ADVANCES IN ECONOMETRICS
VOLUME 20

ECONOMETRIC ANALYSIS OF FINANCIAL AND ECONOMIC TIME SERIES - PART B

THOMAS B. FOMBY
DEK TERRELL
Editors

ECONOMETRIC ANALYSIS OF FINANCIAL AND ECONOMIC TIME SERIES

ADVANCES IN ECONOMETRICS

Series Editors: Thomas B. Fomby and R. Carter Hill

Recent Volumes:

Volume 15:	Nonstationary Panels, Panel Cointegration, and Dynamic Panels, Edited by Badi Baltagi
Volume 16:	Econometric Models in Marketing, Edited by P. H. Franses and A. L. Montgomery
Volume 17:	Maximum Likelihood of Misspecified Models: Twenty Years Later, Edited by Thomas B. Fomby and R. Carter Hill
Volume 18:	Spatial and Spatiotemporal Econometrics, Edited by J. P. LeSage and R. Kelley Pace
Volume 19:	Applications of Artificial Intelligence in Finance and Economics, Edited by J. M. Binner, G. Kendall and S. H. Chen

ADVANCES IN ECONOMETRICS VOLUME 20, PART B

ECONOMETRIC ANALYSIS OF FINANCIAL AND ECONOMIC TIME SERIES

EDITED BY

THOMAS B. FOMBY

Department of Economics, Southern Methodist University, Dallas, TX 75275

DEK TERRELL

Department of Economics, Louisiana State University, Baton Rouge, LA 70803

ELSEVIER
JAI

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

ELSEVIER B.V.
Radarweg 29
P.O. Box 211
1000 AE Amsterdam,
The Netherlands

ELSEVIER Inc.
525 B Street, Suite 1900
San Diego
CA 92101-4495
USA

ELSEVIER Ltd
The Boulevard, Langford
Lane, Kidlington
Oxford OX5 1GB
UK

ELSEVIER Ltd
84 Theobalds Road
London
WC1X 8RR
UK

First edition 2006

British Library Cataloguing in Publication Data
A catalogue record is available from the British Library.

ISBN-10: 0-7623-1273-4
ISBN-13: 978-0-7623-1273-3
ISSN: 0731-9053 (Series)

∞ The paper used in this publication meets the requirements of ANSI/NISO Z39.48-1992 (Permanence of Paper).
Printed in The Netherlands.

CONTENTS

This page intentionally left blank

DEDICATION

Volume 20 of *Advances in Econometrics* is dedicated to Rob Engle and Sir Clive Granger, winners of the 2003 Nobel Prize in Economics, for their many valuable contributions to the econometrics profession. The Royal Swedish Academy of Sciences cited Rob "for methods of analyzing economic time series with time-varying volatility (ARCH)" while Clive was cited "for methods of analyzing economic time series with common trends (cointegration)." Of course, these citations are meant for public consumption but we specialists in time series analysis know their contributions go far beyond these brief citations. Consider *some* of Rob's other contributions to our literature: Aggregation of Time Series, Band Spectrum Regression, Dynamic Factor Models, Exogeneity, Forecasting in the Presence of Cointegration, Seasonal Cointegration, Common Features, ARCH-M, Multivariate GARCH, Analysis of High Frequency Data, and CAViaR. *Some* of Sir Clive's additional contributions include Spectral Analysis of Economic Time Series, Bilinear Time Series Models, Combination Forecasting, Spurious Regression, Forecasting Transformed Time Series, Causality, Aggregation of Time Series, Long Memory, Extreme Bounds, Multi-Cointegration, and Non-linear Cointegration. No doubt, their Nobel Prizes are richly deserved. And the 48 authors of the two parts of this volume think likewise. They have authored some very fine papers that contribute nicely to the same literature that Rob's and Clive's research helped build.

For more information on Rob's and Clive's Nobel prizes you can go to the Nobel Prize website http://nobelprize.org/economics/laureates/2003/index.html. In addition to the papers that are contributed here, we are publishing remarks by Rob and Clive on the nature of innovation in econometric research that were given during the Third Annual *Advances in Econometrics* Conference at Louisiana State University in Baton Rouge, November 5–7, 2004. We think you will enjoy reading their remarks. You come away with the distinct impression that, although they may claim they were "lucky" or "things just happened to fall in place," having the orientation of building models that solve practical problems has been an orientation that served them and our profession very well.

We hope the readers of this two-part volume enjoy its contents. We feel fortunate to have had the opportunity of working with these fine authors and putting this volume together.

Thomas B. Fomby
Department of Economics
Southern Methodist University
Dallas, Texas 75275

Dek Terrell
Department of Economics
Louisiana State University
Baton Rouge, Louisiana 70803

Robert F. Engle III
2003 Nobel Prize Winner
in Economics

Sir Clive W. J. Granger, Knight's designation (KB)
2003 Nobel Prize Winner
in Economics

LIST OF CONTRIBUTORS

Torben G. Andersen	Department of Finance, Kellogg School of Management, Northwestern University, IL, USA
Yong Bao	Department of Economics, University of Texas, TX, USA
Tim Bollerslev	Department of Economics, Duke University, NC, USA
Zongwu Cai	Department of Mathematics & Statistics, University of North Carolina, NC, USA
Ngai Hang Chan	Department of Statistics, The Chinese University of Hong Kong, New Territories, Hong Kong
Rong Chen	Department of Information and Decision Sciences, College of Business Administration, The University of Illinois at Chicago, IL, USA
Francis X. Diebold	Department of Economics, University of Pennsylvania, PA, USA
Robert F. Engle III	Stern School of Business, New York University, NY, USA
Valeriy V. Gavrishchaka	Alexandra Investment Management, LLC, NY, USA
Sir Clive W. J. Granger, KB	Department of Economics, University of California at San Diego, CA, USA
Eric Hillebrand	Department of Economics, Louisiana State University, LA, USA
Gabriel Huerta	Department of Mathematics and Statistics, University of New Mexico, NM, USA

Dennis W. Jansen	Department of Economics, Texas A&M University, TX, USA
Tze Leung Lai	Department of Statistics, Stanford University, CA, USA
Tae-Hwy Lee	Department of Economics, University of California, CA, USA
Hedibert Freitas Lopes	Graduate School of Business, University of Chicago, IL, USA
Wilfredo Palma	Departamento de Estadistica, Pontificia Universidad Catolica De Chile (PUC), Santiago, Chile
Esther Salazar	Instituto de Matemática, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil
Pierre L. Siklos	Department of Economics, Wilfrid Laurier University, Ontario, Canada
Alejandro Villagran	Department of Mathematics and Statistics, University of New Mexico, NM, USA
Zijun Wang	Private Enterprise Research Center, Texas A&M University, TX, USA
Mark E. Wohar	Department of Economics, University of Nebraska–Omaha, NE, USA
Ginger Wu	Department of Banking and Finance, University of Georgia, GA, USA
Haipeng Xing	Department of Statistics, Columbia University, NY, USA
Gawon Yoon	School of Economics, Kookmin University, Seoul, South Korea
Zhengjun Zhang	Department of Statistics, University of Wisconsin-Madison, WI, USA

