SGH Warsaw School of Economics
Collegium of Economic Analysis

Time series events prediction using machine learning methods and the Efficient Markets Hypothesis

Adam Chudziak

Extended abstract of the PhD Thesis written
under the supervision of
dr. hab. Michał Ramsza, prof. SGH

Warsaw, September 2020

1 Motivation and research hypotheses

The Efficient Markets Hypothesis (EMH) has been discussed continuously since its inception and formulation by Fama (Fama 1970). The enduring interest in the question of whether it is possible to anticipate price movements, and exploit this knowledge for profit, is fuelled by the practical needs of traders in financial markets. Thus, it is only rational that whenever new forecasting tools emerge, they are swiftly employed for financial markets predictions. No different are the artificial intelligence and data mining methods, used increasingly frequently in the financial industry. However, the scrutiny the algorithms received has focused often on aspects more relevant for information technology, such as speed or data usage.

The aim of this thesis is to test the usefulness of selected machine learning and data mining methods for financial time series predictions. The main hypothesis states that certain markets do not satisfy the Efficient Markets Hypothesis.

H Certain financial markets do not satisfy the Efficient Markets Hypothesis.

We investigate that claim using two supporting hypotheses:

SH1 Machine learning techniques help forecasting events in financial time series, exposing undetected market inefficiencies.

SH2 Transformation of financial time series with representations using perceptually important points increases the machine learning forecast quality.

The evaluation of the EMH has caused trouble due to its vague statement (Thaler 2009; Guerrien and Gun 2011; Charron 2017). According to Thaler, the statement can be understood in two ways. The first is the "price is right" component — a claim that the prices reflect the true, intrinsic value of a security. The question whether such a value exists, what it is, and whether prices do reflect it, is a discussion we do not engage in. The second com-

ponent — the "no free lunch", or "technical efficiency" states that exploiting price forecasts for profit is impossible. The analysis presented in this thesis concentrates on the technical efficiency in the longer term. Short term EMH violations, including the high-frequency trading possibilities, are well known and out of the scope of the presented analysis.

It is common for techniques developed in computer science to be focused on different measures than those used in economics. Often they are focused on the algorithm speed, memory usage and a basic "does the algorithm perform the task" measures. This thesis aims to analyse the viability of using artificial neural networks combined with landmark representations as a forecasting tool for monthly stock price movements, using only past price data as predictor variables. The analysis is exhaustive as it uses long time series, which is rare in the relevant literature.

2 Thesis outline

Apart from the introduction and conclusions, the thesis consists of three chapters and two appendices.

The second chapter presents the Efficient Markets Hypothesis and showcases different angles of its assessment and critique. The issues with the formulation of the hypothesis are outlined, including the joint hypothesis problem and the issue with distinguishing the "technical" efficiency from the "fundamental" efficiency (Charron 2017). A variety of takes on empirical evaluation of the hypothesis is recounted.

The third chapter is concerned with the methodology of the thesis. The forecasting models tested in the thesis use a combination of the *landmark representations* (LR) for time series preprocessing and artificial neural networks for generating predictions. Construction of both the landmark representation algorithms and artificial neural networks are presented. The methods are well-

established, however, the methodology used is novel. Therefore, apart from presenting the methods, the meta-choices made are discussed.

Chapter four presents the results of the thesis and ends with a discussion of them in a theoretical context. It is split into three sections. In the first section, the landmark representations evaluation is presented, discussing the error, compression, and the capability of retaining information analysed with spectral methods. The second section shows the main results of the stock prices forecasting evaluation, including directional accuracy, testing profitability of trading strategies based on forecasts, and statistical tests of predictive abilities. The differences between the performance in periods of economic growth and recession, dated according to NBER data, are also noted and analysed. The chapter concludes with a discussion of the findings.

3 Methodology

As mentioned before, this thesis tests the long-term technical efficiency. The forecast models predict the stock price movement direction for the month ahead, based on a rolling price series window. The data consists of stock prices of NYSE stocks in the 1970–2015 period.

There is a string of research assessing methods popular with practitioners, such as the technical analysis (Lo, Mamaysky, and Wang 2000; Farias Nazário et al. 2017). Although this research is not testing the technical analysis per se, the ANN forecasting models used rely on a conceptually similar idea of recognising patterns in the data. Raw financial time series often exhibit volatility unnecessary and possibly detrimental to forecasting performance. Thus, an optional preprocessing step of transforming the series using landmark representations (LR) is taken (see figure 1).

