

**An Evaluation of the Effectiveness of Momentum
Strategies in Predicting Future Price Movements,
(Performed on NASDAQ and NYSE)**

by

Mohammad Amin Vassefi
Peyman Ajabi Naeini

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
MASTER OF SCIENCE IN FINANCE

In the Faculty of Business Administration
© M. A. Vassefi & P. A. Naeini, 2012

SIMON FRASER UNIVERSITY

SUMMER 2012

All rights reserved. However, in accordance with the *Copyright Act of Canada*, this work may be reproduced, without authorization, under the conditions for *Fair Dealing*. Therefore, limited reproduction of this work for the purposes of private study, research, criticism, review and news reporting is likely to be in accordance with the law, particularly if cited appropriately.

APPROVAL

Name: Mohammad Amin Vassefi
Peyman Ajabi Naeini

Degree: Master of Science in Finance

Title of Thesis: An Evaluation of Momentum Strategies in
Predicting Future Price Movement

Supervisory Committee”

Andrey Pavlov
Senior Supervisor
Associate Professor

Phil Goddard
Secondary Supervisor
Professor

Date Defended/Approved:

Table of Content

TABLE OF CONTENT	3
ABSTRACT	4
INTRODUCTION	5
LITERATURE REVIEW	6
DATA	8
METHODOLOGY	10
PHASE 1: STOCK SELECTION.....	10
PHASE 2: MODEL TRAINING	11
PHASE 3: PORTFOLIO CREATION	15
PHASE 4: TESTING.....	16
THE METHODOLOGY PROCESS	18
RESULTS	19
CONCLUSION	21
AREAS FOR FURTHER RESEARCH	21
REFERENCES	23
APPENDIX	25
DEFINITIONS	25

Abstract

The purpose of this study is to evaluate the predictive power of ARMA/GARCH models through the implementation of a momentum strategy on all stocks traded on the New York Stock Exchange (NYSE) and the NASDAQ stock market. The data series is tested for serial-correlation in their daily stock price returns, followed by several screening and filtering phases. A floating order ARMA/GARCH model is used to capture the signal from the noise in the data, which is used to forecast future prices. It is shown that the tickers that are successfully predictable carry their momentum into the short-term future. A trading strategy is then proposed and tested to validate the above market returns resulted from this predictive behavior of the tickers.

Introduction

Designing and developing tools to predict future prices or returns on financial securities has been a major topic for research for both practical and academic purposes. Such work is witnessed and examined in fields of economics and finance. Many researchers, economists and financiers have tried to demonstrate whether or not it is even possible to forecast and model future returns. Some believe that equity returns are governed by the random walk hypothesis and thus are unpredictable. Others, however, disagree and have utilized various methods and technologies, which supposedly has allowed them to gain future price information. Development of such predictive tools is important due to the simple fact that an accurate model would result in profitable decision-making. Although academics and industry professionals have differing views on market predictability, one thing is for certain; that is the concept of market efficiency, predictability of stock returns and related investment strategies have been one of the main topics in the financial industry.

The initial research done by Fama in 1970 concluded that markets are generally efficient meaning that share prices reflect all available information, thus making it virtually impossible for one to earn excess return. More specifically it was stated that share prices exhibited no serial correlation and that there are no patterns in asset prices (Fama, 1970). This implied that future price movements were determined entirely by information not contained in the price series and therefore followed a random walk. Thus, preventing one from earning excess returns based on historical share prices or returns. However, more recent articles by Lo and Mackinley (1988) concluded otherwise and reported various forms of serial correlations in weekly returns. This was important as the serial correlation in the data could be utilized to forecast future returns; a conclusion that went against the Market Efficiency Hypothesis (MEH) that was originally put forward by Fama.

Yet another example of work against MEH is of Schiller's (1984, 2003), who examined the performance of the U.S stock market since the 1920s, and concluded that the volatility of the stock market is greater than could possibly be explained by any rational view of the future. Other successful investors and industry professional such as Warren Buffet and Peter Lynch have also supported the inefficiency and referring to their successful careers as evidence. While there are different arguments and approaches in support and against the MEH and its ability to forecast future returns, one thing is clear; that is if there exists some form of correlation (serial correlation) between the asset price and some other variable such as trading volume or the previous day's closing prices or such, one can predict future outcomes.

While there are various prediction methodologies, including fundamental analysis, technical analysis and technological methods, this paper will focus on the technical analysis methodology to further examine market predictability. Additionally, the paper seeks to determine the future returns of stocks based solely on the trends of historical prices, an approach referred to as time series analysis. More specifically, the study will utilize the concept of momentum investing and strategies as it applies to finance to provide further analysis on the concept of market efficiency. Additionally, the model will impose several filtering criteria, followed by a trading strategy to demonstrate the ability of above-mentioned models in predicting the direction of change in prices in short term horizons. So, the thesis objective is to evaluate the predictive power of ARMA/GARCH models through the implementation of momentum strategy.

