6th Waterloo Student Conference in Statistics, Actuarial Science and Finance

Book of Abstracts

Day 1

Keynote (Chair: Benjamin Côté)

Name: Dr. Mike Hoy

Title: Welfare Economics Meets Risk Theory

Abstract: This seminar has three principal aims: (1) a short review of "standard" welfare economics (with some extras) that incorporates distributional concerns, (2) a discussion of how distributionally sensitive welfare analysis and inequality measurement have been influenced by risk theory concepts, (3) a few simple examples illustrating applications involving risk that help to connect theory to practice and also illustrate challenges in doing so.

Session 1A: Topics in Statistics (Chair: Madison Cranstone)

Talk 1

Name: Bryn Crandles

Title: Accelerated Majorization-Minimization Algorithms for Estimating Common Eigenvector Matrices in Gaussian Mixture Models

Abstract: Accelerated majorization-minimization (MM) algorithms are developed to estimate the common eigenvector matrix for members of the families of Gaussian parsimonious clustering models (GPCM) and high-dimensional data clustering models (HDDC). These Gaussian mixture models have often been avoided due to the perceived need for a cumbersome algorithm to estimate the common eigenvector matrix. However, simple MM algorithms have been developed that demonstrate improved computational time and reliable convergence for two common orientation GPCM models. MM algorithms are similarly developed for the common orientation HDDC models. Further, the MM algorithms are accelerated through novel mappings back to the space of matrices with an orthonormal constraint. Simulations and real data analyses demonstrate that the

proposed methods outperform previously available approaches, enabling practical implementation of the full families of Gaussian mixture models.

Talk 2

Name: Rohit Gajendragadkar

Title: Comparing continuously generated probability forecasts

Abstract: Continuously updated probability forecasts are commonly seen - for example in sports betting, where the probability of the home team winning a game is continuously updated as the game goes on. Given a set of two continuously updated probability forecasts, the goal is to determine which is the "better" one, with respect to some loss function. This can be done by comparing the two forecasts on average, using a Diebold-Mariano style test for continuously updated forecasts, or can be done by comparing them conditioned on some available information, using appropriately chosen instrument processes. Of interest is also to construct uniform confidence sets of the (unobservable) loss difference process.

Talk 3

Name: Weiyang Yan

Title: Sample Empirical Likelihood Inference for Non-probability Samples

Abstract: In this paper, we develop a sample empirical likelihood framework for estimation and inference on the finite population mean using a nonprobability sample together with a reference probability sample. We propose two approaches that differ in how propensity score information is incorporated. The first introduces propensity-score calibration constraints in addition to the standard model-calibration constraints, while the second uses propensity scores to construct weighted model-calibration constraints. For asymptotic development, we focus on the second approach and derive the limiting distribution of the sample empirical likelihood ratio statistic in the absence of model-calibration constraints. Bootstrap procedures are developed to construct confidence intervals for both approaches. Simulation studies show satisfactory finite-sample performance and provide empirical evidence of double robustness.

Session 1B: Dynamic Models (Chair: Qianyu Chen)

Talk 1

Name: Omar Essakine

Title: Multivariate Compound risk model based on tree-Structured Poisson Markov Random fields: Asymptotic results and risk sharing

Abstract: We establish asymptotic results for aggregate losses in a heterogeneous risk model whose claim counts follow a tree-structured Poisson Markov random field. We first investigate the weak-dependence structure of this model by deriving explicit exponential-decay bounds for the β-mixing under suitable conditions on the tree structure and dependence parameters. We also show that that the stronger ϕ-mixing conditions generally fail. Under suitable structural and moment conditions, these dependence bounds yield limit theorems for aggregate losses, including the strong law of large numbers, the central limit theorem and the law of the iterated logarithm. As an application, we analyze risk sharing in large insurance pools under linear fair risk sharing rules and the conditional mean risk sharing rule. We identify conditions under which individual contributions converge to their pure premiums and their centered fluctuations are asymptotically normal.

