APPLICATIONS IN COMPUTATIONAL FINANCE WITH A FOCUS ON APPROXIMATION OF FINANCIAL TIME SERIES BY NEUROCOMPUTING

Der Wirtschaftswissenschaftlichen Fakultät der

GOTTFRIED WILHELM LEIBNIZ UNIVERSITÄT HANNOVER

zur Erlangung des akademischen Grades

Doktor der Wirtschaftswissenschaften

— Doctor rerum politicarum —

vorgelegte Dissertation

von

Diplom-Ökonom CHRISTIAN VON SPRECKELSEN

[redacted]

2013

Betreuer/Gutachter: PROF. DR. MICHAEL H. BREITNER
Zweitgutachter: PROF. DR. DANIEL RÖSCH
Vorsitz Prüfungskomission: JUN.-PROF. DR. HANS-JÖRG VON METTENHEIM
Wissenschaftliche Mitarbeiterin: DR. UTE LOHSE

Dedicated to Sarah and Luisa

ABSTRACT

My dissertation shows how neural networks can be used in order to achieve more accurate approximation as well as better decision making in financial markets. In order to study its approximation ability for computational finance, I perform different empirical investigations. First, neural networks are suitable for approximating price functions of assets. I present empirical results for pricing and hedging FX options. Second, the usage of neural computing for forecasting financial time series is investigated, where neural networks compete with traditional time series models. I show empirical studies about the maritime freight rates market and the Chinese FX market. Above all mentioned techniques remains the question of neuronal computing application in the financial industry. In a last step I thus propose the implementation and design of a financial decision support system with neural networks. Nevertheless, I also expose limitations and further research topics in the area of neural networks, which could improve neural networks applications in computational economics in the future.

ZUSAMMENFASSUNG

Meine Dissertation zeigt, wie Neuronale Netze für eine bessere Entscheidungsfindung an den Finanzmärkten eingesetzt werden können. Um die Approximationsfähigkeit für den Einsatz in Computational Finance zu analysieren, habe ich verschiedene empirische Untersuchungen durchgeführt. Zunächst eignen sich Neuronale Netze für die Approximation der Preisfunktion von Assets. Ich zeige empirische Ergebnisse für die Preisfindung und Absicherung von FX-Optionen. Zweitens wird der Einsatz von Neural Computing für die Prognose finanzieller Zeitreihen untersucht, wo Neuronale Netze mit traditionellen Zeitreihenmodellen konkurrieren. Dazu zeige ich empirische Analysen über den maritimen Frachtratenmarkt und den chinesischen Devisenmarkt. Über allen erwähnten Techniken bleibt die Frage der Anwendung von Neural Computing in der Finanzmarktindustrie. Ich schlage daher in einem letzten Schritt die Umsetzung und das Design eines Financial Decision Support System mit Neuronalen Netzen vor. Dennoch stelle ich auch Einschränkungen und weitere Forschungsthemen für den Einsatz von Neural Computing im Bereich von Computational Economics in Zukunft verbessern könnte.

ACKNOWLEDGMENTS

It would not have been possible to write this dissertation without the help and support of the kind people of my personal environment. It is hardly possible to mention all of them and their individual support in my acknowledgments, but I really like to let all of them know, how important and special they are for me. Above all, I would like to thank my wife Sarah for her unequivocal support throughout, as always, for which my mere expression of thanks likewise does not suffice. I would also thank my parents and brother for their great patience at all times.

Foremost, I would like to express my sincere gratitude to my advisor Prof. Michael H. Breitner and for the continuous support of my PhD study and research, for his patience, motivation, and knowledge. The good advice, support and friendship of my second supervisor, Assistance Prof. Hans-Jörg von Mettenheim, has been invaluable on both an academic and a personal level, for which I am extremely grateful. His guidance helped me in all times of research and writing of this dissertation.

I also thank Prof. Dr. Daniel Rösch, my second referee, for accepting the examination of my dissertation.

I would like to acknowledge the support of my employer, the Norddeutsche Landesbank Girozentrale, particularly for giving me the opportunity to combine both my work and my dissertation. I firmly believe that it arises useful synergy effects from my dissertation. Many people contributed in different ways to my dissertation. It is impossible to name them all. Without making a claim of completeness I mention some of them. In particular, I thank Christoph Wegener and Tobias Basse for useful comments and discussions in the field of time series analysis. I am most grateful to Jana Wiegreffe and my wife Sarah for proofreading this dissertation.

