

Forecasting Security Returns With Simple Moving Averages

Camillo Lento, Lakehead University, Canada

ABSTRACT

This study examines the ability of simple moving averages to forecast security returns. Five moving average variants are used to develop a forecasting model using OLS regression for the DJIA, NASDAQ, TSX and CAD-US exchange rate. The forecasting model is compared to the random-walk model without a drift and tested out-of-sample. The results suggest that the moving averages have no predictive ability on the four indices at a 1 day lag. However, the moving averages explain approximately 45% to 48% of the variation in the returns in the following 10 days and clearly outperform the random-walk model. Most of the forecasting ability is derived from the MA (5, 150). Hurst Statistic estimation is used to confirm the long-term dependencies in the lag 10 data set.

Keywords: Technical Analysis; Market Efficiency; Forecasting Stock Returns; Hurst Exponent.

1. INTRODUCTION

Technical analysis is considered to be one of the earliest, if not first, forms of investment analysis with its origins dating back to the 1800s. It was one of the first forms of investment analysis mainly because stock prices and volume levels have been publicly available prior to other types of financial information. After a hundred years in practice, technical analysis has yet to lose its luster as a viable and efficient approach to forecasting security movements.

Technical analysis uses past prices to predict future prices and has been popular among academics and financial analysts despite its many criticisms. Many stock brokerages and major daily newspaper analyze and publish technical commentary on the stock market. Various trading rules and indicators have been developed based on identifiable patterns. However, the belief that historical prices can be used to identify patterns that predict security movements violates the random-walk hypothesis and the weak form of market efficiency. According to efficient market theorists, technical analysis should not be able to generate abnormal returns in an efficient market. Much past research has been conducted to determine whether the market is efficient or whether technical analysis can be used to predict security returns.

The vast majority of the literature focuses on testing trading rule variants in different markets across the globe. One question researchers must address is determining the data frequency to be used for the technical analysis (i.e. intraday, daily, weekly, etc data). As such, there are many studies that test trading rules at different scales but no consensus on the optimal scale.

Another question that is not sufficiently addressed in the literature is how many trading days should be utilized to evaluate a trading signal. For example, if a moving average suggests a long position at t_1 , should the success of the trading rule be evaluated by the holding period return to t_2 , t_3 , or t_{10} ? Therefore, this paper tests the moving averages ability to forecast security returns for the following trading day (lag 1) and the following 10 days (lag 10 or two trading weeks). Therefore, this paper contributes to a more comprehensive understanding of the forecasting ability of technical analysis and the time period to be used to evaluate signals.

OLS regression analysis was used to develop a forecasting model based on the signals from five moving averages. The model attempts to forecast movements in equity markets and the foreign exchange market, namely the Dow Jones Industrial Average, the NASDAQ, the Toronto Stock Exchange, and the Canada-U.S. spot exchange rate. Justification for a linear model is discussed in the methodology section.

Deriving a model that accurately forecasts security returns will deny the weak form of the Efficient Market Hypothesis (EMH). Therefore, this research directly tests the EMH. The first 80% of each data set was used for in-sample model estimation, with the remaining 20% used for out-of-sample evaluation.

The results suggest that the moving averages have no predictive ability on any of the four indices with a lag 1. However, the moving average model is able to explain approximately 45% to 48% of the variation in the market returns at a lag 10. The moving average model results in a higher R^2 and lower mean square error, absolute square error, and root mean square error than the random-walk model. The Hurst exponent was estimated for the lag 1 and lag 10 data sets and confirms the long-term dependencies in the lag 10 data.

The remainder of the paper is organized as follows. Section 2 presents a review of the theory and the hypothesis development. Section 3 describes the data. Section 4 discusses the methodology employed. The empirical results of the study are presented and discussed in section 5. Finally the conclusions are presented in section 6.

2. THEORY AND HYPOTHESIS DEVELOPMENT

2.1 Literature

The results of past studies appear to be inconclusive in assessing whether technical analysis is a useful way of extracting information from historical market prices to forecast future returns. Alexander (1961) and Fama and Blume (1966) were two of the first to test technical trading rules. Both of these studies suggest the same result; excess returns could not be realized by making investment decisions based on movements of certain sizes (filter rules).