INTRODUCTION

Thomas B. Fomby and Dek Terrell

The editors are pleased to offer the following papers to the reader in recognition and appreciation of the contributions to our literature made by Robert Engle and Sir Clive Granger, winners of the 2003 Nobel Prize in Economics. Please see the previous dedication page of this volume. The basic themes of this part of Volume 20 of *Advances in Econometrics* are time-varying betas of the capital asset pricing model, analysis of predictive densities of nonlinear models of stock returns, modeling multivariate dynamic correlations, flexible seasonal time series models, estimation of long-memory time series models, the application of the technique of boosting in volatility forecasting, the use of different time scales in Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) modeling, out-of-sample evaluation of the 'Fed Model' in stock price valuation, structural change as an alternative to long memory, the use of smooth transition autoregressions in stochastic volatility modeling, the analysis of the "balancedness" of regressions analyzing Taylor-type rules of the Fed Funds rate, a mixture-of-experts approach for the estimation of stochastic volatility, a modern assessment of Clive's first published paper on sunspot activity, and a new class of models of tail-dependence in time series subject to jumps. Of course, we are also pleased to include Rob's and Clive's remarks on their careers and their views on innovation in econometric theory and practice that were given at the Third Annual *Advances in Econometrics* Conference held at Louisiana State University, Baton Rouge, on November 5–7, 2004.

Let us briefly review the specifics of the papers presented here. In the first paper, "Realized Beta: Persistence and Predictability," Torben Andersen, Tim Bollerslev, Francis Diebold, and Jin Wu review the literature on the one-factor Capital Asset Pricing Model (CAPM) for the purpose of coming to a better understanding of the variability of the betas of such models. They do this by flexibly modeling betas as the ratio of the integrated stock and market return covariance and integrated market variance in a way that allows, but does not impose, fractional integration and/or cointegration.

They find that, although the realized variances and covariances fluctuate widely and are highly persistent and predictable, the realized betas, which are simple nonlinear functions of the realized variances and covariances, display much less persistence and predictability. They conclude that the constant beta CAPM, as bad as it may be, is nevertheless not as bad as some popular conditional CAPMs. Their paper provides some very useful insight into why allowing for time-varying betas may do more harm than good when estimated from daily data. They close by sketching an interesting framework for future research using high-frequency intraday data to improve the modeling of time-varying betas.

Yong Bao and Tae-Hwy Lee in their paper "Asymmetric Predictive Abilities of Nonlinear Models for Stock Returns: Evidence from Density Forecast Comparison" investigate the nonlinear predictability of stock returns when the density forecasts are evaluated and compared instead of the conditional mean point forecasts. They use the Kullback–Leibler Information Criterion (KLIC) divergence measure to characterize the extent of misspecification of a forecast model. Their empirical findings suggest that the out-of-sample predictive abilities of nonlinear models for stock returns are asymmetric in the sense that the right tails of the return series are predictable via many of the nonlinear models while they find no such evidence for the left tails or the entire distribution.

Zongwu Cai and Rong Chen introduce a new class of flexible seasonal time series models to characterize trend and seasonal variation in their paper "Flexible Seasonal Time Series Models." Their model consists of a common trend function over periods and additive individual trend seasonal functions that are specific to each season within periods. A local linear approach is developed to estimate the common trend and seasonal trend functions. The consistency and asymptotic normality of the proposed estimators are established under weak α-mixing conditions and without specifying the error distribution. The proposed methodologies are illustrated with a simulated example and two economic and financial time series, which exhibit nonlinear and nonstationary behavior.

In "Estimation of Long-Memory Time Series Models: A Survey of Different Likelihood-Based Methods" Ngai Hang Chan and Wilfredo Palma survey the various likelihood-based techniques for analyzing long memory in time series data. The authors classify these methods into the following categories: Exact maximum likelihood methods, maximum likelihood methods based on autoregressive approximations, Whittle estimates, Whittle estimates with autoregressive truncation, approximate estimates based on the Durbin–Levinson algorithm, state–space based estimates for the autoregressive fractionally integrated moving average (ARFIMA) models, and estimation of stochastic volatility models. Their review provides a succinct survey of these methodologies as well as an overview of the important related problems such as the maximum likelihood estimation with missing data, influence of subsets of observations on estimates, and the estimation of seasonal long-memory models. Performances and asymptotic properties of these techniques are compared and examined. Interconnections and finite sample performances among these procedures are studied and applications to financial time series of these methodologies are discussed.

In "Boosting-Based Frameworks in Financial Modeling: Application to Symbolic Volatility Forecasting" Valeriy Gavrishchaka suggests that Boosting (a novel ensemble learning technique) can serve as a simple and robust framework for combining the best features of both analytical and data-driven models and more specifically discusses how Boosting can be applied for typical financial and econometric applications. Furthermore, he demonstrates some of the capabilities of Boosting by showing how a Boosted collection of GARCH-type models for the IBM stock time series can be used to produce more accurate forecasts of volatility than both the best single model of the collection and the widely used GARCH(1,1) model.

Eric Hillebrand studies different generalizations of GARCH that allow for several time scales in his paper "Overlaying Time Scales in Financial Volatility Data." In particular he examines the nature of the volatility in four measures of the U.S. stock market (S&P 500, Dow Jones Industrial Average, CRSP equally weighted index, and CRSP value-weighted index) as well as the exchange rate of the Japanese Yen against the U.S. Dollar and the U.S. federal funds rate. In addition to analyzing these series using the conventional ARCH and GARCH models he uses three models of multiple time scales, namely, fractional integration, two-scale GARCH, and wavelet analysis in conjunction with Heterogeneous ARCH (HARCH). Hillebrand finds that the conventional ARCH and GARCH models miss the important short correlation time scale in the six series. Based on a holding sample test, the multiple time scale models, although offering an improvement over the conventional ARCH and GARCH models, still did not completely model the short correlation structure of the six series. However, research in extending volatility models in this way appears to be promising.

The "Fed Model" postulates a cointegrating relationship between the equity yield on the S&P 500 and the bond yield. In their paper, "Evaluating the 'Fed Model' of Stock Price Valuation: An Out-of-Sample Forecasting Perspective," Dennis Jansen and Zijun Wang evaluate the Fed Model as a vector-error correction forecasting model for stock prices and for bond yields. They compare out-of-sample forecasts of each of these two variables from a univariate model and various versions of the Fed Model including both linear and nonlinear vector error correction models. They find that for stock prices the Fed Model improves on the univariate model for longer-horizon forecasts, and the nonlinear vector-error correction model performs even better than its linear version.