Talk 2

Name: Brandon Tam

Title: Pure-Exchange Equilibria with Trade Aversion

Abstract: We study a problem of optimal allocation in a pure-exchange economy, where agents have preferences that are represented by the sum of a distortion risk measure and a Bregman-Wasserstein (BW) divergence. The addition of a BW divergence to traditional distortion risk measure preferences allows us to study agents that are both risk and trade averse. Under this framework, we show that the optimal allocation problem can be rewritten as a constrained convex optimization problem when agent allocations are assumed to be comonotone. We further derive a numerically implementable characterization for the optimal allocation and provide sufficient conditions under which a closed-form characterization exists. We illustrate our results in a three-agent setting and compare to existing characterization results for optimal allocations under distortion risk measure preferences without trade aversion.

Talk 3

Name: Lilian Hu

Title: Efficient and accurate simulation of the stochastic-alpha-beta-rho model

Abstract: We propose an efficient and reliable simulation scheme for the stochastic-alpha-beta-rho (SABR) model. The main challenges are sampling the integrated variance conditional on terminal volatility and the terminal forward price conditional on terminal volatility and integrated variance. We approximate the conditional integrated variance by a moment-matched shifted lognormal distribution and the conditional terminal forward price by a constant-elasticity-of-variance (CEV) distribution. Our CEV approximation preserves the martingale condition and precludes arbitrage, which is a key advantage over Islah's approximation used in most SABR simulation schemes in the literature. We then sample the CEV distribution exactly based a shifted-Poisson mixture Gamma random variable. The proposed method avoids costly Laplace inversion and inefficient root-finding procedures. Numerical results show that the scheme is highly accurate, efficient, and stable across a wide range of parameter settings.

Session 2A: Neural Networks (Chair: Yixin Ding)

Talk 1

Name: Daniel Zhang

Title: Tune Small, Train Big: BNN-μP for Bayesian Neural Network Language Models

Abstract: Language models give a probability to every possible next word, but not always a trustworthy one. A Bayesian neural network (BNN) keeps a distribution over its parameters, so disagreement among plausible networks shows what the model does not know, but settings tuned on a small BNN stop working on a larger one. We introduce BNN-μP, a recipe built on the maximal-update parametrization (μP) that sets how the learning rate, the initial values and the strength of the prior change with width. A learning rate tuned once at 38 million parameters carries over to 1.01 billion, where the loss is lower than that of the standard recipe re-tuned at every width, with calibrated uncertainty that flags unfamiliar text.

Talk 2

Name: Yixing Zhao

Title: Neural Regression Trees for Conditional Distribution Learning under Censoring

Abstract: Regression methods are widely used to characterize relationships between covariates and outcomes, but many approaches focus on conditional means or impose restrictive distributional assumptions. We propose a Neural Regression Tree (NRT) for estimating the full conditional distribution of exact or censored continuous outcomes. A neural network maps covariates to split locations and branch probabilities on a complete binary tree. The resulting multiscale structure captures complex features such as multimodality, while recursive branch probabilities define terminal-node masses and a valid continuous conditional distribution and survival function. The framework accommodates uncensored and mixed-censored outcomes within a single model. Simulation studies evaluate its estimation and prediction performance across censoring settings.

Talk 3

Name: Zihao Zhen

Title: Neural Network Surrogates for Gradient-Based Bayesian Inference in Electrochemical Parameter Estimation

Abstract: Electrochemistry studies chemical reactions coupled to electrical signals and has applications from battery design to chemical manufacturing. In it, one fits physical models such as partial differential equations to noisy measurements to recover unknown constants like reaction rates. Despite the field's maturity, this fitting is prone to identifiability issues arising from non-uniqueness in the constant-to-observation mapping and experimental noise. In such cases, full uncertainty quantification is desired but requires repeatedly solving PDEs that are computationally expensive. To overcome this, we train a neural network surrogate to replace the traditional PDE solver and embed it within a gradient-based Markov chain Monte Carlo sampler to compute posterior distributions of the unknown constants. Through extensive benchmarks, we verify the feasibility of this workflow and find the results promising, pointing toward practical, uncertainty-aware parameter estimation in electrochemistry.

