

# Fourth Workshop on Recent Trends in Machine Learning and Risk Management

November 3-4, 2025



Stony Brook **University**

## Table of Content

Information.....	3
Parking on Campus.....	4
Workshop Schedule.....	5
Abstracts.....	7

## Information

### **Fourth Workshop on Recent Trends in Machine Learning and Risk Management**

**Date:** November 3-4, 2025 (Monday-Tuesday)

**Venue:** Math Tower room S-240, Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY

**Participation:** In-Person

**Website:** <https://uryasev.ams.stonybrook.edu/32469-2/>

### **Workshop Outline**

This workshop will bring together top academics and practitioners in the fields of Reinforcement Learning, Stochastic Optimization, Statistical Learning, and Machine Learning. The objective is to exchange recent ideas and approaches, focusing on numerical methods and applications in:

- Risk Management
- Portfolio Optimization
- High-Frequency Trading
- Derivatives Pricing and Hedging

### **Hotel Information**

- **Block of Rooms:** A block of rooms is reserved for conference participants at the Hilton Garden Inn Stony Brook at a rate of \$179 plus state tax (breakfast is **not** included). Please reserve rooms by Oct. 12, 2025.

### **Organizers:**

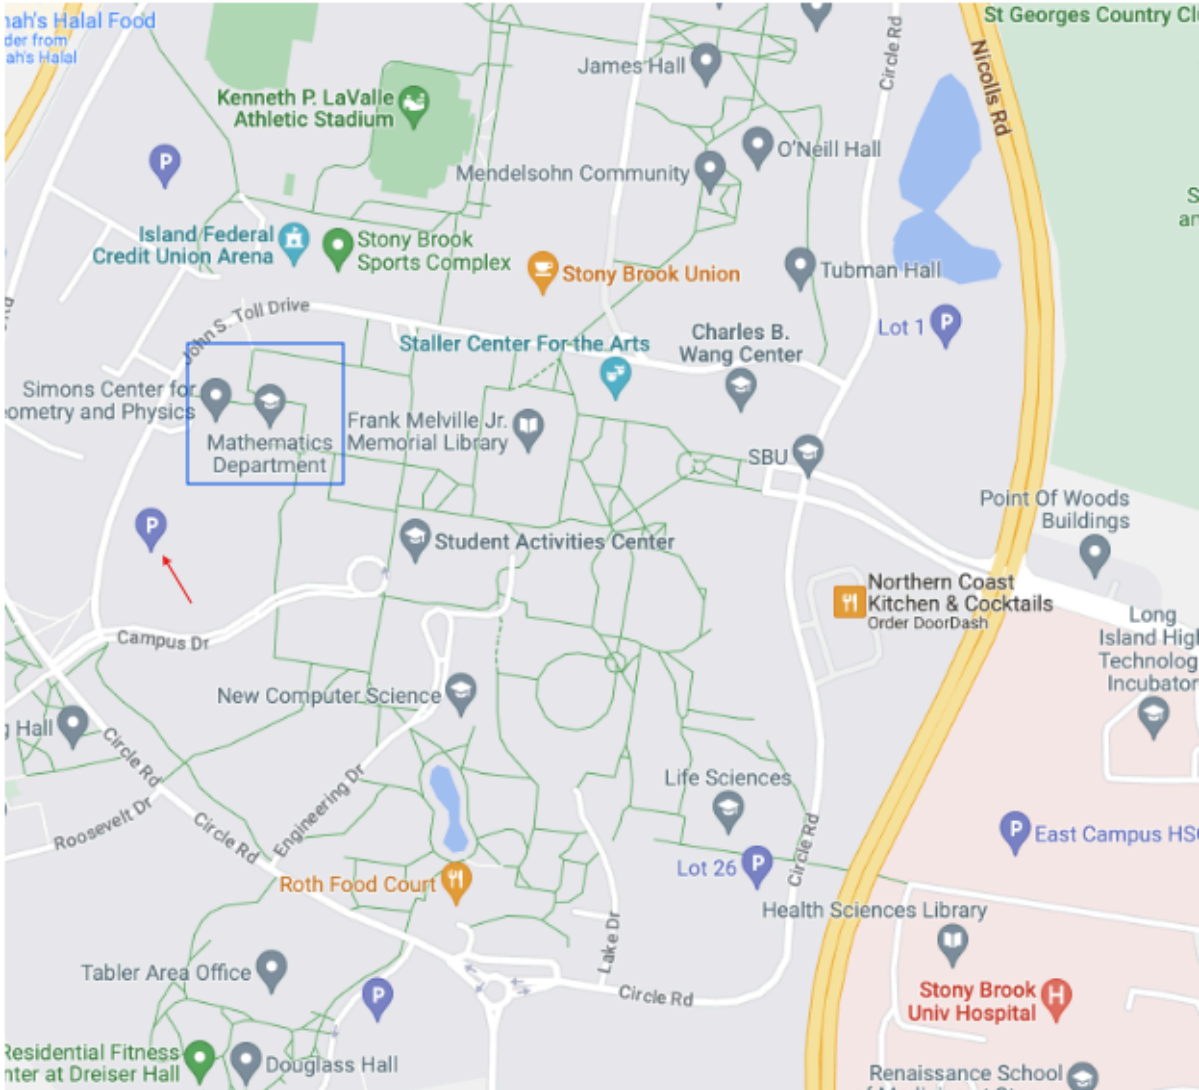
Pawel Polak, Stan Uryasev, Zilin Chen Stony Brook University

