Australian National University

Research School of Finance, Actuarial Studies and Statistics
2026 Summer Research Camp

Tuesday, 1st December

9.00 am **Coffee/tea**

9.20 am **Welcome message** *(University House – Common Room)*

9.30 am **Keynote Session 1** *(University House – Common Room)*

Introduction by Ding Ding (ANU)

Speaker: Stephan Siegel (University of Washington)
Title: Asking Why: Qualitative Research as a Method of Discovery in Finance

10.30 am **Morning tea**

11.00 am **Parallel Session 1**

Finance *(Common Room)*
Chair: Phong Ngo

Speaker: Deniz Okat (UTS)
Title: Debt overhang and strategic communication
Discussant: Le Zhang (ANU)

Statistics *(Drawing Room)*
Chair: Andrew Wood

Speaker: Xinyuan Song (CUHK)
Title: Statistical Learning for Heterogeneous Multivariate Data

Parallel Session 1 (cont.)

Finance *(Common Room)*
Chair: Phong Ngo

Speaker: Kun Li (ANU)
Title: Safe as Houses? Mortgage Finance and Domestic Violence
Discussant: Yi Fan (NUS)

Statistics *(Drawing Room)*
Chair: Andrew Wood

Speaker: Sungkyu Jung (SNU)
Title: Optimal Test-Data Piling in High-Dimensional Classification: A Geometric View of Benign Overfitting

12.30 pm **Lunch**

1.30 pm - 2.30 pm **Keynote Session 2** *(Common Room)*

Introduction by Andrew Wood

Speaker: Matt Wand (University of Technology Sydney)
Title: Practical Sophisticated Semiparametric Regression and Gibbsianity

2.45 pm **Visit to the Botanic Gardens**

6.00 pm **Dinner (Symposium)**

Wednesday, 2nd December

9.00 am **Coffee/tea**

9.15 am **Parallel Session 2**

Finance *(Common Room)*
Chair: Anna von Reibnitz

Speaker: Kate Volkova (U. Melbourne)
Title: Him Too? Analyzing the Effects of Epstein Connections
Discussant: Shasta Shakya (ANU)

Speaker: Hoa Briscoe-Tran (Monash U.)
Title: Disagreeing on DEI: Investors and Workers' Response to Anti-DEI Legislations
Discussant: Takeshi Yamada (ANU)

Speaker: Yi Fan (NUS)
Title: Bricks or Cash? Externalities of Housing Upgrading in High-density Cities
Discussant: Nhan Le (ANU)

Statistics *(Drawing Room)*
Chair: Le Chang

Speaker: Tiandong Wang
Title: Automatic Threshold Selection Leaves Finite-Degree Bias in Network Power-Law Fits

Speaker: Jiazhen Xu
Title: Metric subtraction for non-Euclidean data with applications to time series decomposition and residual analysis

Actuarial Studies *(Function Room)*
Chair: Adam Butt

Speaker: Kwangmin Jung (Pohang University of Science and Technology)
Title: Insurance Stress Testing of Climate Change Risk: The Case of the Korean General Insurance Industry

Speaker: Ramona Meyricke (Taylor Fry)
Title: Heat-Related Risk: From Hidden Workplace Injuries to Complex System Cascades

11.00 am **Morning tea**

11.30 am **Parallel Session 3**

Finance *(Common Room)*
Chair: Kelly Liu

Speaker: Haihao Ouyang
Title: Who Supplies Maturity over the Credit Cycle? Evidence from the Syndicated Loan Market
Discussant: Guangqian (Isaac) Pan (U. Sydney)

Speaker: Muzhao Yang
Title: Beta anomaly and mispricing factor
Discussant: Bruce Grundy (ANU)

Statistics *(Drawing Room)*
Chair: Yanrong Yang

Speaker: Jia Chen
Title: IV Estimation of Heterogeneous Spatial Dynamic Panel Models with Interactive Effects

Speaker: Bradley Rava
Title: Selective Decision-Making for Fair and Reliable Human–AI Collaboration

Actuarial Studies *(Function Room)*
Chair: Ning Wang

Speaker: Xiaobai Zhu (Chinese University of Hong Kong)
Title: Aging by Design: Deriving Mortality Models from Reliability Theory

Speaker: Colin Yellowlees (MNC Consulting)
Title: Uses of Bayesian Modelling in Life Insurance

1.00 pm **Lunch + school-wide thanks**

2.00 pm **Parallel Session 4**

Finance *(Common Room)*
Chair: Dean Katselas

Speaker: Raymond Liu (ANU)
Title: Wrong for rights? Short sales in the global rights issues market
Discussant: Deniz Okat (UTS)

