THE UNIVERSITY OF SYDNEY
Business School

# Time Series and Forecasting Symposium TSF2025


27-28 November 2025

**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.*

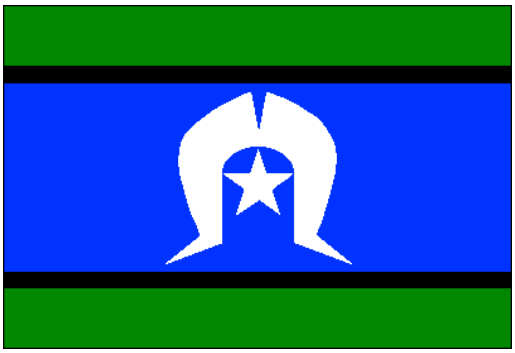

# Location Guide

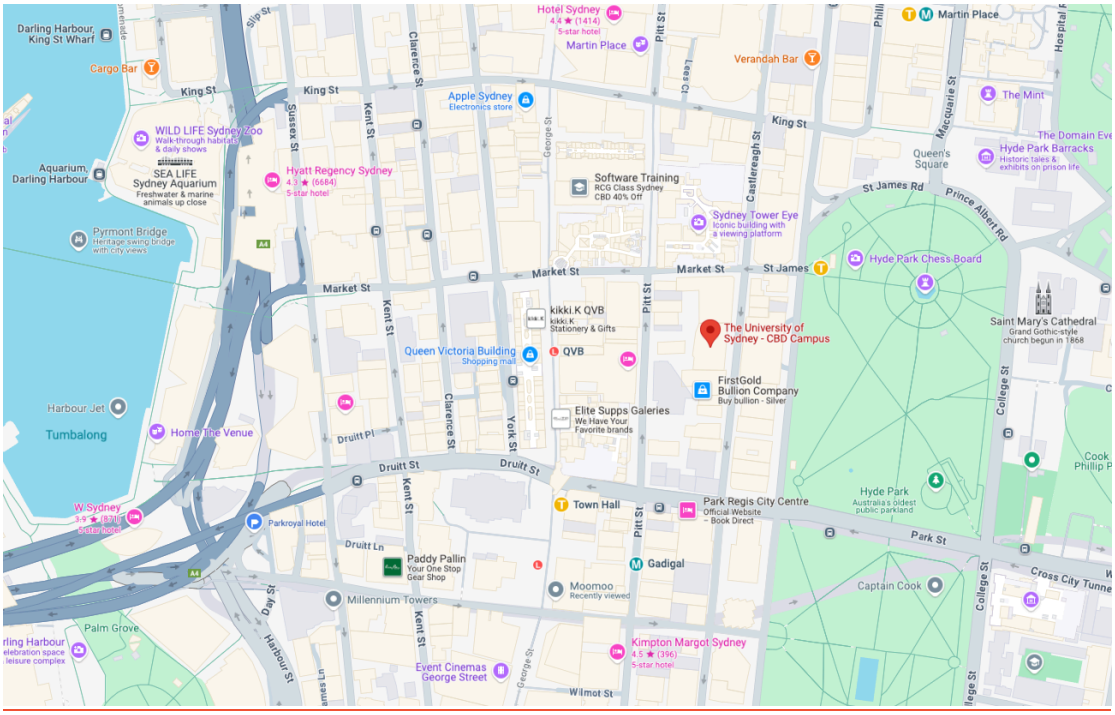


## The University of Sydney, CBD Campus

**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 metro**
The nearest metro stations are Gadigal Station and Martin Place 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, University of Sydney

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

**Jackie Li**, Department of Econometrics and Business Statistics, Monash University

**Xuan (Winny) Li**, School of Computer Science and Engineering, University of New South Wales

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

**Jae Kyung Woo**, School of Risk and Actuarial Studies, University of New South Wales

**Kathy Zhou**, Discipline of Business Analytics, University of Sydney

## Registration

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

## Presentation Format

Keynote talks are 50 minutes each including Q&A
Invited talks are 40 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