Literature Review

Over the past several decades substantial amount of academic and professional articles have been published in relation to momentum strategies and return predictability. Momentum, which refers to the predictable patterns of stock returns states that stocks with an above average return in previous time periods have the tendency to outperform

other stocks in subsequent time periods. Such strategies involve buying stocks that have performed well and selling those that have underperformed previously.

Lee and Swaminathan (2000) used data from all companies traded on the NYSE and AMEX to perform an empirical study. They discovered that the price momentum effect finally reverses and the timing is foreseeable based on trading volume. According to their research, past trading “predicts both the scale and the persistence of future price momentum”. It was concluded that stocks that have gone up accompanied with high volume experience faster momentum and reversals than low volume losers.

Connolly and Stivers (2003) have found significant amount of momentum in consecutive weekly returns when the latter week has unexpectedly high turnover. They also found reversals in consecutive weekly returns when the latter week had an unexpected low turnover. Their research was based on empirical studies of weekly returns of large and small firm portfolio in the U.S., Japanese and U.K. It was reported that the first-order autoregressive coefficient “increases around 0.80 as the turnover shock moves from its 5th to its 95th percentile”

Substantial evidence also indicates that past returns over periods shorter than one year are useful for forecasting future returns (Dijk et al). Jagadeesh and Titman (2001) reported that over a period of 3-12 months, on average, past winners are future winners. Likewise, past losers are the loser of the next period. Rouwenhorst (1999) who also performed a similar research to that of Jegadeesh and Titmat on the European market also reported comparable findings.

With regards to other financial instruments, Jacobs (2000) argued that option-replication strategies are a potential cause of price momentum. That is, investors buy stocks as prices increase to capture the upside of price momentum and sell as they go down to limit downside risk. This is exemplified as portfolio insurance that attempts to replicate a protective put option offering upside gains with a floor against downside losses.

Data

The data used in this model is based on daily closing prices of stocks. The time series record goes from December 31, 2008 to July 1, 2012. After converting the prices to returns and analyzing the data for serial correlation, the tickers that pass the tests are listed from both NASDAQ and NYSE. Then they are sorted by market cap.

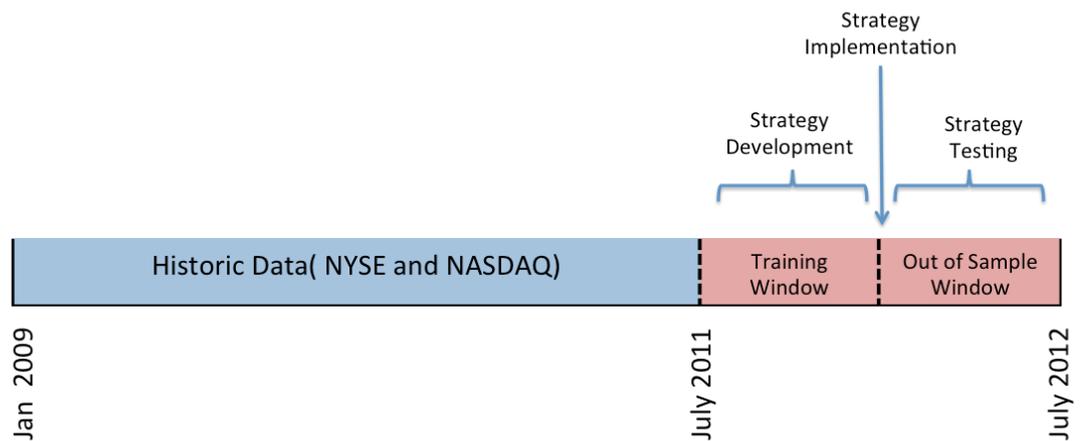
A list of 2754 tickers traded in NASDAQ and 3245 tickers traded in NYSE are passed to an in-house developed MATLAB code, which reads the ticker prices from Yahoo Finance, calculates the returns, and checks each one for serial correlation. The tickers, which pass the tests, are stored in a new list. These tickers include 193 companies from NASDAQ and 33 companies from the NYSE. The market cap ranges from \$19.6 billion to 24.8 million for NYSE, and from \$137 billion to \$1.6 million dollars in NASDAQ.

Sorted by market capitalization, 30 stocks (15 from each Exchange) are selected from the top of the list, due to the fact that finding serial correlation in more liquid tickers is more valuable. Market cap has positive correlation with liquidity, and can be used as a proxy. According to Agrawal and Clark (2009) there is a 5-factor model to determine how liquid a ticker or an ETF is, and market cap is one of those factors. Since they have reported a positive correlation between market cap and liquidity, this paper uses this measure as our criteria for selection of stocks that are “liquid enough”. For tickers that are not liquid, there is a high spread between bid and ask prices, which can result in a false observation of serial correlation.

Noteworthy is the fact that the model uses data following the 2008 financial crises. This helped us not to include data from during and before the crisis, since we had access to 3 years of data, and a rolling 240 days (about a year of working days) to back-test the model. The trends and correlations that is found represent the after crisis global economy, and is not a result of analysis of data which is fragmented in time. The results are

therefore more consistent with recent patterns and are considered more reliable to use for forecasting and predictability purposes.

