Ashwood, Andy & Basu, Anup
(2014)
The quest for alpha: Can artificial neural networks help?
JASSA, *2014*(1), pp. 13-18.

The Quest for Alpha: Can Artificial Neural Networks Help?

Andrew J. Ashwood
Graduate School of Business
Queensland University of Technology
Brisbane QLD 4001
e-mail: andrew.ashwood@qut.edu.au

Anup K. Basu #
School of Economics and Finance
Queensland University of Technology
Brisbane QLD 4001
e-mail: a.basu@qut.edu.au

Abstract

The application of artificial neural networks (ANN) in finance is relatively new area of research. We employed ANNs that used both fundamental and technical inputs to predict future prices of widely held Australian stocks and used these predicted prices for stock portfolio selection over a 10-year period (2001-2011). We found that the ANNs generally do well in predicting the direction of stock price movements. The stock portfolios selected by the ANNs with median accuracy are able to generate positive alpha over the 10-year period. More importantly, we found that a portfolio based on randomly selected network configuration had zero chance of resulting in a significantly negative alpha but a 27% chance of yielding a significantly positive alpha. This is in stark contrast to the findings of the research on mutual fund performance where active fund managers with negative alphas outnumber those with positive alphas.

Keywords: Neural networks, portfolio selection, excess returns, alpha.

Corresponding author.

The authors would like to thank Mike Drew, Scott Boland, Robert Bianchi, and Min Zhu for their helpful suggestions with this research. Comments from the participants in Griffith Industry Academia Symposium and QUT School of Economics and Finance Seminar are acknowledged. We are grateful to the High Performance Computing and Research Support (HPC) at QUT for their excellent assistance with advanced computing issues.

The Quest for Alpha: Can Artificial Neural Networks Help?

Australians have a significant portion of their personal wealth invested in the stock market. A large proportion of this wealth is invested with active fund managers who try to beat the market and earn excess risk-adjusted returns (alpha) by actively selecting stocks to buy and sell. However, there is a plethora of research evidence showing that most active managers underperform the broad market index on a risk-adjusted basis. In Australia, 7 out of 10 actively managed retail funds underperformed the market index over both a one and three year horizon (Karaban & Maguire, 2012). This evidence is consistent with the efficient markets hypothesis, where current prices tend to reflect all available information. Yet the attempt to predict the future course of stock prices and earn excess returns has remained a persistent endeavour for many investors (Malkiel, 2011).

Advances in computing power in combination with the widespread availability of historical datasets have provided investors with increased opportunity to test markets for predictable returns. In this paper, we present an artificial neural network (ANN) model that utilises a combination of technical and fundamental input data to predict future prices of widely held Australian stocks and use these predicted prices for stock portfolio selection. We present evidence on whether such portfolios can earn positive alpha for investors.

An ANN is a mathematical model that is inspired by the structure and function of biological nervous systems, such as the brain, in processing information. The brain continually receives input information from receptors, processes the information, and makes decisions. Like the biological nervous system, the ANN is composed of a large number of highly interconnected processing elements working together to solve specific problems. ANNs, just like human brains, learn by example.[1]

[1]There are number of books and articles that explain neural networks from a beginner's perspective. See for example, Coolen (1998) or Garson (1998). Many articles are also available on the World Wide Web.

Financial research utilising artificial neural networks is a relatively new area with published research in the field only going back to a little over two decades. Over this period, the majority of the studies have focussed on US stock prices and indexes. Very limited research has been undertaken in the Australian context. Among them, Tan (1997) found that statistical auto-regressive models when combined with ANN produce superior forecast and profitability in forex trading (AUD/USD) market than when used in isolation. Ellis & Wilson (2005) applied ANN modelling techniques using seven fundamental indicators to construct portfolios in the Australian property sector that outperformed both DS Australian Real Estate Index and S&P/ASX Property Index on a risk-adjusted basis. Finally, Vanstone, Finnie and Hahn (2010) used four fundamental indicators as inputs to devise a trading rule based on ANN for stock selection in the Australian market. They found that an ANN based rule produced higher returns, albeit with higher volatility, compared with a buy and hold approach and a filter rule based on the same fundamental variables as inputs. Unlike the above studies, which use either technical or fundamental indicators as inputs to ANN, we investigate the stock selection performance of neural networks in the Australian market using a range of fundamental and technical indicators simultaneously as inputs.