Speaker: Jian Zhang (HKU)
Title: A Tale of Two Zoos: Machine Learning Insights on Fifteen Million Retail Investors
Discussant: Qiaoqiao Zhu (ANU)

Statistics *(Drawing Room)*
Chair: Le Chang and Yanrong Yang

Speaker: Zhi Yang Tho (ANU)
Title: Bias-Adjusted Attribution Estimation for Rainfall Enhancement Trials

Speaker: Bao Anh Vu (ANU)
Title: Bioregionalisation with a random effect Region of Common Profile model

Speaker: Adam Nie (ANU)
Title: Multi-task Learning on Principal Subspaces for Heterogeneous Time Series

Speaker: Santiago Marin (ANU)
Title: Adaptive Shrinkage with a Nonparametric Bayesian Lasso

Speaker: Chuang Xu (ANU)
Title: Fairness-regularized Principal Component Analysis

Actuarial Studies *(Function Room)*
Chair: William Lim

Speaker: Patrick Laub (UNSW)
Title: Actuarial Neural Networks: Uncertainty and Interpretability in Deep Learning

Speaker: Oscar Chen (Colonial First State)
Title: Always-On Intelligence: Embedding AI in Retirement Guidance

3.30 pm **End**

PRESENTATION TITLES AND ABSTRACTS
(in order of appearance)

FINANCE

Keynote speaker: Stephan Siegel (U. Washington)
Title: Asking Why: Qualitative Research as a Method of Discovery in Finance
Abstract: Why do people do what they do? Finance typically approaches this question through models: researchers formulate explanations and test their implications using observed behavior. But people can also tell us how they understand their decisions, revealing beliefs and mechanisms that established models may overlook. I will make the case for qualitative research as a method of scientific discovery in finance by discussing how interviews can generate hypotheses and inform quantitative measures and experiments, and how large language models can support qualitative research. As an application, I study limited stock market participation in Germany. Interviews reveal a widespread belief that successful investing requires identifying "safe" stocks, avoiding "bad" ones, and monitoring and trading frequently. These beliefs raise perceived participation costs. A representative survey of more than 7,000 individuals documents their prevalence, and an experiment shows that informing individuals about market efficiency increases their willingness to invest in stocks. Together, the three methods form a sequence: interviews generate hypotheses, surveys establish prevalence, and experiments provide causal evidence.

Speaker: Deniz Okat (UTS)
Title: Debt Overhang and Strategic Communication
Abstract: We study how lenders should communicate information to borrowers when disclosure affects effort incentives. Favorable information can exacerbate debt overhang: a borrower who learns that a project is likely to succeed without further effort may shirk because debt absorbs part of the return to effort. Optimal communication may therefore involve selective opacity, pooling states with high effort responsiveness behind an effort recommendation while recommending no effort in less responsive states. Pooling also increases expected debt recovery, lowering the competitive face value of debt and further strengthening effort incentives. Under single-crossing and regularity conditions, intermediate disclosure maximizes credit capacity, potentially financing projects that neither full transparency nor full opacity can support. State-contingent debt can eliminate this advantage, while contractible effort removes the incentive role of disclosure.

Speaker: Kun Li (ANU)
Title: Safe as Houses? Mortgage Finance and Domestic Violence
Abstract: Mortgage contracts affect both the cost of staying in a shared home and the cost of financing a home after separation. In U.S. county data for 2005-2023, larger increases in scheduled payments on adjustable-rate mortgages are associated with more reported current-partner incidents at home. During the seven quarters of tightening in 2022-2023, a $10 higher monthly cost of replacing fixed-rate debt per household is associated with 0.86% more such incidents among continuously reporting agencies in 1,187 counties. Higher replacement costs are also associated with more former-partner incidents. We use a dynamic household model to compare mortgage payment relief with a grant available upon separation. Across historical starting dates, payment relief lowers simulated first-quarter incidents by 6.0%. In a separate simulation beginning in 2004, relief reduces separation.

Longer relationship duration offsets most of the reduction in incident risk, yielding 0.55% fewer cumulative incidents within active relationships over 2005-2023. The separation grant reduces cumulative incidents primarily by facilitating exit.