CONTENTS

EXECUTIVE SUMMARY

Estimating an underlying relationship from a given finite input-output data set - or more precisely: function approximation involves - has been the fundamental problem for a variety of applications in financial engineering. Nowadays, feedforward neural networks such as Multilayer perceptron (MLP) have been widely used as an alternative approach to function approximation since they provide a generic functional representation. They have been shown to be capable of approximating any continuous function with arbitrary accuracy.

This dissertation shows how neural networks can be used in order to achieve more accurate approximation as well as better decision making in financial markets. The importance of better market price approximation or synthesis, forecasting, and the relationship between spot and derivative markets for better decision making, in the light of increasing financial market volatility and internationalized capital flows, cannot be over exaggerated. In order to study its approximation ability for computational economics, I perform different empirical investigations. Figure 0.1 summarizes the organization of my dissertation.

NETWOK APPROXIMATION BY THEORY

The universal approximation theorem of Cybenko (1989) and Hornik (1989) provides the latent basis of my empirical studies. Artificial neural networks can be mathematically shown to be universal function approximators. This means that NNs can automatically approximate whatever functional form characterizes the data best. Since it is my goal to extract an alternative option pricing function by market observations, I focus on MLP that are applicable to non-linear regression problems. I follow the argumentation of Hornik (1989), that feedforward networks with only one hidden layer and a linear output unit are able to approximate simultaneously its unknown derivatives up to an arbitrary degree of accuracy. This characteristic is substantial since the partial derivatives of a pricing formula are needed for the hedging of option positions.

I perform my network training with the Fast Approximation with Universal Neural Networks (FAUN) neurosimulator. As described in Mettenheim and Breitner (2010) two reasons make FAUN suitable for HFT. First, the FAUN neurosimulator uses fine-grained parallelization. This allows easily achieved speedups on dual and quad core CPUs. FAUN also features coarse-grained parallelization using an easy to install grid computing client. It is possible to use clusters of heterogeneous

Figure 0.1: Overview

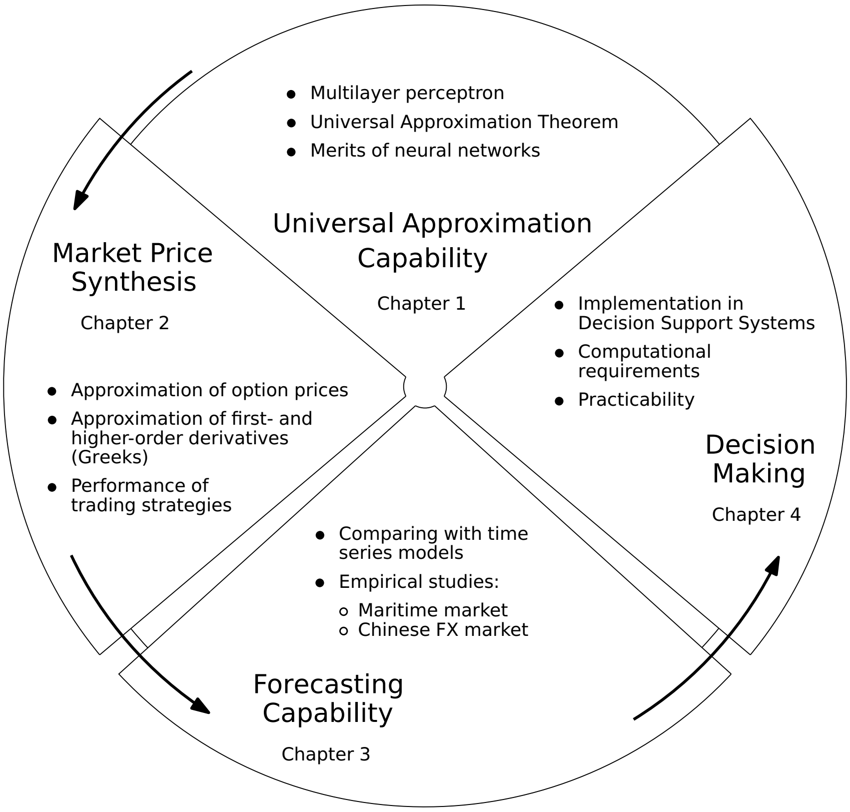

workstations. Second, using reverse accumulation and matrix algorithms allow a very efficient computation.