Thursday, 27 November 2025, 6:00pm
Venue: The Palace Chinese Restaurant
(Shop 38, Level 1 Piccadilly Tower, 133/145 Castlereagh St, 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: Thursday 27 November 2025

**9:00 – 10:00** **Morning Tea and Registration**

| | |
|---|---|
| 10:00 – 10:10 | **Welcome**<br><br>**Boris Choy**<br>*Co-Chair, TSF2025 Organising Committee, The University of Sydney Business School* |
| **Session 1**<br>*Session Chair:* **Boris Choy** | |
| 10:10 – 11:00 | **Keynote Talk**<br><br>**Hoi Ying Wong**, *Chinese University of Hong Kong*<br>Deep learning with Probability and Statistics in Finance: Pairs trading and optimal stopping |
| 11:00 – 12:00 | **Contributed Talks – forecasting**<br><br>**Hanlin Shang**, *Macquarie University*<br>Temporal hierarchical functional time series forecasting: An application to particulate number size distribution<br><br>**Nuttanan Wichitaksorn**, *Auckland University of Technology*<br>TSGBiNet: An End-to-End Spatio-Temporal Graph Neural Network for Multi-Site PV Power Forecasting with Missing Values |

**12:00 – 13:20** **Lunch**

| | |
|---|---|
| **Session 2**<br>*Session Chair:* **Hanlin Shang** | |
| 13:20 – 14:00 | **Invited Talk**<br><br>**Junbin Gao**, *The University of Sydney*<br>Operator-based Implied Volatility Smoothing for Options Trading |
| 14:00 – 15:00 | **Contributed Talks – financial time series**<br><br>**Ken Siu**, *Macquarie University*<br>Bayesian Lower and Upper Estimates for Ether Option Prices with Conditional Heteroscedasticity and Model Uncertainty<br><br>**Simon Kwok**, *The University of Sydney*<br>Riding a Bubble: A Study of Market-Timing Trading Strategies |

**15:00 – 15:30** **Afternoon Tea**

| Session 3<br>*Session Chair:* **Ken Siu** | |
|---|---|
| 15:30 – 17:00 | **Contributed Talks - macroeconomics**<br><br>**Martin McCarthy**, *Reserve Bank of Australia*<br>Forecasts of Period-Average Exchange Rates: Insights from Real-Time Daily Data<br><br>**Pär Österholm**, *NSW Treasury and Örebro University*<br>On the Stability of Macroeconomic Relationships in Australia<br><br>**Williem Sijp**, *Neoval - Ray White and UTS*<br>Market Sensitivities and Growth Differentials Across Australian Housing Markets |

**18:00 – 20:30** **Symposium Dinner at The Palace Chinese Restaurant**
(Shop 38 Level 1 Piccadilly Tower, 133/145 Castlereagh St, Sydney)

# Day 2: Friday 28 November 2025

**9:00 – 10:10** **Morning Tea and Poster Display**

| **Session 4**<br>*Session Chair:* **Richard Gerlach** | |
|---|---|
| 10:10 – 11:00 | **Keynote Talk**<br><br>**Brendan Beare**, *The University of Sydney*<br>The natural components of a regular linear system |
| 11:00 – 12:00 | **Contributed Talks – actuarial applications**<br><br>**Eric Cheung**, *University of New South Wales*<br>Modelling discrete common-shock risks through matrix distributions<br><br>**Jae Kyung Woo**, *University of New South Wales*<br>Cointegration analysis of crop yields and extreme weather factors using Actuaries Climate Index with application of bonus-malus system |