### **Sponsors:**

Applied Mathematics and Statistics Dept. at the Stony Brook University

# Parking on Campus

There are a large number of paid parking spots on campus. The closest one to the venue is the large parking lot on John S. Toll Drive (see the map below). One day of parking (7am-7PM) costs \$17.50.



## Workshop Schedule

	Monday	Speaker	Title
	<b>8:30-8:50</b>	Registration & Coffee	
	<b>8:50-9:00</b>	Welcome session	
<b>Session I</b>	<b>9:00-9:30</b>	Stan Uryasev	Risk Quadrangle and Applications in Statistics, Data Mining, and Portfolio Optimization
	<b>9:30-10:00</b>	Ionut Florescu	Applications of DeFi concepts to Traditional Finance
	<b>10:00-10:30</b>	Coffee break	
<b>Session II</b>	<b>10:30-11:00</b>	Ben Feng	How well can machine learn with noisy data? -- Deep Learning for High Dimensional Nested Simulation
	<b>11:00-11:30</b>	Yingying Fan	LLM-Powered Prediction Inference with Online Text Time Series
	<b>11:30-12:00</b>	Akshay Krishnamurthy	The coverage principle in language models: From pre-training to test-time scaling
	<b>12:00-12:30</b>	Ruslan Tepelyan	Modeling 1-Minute VWAP Returns Using the Timing Information in Bloomberg's OHLC Data
	<b>12:45-1:30</b>	Lunch (lunch talk with Robert Frey)	A Study of Equity Market Structure Over Two Centuries
<b>Session III</b>	<b>1:30-2:00</b>	Luca Capriotti	Path Integral Semi-Analytical Methods for General Default Intensity Models
	<b>2:00-2:30</b>	James Glimm Ryan Kaufman	A Time of Day Risk Model
	<b>2:30-3:00</b>	Alexander Melnikov	Market Completions and Machine Learning For Option Pricing
	<b>3:00-3:30</b>	Eugene Feinberg	Sequential Optimization of Dynamically Augmented CVaR
	<b>3:30-4:00</b>	Coffee break	
<b>Session IV</b>	<b>4:00-4:30</b>	Paul Glasserman	Differential ML with a Difference
	<b>4:30-5:00</b>	Artem Prokhorov	Vector Dependence and Optimal Transport for Measuring Systemic Financial Risk
	<b>5:00-5:30</b>	Pawel Polak	Scaling Conditional Autoencoders for Portfolio Optimization via Uncertainty-Aware Factor Selection
	<b>5:30-7:00</b>	Reception	

