

## The Usefulness of Artificial Neural Networks in Forecasting Exchange Rates

**Blerina Vika**

*Statistics & Applied Informatics Department, Faculty of Economy, University of Tirana*

**Kozeta Sevrani**

*Statistics & Applied Informatics Department, Faculty of Economy, University of Tirana*

**Ilir Vika**

*Research Department, Bank of Albania*

Doi:10.5901/ajis.2016.v5n1p73

### Abstract

*This article contributes to the neural network literature by demonstrating how potent and useful they can be as a tool in the process of economic and financial decision makings. We probe into the usefulness of Nonlinear Autoregressive Networks (NAR) in comparison to the ARIMA models that are commonly used as a benchmark for forecasting exchange rates. To demonstrate it we chose the USD/EUR exchange rate, as a considerably volatile and a highly transacted asset in the international financial market, yet very disputed in academic works due to its often large divergences from the fundamental levels suggested by economic theories. Although through a modest application, our findings show that neural network models can add value and possibly outperform traditional models used to forecast exchange rates. The results were affirmative that the nonlinear autoregressive net consistently beat the ARIMA (and the random walk) static forecasts of the USD/EUR exchange rate.*

**Keywords:** artificial neural network; ARIMA model; exchange rate forecasting.

### 1. Introduction

Time series forecasting is increasingly used as a valuable tool to provide information in the decision making process. It is important for fund managers, corporate treasurers, global traders and policy makers, to mention a few. However, forecasting of financial series has proven a very challenging task, especially for volatile time series such as exchange rates. Being the largest and the most liquid market with trillions of US dollars transacted every day, the empirical literature based on theoretical models has often found it difficult to beat forecasts from more naive random walk processes.

Because exchange rates are influenced by many economic, political and psychological factors, it has been hard to identify a unique economic model that can provide reliable forecasts. Some authors state that "the poor explanatory power of existing theories of the exchange rate is most likely the major weakness of international macroeconomics" (Bacchetta & Wincoop, 2006), and that empirical exchange rate models "...generally fail badly in out-of-sample forecasting tests in the sense that they fail to outperform a random walk" (Sarno & Taylor, 2002). Even if exchange rate theories look fundamentally sound, many researchers blame the empirical implementation as a linear statistical model for the dismal forecast performance. Thus, they propose the use of time-varying or non-linear methods, which could help to better capture the exchange rate adjustment toward its long-term equilibrium in a nonlinear fashion (Kilian & Taylor, 2001).

As a matter of fact, exchange rate series often display signs of nonlinearity, which traditional linear forecasting techniques are ill equipped to handle, often producing unsatisfactory results (Philip, Taofiki, & Bidem, 2011). Researchers confronted with these characteristics turn to techniques that are heuristic and nonlinear. In this rather short article, we first estimate a linear autoregressive model that resembles the random walk to forecast the monthly USD/EUR exchange rate, and then generate a nonlinear autoregressive artificial neural network in an attempt to improve the traditional univariate model forecasting. Both types of models will finally be compared on their forecast performance, where the focus is on minimizing out-of-sample forecast error rather than maximizing in-sample 'goodness of fit'.

Indeed, the random walk modeling (of the form  $y_t = y_{t-1} + \epsilon_t$ , where  $\epsilon_t$  is the random error term) may not effectively

handle the uncertainty or instability nature that characterizes exchange rate movements. Instead, we employ the autoregressive integrated moving average (ARIMA) model, which is widely used as a benchmark in time series forecasting and analysis. ARIMA is a specific subset of univariate modeling, which assumes that the historical characteristics of a time series (i.e. its systematic structural features) will be present in the future; therefore they can be convenient for forecasting purposes.

On the other side, artificial neural networks (ANNs) have in the past decade emerged as a powerful alternative method for time series forecasting due to their higher capabilities. ANNs are nonparametric data driven approaches that can capture nonlinear data structures without prior assumption about the underlying relationship in a particular problem. ANNs can learn from examples and demonstrate some capability for generalization, beyond the training data. A large number of papers that make use of ANNs for prediction can be found in the literature. Nevertheless, the results reported are often mixed.