**12:00 – 13:20** **Lunch**

| **Session 5**<br>*Session Chair:* **Simon Kwok** | |
|---|---|
| 13:20 – 14:00 | **Invited Talk**<br><br>**Tonghui Wang**, *New Mexico State University*<br>Family of Skew Normal Distributions --- Recent Developments and Applications in Econometrics, Business, Agriculture, and Healthcare etc. |
| 14:00 – 14:30 | **Contributed Talk – stochastic process control**<br><br>**Dan Wang**, *La Trobe University, Northwest University (China)*<br>A New Self-Starting EWMA Chart for Monitoring Multivariate Serially Correlated Processes |

| **Session 6**<br>*Session Chair:* **Winny Li** | |
|---|---|
| 14:30 – 15:30 | **3-Minute Talks**<br><br>**Zhengyang Chi**, *The University of Sydney (Business School)*<br>Graph Signal Processing for Global Stock Market Realized Volatility Forecasting<br><br>**Daniel Fynn**, *The University of Wollongong*<br>Neural Inference Function for Margins for Time Series Copula Models |

| | |
|---|---|
| | **Leonard Mushunje**, *The University of Sydney (Maths & Stat)*<br>Graphical Causal Models for Economic Networks: A Machine Learning and Econometric Approach<br><br>**Jinbo Cai**, *University of New South Wales*<br>(Re-)Imag(in)ing and Graphing Correlation Risks<br><br>**Yuji Deng**, *The University of Sydney (School of Economics)*<br>Volatility Forecasting of CSI 300 ETF: A Comparison between HAR-Type Models and Neural Network Approaches<br><br>**Zhixiang (Elvis) Yang**, *Monash University*<br>Measuring Climate Change Concern in the Media: Economic and Financial Impacts in Australia |

**15:30 – 16:00** **Afternoon Tea**

| | |
|---|---|
| 16:00 – 16:30 | **Award Presentation**<br><br>**Keynote and Invited Speakers**<br><br>**Closing Ceremony**<br><br>**Boris Choy**<br>*Co-Chair, TSF2025 Organising Committee, The University of Sydney*<br><br>**Simon Kwok**<br>*Co-Chair, TSF2025 Organising Committee, The University of Sydney*<br><br>**Richard Gerlach**<br>*Co-Lead, Time Series and Forecasting Research Group, The University of Sydney* |

# List of Abstracts

## Keynote Talks

Deep learning with Probability and Statistics in Finance: Pairs trading and optimal stopping

**Hoi Ying Wong**
*Chinese University of Hong Kong*
hywong@cuhk.edu.hk

**Abstract:** Probability and statistics are the building blocks of quantitative finance, even in the rapidly growing era of artificial intelligence. In this presentation, two examples illustrate how deep learning can be integrated with probability and statistical theories for applications in finance.
The first example focuses on pairs trading based on a nonlinear cointegration relationship. As the profitability of traditional linear cointegrated pairs trading on stocks has diminished, a new proposal explores nonlinear cointegration relationships and forms a stock and options portfolio to generate statistical arbitrage profits. Using high-frequency data, a nonlinear cointegration relationship is approximated by a deep neural network (DNN) trained to maximize a cointegration test statistic. The DNN is then calibrated using available option prices, resulting in a strategy that comprises a portfolio of stocks and options. An empirical example based on double-listed stocks in Hong Kong and Mainland China is provided to illustrate this strategy.
The second example pertains to American-style options on multiple assets. Using the martingale representation theorem, we construct DeepMartingale to compute the upper bound of high-dimensional American option prices. We further demonstrate that the architectural design of DeepMartingale can be adjusted to overcome the curse of dimensionality. Specifically, we clarify that the size of the DNN is related to a power of the problem's dimension. Numerical examples of high-dimensional American option prices show the superiority of DeepMartingale.

The natural components of a regular linear system

**Brendan Beare**
*The University of Sydney*
brendan.beare@sydney.edu.au

**Abstract:** The analysis of a finite-dimensional regular linear system may be simplified by separating the system into its natural components. The natural components are smaller linear systems on separate subspaces whose dimensions sum to the dimension of the original linear system. Each natural component has a single eigenvalue and may be studied in isolation. By reassembling the natural components we obtain a general characterization of the set of all solutions to the original linear system which contains the core content of the Granger-Johansen theorem.