Network model

We used a walk forward testing approach, which is considered to be the best method for prediction for time series data. It simulates the real-life trading situation where the model is regularly retrained with new data as it becomes available (an implicitly Bayesian approach). The frequent re-training is time-consuming but allows the network to adapt to changing market conditions. We employed four different training periods: 3, 6, 12, and 24 months. For each walk forward testing (rolling) window, the validation period (containing the data set used for monitoring the error during training) was fixed at 6 months and the out-of-sample testing period at 12 months. Figure 1 diagrammatically shows the walk forward testing approach using an example of a 24-month long training period. All neural network input and output data was pre-processed. Several pre-processing algorithms were adopted in

Comment [KD1]: What happens in the validation period?

Comment [AJA2]: Training and Validation steps added in next page to address this comment.

order to ensure that the neural network learned quickly and provided better performance. The training and validation process comprises the following steps:

1. The training data is presented to the network;
2. The network computes outputs;
3. The network outputs are compared with the desired outputs and error is calculated;
4. Network weights are updated based on the error calculation;
5. Process repeats until the error reaches a pre-defined level or the maximum number of epochs has occurred.

Figure 1: Walk Forward (Rolling) Testing Window

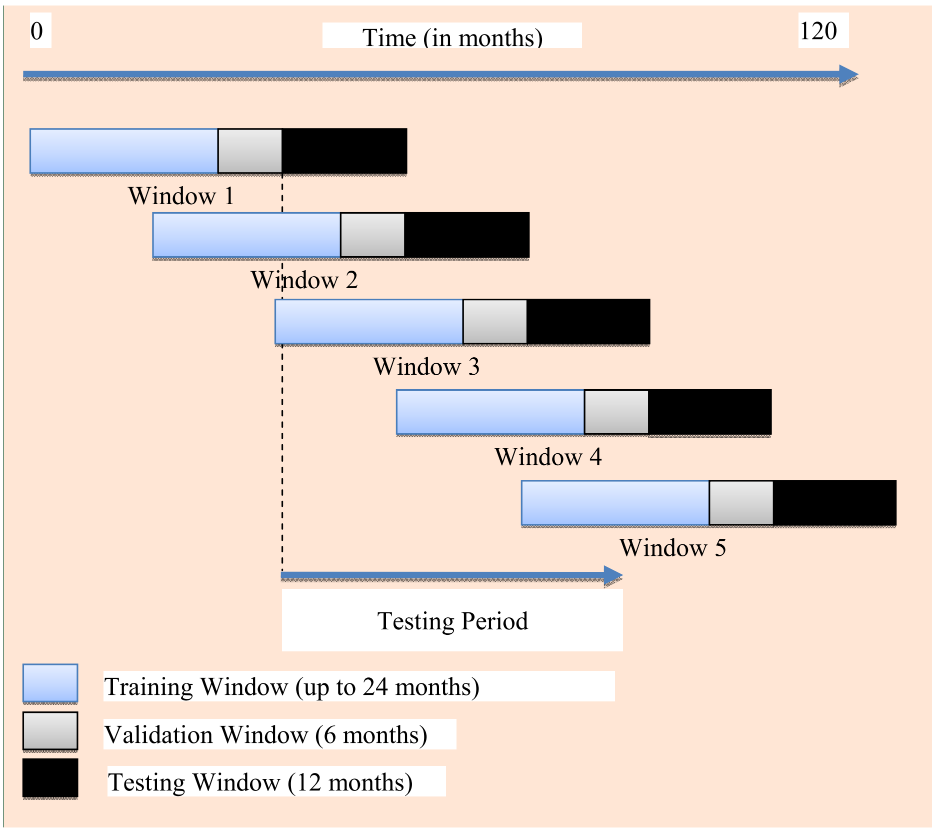

Comment [KD3]: The diagram suggests that there is a 6 month gap between the start of each window? Is that the case? The testing period in the diagram (shown as an arrow) relates to window 1 or all windows – if the latter, the data in months 31-36 enters the tests once, that in 37 to 42 twice etc?

Comment [AJA4]: The diagram is not drawn to scale. There is a 12mth gap between testing windows. Diagram amended to address this issue.

Comment [KD5]: Can you describe in some simple words what goes on in the training window. The reader needs some idea of how the black box operates even if the contents are black.

Comment [AJA6]: Additional text added to address this comment.

For this study, we used weekly price data on 20 randomly selected stocks from ASX50 (which comprises of 50 largest and widely held stocks in the Australian market) between January 1997 and December 2011. The network input data inputs consisted of 59 indicators for each of these stocks. Of these, 18 were fundamental indicators (like return on assets, profit margin, sales growth etc.) and 41 were technical indicators (different momentum indicators for price movements and volume like 20, 30, 50 day moving averages, relative strength index etc.).[2] The values for the fundamental indicators were calculated from the data extracted from half-yearly financial reports published by the companies.[3] The technical indicators were available as a weekly data series.[4] All data was obtained from Bloomberg. All analysis was undertaken using Neural Network toolbox in Matlab® (version R2012a) software program developed by The MathWorks, Inc. The parameter specifications of the neural network configuration is summarised in Table 1.[5]