The figure below demonstrates the data used for validating and testing the model, in a timeline format.



Historic Data: Historical prices/returns used to forecast and test the model

Training Window: Period in which the model tests potential company for presence of recent momentum, which has resulted in a successful forecast.

Out of Sample Window: Period in which the model is testing the effectiveness of the proposed strategy.

Figure 1: Demonstration data usage and steps in forecasting,

Methodology

The following section will present how the study is conducted. For simplification purposes this segment is divided into four phases: Stock Selection, Model Training, Portfolio Creation, and Testing. To demonstrate more effectively, an example of work performed in the actual model is presented in each sub-section.

Stock Selection	Model Training	Portfolio Creation	Model Testing
<ul style="list-style-type: none">• Search NYSE & NASDAQ• For Serial Correlation• LBQ Test• Auto-Correlation Test	<ul style="list-style-type: none">• ARMA/GARCH• Test different orders• AIC: best fit• Forecast future prices• Calculate Success Rate	<ul style="list-style-type: none">• Strategy Development• Choose Prospective Winners: Predictable Tickers• Strategy Implementation	<ul style="list-style-type: none">• Take Same Steps as Model Training Phase• Calculate Success Rate• Propose Trading Strategy• Expected Return• Compare to benchmark

Figure 2: Methodology phases

Phase 1: Stock Selection

Implied by the name, the main purpose of this phase is to filter and create a list of stock tickers that have serial correlation with a 95% confidence. As previously mentioned, serial correlation refers to the relationship between a given variable (stock return) and itself over various time intervals. In order to identify such tickers, a MATLAB code is developed which examines all companies in the NYSE and NASDAQ.

The model performs two major tests for each company: the Ljung-Box Test and Autocorrelation test. Ljung-Box Test, which is a type of statistical test, indicates whether any of a group of autocorrelations of a time series is significantly different from zero. Similarly, Autocorrelation

test finds the points that are out of bound in an “autocorr graph”. If both of these tests support serial correlation, the company’s name is passed to the list of prospects. One condition worth noting is that since penny-stocks may behave differently from the norm, the model has incorporated a 4 dollars minimum price limitation on the prospects to filter such abnormalities. So, after a thorough search in all the listed companies in both exchanges, a list of tickers with a potential for momentum strategy is at hand.

A sample of 30 tickers is selected for the purpose of this study. This list of prospects is then sorted by market cap in both markets. Market cap is one of the indicators of liquidity, so it’s assumed that finding serial correlation and proving the effectiveness of an alpha-generating strategy on high market cap prospects, would be more difficult, and also more valuable in terms of costs of trade and bid-ask spread. Following that, the model selects the top 15 tickers from both exchanges. These tickers are then passed to the MATLAB code to be evaluated for the best orders of the fit and to be tested in two consecutive 120 day windows. Below chart illustrates a list 6 stocks that have passed the stock selection phase and therefore have serial correlation.

Stock Selection: Serially Correlated Companies			
Company Ticker	LBQ Test P-Value	Market Cap (\$)	Serial Correlation
MKTX (NASDAQ)	4.1033e-6	1.22B	PASS
AWH (NYSE)	1.8939e-4	7.53B	PASS
CNK (NYSE)	0.0084	2.75B	PASS
TLAB (NASDAQ)	0.0331	1.23B	PASS
FFIN (NASDAQ)	9.6405e-12	1.11B	PASS
CATO (NYSE)	0.0015	864.97M	PASS

Figure 3: Stock Selection

Phase 2: Model Training

The purpose of this phase is to test the tickers with serial-correlation in a time period of several months to observe if the

ARMA/GARCH model is able to predict future returns. ARMA, which stands for Autoregressive-Moving-Average, provides an analysis of a weakly stationary stochastic process in terms of two polynomials, one for the autoregression and the second for moving average.

Autoregressive-Moving-Average (ARMA) model defined as:

$$R_t = C + \sum_{i=1}^r \phi_i R_{t-i} + \sum_{j=1}^m \theta_j \alpha_{t-j} + \alpha_t$$

And the ARCH/GARCH model is defined as:

$$\sigma^2_t = K + \sum_{i=1}^p G_i \sigma^2_{t-i} + \sum_{j=1}^q A_j \varepsilon^2_{t-j}$$

where: $\alpha_t = \sigma_t \varepsilon_t$ and $\varepsilon_t \sim SWN(0,1)$

As shown in the above-mentioned formula the ARMA model can be altered to reflect additional information and serial correlation. The key idea here is to be able to identify the orders “r” and “m” of the ARMA model.

If ARMA model is assumed for the error variance, the model is a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) (Bollerslev,1986). The primary use of the GARCH (p,q) model is to provide a measure for volatility that can be used for forecasting and portfolio selection (Engle, 2001). The “p” is the order of the GARCH terms σ^2 and “q” is the order for the ARCH terms ε^2 .

With taking into account the time frame needed to test the model, the study has taken the last 240 days of daily returns out of the data series. The model is then trained for each ticker in the first 120 days to find the orders (r, m, p and q) for the ARMA/GARCH fit and find the companies that show predictability.