The landmark representations can be defined as time series data mining algorithms which construct a piecewise-linear representation of a time series based on "importance" of particular points in that series. The roots of this

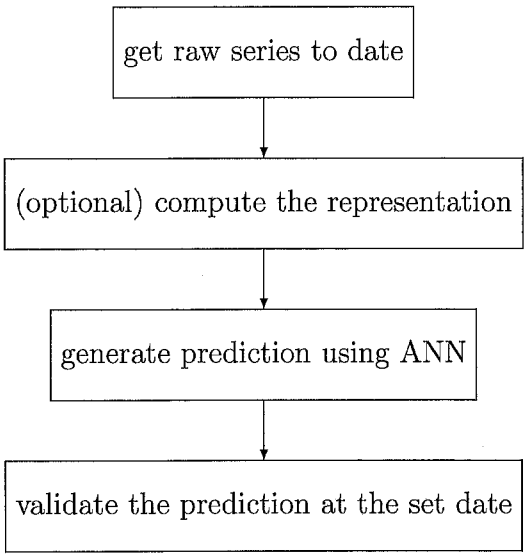

Figure 1: Schematic of the predictive process.

idea come from psychology, as it was observed that humans organise their memory around events of particular significance (Humphreys, Wiles, and Dennis 1994; Glenberg 1997). The LR algorithms operationalise this concept by either moving linearly through the time series and using predefined criteria deciding whether each point is important or not, or by creating a global order of points importance. The result is a subset of original observations, between which linear interpolation is performed. Figure 2 shows sample representations computed for the IBM stock price. Four algorithms are analysed: the Landmark Model (MDPP, Perng et al. 2000), the Major Maxima and Minima algorithm (MMM, Fink, Pratt, and Gandhi 2003), the Perceptually Important Points algorithm (PIP, Fu et al. 2008), and the the Critical Point Model (CPM, Bao 2008).

Artificial neural networks achieved major success in various fields. They are particularly capable in tasks requiring pattern recognition, such as computer vision. The pattern recognition capability makes them a good choice for testing predictability based on past stock prices. Stock market forecasting using neural networks has gathered a considerable attention, e.g., (White 1988; Chong, Han, and Park 2017; Sezer, Ozbayoglu, and Dogdu 2017). However,

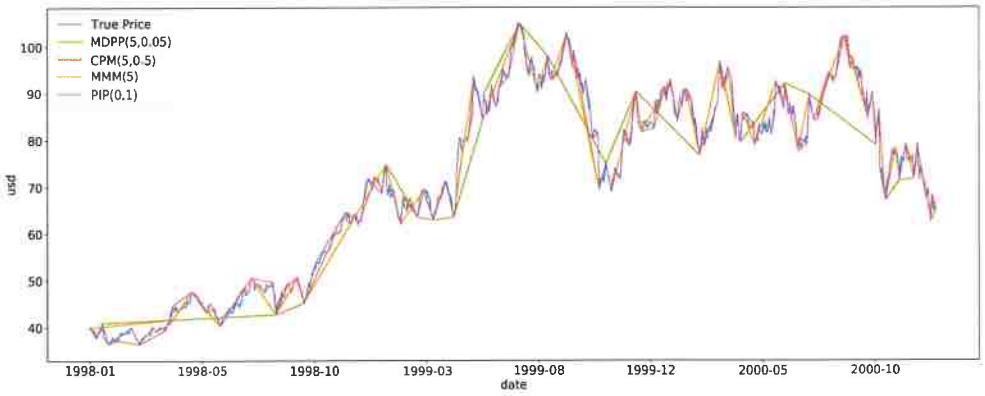

Figure 2: IBM stock adjusted daily closing price transformed via landmark representations: the Landmark Model (MDPP), the Critical Point Model (CPM), the Major Maxima and Minima algorithm (MMM) and the Perceptually Important Points algorithm (PIP). The parameters can be found in brackets.

many studies analyse short periods and do not evaluate the strength of the results. This thesis presents an analysis performed for a long period.

In forecasting, we use three types of neural network topologies: the multilayer perceptron (Rosenblatt 1958), the simple recurrent, and long-short term memory networks (Hochreiter and Schmidhuber 1997). There is a plethora of possible network choices. Even within a selected network topology, several meta-parameters, such as the number of layers and the number of neurons in each layer, have to be considered. The topologies chosen in this dissertation are among the most popular due to their successful deployment in many applications. Various combinations of ANNs and LRs are tested.