MARKET PRICE SYNTHESIS

Neural networks are information processing tools commonly used for function approximation and classification. They offer an alternative way of developing option pricing and hedging models. Their particular strength lies in their ability to approximate highly non-linear and multivariate relationships without the restrictive assumptions implicit in parametric approaches. This property of neural networks makes them attractive for problems such as pricing and hedging options. Moreover, they are adaptive and respond to structural changes in the financial markets. The drawback of this approach is that it is highly data driven, requiring large quantities of historical prices.

I present empirical results for pricing and hedging FX options. The empirical results confirm the ability of neural networks for universal approximation. Subsequent studies mostly investigated daily equity index options data for option pricing approximations. Despite the high liquidity of FX options markets, there is no noticeable investigation about pricing FX options with neural networks in a HFT-context.

Hence, I build on prior investigations, but I extend my studies paper B and paper E with a run-time trading process in order to uncover special characteristics of high-frequency data. In particular, I pose the following challenge: If option prices were truly determined by the theoretical model exactly, can the closed-form formula be estimated by learning networks with a sufficient degree of accuracy to be of practical use? Furthermore, can both models be implemented in an automatic HFT trading process, in which a signal must be precise enough to trigger trades in a fraction of a second?

To assess the approximation capability I use two big data sets. On the one hand there is a full high-frequency data set of cleared 118,291 quotes of an EUR/USD option on currency futures with various strike prices available. On the other hand I generate more than 20,000 simulated intra-day option prices to get a broader range of data.

To assess the potential value of network pricing formulas in HFT, I implement two different investigations: First, paper B and paper E perform a rolling 15 minutes out-of-sample interval for each trading day to assess the models pricing ability. The derived approximation function is then used to perform a delta-hedging examination. All results are benchmarked using a theoretical closed-form model for pricing options on futures. Second, in order to carry out the approximation capability of the network function and its partial derivatives the network in paper A trains on a simulated data set without any rolling-window technique in order to investigate the numerical approximation of option price functions and their derivatives. I am also interested in the question of whether the data availability is crucial for a better approximation.

FORECASTING CAPABILITY

The usage of neural computing for forecasting financial time series is investigated, where neural networks compete with conventional time series models. Theoretically, the efficient market hypothesis implies that in an efficient market, it is impossible to obtain better predictions using forecasting methods because the observable price already reflects all available information and price fluctuations that will occur in the future randomly. In reality, however, systematic patterns might be found in financial time series.

First, I show empirical studies about the maritime spot and derivatives freight rates market. In paper F, paper G and paper H I perform several forecasting techniques in order to examine the forecasting ability of freight rates. I find a lack of jointly spot and forward forecasting investigations with neural networks. Thus, I extend my study on freight derivatives and a wider range of time series models. The main objective of this paper is to investigate neural networks prediction ability for maritime business forecasting and provide a practical framework for actual forecasting and trading applications of neural computing.

I sample daily prices of the International Maritime Exchange (Imarex) TD3 and TD5 freight forward contracts. These contracts are written on daily spot rates for TD3 and TD5 published by the Baltic Exchange. The spot and Forward Freight Agreement (FFA) data is available from 5 April 2004 to 1 April 2011. I investigate short-term forecasts of spot and FFA prices in the market in order to make inferences about the efficiency and usefulness of FFA rates. The question arises: Are forward rates expectations of spot rates? I consider both univariate and multivariate model specifications fitted with lagged spot freight rates returns $\Delta\hat{S}_t$ and forward rates returns $\Delta\hat{F}_t$.

Another interesting research object is the very unique Chinese FX market, which exhibits a dual characteristic of the market. The uniqueness comes from the two separated markets for the Renminbi (RMB), namely the onshore Yuan (CNY) and offshore Yuan (CNH) market. The main goal of paper C is to gain insights in the comparatively new market for offshore RMB and to detect first indications for feasible forecasting models for the onshore RMB respectively to improve CNY spot forecasts. I employ a simple GARCH model as well as neural networks. I do also analyze the somewhat older NDF market for which Ding et al. (2012) found a strong relationship with the CNY spot rate. As their work deals with the three RMB markets until June 2011 and since then the CNH market grew quite rapidly and seems to be replacing the NDF market, I lay our main focus on the CNH market.