Speaker: Kate Volkova (U. Melbourne)
Title: Him Too? Analyzing the Effects of Epstein Connections
Abstract Who captures the value of elite connections? Private ties among elites are rarely observable. We use more than a million documents released by the U.S. Department of Justice to study contacts between Jeffrey Epstein and directors of U.S. public companies. Firms with connected directors face more misconduct incidents. Misconduct rises after the first contact, both within firms and within directors, and more so when emails indicate closer personal ties. This increase ends after Epstein's death. When the connections became public in January 2026, investors penalized firms identified in news coverage. We find no evidence of benefits to firms or professional gains for directors. Instead, the apparent benefits are nonpecuniary: emails document access to information, hospitality, status, and help with personal matters. Connected directors also hold more non-profit board seats and gain access to a broader social circle. They attend more social events and, increasingly, appear at these events together. Elite connections thus appear to serve directors' private interests, at a cost to the firms they oversee.

Speaker: Hoa Briscoe-Tran (Monash U.)
Title: Disagreeing on DEI: Investors and Workers' Response to Anti-DEI Legislations
Abstract: How do investors and workers respond to restrictions on corporate Diversity, Equity, and Inclusion (DEI) initiatives? We exploit Florida's December 2021 announcement of the Stop WOKE Act, which restricts mandatory workplace DEI training and allows employees to sue employers over prohibited content. Florida-headquartered firms experience cumulative abnormal returns approximately 74 basis points below those of other U.S. firms, with the effect reversing when a federal court blocks the workplace provisions in August 2022. These average effects mask substantial heterogeneity. Announcement returns are especially negative for smaller, less established firms, firms with more pro-social investors, and firms with more Republican-leaning employees, while some firms experience positive returns. Analysts revise earnings forecasts downward, particularly for firms with more Republican-leaning employees, consistent with increased litigation risk. Investors with stronger social preferences also reduce their holdings of Florida firms. By contrast, employee ratings of diversity and inclusion improve on average, especially in firms with predominantly Republican-leaning workforces, and worker departures decline overall. Together, the findings suggest that anti-DEI legislation can curb agency costs at some firms while increasing litigation risk and conflicting with investors' social preferences at others. Such legislation redistributes value and legal rights across firms and workers, with effects that depend on investor preferences and workforce political composition.

Speaker: Yi Fan (NUS)
Title: Bricks or Cash? Externalities of Housing Upgrading in High-density Cities
Abstract: We estimate housing externalities in a high-density city, exploiting the staggered rollout of Singapore's nationwide Main Upgrading Programme for public housing. Controlling for non-random neighborhood exposure, we find that upgrading raises treated buildings' prices by 11.5% upon completion and neighboring buildings' resale prices by about 2% within 500 meters, decaying to zero beyond. A model with distance-decaying externalities shows that in dense settings spillovers justify the distortions of in-kind provision; this advantage diminishes and reverses at lower densities. Administrative data on over 2 million residents show that upgrading disproportionately retains older incumbents, suggesting age-specific amenities as an underexplored externality channel.

Speaker: Haihao Ouyang (ANU)
Title: Who Supplies Maturity over the Credit Cycle? Evidence from the Syndicated Loan Market
Abstract: This study examines whether maturity adjusts differently across bank and non-bank credit when credit tightens. Using U.S. syndicated loans from 2000–2025, we distinguish Term B products from non-bank participation and measure exposure in committed dollar-years, combining facility amount and maturity. A one-standard-deviation increase in the Excess Bond Premium shortens TLA-type maturity by about 15%, but Term B maturity by only 5%. This relative maturity resilience does not consistently extend to committed exposure. Greater non-bank participation predicts less contraction across deals, although the differential weakens within deals. Credit cyclicality therefore depends on commitment size, horizon, and lender composition.

Speaker: Muzhao Yang (ANU)
Title: Beta Anomaly and Mispricing Factor
Abstract: I develop a theoretical link between mispricing and the beta anomaly using the mispricing framework of Kozak, Nagel, and Santosh (2018). When the market portfolio contains an undiversifiable mispricing component, CAPM beta reflects both systematic market risk and systematic mispricing. The model predicts that controlling for systematic mispricing beta explains the beta anomaly and restores an upward-sloping Security Market Line, with the slope varying systematically with the level of mispricing beta. To test these predictions, I proxy the mispricing component using characteristic-based factors under the assumption that investor belief distortions are systematically related to firm characteristics. Among the factors examined, the MGMT factor of Stambaugh and Yuan (2017) provides the strongest support for the model. Controlling for the MGMT component subsumes not only the unconditional beta anomaly but also its conditional variation across investor sentiment, economic uncertainty and the volatility puzzle documented by Barroso et al. (2025). I further provide complementary evidence from mutual fund holdings and fund-level behaviour, showing that institutional demand is systematically related to the mispricing component in a manner consistent with the model's mechanism.

