THE UNIVERSITY OF SYDNEY
Business School

Time Series and Forecasting Symposium TSF2022

1-2 December 2022

The University of Sydney, CBD Campus
Level 17, 133 Castlereagh Street, Sydney, Australia

Published by:
Time Series and Forecasting Research Group
University of Sydney Business School
https://www.sydney.edu.au/business/our-research/research-groups/time-series-and-forecasting.html

V: November 29, 2022

Contents

The University of Sydney-CBD Campus Map

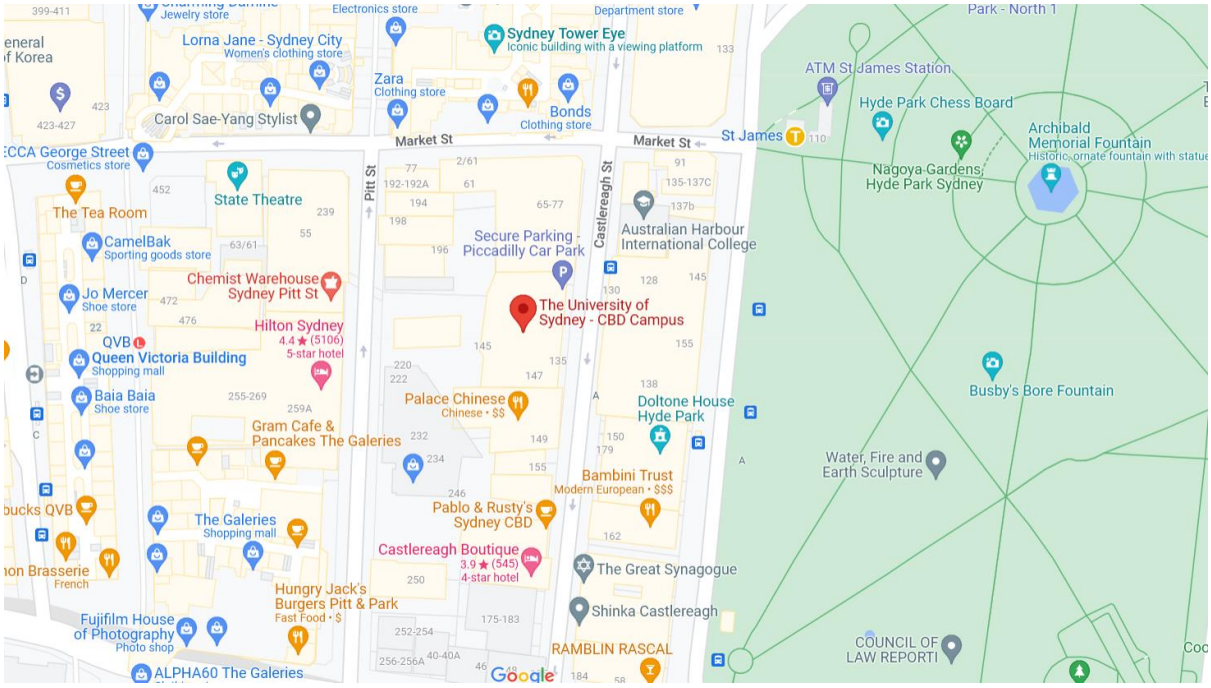

Address:

133 Castlereagh St (C13B), Sydney NSW 2000

Symposium Information

Time Series and Forecasting Research Group Leaders

Richard Gerlach
Discipline of Business Analytics, University of Sydney Business School
Boris Choy
Discipline of Business Analytics, University of Sydney Business School

Local Organising Committee

Boris Choy (Co-Chair)
Discipline of Business Analytics, University of Sydney Business School
Nuttanan Wichitaksorn (Co-Chair)
Department of Mathematical Sciences, Auckland University of Technology
Eric Cheung
School of Risk and Actuarial Studies, University of New South Wales Business School
Jianxin Wang
Finance Discipline Group, University of Technology, Sydney
Chao Wang
Discipline of Business Analytics, University of Sydney Business School
Shiying (Alice) Gao
Discipline of Business Analytics, University of Sydney Business School

Location

The University of Sydney Business School **CBD Campus**, Level 17, 133 Castlereagh Street, Sydney, NSW 2000

Registration

All in-person registrations include refreshments, lunches and symposium dinner.

Presentations

Keynote talks will be 50 minutes including Q&A, invited talks will be 40 minutes including Q&A, contributed talks will be 25 minutes including Q&A, and student talks will be 20 minutes including Q&A.

Dinner

Thursday, 1 December 2022, 6:30pm. Venue: Sky Phoenix, Level 6, Westfield Sydney, 188 Pitt Street, Sydney , NSW 2000

Further Queries

If you have any queries please do not hesitate to contact: tsf.symposium@sydney.edu.au or boris.choy@sydney.edu.au.