## Invited Talks

Family of Skew Normal Distributions --- Recent Developments and Applications in Econometrics, Business, Agriculture, and Healthcare etc.

**Tonghui Wang**
*New Mexico State University*
twang@nmsu.edu

**Abstract:** In this talk, the family of skew normal distributions is introduced, which is an extension of normal family. Properties, such as moment generating function, linear functions, skew chi-square, and skew F distribution are discussed. Real data applications in various fields, such as econometrics, business, agriculture, and healthcare are studied. Current research developments are provided including (i) The a Priori Procedure (APP) for finding minimum sample sizes in parameter estimations, (ii) Gain Probability (GP) Analysis, and (iii) Stochastic frontier model analysis.

Operator-based Implied Volatility Smoothing for Options Trading

**Junbin Gao**
*The University of Sydney*
Junbin.gao@sydney.edu.au

**Abstract:** Options trading volume has grown rapidly in recent years, particularly for index options on the S&P 500. With this surge, market stability and accurate volatility estimation have become more important than ever. Accurate construction of the implied volatility surface (IVS) is essential for derivative pricing and financial risk management. While traditional parametric models such as SVI and SSVI offer arbitrage-free structures, they require frequent recalibration and often perform poorly in the presence of sparse option data. Recent advances in neural operator learning, particularly Graph Neural Operators (GNO), provide a promising data-driven alternative by enabling flexible interpolation without restrictive structural assumptions. However, GNOs involve high-dimensional intermediate representations, incur substantial computational costs, and require large volumes of high-frequency training data, which are often difficult to obtain in practice. This research proposes enhancing the GNO framework by integrating B-spline interpolation into the input representation, significantly improving computational efficiency. By enhancing structural information prior to learning, the method is expected to improve robustness and accuracy in constructing smooth, arbitrage-consistent volatility surfaces. The proposed framework will be evaluated using index option data, contributing to the development of more data-efficient, operator-based approaches in financial modelling.

## Contributed Talks

Temporal hierarchical functional time series forecasting: An application to particulate number size distribution

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

**Abstract:** Particulate matters are often measured at different time scales, such as hour, day, week, or fortnight. Forecasting particulate number size distribution (PNSD) at different scales plays an important role in developing environmental policy from hourly to fortnightly scale. However, independent forecasts at the hourly scale may not add up to the forecasts at the fortnight scale. To address this issue, we consider reconciling forecasts of PNSD, extending the methods of temporal hierarchical aggregation to functional time series, where PNSD is treated as a continuum and measured densely and irregularly spaced at 51 different sizes. The temporal hierarchical functional time series forecasting methods are proposed to produce point forecasts of the PNSD that are aggregated appropriately across different time scales. For evaluating forecast uncertainty, we propose a model-agnostic and distribution-free method for reconciling interval forecasts. Using the PNSD data from 3/January/2011 to 28/October/2018 in London, we investigate and compare the forecast accuracy between the independent and reconciled forecasts. The proposed methods are shown to be useful for reconciling forecasts of the PNSD at various time scales. Because of the temporal hierarchy, one can achieve improved forecast accuracy averaged over different time scales.

TSGBiNet: An End-to-End Spatio-Temporal Graph Neural Network for Multi-Site PV Power Forecasting with Missing Values

**Nuttanan Wichitaksorn**
*Auckland University of Technology*
nuttanan.wichitaksorn@aut.ac.nz

**Abstract:** Accurate spatio-temporal forecasting of photovoltaic (PV) power is critical for the effective management of multi-site PV systems. However, existing approaches often exhibit limitations in accurately modeling spatial dependencies in the absence of explicit graph structures, and typically presume fully observed datasets, an assumption that is frequently violated in practical applications due to sensor malfunctions or the high cost of data acquisition. To address these challenges, we propose a novel end-to-end and scalable framework for multi-site PV power forecasting, capable of effectively handling spatio-temporal data incompleteness. The proposed model, termed TSGBiNet, adopts a hierarchical architecture integrating the Multi-Scale Missing-Aware Temporal Attention

