THE UNIVERSITY OF SYDNEY
Business School

Time Series and Forecasting Symposium TSF2023

20-21 November 2023

The University of Sydney, CBD Campus

Level 17, 133 Castlereagh Street, Sydney, Australia

Acknowledgement of Country

We acknowledge the traditional owners of the lands on which the University of Sydney is located, the Gadigal people of the Eora Nation, and we pay our respect to the knowledge embedded forever within the Aboriginal Custodianship of Country.

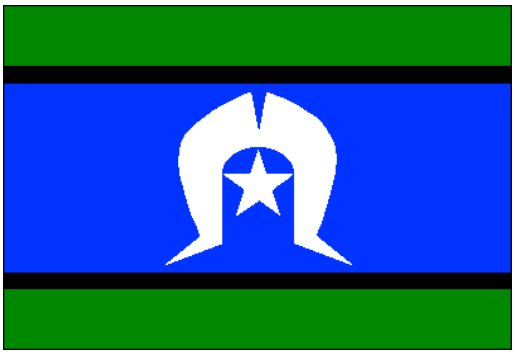

Location Guide

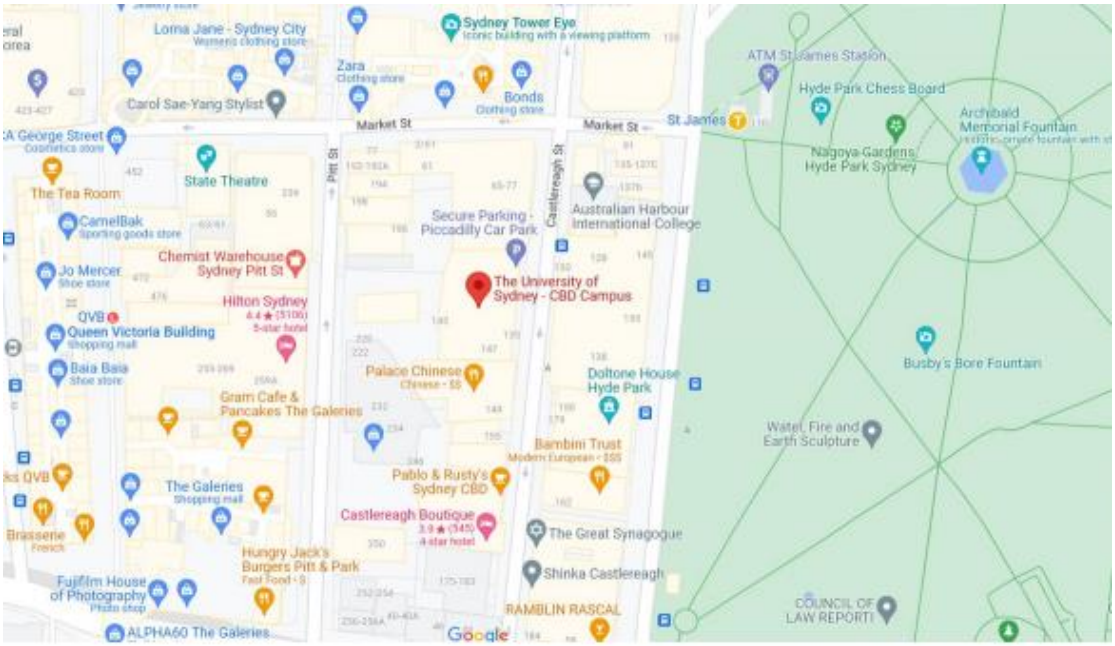

The University of Sydney, CBD Campus

Address: Level 17, 133 Castlereagh Street, Sydney, Australia

Getting There

By train
The nearest train stations are St James Station and Town Hall Station.

By bus
The nearest bus stops are located on Castlereagh Street or on Elizabeth Street.

By car
There is no car parking available onsite. However, Piccadilly Secure Parking (137 Castlereagh St) provides parking spaces. We suggest pre-booking online to avoid a high cost.

Symposium Details

The Time Series & Forecasting Symposium is an annual research event of the Time Series and Forecasting Research Group at the University of Sydney Business School. This symposium aims to promote time series analysis and forecasting in business and other areas. It consists of oral and poster presentations in all areas related to time series and forecasting. The main themes are: time series econometrics, volatility modelling and risk forecasting, risk assessment and management, high-dimensional modelling and forecasting, computational methods, robust inferences, machine learning, and deep learning.

Time Series and Forecasting Research Group Leaders

Richard Gerlach, Discipline of Business Analytics, The University of Sydney

Boris Choy, Discipline of Business Analytics, The University of Sydney

Organising Committee

Boris Choy (Co-Chair), Discipline of Business Analytics, The University of Sydney

Simon Kwok (Co-Chair), School of Economics, The University of Sydney

Alex Lam, Discipline of Business Analytics, The University of Sydney

Shelton Peiris, School of Mathematics and Statistics, The University of Sydney

Ken Siu, Department of Actuarial Studies and Business Analytics, Macquarie University

Nuttanan Wichitaksorn, Department of Mathematical Sciences, Auckland University of Technology

Registration

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

Presentation Format

Keynote and invited talks are 50 minutes each including Q&A
Contributed talks are 30 minutes each including Q&A
Student talks are 3 minutes each, and poster presentations are held throughout the two days