Symposium Schedule

Day 1: Thursday, 1 December 2022

8:45-9:00	**Registration**
9:00-9:10	**Welcome – Richard Gerlach (Co-Lead of TSFRG) and Boris Choy (Co-Chair of TSF2022)**
Morning Session 1	**Session Chair: Boris Choy**
9:10-10:00	**Keynote Talk 1** *Loss-Based Bayesian prediction* **Gael Martin**, Monash University
10:00-10:50	**Contributed Talk 1** *GMM estimation for moment condition models with time-varying parameters* **Yu Bai**, Monash University *Forecasting with machine learning models that are aware of known properties of the time series* **Pablo Montero-Manso**, University of Sydney
10:50-11:10	Morning Tea
Morning Session 2	**Session Chair: Jennifer S.K. Chan**
11:10-12:50	**Invited Talk 1** *Is dimensionality reduction a curse? Bayesian analysis of the mean-volatility dynamic factor model* **Mengheng Li**, University of Technology Sydney
11:50-12:40	**Contributed Talk 2** *A machine learning attack on illegal trading* **Artem Prokhorov**, University of Sydney *A mixture modeling framework for temporal point processes with memory* **Xiaotian Zheng**, University of Wollongong
12:40-13:40	Lunch
Afternoon Session 1	**Session Chair: Anastasios Panagiotelis**
13:40-14:20	**Invited Talk 2** *Threshold AR nearest-neighbour models for claims reserving* **Ken Siu**, Macquarie University
14:20-15:10	**Contributed Talk 2** *Intraday foreign exchange rate volatility forecasting: univariate and multilevel functional GARCH models* **Hanlin Shang**, Macquarie University *Optimal reconciliation with immutable forecasts* **Anastasios Panagiotelis**, University of Sydney
15:10-15:30	Afternoon Tea
Afternoon Session 2	**Session Chair: Mahdi Abolghasemi**
15:30-16:10	**Invited Talk 3** *Long monthly European temperature series and the North Atlantic Oscillation* **Timo Terasvirta**, Aarhus University
16:10-17:00	**Contributed Talk 3** *Optimal forecasting for optimising decisions* **Mahdi Abolghasemi**, University of Queensland *Building multivariate time-varying smooth transition correlation GARCH models, with an example to the four largest Australian banks* **Annastiina Silvennoinen**, Queensland University of Technology
18:30	**Banquet/Dinner** Venue: Sky Phoenix, Level 6, Westfield Sydney, 188 Pitt Street, Sydney, NSW 2000

Day 2: Friday, 2 December 2022

Morning Session 1	**Session Chair: Nate Wichitaksorn**
9:00-9:50	**Keynote Talk 2** *Singular VARs* **Rodney Strachan**, University of Queensland
9:50-10:40	**Contributed Talk 4** *Partial identification of heteroskedastic structural VARs: Theory and Bayesian inference* **Tomasz Woazniak**, Monash University *Electricity price spike clustering: A zero-inflated PARX approach* **Ye Lu**, University of Sydney
10:40-11:00	Morning Tea
Morning Session 2	**Session Chair: Eric Cheung**
11:00-11:40	**Invited Talk 4** *Optimal relativities, profitability, and efficiency in a modified Bonus-Malus system* **Jae Kyung Woo**, University of New South Wales
11:40-12:05	**Contributed Talk 5** *A bivariate Laguerre expansions approach for joint ruin probabilities in a two-dimensional insurance risk process* **Eric C.K. Cheung**, University of New South Wales
12:05-12:45	**Invited Talk 5** *Predicting volatilities, correlations and returns of stock indices using multivariate conditional autoregressive range and return models* **Jennifer S.K. Chan**, University of Sydney
12:45-13:45	Lunch
Afternoon Session 1	**Session Chair: Artem Prokhorov**
13:45-14:25	**Invited Talk 6** *Bespoke realized volatility: Tailored measures of risk for volatility prediction* **Andrew J. Patton**, Duke University
14:25-15:15	**Contributed Talk 6** *Statistical inference for aggregation of Malmquist productivity indices* **Valentin Zelenyuk**, University of Queensland *Empirical dynamic modeling: Automatic causal inference & forecasting for time series* **Patrick Laub**, University of New South Wales
15:10-15:35	Afternoon Tea
Afternoon Session 2	**Session Chair: Shiying (Alice) Gao**
15:35-17:35	**Student Talk** *A sparse dynamic factor model for clustered high-dimensional time series* **Jianjie Shi**, Monash University *Bayesian value at risk forecast using realized conditional autoregressive expectile model with an application of cryptocurrency* **Niya Chen**, University of Sydney *Sama circular model, SARIMA and decomposition techniques in forecasting women's economic empowerment in NSW Australia* **Samanthi Konarasinghe**, Western Sydney University *Stochastic volatility model with correlation model, a generalization of DCC-GARCH model* **Zheng Lyu Huang**, University of Sydney *Distribution vector autoregression: Eliciting macro and financial dependence* **Yunyun Wang**, Monash University *Machine learning methods for forecasting accuracy of crude oil prices: A comparative study* **Adel Gadhi**, University of Sydney
17:35-17:45	Short Break
17:45	**Best Paper Award and Closing Ceremony**