To give a brief view on the results, the forecast evaluation based on the ME, MAE and RMSE indicates the neural network model has been able to provide better one-month ahead forecasts than the ARIMA model throughout the 2014. These findings can lend support to the ANN literature for its applicability in exchange rate forecasting techniques. In what follows, we begin by describing the basic facts about the two methodologies we used in the forecasting process, and then we present and compare the respective results.

## 2. Competing Methodologies

The univariate ARIMA and ANN models are applied on the monthly USD/EUR exchange rate, as published online in the Thomson Reuters website. The data analysis starts from January 1994 till December 2014 – a total of 252 data points. The EViews 7 software was used for the ARIMA modeling. Whereas all partial measurements for the ANNs models are implemented directly in the core of MATLAB 2010b environment, which was used for the experiments. Below we describe separately the two methodologies applied in this article and show the preferred models in each of them. The chosen models are then compared on the basis of their one-month-ahead forecast performance.

### 2.1 ARIMA modeling

Contrary to multivariate models with explanatory variables that are based on economic theories, ARIMA is a purely statistical model, in which a time series is regressed on its own past values (the autoregressive component) plus current and lagged values of a 'white noise' error term (the moving average component) (Meylar, Kenny, & Quinn, 1998). So by their nature, ARIMA models offer very little economic logic; future predictions are formed exclusively from the information contained on past movements and forecast errors. But, despite of being "backward-looking" and not very good at predicting turning points, ARIMA models often outperform multivariate model predictions, particularly in the short term.

A general notation of ARIMA models is ARIMA ( $p, d, q$ ), where  $p$  is the number of autoregressive terms,  $d$  denotes the difference operator (the number of differences needed to convert the series to a stationary level), and  $q$  is the order of moving average terms in the model. As many economic data, including exchange rates, contain changeable and unclear seasonality it could be worthy to test the standard ARIMA representation with additional seasonal dummy variables.

In practice, we have followed these steps to determine the best forecast ARIMA model. We begin by testing the order of integration, whether the data is stationary or if differencing the series is required. Next we run various regressions until the best ARMA combination is identified based on certain diagnostic tests. In this step, particular attention is paid not only to the models' good in-sample characteristics, but also to their forecast performance and to the number of parameters that appear in the model.

Figure 1 plots the USD/EUR exchange rate from January 1994 to December 2014. It can be seen that the euro currency followed a depreciating trend until 2000; it then gradually reversed its losses till 2007; and it has fluctuated around that level thereafter. In another word, the euro exchange rate averaged at about \$1.33 in the first half of this decade, as compared to the much lower level of \$1.19 in the previous 16 years. Meanwhile, its coefficient of variation has been fairly in control at only 4.1 during the post-global crisis period, which is nearly four times lower than the earlier period.

Likewise, the unit root tests indicate that the USD/EUR exchange rate is non-stationary in level. The ADF and Phillips-Perron tests show, however, that the series can become stationary around a constant (and trend) if differenced once. The different test results hold for the whole sample period, as well as for the sub-period starting from the year 2000. Therefore, the euro exchange rate will enter the ARIMA model in the first difference.

Having determined the order of integration,  $I(1)$ , we now try to identify the appropriate AR and MA components in our model based on formal assessments of certain diagnostic checking and the forecasting ability for every competing form. Out of a number of alternative identification methods to determine parameters  $p$  and  $q$  we follow the penalty function criteria, which although offer no theoretical guidelines for choosing the maximum order of ARIMA they are asymptotically consistent and not based on subjective interpretation. Penalty function criteria such as Akaike, Schwarz and Hannan-Quinn statistics help in selecting a model with minimized sum of squared residuals. Nevertheless, our emphasis was finally put on the forecast performance, which suggests more focus on minimizing out-of-sample forecast errors than on maximizing the in-sample 'goodness of fit' – or the adjusted  $R^2$  (Meylar, Kenny, & Quinn, 1998).