Comment [KD7]: This won't help the reader at all and should either go as an appendix or "on available from authors". And there look to be some weird typoes in the variable parameters section

Comment [AJA8]: We acknowledge this comment. We originally included these details to give the reader a feel for the difference combinations networks tested and also to provide some transparency around parameters. We have no problem with adding as an appendix or mentioning this is 'available from authors'.
Sorry we don't see any typos. If you are concerned about 'x' used to denote multiplication symbol, we can use '*' instead.

Table 1: Network Configuration

[2] The full list of fundamental and technical indicators can be obtained from the authors on request.
[3] It can be argued that there is usually a delay of few months between the date of the financial statements and the date of their actual release. However, we have not adjusted for such time lags in our study. If we believe markets are by and large informationally efficient with many well-informed analysts closely following companies, the public announcement of half-yearly results would be well anticipated by the market in most cases, more so for the large companies.
[4] Early data (1997-1999) was used to compute the values for the technical indicators.
[5] Further information regarding network parameters can be found in Beale, M. H., Hagan, M. T., & Demuth, H. B. (2011).

Fixed Parameters

Architecture	3-layer feed-forward network Each layer fully connected to its adjacent layers only (no connections between input and output layer)
Initialisation algorithm	Nguyen-Widrow
Validation period	6 months
Learning algorithm	Gradient descent with momentum
Learning rate	0.1
Momentum factor	0.9
Transformation function	Hyperbolic tangent
Max. training epochs	100,000
Validation stop	50 iterations
Evaluation criteria	Mean square error

Variable Parameters

Input type	Price OR Technical indicators only OR Fundamental inputs only OR Price + Technical OR Price + Fundamental OR Price + Technical + Fundamental (6 options)
Lookback window	4, 8, 12, 16, 20 periods (5 options)
Hidden layer size	30, 60, 90, 120, 150 nodes (5 options)
Training period length	3, 6, 9, 12 months (4 options)

600 different neural networks run for each stock/portfolio for each of the 10 one-year testing periods:

6 input types x 5 lookback windows x 5 hidden layer size x 4 training period lengths
----
= 600 network models

A schematic diagram of the ANN as implemented is shown in Figure 2.

