ACADEMIC PRESS ADVANCED FINANCE SERIES

ELSEVIER
ACADEMIC PRESS

Neural Networks in Finance

GAINING PREDICTIVE EDGE IN THE MARKET

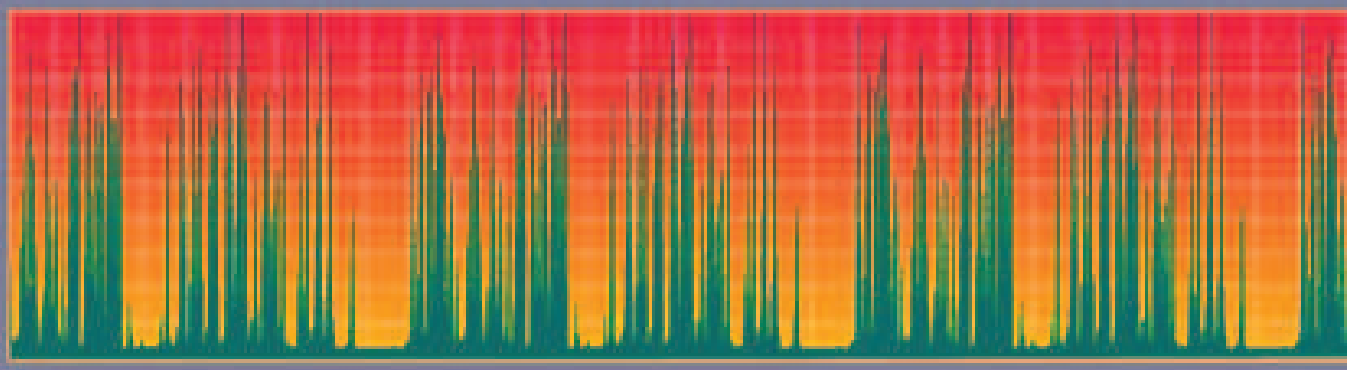

PAUL D. McNELIS

Neural Networks in Finance: Gaining Predictive Edge in the Market

Neural Networks in Finance: Gaining Predictive Edge in the Market

Paul D. McNelis

ELSEVIER
ACADEMIC
PRESS

Amsterdam • Boston • Heidelberg • London • New York • Oxford
Paris • San Diego • San Francisco • Singapore • Sydney • Tokyo

Elsevier Academic Press
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
525 B Street, Suite 1900, San Diego, California 92101-4495, USA
84 Theobald's Road, London WC1X 8RR, UK

This book is printed on acid-free paper.

Library of Congress Cataloging-in-Publication Data
McNelis, Paul D.
Neural networks in finance : gaining predictive edge in the market / Paul D. McNelis.
p. cm.
1. Finance–Decision making–Data processing. 2. Neural networks (Computer science) I. Title.
HG4012.5.M38 2005
332′.0285′632–dc22
2004022859

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN: 0-12-485967-4

For all information on all Elsevier Academic Press publications
visit our Web site at www.books.elsevier.com

Printed in the United States of America
04 05 06 07 08 09 9 8 7 6 5 4 3 2 1

Contents

Preface

> Adjusting to the power of the Supermarkets and the Electronic Herd requires a whole different mind-set for leaders ...
>
> Thomas Friedman, *The Lexus and the Olive Tree*, p. 138

Questions of finance and market success or failure are first and foremost *quantitative*. Applied researchers and practitioners are interested not only in predicting the *direction of change* but also *how much* prices, rates of return, spreads, or likelihood of defaults will change in response to changes in economic conditions, policy uncertainty, or waves of bullish and bearish behavior in domestic or foreign markets. For this reason, the premium is on both the *precision* of the estimates of expected rates of return, spreads, and default rates, as well as the *computational ease* and *speed* with which these estimates may be obtained. Finance and market research is both empirical and computational.

Peter Bernstein (1998) reminds us in his best-selling book *Against the Gods,* that the driving force behind the development of probability theory was the *precise* calculation of odds in games of chance. Financial markets represent the foremost "games of chance" today, and there is no reason to doubt that the precise calculation of the odds and the risks in this global game is the driving force in quantitative financial analysis, decision making, and policy evaluation.

Besides precision, speed of computation is of paramount importance in quantitative financial analysis. Decision makers in business organizations or in financial institutions do not have long periods of time to wait before having to commit to buy or sell, set prices, or make investment decisions.