Referred to as AIC and BIC, these two methods find the best order for the polynomials above. AIC tries to select a model that best describes

the unknown, high dimensional reality. This means that reality is never in the set of candidate models that are being considered. On the contrary, BIC tries to find the true model among the set of candidates. So, AIC assumes that there is no true model and therefore the best fit with the least innovations is a suitable approximation of what is to be estimated, but BIC assumes that there is a true model, so it prevents over-fitting. There is no preferred method between the two, and the selection of one over the other is still the subject for debates. Minimum AIC is what is used here as the indicator for the goodness of the ARMA/GARCH fit. The orders are changed in a structure of 4 integrated loops and the AIC is calculated and stored. Minimum AIC is then found and the respective orders are then reported as the orders of the best ARMA/GARCH fit. As demonstrated in the MATLAB code shown below in figure 3, by putting a floor of 0 and a cap of 5 on AR and MA orders and a floor of 0 and a cap of 2 on ARCH and GARCH orders, the 4 integrated loops store all values for AIC and then the minimum indicates the best orders for the fit.

```

AIC = NaN(nAR+1,nMA+1,nGARCH+1,nARCH+1);
BIC = NaN(nAR+1,nMA+1,nGARCH+1,nARCH+1);

for Q=1:nARCH+1
    for P=1:nGARCH+1
        if Q-1==0 && P-1~=0
            break
        end
        for M=1:nMA+1
            for R=1:nAR+1
                spec = garchset('R',R-1,'M',M-1,'C',1,'P',P-1,'Q',Q-1,'K',1, ...
                    'Display','off');
                [coeff,errors,LLF,Innovations,Sigmas] = garchfit(spec,(Returns));
                [AIC(R,M,P,Q),BIC(R,M,P,Q)] = aicbic(LLF,garchcount(coeff), ...
                    length>Returns));
            end
        end
    end
end

[minNum, minIndex] = min(AIC(:));
[r, m, p, q] = ind2sub(size(AIC), minIndex);
r = r-1;    m = m-1;    p = p-1;    q = q-1;

```

Figure 4: MATLAB Code to demonstrate the use of AIC

With step increments of 4 days in 120 days of data, the model ends up with 30 data-points. At each of these dates we run the model to find the

orders of the ARMA/GARCH fit. Following that, the model develops a forecast returns vector for the next 20 days. These returns are then used to forecast prices in 10, 15 and 20 days (two, three and four weeks). The reason for such (10, 15 and 20 days) windows is to give the noise a chance to cancel out and to observe if the model is able to capture the drift. Also by taking 2 out of 3 successes in predicting price movements as an indicator of success in prediction, the model is effectively reducing the impact of noise on those certain dates. Figure 4 illustrates a snapshot of the data points for the ticker MKTX.

Example: Prediction Result for Ticker MKTX				
MKTX	Pass or Fail 1 or 0	Pass or Fail 1 or 0	Pass or Fail 1 or 0	Total Pass is => 2
Day 0	1	1	0	1
Day 4	1	1	1	1
⋮	⋮	⋮	⋮	⋮
Day 116	0	0	1	0
Day 120	1	0	1	1
Average Success Rate(%)	51.6	51.6	74.2	54.8
Demonstrating whether the model has successfully predicted the direction of change in price in increments of 4 days				

Figure 5: MKTX data snapshot of forecasts

It should be noted that the testing phase is performed on the last 120 days of data available (most recent). The same step of 4 days to process the 120 days of data is selected and thus end-up with 30 results. Each result consists of 3 success rates for 10, 15 & 20 days, which we combine into one indicator of ticker's prediction success rate.

Phase 3: Portfolio Creation

This phase represents a strategy in trading or a way one can create a portfolio based on the stocks selected to create profit. Following phase 2, every ticker has a total of 30 dates in which the model has managed to capture the price movements in increments of 10, 15 and 20 days. If the model has successfully been able to capture the direction of the movement 2 out of 3 times, it is considered a success case. In order to prevent the weekday effect, the model works with a 4 days step between start dates, which is not a multiplication of 5 (a working week). As a result the 30-start date sample out of 120 training days can be considered to represent the data. At this point each ticker is assigned a number representing the average success rate of the model predicting that ticker for 30 sample dates.

Training Window: Sample Ticker Results				
Company Name	10 Day Forecast(%)	15 Day Forecast(%)	20 Day Forecast(%)	Success Rate (%)
MKTX (NASDAQ)	51.6	51.6	74.2	54.8
AWH (NYSE)	51.6	58.1	64.5	61.3
CNK (NYSE)	45.2	38.7	35.5	35.5
TLAB (NASDAQ)	38.7	54.8	38.7	48.4
FFIN (NASDAQ)	32.3	32.3	29.0	38.7
CATO (NYSE)	48.4	48.4	48.4	45.2

Figure 6: Training Window Results

The following phase will introduce the selection criteria to the results. The tickers that have a success rate of more than 50% are to be tested in the next phase.