Dinner

Tuesday, 21 November 2023, 6pm
Venue: Sky Phoenix Chinese restaurant (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

Program

Day 1: Monday 20 November 2023

8.45 – 9.00 **Registration**

9:00 – 9:10	**Welcome** **Richard Gerlach** *Co-Lead, Time Series and Forecasting Research Group, Business School, University of Sydney* **Boris Choy** *Co-Lead, Time Series and Forecasting Research Group, Business School, University of Sydney* **Simon Kwok** *School of Economics, University of Sydney*
Morning Session 1 *Session Chair*: **Simon Kwok**	
9:10 – 10:00	**Keynote Talk 1** **Yongmiao Hong,** *University of Chinese Academy of Sciences* Forecasting Inflation Rates: A Large Panel Micro-Level Data Approach
10.00 – 10.50	**Invited Talk 1** **Peter C. B. Phillips,** *Yale University* 1974-1984 Revisited with Curved Autoregression

10:50 – 11:10 **Morning Tea**

Morning Session 2 *Session Chair*: **Richard Gerlach**	
11:10 – 12:10	**Contributed Talk 1** **Hanlin Shang,** *Macquarie University* Estimating Covariance Functions for High-Dimensional Functional Time Series with Dual Factor Structures **Won-Ki Seo,** *University of Sydney* Optimal Linear Prediction with Functional Observations: Why You Can Use a Simple Post-Dimension Reduction Estimator

12:10 – 13:10 **Lunch**

Afternoon Session 1 *Session Chair*: **Boris Choy**	
13:10 – 14:00	**Keynote Talk 2** **Rob Hyndman,** *Monash University* Forecast Reconciliation: A Brief Overview
14:00 – 15:00	**Contributed Talk 2** **Yu Bai,** *Monash University* Optimal Forecasting under Parameter Instability **Mahdi Abolghasemi,** *University of Queensland* Forecast and Optimise

15:00 – 15:20 **Afternoon Tea**

Afternoon Session 2 *Session Chair*: **Ken Siu**	
15:20 – 16:10	**Invited Talk 2** **Shuping Shi,** *Macquarie University* Realized Drift
16:10 – 17:10	**Contributed Talk 3** **Adam Clements,** *Queensland University of Technology* Modelling and Forecasting Intraday Spot Volatility **Paul Labonne,** *BI Norwegian Business School* Asymmetric Uncertainty: Nowcasting using Skewness in Real-Time Data

Day 2: Tuesday 21 November 2023

Morning Session 1 *Session Chair*: **Nuttanan Wichitaksorn**	
9:00 – 10:00	**Contributed Talk 4** **Pär Österholm,** *Örebro University* The Relation between Treasury Yields and the Corporate Bond-Yield Spread -- Stable or Time-Varying? **Klaus Ackermann,** *Monash University* Predictions in Shared Markets: A Global Forecasting Approach with Deep Learning and Spillover Considerations
10.00 - 10.50	**Keynote Talk 3** **Jiti Gao,** *Monash University* Estimation and Inference for Vector Time Series Models with Time-Varying Parameters

10:50 – 11:10 **Morning Tea**

Morning Session 2 *Session Chair*: **Hanlin Shang**	
11:10 – 12:10	**Contributed Talk 5** **Minh-Ngoc Tran,** *University of Sydney* Particle-based Variational Bayes: Towards Scalable and Accurate Bayesian Computation **Matthew Read,** *Reserve Bank of Australia* Bayesian Inference in SVARs Using 'Soft' Sign Restrictions

12:10 – 13:10 **Lunch**

Afternoon Session 1 *Session Chair*: **Minh-Ngoc Tran**	
13:10 – 14:00	**Invited Talk 3** **Timo Teräsvirta,** *Aarhus University* The Effect of the North Atlantic Oscillation on Monthly Precipitation in Selected European Locations: A Nonlinear Time Series Approach
14:00 – 15:00	**Contributed Talk 6** **Andrey Vasnev,** *University of Sydney* The Role of Data and Priors in Estimating Climate Sensitivity **Mengheng Li,** *University of Technology Sydney* Composite Inference on a Common Equilibrium Climate Sensitivity

15:00 – 15:20 **Afternoon Tea**

Afternoon Session 2 *Session Chair*: **Shelton Peiris**	
15:20 – 16:20	**Contributed Talk 7** **Anastasios Panagiotelis,** *University of Sydney* Boosting Domain-Specific Models with Shrinkage: An application in Mortality Forecasting **Gnanadarsha Dissanayake,** *NSW Ministry of Health* Forecasting Periodic Incidence of Australian Seasonal Influenza after Assessing Persistence with and without Machine Learning
Afternoon Session 3 *Session Chair*: **Alex Lam**	
16:20 – 17:10	**3-Minute Talk** **Chen Liu,** *University of Sydney* DeepVol: A Pre-Trained Universal Asset Volatility Model **Lin Deng,** *University of Melbourne* Efficient Variational Inference for Large Skew-t Copulas with Application to Intraday Equity Returns **Lingfeng Lyu,** *University of New South Wales* Hierarchical House Price Model incorporating Geographical and Macroeconomic Factors **Ye Shiqi,** *Xiamen University and Chinese Academy of Sciences* Multi-Matrix Autoregressive Models with an Application to Multi-Modal Network **Adel Gadhi,** *University of Sydney* Enhancing Financial Forecasting: A Comparative Study of Hybrid Deep Learning Models and Traditional Approaches in Oil Price Prediction **Taiga Saito,** *Monash University* Estimation and Inference of Average Treatment Effects using Deep Neural Networks **Arezoo Orooji,** *Macquarie University* Functional Principal Component Logit Regression for the Subarachnoid Hemorrhage Data **Rajan Shankar,** *University of Sydney* Robust Autoregressive Modelling using a Penalised Approach
17:10 – 17:25	Talk by **Susquehanna International Group (SIG)**
17:25 – 17:30	**Student Award Presentation**