	Tuesday	Speaker	Title
	<b>8:30-9:00</b>	Registration & Coffee	
<b>Session I</b>	<b>9:00-9:30</b>	Charalampos Markakis	Discontinuous Collocation Methods for Option Pricing
	<b>9:30-10:00</b>	Farid Aitsahlia	Liquidity and Risk Management: Crossing Networks Impact
	<b>10:00-10:30</b>	Martin Larsson	Testing by betting: the numeraire e-variable and reverse information projection
	<b>10:30-11:00</b>	Coffee break	
<b>Session II</b>	<b>11:00-11:30</b>	Fabio Mercurio	Machine Learning for Interest Rates: Using Auto-Encoders for the Risk-Neutral Modeling of Yield Curves
	<b>11:30-12:00</b>	Cao Zheng	Hyped Log-Periodic Power Law Model for Financial Bubbles
	<b>12:00-12:30</b>	Michael Zabarankin Volodymyr Volchenko	Reaching Market Equilibrium in the Framework of Mean-deviation Analysis
	<b>12:30-1:30</b>	Lunch	
	<b>1:45- 16:30</b>	Cultural program	Walk in the Avalon Park (generous donation of Simon's family)

## Abstracts

### Monday November 3, Session I

**Title:** Risk Quadrangle and Applications in Statistics, Data Mining, and Portfolio Optimization

**Presenter:** Stan Uryasev

**Abstract:** The paper reviews recent results obtained with the Risk Quadrangle Framework. A quadrangle includes four functions quantifying uncertainty: Risk, Deviation, Regret, and Error. These functions are interconnected with one more function called Statistic. We consider several quadrangles: Expectile, Superquantile Norm, Symmetric Quantile Average, f-Divergence. The quadrangle framework results in many new analytical results. For instance, we show that Support Vector Regression is an asymptotically unbiased estimator of the average of two symmetric conditional quantiles and that Conditional Value-at-Risk, Expectile, and Omega portfolio optimization are equivalent.

This is a joint paper with Anton Malandii.

**Title:** Applications of DeFi concepts to Traditional Finance

**Presenter:** Ionut Florescu

**Abstract:** We will discuss several applications to traditional finance which were born from studying Automated Market Makers - a DeFi concept. We will discuss the central concept of the Liquidity AMM which generalizes the existing various implementations. Then we will present applications such as Target Weighted ETF's, decentralized trading of liquid assets using Frequent Batch Auctions and potentially others.

**Bio:** Ionut Florescu is a Research Professor in Financial Engineering at Stevens Institute of Technology. He serves as Director of the Hanlon Financial Systems Laboratories and Director of the Financial Analytics Master program at Stevens Institute of Technology. He has a Ph.D. in Statistics from Purdue University and a bachelor's in mathematics from the University of Bucharest, Romania. Florescu research spans many areas of probability, statistics, and stochastic processes and their application to other areas. He authored 4 books and edited 4 other books. His work had been supported by several grants from NSF, and other agencies, including private companies totaling over \$700k. He published over 50 peer research journal articles, and many conference papers. He owns a patent in computer vision and serves as associate editor in the editorial board for 2 journals. He served as co-Editor in Chief for the High Frequency Journal and has organized a series of 9 conferences in High Frequency Finance. He served as chairman or co-chairman of the dissertation committee for 13 Ph.D. students as well as many Master dissertation committees. His current research is focused on the use of Automated Market Makers (AMM's) in traditional Finance, as well as investigating Large Language Model (LLM) implementations.

## Monday November 3, Session II

**Title:** How well can machine learn with nosy data? -- Deep Learning for High Dimensional Nested Simulation

**Presenter:** Ben Feng

**Abstract:** Deep learning models have achieved remarkable success across a wide range of applications, from autonomous driving to generative AI. Their potential in financial and actuarial fields has captured significant interest among researchers, practitioners, and regulators. However, the lack of transparency and interpretability in these models has raised concerns about their reliability and resilience, which are critical factors for ensuring financial stability and fulfilling insurance obligations. In this study, we use stochastic simulation as a controlled data generation framework to explore fundamental questions about deep learning models, such as "How effectively do deep learning models learn from noisy data?" Our findings reveal intriguing insights into their behavior and capabilities. Building on these insights, we introduce an innovative nested simulation procedure that employs deep learning models as proxies to efficiently estimate tail risk measures for variable annuity hedging errors. This approach not only optimizes simulation resources by focusing on tail scenarios but also maintains transparency in the estimation process.

