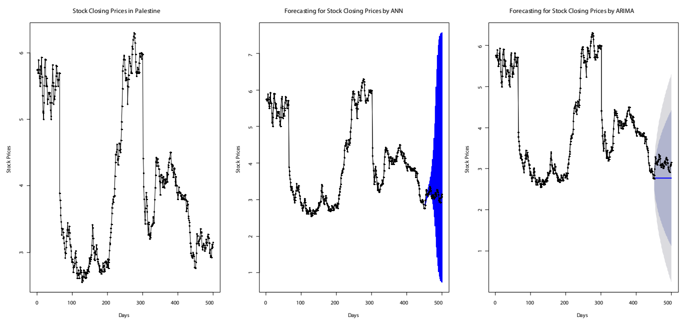

Figure 2: Predictions for Moderate Data

RMSE, MAE, and MAPE equal 0.1333, 0.1121, and 2.0524, respectively. As expected these are larger than observed for the moderate and long term series. However, the increase from moderate to short term is not as large as seen between the daily and the moderate data set.

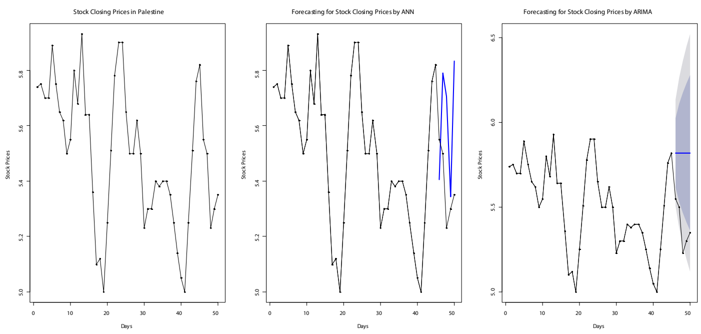

Figure 3: Predictions for Short-Term Data

The best fit ARIMA model is ARIMA(0,1,0). Hence, the differences in prices in successive data points represent white noise. In Table 3, we see that the values for RMSE, MAE, and MAPE equal 0.1579, 0.1228, and 2.2319, respectively, for this model.

Considering RMSE, the relative efficiency of ANN to ARIMA equals 0.8446. This result indicates that RMSEs for ANNs equal 84.46% of ARIMA models. As before, the results based on MAE and MAPE mimic those for RMSE. Note, that the performance of ANN relative to ARIMA is improving. Therefore, the ANN is much more efficient than ARIMA for short-term data.

Table 3: Accuracy Measures and Relative Efficiencies: Short-Term Data

RMSE	MAE	MAPE	Method
0.1333	0.1121	2.0524	ANN
0.1579	0.1228	2.2319	ARIMA
0.8446	0.9128	0.9195	ANN/ARIMA

4.4 Comparison of ANN Performance across Different Datasets

As we have mentioned in Section 3.3, MAPE is used to compare the accuracy of the same or different methods for different time series data with different scales. In this section, we compare the performance of ANN across daily, moderate and short-term stock data sets. Tables 1 – 3 show that the values of MAPE equal 1.0416%, 2.2129%, and 2.0524% for daily, moderate and short-term stock data sets, respectively. Using the MAPE for the daily data as a benchmark, this result indicates that the ratios of MAPEs for daily data equal 47.07% and 50.75%, respectively with respect to moderate and short-term data. Taking a similar approach for the other forecasting method, ARIMA, we see the values of MAPE equal 1.0624%, 2.2177%, and 2.2319% for daily, moderate and short-term stock data sets, respectively using the ARIMA model. This result indicates that the ratios of MAPEs for daily data equal 47.90% and 47.60%, respectively with respect to moderate and short-term data. Using MAPE as a measure of comparison, these results reveal that ANN and ARIMA perform better for the larger daily data set than the moderate and short-term data sets.

5 Conclusion and Future Research

In this paper, we compared the forecast accuracy for two methods using data from the Palestine stock market. In addition to comparing the methods, three levels of length of series were used: daily, moderate and short-term. As is common for stock prices, the data exhibits considerable variability, nonlinearity and non-stationarity. The results indicate that the ANN models produced the most accurate forecasts at each level of granularity. Furthermore, the forecast for ANN and ARIMA models based on larger and finer data sets were more accurate than those on the smaller data sets. These results add to the growing body of literature that recommends the use of ANNs to forecast economic data. In addition, the results indicate that the ANNs will become more accurate as the more information is fed into the model (i. e. larger data sets). ANN may often be more preferable than assuming an ARIMA model when the actual model is non-linear. In other words, it is sometimes better to ignore the complexity of time series models and use the ANN technique rather than to incorrectly assume the model is an ARIMA.