18:00 **Symposium Dinner**

Venue: Sky Phoenix (Level 6, Westfield Sydney, 188 Pitt Street, Sydney)

List of Abstracts

Keynote and Invited Talks

Forecasting Inflation Rates: A Large Panel Micro-Level Data Approach

Yongmiao Hong
University of Chinese Academy of Sciences
yh20@cornell.edu

Abstract: Economists and econometricians typically use aggregate macroeconomic and financial data for inflation prediction. However, aggregation often results in a loss of valuable information, diminishing key features like heterogeneity, interactions, nonlinearity, and structural breaks. We propose a novel microeconometric approach to forecasting inflation rates, making use of a large panel of individual stock prices. By employing machine learning algorithms, we can effectively exploit this micro-level information to achieve substantially more accurate inflation forecasts. Our empirical findings highlight the advantages and potential of utilizing micro-level data for macro prediction, diverging from conventional macro-forecasting approaches that rely on aggregate data for forecasting macro variables.

1974-1984 Revisited with Curved Autoregression

Peter C. B. Phillips
Yale University
peter.phillips@yale.edu

Abstract: New asymptotic theory is reported for estimation, inference and prediction in autoregression with curved cross section time series, providing a new approach to dynamic panel regression with high dimensional dependent cross section data.

Forecast Reconciliation: A Brief Overview

Rob Hyndman
Monash University
rob.hyndman@monash.edu

Abstract: Collections of time series that are formed via aggregation are prevalent in many fields. These are commonly referred to as hierarchical time series and may be constructed cross-sectionally across different variables, temporally by aggregating a single series at different frequencies, or may even be generalised beyond aggregation as time series that respect linear constraints. When forecasting such time series, a desirable condition is for forecasts to be coherent; that is, to respect the constraints. The past decades have seen substantial growth in this field with the development of reconciliation methods that not only ensure coherent forecasts but can also improve forecast accuracy. This talk provides an overview of recent work on forecast reconciliation.

Realized Drift

Shuping Shi
Macquarie University
shuping.shi@mq.edu.au

Abstract: Drift and volatility are two mainsprings of asset price dynamics. While volatilities have been studied extensively in the literature, drifts are commonly believed to be impossible to estimate and largely ignored in the literature. This paper shows how to detect drift using realized autocovariance implemented on high-frequency data. We use a theoretical treatment in which the classical model for the efficient price, an Ito semimartingale possibly contaminated by microstructure noise, is enriched with drift and volatility explosions. Our theory advocates a novel decomposition for realized variance into a drift and a volatility component, which leads to significant improvements in volatility forecasting.

Estimation and Inference for Vector Time Series Models with Time-Varying Parameters

Jiti Gao
Monash University
jiti.gao@monash.edu

Abstract: Vector time series models are widely used in practical studies, e.g., forecasting, modelling policy transmission mechanism, and measuring connection of economic agents. To better capture the dynamics, I will introduce a new class of time-varying vector time series models in which the coefficients and covariance matrices of the error innovations are allowed to change smoothly over time. I will focus on time-varying vector auto-regressive (TV-VAR) models. Discussion will first be on a summary of time-varying impulse response analyses, an information criterion to select the optimal lag, and a Wald-type test to determine the constant coefficients. Simulation studies will then be given, and real data examples will also be used to demonstrate the empirical relevance and usefulness of the proposed models and methods. If time permits, I will also briefly discuss time-varying vector error correction models (TV-VECMs).

The Effect of the North Atlantic Oscillation on Monthly Precipitation in Selected European Locations: A Nonlinear Time Series Approach

Timo Teräsvirta
Aarhus University
tterasvirta@econ.au.dk

Abstract: In this presentation, the relationship between the monthly precipitation in 30 European cities and towns, and two Algerian ones, and the North Atlantic Oscillation (NAO) index is characterised 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 considered before moving on to discussing data and the empirical results. The results, based on monthly time series from 1851 up until 2022, include shifting monthly means for the rainfall series and, even more importantly, the estimated coefficients of the exogenous NAO variable.
The empirical results suggest that in the north and the west, the amount of rain in the boreal winter months has increased or stayed the same during the observation period, whereas in the Mediterranean area there have been some decreases. The results on the NAO agree with previous ones in the literature in that the NAO has its strongest effect on precipitation during the winter months. The (negative) effect is particularly strong in Western Europe, Lisbon and the Mediterranean rim. The effect in northern locations is positive for the winter months and as such opposite to the corresponding effect in the west and the Mediterranean. There is plenty of individual variation, however.
The model also contains a time-varying error covariance matrix that is decomposed to time-varying variances and correlations. Constancy of the error variances is tested and the results reported.