**Bio:** Ben Mingbin Feng is an Associate Professor and Director of the Master of Actuarial Science Program at the University of Waterloo. He holds a PhD in Industrial Engineering and Management Sciences from Northwestern University. He is an Associate of the Society of Actuaries (ASA) and a Certified Analytics Professional (CAP). His research focuses on quantitative risk management, financial engineering, Monte Carlo simulation, and nonlinear optimization, with particular interest in efficient simulation algorithms for risk measurement and AI applications in actuarial science.

**Title:** LLM-Powered Prediction Inference with Online Text Time Series

**Presenter:** Yingying Fan

**Abstract:** Time series prediction inference is an important yet challenging task in economics and business, where existing approaches often rely on low-frequency, survey-based data. With the recent advances of large language models (LLMs), there is growing potential to leverage high-frequency online text data for improved time series prediction, an area still largely unexplored. This paper proposes LLM-TS, an LLM-based approach for time series prediction inference incorporating online text data. The LLM-TS is based on a joint time series framework that combines survey-based low-frequency data with LLM-generated high-frequency surrogates. The framework relies only on an error correlation assumption, combining a text-embedding-augmented ARX model for the observed gold-standard measurements with a VARX model for the LLM-generated surrogates. LLM-TS employs LLMs such as ChatGPT and the trained BERT models to construct LLM surrogates. Online text embeddings are extracted via LDA and BERT. We establish the asymptotic properties of the method and provide two forms of constructed prediction intervals. To demonstrate the practical power of LLM-TS, we apply it to a critical real-world example: inflation forecast. We collect a large set of high-frequency online texts from a widely used Chinese social media platform and employ LLMs to construct inflation labels for posts that are related to inflation. The finite-sample performance and practical advantages of LLM-TS are illustrated through simulations and this noisy real data example, highlighting its potential to improve time series prediction in economic applications.

**Bio:** Yingying Fan is Associate Dean for the PhD Program, Centennial Chair in Business Administration, Professor in Data Sciences and Operations Department of the Marshall School of Business at the University of Southern California. Her research interests include statistics, data science, artificial intelligence, machine learning, and economics and business applications. Her latest works have focused on inference for AI models, networks, and texts empowered by some most recent developments in learning theory and random matrix theory. She has served as a Co-Editor of *Journal of Business & Economic Statistics* (2024-2027), IMS-CUP Coordinating Editor (2024-2027), and IMS Editor of *Statistics Surveys* (2023-2025).

**Title:** The coverage principle in language models: From pre-training to test-time scaling

**Presenter:** Akshay Krishnamurthy

**Abstract:** Test-time compute has emerged as a new axis for scaling language model capabilities, yet we lack a principled understanding of this paradigm. What are the right algorithms and trade-offs for test-time scaling? What properties of the pre-trained model enable it? And can we better align pre-training recipes for test-time success?

This talk addresses these questions through a unified lens of coverage. We first show that test-time scaling strategies like best-of-N sampling succeed if and only if the pre-trained model has coverage over high-quality responses. We then demonstrate that coverage, and hence best-of-N performance, can be improved through deliberate exploration, either purely at test time or via RL-style post-training. Finally, we ask why pre-training via next-token prediction yields models with good coverage in the first place. We uncover a rich theoretical landscape driven by an implicit bias of the next-token prediction objective, while also identifying a fundamental misalignment between next-token prediction and coverage, raising the possibility of future algorithmic innovations.

**Bio:** Akshay Krishnamurthy is a senior principal research manager at Microsoft Research, New York City. Previously, he spent two years as an assistant professor in the College of Information and Computer Sciences at the University of Massachusetts, Amherst and a year as a postdoctoral researcher at Microsoft Research, NYC. Before that, he completed his PhD in the Computer Science Department at Carnegie Mellon University. He is broadly interested in foundational aspects of machine learning with a focus on interactive decision making, reinforcement learning, and, more recently, language modeling and generative AI.