This paper focuses on forecasts of the future of univariate time series based on past and present values. It is known that many economic indicators are correlated and that it is possible to improve forecasts by using multivariate models including multiple indicators simultaneously. A logical next step is to compare the performance of more traditional methods of multivariate regression with multivariate ANNs in the context of economic data. Naturally, as the number of indicators used increases, the models become more complicated, so it will be important to compare the complexity of the resulting models. Additionally, comparing ANN with nonlinear time series models.

Acknowledgment
We would like to express the deepest appreciation to the referees for their valuable comments, suggestions, and review on earlier draft of this paper. Also, we are grateful for Ramzy Khalifa and Said El-Abadllah, Palestine Exchange for providing us with the data sets.

References

Abhyankar, A., Copeland, L. S., and Wong, W. (1997). Uncovering nonlinear structure in real-time stock-market indexes: the s&p 500, the dax, the nikkei 225, and the ftse-100. *Journal of Business & Economic Statistics*, 15(1):1–14.

Aksoy, H. and Dahamsheh, A. (2009). Artificial neural network models for forecasting monthly precipitation in jordan. *Stochastic Environmental Research and Risk Assessment*, 23(7):917–931.

Beccali, M., Cellura, M., Brano, V. L., and Marvuglia, A. (2008). Short-term prediction of household electricity consumption: Assessing weather sensitivity in a mediterranean area. *Renewable and Sustainable Energy Reviews*, 12(8):2040–2065.

Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.

Cryer, J. and Chan, K.-S. (2008). *Time Series Analysis with Applications in R*. Springer, New York.

Darbellay, G. A. and Slama, M. (2000). Forecasting the short-term demand for electricity: Do neural networks stand a better chance? *International Journal of Forecasting*, 16(1):71–83.

Desai, V. S. and Bharati, R. (1998). A comparison of linear regression and neural network methods for predicting excess returns on large stocks. *Annals of Operations Research*, 78:127–163.

Dhamija, A. and Bhalla, V. (2011). Exchange rate forecasting: comparison of various architectures of neural networks. *Neural Computing and Applications*, 20(3):355–363.

El-Sharkawi, M. and Niebur, D. (1996). Short-term load forecasting with artificial neural networks: the international activities. *IEEE power engineering society: tutorial course on artificial neural networks with applications to power systems*, pages 90–103.

Grant, J. L. (2014). *Short-Term Peak Demand Forecasting Using an Artificial Neural*

Network with Controlled Peak Demand Through Intelligent Electrical Loading. PhD thesis, "University of Miami.

Hippert, H. S., Pedreira, C. E., and Souza, R. C. (2001). Neural networks for short-term load forecasting: A review and evaluation. *Power Systems, IEEE Transactions on*, 16(1):44–55.

Hyndman, R. J. (2012). New in forecast 4.0 [weblog post]. retreived from. https://urldefense.proofpoint.com/v2/url?u=http-3A__robjhyndman.com_hyndsight_forecast4_&d=BQIGaQ&c=OrYO-caJHQE1g_AJU3az1awi55It-bjDIQrtRiZ6WBk&r=WZUj6zQTvPZOGR7Wh0snaA&m=9nCAUNq55kxazLhzlg-7tNneKmx26Wd7mo_RlZ-zd-0&s=HDaE62nbcBgAsKDBFPnG29BXRzeqQJD1hTIHtN9cS2U&e=.

Hyndman, R. J. (2015). *forecast: Forecasting functions for time series and linear models*. R package version 6.2.

Hyndman, R. J. and Athanasopoulos, G. (2014). *Forecasting: principles and practice*. OTexts.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International journal of forecasting*, 22(4):679–688.

Kandil, N., Wamkeue, R., Saad, M., and Georges, S. (2006). An efficient approach for short term load forecasting using artificial neural networks. *International Journal of Electrical Power & Energy Systems*, 28(8):525–530.

Kara, Y., Boyacioglu, M. A., and Baykan, Ö. K. (2011). Predicting direction of stock price index movement using artificial neural networks and support vector machines: The sample of the istanbul stock exchange. *Expert systems with Applications*, 38(5):5311–5319.

Kohzadi, N., Boyd, M. S., Kermanshahi, B., and Kaastra, I. (1996). A comparison of artificial neural network and time series models for forecasting commodity prices. *Neurocomputing*, 10(2):169–181.