Figure 1: The Daily USD-EUR Exchange Rate

Table 1: Unit Root Tests on the USD/EUR exchange rate

Variables	ADF test results				Phillips-Perron test results			
	Null Hypothesis: Unit root							
	Levels		1st difference		Levels		1st difference	
	Prob.	Lag <sup>a</sup>	Prob.	Lag <sup>a</sup>	Prob.	Bandwidth <sup>b</sup>	Prob.	Bandwidth <sup>b</sup>
<u>1994M01:2014M12</u>								
With constant	0.3000	1	0.0000	0	0.3854	5	0.0000	1
With constant and trend	0.5571	1	0.0000	0	0.6542	5	0.0000	1
None	0.6590	1	0.0000	0	0.6672	5	0.0000	1
<u>2000M01:2014M12</u>								
With constant	0.3948	1	0.0000	0	0.4649	4	0.0000	1
With constant and trend	0.6792	1	0.0000	0	0.7898	4	0.0000	1
None	0.7277	1	0.0000	0	0.7443	4	0.0000	2

<sup>a</sup>Automatic selection of lags based on SIC; <sup>b</sup>Newey-West bandwidth selection using Bartlett kernel

In practice, we estimated numerous ARIMA forms with dozens of AR and MA lags and seasonal dummies, which were constantly tested for normal distribution, serial correlation and heteroskedasticity in the residuals. The sample estimation period covered January 1994 to December 2013, whereas observations in 2014 were saved to evaluate the out-of-sample forecast performance. It turned out that overfitting the model with too many parameters increased the in-sample explanatory power but weakened the diagnostic tests and/or the out-of-sample forecasting ability. Thus we eventually retained a more parsimonious model with the smallest possible number of parameters. The constant term as well as the seasonal dummies stayed in the model because, although hardly statistically significant, they were found to help to satisfy the aforementioned criteria. The structure of our preferred model is seasonal ARIMA(7,1,48), estimated as a parsimonious ARMA Conditional Least Squares (Marquardt) method in the following specification:

$$\begin{aligned} I(\text{USDEUR}) &= 0.35 \cdot \text{AR}(1) - 0.16 \cdot \text{AR}(2) + 0.08 \cdot \text{AR}(3) - 0.12 \cdot \text{AR}(7) - 0.28 \cdot \text{MA}(48) + 0.01 \cdot c + \text{seas} \\ \text{Coefs prob.} & \quad 0.00 \quad 0.02 \quad 0.23 \quad 0.07 \quad 0.00 \quad 0.16 \quad 0.36 \\ \text{Adj. } R^2 &= 0.12 \quad \text{S.E. of regression} = 0.0277 \\ \text{Normal distribution: JB test prob.} &= 0.15 \\ \text{Serial correlation: LM test, probability Chi-sq}(12) &= 0.55 \end{aligned}$$

Heteroskedasticity: BPG test, probability Chi-sq(11) = 0.53

where  $\nabla$  represents the change operator;  $c$  indicates the constant; and  $\text{seas}$  stands for the eleven seasonal dummies included in the specification. Most of the selected autoregressive and moving average terms are statistically significant, as shown by the probabilities below each of the coefficients. Also, the diagnostic tests suggest that errors are normally distributed and do not have serial correlation or heteroskedasticity. In what follows, we describe the nonlinear autoregressive model as part of the neural network and then compare the forecasting ability between ARIMA and NAR approaches.

## 2.2 Neural Networks Modeling

This section gives a brief description of the artificial neural networks methodology, starting with basic issues in ANNs for time series forecasting and continuing with the nonlinear autoregressive construction (NAR) to predict the USD/EUR exchange rate. As previously mentioned this technique is a data driven approach, and it is nonparametric in the sense that is not necessary to know any prior information regarding the process that generates the signal. In an ARIMA model, we forecast future observation by using a certain function of past observations. Whereas, a network is trained through general-purpose algorithms based on time-series data and focusing on the computation of weighted neuron connections in a feedforward network to accomplish a desired input-output mapping (Zhang & Hu, 1998).

The common feedforward architecture of a neural network (NN) is organized into several layers of nodes. The first layer is the input layer; then comes the number of nodes in this layer that corresponds to the lagged data observations; and the last layer, also known as output layer are the forecasting values. Between the input and the output layer, we put one or more hidden layers. The layers have unidirectional connection between them (Janil & Mao, 1996), that is, the information must flow from input to output in only one direction with no back-loops. Each connection has a numeric weight which signifies its strength. Many authors have rigorously demonstrated that a three-layer neural network with a logistic activation function in the hidden units is a universal approximation (Gonzalez, Steven; Canada Economic and Fiscal Policy Branch, 2000).