Lo and MacKinlay (1988) wrote a seminal paper on technical trading rules that provides evidence that the random-walk hypothesis is strongly rejected by using variance ratio statistics for the sample period 1962 – 1985 and for different sub-periods for a variety of aggregate U.S. indices and for size-sorted portfolio. The major factor leading to the rejection of the random walk was the behaviour of the small stocks.

Brock, Lakonishok, and LeBaron (1992) used bootstrapping techniques and simple, yet popular, trading rules to reveal strong evidence in support of technical analysis's predictive nature. Brock et. al. suggest that the patterns uncovered by technical rules cannot be explained by first order correlations or by the potential for changing expected returns caused by changes in volatility. The profits generated from the technical trading rules are not consistent with a random-walk model.

Further empirical insight and evidence for the profitability of filter rules and MAC-O rules is provided by Levich and Thomas (1993). Consistent with the previous studies, moving average rules and filter rules produced profitable trading signals. The profits from the moving average rules were higher than those of the filter rules.

Gencay (1998) tested the daily Dow Jones Industrial Average Index from 1897 to 1988 for the linear (and nonlinear) predictability of stock market returns with simple technical trading rules. The technical trading rules used in the study were very simple and popular. The moving average rules studied provided at least a 10 percent forecast improvement in volatile years.

The effectiveness of simple trading rules in relation to the size of firm was investigated by Bokhardi et. al. (2005). Bokhardi et. al. segregated a number of companies based on their size and applied trading rules on their past

prices for a sample period of 1987 – 2002. The results indicate that trading rules are more effective at predicting the future price movements of firms with smaller capitalization.

A novel combined signal approach to individual technical trading rules was applied by Lento and Gradojevic (2007). The combined signal approach improved the profitability of individual technical trading rules, thereby suggesting that the individual signals, when combined, result in a more effective trading signal.

2.2 Moving Average Trading Rules

Although there is a host of trading rules, this study focuses solely on moving averages. A moving average cross-over (MAC-O) rule compares a short moving average to a long moving average. The MAC-O rule tries to identify a change in a trend. There are two categories of the MAC-O rule: variable length moving average (VMA) and fixed length moving average (FMA). The FMA stresses that the returns for a few days following the crossing of the moving averages should be abnormal. The VMA generates a buy (sell) signal whenever the short moving average is above (below) the long moving average. Moving averages can provide useful trading information if the difference between a short and long moving averages can forecast the future security returns. The following equation is used to calculate the difference between short and long moving averages:

Equation 1 – VMA Trading Signal

$$\frac{\sum_{s=1}^S R_{i,t}}{S} - \frac{\sum_{l=1}^L R_{i,t-1}}{L} = Buy, \tag{1}$$

where $R_{i,t}$ is the log return given the short period of S, and $R_{i,t-1}$ is the log return over the long period L

2.3 Random-Walk Model

The concept of an efficient market is associated with the idea of a random-walk. A random-walk has traditionally been used in finance literature to characterize a price series where all subsequent price changes represent random departures from previous prices.

The logic of the random-walk idea is that if the flow of information is unimpeded and information is immediately reflected in stock prices, then tomorrow’s price change will reflect only tomorrow’s news and will be independent of the price changes today. But news is by definition unpredictable, and, thus, resulting price changes must be unpredictable and random. As a result, prices fully reflect all known information (Malkeil 2003).

The following equation describes a time series of security returns where each current value (S_t) depends on an earlier value (S_{t-1}) plus a random error term (u_t).

Equation 2 – Random-Walk Model

$$r_t = \Phi r_{t-1} + u_t, \tag{2}$$

When the parameter $\Phi = 1$, the security prices evolves as a random-walk without a drift. Expressed differently, the change in security price is essentially the error term.

2.4 Hypothesis Development

The efficient market theory suggests that technical trading rules, and specifically the moving-average rules in this study, should not have the ability to forecast security prices. However, past research on technical analysis is mixed at best. Therefore, this study aims to provide further research in the area of technical analysis by testing the following hypothesis:

H₁: Moving averages can be used to develop a forecasting model that outperforms the random-walk model.

Another question that is not sufficiently addressed in the literature is how many trading days should be utilized to evaluate a trading signal. For example, if a moving average suggests a long position at t_1 , should the success of the trading rule be evaluated by the resulting results at t_2 , t_3 , or t_{10} ? This research will provide insight into the optimal forecasting assessment period given that this paper tests the moving averages ability to forecast security returns for the following trading day (lag 1) and the following 10 days (lag 10).