4 Results

This thesis analyses whether a forecasting model combining the landmark representations with the artificial neural networks can be successfully used for price movements predictions in a stock market. The results are mixed. We found a limited, varying over time predictability, which, however, we consider difficult to exploit for profit, due to model choice issues.

There are periods in which the predictability seems higher. The occasional rises of directional accuracy of predictions could constitute an opportunity window for traders. However, a method of identification of heightened predictability period and the models which allow to discover them is unclear, and further research can be pursued in this direction, see e.g., (Harvey et al. 2018).

Also, there are differences in the quality of forecasts between the tested stocks. For some of the stocks, none of the forecasting models could achieve satisfactory performance. This result suggests differences in the micro-level efficiency on the market.

One of the benchmark measures we used was the return on investment of a buy-and-hold trading strategy. In general, the proposed approach does not allow to outperform it. On the other hand, using the models leads to reduced risk of the held portfolios, measured by their variance. This concurs with the theoretical standpoint, see e.g., (Merton 1973).

We found that during the recession periods the models perform better, compared to the respective benchmarks, than in expansions. However, one has to bear in mind that since the US economy has been growing for the majority of the last half of century, the samples are imbalanced. We have much more data regarding expansions, and thus this result has to be taken with caution.

The landmark representations generally fulfil their task of denoising the series, while preserving their shape. Alas, the outcomes rely on the parameters and the optimal choice varies depending on the series. Out of the analysed representations one was discarded, as it was found in the tests that it does not preserve the frequency it theoretically should with given parameters.

The evidence regarding the effect of preprocessing using landmark representations on predictions is mixed. Depending on the stock, there was a varying degree of difference in predictions between a model which did not use landmark representation preprocessing and the models which did. The differences

between forecasts were tested using the Diebold-Mariano test (Diebold and Mariano 1995), and not all were statistically significant. Certain representations enhanced the predictions. However, no single representation was identified as the best one.

The results of the predictive ability statistical tests, notably the Pesaran-Timmermann test (Pesaran and Timmermann 1992), are also mixed. A statistically significant predictability coupled with outperforming the benchmark buy-and-hold return on investment in simulated portfolios occurred for a few models out of a wide array of models tested. However, it is unclear how these models could be identified beforehand, as the setup does not allow extensive validation.

The additional difficulty is posed by the differences in the predictability of various time series. Furthermore, the changing quality of forecasts over time suggest that predictability is not a constant characteristic of price series. It is in accordance with (Timmermann and Granger 2004), suggesting that a technique is likely to gain and lose predictive power, rather than perform on a constant level.

The quality of model performance can be at least partially attributed to the length of the forecast horizon. There might simply not be enough information in the rolling window of price series to perform a successful monthly forecast. This is in line with a view, that the predictable horizon for stock prices decreases (Schulmeister 2009).

We found limited predictability, which is hardly exploitable for profit. Even though several models surpassed the benchmarks, the model choice remains the problem in the proposed approach. Thus, we believe that no sufficient evidence to support the hypotheses was found.

A few possibilities to build on the research presented in this thesis can be outlined. The rapidly growing machine learning toolbox makes the refinement of the models one possibility. Investigating alternative ideas, such as

the use of image-like data preprocessing or attention mechanism in statistical learning, is an exhilarating perspective. The differences between the quality of forecasts among different stocks make a comparison of performance with another, possibly less efficient, developing market an interesting rout for further research.

5 Contributions

This thesis undertakes an ever actual problem of whether technological advances influence market efficiency. A novel methodology of combining the landmark representations and the artificial neural networks was introduced. The behaviour of the LR was analysed and their sensitivity to parameter choices was shown. The forecasting capabilities of the ANN were tested systematically. The found predictability does not, in the author's view, imply inefficiency. The results have highly applicable consequences for practitioners. From their point of view, this dissertation shows that over the long run any particular method does not systematically lead to profit in isolation of other tools. The key question then concerns the method of selecting the right model for the right asset at the right time. Furthermore, in the light of the findings, the generality of selective results published in the neural network literature becomes questionable.

References

Bao, Depei (2008). "A generalized model for financial time series representation and prediction". In: *Applied Intelligence* 29.1, pp. 1–11. ISSN: 0924-669X, 1573-7497.

Charron, Jacques-Olivier (2017). "Inefficient Debate. The EMH, the "Remarkable Error" and a Question of Point of View". In: *Accounting, Economics, and Law: A Convivium* 7.3. ISSN: 2152-2820.

Chong, Eunsuk, Han, Chulwoo, and Park, Frank C. (2017). "Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies". In: *Expert Systems with Applications* 83, pp. 187–205. ISSN: 0957-4174.