Contributed Talks

Estimating Covariance Functions for High-Dimensional Functional Time Series with Dual Factor Structures

Hanlin Shang
Macquarie University
hanlin.shang@mq.edu.au

Abstract: We consider high-dimensional stationary functional time series with dual functional factor model structures. A high-dimensional fully functional factor model is imposed on the observed functional processes whereas a low-dimensional one is assumed for the latent functional factors via the series approximation. Our primary interest is to estimate the large matrix of covariance functions which satisfies the so-called functional "low-rank plus sparse" structure. We extend the classic principal component analysis technique to functional time series and estimate the functional factor loadings, common factors and covariance of functional common components. A functional generalised shrinkage is subsequently applied to the estimated idiosyncratic covariance functions. Under some regularity conditions, we derive the large sample theory of the developed estimates, including the consistency of the estimated factors and functional factor loadings and the uniform convergence rates of the estimated covariance functions. Both the simulation and empirical studies are provided to demonstrate reliable finite-sample performance of the developed model and estimation methodology.

Optimal Linear Prediction with Functional Observations: Why You Can Use a Simple Post-Dimension Reduction Estimator

Won-Ki Seo
University of Sydney
won-ki.seo@sydney.edu.au

Abstract: We study the optimal linear prediction of a random function, assuming it takes values in an infinite dimensional Hilbert space. We begin by characterizing the mean square prediction error (MSPE) associated with a linear predictor and discussing the minimal achievable MSPE. This analysis reveals that, in general, there are multiple non-unique linear predictors that minimize the MSPE, and even if a unique solution exists, consistently estimating it from finite samples is generally impossible. Nevertheless, we can define asymptotically optimal linear operators whose empirical MSPEs approach the minimal achievable level as the sample size increases. Interestingly, we show that standard post-dimension reduction estimators, which have been widely used in the literature, attain such asymptotic optimality under minimal conditions.

Optimal Forecasting under Parameter Instability

Yu Bai
Monash University
yu.bai1@monash.edu

Abstract: This paper considers the problem of using local estimator in a forecasting model which is affected by parameter instability. We first show that local estimator is consistent under various types of parameter instability. Then, we analysis the choices of weighting function and tuning parameter associated with the local estimator. We prove the asymptotic optimality of the tuning parameter selection procedure and provide analytical criterion on the choice of weighting function. The theoretical results are examined through an extensive Monte Carlo study and three empirical applications on forecasting inflation, bond returns and real house prices changes.

Forecast and Optimise

Mahdi Abolghasemi
University of Queensland
m.abolghasemi@uq.edu.au

Abstract: Forecasting plays a pivotal role in managerial decisions, particularly when faced with uncertainty. Forecasts are often used as an input in decision making models. This paradigm of research, termed "predict and optimise," uses forecasts as inputs to derive optimal decisions and have found numerous applications in real world, e.g., forecasting sales and determining optimal inventory. There are studies that suggest more accurate forecasts do not necessarily lead to better decisions. While literature shows there may be a correlation between them, it is not clear what their correlation is and how we can improve upon their performance. I will walk you through problems in which we have used both forecasting and optimisation for decision making. I will further delve into recent advances in merging forecasting with optimisation via machine learning.

Modelling and Forecasting Intraday Spot Volatility

Adam Clements
Queensland University of Technology
a.clements@qut.edu.au

Abstract: This paper proposes a HAR-style regression method for modelling and forecasting intraday spot volatility. Under this approach, each intraday interval is treated as an individual time-series, as opposed to most methods that treat the data as one continuous sample. The empirical analysis is based on a sample of large U.S. individual equities, with the OK candlestick (OHLC) volatility estimator of Li, Wang and Zhang (2021) used as an observed measure of intraday high-frequency spot volatility. It is shown that the proposed framework successfully captures the complex dynamics in intraday volatility. Improvements in forecasts relative to competing regression methods and popular tree-based machine learning forecasts are also observed. Improvements in forecast performance carry over to logarithmic models and forecasts over multiple horizons. Overall, the results indicate that the parameter flexibility in the proposed multiple-regression framework is beneficial. This flexibility comes without undue computational burden.

Asymmetric Uncertainty: Nowcasting using Skewness in Real-Time Data

Paul Labonne
BI Norwegian Business School
paul.labonne@bi.no

Abstract: This paper presents a new way to account for downside and upside risks when producing density nowcasts of GDP growth. The approach relies on modelling location, scale and shape common factors in real-time macroeconomic data. While movements in the location generate shifts in the central part of the predictive density, the scale controls its dispersion (akin to general uncertainty) and the shape its asymmetry, or skewness (akin to downside and upside risks). The empirical application is centred on US GDP growth and the real-time data come from Fred-MD. The results show that there is more to real-time data than their levels or means: their dispersion and asymmetry provide valuable information for nowcasting economic activity. Scale and shape common factors (i) yield more reliable measures of uncertainty and (ii) improve precision when macroeconomic uncertainty is at its peak.

The Relation between Treasury Yields and the Corporate Bond-Yield Spread – Stable or Time-Varying?