Figure 2: Schematic Diagram of ANN

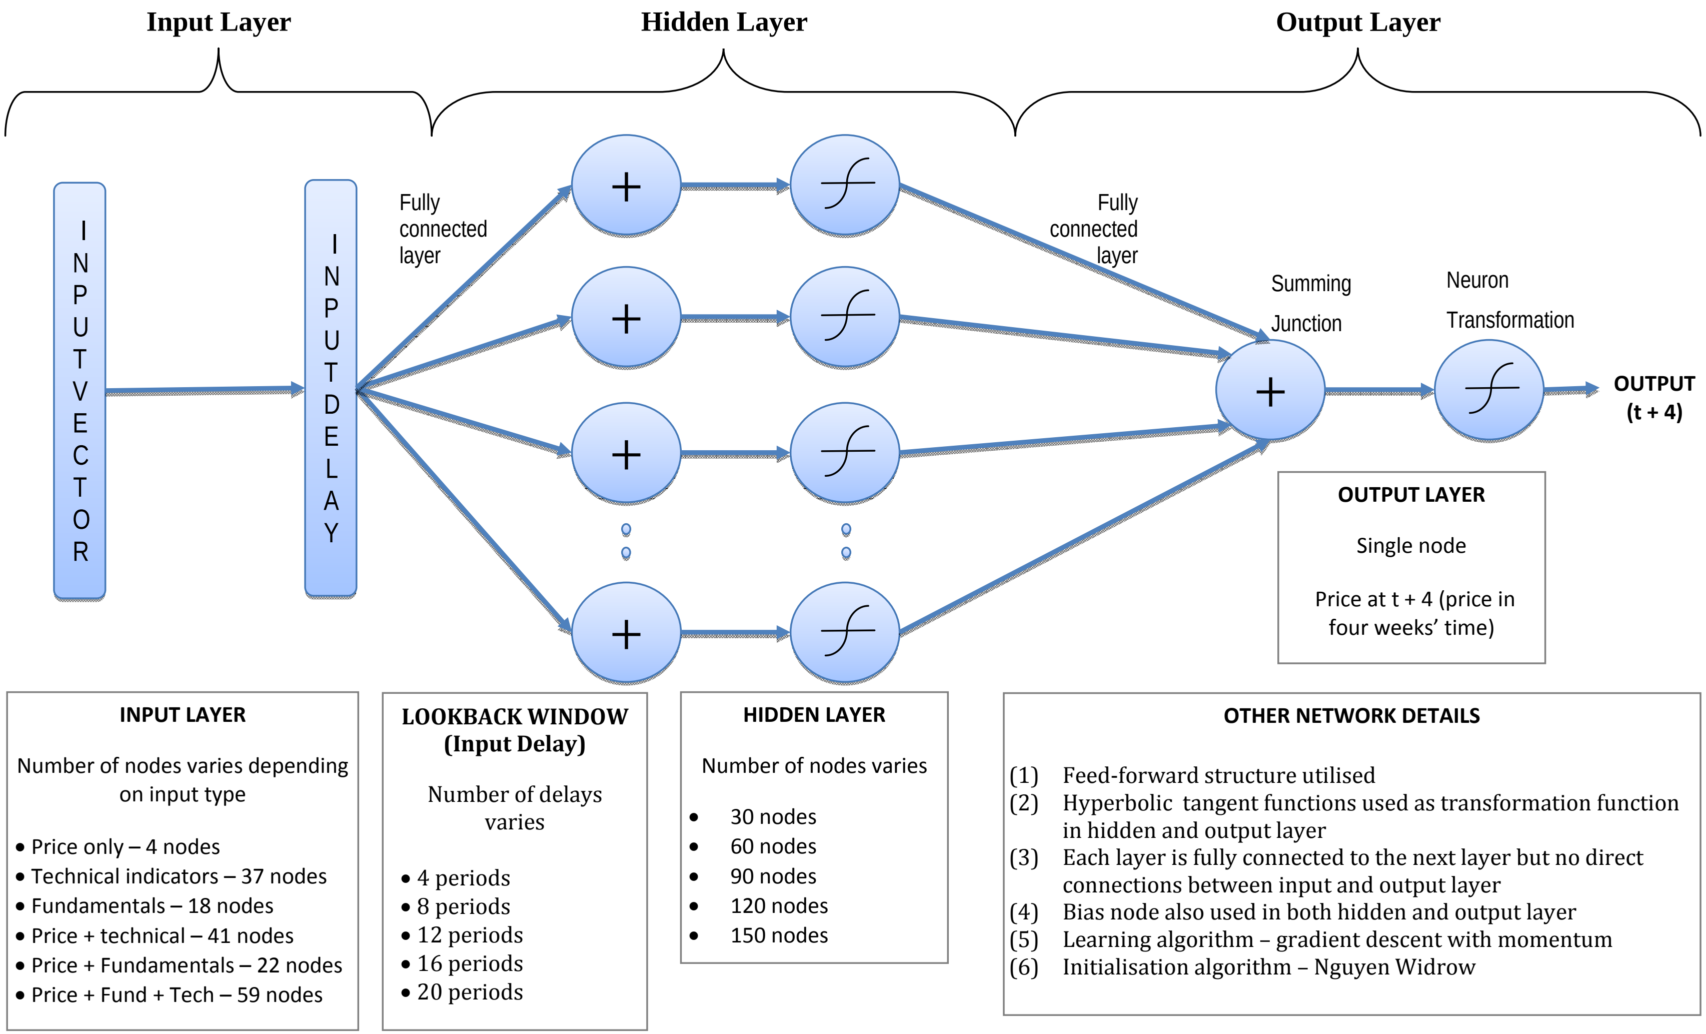

Predictive ability

While a network model that correctly predicts stock price movement could be used to achieve superior returns, from a practitioner's perspective (and depending on the trading system implemented) knowing the precise quantum of the price movement may be less important than correctly predicting the direction of the price movement. For example, few investors would be unhappy with a situation whereby their network model predicted a +15% price movement and the actual price movement was +5%. Whilst the prediction error was 10% but the trade would be still profitable as the directional movement was accurately predicted. Contrast this with a trading situation where the network model predicted a +2% price move and the actual price movement was -3%. In this case the prediction error was 5% (far more accurate than the previous example) but the trade was not profitable as the directional movement was incorrectly predicted.

Comment [KD9]: I don't think this is correct for long-short portfolios, because long-short weights will be based on predicted price movements. And you r results seem to indicate this.

Comment [AJA10]: You are right but this example is about a long only investment. We are merely trying to make the point that a network with a large error magnitude could still be highly profitable if it could reliably predict directional movement.

The neural network in our study performed reasonably well at predicting the correct directional movement of stock prices. For all stocks (with one exception), direction of price changes were accurately predicted at least half the time. Though most of the success rates were only marginally above 50% (generally ranging from 50 to 55%) there were better performers such as 65%. The large number of observations for each stock and portfolio (n = 521 i.e. 521 four-week ahead predictions over the 10-year testing period) denoted that any directional prediction accuracy above 53% was significant at least at the 5% level. Overall, 13 of the 20 stocks achieved directional movement accuracy with statistical significance at the 5% level. 5 of the stocks achieved directional movement accuracy with statistical significance at 0.1% level.

Comment [KD11]: Four-week-ahead?

Comment [AJA12]: Yes – text amended. Thank you.