Phase 4: Testing

The tickers selected in the previous phase can be used to form an identically weighted portfolio of investments. To avoid dealing with portfolio optimisation tools, the investments are purposely chosen to be equally weighted. To test the portfolio, we forecast the prices for 10, 15 and 20 days and then check the portfolio return for a 10, 15 and 20 day windows. If the selected tickers generate positive alpha in comparison to the market return (benchmark) then we conclude that the strategy is effective.

The testing phase is performed on the last 120 days of data available, referred to as out-of-sample window. With the same step of 4 days we end-up with 30 sample dates to test the portfolio. The criteria for success for such dates are the overall ability of the model to forecast the direction of price movements after 10, 15 & 20 days.

By counting the success cases for each time horizon, 3 success rates are calculated for each ticker, leading to the calculation of portfolio success rate for each time horizon (2, 3 and 4 weeks). These numbers, if greater than the market return, can be used to form a portfolio of digital options, since we are only interested in the direction of price movement.

To compare the overall success rate, a measure of 2 out of 3 successful predictions of direction of price movement is again used. This number then can be compared with the success rate of the selected tickers in the previous phase and be an indicator of existence of momentum in returns.

An example of the way the model performs the abovementioned filtering in the out of sample window is provided in the below chart:

Example: Out of Sample Window (Prediction Vs. Actual)							
	Actual			Forecasted			
	Initial Price	End-of-Period Price	Movement	Initial Price	Forecasted Price	Movement	Success Case
Sample Day 1	25.4	23.4	-2	25.4	27.4	+2	0
Sample Day 2	25.4	26.9	+1.5	25.4	27.4	+2	1

Figure 7: Arbitrary example to show how the model decides whether it has generated the correct forecast. If the actual movement and forecasted movement are the same it is considered a success.

The model utilizes price data from the end of 2008 to July 1st 2012. A window of 240 days is subtracted from the end of the vector (more recent prices) and the remaining data is used to determine the order of the ARMA/GARCH fit and to calculate coefficients for the model. The model is then used to forecast returns for 10 day (2 weeks), 15 day (3 weeks) and 20 day (4 weeks) horizons, and these forecasted returns are used in evolution of the prices to see if the model can predict the direction of price changes. This process is done every 4 days for the last 120 days and the returns are compared (as shown in figure 8). This setting provides 30 test date points in the 120 recent dates (about half a year). The probability of a successful prediction in at least 2 out of 3 time horizons is calculated and stored.

Out of Sample Window Results					
Company Ticker	10 Days (%)	15 Days (%)	20 Days (%)	Success Rate (Training Window)	Success Rate (Out of Sample Window)
MKTX	60.0	63.3	60.0	54.8	63.3
AWH	66.7	80	83.3	61.3	86.7

Figure 8: Out of sample window results

By comparing the results for the selected tickers from the first 120 days (training) and the second 120 days (testing) the model is able to show

the effectiveness of momentum strategies using a floating order ARMA/GARCH model.