Pär Österholm
Örebro University
par.osterholm@sydney.edu.au

Abstract: In this paper we assess whether the relation between US Treasury yields and the corporate bond yield spread has been stable over time. This is done by estimating hybrid time-varying parameter Bayesian VAR models with stochastic volatility where we – as a methodological contribution – also allow for disturbances with fat tails. We analyse monthly data from April 1953 to February 2023 both within- and out-of-sample. Our results indicate that the relation has not been stable; more specifically, there is evidence that the equation of the corporate bond yield spread is subject to time-variation in its parameters. We also find that an increase in the corporate bond yield spread decreases the risk free rate. Finally, we note that while allowing for fat tails receives a fair amount of support within sample, it appears to be of more limited importance from a forecasting perspective.

Predictions in Shared Markets: A Global Forecasting Approach with Deep Learning and Spillover Considerations

Klaus Ackermann
Monash University
klaus.ackermann@monash.edu

Abstract: We introduce a novel forecasting method using global deep learning models to estimate the causal effects of an intervention on multiple treated units, considering the counterfactual and synthetic control approach for policy evaluation. The method is designed for applications in shared markets, where potential spillover effects would otherwise bias the results and rely on time series properties for identification. We recast the causal effect estimation problem as a counterfactual prediction outcome of the treated units without policy intervention. In the first stage, a counterfactual outcome is estimated based on a high-dimensional collection of treated and untreated related time series based on time series forecasting. The time series forecasting takes advantage of the series cross-correlation, allowing us to estimate an autoregressive recurrent neural network that employs parameter sharing across all series. In the second stage, the average treatment effect on the dependent variable of interest is estimated and tested to verify its statistical significance. The method functioning is illustrated by simulation and empirical studies. In the empirical example, we propose estimating the marketing effect of offering deals on overall retail sales in the US grocery stores market. We contribute to the literature using only pre-treatment data for the effect estimation when the no-interference assumption of control units is violated.

Particle-based Variational Bayes: Towards Scalable and Accurate Bayesian Computation

Minh-Ngoc Tran
The University of Sydney
minh-ngoc.tran@sydney.edu.au

Abstract: Variational Bayes (VB) is widely recognised as a highly efficient and scalable technique for Bayesian inference. However, classical VB imposes restrictions on the space of variational distributions, typically restricting it to a specific set of parametric distributions or factorized distributions. This talk explores ways to relax these restrictions by traversing a set of particles to approximate the target distribution. The theoretical basis of the new particle VB method is established using Optimal Transport theory, which equips the space of probability measures with useful calculus tools. This paves the way for new research avenue, enabling precise Bayesian inference even in intricate, high-dimensional scenarios.

Bayesian Inference in SVARs Using 'Soft' Sign Restrictions

Matthew Read
Reserve Bank of Australia
ReadM@rba.gov.au

Abstract: We propose algorithms for conducting Bayesian inference in structural vector autoregressions that are identified using sign restrictions. The key feature of our approach is a sampling step based on 'soft' sign restrictions. This step draws from a target density that smoothly penalises parameter values violating the restrictions, facilitating the use of computationally efficient Markov Chain Monte Carlo sampling algorithms. An importance-sampling step yields draws from the desired distribution conditional on the 'hard' sign restrictions. Relative to standard accept-reject sampling, our approach substantially improves computational efficiency when identification is 'tight'. It also greatly reduces the computational burden of 'prior-robust' Bayesian methods. We illustrate the broad applicability of our approach in a model of the global oil market identified using a rich set of sign, elasticity and narrative restrictions.

The Role of Data and Priors in Estimating Climate Sensitivity

Andrey Vasnev
The University of Sydney
andrey.vasnev@sydney.edu.au

Abstract: In Bayesian theory, the data together with the prior produce a posterior. We show that it is also possible to follow the opposite route, that is, to use data and posterior information (both of which are observable) to reveal the prior (which is not observable). We then apply the theory to equilibrium climate sensitivity as reported by the Intergovernmental Panel on Climate Change in an attempt to get some insight into the prior beliefs of the IPCC scientists. It appears that the data contain much less information than one might think, due to the presence of correlation. We conclude that the prior in the fifth IPCC report was too low, and in the sixth report too high.

Composite Inference on a Common Equilibrium Climate Sensitivity

Mengheng Li
University of Technology Sydney
Mengheng.Li@uts.edu.au

Abstract: The equilibrium climate sensitivity (ECS) stands as a pivotal parameter in climate science with far-reaching implications for addressing the economic and social costs of climate change. This study estimates ECS as a well-defined physical parameter that is common across energy balance models (EBMs), derived from a CO2 quadrupling experiment conducted by 16 climate models. We interpret EBMs as physics-consistent time series models and introduce a Bayesian composite likelihood approach to simultaneously integrate and estimate all the constituent EBMs. In contrast to existing methods commonly employed by climate scientists, our econometric alternative provides a data-driven ECS estimator with an intuitive probabilistic interpretation. We find an ECS estimate of 2.95K, characterized by a unimodal posterior distribution that facilitates uncertainty quantification. Our approach also yields a 95% error band of [2.73K, 3.34K], consistent with, yet notably tighter than, the range of ECS values previously reported.