List of Authors and Titles

Optimal forecasting for optimising decisions

Mahdi Abolghasemi
School of Mathematics and Physics, University of Queensland
m.abolghasemi@uq.edu.au

Forecasting is the base for many managerial decisions that involves uncertainty. Forecasting per se is not the end goal but often a means to a final goal which may involve making decision for a problem, e.g., determining the inventory level, scheduling some activities, production planning. From a forecasting perspective we are interested to know how does forecasting accuracy translate into decisions in downstream. The literature suggest that more accurate forecasts do not necessarily lead to better decisions but there may be a positive correlation between them. Nevertheless, it is not evident how forecasts can impact the final decisions, and the process of incorporating useful information from downstream decisions to the prediction model is still open for researchers and practitioners. I will walk you through our winning method for IEEE 3rd technical challenge where the goal was to forecast energies and accordingly solve a timetabling problem to minimise the electricity cost. I will discuss how the forecast accuracy translates to better decisions and provide insights on how we can generate optimal forecasts by considering decisions.

GMM estimation for moment condition models with time-varying parameters

Yu Bai
Monash University
yu.bai1@monash.edu

I develop time-varying continuously updated GMM estimation and inferential theory for models whose parameters vary smoothly over time. Pointwise consistency and asymptotic normality are established. It is shown that identification requirements could be weaker than the constant coefficient case in some applications. Time-varying version of the overidentification test is also proposed. I illustrate the new method using simulated and real data.

Predicting volatilities, correlations and returns of stock indices using multivariate conditional autoregressive range and return models

Jennifer S.K. Chan
School of Mathematics and Statistics, University of Sydney
jchan@maths.usyd.edu.au

This paper extends the conditional autoregressive range (CARR) model to multivariate CARR (MCARR) model and further to the two-stage MCARR-return model to model and forecast volatilities, correlations and returns of multiple financial assets. The first stage model fits the scaled realised Parkinson volatility measures using individual series and their pairwise sums of indices to the MCARR model to obtain in-sample estimates and forecasts of volatilities for these individual and pairwise sum series. Then covariances are calculated to construct the fitted variance-covariance matrix of returns which are imputed into the stage-two return model to capture the heteroskedasticity of assets' returns. We investigate different choices of mean functions to describe the volatility dynamics. Empirical applications are based on the Standard and Poor 500, Dow Jones Industrial Average and Dow Jones United States Financial Service Indices. Results show that the stage-one MCARR models using asymmetric mean functions give better in-sample model fits than those based on symmetric mean functions. They also provide better out-of-sample volatility forecasts than those using CARR models based on two robust loss functions with the scaled realised open-to-close volatility measure as the proxy for the unobserved true volatility. We also find that the stage-two return models with constant means and multivariate Student-t errors give better in-sample fits than the BEKK-GARCH models. The estimates and forecasts of value-at-risk (VaR) and conditional VaR based on the best MCARR-return models for each asset are provided and tested using Kupiec test to confirm the accuracy of the VaR forecasts.

Bayesian Value At Risk Forecast Using Realized Conditional Autoregressive Expectiel Mdodel With An Application of Cryptocurrency*

Niya Chen
School of Mathematics and Statistics, University of Sydney
nche6444@uni.sydney.edu.au

In the financial market, risk management helps to minimize potential loss and maximize profit. There are two ways to assess risks, first way is to calculate the risk directly based on the volatility. The most common risk measurements are Value at Risk (VaR), sharp ratio, and beta. Alternatively, we could look at the quantile of the return to access the risk. Popular return models such as GARCH and stochastic volatility (SV) focus on modeling the mean of the return distribution via capturing the volatility dynamics, however, the quantile/expectile method will give us an idea of the distribution with the extreme return value. It will allow us to forecast VaR using return which is direct information. The advantage of using these non-parametric methods is that it is not bounded by the distribution assumptions from the parametric method. But the difference between them is that expectile uses a second-order loss function while quantile regression uses a first-order loss function. We consider several quantile functions, different volatility measures, and estimates from some volatility models. To estimate the expectile of the model, we use Realized Conditional Autoregressive Expectile (CARE) model with the bayesian method to achieve this. We would like to see if our proposed models outperform existing models in cryptocurrency and we will test it by using Bitcoin mainly as well as Ethereum.

A bivariate Laguerre expansions approach for joint ruin probabilities in a two-dimensional insurance risk process

Eric C.K. Cheung, Hansjoerg Albrecher, Haibo Liu, Jae Kyung Woo
School of Risk and Actuarial Studies, University of New South Wales
eric.cheung@unsw.edu.au

In this paper, we consider a two-dimensional insurance risk model where each business line faces not only stand-alone claims but also common shocks that induce dependent losses to both lines simultaneously. The joint ruin probability is analyzed, and it is shown that under certain model assumptions it can be expressed in terms of a bivariate Laguerre series. For computational purposes, the bivariate Laguerre series needs to be truncated, and the corresponding Laguerre coefficients can be obtained through a system of linear equations. The computational procedure is easy to implement, and our numerical examples illustrate its excellent performance. The results are also applied to address a related capital allocation problem. This part is joint work with Hansjoerg Albrecher, Haibo Liu and Jae-Kyung Woo.