Speaker: Wai-Man (Raymond) Liu (ANU)
Title: Wrong for Rights? Short Sales in the Global Rights Issues Market
Abstract: Rights issues are designed to protect shareholder wealth during capital raisings, but paradoxically, they attract low shareholder participation. We provide theory and evidence indicating that rights holders are vulnerable to strategic short-selling activity, which depresses stock prices and take-up rates. Our comprehensive cross-country data reveal a one-third decline in rights values during issue execution periods, coinciding with increased short volumes and greater involvement by large short-sellers. Regression discontinuity analysis further suggests that this relationship is likely causal. Retail investors, facing higher subscription costs, are disproportionately impacted. Overall, there is a case for tightening short-selling restrictions on _rms conducting rights issues.

Speaker: Jian Zhang (HKU)
Title: A Tale of Two Zoos: Machine Learning Insights on Fifteen Million Retail Investors
Abstract: Both stock characteristics (the "factor zoo") and behavioral biases (the "bias zoo") impact retail investors' returns. But which factors/biases are more important? We address this question by utilizing machine-learning tools to analyze 15 million retail investor accounts in India. We observe that Neural Networks outperform other algorithms in uniquely predicting both good and bad out-of-sample performance, with (under)diversification, portfolio turnover, and momentum being the leading factors to influence investors' total returns. Neural Network-identified smart investors also help improve market efficiency. Our results provide insights into the role of retail investors and the plausible asset-pricing framework that explains their returns.

STATISTICS

Keynote speaker: Matt Wand (UTS)
Title: Practical Sophisticated Semiparametric Regression and Gibbsianity
Abstract: The speaker's main research area is semiparametric regression and he is regularly involved in developing methodology and software for various sophistications. Entwined with this work is the considerable effort that goes into making the methodology practical. This involves aspects such as orthogonal spline bases, lower-level language coding and public software dissemination. For Bayesian approaches the notion of Gibbsianity is also important, and is a central theme of this presentation. A highlight is an October 2026 development that results in Gibbsian algorithms for Bayesian semiparametric regression with count responses.

Speaker: Xinyuan Song (Chinese University of Hong Kong)
Title: Statistical Learning for Heterogeneous Multivariate Data
Abstract: Modern scientific studies routinely generate heterogeneous multivariate data with complex dependence across variables, individuals, and time, posing substantial challenges for statistical learning. In this talk, we propose a unified framework for modeling such data by integrating latent variable models, mixed membership structures, and longitudinal trajectory analysis. We first introduce a mixed membership latent variable model that combines exploratory factor analysis with flexible subpopulation modeling, enabling simultaneous learning of latent factors, factor structures, and profile memberships while capturing individual heterogeneity and addressing multicollinearity. We then present a multi-label growth curve change-point model for multivariate longitudinal data, which accommodates between-individual, between-variable, and within-subject temporal heterogeneity through domain-specific latent classes and subject-specific change points, allowing individuals to exhibit distinct progression patterns across multiple outcome domains. Together, these approaches provide a general paradigm for statistical learning in heterogeneous multivariate data, emphasizing interpretable latent representations, partial memberships, and dynamic structures. Theoretically, we show that the proposed methods accurately recovers the underlying heterogeneity structure. Applications to Parkinson's disease data demonstrate their ability to uncover meaningful latent traits, domain-specific progression patterns, and refined disease subtypes, offering valuable insights for complex data analysis and precision medicine.

Speaker: Sungkyu Jung (Seoul National University)
Title: Optimal Test-Data Piling in High-Dimensional Classification: A Geometric View of Benign Overfitting
Abstract: High-dimensional classification can exhibit seemingly paradoxical behavior: a classifier may fit the training data perfectly and still generalize well. In this talk, I discuss this phenomenon through the geometry of data piling, where high-dimensional observations collapse onto a small number of points after projection. When the dimension exceeds the sample size, many directions produce perfect class-wise piling of the training data, but such interpolation alone does not guarantee good out-of-sample performance. The key question is whether independent test observations also pile, and which direction maximizes the separation between the resulting test-data piles. I focus on classification under heterogeneous spiked covariance models. We characterize directions that yield asymptotic test-data piling and identify the Second Maximal Data Piling direction, which maximizes asymptotic test separation. The results show that substantial covariance heterogeneity can still be compatible with perfect classification, while imbalance in the tail eigenvalues plays a critical role. These findings provide a geometric perspective on benign overfitting in high-dimensional classification: perfect training fit is abundant, but generalization depends on which interpolating direction is selected.