Boosting Domain-Specific Models with Shrinkage: An Application in Mortality Forecasting

Anastasios Panagiotelis
The University of Sydney
anastasios.panagiotelis@sydney.edu.au

Abstract: This paper extends the technique of gradient boosting with a focus on using domain-specific models instead of trees. The domain of mortality forecasting is considered as an application. The two novel contributions are to use well-known stochastic mortality models as weak learners in gradient boosting rather than trees, and to include a penalty that shrinks the forecasts of mortality in adjacent age groups and nearby geographical regions closer together. The proposed method demonstrates superior forecasting performance based on US male mortality data from 1969 to 2019. The proposed approach also enables us to interpret and visualize the results. The boosted model with age-based shrinkage yields the most accurate national-level mortality forecast. For state-level forecasts, spatial shrinkage provides further improvement in accuracy in addition to the benefits achieved by age-based shrinkage. This additional improvement can be attributed to data sharing across states with both large and small populations in adjacent regions, as well as states which share common risk factors.

Forecasting Periodic Incidence of Australian Seasonal Influenza after Assessing Persistence with and without Machine Learning

Gnanadarsha Dissanayake
NSW Ministry of Health
gnanadarshad@gmail.com

Abstract: Seasonal influenza is today a health burden on the Australian Public health sector. Accurate forecasting of the disease incidence rate enables health service planners and health economists to allocate requisite resources and funding with confidence. To minimise errors in forecasting the impact of different factors on the available data needs to be assessed. In this research endeavour a significant factor in time series-based forecasting methodology known as persistence is assessed using feasible and viable theoretical statistical testing mechanisms in the form of two established hypothesis tests on serial correlation. It is done by examining serial correlation among observations in the data to select the best fit forecasting time series model and subsequently provide out of sample predictions. A validation and a meta-analysis is done thereafter using machine learning paradigms in assessing the forecasting mechanism.

3-Minute Talk

DeepVol: A Pre-Trained Universal Asset Volatility Model

Chen Liu
University of Sydney
liu.chen@sydney.edu.au

Abstract: This paper introduces DeepVol, a pre-trained deep learning volatility model that is more general than traditional econometric models. DeepVol leverage the power of transfer learning to effectively capture and model the volatility dynamics of all financial assets, including previously unseen ones, using a single universal model. This contrasts to the usual practice in the econometrics literature, which trains a separate model for each asset. The introduction of DeepVol opens up new avenues for volatility modeling in the finance industry, potentially transforming the way volatility is predicted.

Efficient Variational Inference for Large Skew-t Copulas with Application to Intraday Equity Returns

Lin Deng
University of Melbourne
L.Deng@mbs.edu

Abstract: Large skew-t factor copula models are attractive for the modeling of financial data because they allow for asymmetric and extreme tail dependence. We show that the copula implicit in the skew-t distribution of Azzalini and Capitanio (2003) allows for a higher level of pairwise asymmetric dependence than two popular alternative skew-t copulas. Estimation of this copula in high dimensions is challenging, and we propose a fast and accurate Bayesian variational inference (VI) approach to do so. The method uses a conditionally Gaussian generative representation of the skew-t distribution to define an augmented posterior that can be approximated accurately. A fast stochastic gradient ascent algorithm is used to solve the variational optimization. The new methodology is used to estimate copula models for intraday returns from 2017 to 2021 on 93 U.S. equities. The copula captures substantial heterogeneity in asymmetric dependence over equity pairs, in addition to the variability in pairwise correlations. We show that intraday predictive densities from the skew-t copula are more accurate than from some other copula models, while portfolio selection strategies based on the estimated pairwise tail dependencies improve performance relative to the benchmark index.

Hierarchical House Price Model incorporating Geographical and Macroeconomic Factors

Lingfeng Lyu
University of New South Wales
lingfeng.lyu@unsw.edu.au

Abstract: This paper presents a tri-level hierarchical approach to house price modelling at the postcode level, which is considered the most granular geographical scale, incorporating macroeconomic influences from the national level and integrating data from the largest sub-state level (SA4). By employing a Risk Premium - Principal Component Analysis (RP-PCA) for SA4-level risk factors and combining these with national-level risk factors, a vector autoregressive (VAR) model is developed. This geographically conditional multi-factor model with a hierarchical structure offers enhanced short-term prediction accuracy while maintaining long-term forecasting capabilities. The model's predictive accuracy is further enhanced by introducing an empirical copula to describe the dependence structure of one-step residuals across various suburbs. This methodology grants a dynamic and granular view of housing price trends in Australia. Key determinants like interest rate shifts, GDP growth, and exchange rate variances play a crucial role, particularly in urban areas in metropolitan cities. The analysis of economic and demographic factors on the SA4 level indicates that elements such as home debt increments, wage fluctuations, and population shifts are pivotal in shaping housing prices, underscoring the significance of a granular regional analysis.

Multi-Matrix Autoregressive Models with an Application to Multi-Modal Network

Ye Shiqi
Xiamen University and Chinese Academy of Sciences
shiqi.ye.c@gmail.com

Abstract: Matrix time series data have become increasingly prevalent across diverse fields, including economics, finance, computer science, engineering, and signal processing. This study introduces a novel multi-matrix autoregressive (MMAR) model designed to jointly model matrix time series with varying structures. Notably, the well-known matrix-valued autoregressive model and three-order tensor autoregressive model are special cases of the proposed model. We present three distinct estimation methods for the MMAR model, investigate their statistical properties, and provide numerical simulations to corroborate them. Moreover, we integrate the MMAR model with connectedness network analysis to concurrently model the macroeconomic matrix time series of China's 31 provinces and a vector time series comprising the economic policy uncertainty index, trade policy uncertainty index, and China's geopolitical risk. By constructing multi-modal connectedness networks, we delve into the intricate interrelationships between China's regional economy and macroeconomic regulation. The findings of this study provide valuable insights for further research and policy-making in the relevant domains.