Machine learning Methods for Forecasting Accuracy of Crude Oil Prices: A Comparative Study*

Adel Gadhi
School of Mathematics and Statistics, University of Sydney
agad0777@uni.sydney.edu.au

The difficulties faced in the prediction of the global oil market have recently attracted a lot of attention by the researchers as the movementsin recent prices are influenced bymany external factors. The nonlinear and chaotic behavior of the prices make it difficult to predict the future and this, in turn affect the overall economic activities in many countries, especially in developing countriesand brings risk to all oil related enterprises. In order to mitigate the risk and uncertainty in oil price fluctuations, the oil related enterprises need to make better hedging, future investments, and proper evaluations in decision making for better forecasting of trends in prices (usually undertaken in U.S dollars currency terms). In this study

four predictive methods are compared to find the best approach inpredicting daily crude oil prices during the period of 23 July 2007to 27 June 2022 using a total of 2643 observations. The models applied are the Autoregressive Integrated Moving Average (ARIMA), Random Forests (RF), Decision Trees (DT) and Support Vector Regressions (SVR). The performances are evaluated based on three statistical error measures, namely the Mean Absolute Error (MAE), the Mean Square Error (MSE) and the Coefficient of Determination (R^2),in order to identify the best performing model. The results of the analysis show that the SVR performs the best in forecasting crude oil‘ when compared with the other two methods.

Stochastic volatility model with correlation model, a generalization of DCC-GARCH model*

Zheng Lyu Huang
School of Mathematics and Statistics, University of Sydney
zhua8113@uni.sydney.edu.au

This paper proposes to model some time series of observed covariance matrices using a matrix distribution such as Wishart distribution. We adopt the idea of DCC-GARCH model to model the mean matrix of the distribution via modeling the correlation matrix as well as the volatility of each component. We use Rstan to perform model fitting under the Bayesian approach. We perform simulation studies to assess the accuracy of estimation under two scenarios, with varying shape parameters and long-term correlation coefficient. We apply the proposed models to some time series of observed covariance matrices for Bitcoin and Ethereum in the crypto market. We also extend the mean matrix of the covariance matrix model by incorporating leverage effect.

Sama Circular Model, SARIMA and Decomposition Techniques In Forecasting Women's Economic Empowerment in NSW Australia*

W. G. Samanthi Konarasinghe
Western Sydney University
sinasisi@gmail.com

Economic empowerment is the capacity of women and men to participate in, contribute to and benefit from growth processes in ways that recognise the value of their contributions, respect their dignity and make it possible to negotiate a fairer distribution of the benefits of growth. Economic empowerment of women is important to the society in various ways. Economically empowered women contribute to the economy of the family, health care of family members, education of children and more. The study is aimed to forecast the economic empowerment of females in New South Wales (NSW), Australia. The employment data were obtained from the Australian bureau of Statistics. The data series follows an irregular wave like pattern. The Decomposition techniques, Seasonal Auto Regressive Integrated Moving Average (SARIMA) and Sama Circular Model (SCM) were tested for forecasting. Time Series plots and Auto Correlation Functions were used for pattern recognition. The Auto Correlation Functions (ACF) of residuals and Ljung-Box Q statistics (LBQ) were used to test the independence of residuals. The Anderson Darling test was used to test the normality of residuals. Forecasting ability of the models was assessed by Mean Square Error (MSE) and Mean Absolute Deviation (MAD). Results revealed that the SARIMA and SCM are suitable for the purpose. However, the pattern of SARIMA forecasts does not follow the actual data whilst the SCM forecasts does, hence the SCM is superior to SARIMA in this context. It is recommended to foresee the women's economic empowerment in the other states of Australia

Empirical Dynamic Modeling: Automatic Causal Inference & Forecasting for Time Series

Patrick Laub, Jinjing Li, Michael Zyphur, George Sugihara
School of Risk and Actuarial Studies, University of New South Wales
p.laub@unsw.edu.au

How can social and health researchers study complex dynamic systems that function in nonlinear and even chaotic ways? Common methods, such as experiments and equation-based models, may be ill-suited to this task. To address the limitations of existing methods and offer nonparametric tools for characterising and testing causality in nonlinear dynamic systems, we created empirical dynamic modeling (EDM) packages for Stata, R, and Python. The packages implement the key EDM methods for time series and panel data. In particular, it implements convergent cross-mapping, which offers a nonparametric approach to modeling causal effects. We can observe these algorithms in action on simulated data, and on real daily Chicago temperature and crime, showing an effect of temperature on crime but not the reverse. This is joint work with Jinjing Li, Michael Zyphur, and George Sugihara.