3. DATA

The investigations are based on the daily closing prices of the Dow Jones Industrial Average, S&P 500, Toronto Stock Exchange, and Canada-U.S. spot exchange rate.

Trading rules can be calculated at various data frequencies. Investors can use high-frequency data (e.g. intra-day data) or longer horizons (e.g. weekly or yearly) as inputs for the trading rules. The data frequency selected by a technical investor depends on many different factors and personal preferences. Daily data will be used because a typical off floor trader will most likely use daily data (Kaastra and Boyd, 1996). Furthermore, intraday time series can be extremely noisy. The number of daily observations in each data set provides a sufficient number of daily observations to allow for the formation, recurrence and investigation of the moving average signals.

The use of daily data in the stock and foreign exchange market has many problems as movements are generally non-stationary and quite random in nature and therefore not suitable for forecasting purposes (Mehta, 1995). The market index series are transformed into rates of return to overcome these problems. Given the price level $P_1, P_2, \dots P_t$, the rate of return at time t is transformed by:

Equation 3 – Daily Holding Period Return

$$r_t = \log(p_t) - \log(p_{t-1}) \tag{3}$$

where p_t denotes the market price. Another advantages of using a return series is that it helps in making the time series stationary (Dunis and Williams, 2003). Formal confirmation that the time series are stationary is confirmed at the 10% level of significance by the Augmented Dickey-Fuller (ADF) test statistic.

All of the time series range from 1995 to 2008, approximately thirteen years. The data is divided into two periods. The first period is used for model estimation and is classified as in-sample. The second period is reserved for out-of-sample forecasting and model evaluation. The exact dates and observations for both periods are presented in Table 1. Table 1 also presents the descriptive statistics of the data sets, highlighting the skewness and kurtosis.

The Jacque-Bera statistic confirms that the data sets are not normal at the 95% confidence interval. Such features are common in high frequency financial time series data (Gencay, 1998).

Table 1 – Descriptive Statistics

Data Set	n	Mean	Std. Dev.	Skew.	Kurt.	Period Tested	In-sample	Out-of-sample	Jacque-Bera stat p-value	Augmented Dickey-Fuller stat
DJIA_lag 10	3,307	0.00264	0.03193	-0.5564	3.115	15/05/95 – 02/03/08	2,430	607	0.0000	-12.4933
DJIA_lag 1	3,307	0.00027	0.01087	-0.2267	3.933	15/05/95 – 02/03/08	2,430	607	0.0000	-19.7495
NAQ_lag 10	2,812	0.00264	0.05332	-0.5535	2.390	23/02/96 – 02/04/08	2,250	562	0.0000	-13.2139
NAQ_lag 1	2,812	0.00027	0.01766	0.0354	4.065	23/02/96 – 02/04/08	2,250	562	0.0000	-39.1629
TSX_lag 10	2,940	0.00340	0.03174	-0.6891	1.736	28/02/96 – 31/12/07	2,352	588	0.0000	-13.1014
TSX_lag 1	2,940	0.00035	0.0099	-0.6787	5.393	28/02/96 – 31/12/07	2,352	588	0.0000	-37.9414
CAD-US lag 10	2,220	-0.00054	0.0122	-0.0683	0.813	27/02/96 – 31/12/04	1,776	444	0.0000	-10.7125
CAD-US lag 1	2,220	-0.00006	0.0039	-0.0409	1.688	27/02/96 – 31/12/04	1,776	444	0.0000	-33.1419

4. RESEARCH DESIGN

4.1 Methodology

This research study utilizes a combination of technical trading rules and the daily closing prices of the DJIA, Nasdaq, TSX, and the Canada-U.S. spot exchange rate. It is important to note that the purpose of this study is not to determine the profitability of technical trading rules, but rather to evaluate the forecasting effectiveness of trading rules. There are significant differences in the methodology although the differences may appear to be subtle. Since we are evaluating the forecasting ability of trading rules, as opposed to determining their profitability, the methodology does not consider trading costs and the bid-ask spread.