In their paper, "Structural Change as an Alternative to Long Memory in Financial Time Series," Tze Leung Lai and Haipeng Xing note that volatility persistence in GARCH models and spurious long memory in autoregressive models may arise if the possibility of structural changes is not incorporated in the time series model. Therefore, they propose a structural change model that allows changes in the volatility and regression parameters at unknown times and with unknown changes in magnitudes. Their model is a hidden Markov model in which the volatility and regression parameters can continuously change and are estimated by recursive filters. As their hidden Markov model involves gamma-normal conjugate priors, there are explicit recursive formulas for the optimal filters and smoothers. Using NASDAQ weekly return data, they show how the optimal structural change model can be applied to segment financial time series by making use of the estimated probabilities of structural breaks.

In their paper, "Time Series Mean Level and Stochastic Volatility Modeling by Smooth Transition Autoregressions: A Bayesian Approach," Hedibert Lopes and Esther Salazar propose a Bayesian approach to model the level and variance of financial time series based on a special class of nonlinear time series models known as the logistic smooth transition autoregressive (LSTAR) model. They propose a Markov Chain Monte Carlo (MCMC) algorithm for the levels of the time series and then adapt it to model the stochastic volatilities. The LSTAR order of their model is selected by the three information criteria Akaike information criterion (AIC), Bayesian information criterion (BIC), and Deviance information criteria (DIC). They apply their algorithm to one synthetic and two real-time series, namely the Canadian Lynx data and the SP500 return series, and find the results encouraging when modeling both the levels and the variance of univariate time series with LSTAR structures.

Relying on Robert Engle's and Clive Granger's many and varied contributions to econometrics analysis, Pierre Siklos and Mark Wohar examine some key econometric considerations involved in estimating Taylor-type rules for U.S. data in their paper "Estimating Taylor-Type Rules: An Unbalanced Regression?" They focus on the roles of unit roots, cointegration, structural breaks, and nonlinearities to make the case that most existing estimates are based on unbalanced regressions. A variety of their estimates reveal that neglecting the presence of cointegration results in the omission of a necessary error correction term and that Fed reactions during the Greenspan era appear to have been asymmetric. They further argue that error correction and nonlinearities may be one way to estimate Taylor rules over long samples when the underlying policy regime may have changed significantly.

Alejandro Villagran and Gabriel Huerta propose a Bayesian Mixture-of-Experts (ME) approach to estimating stochastic volatility in time series in their paper "Bayesian Inference on Mixture-of-Experts for Estimation of Stochastic Volatility." They use as their "experts" the ARCH, GARCH, and EGARCH models to analyze the stochastic volatility in the U.S. dollar/German Mark exchange rate and conduct a study of the volatility of the Mexican stock market (IPC) index using the Dow Jones Industrial (DJI) index as a covariate. They also describe the estimation of predictive volatilities and their corresponding measure of uncertainty given by a Bayesian credible interval using the ME approach. In the applications they present, it is interesting to see how the posterior probabilities of the "experts" change over time and to conjecture why the posterior probabilities changed as they did.

Sir Clive Granger published his first paper "A Statistical Model for Sunspot Activity" in 1957 in the prestigious *Astrophysical Journal*. As a means of recognizing Clive's many contributions to the econometric analysis of time series and celebrating the near 50th anniversary of his first publication, one of his students and now professor, Gawon Yoon, has written a paper that provides a modern time series assessment of Clive's first paper. In "A Modern Time Series Assessment of 'A Statistical Model for Sunspot Activity' by C. W. J. Granger (1957)," Yoon reviews Granger's statistical model of sunspots containing two parameters representing an amplitude factor and the occurrence of minima, respectively. At the time Granger's model accounted for about 85% of the total variation in the sunspot data. Interestingly, Yoon finds that, in the majority, Granger's model quite nicely explains the more recent occurrence of sunspots despite the passage of time. Even though it appears that some of the earlier observations that Granger had available were measured differently from later sunspot numbers, Granger's simple two-parameter model still accounts for more than 80% of the total variation in the extended sunspot data. This all goes to show (as Sir Clive would attest) that simple models can also be useful models.

With respect to Yoon's review of Granger's paper, Sir Clive was kind enough to offer remarks that the editors have chosen to publish immediately following Yoon's paper. In reading Clive's delightful remarks we come to know (or some of us remember, depending on your age) how difficult it was to conduct empirical analysis in the 1950s and 1960s. As Clive notes, "Trying to plot by hand nearly two hundred years of monthly data is a lengthy task!" So econometric researchers post-1980 have many things to be thankful for, not the least of which is fast and inexpensive computing. Nevertheless, Clive and many other young statistical researchers at the time were undaunted. They were convinced that quantitative research was important and a worthwhile endeavor regardless of the expense of time and they set out to investigate what the naked eye could not detect. You will enjoy reading Clive's remarks knowing that he is always appreciative of the comments and suggestions of colleagues and that he is an avid supporter of best practices in statistics and econometrics.

In the final paper of Part B of Volume 20, "A New Class of Tail-Dependent Time Series Models and Its Applications in Financial Time Series," Zhengjun Zhang proposes a new class of models to determine the order of lag-k tail dependence in financial time series that exhibit jumps. His base model is a specific class of maxima of moving maxima processes (M3 processes). Zhang then improves on his base model by allowing for possible asymmetry between positive and negative returns. His approach adopts a hierarchical model structure. First you apply, say GARCH(1,1), to get estimated standard deviations, then based on standardized returns, you apply M3 and Markov process modeling to characterize the tail dependence in the time series. Zhang demonstrates his model and his approach using the S&P 500 Index. As he points out, estimates of the parameters of the proposed model can be used to compute the value at risk (VaR) of the investments whose returns are subject to jump processes.

GOOD IDEAS

Robert F. Engle III

The Nobel Prize is given for good ideas – very good ideas. These ideas often shape the direction of research for an academic discipline. These ideas are often accompanied by a great deal of work by many researchers.

Most good ideas don't get prizes but they are the centerpieces of our research and our conferences. At this interesting *Advances in Econometrics* conference hosted by LSU, we've seen lots of new ideas, and in our careers we have all had many good ideas. I would like to explore where they come from and what they look like.

When I was growing up in suburban Philadelphia, my mother would sometimes take me over to Swarthmore College to the Physics library. It was a small dusty room with windows out over a big lawn with trees. The books cracked when I opened them; they smelled old and had faded gold letters on the spine. This little room was exhilarating. I opened books by the famous names in physics and read about quantum mechanics, elementary particles and the history of the universe. I didn't understand too much but kept piecing together my limited ideas. I kept wondering whether I would understand these things when I was older and had studied in college or graduate school. I developed a love of science and the scientific method. I think this is why I studied econometrics; it is the place where theory meets reality. It is the place where data on the economy tests the validity of economic theory.