Speaker: Tiandong Wang (Fudan University)
Title: Automatic Threshold Selection Leaves Finite-Degree Bias in Network Power-Law Fits
Abstract: Minimum-distance fitting selects a degree threshold and a power-law exponent from the same network. We study the procedure that fits a continuous Pareto law with the Hill estimator and minimizes the Kolmogorov--Smirnov distance over retained ranks. For homogeneous linear preferential attachment and Norros--Reittu graphs with Pareto weights, we derive the joint limiting law of the selected rank, retained fraction of the threshold tie, and exponent error. The selector retains of order $n^{2/(\alpha+2)}$ of n vertices, where α is the survival exponent. At this rank, integer spacing and finite-degree departures from a Pareto tail have the same order as sampling fluctuations. The exponent estimate is consistent and retains a shift on the scale used for confidence intervals. For integer degree laws whose tail sums have a first-order inverse-degree expansion, adding the retained tie fraction divided by the threshold to the Hill statistic removes the leading contribution of the tie. Under homogeneous preferential attachment with a fixed known number of edges per new vertex, removing the remaining model-specific term and calibrating for selection gives confidence intervals with asymptotically exact coverage. Observed networks show sensitivity of the fitted exponent and the classification of degree variance to threshold ties. Simulations assess interval coverage and show improved pooled connectivity predictions after correction under both specified graph laws. Changes to the Norros--Reittu weight distribution can reverse the prediction gain.

Speaker: Jiazhen Xu (Macquarie University)
Title: Metric subtraction for non-Euclidean data with applications to time series decomposition and residual analysis
Abstract: Subtraction plays a fundamental role in time series decomposition and regression diagnostics, yet a natural subtraction operation is generally unavailable for non-Euclidean data. This talk introduces metric subtraction as a geometric framework for removing systematic structure while retaining observations in their original object space. The central idea is to transport geodesic to a common reference object, extending the interpretation of Euclidean subtraction to non-Euclidean settings. We develop this framework through two applications. First, we introduce a decomposition model with a smooth trend, an unknown periodic component for spherical time series. Second, for model with object-valued responses, we develop diagnostic and inferential tools for residual analysis. An application to U.S. electricity generation compositions illustrates the utility of metric subtraction for interpretable decomposition, model assessment, and prediction.

Speaker: Jia Chen (University of Macau)
Title: IV Estimation of Heterogeneous Spatial Dynamic Panel Models with Interactive Effects
Abstract: We develops a Mean Group Instrumental Variables (MGIV) approach for spatial dynamic panel data models with interactive effects, under large N and T asymptotics. Unlike existing approaches that typically impose slope-parameter homogeneity, MGIV accommodates cross-sectional heterogeneity in slope coefficients. The proposed estimator is linear, making it computationally efficient and robust. Furthermore, it enables asymptotically valid inferences without requiring bias correction. The Monte Carlo experiments indicate strong finite-sample performance of the MGIV approach across various sample sizes and parameter configurations. The practical utility of the methodology is illustrated through an application to regional economic growth in Europe. Our results provide evidence of conditional convergence in regional growth dynamics. Spillovers are shown to play a dominant role, particularly for investment rates where nearly four-fifths of the total effect on GDP per capita growth originates from neighboring regions.

Speaker: Bradley Rava (University of Sydney)
Title: Selective Decision-Making for Fair and Reliable Human–AI Collaboration
Abstract: As AI systems are increasingly used to support consequential decisions, an important question is not simply how to automate decisions accurately, but when decisions should be automated at all. In many settings, allowing an algorithm to abstain and defer difficult cases to a human can provide a principled foundation for human–AI collaboration, while enabling stronger guarantees on accuracy, error control, and fairness.
In this talk, I will present three approaches to selective decision-making. First, motivated by algorithmic fairness, I introduce a framework for selectively automating decisions while controlling errors across protected groups, illustrating how abstention can help balance automation with equitable performance. I then study the fundamental role of indecisions in classification and hypothesis testing, characterizing how much abstention is necessary to achieve a desired level of accuracy and showing that selective decisions can attain error rates beyond what is possible when every observation must be classified. Finally, I consider Neyman–Pearson classification when both Type I and Type II errors must be controlled simultaneously. When these requirements cannot be satisfied by a fully automated classifier, allowing indecisions restores feasibility and enables finite-sample, high-probability control of both selective error rates.
Together, these results provide a statistical framework for deciding when an AI system should act and when it should defer. More broadly, they demonstrate how selective automation can support reliable and fair human–AI collaboration by reserving automated decisions for cases where meaningful statistical guarantees can be provided.
Papers
1. A Burden Shared is a Burden Halved: A Fairness Adjusted Approach to Classification
   Journal of the American Statistical Association
   https://www.tandfonline.com/doi/full/10.1080/01621459.2026.2634937
2. Ask for More than Bayes Optimal: A Theory of Indecisions for Selective Hypothesis Testing
   https://arxiv.org/abs/2412.12807
3. Neyman–Pearson Classification with Indecisions: Simultaneous Control of Selective Type I and Type II Errors
   https://arxiv.org/abs/2610.04233