Network (MS-MTAN) module and the Graph Bias-injected Network (GBiN) module. In particular, MS-MTAN is designed to capture multi-scale temporal dependencies that are often disrupted by missing values within each individual PV site. Meanwhile, GBiN learns hidden spatial correlations across sites while explicitly modelling missing data patterns through an adaptive bias-injection mechanism. Ablation studies validate the framework's effectiveness in uncovering intricate spatio-temporal relationships, even under severe data incompleteness. Extensive experiments on real-world datasets demonstrate that TSGBiNet consistently improves forecasting accuracy and robustness.

## Bayesian Lower and Upper Estimates for Ether Option Prices with Conditional Heteroscedasticity and Model Uncertainty

**Ken Siu**
*Macquarie University*
ken.siu@mq.edu.au

**Abstract:** In this talk, we shall discuss the use of Bayesian nonlinear expectations to construct lower and upper estimates for prices of Ether options with conditional heteroscedasticity and model uncertainty. Specifically, a GARCH model is used to incorporate conditional heteroscedasticity in the logarithmic returns of Ethereum. Bayesian nonlinear expectations are adopted to introduce uncertainty about the conditional mean and volatility of the logarithmic returns of Ethereum. Extended Girsanov's principle is adopted to change probability measures and introduce a family of alternative GARCH models. The Bayesian credible intervals for "uncertain" drift and volatility parameters obtained from conjugate priors and residuals of the estimated GARCH model are used to construct Bayesian superlinear and sublinear expectations giving the Bayesian lower and upper estimates for the price of an Ether option, respectively. Empirical and simulation studies are provided using real data on Ethereum in AUD.

This talk is based on the following paper: Siu, Tak Kuen. 2024. Bayesian Lower and Upper Estimates for Ether Option Prices with Conditional Heteroscedasticity and Model Uncertainty; Journal of Risk and Financial Management 17, 10: 436. https://doi.org/10.3390/jrfm17100436.

## Riding a Bubble: A Study of Market-Timing Trading Strategies

**Simon Kwok**
*The University of Sydney*
simon.kwok@sydney.edu.au

**Abstract:** This paper characterizes the probability distribution of local martingale price bubble processes in order to construct profitable trading strategies which exploit price bubble dynamics. Various such market-timing trading strategies are identified; all are based on the idea of "riding a bubble" by holding a stock during a bubbly period until the time when the bubble's magnitude hits a predetermined barrier. We study these market-timing trading strategies via both simulations and back-tests using historical market prices. We show both in theory and practice that implementing such strategies in the presence of a bubble leads to increased performance relative to the traditional buy-and-hold trading strategy.

## Forecasts of Period-Average Exchange Rates: Insights from Real-Time Daily Data

**Martin McCarthy**
*Reserve Bank of Australia*
mccarthym@rba.gov.au

**Abstract:** Forecasting period-average exchange rates requires high-frequency data to efficiently construct forecasts and test their accuracy against the traditional random walk hypothesis. We construct the first real-time dataset of daily effective exchange rates for all available countries, nominal and real. Forecasts based on daily data improve accuracy by up to 40 percent relative to monthly averages. Unlike bilateral exchange rates, daily effective exchange rates exhibit properties distinct from random walks. Using efficient estimation and testing methods made possible by the data, we find new evidence of real-time predictability for effective exchange rates in up to fifty percent of countries.