Portfolio strategy

To determine if the ANN model could be used to achieve abnormal returns, the following procedure was adopted. The ANN price predictions were undertaken for each stock. For each four-weekly time-step over the 10-year period, 600 ANN models were run to predict future prices. These 600 simulations were the result of testing 6 different input types, 5 lookback window options, 5 hidden layer size options, and 4 training period lengths (refer to Table 1 and Figure 2 for further details). For each four-weekly time-step, the ANN simulated price was compared to the real price and the 600 networks specifications were placed in a rank order based on the size of the error. For the purpose of this analysis, we report the results of two network configurations: the most accurate network (i.e. one that produced the minimum error over each twelve month testing period) and the network with median accuracy.[6]

Comment [KD13]: For each forecast date – ie different ANNs, or the one ANN giving minimum over the forecast period

Comment [AJA14]: Minimum over the testing period of each walk forward testing window i.e. minimum error over each 12 month period. We have included additional words for clarity.

Once the predicted price returns for each stock and time step were calculated for the two configurations described above, two portfolios were constructed for each case. The first portfolio was a long only portfolio in which a long position was taken for all stocks with a positive expected return weighted according to the magnitude of their expected price return. For example, a stock with a 10% expected price return was given double the weighting of a stock with a 5% expected price return. All stocks that had an expected negative price return were assigned a portfolio weight of zero. The second portfolio was a 'long minus short' portfolio. To construct this portfolio, all 20 stocks were weighted according to the absolute value of the magnitude of their expected price return. A long position was taken for all stocks with a positive expected return while a short position was entered for all stocks with a negative expected return.

Comment [AJA15]: Yes

Comment [AJA16]: Yes

Comment [KD17]: Any particular reason – it will give rise to some weird portfolios, eg if there are several stocks with expected returns just above zero and some with say 10%

Comment [AJA18]: We wanted to implement a trading system whereby strong price signals provided by the ANN were weighted more heavily (the intuition being that if the network predicts a large price movement it should be a more reliable).

The portfolio construction occurred at the beginning of each four-weekly time-step over the 10-year period and stocks were held until the end of the time step. While the portfolio construction was based upon the expected return generated by the ANN model, the actual return for the portfolios at the end of four weeks was calculated using the actual price data. The process was then repeated for the entire testing period.

Performance Measurement

[6] Obviously the rankings of these network configurations changed every time step based on performance.

The four-factor model (Carhart, 1997) was used as the performance measurement framework for the ANN models. Excess portfolio returns were regressed against the market, size, value and the momentum factors as:

$$r_{pt} = \propto_{pT} + \ b_{pT} RMRF_t + s_{pT} SMB_t + h_{pT} HML_t \ + \ u_{pT} UMD_t \ + e_{pt} \qquad t = 1, 2, \cdots, T \qquad (1)$$

where:

r_{pt} is the monthly return on portfolio (p) in excess of the 10-year Australian government bond rate in month *t*

$RMRF_t$ is the excess return on the ASX200 index in month *t* over the 10-year Australian government bond rate

SMB_t is the monthly return on the mimicking size portfolio i.e. excess return of 'small' stocks over 'large' stocks in month t

HML_t is the monthly return on the mimicking book-to-market portfolio i.e. 'value' stocks over 'growth' stocks in month t

UMD_t is the monthly return on the mimicking momentum portfolio i.e. excess return of recent 'winner' over recent 'loser' stocks in month t

Comment [KD19]: Am I correct in thinking that there are 130 observations ie 13 four week periods each year for 10 years. You could do a lot more to help the reader fully understand the steps in your approach.

Comment [AJA20]: Yes. We acknowledge the criticism. (It has been very difficult to achieve clarity to our desired level due to the journal's word count requirement).

Table 2: Four Factor Regression Estimates for Portfolio Returns

Coefficients	Most Accurate Network		Median Network	
	Long portfolio	Long-short portfolio	Long portfolio	Long-short portfolio
α	0.016**	0.013***	0.012**	-0.001
RMRF	0.712***	-0.011	0.137	0.003
SMB	-0.317	-0.604**	-0.050	-0.011
HML	-0.416**	-0.306	-0.310	-0.064
WML	-0.011	0.038	-0.326	-0.115

*, **, *** indicates statistical significance at 5%, 1% and 0.1% level

We proxied the *SMB* return by the return difference between monthly returns of ASX Small Ordinaries and ASX 100 index. Monthly HML data for Australia was obtained from Ken French's website. Momentum portfolios were constructed using monthly returns data from CRIF database following Jegadeesh and Titman (1993).