The Methodology Process

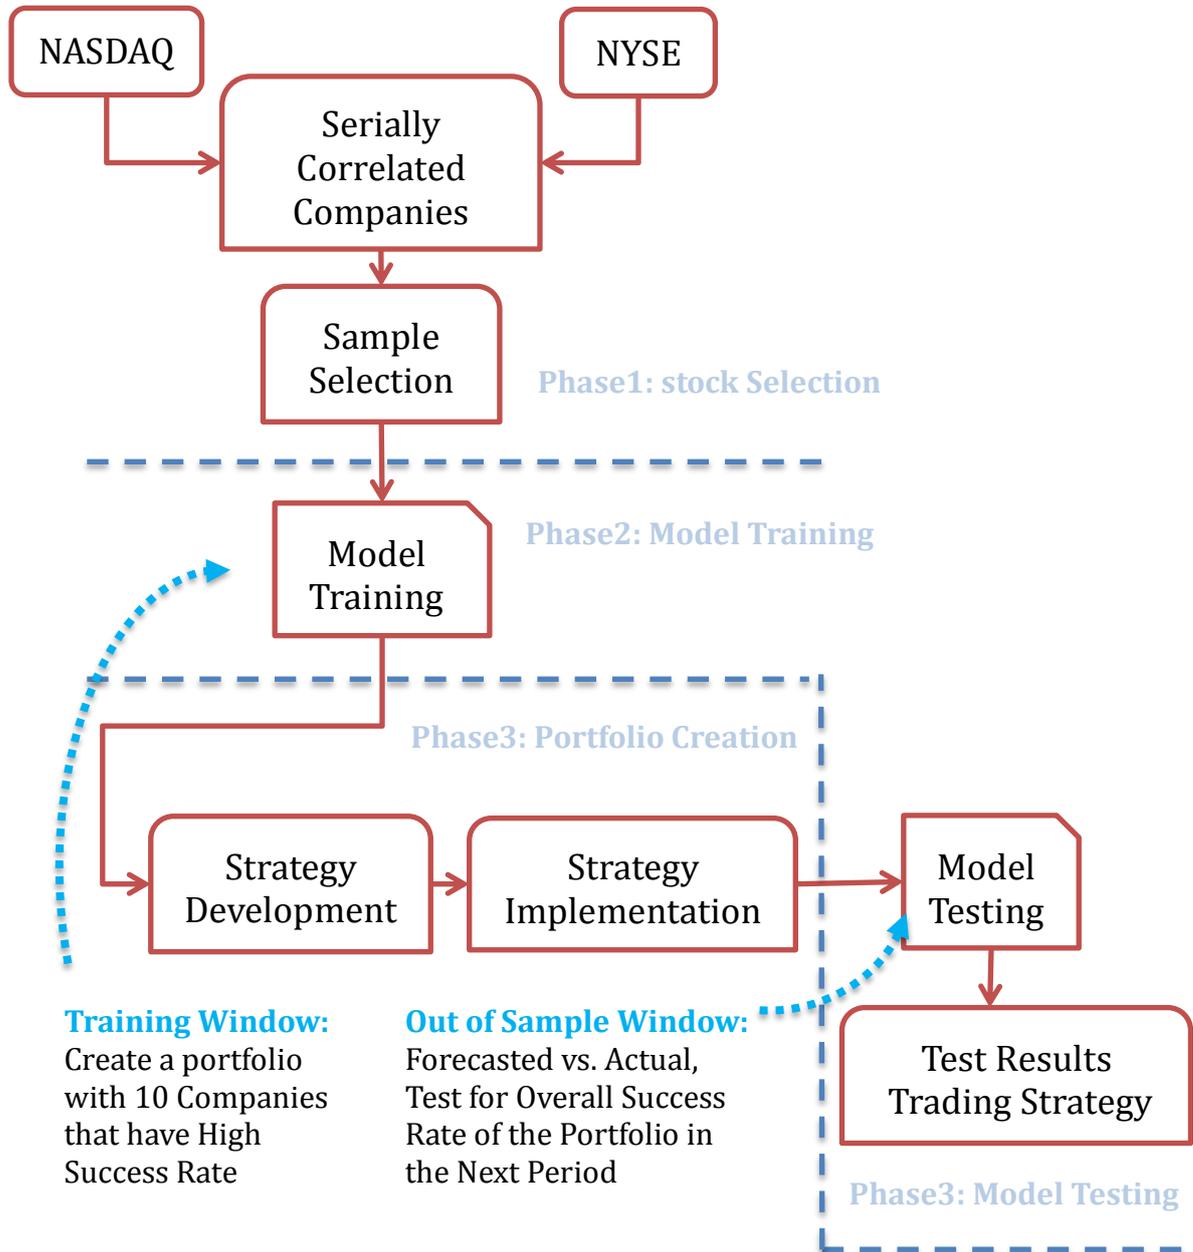


Figure 9: Methodology Process

Results

The average predictability of the direction of price movements for the 30 tickers is about 50 percent. For tickers that have above 0.54 probability of being predicted correctly, the probability of successful prediction in the out-of-sample window shows proof of momentum. These tickers (in our sample of 30, 10 tickers) had overall average success rate of 0.597 in the training window. The same tickers, have an impressive overall success rate of 0.573 in the out-of-sample window.

As we got approximately 50 percent on average for direction prediction, the plain ARMA/GARCH model without screening and filtering criteria fails to produce a winning strategy, but with addition of the proposed selection strategy, the selected tickers can be used to form an alpha-generating portfolio. Since we only care about the direction of the change and not the actual change in prices, this model can be used to pursue winner opportunities in digital derivative trading. There are 10 tickers out of 30 tickers in the sample that pass the criteria for the strategy, 7 of which carry the momentum into the second 120 days. Below graph demonstrate the difference in ability to forecast between the training and out-of-sample windows.