Enhancing Financial Forecasting: A Comparative Study of Hybrid Deep Learning Models and Traditional Approaches in Oil Price Prediction

Adel Gadhi
University of Sydney
agad0777@uni.sydney.edu.au

Abstract: This paper presents an in-depth analysis of the predictive capabilities of hybrid deep learning models, specifically focusing on the integration of Wasserstein Generative Adversarial Networks (WGAN) with an enhanced gradient penalty. We compare these advanced hybrid models against traditional deep learning architectures such as GARCH, ANN, LSTM-ANN, and BLSTM-ANN. Utilizing widely-accepted metrics, we evaluate the forecasting accuracy of these models in predicting daily oil prices. Our findings offer valuable insights into the advantages of combining traditional financial forecasting methods with state-of-the-art deep learning techniques, thereby paving the way for more accurate and reliable financial predictions.

Estimation and Inference of Average Treatment Effects using Deep Neural Networks

Taiga Saito
Monash University
taiga.saito@monash.edu

Abstract: We propose a semiparametric ATE estimator using Deep Neural Networks (DNN), allowing for moderately large dimension covariates. We consider the doubly robust estimator with two nuisance functions, the potential outcomes and the propensity score, nonparametrically approximated through a specific type of DNN, and construct the asymptotic normality of the estimator. Due to the convergence theory developed in the recent literature (Schmidt-Hieber, 2020; Farrell et al., 2021), obtaining the consistency of the nuisance functions and the subsequent doubly robust estimator is relatively straightforward. However, these convergence rates are slower than root N and thus insufficient for constructing asymptotic normality of the second-stage ATE estimators when naively plugged in. To overcome this, we introduce an unbalanced sample-splitting technique that allows us to provide a theoretical guarantee for asymptotic normality at the rate of $N^{-1/2}$, enabling inference such as hypothesis testing and the construction of confidence intervals. Our simulation and an empirical application to a job training program in the National Supported Work Demonstration (NSW) support the theory and highlight the practical utilities of the proposed method in finite samples.

Functional Principal Component Logit Regression for the Subarachnoid Hemorrhage Data

Arezoo Orooji
Macquarie University
arezoo.orooji@mq.edu.au

Abstract: In today's world, with the passage of time and advancements in science and technology, data collection over time or location has become easier, enabling more densely sampled observations over time and space. Various statistical models have been designed to analyze such data, including time series regression, mixed-effects models, and marginal models, among others. One of the analyses that can be used for data collected over time is the functional data analysis. In this study, since the intracranial pressure was measured over time in patients with subarachnoid hemorrhage due to aneurysm, and we were also interested in the association between intracranial pressure and mortality, Functional principal component logit regression was employed. This approach serves a dual purpose: first, it reduces dimensionality, enhancing parameter estimation, and second, it effectively addresses issues related to multicollinearity. This method primarily entails converting the functional logistic regression model into a multiple logistic regression model, where we incorporate a reduced set of Principal Components (PCs) as covariates. In regards to the results of functional principal component logit regression, the best goodness of fit was obtained with 4 principal components. the correct classification rate and AUC were 79% and 0.82 respectively. Based on the results we've acquired; this model exhibits a reasonably strong capability to predict the patients' status.

Robust Autoregressive Modelling using a Penalised Approach

Rajan Shankar
University of Sydney
rsha5949@uni.sydney.edu.au

Abstract: Classical methods for estimating the parameters of an autoregressive lag-p (AR(p)) time series process are sensitive to additive outliers in the data. Such outliers are especially bad for time series modelling since each outlier contaminates the following p rows of data. More specifically, this causes the next p residuals to be unreasonably small, which can lead to 'outlier masking'. By representing the AR(p) process as a state-space process, we can apply robust filtering techniques to obtain 'robust residuals'. Then, we can maximise a robust likelihood of these robust residuals to obtain robust parameter estimates. While there exist other approaches to construct the robust likelihood, we propose a novel approach where we add a penalty term of 'shifts' of the robust residuals to the Gaussian likelihood. With this comes the double-edged sword of being able to select the number of outliers via a tuning parameter. We also show simulations to benchmark how the method performs against existing methods and empirically investigate its robustness properties.