Talk 4

Name: Taj Jones-McCormick

Title: Towards Smarter LLM Powered Prediction Systems

Abstract: Recent research has begun incorporating Large Language Models (LLMs) into various predictive modeling tasks. LLMs possess both a plethora of world knowledge from pre-training as well as the ability to receive and use new information in context. This makes them an attractive candidate for next event prediction (Temporal Point Process Forecasting) as they often possess significant prior knowledge on the domain of interest and allow for flexible and non-lossy incorporation of predictive features. In this work we accomplish two main goals. The first is to demonstrate that some of this world knowledge which gives LLMs an advantage over traditional models can be distilled into more lightweight 'custom' traditional models by using an LLM as a model builder. The second is to directly compare various LLM based approaches for predictive modeling in order to understand their various strengths and weaknesses.

Session 2B: Modelling, pricing, and reserving (Chair: Maria Andraos)

Talk 1

Name: Yan Yu

Title: Variable Annuity Pricing with Limited Attention

Abstract: We study the role of limited attention in the pricing of variable annuities. In contrast to the existing literature, which typically assumes continuous monitoring, we introduce an attention friction where policyholders observe their account only intermittently. At each decision date, observation occurs with an exogenous probability p, and surrender decisions are made only when observation takes place. We show that limited attention reduces the frequency of policyholder surrender behaviour, which in turn lowers the fair insurance fee and helps reconcile the gap between standard model-implied fees and those observed in the market. Furthermore, by allowing the attention probability to vary with time, our model generates spike-shaped surrender patterns consistent with the market observation, highlighting limited attention as a potential mechanism behind this empirical feature.

Talk 2

Name: Kathleen Miao

Title: Consensus Copulas: Kullback-Leibler Barycentres for Model Aggregation

Abstract: Copula models are widely used to represent nonlinear and tail dependence in multivariate data, yet selecting a single specification is difficult when several models are empirically plausible. We address this source of model uncertainty by introducing a consensus copula defined as a confidence-weighted Kullback-Leibler barycentre. The proposed construction minimises the aggregate divergence from the candidate models while satisfying the uniform-marginal constraints required of a copula. We characterise the structure of the consensus copula, establish conditions for its existence and uniqueness, and examine how it changes as the model weights vary. We also develop numerical procedures for computing the consensus copula and generating observations from it, and compare its behaviour with that of a traditional mixture-based aggregation.

Talk 3

Name: Tianxing Yan

Title: Event-structured dependence in frequency-severity modelling of storm losses

Abstract: Severe convective storm events produce highly heterogeneous insurance loss outcomes. Some lead to large numbers of relatively small claims, while others generate concentrated catastrophic losses, requiring substantially different operational responses. We develop an event-based frequency-severity model that captures latent storm heterogeneity through correlated storm-level effects that jointly govern claim propensity and the distribution of individual loss amounts. The latent structure induces within-storm dependence across policies and allows the relationship between frequency and severity to differ systematically across storms, providing an interpretable characterization of heterogeneous loss patterns across events. Applying a two-stage estimation procedure to replicated storm loss data, we demonstrate that incorporating heterogeneous frequency-severity dependence enhances insurers' post-event risk assessment.

Talk 4

Name: Zhiyang Cheng

Title: Neural Bayes Estimation for GLM-Based Loss Reserving

Abstract: Predicting outstanding claims is a central task in property and casualty loss reserving. Generalized linear models (GLMs) provide an interpretable framework, but parameter estimation can be challenging when only aggregate loss triangles are available. We propose an amortized neural estimation framework that retains the parameter

structure of a Tweedie compound Poisson GLM while using neural networks for parameter estimation. The networks are trained on simulated loss triangles and can estimate parameters for new triangles without repeated optimization. We consider a multilayer perceptron (MLP) and a DeepSets extension based on replicated simulations. Simulation results show improved estimation accuracy and stability over maximum likelihood estimation. Experiments on real-world data also demonstrate strong predictive performance. The proposed approach offers an efficient alternative for parameter estimation while preserving the interpretability of the underlying reserving model.