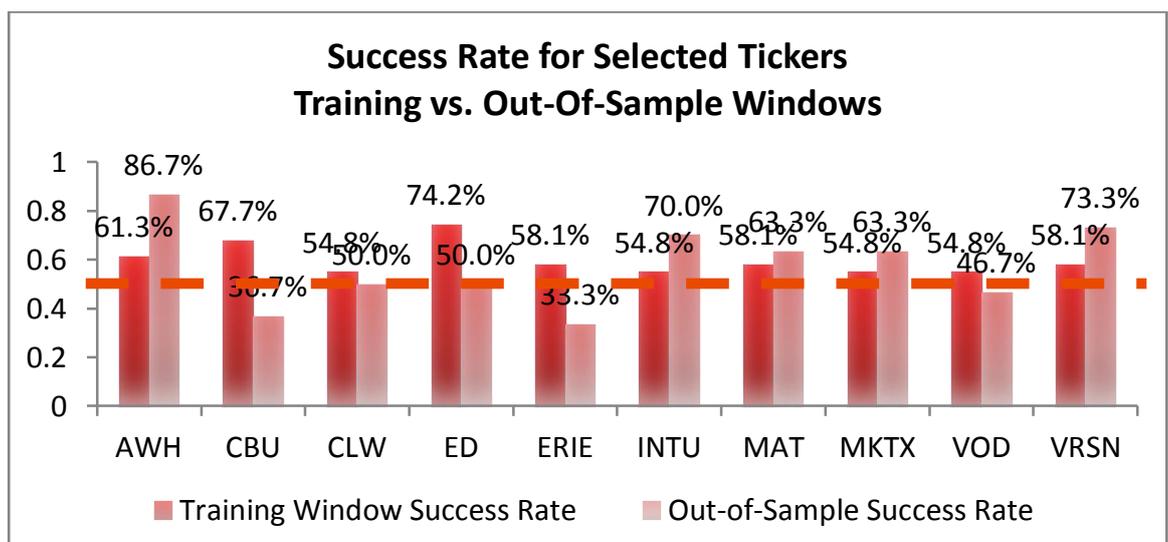


Figure 10: Success rate chart

From a trader's point of view, the same strategy can be implemented when the model fails. For example for a ticker that can be predicted successfully only 30 percent of the time using this model a strategy can be implemented reversing the strategy proposed by the model.

To implement a digital derivative trading strategy, the trader should buy or sell digital options whenever the model proposes so (on the tickers that are selected in the training window). In the case of this research if a trader follows the predictions of the model, 57 percent of the options will end-up in the money on average. If digital double or nothing options are to be used, $0.57 \times 2 + 0.43 \times 0 = 1.14$, which is 14 percent return on average for a set of executed trades. Of course commissions, fees, and transaction costs have to be taken into account, but the margin here (although risky) surpasses those costs with an acceptable spread.

These are the steps of the proposed trading Strategy:

- Proposed Trading Strategy:
 - Find tickers that have a success rate of more than a defined threshold (e.g. 54 percent) in the last 120 days
 - Form a portfolio of digital double-or-nothing options
 - For dates that model predicts positive price movement buy a digital call, for the ones with a negative prediction buy a digital put
 - Do the same for 120 days
 - The model suggests that you're holdings will end-up in the money more than 50 percent of the time (in our case 57 percent)
 - This is a 14 percent return in 120 days

Conclusion

Momentum strategies have been the subject of a lot of research, and have been tested thoroughly. They are considered common indicators of market efficiency. Yet there are cases (and in our case, tickers) that show positive results when the returns are evaluated for serial correlation. The signal to noise ratio is low in most cases, and the predictability greatly suffers from it. Even the majority of tickers that pass the tests for serial correlation, when tested intensely; fail to predict the direction of price change with an overall success rate of more than 50 percent.

Our approach consisted of 4 basic modules. The first module screens all the tickers and generates a list of prospects. The second module goes deeper into the analysis of data and uses some kind of pattern recognition or regression tool to distinguish the signal from the noise. The third module devises a screening criterion to choose the tickers that show hints of serial correlation and predictability in recent dates. The fourth module is about testing the tickers chosen. The ARMA/GARCH “fit” is then tested in recent months to see how accurate the model can predict the subjects of the test.

After these 4 steps are taken, by analysing the results we can come up with trading or investing strategies in cases that the used model can predict actual results with high success rate.

Devising a well-established screening in combination with a trading strategy and an exit strategy is of key importance, especially when we are dealing with risky models like the one we presented here. These models can be used as part of a diverse portfolio of investments, as a means to generate returns alongside other less risky components of the portfolio.

Areas for Further Research

Each of the steps taken in our methodology can be a subject for change, improvement, and further research.

The regression tool in the 2nd module is the core characteristic that separates different models from each other. This part is so important that it gives models their name. Further researches can use other methods like VARMA, EGARCH , etc. as the regression tool and then the first module should change to find prospects consistent with the new setup.

Another major opportunity for further research is to change in the nature or frequency of the data series. The model can be tested in other areas like volatility or other interesting criteria, and with incorporating data with other frequencies from minute to seasonal with respect to the subject and the nature of the phenomenon that is going to be predicted.

Testing the same model and strategy during times of financial crisis to see if there is ability to capture momentum and implement the same strategy can be another area for further research.

Economics of these predictive tools, and determination or interpretation of both the tools and the results is another major area for further research. And for more practical purposes, development of relevant trading strategies for each case needs design and refinement, and can be a subject of further investigation.

References

Blake LeBaron. Some Relations Between Volatility and Serial Correlations in Stock Market Returns. *The Journal of Business* , Vol. 65, No. 2 (Apr., 1992), pp. 199-219

Bollerslev, Tim (1986). "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31:307-327

Campbell, John Y. "Stock Returns and the Term Structure." *Journal of Financial Economics*, June 1987, 18(2), pp. 373-99.

Covrig, V., Lilian, N., 2004. "Volume Autocorrelation, Information and Investor Trading." *Journal of Banking & Finance*, 28, 2155-2174.

Francis X. Diebold. Serial Correlation and the Combination of Forecasts. *Journal of Business & Economic Statistics* , Vol. 6, No. 1 (Jan., 1988), pp. 105-111

Francis X. Diebold and Jose A. Lopez. *Forecast Evaluation and Combination*. University of Pennsylvania, June 1996.