Is dimensionality reduction a curse? Bayesian analysis of the mean-volatility dynamic factor model

Mengheng Li
Economics Discipline Group, University of Technology Sydney
mengheng.li@uts.edu.au

The factor stochastic volatility (FSV) model is a powerful dimensionality reduction device and has gained much popularity due to its parsimonious structure in modelling both time-varying mean and covariance matrix for multivariate time series. We document the failure of FSV models by observing a strong common volatility component left in the residuals, irrespective of the chosen number of factors. To adequately model co-movement in mean and volatility, we introduce the mean-volatility dynamic factor model which assumes separate factor structures for the first and the second moment of a high-dimensional vector time series. We identify and extract the mean factors and volatility factors via a Bayesian variable selection technique that pins down zeros in associated loading matrices and thus the factor space. Such identification scheme is order-invariant and gives economically meaningful factors. We also propose a computationally efficient multi-move sampler that samples all volatility series in parallel to speed up estimation. In the empirical study, we fit the model to the Fred-MD data of 128 monthly macroeconomic variables and find group-specific mean factors and economy-wide volatility factors.

Electricity Price Spike Clustering: A Zero-Inflated PARX Approach

Ye Lu
School of Economics, University of Sydney
ye.lu1@sydney.edu.au

Sudden dramatic rises in electricity prices, known as electricity "price spikes", are ubiquitous in electricity spot markets worldwide. These price spikes often cluster. Energy retailers in many countries, including Australia, purchase electricity in an unregulated spot market and sell it to consumers at a heavily regulated price. As a result, the occurrence and clustering of spikes in the spot electricity price is particularly hazardous for retailers, and effective modelling and forecasting of these spikes is of foremost importance for effective risk management. This paper proposes a zero-inflated Poisson autoregression model with exogenous covariates (ZIPARX) for modelling the time series count of price spikes. This model is shown to well capture the salient features of electricity price spikes and can provide effective rolling probability forecasts of price spike occurrence. We apply the proposed approach to analyze the spot prices in four major regions in the Australian National Electricity Market (NEM) over the period April 2016 to February 2022. Results show that price spike clustering is captured very well using our model, and the persistence in price spike occurrence remains even after controlling for the relevant exogenous covariates. We also find that, in most of our sample period, the increasing use of renewable generation dampens the price spike intensity in all regions except South Australia.

Loss-Based Bayesian Prediction

Gael Martin
Department of Econometrics and Business Statistics, Monash University
gael.martin@monash.edu

The conventional paradigm for Bayesian prediction is underpinned by the assumption that the true data generating process is either equivalent to the predictive model adopted, or spanned by a finite set of models over which we average. Such a heroic assumption is clearly at odds with reality; in particular in the realm of the social and economic sciences where statistical data arise through complex processes that we can only ever hope to approximate. In response to this limitation of the conventional paradigm, we propose an alternative approach to Bayesian prediction. We begin by defining a predictive model that we believe to be a 'plausible' mechanism for generating predictions. A prior is placed over the parameters of this model, and the prior then updated to a posterior via a criterion function that captures a user-specified measure of predictive accuracy – or predictive 'loss'. By construction, the posterior that results gives high mass to points in the parameter space that yield predictive accuracy according to the specified measure. In turn, the usual Bayesian calculus produces a predictive distribution that is tailored to yield predictions that are accurate in the given measure – despite the misspecification of the model. The theoretical validity of this loss-based approach has been established, and a form of optimality demonstrated, including when an approximation to the loss-based predictive distribution is produced via variational methods. In extensive simulation experiments and empirical explorations using a range of predictive models – including both high-dimensional and state space models – we find notable gains in predictive accuracy relative to conventional likelihood-based Bayesian prediction.

Forecasting with Machine Learning models that are aware of known properties of the time series

Pablo Montero-Manso
Discipline of Business Analytics, University of Sydney
pablo.monteromanso@sydney.edu.au

Machine Learning models are notoriously overly flexible, particularly for time series. This is due to two main reasons: (1) lack of data, many time series are a few dozens of observations long. (2) Many relevant time series processes are known to be not complex, for example exponential or sigmoid growth curves or cyclical patterns. The novel paradigm of Data-pooling/ Global models /Cross-Learning mitigates the sample size issue for time series. However, there is very little research on the issue of simplicity. Even in current times, the amount of data is limited, and the modern general view in Machine Learning and AI is to impose some 'known properties' from the problem we try to solve into the model, creating a restricted or simpler model. For example, convolutional neural networks are a simpler version of a multilayer perceptron because it imposes translation invariance(a picture of a cat can be classified as 'cat' wherever the cat appears). Similar properties have been imposed in time series, for example autoregressions, but there is much more structure that we could be adding to the models to make them better. We will cover a few of these properties, how they can be imposed into deep learning models and show experimental results. Some of the properties are: Positive only (e.g. count data), Always Increasing (e.g. cumulative data), Integer only (count), invariant to noise (measurement error), multi-scale (temporal aggregation), time warping invariance, shift equivariance, scale equivariance, and even linearity.