The regression results are reported in Table 2. The long portfolio created using the most accurate ANN produced a monthly alpha of 1.6% while the median ANN achieved a positive alpha of 1.2%. Both these estimates were significant at the 1% level. The alpha estimates for the long minus short portfolios were less encouraging. Whilst the most accurate ANN portfolio's alpha was still significant, albeit at a diminished 1.3%, for the median ANN, the alpha dissipated. These results suggest that network predicted the positive price movements more successfully than it predicted the negative price movements. Among the other regression coefficients, the market (RMRF) and the value (HML) factors were significantly related to the returns of the long portfolio created using the most accurate ANN but none of the coefficients were significant for the median network's portfolio.

Comment [KD21]: Can't this simply be a result of the particular weighting scheme you used for stocks included in the portfolios.

Comment [AJA22]: We have applied the same weighting mechanism to the long and long/short portfolios. Therefore, one would not expect any differential influence of the weighting scheme itself on the predictive ability of the two portfolios.

We need to caution the reader here that there is no way of predicting ex-ante the accuracy of a particular network and therefore, its ability to result in a positive alpha. In other words, a randomly selected network specification may result in an outcome that is far inferior to that achieved by the two networks in Table 2. In order to gain an understanding of the distribution of alphas generated by the different ANN specifications selected without any ex-ante knowledge about their accuracy, the different network parameter combinations were applied uniformly to all stocks over the entire 10-year period. Long portfolios based on expected returns were formed and their returns regressed against the four-factor model resulting in 600 estimates of alpha. This analysis mimics the real world situation where the practitioner does not have prior knowledge of which network specification to apply in order to achieve maximum predictive capability and randomly selects a particular network specification for application to all stocks in her investment universe.

Figure 3: Frequency Distribution of Four-factor Alphas

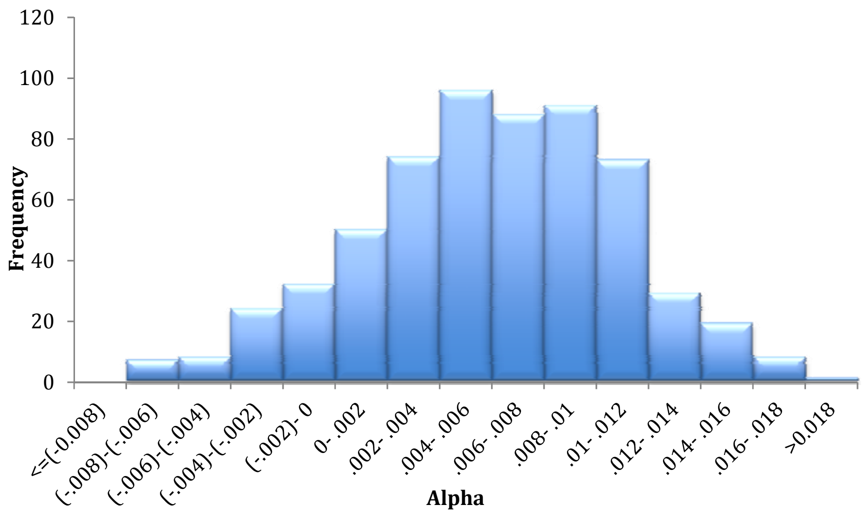
120
100
80
60
40
20
0
Frequency
<=(-0.008)
(-.008)-(-.006)
(-.006)-(-.004)
(-.004)-(-.002)
(-.002)- 0
0- .002
.002- .004
.004- .006
.006- .008
.008- .01
.01- .012
.012- .014
.014- .016
.016- .018
>0.018
Alpha

Figure 3 presents the distribution of alphas. Out of the 600 portfolios, 529 (88%) achieved positive alphas. The alpha value was significantly positive for 144 portfolios (27% of all portfolios) at the 5% level. The distribution of alphas peaked between 0.4% and 1%. The distribution had a slight left skew but there were very few alphas below -0.4%. Remarkably none of the negative alphas were statistically significant. This finding is extremely important as it demonstrates that even without any knowledge about the predictive capability of the networks ex-ante, the practitioner would still have a much higher likelihood of generating positive alpha relative to negative alpha.