On the Stability of Macroeconomic Relationships in Australia

**Pär Österholm**
*NSW Treasury & Örebro University*
par.osterholm@sydney.edu.au

**Abstract:** In this paper, we analyse whether two key macroeconomic relationships in Australia – Okun's law and the Phillips curve – have been stable over time. This is done by estimating hybrid time-varying parameter Bayesian VAR models with stochastic volatility using quarterly data from 1978 to 2024. Model comparison based on marginal likelihoods indicates that Okun's law has been stable, whereas the Phillips curve has not. However, our results do not lend any support to a flattening of the Australian Phillips curve – an otherwise common finding in the international literature. If anything, we find evidence of the contrary. Using the preferred specification of the BVAR for the unemployment rate and inflation, we also calculate trend values for both variables. The model's trend unemployment rate at the end of the sample is approximately five percent; estimated trend inflation at the same point in time is close to the Reserve Bank of Australia's inflation target.

Market Sensitivities and Growth Differentials Across Australian Housing Markets

**Williem Sijp**
*Neoval - Ray White and UTS*
wim.sijp@gmail.com

**Abstract:** Australian house prices show persistent cross-city growth differences: major capitals rose about 5.5x - 10.6x over 1995 - 2024. Raw growth conflates long-run sensitivity to the national market with mean-reverting cycles. I decompose each regional index into (i) a national Market factor, (ii) two stationary spreads: the Mining spread (Perth minus Sydney), capturing resource-cycle swings, and the Lifestyle spread, tracking amenity-driven coastal migration, and (iii) a city-specific stationary remainder. The market loading ("beta") isolates fundamental sensitivity: regional growth scales with national growth approximately as "national growth raised to beta." Expanding-window ARIMAX models estimated in levels show beta is stable despite major shocks (GFC, mining boom, COVID-19). Melbourne exhibits the highest amplification (beta ≈ 1.22; a national doubling maps to ~2.33x), Sydney tracks the market (beta ≈ 1.01), and many regional areas dampen (beta ≈ 0.70 - 0.88). Non-market volatility in Sydney and Perth is dominated by the Mining spread (opposite signs), while Melbourne's is more tied to the Lifestyle spread. The spreads widen 10-year forecast bands without overturning the beta rankings. Methodologically, I use observable index differentials (not latent extraction), estimate ARIMAX(p,0,q) with a seasonal MA(1) in levels to preserve long-run co-movement, and validate stability across expanding windows.

Modelling discrete common-shock risks through matrix distributions

**Eric Cheung**
*University of New South Wales*
eric.cheung@unsw.edu.au

**Abstract:** In this presentation, we introduce a novel class of bivariate common-shock discrete phase-type (CDPH) distributions to describe dependencies in loss modelling, with an emphasis on those induced by common shocks. By constructing two jointly evolving terminating Markov chains that share a common evolution up to a random time corresponding to the common shock component, and then proceed independently, we capture the essential features of risk events influenced by shared and individual-specific factors. We provide explicit expressions for the joint distribution of the termination times and demonstrate various class and distributional properties such as closure under mixtures and sums. We also consider random sums where aggregate claims are sums of continuous phase-type random variables with counts determined by these termination times, and show that their joint distribution belongs to the multivariate phase-type or matrix-exponential class. The applicability of our models is illustrated through simulation studies and an application to bivariate insurance claim frequency data. In particular, via an EM fitting algorithm we are able to estimate the latent common shock component present in a bivariate distribution or dataset. This is joint work with Martin Bladt, Oscar Peralta and Jae-Kyung Woo.

Cointegration analysis of crop yields and extreme weather factors using Actuaries Climate Index with application of bonus-malus system

**Jae-Kyung Woo**
*University of New South Wales*
j.k.woo@unsw.edu.au

**Abstract:** The Actuaries Climate Index results from joint research funded by the actuarial professions of the US and Canada. It serves as a valuable source of the frequency of extreme weather events and the extent of sea level change in the US and Canada. Using extreme weather data from the ACI and an Error Correction Model, we analyze the long-term comovement of these variables with crop yields. Our analysis suggests that significant weather variables can serve as trigger parameters in the pricing framework of weather-index crop insurance. To address the challenges of weather-index crop insurance while maintaining the advantages of a bonus-malus system (BMS), we propose a transition rule that differentiates between damages caused by severe weather and those arising from policyholder decisions. We further explore the challenges of implementing such a hybrid BMS for crop insurance, where extreme weather outcomes are incorporated into the classical BMS framework. This is joint work with Cheung, E.C.K., Ip, R.H.L., and Tam, H.O.