Optimal reconciliation with immutable forecasts

Anastasios Panagiotelis, Bohan Zhang, Yanfei Kang, Feng Li
Discipline of Business Analytics, University of Sydney
anastasios.panagiotelis@sydney.edu.au

The practical importance of coherent forecasts in hierarchical forecasting has inspired many studies on forecast reconciliation. Under this approach, base forecasts are produced for every series in the hierarchy and are subsequently adjusted to be coherent in a second reconciliation step. Reconciliation methods have been shown to improve forecast accuracy but will generally adjust the base forecast of every series. However, in an operational context, it is sometimes necessary or beneficial to keep forecasts of some variables unchanged after forecast reconciliation. In this paper, we formulate a reconciliation methodology that keeps forecasts of a pre-specified subset of variables unchanged or "immutable". In contrast to existing approaches, these immutable forecasts need not all come from the same level of a hierarchy, and our method can also be applied to grouped hierarchies. We prove that our approach preserves unbiasedness in base forecasts. Our method can also account for correlations between base forecasting errors and ensure the non-negativity of forecasts. We also perform empirical experiments, including an application to a large-scale online retailer's sales, to assess our proposed methodology's impacts.

Bespoke Realized Volatility: Tailored Measures of Risk for Volatility Prediction

Andrew J. Patton, Haozhe Zhang
Department of Economics, Duke University
andrew.patton@duke.edu

This paper presents a method to tailor the construction of an estimate of asset price volatility to the forecasting model in which it is subsequently employed. Existing estimators of volatility, such as realized volatility (RV), see Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002), as well as numerous extensions and refinements, measure volatility independent of how the measure is later used. We exploit machine learning methods to tailor the measure of volatility to the application and, using data on 886 US stocks over the period 1995-2019, show that our "bespoke RV" leads to significantly improved forecasts. We find that the forecast improvements from using bespoke RVs arise from because they are more responsive to recent news, and they are able to re-weight the underlying high frequency returns to capture an information effect and a news-announcement effect.

A Machine Learning Attack on Illegal Trading

Artem Prokhorov, Robert James, Henry Leung
Discipline of Business Analytics, University of Sydney
artem.prokhorov@sydney.edu.au

We design an adaptive framework for the detection of illegal trading behavior. Its key component is an extension of a pattern recognition tool, originating from the field of signal processing and adapted to modern electronic systems of securities trading. The new method combines the flexibility of dynamic time warping with contemporary approaches from extreme value theory to explore large-scale transaction data and accurately identify illegal trading patterns. Importantly, our method does not need access to any confirmed illegal transactions for training. We use a high-frequency order book dataset provided by an international investment firm to show that the method achieves remarkable improvements over alternative approaches in the identification of suspected illegal insider trading cases.

Intraday foreign exchange rate volatility forecasting: univariate and multilevel functional GARCH models

Hanlin Shang
Department of Actuarial Studies and Business Analytics, Macquarie University
hanlin.shang@mq.edu.au

This paper seeks to predict conditional intraday volatility in foreign exchange (FX) markets using functional Generalized AutoRegressive Conditional Heteroscedasticity (GARCH) models. We contribute to the existing models by accounting for the stylised features of long-range conditional heteroscedasticity and cross-dependence in major FX currencies through long-range dependent and multi-level functional principal component basis functions in estimating the functional GARCH-type models, as well as incorporating the intraday bid-ask spread microstructure information. Overall, we demonstrate the statistical and economic superiority of appropriately modelling FX volatility using various functional GARCH based models. Remarkably, we find that taking account of cross-dependency dynamics between the major currencies can significantly improve intraday conditional volatility forecasting in the FX market. Intraday risk management benefits and inter-daily asset allocation applications are presented to highlight the practical benefits of our proposed approaches.

A Sparse Dynamic Factor Model for Clustered High-dimensional Time Series*

Jianjie Shi, Catherine Forbes, Dan Zhu
Department of Econometrics and Business Statistics, Monash University
jianjie.shi@monash.edu

Dynamic factor models are widely-used to model high-dimensional time series data. While many applications are based on reducing the model into a minimal number of factors, empirical evidence suggests many factors may be required to capture the complex dynamics observed in the data. However, permititng a large number of factors will often lead to statistical problems, particularly in limited data contexts. In this article, we develop a new sparse dynamic factor model that substantially reduces the number of parameters, and caters for high-dimensional time series data containing heterogeneous patterns including clustering. In particular, our method provides a data-driven classification of the clustered high-dimensional time series witha single-factor model within each class. While each factor drives its own intra-cluster co-movement, the joint evolution of the factors induces inter-cluster dependency. In an application to French mortality data, our proposed model demonstrates superior forecasting performances when compared to a class of dynamic factor models, providing valuable insights for longevity risk management.