Conclusion

Artificial intelligence is increasingly used in different fields of human endeavours mainly due to its predictive abilities based on pattern recognition and learning. The data-rich environment of stock price movements offers fertile ground for testing these capabilities. The ANN model presented in this paper provides encouraging results for investors. We found that the ANNs generally do well in predicting the direction of stock price movements. The portfolio selected by the ANNs with median accuracy every 1-year testing period was able to generate positive alpha over a 10-year period. More importantly, we found that practitioners can improve the likelihood of generating positive alphas using neural networks even without any ex-ante knowledge about their accuracy as many of the network configurations resulted in positive alphas while none resulted in a negative alpha with statistical significance. This is in stark contrast to the findings of the research on mutual fund performance, which show that funds with negative alphas outnumber those with positive alphas. It is also important to note that we have considered only price returns in this study. Total returns inclusive of dividends would certainly be higher; hence the alphas for the portfolios selected by the ANN are also likely to be higher.

Although the portfolios derived using ANN model produced positive alphas in many cases, it remains unclear what kind of pricing inefficiencies or risk exposures the network might be exploiting. In fact, returns for many of the ANN portfolios had no relationship with the known risk factors. At this point, given the ‘black box’ nature of the ANN, it is difficult to offer any explanation beyond the well-known ability of the ANN to capture ‘hidden’ relationship between inputs and outputs. It is not beyond the realms of possibility that ANN’s artificial intelligence is able to detect patterns in stock price movements which are not obvious to human intelligence and commonly dismissed as ‘noise’. We hope that future research in the fields of both asset pricing and artificial intelligence would be able to offer more insight.

References

Beale, M. H., Hagan, M. T., & Demuth, H. B. (2011). *Neural Network Toolbox (TM) 7 User's Guide*: The MathWorks Inc, Natick, MA.

Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *The Journal of Finance, 52*(1), 57-82.

Coolen, A.C.C. (1998). A Beginner's Guide to the Mathematics of Neural Networks. In `*Concepts for Neural Networks - A Survey,* Eds. Landau, L.J. and Taylor, J.G.: Springer 13-70.

Ellis, C., & Wilson, P. J. (2005). Can a Neural Network Property Portfolio Selection Process Outperform the Property Market. *Journal of Real Estate Portfolio Management, 11*(2), 105-121.

Garson, D.G. (1998). *Neural Networks: An Introductory Guide For Social Scientists*: Sage Publications, London.

Karaban, S., & Maguire, G. (2012). S&P Indices Versus Active Funds Scorecard (SPIVA Australia Scorecard), *S&P Dow Jones Indices*. Australia: S&P Dow Jones Indices LLC.

Malkiel, B. G. (2011). *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing*: WW Norton & Company, New York.

Tan, C. N. (1997). An Artificial Neural Networks Primer with Financial Applications Examples in Financial Distress Predictions and Foreign Exchange Hybrid Trading System: School of Information Technology, Bond University.

Vanstone, B., Finnie, G., & Hahn, T. (2010). Stockmarket trading using fundamental variables and neural networks. Paper presented at the *ICONIP 2010: 17th International Conference on Neural*

Inforamtion Processing. Sydney, Australia: ePublications@bond. from http://epublications.bond.edu.au/infotech_pubs/153

Formatted: PhD Normal, Tab stops: Not at 2.31 cm