The most challenging task is how to design a network of appropriate size for capturing the underlying patterns in the training data. Ultimately, for a network model to be useful it should have generalization or forecasting capability. Hidden nodes are used to capture the nonlinear structures in a time series. Determination of how many hidden nodes to use is another difficult issue in the ANN model construction process. Since no theoretical basis exists to guide the selection, in practice the number of input and hidden nodes is often chosen through experimentation or by trial-and-error (Zhang, Patuwo, & Hu, 1998). The NN learns by adjusting the weights. Various learning algorithms have been used to train the network. After the nets are trained, for each ANN model we choose the best architecture based on a certain performance criterion.

The general process used for training the network is done in three basic steps:

1. Inputting the training data and the target data
2. Learning the 'rules' from the given data collection
3. Improving the network performance by iteratively adjusting the weights.

The neural networks can be trained to predict future values of the exchange rate by relying on the own past values. But different from the linear estimation in the ARIMA method, the architectural approach to construct our neural network will be based on the "Nonlinear Autoregressive models (NAR)". Our main interest is to systematically examine how the forecasting performance of the neural network is affected by various factors, where a number of input nodes and hidden layer nodes are selected for the experimental process.

As in the ARIMA regression, the NAR analysis uses average monthly data of the USD/EUR exchange rate, spanning from January 1994 to December 2014, a total of 252 observations. The data is divided into two periods: we use the period from January 1994 to December 2013 (240 observations) to estimate and evaluate the model, and retain the twelve months in 2014 (5% of the date) for out-of-sample forecast analysis. Also, the in-sample data used for model estimation and evaluation is divided into another three parts: the training period consists of 70 percent of the total, while the rest is divided equally between the validation (15%) and the testing (15%) periods. The NAR networks are trained for 1,000 epochs. The Levenberg-Marquardt optimization was employed to train these networks, in which the weights and bias values are updated after each epoch in order to find the best configuration of the weights.

To determine the best architecture of the NAR model we examined a number of neurons in the input and the hidden layers. The delayed data used as input neurons in the input layer for our analysis are in the range of one to twelve; the number of neurons in the hidden layer is similarly tried one through twelve; so that 144 different architectures

(12x12) are examined in the process. For each network architecture the training was repeated ten times using different starting values for the weights; they were randomly assigned in order to find the global minimum. Finally, each architectures is evaluated on the basis of their predictive power, and the one with lowest forecast errors in the out-of-sample forecast evaluation is selected.

Table 2 shows the 144 estimated NN structures and their respective one-month ahead forecast performance during 2014. It appears that the NN structure 11—12—1, i.e. involving eleven time lags of the variable and twelve hidden layers, is the finest network in terms of providing the most accurate exchange rate forecast for the next month. Both measures of forecast evaluation, the RMSE and MAE, indicate that the forecast errors derived by the 11—12—1 net are the lowest in comparison to the rest of the NAR networks. The next discussion reviews how this compares with the forecast performance of the ARIMA(7,1,48) approach that we developed above.