An Ordinary Least Squares (OLS) multiple regression will be used to test the forecasting ability of the technical trading rules. Although many research studies have argued the benefits of using non-linear statistical techniques, a linear technique is tested for the following two reasons. Firstly, the data inputs used in this study are by definition non-linear. The data inputs include the daily closing prices, the log percentage return, and the technical trading rules. The non-linear dynamics of the time series and the trading rules is accounted for within the data inputs as opposed to the statistical techniques utilized. Therefore, the use of a non-linear model is justified. Secondly, there are a number of research studies that have utilized linear models to test the effectiveness of technical trading rules. The results of these studies are mixed. By utilizing a linear model, the results of this study can then be compared to past studies to provide further evidence on the effectiveness of linear models. The study will serve to further corroborate the notion that non-linear techniques may be superior if this linear model is not effective.

4.2 Forecasting Accuracy

The moving average model will be compared to the random-walk model based on the resulting R^2 and the out of sample forecasting error estimates. The first 80% of each data set will be used for in-sample model estimation, while the remaining 20% of each data set will be used for out-of-sample forecasting accuracy estimation. Several measurements have been developed to make comparison between the forecasting ability of the benchmark and the trading rule model. The forecasting performance will be determined by the Mean Squared Error, Root Mean Squared Error, and Mean Absolute Error statistic. Refer to Hanke and Reitsch (1998) for a full discussion of these and other performance measures.

4.3 Empirical Model

The following specification is employed to examine the forecasting ability of the trading rules on the market indices:

$$r' = \alpha + \beta_1 \text{MATR}_1 + \beta_2 \text{MATR}_2 + \dots + \beta_n X_n$$

Where: r' is the forecasted security price,

α is the intercept term

β_1 is the partial regression coefficient

MATR is the technical trading signal

Five different moving average rules will be included in the empirical model. Each moving average rule is defined by the short (S) and long (L) moving average as follow: MA (S, L). The moving average rules are the (1, 50), (1, 150), (5, 150), (1, 200), and (2, 200). These moving averages (independent variables) are utilized in an attempt to generate a model that can forecast security returns (dependent variable) as measured by the combined R^2 .

The theory would expect that the trading rules will be positively correlated to the log returns at the lag 1 and lag 10. For example, the difference will be positive if the short moving average is greater than the long moving average. This will also suggests a long position whereby the log return should be positive. Table 2 presents the Pearson correlation matrix for the time series returns and the moving-average rules.

Table 2 – Correlations Matrix

DJIA - Pearson Correlation Matrix							
Returns							
	Lag 1	Lag 10	MA(1,50)	MA(1,150)	MA(5,150)	MA(1,200)	MA(2,200)
Returns	1.0000	1.0000					
MA (1,50)	-0.0140	0.2571	1.0000				
MA (1,150)	-0.0169	0.2825	0.9952	1.0000			
MA (5, 150)	-0.0330	0.6567	0.4163	0.4363	1.0000		
MA (1,200)	-0.0181	0.2834	0.9949	0.9994	0.4349	1.0000	
MA (2,200)	-0.0315	0.3980	0.6879	0.6982	0.6140	0.6990	1.0000

NASDAQ - Pearson Correlation Matrix							
Returns							
	Lag 1	Lag 10	MA(1,50)	MA(1,150)	MA(5,150)	MA(1,200)	MA(2,200)
Returns	1.0000	1.0000					
MA (1,50)	0.0046	0.2527	1.0000				
MA (1,150)	0.0036	0.2847	0.9936	1.0000			
MA (5, 150)	-0.0274	0.6506	0.4105	0.4344	1.0000		
MA (1,200)	0.0032	0.2886	0.9931	0.9992	0.4343	1.0000	
MA (2,200)	-0.0321	0.4042	0.6926	0.7049	0.6136	0.7067	1.0000

TSX - Pearson Correlation Matrix							
Returns							
	Lag 1	Lag 10	MA(1,50)	MA(1,150)	MA(5,150)	MA(1,200)	MA(2,200)
Returns	1.0000	1.0000					
MA (1,50)	0.06558	0.2621	1.0000				
MA (1,150)	0.06612	0.3033	0.9918	1.0000			
MA (5, 150)	0.00083	0.6628	0.4019	0.4323	1.0000		
MA (1,200)	0.06599	0.3072	0.9912	0.9990	0.4316	1.0000	
MA (2,200)	0.03046	0.4171	0.7111	0.7266	0.6247	0.7284	1.0000