A New Self-Starting EWMA Chart for Monitoring Multivariate Serially Correlated Processes

**Dan Wang**
*La Trobe University, Northwest University (China)*
wangdan@nwu.edu.cn

**Abstract:** Statistical process control (SPC) charts are applied to various sequential processes to detect shifts in their distributions. Traditional SPC control charts are studied under the assumption that the process observations are independent and identically distributed. Furthermore, it is typically assumed that the observations obtained under the in-control (IC) process follow a parametric distribution. In practice, it is difficult to satisfy these assumptions. If the assumptions are invalid, the monitoring performance of traditional SPC charts will be inaccurate. In this paper, we present a new multivariate chart based on data decorrelation, combined with Gaussian kernel density estimation (KDE). In contrast to recent nonparametric control charts that categorize or rank the observed data, our new method only uses the observed data to fit the underlying distribution. Applying Gaussian KDE and data transformation of the process observations, after data decorrelation, a process that is approximately standard normal is obtained. The control chart is also self-starting and all relevant IC quantities can be recursively updated. Numerical studies and real data examples are presented to demonstrate the reliability and effectiveness of this method for online monitoring in different situations.

**Poster and 3-Minute Presentation**

Graph Signal Processing for Global Stock Market Realized Volatility Forecasting

**Zhengyang Chi**
*The University of Sydney*
zhengyang.chi@sydney.edu.au

**Abstract:** This paper introduces an innovative realized volatility (RV) forecasting framework that extends the conventional Heterogeneous autoregressive (HAR) model via integrating Graph Signal Processing (GSP). The study first evaluates various constructions of volatility-interrelationship networks by analysing how the associated graph signal energy tracks global financial market volatility. Volatility spillovers are subsequently embedded into the proposed framework, which employs the graph Fourier transform (GFT) and its inverse to effectively capture global stock market dynamics in both the spectral and spatial domains. The framework not only provides a global context for modelling the volatility interrelationships, but also captures the nonlinearity and directionality of the volatility spillover effect. The empirical study using RV data of 24 global stock market indices compares short-, mid- and long-term RV forecasts with various HAR-type benchmarks and a graph neural network-based HAR model. The proposed model consistently outperforms all comparators, demonstrating the effectiveness of integrating GSP into the HAR model for RV forecasting.

Neural Inference Function for Margins for Time Series Copula Models

**Daniel Fynn**
*The University of Wollongong*
df390@uowmail.edu.au

**Abstract**: opula models are widely used for multivariate time series data because they allow the modelling of marginal distributions separately from the dependence structure, which is captured by the copula function. Implicit copulas are the most commonly used approach for modelling dependence in high-dimensional settings. This broad class includes various types of copulas, such as elliptical copulas and skew t-copulas. Estimation of these copula models in high dimensions is challenging. Motivated by the classical inference function for margins approach, which first estimates the parameters of the marginal models and then estimate the copula parameters, we propose a new approach based on neural amortised inference, which we call neural inference function on margins. In the first step, the method uses neural networks to obtain an approximate posterior distribution of the marginal models. Given the estimated marginal parameters, we use another neural network model to obtain the approximate posterior distribution of the copula parameters in the second step. This approach provides rapid parameter estimation given new data, fast sequential prediction and efficient model comparisons using time-series cross validation. The proposed approach is evaluated using simulated and real datasets and compared to the popular Hamiltonian Monte Carlo method.