**Table 2:** Statistical performance of NAR networks for the USD/EUR exchange rate

NR	NN Structure	RMSE	MAE	NR	NN Structure	RMSE	MAE	NR	NN Structure	RMSE	MAE	NR	NN Structure	RMSE	MAE
1	1-1-1	0.0157	0.0123	37	4-1-1	0.0162	0.0136	73	7-1-1	0.0123	0.0093	109	10-1-1	0.0147	0.0119
2	1-2-1	0.0178	0.0144	38	4-2-1	0.0149	0.0122	74	7-2-1	0.0139	0.0113	110	10-2-1	0.0125	0.0096
3	1-3-1	0.0166	0.0133	39	4-3-1	0.0125	0.0107	75	7-3-1	0.0137	0.0111	111	10-3-1	0.0140	0.0121
4	1-4-1	0.0160	0.0128	40	4-4-1	0.0138	0.0113	76	7-4-1	0.0120	0.0101	112	10-4-1	0.0130	0.0109
5	1-5-1	0.0176	0.0141	41	4-5-1	0.0142	0.0117	77	7-5-1	0.0128	0.0101	113	10-5-1	0.0121	0.0093
6	1-6-1	0.0169	0.0133	42	4-6-1	0.0130	0.0113	78	7-6-1	0.0132	0.0100	114	10-6-1	0.0131	0.0114
7	1-7-1	0.0165	0.0133	43	4-7-1	0.0138	0.0109	79	7-7-1	0.0147	0.0135	115	10-7-1	0.0120	0.0098
8	1-8-1	0.0178	0.0144	44	4-8-1	0.0129	0.0101	80	7-8-1	0.0159	0.0124	116	10-8-1	0.0141	0.0099
9	1-9-1	0.0192	0.0150	45	4-9-1	0.0138	0.0115	81	7-9-1	0.0130	0.0105	117	10-9-1	0.0118	0.0091
10	1-10-1	0.0178	0.0147	46	4-10-1	0.0128	0.0101	82	7-10-1	0.0138	0.0100	118	10-10-1	0.0118	0.0093
11	1-11-1	0.0175	0.0150	47	4-11-1	0.0141	0.0113	83	7-11-1	0.0117	0.0094	119	10-11-1	0.0121	0.0095
12	1-12-1	0.0181	0.0152	48	4-12-1	0.0135	0.0109	84	7-12-1	0.0136	0.0107	120	10-12-1	0.0113	0.0083
13	2-1-1	0.0149	0.0123	49	5-1-1	0.0154	0.0128	85	8-1-1	0.0133	0.0102	121	11-1-1	0.0129	0.0104
14	2-2-1	0.0137	0.0104	50	5-2-1	0.0140	0.0113	86	8-2-1	0.0133	0.0102	122	11-2-1	0.0135	0.0102
15	2-3-1	0.0140	0.0114	51	5-3-1	0.0140	0.0114	87	8-3-1	0.0122	0.0087	123	11-3-1	0.0106	0.0083
16	2-4-1	0.0139	0.0114	52	5-4-1	0.0129	0.0109	88	8-4-1	0.0127	0.0103	124	11-4-1	0.0132	0.0105
17	2-5-1	0.0123	0.0090	53	5-5-1	0.0138	0.0114	89	8-5-1	0.0124	0.0095	125	11-5-1	0.0126	0.0090
18	2-6-1	0.0125	0.0091	54	5-6-1	0.0116	0.0090	90	8-6-1	0.0122	0.0098	126	11-6-1	0.0115	0.0096
19	2-7-1	0.0129	0.0102	55	5-7-1	0.0127	0.0104	91	8-7-1	0.0175	0.0146	127	11-7-1	0.0132	0.0106
20	2-8-1	0.0132	0.0102	56	5-8-1	0.0138	0.0114	92	8-8-1	0.0120	0.0097	128	11-8-1	0.0126	0.0114
21	2-9-1	0.0125	0.0106	57	5-9-1	0.0122	0.0101	93	8-9-1	0.0152	0.0122	129	11-9-1	0.0134	0.0114
22	2-10-1	0.0144	0.0117	58	5-10-1	0.0125	0.0100	94	8-10-1	0.0138	0.0112	130	11-10-1	0.0111	0.0087
23	2-11-1	0.0129	0.0109	59	5-11-1	0.0127	0.0108	95	8-11-1	0.0133	0.0100	131	11-11-1	0.0149	0.0111
24	2-12-1	0.0125	0.0094	60	5-12-1	0.0135	0.0113	96	8-12-1	0.0132	0.0102	<b>132</b>	<b>11-12-1</b>	0.0105	0.0079
25	3-1-1	0.0157	0.0130	61	6-1-1	0.0154	0.0128	97	9-1-1	0.0150	0.0127	133	12-1-1	0.0157	0.0135
26	3-2-1	0.0144	0.0115	62	6-2-1	0.0150	0.0123	98	9-2-1	0.0134	0.0108	134	12-2-1	0.0156	0.0135
27	3-3-1	0.0143	0.0116	63	6-3-1	0.0130	0.0104	99	9-3-1	0.0147	0.0122	135	12-3-1	0.0135	0.0105
28	3-4-1	0.0163	0.0139	64	6-4-1	0.0126	0.0099	100	9-4-1	0.0146	0.0108	136	12-4-1	0.0129	0.0116
29	3-5-1	0.0146	0.0119	65	6-5-1	0.0134	0.0107	101	9-5-1	0.0157	0.0128	137	12-5-1	0.0126	0.0086
30	3-6-1	0.0143	0.0116	66	6-6-1	0.0119	0.0100	102	9-6-1	0.0126	0.0101	138	12-6-1	0.0116	0.0093
31	3-7-1	0.0131	0.0099	67	6-7-1	0.0132	0.0107	103	9-7-1	0.0123	0.0099	139	12-7-1	0.0120	0.0091
32	3-8-1	0.0139	0.0112	68	6-8-1	0.0124	0.0099	104	9-8-1	0.0156	0.0128	140	12-8-1	0.0133	0.0095
33	3-9-1	0.0145	0.0120	69	6-9-1	0.0137	0.0110	105	9-9-1	0.0177	0.0138	141	12-9-1	0.0181	0.0151
34	3-10-1	0.0146	0.0120	70	6-10-1	0.0140	0.0118	106	9-10-1	0.0135	0.0108	142	12-10-1	0.0138	0.0100
35	3-11-1	0.0126	0.0103	71	6-11-1	0.0113	0.0090	107	9-11-1	0.0120	0.0099	143	12-11-1	0.0126	0.0086
36	3-12-1	0.0154	0.0132	72	6-12-1	0.0125	0.0103	108	9-12-1	0.0126	0.0102	144	12-12-1	0.0159	0.0135