While the development of faster and faster computer hardware has helped to minimize this problem, the specific way of conceptualizing problems continues to play an important role in how quickly reliable results may be obtained. Speed relates both to computational hardware and software.

Forecasting, classification of risk, and dimensionality reduction or distillation of information from dispersed signals in the market, are three tools for effective portfolio management and broader decision making in volatile markets yielding "noisy" data. These are not simply academic exercises. We want to forecast more accurately to make better decisions, such as to buy or sell particular assets. We are interested in how to measure risk, such as classifying investment opportunities as high or low risk, not only to rebalance a portfolio from more risky to less risky assets, but also to price or compensate for risk more accurately.

Even in a policy context, decisions have to be made in the context of many disparate signals coming from volatile or evolving financial markets. As Othmar Issing of the European Central Bank noted, "disturbances have to be evaluated as they come about, according to their potential for propagation, for infecting expectations, for degenerating into price spirals" [Issing (2002), p. 21].

How can we efficiently distill information from these market signals for better diversification and effective hedging, or even better stabilization policy? All of these issues may be addressed very effectively with neural network methods. Neural networks help us to approximate or "engineer" data, which, in the words of Wolkenhauer, is both the "art of turning data into information" and "reasoning about data in the presence of uncertainty" [Wolkenhauer (2001), p. xii]. This book is about predictive accuracy with neural networks, encompassing forecasting, classification, and dimensionality reduction, and thus involves data engineering.[1]

The benchmark against which we compare neural network performance is the time-honored linear regression model. This model is the starting point of any econometric modeling course, and is the standard workhorse in econometric forecasting. While there are doubtless other nonlinear methods against which we can compare the performance of neural network methods, we choose the linear model simply because it is the most widely used and most familiar method of applied researchers for forecasting. The neural network is the nonlinear alternative.

Most of modern finance theory comes from microeconomic optimization and decision theory under uncertainty. Economics was originally called the "dismal science" in the wake of John Malthus's predictions about the relative rates of growth of population and food supply. But economics can be dismal in another sense. If we assume that our real-world observations

[1] Financial engineering more properly focuses on the design and arbitrage-free pricing of financial products such as derivatives, options, and swaps.

come from a linear data generating process, that most shocks are from an underlying normal distribution and represent small deviations around a steady state, then the standard tools of classical regression are perfectly appropriate. However, making use of the linear model with normally generated disturbances may lead to serious misspecification and mispricing of risk if the real world deviates significantly from these assumptions of linearity and normality. This is the dismal aspect of the benchmark linear approach widely used in empirical economics and finance.

Neural network methods, coming from the brain science of cognitive theory and neurophysiology, offer a powerful alternative to linear models for forecasting, classification, and risk assessment in finance and economics. We can learn once more that economics and finance need not remain "dismal sciences" after meeting brain science.

However, switching from linear models to nonlinear neural network alternatives (or any nonlinear alternative) entails a cost. As we discuss in succeeding chapters, for many nonlinear models there are no "closed form" solutions. There is the ever-present danger of finding locally optimal rather than globally optimal solutions for key problems. Fortunately, we now have at our disposal evolutionary computation, involving the use of genetic algorithms. Using evolutionary computation with neural network models greatly enhances the likelihood of finding globally optimal solutions, and thus predictive accuracy.

This book attempts to give a balanced critical review of these methods, accessible to students with a strong undergraduate exposure to statistics, econometrics, and intermediate economic theory courses based on calculus. It is intended for upper-level undergraduate students, beginning graduate students in economics or finance, and professionals working in business and financial research settings. The explanation attempts to be straightforward: what these methods are, how they work, and what they can deliver for forecasting and decision making in financial markets. The book is not intended for ordinary M.B.A. students, but tries to be a technical exposé of a state-of-the-art theme for those students and professionals wishing to upgrade their technical tools.

Of course, readers will have to stretch, as they would in any good challenging course in statistics or econometrics. Readers who feel a bit lost at the beginning should hold on. Often, the concepts become much clearer when the applications come into play and when they are implemented computationally. Readers may have to go back and do some further review of their statistics, econometrics, or even calculus to make sense of and see the usefulness of the material. This is not a bad thing. Often, these subjects are best learned when there are concrete goals in mind. Like learning a language, different parts of this book can be mastered on a need-to-know basis.