CAD-US XRate - Pearson Correlation Matrix							
Returns							
	Lag 1	Lag 10	MA(1,50)	MA(1,150)	MA(5,150)	MA(1,200)	MA(2,200)
Returns	1.0000	1.0000					
MA (1,50)	-0.0437	0.2464	1.0000				
MA (1,150)	-0.0435	0.2818	0.9941	1.0000			
MA (5, 150)	-0.0109	0.6676	0.4234	0.4478	1.0000		
MA (1,200)	-0.0433	0.2853	0.9932	0.9993	0.4470	1.0000	
MA (2,200)	-0.0102	0.4171	0.6746	0.6878	0.6397	0.6894	1.0000

The Pearson correlation matrix reveals that all five moving average variants are positive and significant across all four data sets at the 5 percent level of significance for the lag of 10. This implies that the larger the moving average difference as calculated in Equation (1), the larger the log holding period return in the following ten days.

The correlation analysis results for the lag 1 are not consistent with the theory and vary significantly across the data set. All of the moving averages are negatively correlated with the log return on the Dow Jones Industrial Average and Canada-U.S. spot exchange rate. However, all of the moving averages are positively correlated to the log return on the Toronto Stock Exchange while three of the five moving averages are positive and significant on the NASDAQ.

Table 3 – In-sample Forecasting Model Estimation: All MA Variables

	<i>DJIA</i>		<i>NASDAQ</i>		<i>TSX</i>		<i>CAD/US</i>	
	(lag 1)	(lag 10)	(lag 1)	(lag 10)	(lag 1)	(lag 10)	(lag 1)	(lag 10)
Coefficient (p-value)								
MA (1,50)	0.2792 (0.181)	-3.8200* (0.0000)	0.0531 (0.781)	-4.0958* (0.000)	-0.1196 (0.459)	-4.6233* (0.000)	0.0207 (0.930)	-5.4022* (0.000)
MA (1,150)	0.8966 (0.146)	-1.4361 (0.2904)	0.1639 (0.770)	-4.8515* (0.000)	0.3119 (0.506)	-2.7821 (0.010)	0.4720 (0.512)	-1.4149 (0.372)
MA (5, 150)	-0.0368 (0.428)	4.4350* (0.0000)	-0.0235 (0.710)	4.3834* (0.000)	-0.0760 (0.196)	4.2992* (0.000)	-0.0882 (0.216)	4.3381* (0.000)
MA (1,200)	-0.0426 (0.484)	5.2801* (0.0000)	-0.1606 (0.767)	8.9620* (0.000)	-0.0816 (0.857)	7.5512* (0.000)	-0.5205 (0.451)	6.8021* (0.000)
MA (2,200)	-1.1551 (0.053)	-0.0603 (0.5559)	-0.0880 (0.068)	-0.0247 (0.819)	-0.0323 (0.498)	-0.1391 (0.211)	0.0677 (0.219)	-0.0069 (0.954)
Intercept	0.0002	0.0029	0.0003	0.0032	0.0003	0.0029	0.0000	0.0006
Adj. R ²	0.00134	0.452	0.0026	0.456	0.0059	0.478	0.0015	0.477
N	2,430	2,430	2,250	2,250	2,352	2,352	1,776	1,776
F-stat	1.655	402.8*	1.189	378.9*	3.8	432.2*	0.5	325.8*
Random Walk Model								
Adj. R ²	0.0003	0.0003	-0.0003	0.0001	0.0041	-0.0003	0.0013	-0.0004
N	3,306	3,306	2,811	2,811	2,939	2,939	2,219	2,219
Hurst Statistic	0.527	0.673	0.599	0.729	0.584	0.746	0.584	0.714

* Significant at the 1 percent level.

5. RESULTS

Table 3 reports the results of the least squares estimation of the moving average model specifications. The dependent variable is the log return of each index at a lag of one and ten days and the explanatory variables are the five moving average variants.

The regression results that accounts for all the considered explanatory variables indicates that the MA (1,50), MA (5,150), and MA (1,200) are the only estimates statistically different from zero across all four data sets at a lag of 10. However, the sign of the MA (1,50) estimate is negative and therefore contradicts the theory and previous evidence on the explanatory power of the moving average. The R^2 for the trading models are in excess of 45 percent for all four data sets, far exceeding the random-walk model. Given this explanatory power, it is very likely that the MA (1,50) coefficient is negative due to multicollinearity (i.e. the coefficient may not be precise but the model as a whole is accurate). In addition, the p-value supports the notion that the MA (1,50) variable is informative for forecasting security returns.