Fundamentally I think good ideas are simple. In Economics, most ideas can be simplified until they can be explained to non-specialists in plain language. The process of simplifying ideas and explaining them is extremely important. Often the power of the idea comes from simplification of a collection of complicated and conflicting research. The process of distilling out the simple novel ingredient is not easy at all and often takes lots of fresh starts and numerical examples. Discouragingly, good ideas boiled down to their essence may seem trivial. I think this is true of ARCH and Cointegration and many other Nobel citations. But, I think we should not be offended by this simplicity, but rather we should embrace it. Of course it is easy to do this after 20 years have gone by; but the trick is to recognize good ideas early. Look for them at seminars or when reading or refereeing or editing.

Good ideas generalize. A good idea, when applied to a new situation, often gives interesting insights. In fact, the implications of a good idea may be initially surprising. Upon reflection, the implications may be of growing importance. If ideas translated into other fields give novel interpretations to existing problems, this is a measure of their power.

Often good ideas come from examining one problem from the point of view of another. In fact, the ARCH model came from such an analysis. It was a marriage of theory, time series and empirical evidence. The role of uncertainty in rational expectations macroeconomics was not well developed, yet there were theoretical reasons why changing uncertainty could have real effects. From a time series point of view a natural solution to modeling uncertainty was to build conditional models of variance rather than the more familiar unconditional models. I knew that Clive's test for bilinearity based on the autocorrelation of squared residuals was often significant in macroeconomic data, although I suspected that the test was also sensitive to other effects such as changing variances. The idea for the ARCH model came from combining these three observations to get an autoregressive model of conditional heteroskedasticity.

Sometimes a good idea can come from attempts to disprove proposals of others. Clive traces the origin of cointegration to his attempt to disprove a David Hendry conjecture that a linear combination of the two integrated series could be stationary. From trying to show that this was impossible, Clive proved the Granger Representation theorem that provides the fundamental rationale for error correction models in cointegrated systems.

My first meeting in Economics was the 1970 World Congress of the Econometric Society in Cambridge England. I heard many of the famous economists of that generation explain their ideas. I certainly did not understand everything but I wanted to learn it all. I gave a paper at this meeting at a session organized by Clive that included Chris Sims and Phoebus Dhrymes. What a thrill. I have enjoyed European meetings of the Econometric Society ever since.

My first job was at MIT. I had a lot of chances to see good ideas; particularly good ideas in finance. Myron Scholes and Fischer Black were working on options theory and Bob Merton was developing continuous time finance. I joined Franco Modigliani and Myron on Michael Brennan's dissertation committee where he was testing the CAPM. Somehow I missed the opportunity to capitalize on these powerful ideas and it was only many years later that I moved my research in this direction.

I moved to UCSD in 1976 to join Clive Granger. We studied many fascinating time series problems. Mark Watson was my first PhD student at UCSD. The ARCH model was developed on sabbatical at LSE, and when I returned, a group of graduate students contributed greatly to the development of this research. Tim Bollerslev and Dennis Kraft were among the first, Russ Robins and Jeff Wooldridge and my colleague David Lilien were instrumental in helping me think about the finance applications. The next 20 years at UCSD were fantastic in retrospect. I don't think we knew at the time how we were moving the frontiers in econometrics. We had great visitors and faculty and students and every day there were new ideas.

These ideas came from casual conversations and a relaxed mind. They came from brainstorming on the blackboard with a student who was looking for a dissertation topic. They came from "Econometrics Lunch" when we weren't talking about gossip in the profession. Universities are incubators of good ideas. Our students come with good ideas, but they have to be shaped and interpreted. Our faculties have good ideas, which they publish and lecture on around the world. Our departments and universities thrive on good ideas that make them famous places for study and innovation. They also contribute to spin-offs in the private sector and consulting projects. Good ideas make the whole system work and it is so important to recognize them in all their various forms and reward them.

As a profession we are very protective of our ideas. Often the origin of the idea is disputable. New ideas may have only a part of the story that eventually develops; who gets the credit? While such disputes are natural, it is often better in my opinion to recognize previous contributions and stand on their shoulders thereby making your own ideas even more important. I give similar advice for academics who are changing specialties; stand with one foot in the old discipline and one in the new. Look for research that takes your successful ideas from one field into an important place in a new discipline.

Here are three quotations that I think succinctly reflect these thoughts.

- *"The universe is full of magical things, patiently waiting for our wits to grow sharper"*
  Eden Philpotts
- *"To see what is in front of one's nose requires a constant struggle."*
  George Orwell
- *"To select well among old things is almost equal to inventing new ones"*
  Nicolas Charles Trublet

There is nothing in our chosen career that is as exhilarating as having a good idea. But a very close second is seeing someone develop a wonderful new application from your idea. The award of the Nobel Prize to Clive and me for our work in time series is an honor to all of the authors who contributed to the conference and to this volume. I think the prize is really given to a field and we all received it. This gives me so much joy. And I hope that someone in this volume will move forward to open more doors with powerful new ideas, and receive her own Nobel Prize.

Robert F. Engle III
Remarks Given at Third Annual Advances in Econometrics Conference
Louisiana State University
Baton Rouge, Louisiana
November 5–7, 2004

THE CREATIVITY PROCESS

Sir Clive W. J. Granger, KB

In 1956, I was searching for a Ph.D. topic and I selected time series analysis as being an area that was not very developed and was potentially interesting. I have never regretted that choice. Occasionally, I have tried to develop other interests but after a couple of years away I would always return to time series topics where I am more comfortable.

I have never had a long-term research topic. What I try to do is to develop new ideas, topics, and models, do some initial development, and leave the really hard, rigorous stuff to other people. Some new topics catch on quickly and develop a lot of citations (such as cointegration), others are initially ignored but eventually become much discussed and applied (causality, as I call it), some develop interest slowly but eventually deeply (fractionally integrated processes), some have long term, steady life (combination of forecasts), whereas others generate interest but eventually vanish (bilinear models, spectral analysis).

The ideas come from many sources, by reading literature in other fields, from discussions with other workers, from attending conferences (time distance measure for forecasts), and from general reading. I will often attempt to take a known model and generalize and expand it in various ways. Quite frequently these generalizations turn out not to be interesting; I have several examples of general I(d) processes where d is not real or not finite. The models that do survive may be technically interesting but they may not prove useful with economic data, providing an example of a so-called "empty box," bilinear models, and I(d), d non-integer could be examples.