I collect daily exchange rate data for onshore spot CNY, offshore spot CNH, one-month offshore NDF and CNH forward rates from Bloomberg. The sample period spans 08 September 2010 to 20 March 2013. All forecast models are separated in univariate and multivariate classes: The univariate models consist of single series of CNY, CNH and their spread. I exclusively analyze the CNY in a multivariate way by incorporation of the one-month forward rates NDF and CNH respectively.

DECISION MAKING

Above all mentioned techniques remains the question of neuronal computing application in the financial industry. In a last step I thus propose the implementation and design of a financial decision support system with neural networks, which is a more business informatics oriented discussion. The merits of neural networks especially for high-dimensional problems are shown.

I present steps towards a model-driven DSS to pricing option on currency futures, which can be embedded in a high-frequency trading process. In order to develop an appropriate DSS, I use the design science methodology of Hevner et al. (2004). Efficient implementation of trading algorithms is crucial, because a vast amount of data has to be processed in very short time.

MAIN CONTRIBUTIONS

In summary, I have attempted to provide empirical evidence for neural networks capability to approximate financial time series. Main contributions are:

- Model option prices derived from NN can synthesize HFT option market prices in a similar manner, but in a simultaneous way and with a more parsimonious input specification. There is e.g. no need of volatility or interest estimation.
- If market liquidity exists, which is equivalent to full data availability in a particular state space, learning networks are capable to approximate first- and higher-order partial derivatives with a sufficient accuracy. But the approximation accuracy decreases with higher-order partial derivatives.
- However, I can not confirm the hypothesis that once a predominant network approximation is found for pricing purposes, the same could be applied for hedging. I have to notice that it is an exhausting balancing act for learning systems to apply the delivered pricing approximation function on unknown hedge parameters.
- In case of forecasting financial time series neural network results are comparable to those of the other models. Some regularities from two different financial markets:
- Tanker freight rates market: Changes in spot rates are explained by autocorrelation and by changes in the forward rates; but: changes in forward rates are not explained by past changes in spot rates. There is, however, a highly significant autocorrelation in forward rates that is difficult to conciliate with efficient markets. These results imply that the futures prices contain valuable information about future spot rates.
- Chinese FX market: Our results do not support our assumption of a parity between the CNY and CNH. On the one hand the fact that the used forecasting methods do not outperform the naïve RW forecasts points to the direction that the price movements in the Chinese FX markets are similar to the movements in developed economies' FX markets, which are said to be rather efficient. On

the other hand I found strong evidence that structural breaks do exist in the RMB markets.

- Neural networks are a suitable core engine for a model-driven DSS embedded in a high-frequency trading process and can support trading decisions.

Hence, this dissertation provides empirical evidence that neural networks may be put to work for more accurate approximation and for better decision making in financial markets.

EVALUATION CRITERIA

In evaluating my empirical studies, there are still some questions left: First, can the empirical results be generalized? Second, are there any restrictions to a practical implementation, which have not been taken into account? For this purpose, I have identified three assessment criteria as shown in figure 0.2. I will give answers in detail to the two questions mentioned above in chapter 5.

Figure 0.2: Assessment criteria

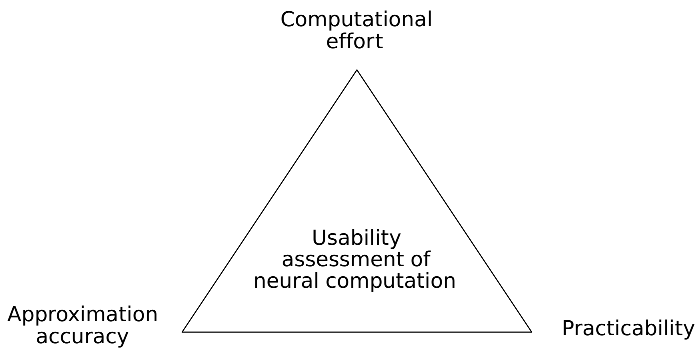

In summary, it can be stated that:

- All empirical investigations in each case refer only to certain time periods and assets. There is a need for further evidence to confirm a generalization or robustness of the models.

- The approximation of neural networks suffer from inhomogeneous data density, in particular when trainable data is rare.

- To implement large and effective software neural networks, much processing and storage resources need to be committed. Neural network systems will often need to simulate the transmission of signals through many of these connections and their associated neurons - which must often be matched with incredible amounts of CPU processing power and time.

The good news are: I also expose further research topics in doing with neural networks, which could improve neural networks applications in computational economics in future.