Session 3A: Meta-Analysis (Chair: Xamy López Córtes)

Talk 1

Name: Diribsa Tsegaye Bedada

Title: Correcting Verification Bias in Individual Participant Data Meta-Analysis of Diagnostic Test Accuracy Studies: A Joint Likelihood Approach

Abstract: Verification by a reference standard is often incomplete in diagnostic accuracy studies and may depend on index test results and participant characteristics, leading to biased estimates of sensitivity and specificity. We propose a one-stage joint likelihood framework for correcting verification bias in individual participant data meta-analysis of diagnostic test accuracy studies. The framework jointly models disease status, index-test outcomes, and verification while incorporating participant-level covariates and correlated study-specific random effects. For unverified participants, disease status is treated as latent and marginalized in the likelihood. Estimation uses Template Model Builder with automatic differentiation and Laplace approximation. Simulations indicate that verified-only analyses can be substantially biased under selective verification, whereas the proposed method substantially reduces bias and provides reliable estimation of accuracy and between-study heterogeneity.

Talk 2

Name: Carmen Alicia Liang Zhen

Title: Visualizing Joint Ranking Distributions in Network Meta-Analysis with Hammock Plots

Abstract: In network meta-analysis (NMA), treatment ranking summarizes relative effectiveness of competing treatments. Existing graphics focus on marginal ranking distributions or treatment hierarchies with limited insight into their uncertainty. We introduce hammock plots to visualize the joint posterior distribution of treatment rankings. Separate hammock plots show relationships between (a) adjacent treatments or (b) adjacent ranks. Using highlighting, we show how to assess posterior support for selected treatment hierarchies, evaluate hierarchies identified by ranking metrics against the posterior distribution, and examine uncertainty within selected regions of a hierarchy. We apply these methods to an NMA of eight thrombolytic treatments, revealing patterns in their posterior ranking distribution and the location of uncertainty within the hierarchy that were hard to identify using existing ranking graphics. The methods are implemented in a web app (https://hammock-nma.streamlit.app/).

Talk 3

Name: Olana Angesa Dabi

Title: A Meta-Analytic Model for Diagnostic Tests with Repeated Measurements

Abstract: Meta-analyses of diagnostic test accuracy (DTA) often ignore within-study replicate testing, which can bias sensitivity (Se) and specificity (Sp) estimates. We propose a bivariate repeated-measures GLMM (BRGLMM), fitted via TMB with study- and replicate-level random effects, and compare it to a common strategy that collapses replicates (mean) before fitting a standard bivariate GLMM (BGLMM). A simulation study will be conducted to evaluate the proposed model through bias, RMSE, CI width, and coverage for Se, Sp, and variance-covariance parameters. BRGLMM showed lower bias and RMSE than the collapsing method, especially as the number of replicates increased, while maintaining near-nominal coverage; collapsing approaches underestimated fixed effect parameters and produced less accurate intervals, especially when there was sparse data. Results support using repeated-measures models for DTA meta-analysis when replicated data are available.

Session 3B: Risk Measures (Chair: Qianyu Chen)

Talk 1

Name: Muqiao Huang

Title: Lambda Expected Shortfall

Abstract: Lambda Value-at-Risk extends Value-at-Risk by allowing the confidence level to depend on the loss level, offering greater flexibility in risk assessment. Motivated by the classical relationship between Value-at-Risk and Expected Shortfall, we introduce Lambda Expected Shortfall as a flexible generalization of Expected Shortfall and a natural counterpart to Lambda Value-at-Risk. We show that Lambda Expected Shortfall is the smallest law-invariant and quasi-convex risk measure dominating Lambda Value-at-Risk under mild conditions. We further study its structural properties, derive a dual representation, and develop a Rockafellar–Uryasev-type formulation that connects the new risk measure to optimization problems.