In developing these models one is playing a game. One can never claim that a new model will be relevant, only that it might be. Of course, when using the model to generate forecasts, one has to assume that the model is correct, but one must not forget this assumption. If the model is correct, the data will have certain properties that can be proved, but it should always be remembered that other models may generate the same properties, for example I(d), d a fraction, and break processes can give similar ‘‘long memory’’ autocorrelations. Finding properties of data and then suggesting that a particular model will have generated the data is a dangerous game.

Of course, once the research has been done one faces the problem of publication. The refereeing process is always a hassle. I am not convinced that delaying an interesting paper (I am not thinking of any of my own here) by a year or more to fix a few minor difficulties is actually helping the development of our field. Rob and I had initial rejections of some of our best joint papers, including the one on cointegration. My paper on the typical spectral shape took over three and a half years between submission and publication, and it is a very short paper.

My favorite editors’ comment was that ‘‘my paper was not very good (correct) but is was very short,’’ and as they just had that space to fill they would accept. My least favorite comment was a rejection of a paper with Paul Newbold because ‘‘it has all been done before.’’ As we were surprised at this we politely asked for citations. The referee had no citations, he just thought that must have been done before. The paper was published elsewhere.

For most of its history time series theory considered conditional means, but later conditional variances. The next natural development would be conditional quantiles, but this area is receiving less attention than I expected. The last stages are initially conditional marginal distributions, and finally conditional multivariate distributions. Some interesting theory is starting in these areas but there is an enormous amount to be done.

The practical aspects of time series analysis are rapidly changing with improvements in computer performance. Now many, fairly long series can be analyzed jointly. For example, Stock and Watson (1999) consider over 200 macro series. However, the dependent series are usually considered individually, whereas what we are really dealing with is a sample from a 200-dimensional multivariate distribution, assuming the processes are jointly stationary. How to even describe the essential features of such a distribution, which is almost certainly non-Gaussian, in a way that is useful to economists and decision makers is a substantial problem in itself.

My younger colleagues sometimes complain that we old guys solved all the interesting easy questions. I do not think that was ever true and is not true now. The higher we stand the wider our perspective; I hope that Rob and I have provided, with many others, a suitable starting point for the future study in this area.

REFERENCE

Stock, J. H., & Watson, M. W. (1999). A Comparison of linear and nonlinear univariate models for forecasting macroeconomic time series. In: R. F. Engle & H. White (Eds), *Cointegration, causality, and forecasting: A festschrift in honour of Clive W. J. Granger*. Oxford: Oxford University Press.

Sir Clive W. J. Granger, KB
Remarks Read at Third Annual Advances in Econometrics Conference
Louisiana State University
Baton Rouge, Louisiana
November 5–7, 2004

This page intentionally left blank

REALIZED BETA: PERSISTENCE AND PREDICTABILITY[☆]

Torben G. Andersen, Tim Bollerslev,
Francis X. Diebold and Ginger Wu

ABSTRACT

A large literature over several decades reveals both extensive concern with the question of time-varying betas and an emerging consensus that betas are in fact time-varying, leading to the prominence of the conditional CAPM. Set against that background, we assess the dynamics in realized betas, vis-à-vis the dynamics in the underlying realized market variance and individual equity covariances with the market. Working in the recently popularized framework of realized volatility, we are led to a framework of nonlinear fractional cointegration: although realized variances and covariances are very highly persistent and well approximated as fractionally integrated, realized betas, which are simple nonlinear functions of those realized variances and covariances, are less persistent and arguably best modeled as stationary I(0) processes. We conclude by drawing implications for asset pricing and portfolio management.

[☆]We dedicate this paper to Clive W. J. Granger, a giant of modern econometrics, on whose broad shoulders we are fortunate to stand. This work was supported by the National Science Foundation and the Guggenheim Foundation.

Econometric Analysis of Financial and Economic Time Series/Part B
Advances in Econometrics, Volume 20, 1–39

ISSN: 0731-9053/doi:10.1016/S0731-9053(05)20020-8

1. INTRODUCTION

One of the key insights of asset pricing theory is also one of the simplest: only systematic risk should be priced. Perhaps not surprisingly, however, there is disagreement as to the sources of systematic risk. In the one-factor capital asset pricing model (CAPM), for example, systematic risk is determined by covariance with the market (Sharpe, 1963; Lintner, 1965a, b), whereas, in more elaborate pricing models, additional empirical characteristics such as firm size and book-to-market are seen as proxies for another set of systematic risk factors (Fama & French, 1993).[1]

As with most important scientific models, the CAPM has been subject to substantial criticism (e.g., Fama & French, 1992). Nevertheless, to paraphrase Mark Twain, the reports of its death are greatly exaggerated. In fact, the one-factor CAPM remains alive and well at the frontier of both academic research and industry applications, for at least two reasons. First, recent work reveals that it often works well – despite its wrinkles and warts – whether in traditional incarnations (e.g., Ang & Chen, 2003) or more novel variants (e.g., Cohen, Polk, & Vuolteenaho, 2002; Campbell & Vuolteenaho, 2004). Second, competing multi-factor pricing models, although providing improved statistical fit, involve factors whose economic interpretations in terms of systematic risks remain unclear, and moreover, the stability of empirically motivated multi-factor asset pricing relationships often appears tenuous when explored with true out-of-sample data, suggesting an element of data mining.[2]

In this paper, then, we study the one-factor CAPM, which remains central to financial economics nearly a half century after its introduction. A key question within this setting is whether stocks' systematic risks, as assessed by their correlations with the market, are constant over time – i.e., whether stocks' market betas are constant. And if betas are not constant, a central issue becomes how to understand and formally characterize their persistence and predictability vis-à-vis their underlying components.

The evolution of a large literature over several decades reveals both extensive concern with this question and, we contend, an eventual implicit consensus that betas *are* likely time-varying.[3] Several pieces of evidence support our contention. First, leading texts echo it. For example, Huang and Litzenberger (1988) assert that "It is unlikely that risk premiums and betas on individual assets are stationary over time" (p. 303). Second, explicitly dynamic betas are often modeled nonstructurally via time-varying parameter regression, in a literature tracing at least to the early "return to normality" model of Rosenberg (1973), as implemented in the CAPM by Schaefer Broaley, Hodges, and Thomas (1975). Third, even in the absence of explicit allowance for time-varying betas, the CAPM is typically estimated using moving estimation windows, usually of 5–10 years, presumably to guard against beta variation (e.g., Fama, 1976; Campbell, Lo, & MacKinlay, 1997). Fourth, theoretical and empirical inquiries in asset pricing are often undertaken in conditional, as opposed to unconditional, frameworks, the essence of which is to allow for time-varying betas, presumably because doing so is viewed as necessary for realism.