There are several excellent books on financial time series and financial econometrics, involving both linear and nonlinear estimation and

forecasting methods, such as Campbell, Lo, and MacKinlay (1997); Frances and van Dijk (2000); and Tsay (2002). In additional to very careful and user-friendly expositions of time series econometrics, all of these books have introductory treatments of neural network estimation and forecasting. This work follows up these works with expanded treatment, and relates neural network methods to the concepts and examples raised by these authors.

The use of the neural network and the genetic algorithm is by its nature very computer intensive. The numerical illustrations in this book are based on the MATLAB programming code. These programs are available on the website at Georgetown University, www.georgetown.edu/mcnelis. For those who do not wish to use MATLAB but want to do computation, Excel add-in macros for the MATLAB programs are an option for further development. Making use of either the MATLAB programs or the Excel add-in programs will greatly facilitate intuition and comprehension of the methods presented in the following chapters, and will of course enable the reader to go on and start applying these methods to more immediate problems. However, this book is written with the general reader in mind — there is no assumption of programming knowledge, although a few illustrative MATLAB programs appear in the text. The goal is to help the reader understand the logic behind the alternative approaches for forecasting, risk analysis, and decision-making support in volatile financial markets.

Following Wolkenhauer (2001), I struggled to impose a linear ordering on what is essentially a web-like structure. I know my success in this can be only partial. I encourage readers to skip ahead to find more illustrative examples of the concepts raised in earlier parts of the book in succeeding chapters.

I show throughout this book that the application of neural network approximation coupled with evolutionary computational methods for estimation have a predictive edge in out-of-sample forecasting. This predictive edge is relative to standard econometric methods. I do not claim that this predictive edge from neural networks will always lead to opportunities for profitable trading [see Qi (1999)], but any predictive edge certainly enhances the chance of finding such opportunities.

This book grew out of a large and continuing series of lectures given in Latin America, Asia, and Europe, as well as from advanced undergraduate seminars and graduate-level courses at Georgetown University and Boston College. In Latin America, the lectures were first given in São Paulo, Brazil, under the sponsorship of the Brazilian Association of Commercial Bankers (ABBC), in March 1996. These lectures were offered again in March 1997 in São Paulo, in August 1998 at Banco do Brasil in Brasilia, and later that year in Santiago, Chile, at the Universidad Alberto Hurtado.

In Asia and Europe, similar lectures took place at the Monetary Policy and Economic Research Department of Bank Indonesia, under the sponsorship of the United States Agency for International Development, in

January 1996. In May 1997 a further series of lectures on this subject took place under the sponsorship of the Programme for Monetary and Financial Studies of the Department of Economics of the University of Melbourne, and in March of 1998 a similar course was offered at the Facultat d'Economia of the Universitat Ramon Llull sponsored by the Callegi d'Economistes de Calalunya in Barcelona.

The Center for Latin American Economics of the Research Department of the Federal Reserve Bank of Dallas provided the opportunity in the autumn of 1997 to do some of the initial formal research for the financial examples illustrated in this book. In 2003 and early 2004, the Hong Kong Institute of Monetary Research was the center for a summer of research on applications of neural network methods for forecasting deflationary cycles in Hong Kong, and in 2004 the School of Economics and Social Sciences at Singapore Management University and the Institute of Mathematical Sciences at the National University of Singapore were hosts for a seminar and for research on nonlinear principal components

Some of the most useful inputs for the material for this book came from discussions with participants at the International Joint Conference on Neural Networks (IJCNN) meetings in Washington, DC, in 2001, and in Honolulu and Singapore in 2002. These meetings were eye-openers for anyone trained in classical statistics and econometrics and illustrated the breadth of applications of neural network research.

I wish to thank my fellow Jesuits at Georgetown University and in Washington, DC, who have been my "company" since my arrival at Georgetown in 1977, for their encouragement and support in my research undertakings. I also acknowledge my colleagues and students at Georgetown University, as well as economists at the universities, research institutions, and central banks I have visited, for their questions and criticism over the years. We economists are not shy about criticizing one another's work, but for me such criticism has been more gain than pain. I am particularly grateful to the reviewers of earlier versions of this manuscript for Elsevier Academic Press. Their constructive comments gave me new material to pursue and enhanced my own understanding of neural networks.

I dedicate this book to the first member of the latest generation of my clan, Reese Anthony Snyder, born June 18, 2002.