RSFAS Statistics Speakers

Speaker: Zhi Yang Tho (ANU)
Title: Bias-Adjusted Attribution Estimation for Rainfall Enhancement Trials
Abstract: Model-based analyses of rainfall enhancement trial data typically involve modelling log-transformed rainfall using linear mixed models to assess the effectiveness of enhancement methods under real-world conditions. This approach improves on traditional average-based analyses by allowing explicit control for the effects of meteorological and topographical covariates that may affect precipitation amounts. However, a key issue with such analyses is the bias that arises when back-transforming the log-rainfall to the original scale for estimating attribution, defined as the additional raw-scale rainfall attributable to the enhancement method. To address this issue, we propose a new attribution estimator that incorporates theoretically justified, observation-specific bias-adjustment terms. The proposed estimator improves upon existing estimators that rely on arbitrary adjustments, and satisfies a coherence property that ensures zero estimated attribution for observations without enhancement intervention. A proportional random effect block bootstrap is further used to conduct inference on the attribution quantities. Applying both the proposed and an existing model-based estimator to the Oman rainfall enhancement trial from 2013 to 2018, we find a statistically significant positive effect of the ground-based ionization technology on downwind rainfall at the 5% significance level, with our proposed estimator indicating a smaller effect than the existing method, while average-based analyses produce substantially wider bootstrap intervals and larger p-values. A simulation study further supports the findings based on the proposed estimator, demonstrating its superior estimation accuracy and improved inferential performance of the associated bootstrap confidence intervals, as well as its greater efficiency relative to the average-based estimators.

Speaker: Bao Anh Vu (ANU)
Title: Bioregionalisation with a random effect Region of Common Profile model
Abstract: In conservation and resources planning, it is often useful to divide a spatial region of interest into smaller regions with similar biological content and develop management strategies that are tailored to each region's needs. Region of Common Profile (RCP) models are a class of mixture-of-expert models used for clustering spatial locations into regions with homogeneous species composition and understanding their relationship to environmental factors. Existing RCP models assume that species composition does not vary between sites between the same region and that species within a region are uncorrelated, both of which are restrictive and unrealistic assumptions given the diversity often observed in ecological data. In this work, we propose an extension to the standard RCP models, called random-effect RCP (RERCP), by modelling the mean species response at each site as random effects drawn from a finite Gaussian mixture. The introduction of these latent variables provides a more flexible representation of within-region heterogeneity and further enables examination of residual correlation among species within each region. Model fitting is done via an expectation-maximisation (EM) algorithm, and we assume a low-rank covariance structure to simplify computations when the number of species is large. Simulations show that our proposed RERCP model achieves comparable, often better, classification accuracy and parameter estimation to the standard RCP model. An application to demersal fish data from the Kerguelen Plateau region in the Southern Indian Ocean demonstrates that the proposed RERCP model is effective for determining the number and spatial extent of bioregions as well as interpreting their species composition and environmental characteristics.

Speaker: Adam Nie (ANU)
Title: Multi-task Learning on Principal Subspaces for Heterogeneous Time Series
Abstract: Estimating low-dimensional structure across related time series requires balancing information sharing against population heterogeneity. We develop a multi-task principal component analysis framework that penalizes deviations of task-specific subspace projectors from an estimated common centre, interpolating between individual and fully pooled PCA. In a separable Hilbert space, we establish a unified nonasymptotic error decomposition under partial homogeneity, allowing an unknown subset of structurally unrelated tasks. The decomposition separates individual estimation error, pooled first-order fluctuations, heterogeneity bias, outlier contamination, and higher-order terms. For an admissible penalty range, we prove existence and local uniqueness of the estimator and give conditions under which information sharing improves subspace estimation. Under independent regular tasks, averaging reduces first-order stochastic error by the square root of their number, subject to bias and remainder terms. We specialize the theory to stationary and cointegrated high-dimensional factor models and functional time series, obtaining explicit convergence rates and penalty scales. Simulations illustrate the benefits of partial pooling across different levels of heterogeneity. An application to mortality forecasting across 32 countries reduces overall forecasting mean squared error by 3.5%–9.8% relative to individual PCA and 4.4%–7.2% relative to pooled PCA.