Building Multivariate Time-Varying Smooth Transition Correlation GARCH models, with an example to the four largest Australian banks

Annastiina Silvennoinen, Anthony D. Hall, Timo Teräsvirta
School of Economics and Finance, Queensland University of Technology
silvennoinen@qut.edu.au

This paper looks at changes in the correlations of daily returns between the four major banks in Australia. Revelations from the analysis are of importance to investors, but also to government involvement, due to the large proportion of the highly concentrated financial sector relying on the stability of the Big Four. For this purpose, a methodology for building Multivariate Time-Varying STCC-GARCH models is developed. The novel contributions in this area are the specification tests related to the correlation component, the extension of the general model to allow for additional correlation regimes, and a detailed exposition of the systematic, improved modelling cycle required for such nonlinear models. There is an R-package that includes the steps in the modelling cycle. Simulations evidence the robustness of the recommended model building approach. The empirical analysis reveals an increase in correlations of the Australia's four largest banks that coincides with the stagnation of the home loan market, technology changes, the mining boom, and Basel II alignment, increasing the exposure of the Australian financial sector to shocks.

Singular VARs

Rodney Strachan
School of Economics, University of Queensland
r.strachan@uq.edu.au

There exists a strong rationale for the adoption of multivariate time series that have a strictly singular spectral density. Further, these processes have a well-developed theoretical framework. This paper develops formal methods for the empirical analysis of such singular processes. The assumption of a singular spectral density is consistent with economic theory underlying, for example, Dynamic Stochastic General Equilibrium (DSGE) models in which the number of variables is greater than the number of structural shocks. While this assumption guarantees the existence of a finite order VAR representation under mild regularity conditions, there does not exist a unique probability density function with respect to the Lebesgue measure. We overcome this issue by defining a density on a compact submanifold with respect to the Hausdorff measure. Accordingly, we develop an HMC algorithm that jointly samples model parameters, the VAR lag length, as well as the number of shocks in a fully specified Bayesian framework. The effectiveness of the methodology is demonstrated in an extensive Monte Carlo exercise involving a multi-sector DSGE model. Finally, we use the proposed framework to carry out structural analysis on US macroeconomic data in a sample involving COVID shocks.

Threshold AR Nearest-Neighbour models for claims reserving

Ken Siu
Department of Actuarial Studies and Business Analytics, Macquarie University
ken.siu@mq.edu.au

Motivated by claims reserving in run-off triangles, a class of threshold autoregressive nearest-neighbour (TAR-NN) models extending a major class of parametric nonlinear time series models, namely threshold autoregressive (TAR) models, is introduced. The proposed class of models also introduces a flexible regime-switching mechanism to nearest-neighbour models. Attention is given to a sub-class of TAR-NN models, namely self-exciting threshold autoregressive nearest-neighbour models (SETAR-NN), for uses in claims reserving. The (strict) stationarity and geometric ergodicity of the SETAR-NN model, and more generally, a two-dimensional nonlinear autoregressive random field, are discussed. The conditional least-square (CLS) method is used to estimate the SETAR-NN model and some of its nested models. Simulation studies on the parameter estimates from the CLS method are conducted. Using real insurance claims data and stochastic simulations, the applications of the SETAR-NN model and the nested models for projecting future claims liabilities are discussed. Comparisons of those models with the Bootstrap-Chain-Ladder (BCL) model for claims reserving are provided.

Long Monthly European Temperature Series and the North Atlantic Oscillation

Timo Teräsvirta, Jian Kang, Changli He, Annastiina Silvennoinen
School of Business and Social Sciences, Aarhus University
tterasvirta@econ.au.dk

In this presentation, the relationship between the surface air temperatures in 28 European cities and towns and the North Atlantic Oscillation (NAO) is modelled using the Vector Seasonal Shifting Mean and Covariance Autoregressive model, extended to contain exogenous variables. Central statistical and time series features of the model are discussed before moving on to discussing data and showing empirical results. The model also incorporates season-specific spatial correlations that are functions of latitudinal, longitudinal, and elevation differences of the various locations.

The empirical results, based on long monthly time series, agree with previous ones in the literature in that the NAO is found to have its strongest effect on temperatures during winter months. The transition from the boreal winter to the summer is not monotonic, however. The strength of the error correlations of the model between locations is inversely related to the distance between the locations, with a slower decay in the east-west than north-south direction. Altitude differences also matter but only during the boreal winter half of the year

Distribution Vector Autoregression: Eliciting Macro and Financial Dependence*

Yunyun Wang
Department of Econometrics and Business Statistics, Monash University
yunyun.wang@monash.edu

The global financial crisis prompted an increased focus on the role of financial factors in driving real economic fluctuations. Vector autoregression models are popular in this literature, providing simple yet insightful information such as the impulse response function of different shocks. This paper develops a flexible and robust alternative based on a multivariate Distribution Regression method. The resulting distribution impulse response function provides a more comprehensive picture of the dynamic heterogeneity. In the study of the U.S. GDP growth and financial conditions, the empirical results from our new framework confirm some existing findings in the literature that 1) the tight financial condition creates multimodality in the conditional joint distribution, and 2) restricting the upper tail of financial condition has a noticeable impact on long-term GDP growth. Yet, the extracted information on the effect of restricting the lower tail of GDP during the global financial crisis suggests an alternative conclusion, i.e., negligible impact on financial condition.