### 3. Comparison between Models' Forecast Abilities

Common measures to evaluate and compare the predictive power of a model are the Mean Error, Mean Absolute Error, the Root Mean Squared Error, and the Theil's U statistic. The first measure, ME, hints on the presence and the direction of the bias in the forecasts. A high and positive (negative) ME indicates a tendency to overshoot (undershoot) the actual

exchange rate developments. The next two measures, MAE and RMSE, are useful in examining the size of the forecast errors. The mean absolute error implies that severity of errors increases in proportion to the MAE (e.g. 2 percent error is 2 times as serious a 1 percent error), while RMSE assumes that larger forecast errors are worse than smaller ones. The final measure, Theil's U, is an attractive indicator about relative accuracy that allows us to infer whether the applied forecasting techniques are better, as good as, or worse than simple guessing.

As mentioned before, the forecast evaluation period for the two competing models, ARIMA and NAR net, was purposely done on the out-of sample data. This is a common procedure among forecasters, who presume that in-sample evidence of predictability does not guarantee significant out-of-sample predictability. Therefore to avoid any possibility of spurious conclusions, the comparison of the predictive power of our chosen models will be conducted for the period from January to December 2014, which was not used in the regression estimation or the testing process.

Figure 2 exhibits the USD/EUR exchange rate during 2014 and its 12 one-month ahead forecasts based on the ARIMA (7,1,48) and NAR (11-12-1), as well as a naive prediction that assumes 'no change' in the next month. A quick visual inspection suggests a better performance of the neural network model, since the NAR forecasts appear generally closer than other forecasts values to the actual observations. On the other hand, there is no clear view of the ARIMA predictions to outperform the naive guessing values. Indeed, the ARIMA forecast values tend to alternate their position till July, but remain above actuals (i.e. positive errors) thereafter. This implies that the ARIMA model has provided robust forecasts in a more normal time, but it has consistently overpredicted the USD/EUR exchange rate and was slow to catch up with the new trend in the second half of the year, during which the European Central Bank was rallying the markets with hints on initiating a quantitative easing programme<sup>1</sup>.

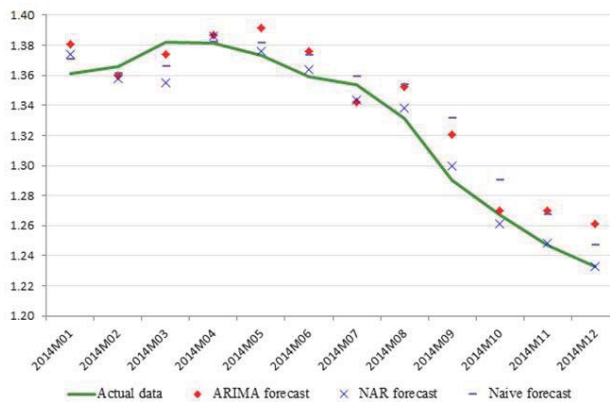


Figure 2: Static forecasts of the USD/EUR exchange rate

Table 3: Evaluation of Static Forecasts of the USD/EUR exchange rate 12 one-month ahead periods, January to December 2014