**Title:** Modeling 1-Minute VWAP Returns Using the Timing Information in Bloomberg's OHLC Data

**Presenter:** Ruslan Tepelyan

**Abstract:** Open, High, Low, Close (OHLC) bar data is a widely used format for representing financial asset prices over time due to its balance of simplicity and informativeness. Bloomberg has recently introduced a new bar data product that includes additional timing information - specifically, the timestamps of the open, high, low, and close prices within each bar. In this presentation, we investigate the impact of incorporating this timing data into models for predicting the changes in one-minute volume-weighted average price (VWAP). First, we introduce a high-level concept called timing surprise and show that such surprises are correlated with changes in the direction of VWAP returns, motivating a deeper investigation of timing data. Second, we present experimental results that show that including timing features in ML models consistently improves predictive performance of VWAP returns across multiple architectures. We observe gains across several key metrics, including log-likelihood, mean squared error (MSE), R squared, conditional variance estimation, and directional accuracy.

**Bio:** Ruslan Tepelyan is a quant researcher with the Quant Research team in the Office of the CTO at Bloomberg. He has worked as an engineer, equity trader, data scientist, and product manager, but he is mainly interested in applications of machine learning techniques in finance. Most recently, he has been working on projects using ML for time series analysis and evaluating ML-based approaches to core finance optimization problems. He holds a bachelor's degree in electrical and computer engineering from Cooper Union and an MBA from Columbia University.

## Monday November 3, Lunch Talk

**Title:** A Study of Equity Market Structure Over Two Centuries

**Presenter:** Robert Frey

**Abstract:** A detailed examination of equity market performance over the past two centuries reveals a remarkable—and often underappreciated—consistency in the duration and magnitude of both up and down markets. Despite profound changes in market mechanisms, regulation, technology, and global participation, the statistical and structural characteristics of market expansions and contractions exhibit enduring regularities.

These findings suggest that key aspects of market dynamics—particularly those governing the onset, duration, and recovery of major drawdowns—remain poorly understood. The persistence of these structures raises important questions about the assumptions underlying long-term investment strategies and their exposure to large tail risks.

Moreover, evidence indicates that prevailing risk modeling techniques, even those employing more representative distributional assumptions, may systematically underestimate both the likelihood and magnitude of market declines. Understanding these persistent patterns is therefore essential for improving portfolio construction, stress testing, and the broader theoretical framework of market behavior.

## Monday November 3, Session III

**Title:** Path Integral Semi-Analytical Methods for General Default Intensity Models

**Presenter:** Luca Capriotti

**Abstract:** Using the path-integral formalism, we develop an accurate and easy-to-compute semi-analytical approximation for a general class of default intensity models. We illustrate the accuracy of the method by presenting results for the Black-Karasinski model for which the proposed approximation provides remarkably accurate results, even in regimes of high volatility and multi-year time horizons. The accuracy and the computational efficiency of the proposed approximation makes it a viable alternative to fully numerical schemes for a variety of applications in econometrics and derivatives pricing, including the computation of XVA for credit products. As a practical example, we consider the pricing of a quanto Credit Default Swap (CDS) under stochastic intensity of default and an FX devaluation model.

**Bio:** Luca is a Managing Director at UBS, where he heads up the quant team for Non-Core and Legacy (ex-Credit Suisse). Luca is also an Adjunct Professor at the Maths and IOR Departments at Columbia University. He was previously Visiting Professor at the Department of Mathematics of University College London. Prior to working in Finance, Luca was a researcher at the Kavli Institute for Theoretical Physics, Santa Barbara, California, working in High-Temperature Superconductivity and Quantum Monte Carlo methods for Condensed Matter systems.

Luca holds an M.S. cum laude in Physics from the University of Florence, and an M.Phil. and a Ph.D. cum laude in Condensed Matter Theory, from the International School for Advanced Studies, Trieste

**Title:** A Time of Day Risk Model

**Presenter:** James Glimm, Ryan Kaufman

**Abstract:** A risk model sensitive to time of day is proposed. For thirty minute trading with large volume assets, the model is confirmed. For 5 minute data sensitivity to time of day as well as trading volume are demonstrated. Several models will be shown to attempt to capture intraday volatility in thinly traded assets and trading times to model this effect.