The motivation for the conditional CAPM comes from at least two sources. First, from a theoretical perspective, financial economic considerations suggest that betas may vary with conditioning variables, an idea developed theoretically and empirically in a large literature that includes, among many others, Dybvig and Ross (1985), Hansen and Richard (1987), Ferson, Kandel, and Stambaugh (1987), Ferson and Harvey (1991), Jagannathan and Wang (1996), and Wang (2003).[4] Second, from a different and empirical perspective, the financial econometric volatility literature (see Andersen, Bollerslev, & Diebold, 2005, for a recent survey) has provided extensive evidence of wide fluctuations and high persistence in asset market conditional variances, and in individual equity conditional covariances with the market. Thus, even from a purely statistical viewpoint, market betas, which are ratios of time-varying conditional covariances and variances, might be expected to display persistent fluctuations, as in Bollerslev, Engle, and Wooldridge (1988). In fact, unless some special cancellation occurs – in a way that we formalize – betas would inherit the persistence features that are so vividly present in their constituent components.

Set against this background, we assess the dynamics in betas vis-à-vis the widely documented persistent dynamics in the underlying variance and covariances. We proceed as follows: In Section 2, we sketch the framework, both economic and econometric, in which our analysis is couched. In Section 3, we present the empirical results with an emphasis on analysis of persistence and predictability. In Section 4, we formally assess the uncertainty in our beta estimates. In Section 5, we offer summary, conclusions, and directions for future research.

2. THEORETICAL FRAMEWORK

Our approach has two key components. First, in keeping with the recent move toward nonparametric volatility measurement, we cast our analysis within the framework of realized variances and covariances, or equivalently, empirical quadratic variation and covariation. That is, we do not entertain a null hypothesis of period-by-period constant betas, but instead explicitly allow for continuous evolution in betas. Our "realized betas" are (continuous-record) consistent for realizations of the underlying ratio between the integrated stock and market return covariance and the integrated market variance.[5] Second, we work in a flexible econometric framework that allows for – without imposing – fractional integration and/or cointegration between the market variance and individual equity covariances with the market.

2.1. Realized Quarterly Variances, Covariances, and Betas

We provide estimates of quarterly betas, based on nonparametric realized quarterly market variances and individual equity covariances with the market. The quarterly frequency is appealing from a substantive financial economic perspective, and it also provides a reasonable balance between efficiency and robustness to microstructure noise. Specifically, we produce our quarterly estimates using underlying daily returns, as in Schwert (1989), so that the sampling frequency is quite high relative to the quarterly horizon of interest, yet low enough so that contamination by microstructure noise is not a serious concern for the highly liquid stocks that we study. The daily frequency further allows us to utilize a long sample of data, which is not available when sampling more frequently.

Suppose that the logarithmic $N \times 1$ vector price process, p_t, follows a multivariate continuous-time stochastic volatility diffusion,

$$\mathrm{d}p_t = \mu_t \mathrm{d}t + \Omega_t \mathrm{d}W_t \tag{1}$$

where W_t denotes a standard N-dimensional Brownian motion, and both the process for the $N \times N$ positive definite diffusion matrix, Ω_t, and the N-dimensional instantaneous drift, μ_t, are strictly stationary and jointly independent of the W_t process. For our purposes it is helpful to think of the Nth element of p_t as containing the log price of the market and the ith element of p_t as containing the log price of the ith individual stock included in the analysis, so that the corresponding covariance matrix contains both the market variance, say $\sigma^2_{M,t} = \Omega_{(NN),t}$, and the individual equity covariance with the market, $\sigma_{iM,t} = \Omega_{(iN),t}$. Then, conditional on the sample path realization of μ_t and Ω_t, the distribution of the continuously compounded h-period return, $r_{t+h,h} \equiv p_{t+h} - p_t$, is

$$r_{t+h,h} | \sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h} \sim N\left(\int_0^h \mu_{t+\tau}\mathrm{d}\tau,\ \int_0^h \Omega_{t+\tau}\mathrm{d}\tau\right) \tag{2}$$

where $\sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h}$ denotes the σ-field generated by the sample paths of $\mu_{t+\tau}$ and $\Omega_{t+\tau}$, for $0 \leq \tau \leq h$. The integrated diffusion matrix $\int_0^h \Omega_{t+\tau} \mathrm{d}\tau$, therefore provides a natural measure of the true latent h-period volatility.[6] The requirement that the innovation process, W_t, is independent of the drift and diffusion processes is rather strict and precludes, for example, the asymmetric relations between return innovations and volatility captured by the so-called leverage or volatility feedback effects. However, from the results in Meddahi (2002), Barndorff-Nielsen and Shephard (2003), and Andersen, Bollerslev, and Meddahi (2004), we know that the continuous-record asymptotic distribution theory for the realized covariation continues to provide an excellent approximation for empirical high-frequency realized volatility measures.[7] As such, even if the conditional return distribution result (2) does not apply in full generality, the evidence presented below, based exclusively on the realized volatility measures, remains trustworthy in the presence of asymmetries in the return innovation–volatility relations.

By the theory of quadratic variation, we have that under weak regularity conditions, and regardless of the presence of leverage or volatility feedback effects, that

$$\sum_{j=1,\ldots,[h/\Delta]} r_{t+j\cdot\Delta,\Delta} \cdot r'_{t+j\cdot\Delta,\Delta} - \int_0^h \Omega_{t+\tau} \mathrm{d}\tau \rightarrow 0 \tag{3}$$

almost surely for all t as the sampling frequency of the returns increases, or $\Delta \rightarrow 0$. Thus, by summing sufficiently finely sampled high-frequency returns, it is possible to construct ex-post *realized* volatility measures for the integrated latent volatilities that are asymptotically free of measurement error. This contrasts sharply with the common use of the cross-product of the h-period returns, $r_{t+h,h} \cdot r'_{t+h,h}$, as a simple ex post (co)variability measure. Although the squared return (innovation) over the forecast horizon provides an unbiased estimate for the integrated volatility, it is an extremely noisy estimator, and predictable variation in the true latent volatility process is typically dwarfed by measurement error. Moreover, for longer horizons any conditional mean dependence will tend to contaminate this variance measure. In contrast, as the sampling frequency is lowered, the impact of the drift term vanishes, thus effectively annihilating the mean.

These assertions remain valid if the underlying continuous time process in Eq. (1) contains jumps, so long as the price process is a special semimartingale, which will hold if it is arbitrage-free. Of course, in this case the limit of the summation of the high-frequency returns will involve an additional jump component, but the interpretation of the sum as the realized h-period return volatility remains intact.