The regression results at the lag of one day do not support to the notion that moving averages are effective at forecasting security returns. The R^2 for the model that considered all explanatory variables indicates that the model has little explanatory power. Furthermore, none of individual moving average estimates are statistically different from zero across all four data sets.

The Hurst statistic (H) was estimated for the four data sets at both the lag 1 and lag 10. The Hurst statistic is a measure of a time series' long run dependencies. A H of 0.50 indicates a series is random. A $0 < H < 0.5$ indicates an anti-persistent series suggesting the data set exhibits mean reverting tendencies (a positive return is followed by a negative return). A $0.5 < H < 1$ indicates a persistent series whereby the data is trend reinforcing (a positive return is followed by a positive return). The strength of the trend increases as H approaches 1. Therefore, technical analysis should be more successful on data sets that exhibit trend reinforcing characteristics or $0.5 < H < 1$.

Table 3 presents the H for all eight data sets and reveals that the lag 10 data sets have a H in the range of 0.67 to 0.72, whereas the lag 1 data sets have a H in the range of 0.52 to 0.59. The Hurst statistic estimation corroborates the regression results that suggest the forecasting ability of the trading rules is much better on the lag 10 data than the lag 1 data.

Based on the results of the regression for a lag of ten days, a new model was developed by utilizing only the MA (1,50), MA (5,150), and MA (1,200), which were the estimates statistically different from zero in the five variable model. The purpose of this refined model is to reduce the potential impacts of multicollinearity. Although the regression suggests that the MA (2, 200) should also be included in the model, the correlation matrix reveals that it is highly correlated with the MA (5, 150) which is the most influential variable; therefore, it can be dropped as it does not provide much further explanatory power in a multiple regression model that includes the MA (5, 150). Table 4 presents the results of the three variable model.

The result of the model suggest that the MA (2, 200) and MA (1, 150) do not provide any additional explanatory powers as the R^2 in the three variable model is virtually unchanged from the model that takes into consideration all five moving averages. The three variable model also far exceeds the forecasting ability of the random-walk model. All three estimates are significant at the 1 percent level of significance; however, the coefficient for the MA (1, 50) remains negative.

Table 4 – In-sample Forecasting Model Estimation: Effective MA Variables

Coefficient (p-value)	DJIA (lag 10)	NASDAQ (lag 10)	TSX (lag 10)	CAD-US (lag 10)
MA (1,50)	-3.944* (0.0000)	-4.528* (0.000)	-4.872* (0.000)	-5.536* (0.000)
MA (5, 150)	4.383* (0.0000)	4.337* (0.000)	4.183* (0.000)	4.322* (0.000)
MA (1,200)	3.947* (0.0000)	4.543* (0.000)	4.965* (0.000)	5.521* (0.000)
Intercept	0.0029	0.0032	0.0030	0.0006
Adj. R ²	0.452	0.453	0.447	0.4781
n	2,430	2,250	2,325	1,776
Durbin-Watson	0.72	0.70	0.66	0.75
F-stat	671.02*	622.91*	716.18*	543.19*
Random Walk Model				
Adj. R ²	0.0003	0.0001	-0.0003	-0.0004
n	3,306	2,811	2,939	2,219

* Significant at the 1 percent level.

Table 5 – Out-of-sample Forecasting Results

	DJIA (lag 10)		NASDAQ (lag 10)		TSX (lag 10)		CAD-US (lag 10)	
	RW	MA	RW	MA	RW	MA	RW	MA
Mean Squared Error (MSE)	0.00048	0.00031	0.00083	0.00047	0.00064	0.00036	0.00029	0.00016
Root Mean Squared Error (RMSE)	0.02212	0.01762	0.02893	0.02184	0.02536	0.01906	0.01717	0.01277
Mean Absolute Error (MAE)	0.01680	0.01350	0.02136	0.01661	0.01966	0.01504	0.01396	0.01013

Table 5 presents the results of the out-of-sample forecasting assessment. The Mean Squared Error, Root Mean Squared Error, and Mean Absolute Error are calculated for the three moving average model and for the random-walk model. All three out-of-sample forecasting error statistics suggest that the moving average model is a more accurate forecasting model than the random-walk model.

6. DISCUSSION AND CONCLUSION

The results of this study are somewhat mixed, suggesting that moving averages are clearly able to outperform the random-walk model at a lag 10 but not at the lag 1. The resulting R^2 of the moving average model at the lag 10 ranges from 44.7 percent to 47.8 percent on the four data sets, whereas the random-walk model's resulting R^2 is negligible. However, there is no evidence that moving averages can outperform the random-walk model at the lag 1.