**Title:** Market Completions and Machine Learning For Option Pricing

**Presenter:** Alexander Melnikov

**Abstract:** In a complete market, every contingent claim can be replicated by creating a portfolio of financial instruments that has the same pay-off as the contingent claim. Furthermore, the arbitrage-free price for every contingent claim is unique. On the other hand, for an incomplete market, there exist contingent claims whose pay-off cannot be achieved exactly by creating an investment portfolio with the available assets. A key characteristic of incomplete markets is that the arbitrage-free price of a contingent claim is not unique, but rather belongs to an interval of non-arbitrage prices. In this paper, we combine computational techniques and machine learning methods to study the problem of market completion for contingent claims in a one-factor jump-diffusion model. A neural-network-based approach is used to improve parameter estimation and adapt pricing dynamics to observed market behaviour. In addition, utility-based indifference pricing is examined as a complementary valuation method. The proposed framework is then applied to equity-linked insurance contracts, demonstrating how financial and actuarial risks can be jointly addressed within a market completion setting.

**Bio:** Alexander Melnikov is a Professor at the University of Alberta (Edmonton, Canada) working in stochastic analysis and its numerous applications in finance, statistics, risk management and insurance. He is the author of nine books and about one hundred fifty research papers in academic journals and venues. He is a Fellow of the Russian Academy of Natural Sciences, a recipient of the Leontiev medal of this academy and the McCalla Professorship of the University of Alberta. In addition to his academic engagements, he held several senior positions in business and professional organizations.

**Title:** Sequential Optimization of Dynamically Augmented CVaR

**Presenter:** Eugene Feinberg

**Abstract:** This talk deals with optimization of the Conditional Value-at-Risk (CVaR) of total discounted costs for Markov Decision Processes (MDPs) with finite state and action sets. It introduces the Dynamically augmented CVaR (DCVaR) risk measure and provides an algorithm for its optimization. This paper investigates a specially defined Robust MDP (RMDP), in which the state space is augmented with the tail risk level. This RMDP, which we call the Dynamically augmented RMDP (DRMDP), was introduced to the literature for calculations of optimal CVaR values by value iteration, but, as was understood later, these value iterations compute lower bounds of minimal static CVaRs. DCVaR is defined as a time consistent version of the static CVaR, and it is a lower bound of the static CVaR. It also can be considered as a dynamic version of the nested CVaR. The correctness of the provided algorithm is proved by studying a special mass transfer problem. We also provide the results on RMDPs needed for this paper.

## Monday November 3, Session IV

**Title:** Differential ML with a Difference

**Presenter:** Paul Glasserman

**Abstract:** Differential ML (Huge and Savine 2020) is a technique for training neural networks to compute option prices in simulation-based models. It uses price sensitivities calculated through pathwise adjoint differentiation to reduce pricing errors. However, for options with discontinuous payoffs, such as digital or barrier options, the pathwise sensitivities are biased, and incorporating them into the loss function can magnify errors. We consider alternative methods for estimating sensitivities and find that they can substantially reduce test errors in prices and in their sensitivities. This is joint work with Siddharth Karmarkar.

**Bio:** Paul Glasserman is the Jack R. Anderson Professor of Business at Columbia Business School, and currently serves as chair of the Financial and Business Analytics Center within Columbia University's Data Science Institute. He is a past recipient of the IAQF Financial Engineer of the Year Award and Risk magazine's Quant of the Year Award. He received the INFORMS Lanchester Prize for his book, Monte Carlo Methods in Financial Engineering.

**Title:** Vector Dependence and Optimal Transport for Measuring Systemic Financial Risk

**Presenter:** Artem Prokhorov

**Abstract:** The talk is about a follow up to our 2025 paper (Medovikov and Prokhorov, Journal of Financial Econometrics, 2015) which introduced a vector version of Hoeffding's  $\Phi^2$ , a nonparametric and copula-based statistic that isolates dependence between vector blocks while remaining invariant to within-block correlation. The follow up generalizes this framework to systemic-risk measurement through Kendall vector copulas and an entropic-optimal-transport index  $\tau$ , which quantifies tail co-monotonicity across global financial regions. Empirical results show  $\tau$  anticipates major crises and improves risk prediction relative to existing metrics. This helps establish a useful connection between optimal transport, vector copula and measures of global systemic vulnerability.