Graphical Causal Models for Economic Networks: A Machine Learning and Econometric Approach

**Leonard Mushunje**
*The University of Sydney*
m.samhan@unsw.edu.au

**Abstract:** Economic systems are increasingly characterized by complex interdependencies and networked interactions among agents, sectors, and markets. Understanding the causal structure underlying these economic networks is vital for effective policy design, risk management, and forecasting. This paper aims to develop and apply novel graphical causal modeling frameworks to infer, quantify, and analyze causal relationships in economic networks, integrating state-of-the-art techniques from machine learning, structural econometrics, and network science.

(Re-)Imag(in)ing and Graphing Correlation Risks

**Jinbo Cai**
*University of New South Wales*
jinbo.cai@student.unsw.edu.au

**Abstract:** In this talk, we shall discuss the use of Bayesian nonlinear expectations to construct lower and upper estimates for prices of Ether options I develop an embeddings–enhanced HAR framework for forecasting correlation risks, which first treats realized correlation matrices as Images (for convolutional neural networks, CNNs) to capture temporal patterns and as Graphs (for graph neural networks, GNNs) to capture cross-asset comovements, then integrates these embedding representations in a partial linear manner: the HAR block enters linearly while the learned high-dimensional embeddings enter flexibly. It serves dual motives. For out-of-sample forecasting, the approach manifests in statistically superior performance, and economically meaningful gains in pairs trading, equity premium predictions, risk targeting and global minimum variance (GMV) portfolios. For inference, valid inference relies on extended Frisch--Waugh--Lovell (FWL) partialling-out and cross-fitted Double Machine Learning (DML) with Neyman-orthogonal scores, yielding debiased estimates of HAR coefficients provided embeddings' incremental contribution. The framework is also interpretable for the embeddings, unveiling what CNNs/GNNs actually learn from the realized-correlations-encoded Images and Graphs.

Volatility Forecasting of CSI 300 ETF: A Comparison between HAR-Type Models and Neural Network Approaches

**Yuji Deng**
*The University of Sydney*
yden0728@uni.sydney.edu.au

**Abstract:** This study examines the predictive performance of HAR-type models and neural network approaches for volatility forecasting of the CSI 300 ETF using high-frequency data. Multiple HAR-type models, including HAR, HAR-J, and their extended variants, are constructed to capture the multi-scale dependency and jump characteristics of realized volatility. Deep learning architectures such as LSTM, GRU, and MLP are then applied to explore nonlinear dynamics. Empirical results reveal that while HAR-type models effectively capture the hierarchical memory structure of volatility, neural networks outperform in modelling complex nonlinear patterns. Furthermore, hybrid models combining HAR structures with neural networks achieve superior forecasting accuracy. These findings offer new insights into volatility modelling and risk management in high-frequency financial markets.

Measuring Climate Change Concern in the Media: Economic and Financial Impacts in Australia

**Zhixiang (Elvis) Yang**
*Monash University*
Zhixiang.Yang@monash.edu

**Abstract:** Climate change not only increases the frequency and severity of extreme events, but also reshapes public perceptions of economic and financial risks, thereby putting pressure on both the real economy and financial markets. Using a dictionary-based textual analysis, we construct novel newspaper-based narrative indices that track media attention to climate change in major Australian outlets. To capture heterogeneity in narratives, we further classify climate articles into thematic categories via Latent Dirichlet Allocation (LDA) and build topic-specific indices. In the empirical analysis, we first study asset-pricing implications using daily equity returns matched to firm-level greenhouse-gas (GHG) intensities. We find that increases in climate concern are associated with lower contemporaneous returns for nearly all firms, with significantly larger (more negative) exposures for high emitters; the interaction of climate concern with log-GHG intensity amplifies these effects and varies over time. At the topical level, transition- and policy-oriented narratives (e.g., carbon pricing, domestic regulation, energy-market transition) are priced in the cross-section, penalising carbon-intensive firms more strongly, whereas science-only or disaster-reporting narratives display weaker and less persistent pricing power. We then examine macroeconomic effects in a VARX framework. An unanticipated rise in climate concern has various impacts on macroeconomic activities as well.