Diebold, Francis X. and Mariano, Roberto S. (1995). "Comparing Predictive Accuracy". In: *Journal of Business & Economic Statistics* 13.3, pp. 253–263. ISSN: 0735-0015.

Fama, Eugene F. (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work". In: *The Journal of Finance* 25.2, pp. 383–417. ISSN: 00221082.

Farias Nazário, Rodolfo Toríbio et al. (2017). "A literature review of technical analysis on stock markets". In: *The Quarterly Review of Economics and Finance* 66, pp. 115–126. ISSN: 1062-9769.

Fink, E., Pratt, K.B., and Gandhi, H.S. (2003). "Indexing of time series by major minima and maxima". In: *SMC'03 Conference Proceedings. 2003 IEEE International Conference on Systems, Man and Cybernetics. Conference Theme - System Security and Assurance (Cat. No.03CH37483).* SMC '03 2003 IEEE International Conference on Systems, Man and Cybernetics. Vol. 3. Washington, DC, USA: IEEE, pp. 2332–2335. ISBN: 978-0-7803-7952-7.

Fu, Tak-chung et al. (2008). "Representing financial time series based on data point importance". In: *Engineering Applications of Artificial Intelligence* 21.2, pp. 277–300. ISSN: 09521976.

Glenberg, Arthur M. (1997). "What memory is for". In: *Behavioral and Brain Sciences* 20.1, pp. 1–19. ISSN: 0140-525X, 1469-1825.

Guerrien, Bernard and Gun, Ozgur (2011). "Efficient Market Hypothesis: What are we talking about?" In: *Real-World Economics Review* 56, p. 12.

Harvey, David et al. (2018). *Detecting Regimes of Predictability in the U.S. Equity Premium.* Essex Finance Centre Working Paper. University of Essex, Essex Business School.

Hochreiter, Sepp and Schmidhuber, Jürgen (1997). "Long Short-Term Memory". In: *Neural Computation* 9.8, pp. 1735–1780. ISSN: 0899-7667, 1530-888X.

Humphreys, Michael S., Wiles, Janet, and Dennis, Simon (1994). "Toward a theory of human memory: Data structures and access processes". In: *Behavioral and Brain Sciences* 17.4, pp. 655–667. ISSN: 0140-525X, 1469-1825.

Lo, Andrew W., Mamaysky, Harry, and Wang, Jiang (2000). "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation". In: *The Journal of Finance* 55.4, pp. 1705–1765. ISSN: 1540-6261.

Merton, Robert C. (1973). "An Intertemporal Capital Asset Pricing Model". In: *Econometrica* 41.5, pp. 867–887. ISSN: 0012-9682.

Perng, C.-S. et al. (2000). "Landmarks: a new model for similarity-based pattern querying in time series databases". In: *Data Engineering, 2000. Proceedings. 16th International Conference on.* IEEE, pp. 33–42.

Pesaran, M. Hashem and Timmermann, Allan (1992). "A Simple Nonparametric Test of Predictive Performance". In: *Journal of Business & Economic Statistics* 10.4, pp. 461–465. ISSN: 0735-0015, 1537-2707.

Rosenblatt, F. (1958). "The perceptron: A probabilistic model for information storage and organization in the brain." In: *Psychological Review* 65.6, pp. 386–408. ISSN: 1939-1471, 0033-295X.

Schulmeister, Stephan (2009). "Profitability of technical stock trading: Has it moved from daily to intraday data?" In: *Review of Financial Economics* 18.4, pp. 190–201. ISSN: 1058-3300.

Sezer, Omer Berat, Ozbayoglu, Murat, and Dogdu, Erdogan (2017). "A Deep Neural-Network Based Stock Trading System Based on Evolutionary Optimized Technical Analysis Parameters". In: *Procedia Computer Science* 114, pp. 473–480. ISSN: 18770509.

Thaler, Richard (2009). "Markets can be wrong and the price is not always right". In: *Financial Times*, p. 4.

Timmermann, Allan and Granger, Clive W. J. (2004). "Efficient market hypothesis and forecasting". In: *International Journal of Forecasting* 20.1, pp. 15–27. ISSN: 0169-2070.

White, Halbert (1988). "Economic prediction using neural networks: the case of IBM daily stock returns". In: *IEEE International Conference on Neural Networks*. Proceedings of 1993 IEEE International Conference on Neural Networks (ICNN '93). San Diego, CA, USA: IEEE, 451–458 vol.2. ISBN: 978-0-7803-0999-9.