Finally, with the realized market variance and realized covariance between the market and the individual stocks in hand, we can readily define and empirically construct the individual equity "realized betas." Toward that end, we introduce some formal notation. Using an initial subscript to indicate the corresponding element of a vector, we denote the realized market volatility by

$$\hat{v}^2_{M,t,t+h} = \sum_{j=1,\ldots,[h/\Delta]} \boldsymbol{r}^2_{(N),t+j\cdot\Delta,\Delta} \tag{4}$$

and we denote the realized covariance between the market and the ith individual stock return by

$$\hat{v}_{iM,t,t+h} = \sum_{j=1,\ldots,[h/\Delta]} \boldsymbol{r}_{(i),t+j\cdot\Delta,\Delta} \cdot \boldsymbol{r}_{(N),t+j\cdot\Delta,\Delta} \tag{5}$$

We then define the associated realized beta as

$$\hat{\beta}_{i,t,t+h} = \frac{\hat{v}_{iM,t,t+h}}{\hat{v}^2_{M,t,t+h}} \tag{6}$$

Under the assumptions invoked for Eq. (1), this realized beta measure is consistent for the true underlying integrated beta in the following sense:

$$\hat{\beta}_{i,t,t+h} \rightarrow \beta_{i,t,t+h} = \frac{\int_0^h \Omega_{(iN),t+\tau} \mathrm{d}\tau}{\int_0^h \Omega_{(NN),t+\tau} \mathrm{d}\tau} \tag{7}$$

almost surely for all t as the sampling frequency increases, or $\Delta \rightarrow 0$.

A number of comments are in order. First, the integrated return covariance matrix, $\int_0^h \Omega_{t+\tau} \mathrm{d}\tau$, is treated as stochastic, so both the integrated market variance and the integrated covariances of individual equity returns with the market over $[t, t+h]$ are ex ante, as of time t, unobserved and governed by a non-degenerate (and potentially unknown) distribution. Moreover, the covariance matrix will generally vary continuously and randomly over the entire interval, so the integrated covariance matrix should be interpreted as the average realized covariation among the return series. Second, Eq. (3) makes it clear that the realized market volatility in (4) and the realized covariance in (5) are continuous-record consistent estimators of the (random) realizations of the underlying integrated market volatility and covariance. Thus, as a corollary, the realized beta will be consistent for the integrated beta, as stated in (7). Third, the general representation here encompasses the standard assumption of a constant beta over the measurement or estimation horizon, which is attained for the degenerate case of the Ω_t process being constant throughout each successive h-period measurement interval, or $\Omega_t = \Omega$. Fourth, the realized beta estimation procedure in Eq. (4)–(6) is implemented through a simple regression (without a constant term) of individual high-frequency stock returns on the corresponding market return. Nonetheless, the interpretation is very different from a standard regression, as the Ordinary Least Square (OLS) point estimate now represents a consistent estimator of the ex post realized regression coefficient obtained as the ratio of unbiased estimators of the average realized covariance and the realized market variance. The associated continuous-record asymptotic theory developed by Barndorff-Nielsen and Shephard (2003) explicitly recognizes the diffusion setting underlying this regression interpretation and hence facilitates the construction of standard errors for our beta estimators.

2.2. Nonlinear Fractional Cointegration: A Common Long-Memory Feature in Variances and Covariances

The possibility of common persistent components is widely recognized in modern multivariate time-series econometrics. It is also important for our analysis, because there may be common persistence features in the underlying variances and covariances from which betas are produced.

The idea of a common feature is a simple generalization of the well-known cointegration concept. If two variables are integrated but there exists a function f of them that is not, we say that they are cointegrated, and we call f the conintegrating function. More generally, if two variables have property X but there exists a function of them that does not, we say that they have common feature X. A key situation is when X corresponds to *persistence*, in which case we call the function of the two variables that eliminates the persistence the *copersistence function*. It will prove useful to consider linear and nonlinear copersistence functions in turn.

Most literature focuses on linear copersistence functions. The huge cointegration literature pioneered by Granger (1981) and Engle and Granger (1987) deals primarily with linear common long-memory $I(1)$ persistence features. The smaller copersistence literature started by Engle and Kozicki (1993) deals mostly with linear common short-memory $I(0)$ persistence features. The idea of fractional cointegration, suggested by Engle and Granger (1987) and developed by Cheung and Lai (1993) and Robinson and Marinucci, (2001), among others, deals with linear common long-memory $I(d)$ persistence features, $0<d<1/2$.

Our interest is closely related but different. First, it centers on *nonlinear* copersistence functions, because betas are ratios. There is little literature on nonlinear common persistence features, although they are implicitly treated in Granger (1995). We will be interested in nonlinear common long-memory $I(d)$ persistence features, $0<d<1/2$, effectively corresponding to nonlinear fractional cointegration.[8]

Second, we are interested primarily in the case of *known* cointegrating relationships. That is, we may not know whether a given stock's covariance with the market is fractionally cointegrated with the market variance, but if it is, then there is a good financial economic reason (i.e., the CAPM) to suspect that the cointegrating function is the *ratio* of the covariance to the variance. This provides great simplification. In the integer-cointegration framework with known cointegrating vector under the alternative, for example, one could simply test the cointegrating combination for a unit root, or test the significance of the error-correction term in a complete error-correction model, as in Horvath and Watson (1995). We proceed in analogous fashion, examining the integration status (generalized to allow for fractional integration) of the realized market variance, realized individual equity covariances with the market, and realized market betas.

Our realized beta series are unfortunately relatively short compared to the length required for formal testing and inference procedures regarding (fractional) cointegration, as the fractional integration and cointegration estimators proposed by Geweke and Porter-Hudak (1983), Robinson and Marinucci (2001), and Andrews and Guggenberger (2003) tend to behave quite erratically in small samples. In addition, there is considerable measurement noise in the individual beta series so that influential outliers may have a detrimental impact on our ability to discern the underlying dynamics. Hence, we study the nature of the long range dependence and short-run dynamics in the realized volatility measures and realized betas through intentionally less formal but arguably more informative graphical means, and via some robust procedures that utilize the joint information across many series, to which we now turn.

3. EMPIRICAL ANALYSIS

We examine primarily the realized quarterly betas constructed from daily returns. We focus on the dynamic properties of market betas vis-à-vis the dynamic properties of their underlying covariance and variance components. We quantify the dynamics in a number of ways, including explicit measurement of the degree of predictability in the tradition of Granger and Newbold (1986).

3.1. Dynamics of Quarterly Realized Variance, Covariances and Betas

This section investigates the realized quarterly betas constructed from daily returns obtained from the Center for Research in Security Prices from July 1962 to September 1999. We take the market return $r_{m,t}$ to be the 30 Dow Jones Industrial Average (DJIA), and we study the subset of 25 DJIA stocks as of March 1997 with complete data from July 2, 1962 to September 17, 1999, as detailed in Table 1. We then construct quarterly realized DJIA variances, individual equity covariances with the market, and betas, 1962:3–1999:3 (149 observations).