Talk 2

Name: Jingcheng Yu

Title: E-backtesting Expected Shortfall without Value-at-Risk forecast

Abstract: Expected Shortfall (ES) is central to financial risk management, but the standard ES e-statistic requires both an ES forecast and its associated Value-at-Risk (VaR) forecast. We study sequential backtesting when only ES is reported.

Talk 3

Name: Chenxi Xia

Title: Rating Criteria, Risk Measures, and Risk Consistency

Abstract: Categorical ratings are widely used by investors to screen financial products. We model rating criteria as discrete transforms of risk measures and study three notions of risk consistency that promote prudent investment decisions. These notions concern risk aversion, the asset pooling effect, and the benefit of portfolio diversification, and can be formulated under multiple probability measures. For the monetary and positively homogeneous (MPH) and state-dependent expected utility (SEU) classes, we establish how these properties translate between rating criteria and their underlying risk measures and identify their hierarchical relationships. These results yield full characterizations of risk-consistent criteria in the Choquet subclass of MPH and in the SEU class. Case studies on collateralized loan obligations and investment funds illustrate how risk-consistent scenario-based criteria can provide alternatives to the probability of default and expected loss criteria used in practice.

Day 2

Keynote (Chair: Laura Bumbulis)

Name: Dr. Erica Moodie

Title: Leveraging partial adherence to improve estimation of individualized treatments

Abstract: Precision medicine is focused on optimizing sequential decision-making; in the statistical literature, this is formalized through dynamic treatment regimes. A flexible approach to estimation relies on inverse probability weighting. However, standard inverse weighted estimators can exhibit high variability. A key reason for this inefficiency arises because of the estimator's requirement of strict alignment between observed and candidate treatment strategies which results in discarding substantial information from individuals with only slight deviations from the candidate strategies. In this talk, I will introduce two novel approaches that relax the strict alignment, resulting in estimators that are asymptotically consistent with the target estimand and offer good finite sample properties.

Session 4A: Machine Learning (Chair: Diribsa Bedada)

Talk 1

Name: Yi Chang

Title: Semiparametric Positive-Unlabeled Learning for Optimal Combination of Biomarkers

Abstract: We develop an optimal linear combination of multiple biomarkers for disease classification using positive-unlabeled data, where only confirmed cases or controls are observed and the remaining subjects are unlabeled. Because the case and control biomarker distributions are not identifiable without additional assumptions, we impose a semiparametric density ratio model that ensures identifiability. We then use empirical likelihood to estimate the optimal combination coefficients and the associated ROC curve, AUC, and Youden index. We establish asymptotic normality of the estimators and propose an efficient, numerically stable algorithm. Simulations show satisfactory finite-sample performance, and an asthma data analysis demonstrates improved diagnostic accuracy when either case or control labels are unavailable.

Talk 2

Name: Nafisat Ibrahim

Title: Identifying Factors Influencing TB Knowledge Among Nigerian Women Using Supervised Learning

Abstract: Tuberculosis (TB) knowledge gaps contribute to delayed care-seeking in Nigeria. This study applies supervised machine learning to the 2024 Nigeria Demographic and Health Survey to identify factors influencing TB knowledge among women aged 15-49. A composite score was built across four domains (awareness, transmission, symptom recognition, curability) using DHS-8's TB module and validated through known-groups analysis. Logistic Regression, Random Forest, and XGBoost models are trained to predict this outcome, accounting for the survey's complex design through survey weights and cluster-aware cross-validation. Model interpretability is assessed using SHapley Additive exPlanations (SHAP) to identify the key drivers of TB knowledge. This work provides insights into the factors shaping TB knowledge in Nigeria and the findings aim to help policymakers better target TB education interventions to the groups most in need.