	ARIMA(7,1,48)	NAR 11-12-1
Mean Error	0.0115	-0.0008
Mean Absolute Error	0.0159	0.0079
Root Mean Squared Error	0.0181	0.0105
Theil's U statistic	0.9921	0.5748

Table 3 presents the forecast evaluation measures for the ARIMA and NAR models. Obviously, the results confirm the superiority of the NAR technique by all measures. Its mean error is close to zero – or more precisely, less than a tenth of a cent – suggesting that NAR has on average neither overshoot, nor undershot the exchange rate in 2014. On the

<sup>1</sup>"Markets rally after Draghi QE hint", article in Financial Times, August 25, 2014. Also see "ECB meeting recap: Draghi hints at more easing action next year", article in blogs.marketwatch.com, December 4, 2014.

contrary, the mean forecast error of the chosen ARIMA model, though moderate, has overpredicted the actuals by 1.15 cent in the same period.

The size of ME in ARIMA is around three-fourth of MAE, indicating the model is not predicting consistently too high. But again, that ratio is only one-tenth for the NAR. Similarly, the magnitudes of MAE and RMSE point out that larger forecast errors are to be taken more seriously in the case of ARIMA, and less for the neural network model.

Finally, the Theil's statistic in the last row of the table shows how both models compare to a naïve prediction (i.e. assuming the exchange rate in the next month will be the same as in the current period). Calculated as the ratio of the RMSE of the ARIMA (NAR) model to the RMSE of the naïve model, a Theil statistic greater than one would suggest our estimated model forecasts perform worse than a random walk model. This statistic for the ARIMA model is close to one, which – likewise the graphical inspection – suggests that predictions generated in 2014 were, on average, no better than the 'no change' method. On the other hand, the NAR predictions appear again to outperform even the naïve model, whose inferences for the next month exchange rate are about twice more erroneous.

#### 4. Concluding Remarks

This article contributes to the neural network literature by demonstrating how potent and useful they can be as a tool in the process of economic and financial decision makings. To demonstrate it we chose the USD/EUR exchange rate, as a considerably volatile and a highly transacted asset in the international financial market, yet very disputed in academic works due to its often large divergences from the fundamental levels suggested by economic theories.

We were able to show that, while through a modest application, neural network models can add value and possibly outperform traditional models used to forecast exchange rates. The results were affirmative that the nonlinear autoregressive net consistently beat the ARIMA (and the random walk) static forecasts of the USD/EUR exchange rate. Nonetheless, in addition to autoregressive techniques and static forecasts, this research analysis can be further extended by employing more advanced models that are based on economic fundamentals as well as dynamic forecasts.

#### References

- Bacchetta, P., & Wincoop, E. v. (2006). Can information heterogeneity explain the exchange rate determination puzzle? *American Economic Review* 96, 552-576.
- Gonzalez, Steven; Canada Economic and Fiscal Policy Branch. (2000). Neural networks for macroeconomic forecasting: a complementary approach to linear regression models.
- Inoue, A., & Kilian, L. (2002). In-sample or out-of-sample tests of predicability: which one should we use? *European Central Bank, Working Paper no. 195*.
- Janil, A. K., & Mao, J. (1996). Artificial Neural Networks: A Tutorial. *IEEE Transactions on Computer*, 29(3 (March)), 31-44; doi:10.1109/2.485891. doi:10.1109/2.485891
- Kilian, L., & Taylor, M. P. (2001). Why is it so difficult to beat the random walk forecast of exchange rates?
- Meylar, A., Kenny, G., & Quinn, T. (1998). Forecasting Irish inflation using ARIMA models. *Munich Personal RePEs Archive (MPRA) paper no 11359 (December)*.
- Philip, A., Taofiki, A., & Bidem, A. (2011). Artificial Neural Network Model for Forecasting Foreign Exchange Rate. 1(No. 3), 110-118.
- Sarno, L., & Taylor, M. P. (2002). *The Economics of Exchange Rates*. Cambridge University Press.
- Zhang, G., & Hu, M. Y. (1998). Neural network forecasting of the British pound/US dollar exchange rate. *Omega* 26.4, 495-506.
- Zhang, G., Patuwo, B., & Hu, M. Y. (1998). Forecasting with artificial neural networks: The state of the art. *International Journal of Forecasting*(14 (1)), 35-62.