In Fig. 1, we provide a time-series plot of the quarterly realized market variance, with fall 1987 included (top panel) and excluded (bottom panel). It is clear that the realized variance is quite persistent and, moreover, that the fall 1987 volatility shock is unlike any other ever recorded, in that volatility reverts to its mean almost instantaneously. In addition, our subsequent computation of asymptotic standard errors reveals that the uncertainty associated with the fall 1987 beta estimate is enormous, to the point of rendering it entirely uninformative. In sum, it is an exceptional outlier with potentially large influence on the analysis, and it is measured with huge imprecision. Hence, following many other authors, we drop the fall 1987 observation from this point onward.

In Figs. 2 and 3, we display time-series plots of the 25 quarterly realized covariances and realized betas.[9] Like the realized variance, the realized covariances appear highly persistent. The realized betas, in contrast, appear noticeably less persistent. This impression is confirmed by the statistics presented in Table 2: the mean Ljung–Box Q-statistic (through displacement 12) is 84 for the realized covariance, but only 47 for the realized beta, although both are of course significant relative to a $\chi^2(12)$ distribution.[10]

The impression of reduced persistence in realized betas relative to realized covariances is also confirmed by the sample autocorrelation functions for the realized market variance, the realized covariances with the market, and the realized betas shown in Fig. 4.[11] Most remarkable is the close correspondence between the shape of the realized market variance correlogram and the realized covariance correlograms. This reflects an extraordinary high degree of dependence in the correlograms across the individual realized

Table 1. The Dow Jones Thirty.

Company Name	Ticker	Data Range
Alcoa Inc.	AA	07/02/1962 to 09/17/1999
Allied Capital Corporation	ALD	07/02/1962 to 09/17/1999
American Express Co.	AXP[a]	05/31/1977 to 09/17/1999
Boeing Co.	BA	07/02/1962 to 09/17/1999
Caterpillar Inc.	CAT	07/02/1962 to 09/17/1999
Chevron Corp.	CHV	07/02/1962 to 09/17/1999
DuPont Co.	DD	07/02/1962 to 09/17/1999
Walt Disney Co.	DIS	07/02/1962 to 09/17/1999
Eastman Kodak Co.	EK	07/02/1962 to 09/17/1999
General Electric Co.	GE	07/02/1962 to 09/17/1999
General Motors Corp.	GM	07/02/1962 to 09/17/1999
Goodyear Tire & Rubber Co.	GT	07/02/1962 to 09/17/1999
Hewlett–Packard Co.	HWP	07/02/1962 to 09/17/1999
International Business Machines Corp.	IBM	07/02/1962 to 09/17/1999
International Paper Co.	IP	07/02/1962 to 09/17/1999
Johnson & Johnson	JNJ	07/02/1962 to 09/17/1999
JP Morgan Chase & Co.	JPM[a]	03/05/1969 to 09/17/1999
Coca–Cola Co.	KO	07/02/1962 to 09/17/1999
McDonald's Corp.	MCD[a]	07/05/1966 to 09/17/1999
Minnesota Mining & Manufacturing Co.	MMM	07/02/1962 to 09/17/1999
Philip Morris Co.	MO	07/02/1962 to 09/17/1999
Merck & Co.	MRK	07/02/1962 to 09/17/1999
Procter & Gamble Co.	PG	07/02/1962 to 09/17/1999
Sears, Roebuck and Co.	S	07/02/1962 to 09/17/1999
AT&T Corp.	T	07/02/1962 to 09/17/1999
Travelers Group Inc.	TRV[a]	10/29/1986 to 09/17/1999
Union Carbide Corp.	UK	07/02/1962 to 09/17/1999
United Technologies Corp.	UTX	07/02/1962 to 09/17/1999
Wal–Mart Stores Inc.	WMT[a]	11/20/1972 to 09/17/1999
Exxon Corp.	XON	07/02/1962 to 09/17/1999

Note: A summary of company names and tickers, and the range of the data are examined. We use the Dow Jones Thirty as of March 1997.
[a]Stocks with incomplete data, which we exclude from the analysis.

covariances with the market, as shown in Fig. 5. In Fig. 4, it makes the median covariance correlogram appear as a very slightly dampened version of that for the market variance. This contrasts sharply with the lower and gently declining pattern for the realized beta autocorrelations. Intuitively, movements of the realized market variance are largely reflected in movements of the realized covariances; as such, they largely ''cancel'' when we form ratios (realized betas). Consequently, the correlation structure across

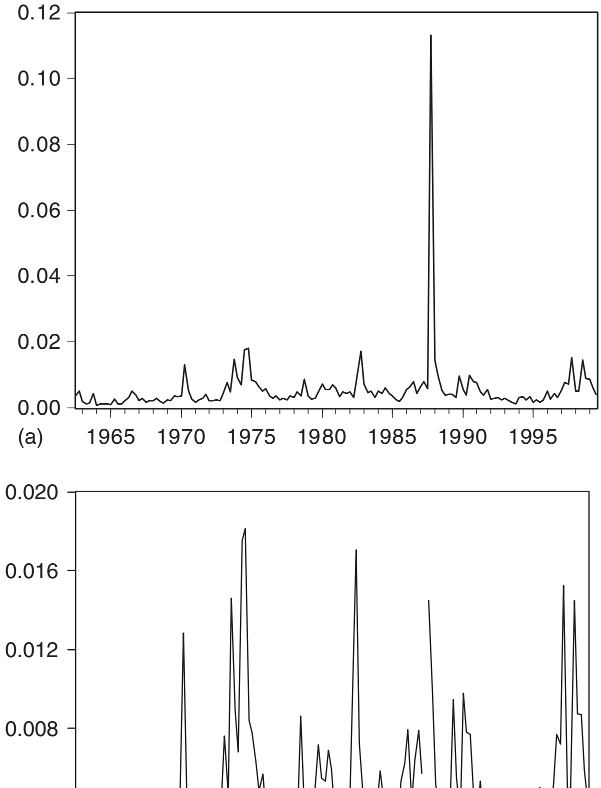

Fig. 1. Time Series Plot of Quarterly Realized Market Variance, Fall 1987 (a) Included (b) Exculded. *Note:* The Two Subfigures Show the Time Series of Quarterly Realized Market Variance, with The 1987:4 Outlier Included (a) and Excluded (b). The Sample Covers the Period from 1962:3 through 1999:3, for a Total of 149 Observations. We Calculate the Realized Quarterly Market Variances from Daily Returns.

the individual realized beta series in Fig. 6 is much more dispersed than is the case for the realized covariances in Fig. 5. This results in an effective averaging of the noise and the point estimates of the median correlation values are effectively zero beyond 10 quarters for the beta series.[12]