Optimal relativities, profitability, and efficiency in a modified Bonus-Malus system

Jae Kyung Woo
School of Risk and Actuarial Studies, University of New South Wales
j.k.woo@unsw.edu.au

In the classical Bonus-Malus System (BMS) in automobile insurance, the premium for the next year is adjusted according to the policyholder's claim history (particularly frequency) in the previous year. Some variations of the classical BMS have been considered by taking more of the driver's claim experience into account to better assess an individual's risk. In this talk, we revisit a modified BMS briefly introduced by Lemaire (1995) and Pitrebois et al. (2003a). Specifically, such a BMS extends the number of Bonus-Malus (BM) levels due to an additional component in the transition rules representing the number of consecutive claim-free years. With the extended BM levels granting a more reasonable bonus to careful drivers, this paper investigates the transition rules more rigorously and provides the optimal BM relativities under various statistical model assumptions, including the frequency random effect model and the dependent collective risk model. Also, a numerical analysis of a real data set is provided to compare the classical BMS and our proposed BMS. Finally, some remarks regarding overall stationary premium and efficiency measures are provided. Most of the work is based on joint work with Ahn, Cheung, and Oh (funded by the Casualty Actuarial Society).

Partial Identification of Heteroskedastic Structural VARs:Theory and Bayesian Inference

Tomasz Woźniak, Helmut Lütkepohl, Fei Shang, Luis Uzeda
Department of Economics, University of Melbourne
tomasz.wozniak@unimelb.edu.au

We propose structural vector autoregressive models in which the structural parameters are identified via a stochastic volatility process for time-varying conditional variances. Our focus is on the question of how many and what shocks are identified via heteroskedasticity. Therefore, we derive a set of parametric restrictions under which the structural matrix is partially or globally unique, and Savage-Dickey density ratios are used to assess the validity of the identification conditions. We propose a shrinkage prior distribution for conditional log-volatilities and variances that is centred on a hypothesis of homoskedasticity, which assures that the evidence for the identification of the structural shocks is providedby the data. We apply identification through heteroskedasticity to estimate the dynamic output effects of unanticipated changes in tax policy that have been identified in previous studies by exclusion restrictions as well as by using narrative measures as proxies or time-varying volatility.

Statistical Inference for Aggregation of Malmquist Productivity Indices

Valentin Zelenyuk, Manh D. Pham, Léopold Simar
School of Economics and Centre for Efficiency and Productivity Analysis, University of Queensland
v.zelenyuk@uq.edu.au

The Malmquist Productivity Index (MPI) has gained popularity amongst studies on the dynamic change of productivity of decision-making units (DMUs). In practice, this index is frequently reported at aggregatelevels (e.g., public and private firms) in the form of simple equally-weighted arithmetic or geometric meansof individual MPIs. A number of studies have emphasized that it is necessary to account for the relativeimportance of individual DMUs in the aggregations of indices in general and of MPI in particular. While more suitable aggregations of MPIs have been introduced in the literature, their statistical properties have not been revealed yet, preventing applied researchers from making essential statistical inferences such as confidence intervals and hypothesis testing. In this paper, we will fill this gap by developing a full asymptotic theoryfor an appealing aggregation of MPIs. On the basis of this, meaningful statistical inferences are proposed, their finite-sample performances are verified via extensive Monte Carlo experiments, and the importance ofthe proposed theoretical developments is illustrated with an empirical application to real data.

A Mixture Modeling Framework for Temporal Point Processes with Memory

Xiaotian Zheng, Athanasios Kottas, Bruno Sansó
School of Mathematics and Applied Mathematics, University of Wollongong
xzheng@uow.edu.au

We propose a constructive approach to building temporal point processes that in-corporate dependence on their history. The dependence is modeled through the condi-tional density of the duration, i.e., the interval between event times, using a mixture of first-order conditional densities for each one of a specific number of lagged durations. Such a formulation for the conditional duration density accommodates high-order, non-Gaussian dynamics, and thus it enables flexible modeling for point processeswith memory. The implied conditional intensity function admits a representation asa local mixture of first-order hazard functions. By specifying appropriate families ofdistributions for the first-order conditional densities that imply different shapes of the associated hazard functions, we can obtain either self-exciting or self-regulating pointprocesses. From the perspective of duration processes, we obtain a tool to set a stationary marginal density. The resulting model, interpreted as a dependent renewal process, introduces high-order Markov dependence among identically distributed durations. Furthermore, we provide extensions to cluster point processes. These can describe duration clustering behaviors attributed to different factors, expanding the scope of the modeling framework to a wider range of applications. Regarding implementation, we develop a Bayesian approach to inference and model checking. We investigate point process model properties analytically, and illustrate the methodology with both simulations and an analysis of the market microstructure in the foreign exchange market.

Index of Authors