The results at a lag of one day may be limited by the fact a linear model was utilized. Gencay (1999) cites many previous studies to support his conclusion that high-order movements of a foreign exchange rate series demonstrate the existent of non-linearities.

The MA (1,50), MA (5,150), and MA (1,200) were individually significant across all four data sets and therefore the study suggests that these trading rules can be used in future research to develop a more robust and comprehensive model that may also incorporate trading rules other than moving averages.

The results also suggest that determining the optimal time period to evaluate technical trading rules is key, yet complex. For example, the trading rules were not successful at a lag of one day but provided useful information for forecasting the holding period returns resulting in the ten days following a signal. These results are consistent with Brock, Lakonishok, and LeBaron (1992) who also suggest that moving averages provide useful trading information at a lag of 10 days. Future research should continue to explore the optimal evaluation period.

Overall, the results of this study contribute to our understanding of the forecasting ability of simple moving averages as the MA (1,50), MA (5,150), and MA (1,200) model was able to explain a significant portion of the variation in the DJIA, NASDAQ, TSX, and Canada-U.S. spot exchange rate returns.

AUTHOR INFORMATION

Camillo Lento is an Accounting Lecturer in the Faculty of Business Administration at Lakehead University. Camillo is a Chartered Accountant (Canada, Ontario) and holds a Master of Science in Management Degree and an Honours Bachelor of Commerce Degree (majors in accounting and finance). Camillo is currently working towards a PhD and has published in various journals such as the Journal of Applied Business Research and Applied Economics Letters. His research interests include earnings management, technical analysis, and fair value accounting.

REFERENCES

1. Alexander, S. 1961. "Price Movements in Speculative Markets: Trends or Random Walks" *Industrial Management Review*, Vol. 2, pp. 7 – 26.
2. Bokhari, J., Cai, C., Hudson, R., and Keasey, K. 2005. "The Predictive Ability and Profitability of Technical Trading Rules: Does Company Size Matter?" *Economics Letters*, Vol. 86, pp. 21 – 27.
3. Brock, W., J. Lakonishok, and B. LeBaron. 1992. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance*, 47 (Dec. 1992), 1931 – 1764.
4. Dunis, C. and Williams, M. 2003. "Applications of Advanced Regression Analysis for Trading and Investment", in C. Dunis, J. Laws and P. Naïm [eds.], *Applied Quantitative Methods for Trading and Investment*, John Wiley.
5. Fama, E. and Blume, M. 1966. "Filter Tests and Stock Market Trading" *Journal of Business*, Vol. 39, pp. 226 – 241.

6. Gencay, R. 1998. "The Predictability of Security Returns with Simple Technical Trading Rules." *Journal of Empirical Finance* 5 (1998) 347 – 359.
7. Gencay, R. 1999. "Linear, Non-Linear and Essential Foreign Exchange Rate Prediction with Simple Technical Trading Rules." *Journal of International Economics* 47 (1999) 91 – 107.
8. Hanke, J.E. & Reitsch, A.G. 1998. "*Business forecasting*" (6th ed.). Englewood Cliffs: Prentice-Hall.
9. Kaastra, I. and M. Boyd 1996. "Designing a neural network for forecasting financial and economic time series." *Neurocomputing* 10(3): 215-236.
10. Lento, C., Gradojevic, N. 2007. The profitability of technical trading rules: a combined signal approach. *Journal of Applied Business Research*, 23(1), 13-27.
11. Levich, R. and L. Thomas. 1993. "The Significance of Technical Trading Rules Profits in the Foreign Exchange Market: A Bootstrap Approach." *Journal of International Money and Finance*, 12 (Oct. 1993), 451 - 474.
12. Lo, A. W. and MacKinlay, A. 1988. "Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test." *Review of Financial Studies*, 1, 41 – 66.
13. Malkiel, B. 2003. "The Efficient Markey Hypothesis and its Critics" *Journal of Economic Perspectives*. Volume 17, Number 1—Winter 2003—Pages 59 – 82
14. Mehta, M. 1995. "Foreign Exchange Markets", 176-198, in A. N. Refenes [ed.], *Neural Networks in the Capital Markets*, John Wiley, Chichester.

NOTES