Talk 3

Name: Elijah French

Title: Bayesian Learning in Graphon Games

Abstract: We study a game with a continuum of agents who repeatedly interact without exact information about their connections. A variation of a canonical linear-quadratic payoff with interactions defined by a graphon, a limiting object for graphs, is used at each stage. We assume the graphon has known eigenfunctions but unknown eigenvalues. Agents hold a common prior belief over these values which is sequentially updated using noisy information from the stage games under the true values, creating a sequence of posteriors. The myopic stage optimal action is shown to be the unique sub-game perfect Nash equilibrium as a result of atomization in these updates. Equilibrium play is then characterized, at each stage of play, solely by current beliefs. In a truncated Gaussian prior and Gaussian agent noise setting, convergence of the posteriors is shown. This is compared with the standard consistency present in the single repetition game where samples are accumulated across agents rather than stages.

Session 4B: Portfolio Selection & Optimal Reinsurance (Chair: Maria Andraos)

Talk 1

Name: Yuqi Jing

Title: Adaptive Combination of Portfolio Rules

Abstract: We combine two fully-invested portfolio rules and estimate the combination coefficient via a mean-variance ratio with conditioning information: the conditional mean of the rules' return difference over its conditional variance. We model the variance as a predictor-driven level times a stationary short-run component, estimated by quasi-likelihood, while holding the hard-to-estimate risk-premium numerator fixed and calibrating it from past excess returns; the coefficient is bounded to [0,1]. Across seven Fama-French datasets (\f{EvalStart}–\f{EvalEnd}, net of costs), combining 1/N with sample or nonlinearly-shrunk GMV improves Sharpe ratio and certainty-equivalent returns over both constituents, with significant gains over 1/N in 29/42 cases. Across thirty long-only vs. unconstrained rule pairs, the combination beats the unconstrained rule on all three measures, significant in ~40% of cases.

Talk 2

Name: Yunqi Zhang

Title: Optimal Reinsurance to Minimize the Drawdown Risk: A Probabilistic Approximation and Policy Iteration (PAPI) Approach

Abstract: This paper studies optimal reinsurance that maximizes the probability of reaching a prescribed surplus target before drawdown under a diffusion-perturbed Cramér–Lundberg model. The presence of claim jumps, together with the infinite-dimensional control space induced by general reinsurance contracts, makes closed-form solutions generally unattainable and standard numerical methods difficult to apply. We develop a probabilistic approximation and policy iteration (PAPI) approach that retains the claim-jump structure while allowing optimization over general state-dependent retention functions. The numerical results provide evidence of convergence and show that the optimal retention strategy under the original jump model depends on both the current surplus and its running maximum and generally differs from conventional forms, such as proportional or excess-of-loss reinsurance.

Talk 3

Name: Brian Ceco

Title: Active Portfolio Management in Concentrated Equity Markets

Abstract: The equal-weighted portfolio has historically proven to be a difficult passive strategy to outperform, delivering higher returns than the capitalization-weighted "market" benchmark across many markets and periods. Stochastic portfolio theory (SPT) reveals that this relative performance is regime dependent, with the equal-weighted portfolio underperforming during periods of increasing market concentration and high correlations, particularly market bubbles. These observations have motivated us to formulate and solve a stochastic control problem in which an investor actively allocates between the equal-weighted and market portfolios. We derive the optimal strategy via a linear forward–backward SDE and obtain an explicit "aiming in front of a moving target'' representation of the optimal trading rate. Extensive experiments on S&P 500 data show that these optimal strategies outperform the market, even as the equal-weighted portfolio has underperformed.

Session 5A: E-Values (Chair: Yixin Ding)

Talk 1

Name: Guneet Dhillon

Title: E-Scores for (In)Correctness Assessment of Generative Model Outputs

Abstract: While generative models, especially large language models (LLMs), are ubiquitous in today's world, principled mechanisms to assess their (in)correctness are limited. Using the conformal prediction framework, previous works construct sets of LLM responses where the probability of including an incorrect response, or error, is capped at a user-defined tolerance level. However, since these methods are based on p-values, they are susceptible to p-hacking, i.e., choosing the tolerance level post-hoc can invalidate the guarantees. We therefore leverage e-values to complement generative model outputs with e-scores as measures of incorrectness. In addition to achieving the guarantees as before, e-scores further provide users with the flexibility of choosing data-dependent tolerance levels while upper bounding size distortion, a post-hoc notion of error.