**Title:** Scaling Conditional Autoencoders for Portfolio Optimization via Uncertainty-Aware Factor Selection

**Presenter:** Pawel Polak

**Abstract:** Conditional Autoencoders (CAEs) offer a flexible, interpretable approach for estimating latent asset-pricing factors from firm characteristics. However, existing studies usually limit the latent factor dimension to around  $K = 5$  due to concerns that larger  $K$  can degrade performance. To overcome this challenge, we propose a scalable framework that couples a high-dimensional CAE with an uncertainty-aware factor selection procedure. We employ three models for quantile prediction: zero-shot Chronos, a pretrained time-series foundation model (ZS-Chronos), gradient-boosted quantile regression trees using XGBoost and RAPIDS (Q-Boost), and an I.I.D. bootstrap-based sample mean model (IID-BS). For each model, we rank factors by forecast uncertainty and retain the top- $k$  most predictable factors for portfolio construction, where  $k$  denotes the selected subset of factors. This pruning strategy delivers substantial gains in risk-adjusted performance across all forecasting models. Furthermore, due to each model's uncorrelated predictions, a performance-weighted ensemble consistently outperforms individual models with higher Sharpe, Sortino, and Omega ratios.

## Tuesday November 4, Session I

**Title:** Discontinuous Collocation Methods for Option Pricing

**Presenter:** Charalampos Markakis

**Abstract:** We present a generalization of the Crank–Nicolson scheme, commonly used in finite difference discretizations of the Black–Scholes equation. The method is unconditionally stable and second order in both space and time. Symmetric Padé approximants achieve arbitrarily high temporal order while maintaining unconditional stability. Within a method of lines framework, spatial derivatives are approximated using high-order collocation (finite-difference or pseudospectral) methods. To model non-smooth initial data, such as jump discontinuities in the payoff or evolving solution, we introduce a discontinuous collocation method. Numerical experiments demonstrate its accuracy in capturing discontinuities without loss of stability or order.

**Title:** Liquidity and Risk Management: Crossing Networks Impact

**Presenter:** Farid Aitsahlia

**Abstract:** It is generally acknowledged that spiraling feedback between risk management and liquidity can arise. Tighter risk management leads to more restricted portfolios, resulting in illiquidity due to the need to execute increasingly large transactions. In turn, the price impact of these trades induces further changes in value-at-risk (VaR). As an alternative trading platform with lower price impact risk, *crossing networks* (also known as *dark pools*) have experienced rapid growth over the past few years. Although one of their disadvantages is uncertain execution, empirical evidence suggests that they are used particularly when bid–ask spreads are very large.

In this paper, I present a stochastic control model that determines critical bid and ask price levels beyond which VaR-related transactions are routed through crossing networks that use the midpoints of the National Best Bid and Offer (NBBO) prices. I also show the existence of a critical probability value below which no orders are routed through crossing networks.

**Bio:** Farid AitSahlia is on the faculty of the finance department in the Warrington College of Business at the University of Florida. Prior to re-joining academia, he spent a few years in Silicon Valley, working at HP Labs and startups Financial Engines (now part of Eldman Financial Engines) and DemandTec (now part of IBM). Farid is also the current editor-in-chief of the Journal of Risk.

**Title:** Testing by betting: the numeraire e-variable and reverse information projection

**Presenter:** Martin Larsson

**Abstract:** Testing by betting is a recent approach to sequential (or online) statistical inference drawing inspiration from finance. This approach naturally leads to the concept of an e-variable: a nonnegative sample statistic whose expected value is at most one if a given null hypothesis is true. This approach has been found to produce strong statistical error bounds and high statistical power in a wide range of applications. In this talk I will introduce testing by betting, and discuss how classical ideas from mathematical finance, in particular the numeraire portfolio, enables an optimality theory for e-variables that dramatically generalizes earlier work. Our results also lead to a duality theory which yields the so-called reverse information projection in complete generality. (Joint work with Aaditya Ramdas and Johannes Ruf.)