List of Participants

Last Name	First Name	Affiliation	Email Address
Abolghasemi	Mahdi	University of Queensland	m.abolghasemi@uq.edu.au
Ackermann	Klaus	Monash University	klaus.ackermann@monash.edu
Bai	Yu	Monash University	yu.bai1@monash.edu
Beare	Brendan	University of Sydney	brendan.beare@sydney.edu.au
Bromell	Josh	Reserve Bank of New Zealand, New Zealand	josh.bromell@rbnz.govt.nz
Bu	Dingliu	University of Sydney	dibu6103@uni.sydney.edu.au
Chen	Mingdi	University of Sydney	mingdi.chen@sydney.edu.au
Choy	Boris	University of Sydney	boris.choy@sydney.edu.au
Clements	Adam	Queensland University of Technology	a.clements@qut.edu.au
Deng	Lin	University of Melbourne	l.deng@mbs.edu
Deutscher	Sherri	Martin Brower	SDeutscher@martinbrower.com
Dissanayake	Gnanadarsha	NSW Ministry of Health & University of Sydney	gnanadarshad@gmail.com
Exterkate	Peter	Susquehanna International Group (SIG)	Peter.exterkate@sig.com
Fairley	Andrew	Martin Brower	afairley@martinbrower.com
Gadhi	Adel	University of Sydney	agad0777@uni.sydney.edu.au
Gao	Jiti	Monash University	jiti.gao@monash.edu
Gerlach	Richard	University of Sydney	richard.gerlach@sydney.edu.au
Gezdur	Arda	University of Sydney	arda.gezdur@sydney.edu.au
Hartigan	Luke	University of Sydney	luke.hartigan@sydney.edu.au
Hong	Yongmiao	University of Chinese Academy of Sciences, China	yh20@cornell.edu
Hunt	Richard	University of Sydney	richard@huntemail.id.au
Hyndman	Rob	Monash University	rob.hyndman@monash.edu
Jayathunga	Naotunnage	University of Sydney	njay0752@uni.sydney.edu.au
Kwok	Simon	University of Sydney	simon.kwok@sydney.edu.au
Labonne	Paul	BI Norwegian Business School, Norway	paul.labonne@bi.no
Lam	Alex	University of Sydney	fong.lam@sydney.edu.au
Laurent	Sébastien	Aix-Marseille Université	sebastien.laurent@iae-aix.com
Li	Haley	Susquehanna International Group (SIG)	Haley.li@sig.com
Li	Winny	University of New South Wales	winny.li@unsw.edu.au
Li	Mengheng	University of Technology Sydney	mengheng.li@uts.edu.au
Liao	Moyu	University of Sydney	moyu.liao@sydney.edu.au
Liu	Chen	University of Sydney	liu.chen@sydney.edu.au
Lu	Ye	University of Sydney	ye.lu1@sydney.edu.au
Lyu	Lingfeng	University of New South Wales	lingfeng.lyu@unsw.edu.au
Madhavani	Fal	Susquehanna International Group (SIG)	Fal.madhavani@sig.com
Nallabolu	Ranjani	Susquehanna International Group (SIG)	Ranjani.nallabolu@sig.com
Orooji	Arezoo	Macquarie University	arezoo.orooji@mq.edu.au
Österholm	Pär	Örebro University, Sweden	par.osterholm@oru.se
Panagiotelis	Anastasios	University of Sydney	anastasios.panagiotelis@sydney.edu.au
Peiris	Shelton	University of Sydney	shelton.peiris@sydney.edu.au
Perrens	Sarah	Susquehanna International Group (SIG)	Sarah.perrens@sig.com

Phillips	Peter	Yale University, USA	peter.phillips@yale.edu
Read	Matthew	Reserve Bank of Australia	readm@rba.gov.au
Saito	Taiga	Monash University	taiga.saito@monash.edu
Seo	Dennis	Susquehanna International Group (SIG)	Dennis.seo@sig.com
Seo	Won-Ki	University of Sydney	won-ki.seo@sydney.edu.au
Shang	Hanlin	Macquarie University	hanlin.shang@mq.edu.au
Shang	Tony	University of Sydney	ce.shang@sydney.edu.au
Shankar	Rajan	University of Sydney	rsha5949@uni.sydney.edu.au
Shi	Shuping	Macquarie University	shuping.shi@mq.edu.au
Sijp	Williem	Neoval - Ray White Group	will@neoval.io
Siu	Ken	Macquarie University	ken.siu@mq.edu.au
Teräsvirta	Timo	Aarhus University, Denmark	tterasvirta@econ.au.dk
Tran	Minh-Ngoc	University of Sydney	minh-ngoc.tran@sydney.edu.au
Vasnev	Andrey	University of Sydney	andrey.vasnev@sydney.edu.au
Wang	Chao	University of Sydney	chao.wang@sydney.edu.au
Wichitaksorn	Nuttanan	Auckland University of Technology, New Zealand	nuttanan.wichitaksorn@aut.ac.nz
Yang	Yang	University of Newcastle	yang.yang10@newcastle.edu.au
Yao	Juan	University of Sydney	Juan.yao@sydney.edu.au
Ye	Shiqi	Xiamen University, China	shiqi.ye.c@gmail.com
Zhang	Weiben	University of Melbourne	w.zhang@mbs.edu

Acknowledgements

Time Series and Forecasting Research Group, The University of Sydney Business School

Discipline of Business Analytics, The University of Sydney Business School

School of Economics, Faculty of Arts and Social Sciences

Exclusive Industrial Sponsor

We are delighted to have the support of Susquehanna International Group (SIG) as the exclusive industrial sponsor of the 2023 Time Series and Forecasting Symposium. SIG is a global quantitative trading firm founded with a growth mindset and an analytical approach to decision making.

As one of the largest proprietary trading firms in the world, SIG benefits the financial markets by providing liquidity and ensuring competitive prices for buyers and sellers. SIG brings together the brightest minds and the best technology, to design and implement quantitative trading strategies that make them leaders in the financial markets. SIG will be providing the awards for the best poster and the best student presentation, and SIG representatives will be available for 1-2-1 chats at their conference booth to discuss postdoc opportunities in quantitative research at their Sydney headquarters in Barangaroo.