Grant McQueen and Steven Thorley. Are Stock Returns Predictable? A Test Using Markov Chains *The Journal of Finance* , Vol. 46, No. 1 (Mar., 1991), pp. 239-263

Guo, Hui. *Stock Market Returns, Volatility, and Future Output*. The Federal Reserve Bank of St. Louis. 2002

Lee, C.M.C., Swaminathan, B., 2000. "Price Momentum and Trading Volume." *Journal of Finance*, 55, 2017-2069.

Lo, Andrew and MacKinlay, A. Craig. *The Econometrics of Financial Markets*. Princeton: Princeton University Press, 1997.

Lilien, David M. "Sectoral Shifts and Cyclical Unemployment." *Journal of Political Economy*, August 1982, 90(4), pp. 777-93.

Polinsky, David and Zhang, Sisi. *A High Frequency, Volumetric Trading Strategy*. Simon Fraser University, 2010

Ronald J. Balvers and Douglas W. Mitchell. Autocorrelated Returns and Optimal Intertemporal Portfolio Choice. *Management Science*, Vol. 43, No. 11 (Nov., 1997), pp. 1537-1551

Sarno, Lucio and Valente, Giorgio. *Modelling and Forecasting Stock Returns: Exploiting the Futures Market, Regime Shifts and International Spillovers*. University of Warwick and Center for Economic Policy Research, November 2002.

Schwert, G. William. "Business Cycles, Financial Crises, and Stock Volatility." *Carnegie-Rochester Conference Series on Public Policy*, Autumn 1989a, 31

Shiller, R. J. (1984): "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, 2, 457---498.

Shiller, R. J. (2003): "From Efficient Markets Theory to Behavioral Finance," *Journal of Economic Perspectives*, 17, 83---104.

Appendix

Definitions

1 Price to Return Conversion - Once the data has been gathered and populated from NYSE and NASDAQ it is converted to log returns in order to end-up with a stationary (more accurately: a zero mean) data series.

$$r_t = \ln P_t - \ln P_{t-1}$$

Mean - The returns are then checked for zero and constant mean by performing a t-test using test statistic. The null hypothesis of constant and zero mean can then be rejected (fail to be rejected) at $\alpha=5\%$ significance level if the condition below is satisfied(not satisfied).

$$|t| = \frac{x_t^c}{\sqrt{S_x^2 / T}} \stackrel{3}{\approx} Z_{\alpha/2, T} \gg 1.96$$

Variance – The fact that sample variance for independent and identical distributed process has a Chi-Squared distribution can be utilized as the basis for a hypothesis test of constant variance. Due to the fact that almost all financial time series illustrate time-varying volatility, referred to as “heteroscedasticity”, it is preferred to identify constant variance graphically and search for “volatility clustering”. In addition to testing for constant variance visually, our model also performs a chi-square test, which checks for constant variance.

Serial Correlations – While there are number of ways to detect the presence of serial correlation, our model utilizes Ljung-Box test statistics, and rejects (fail to reject) the null hypothesis of no serial correlation if the condition $|Q(m)| \geq \chi_m^2(\alpha)$ is satisfied. This is done through the MATLAB functions `lbqtest` and `autocorr`.

$$Q(m) = T(T+2) \sum_{k=1}^m \frac{\rho_k^2}{T-k}$$

Auto Correlations - the correlation between two random variables X and Y is defined as $\rho_{xy} = \frac{Cov(X,Y)}{\sqrt{Var(X)Var(Y)}}$. When the linear dependence between returns today (r_t) and its past value (r_{t-l}) is of interest, the concept of correlation is referred to as autocorrelation. Furthermore the correlation coefficient between (r_t) and its past value (r_{t-l}) is called lag-l autocorrelation of r_t . Denoted as ρ_l , the autocorrelation for a weakly stationary series is defined as $\rho_l = \frac{Cov(r_t, r_{t-l})}{Var(r_t)}$ where the property $Var(r_t) = Var(r_{t-l})$ for a weakly stationary series is used. A weakly stationary series r_t is not serially correlated if and only if $\rho_l = 0$ for all $l > 0$.

Partial Autocorrelation Function (PACF) – The partial autocorrelation function plays a significant role when it comes to data analyses at identifying the extent of the lag in an AR model. Once the order has been determined, the model checks whether it is below or above 5 lags. Our model is designed to have a maximum of 5 lags even if the PACF function determines a higher order.

GARCHFIT – At this point we pass the orders we have obtained from the previous two steps and place into GARCHFIT. Given an observed univariate return series, garchfit estimates the parameters of a conditional mean specification and conditional variance specification of GARCH. The estimation process infers the innovation (its residuals) from the returns series. It then fits the model specification to the return series by maximum likelihood. Using the likelihood values as input, the model utilizes the Akaike (AIC) information criteria and computes the best possible orders for AR and MA model and therefore a superior fit. The same process is then repeated in order to find the best orders to the GARCH function to estimate and forecast volatility.