Talk 2

Name: Liulei Sun

Title: Admissibility and Complete Classes for False Discovery Rate Control with E-values

Abstract: The false discovery rate (FDR) is the most widely used error metric in modern multiple testing. We provide a comprehensive analysis of the admissibility of e-value-based procedures with FDR control. We consider both simultaneous and point procedures and introduce strong and weak notions of dominance. Every simultaneous procedure is strongly, and hence weakly, dominated by an admissible weighted-mean closed eBH procedure, and thus weighted-mean closed eBH procedures form a complete class. Every constant-free weighted-mean closed eBH procedure is admissible at every level, and we propose two ways for choosing constant-free weights based on side information. Within the symmetric class, the mean closed eBH procedure is the largest element iff the FDR level is small enough. We also obtain results on the admissibility of symmetric weighted-mean closed eBH procedures with constant terms. Point e-testing procedures have a parallel theory of admissibility.

Talk 3

Name: Chenxin Lyu

Title: The Simes inequality and false discovery rate under Gaussian dependence

Abstract: We study the false discovery rate (FDR) of the Benjamini–Hochberg procedure for Gaussian-dependent p-values. Two questions are whether the global-null rejection probability is close to the nominal level q and whether the global null is least favorable. For each fixed dimension, the largest global-null rejection probability is asymptotic to q, uniformly over all correlation matrices, including singular ones. We also give a dimension-free finite-level bound for nonpositive correlations. In the mixed-null problem with monotone alternative margins, the global null is least favorable in dimensions at most three, but not in dimension eight at sufficiently small levels. To understand the latter phenomenon, we first show that a fixed positive excess of FDR/q above m_0/K at small levels requires near-perfect absolute correlation involving a true null. We then solve the rank-one problem: for $0 < q < 1/2$, the maximum FDR is q through dimension seven and $21q/20$ in dimension eight, and the sharp factor over all dimensions is $2W(1)$. These rank-one conclusions also hold when only the true-null block has rank one. The proofs combine Gaussian projection and tail bounds with distribution-free inequalities for two counter-monotone blocks.

Session 5B: Causal Inference and Precision Medicine (Chair: Madison Cranstone)

Talk 1

Name: Gujie Fu

Title: Causal Inference via Balanced Representation Learning

Abstract: We propose an autoencoder framework with integral probability metrics to balance high-dimensional covariates in a low-dimensional latent space. Individual treatment effects are estimated by two-arm nearest-neighbour smoothing. An augmented inverse probability weighting estimator of the average treatment effect combines latent-space outcome estimates with a propensity score fitted on the original covariates. Under identification and regularity conditions, we establish pointwise consistency and asymptotic normality of the nearest-neighbour estimators, and asymptotic normality with consistent variance estimation for the average effect estimator. Staged hyperparameter selection balances reconstruction fidelity and distributional balance. Simulations and an analysis of Infant Health and Development Program data assess finite-sample performance.

Talk 2

Name: Xamy López Córtes

Title: Distributional Effects for Compliers under Instrumental Variables with Treatment Misclassification

Abstract: Instrumental variable (IV) methods are useful for identifying causal effects when treatment assignment is not randomized, but treatment misclassification can weaken the instrument, producing unstable estimators. This project extends a nonparametric plug-in IV framework to account for treatment misclassification using a mixture representation of the latent compliance classes. We theoretically show that treatment misclassification weakens the instrumental variable. Additionally, we develop a four-group bootstrap procedure for testing equality, first-order stochastic dominance, and second-order stochastic dominance of complier outcome distributions. Simulation results show that the proposed mixture approach improves the power of the test under weak instruments when the outcome distributions differ. An application to the Oregon Health Insurance Experiment illustrates the improvements of accounting for treatment misclassification with our proposed mixture approach.