Lee, T.-S. and Chen, I.-F. (2005). A two-stage hybrid credit scoring model using artificial neural networks and multivariate adaptive regression splines. *Expert Systems with Applications*, 28(4):743–752.

Li, L.-K., Pang, W.-K., Yu, W.-T., and Troutt, M. D. (2004). Forecasting short-term exchange rates: A recurrent neural network approach. *Neural Networks in Business Forecasting*, page 195.

Lyhagen, J., Ekberg, S., and Eidestedt, R. (2015). Beating the var: Improving swedish gdp forecasts using error and intercept corrections. *Journal of Forecasting*, 34(5):354–363.

Masoud, N. (2014). Predicting direction of stock prices index movement using artificial neural networks: The case of Libyan financial market. *British Journal of Economics, Management & Trade*, 4(4):597–619.

McKenzie, J. (2011). Mean absolute percentage error and bias in economic forecasting. *Economics Letters*, 113(3):259–262.

Mohammadi, K., Eslami, H., and Dardashti, S. D. (2005). Comparison of regression, arima and ann models for reservoir inflow forecasting using snowmelt equivalent (a case study of karaj). *J. Agric. Sci. Technol*, 7:17–30.

Potočnik, P., Strmčnik, E., and Govekar, E. (2015). Linear and neural network-based models for short-term heat load forecasting. *Strojniški vestnik-Journal of Mechanical Engineering*, 61(9):543–550.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rodrigues, F., Cardeira, C., and Calado, J. M. F. (2014). The daily and hourly energy consumption and load forecasting using artificial neural network method: a case study using a set of 93 households in portugal. *Energy Procedia*, 62:220–229.

Rumelhart, D. E., Smolensky, P., McClelland, J. L., and Hinton, G. E. (1986). Schemata and sequential thought processes in PDP models. In McClelland, J. L., Rumelhart, D. E., and PDP Research Group, C., editors, *Parallel Distributed Processing*, pages 7–57. MIT Press, Cambridge, MA, USA.

Safi, S. K. (2013). Artificial neural networks approach to time series forecasting for electricity consumption in gaza strip. *IUG Journal of Natural and Engineering Studies*, 21(2):1–22.

Shrivastava, G., Karmakar, S., Kowar, M. K., and Guhathakurta, P. (2012). Application of artificial neural networks in weather forecasting: a comprehensive literature review. *International Journal of Computer Applications*, 51(18).

Szkuta, B., Sanabria, L., and Dillon, T. (1999). Electricity price short-term forecasting using artificial neural networks. *Power Systems, IEEE Transactions on*, 14(3):851–857.

Taylor, J. W. (2012). Short-term load forecasting with exponentially weighted methods. *Power Systems, IEEE Transactions on*, 27(1):458–464.

Thawornwong, S. and Enke, D. (2003). Forecasting stock returns with artificial neural networks. In Zhang, G. P., editor, *Neural Networks in Business Forecasting*, pages 47–74. Idea Group Inc, Hershey, PA, USA.

Toth, E., Brath, A., and Montanari, A. (2000). Comparison of short-term rainfall prediction models for real-time flood forecasting. *Journal of Hydrology*, 239(1):132–147.

Valipour, M., Banihabib, M., and Behbahani, S. (2012). Monthly inflow forecasting using autoregressive artificial neural network. *Journal of Applied Sciences*, 12(20):2139.

Wei, W. W.-S. (2006). *Time Series Analysis Univariate and Multivariate Methods*. Second edition, Pearson Education, Inc.

White, A. K. and Safi, S. K. (2016). The efficiency of artificial neural networks for forecasting in the presence of autocorrelated disturbances. *International Journal of Statistics and Probability*, 5(2):51.

Zhang, G. P. and Kline, D. M. (2007). Quarterly time-series forecasting with neural networks. *Neural Networks, IEEE Transactions on*, 18(6):1800–1814.

Appendix: Tables of Estimates

Table A1: Maximum Likelihood Estimates: Long Term Daily Stock Prices

	ar1	ar2	ar3	ar4	ma1	ma2
Coefficient	-0.744	-0.73	0.142	0.029	0.849	0.887
SE	0.063	0.145	0.030	0.030	0.059	0.144
T	-11.791	-5.028	4.673	0.967	14.345	6.148

Table A2: Maximum Likelihood Estimates: Moderate Term Daily Stock Prices

	ar1	ma1
Coefficient	0.153	-.983
SE	0.062	0.012
T	2.479	-83.271