Speaker: Santiago Marin (ANU)
Title: Adaptive Shrinkage with a Nonparametric Bayesian Lasso
Abstract: Modern approaches to perform Bayesian variable selection rely mostly on the use of shrinkage priors. That said, an ideal shrinkage prior should be adaptive to different signal levels, ensuring that small effects are ruled out, while keeping relatively intact the important ones. With this task in mind, we develop the nonparametric Bayesian Lasso, an adaptive and flexible shrinkage prior for Bayesian regression and variable selection, particularly useful when the number of predictors is comparable or larger than the number of available data points. We build on spike-and-slab Lasso ideas and extend them by placing a Dirichlet process prior on the shrinkage parameters. The result is a prior on the regression coefficients that can be seen as an infinite mixture of Laplace distributions, all offering different amounts of regularization, ensuring a more adaptive and flexible shrinkage. We also develop an efficient Markov chain Monte Carlo algorithm for posterior inference. Through extensive simulation studies and real-world data analyses, we illustrate that our proposed method leads to coefficient recovery, variable selection accuracy, and out-of-sample predictions that are comparable to or better than those from state-of-the-art shrinkage priors, highlighting the benefits of the nonparametric Bayesian Lasso over existing methods.

Speaker: Chuang Xu (ANU)
Title: Fairness-regularized Principal Component Analysis
Abstract: Principal component analysis (PCA) minimizes aggregate reconstruction risk but may represent heterogeneous groups unevenly. We study fairness-regularized PCA, which penalizes the largest group-specific reconstruction regret relative to each group's optimal rank-r representation. Under spectral nondegeneracy, a small fairness penalty reduces worst-group regret to first order while increasing aggregate risk only to second order. We characterize when the regularization path is locally constant and derive a minimax spectral representation as PCA with endogenously adjusted group weights, establishing convex relaxation exactness under an eigengap condition. A uniform eigengap over the possibly nonunique dual-optimal set yields polyhedral identification, quadratic growth, and principal-subspace convergence rates allowing the dimension, rank, and number of groups to grow. We also establish uniform path recovery and guarantees for penalty calibration using preferences or fairness tolerances. Simulations

and an NHANES accelerometry application illustrate the trade-off: a moderate penalty reduces median test worst-group regret by 18.54, with a median aggregate risk increase of only 0.053. These results provide a principled statistical foundation for selecting shared low-dimensional representations that balance aggregate reconstruction accuracy with group-specific representation quality.

ACTUARIAL STUDIES

Speaker: Kwangmin Jung (Pohang University of Science and Technology)
Title: Insurance Stress Testing of Climate Change Risk: The Case of the Korean General Insurance Industry
Abstract: Climate change risk has emerged as a critical concern for Property & Casualty (P&C) insurers, as increasingly frequent and severe extreme weather events threaten underwriting performance and solvency. However, the inherent difficulty of quantifying an evolving climate impedes the development of reliable stress-testing practices. This talk presents an empirically grounded, liability-side stress-testing framework that employs the Actuaries Climate Index (ACI), localized to South Korea, to assess the impact of physical climate risk on insurers' financial resilience. We model the dependence between climate conditions and line-level incurred losses using a conditional copula approach and translate projected ACI paths under SSP scenarios into forward-looking incremental capital measures within the Korean Insurance Capital Standard (K-ICS). The results reveal heterogeneity across insurers and lines of business, suggesting that portfolio structure matters for the magnitude and horizon profile of capital needs. Beyond the empirical findings, the talk discusses how supervisory climate stress-testing frameworks can integrate long-term climate trends into annual capital adequacy cycles, and what this implies for regulators seeking to balance internal model flexibility with supervisory consistency in managing climate-related solvency risks.

Speaker: Ramona Meyricke (Taylor Fry)
Title: Heat-Related Risk: From Hidden Workplace Injuries to Complex System Cascades
Abstract: Australia is getting hotter, yet the health and economic consequences of extreme heat remain systematically underestimated. In workers' compensation data, heat often contributes to illness and injury without being recorded as the primary cause. Similar attribution challenges affect mortality and hospitalisation statistics, while wider heat impacts can propagate through interconnected infrastructure, health, and emergency-response systems. This presentation introduces the Heat-Related Workforce Injury Risk Index, a metric integrating meteorological, workforce, and compensation data to quantify and compare heat-related injury risk across Australian industries and occupations. Using Safe Work Australia's National Data Set for Compensation-based Statistics (NDS), it links temperature exposure directly to observed injury patterns to identify priority risks, target prevention, and inform climate adaptation investment. We also demonstrate how machine learning can identify NDS claims where heat likely contributed despite not being explicitly coded. Applied to over a decade of national data, this method provides a fuller picture of the burden and distribution of heat-related workplace injury. Finally, we examine complex, cascading heat risks driven by interactions between hazards, exposed systems, vulnerabilities, and responses. Because heatwaves span wide areas and simultaneously stress electricity, buildings, healthcare, and emergency services - for instance, driving up power demand while reducing generation and transmission capacity - they risk triggering widespread power disruptions. We review recent research on heat-risk cascades relevant to Australia and outline future research priorities.