**Bio:** Martin Larsson is a professor in the Department of Mathematical Sciences at Carnegie Mellon University working in mathematical finance, probability theory, stochastic analysis, and statistics. His research has applications to term structure modeling, stochastic portfolio theory, equilibrium analysis, and the statistical analysis of sequentially observed financial data. Before joining CMU in 2019, he was an Assistant Professor of Mathematical Finance at the Department of Mathematics at ETH Zurich. He holds a PhD in Operations Research and Information Engineering from Cornell University, and was a postdoctoral fellow at the Swiss Finance Institute at EPFL in Lausanne. He represents the Department of Mathematical Sciences on the Steering Committee of the Master of Science in Computational Finance (MSCF) program at CMU.

## Tuesday November 4, Session II

**Title:** Machine Learning for Interest Rates: Using Auto-Encoders for the Risk-Neutral Modeling of Yield Curves

**Presenter:** Fabio Mercurio

**Abstract:** In this talk, we use autoencoders (AE) to capture the historical dependence structure of interest rates and introduce risk-neutral forward rate dynamics that are consistent with a given AE curve manifold. By allowing a small convexity-driven deviation from the AE curve manifold, we derive a risk-neutral modeling framework that is arbitrage-free and incorporates the information built into the AE (low-dimensional) latent space. Finally, we showcase numerical results based on historical market swap data for multiple currencies, visualizing graphically some of our key concepts.

**Bio:** Fabio is global head of Quantitative Analytics at Bloomberg LP, New York. His team is responsible for the research on and implementation of cross-asset models for derivatives pricing, XVA and risk analytics. Fabio is also an adjunct professor at NYU. He has jointly authored the book "Interest rate models: theory and practice" and published extensively in books and international journals, including 21 cutting-edge articles in Risk Magazine.

Fabio is the recipient of the 2020 Risk quant of the year award.

**Title:** Hyped Log-Periodic Power Law Model for Financial Bubbles

**Presenter:** James Zheng Cao

**Abstract:** We propose a novel model, the Hyped Log-Periodic Power Law Model (HLPPL), to the problem of quantifying and detecting financial bubbles, an ever-fascinating one for academics and practitioners alike. Bubble labels are generated using a Log-Periodic Power Law (LPPL) model, sentiment scores, and a hype index we introduced in previous research on NLP forecasting of stock return volatility. Using these tools, a dual-stream transformer model is trained with market data and machine learning methods, resulting in a time series of confidence scores as a Bubble Score. A distinctive feature of our framework is that it captures phases of extreme overpricing and underpricing within a unified structure.

We achieve an average yield of 34.13 percentage annualized return when backtesting U.S. equities during the period 2018 to 2024, while the approach exhibits a remarkable generalization ability across industry sectors. Its conservative bias in predicting bubble periods minimizes false positives, a feature which is especially beneficial for market signaling and decision-making. Overall, this approach utilizes both theoretical and empirical advances for real-time positive and negative bubble identification and measurement with HLPPL signals.

**Bio:** Zheng Cao (James) is a Quant Researcher with the Quant Research team in the Office of the CTO at Bloomberg. His work lies at the intersection of quantitative finance, NLP, and machine learning, with recent projects on hype index, financial bubbles, and changes of probability measure. He holds an M.S. in Financial Mathematics from Johns Hopkins University (supervised by Helyette Geman) and a B.S. with Honors in Mathematics from the University of Washington, Seattle (supervised by Zhen-Qing Chen).

**Title:** Reaching Market Equilibrium in the Framework of Mean-deviation Analysis

**Presenter:** Michael Zabarankin, Volodymyr Volchenko

**Abstract:** Market equilibrium is a fundamental notion in economics and finance — when it exists and how it can be reached are two main questions that have long been studied under various conditions and assumptions. A typical approach is to specify individual investor investment strategy and a mechanism for adjusting asset prices based on asset supply/demand (in a tatonnement process, no trades occur until the market-clearing prices are discovered). Rockafellar, Uryasev, and Zabarankin [Journal of Banking & Finance, vol. 31, 2007] showed that there exists a market equilibrium with investors using a Markowitz-type portfolio problem with different deviation measures. The aim of this work is to investigate mechanisms of how such equilibrium can be reached in the framework of mean-deviation analysis, which includes mean-variance analysis as a particular case.

**Bio:** Volodymyr Volchenko is a doctoral student at Stevens Institute of Technology.