Speaker: Xiaobai Zhu (Chinese University of Hong Kong)
Title: Aging by Design: Deriving Mortality Models from Reliability Theory
Abstract: The Cairns-Blake-Dowd (CBD) model is a workhorse mortality model in actuarial science. Despite its seminal role in actuarial applications, its logit-linear-in-age structure rests on statistical assumption rather than biological mechanism. Using reliability theory of aging, this paper provides the first mechanistic derivation of the CBD class of mortality models. Specifically, its defining structure emerges as an exact mathematical consequence of a natural aging process in which an organism's biological subsystems fail through a self-amplifying cascade. We then progressively enrich the reliability parameter structure to derive a new hierarchy of mortality models, where cohort heterogeneity in mortality is attributed to initial frailty and aging intensity. Calibrated on national mortality data, the reliability-based models consistently outperform their actuarial counterparts at every structural level of the hierarchy. We also demonstrate the practical value of grounding mortality models in an interpretable biological aging process.

Speaker: Colin Yellowlees (MNC Consulting)
Title: Uses of Bayesian Modelling in Life Insurance
Abstract: Bayesian methods can allow actuaries to solve actuarial problems in ways that are often difficult under traditional approaches. For example, they allow actuaries to quantify the full distribution of outcomes rather than relying solely on point estimates or asymptotic approximations. They also adapt well to situations with sparse data, shifting business practices, or environments where existing knowledge can add meaningful insights. These strengths can improve decision-making in contexts such as pricing, reserving, and capital modelling, where understanding the range and drivers of uncertainty is critical. Historically Bayesian methods remain underused in life insurance practice. This presentation will introduce Bayesian approaches, highlighting how they differ from more traditional methods and why they can be useful. We will discuss their flexibility in handling complex models, their natural treatment of uncertainty through posterior distributions, and their ability to incorporate expert judgment or external information via priors. The presentation will present a couple of case studies. One will examine Bayesian approaches to experience investigations, showing how Bayesian modelling can address sparse data and credibility challenges. The other will look at Incurred But Not Reported (IBNR) reserving and how the Chain Ladder Method can be extended.

Speaker: Patrick Laub (UNSW)
Title: Actuarial Neural Networks: Uncertainty and Interpretability in Deep Learning
Abstract: A typical neural network takes in some input, applies many indecipherable transformations, then produces a single number as the output. As the internal logic is so opaque, it is not obvious how the network arrived at its prediction, nor whether it is making a confident prediction or is just wildly speculating. This talk covers two pieces of work which each tackle one of these problems. In the first part of the talk, I will touch on distributional regression, where we ask the networks to give us predictions in the form of probability distributions, and specifically our Distributional Refinement Network (DRN). The DRN starts from a trusted baseline model (typically a generalised linear model) and augments it with some neural network flexibility. We have applied these networks to claim severity modelling and to individual claim reserving. In the second part, I will introduce the Actuarial Neural Additive Model (ANAM), an inherently interpretable deep learning model for general insurance pricing. Each rating factor (and selected pairwise interactions) enters the model through its own shape function, which can be plotted and inspected directly, and the model enforces properties which pricing actuaries need, such as smoothness and monotonicity.

Speaker: Oscar Chen (Colonial First State)
Title: Always-On Intelligence: Embedding AI in Retirement Guidance
Abstract: As superannuation funds evolve to meet the Retirement Income Covenant and rising member expectations, artificial intelligence offers a promising avenue to enhance retirement guidance and engagement. This presentation explores an AI-enabled support system that combines predictive analytics with human expertise to deliver timely, personalised nudges and future-self scenarios. Drawing on a holistic view of member data, including transactional, behavioural and family wealth information, the system aims to anticipate financial shifts and connect members to appropriate support at key decision points. The session will outline the design principles and key features of this AI-human hybrid model, including dynamic net worth analytics, predictive alert mechanisms and seamless adviser handovers. It will consider how such systems can align with regulatory expectations around effective communication and member-centric engagement, including the principles of clarity, relevance and timing in retirement communications outlined by Treasury and ASIC. By framing AI as a strategic enabler within retirement ecosystems, the presentation invites audience and industry practitioners to consider how intelligent systems can be responsibly embedded to support better outcomes. It offers a forward-looking perspective on how predictive intelligence, paired with human advice, can transform retirement planning from a static process into a dynamic, lifelong journey.