Modeling, Learning and Understanding: Modern Challenges between Financial Mathematics, Financial Technology and Financial Economics

Christoph Frei (University of Alberta)
Ruimeng Hu (University of California, Santa Barbara)
Emma Hubert (Princeton University)
Antonis Papapantoleon (TU Delft | NTU Athens | FORTH)

November 10–15, 2024

1 Overview of the Field

The intersection of financial mathematics, financial technology, and financial economics represents a rapidly evolving landscape where advanced quantitative methods meet practical challenges in finance. This field addresses key issues such as systemic risk, market microstructure, and the integration of new technologies like blockchain and machine learning into financial markets. With the increasing complexity and interconnectivity of global financial systems, as well as the advances in the digitization of financial systems, there is a growing demand for innovative models and algorithms to analyze risks, optimize strategies, and ensure financial stability.

One central focus of the field is the study of systemic risk, which seeks to understand and mitigate cascading failures within interconnected financial institutions. Researchers utilize tools like stochastic network models and mean-field games to analyze these risks, highlighting the importance of mathematical frameworks for effective risk management [1, 2]. Additionally, advancements in computational methods, particularly in machine learning and artificial intelligence, have significantly enhanced financial modelling, enabling more accurate predictions and robust decision-making [5].

Another critical area is the development of financial technologies (FinTech), which has transformed traditional banking, trading, and investment practices. Innovations such as open banking, high-frequency trading, and cryptocurrency markets present both opportunities and challenges, necessitating novel approaches to regulation and risk assessment [3]. As FinTech continues to reshape the financial landscape, interdisciplinary collaboration between mathematicians, computer scientists, and economists becomes increasingly essential.

This workshop provided a platform for researchers to address these pressing issues, fostering collaboration and innovation at the intersection of mathematics and finance.

2 Recent Developments and Open Problems

Recent advancements in financial mathematics and technology have paved the way for significant progress, yet numerous challenges and open questions remain. Key developments include:

Advancements in Systemic Risk and Market Models

- Researchers have developed sophisticated models for systemic risk, including dynamic stochastic networks and game-theoretical frameworks, to analyze financial interdependencies and market stability. These models address real-world concerns, such as the impact of interconnected failures in banking systems [2].
- Advances in portfolio optimization, such as those incorporating path constraints and Wasserstein penalizations, enable more refined approaches to investment strategies, reflecting recent innovations in market theory [8].

Emerging Trends in Financial Technology

- The integration of machine learning and artificial intelligence into financial markets has revolutionized trading strategies and risk management. For instance, reinforcement learning techniques now address complex multi-agent environments, enhancing decision-making capabilities in high-frequency trading and market equilibrium [5].
- Open banking initiatives have introduced dynamic competition between traditional financial institutions and FinTech companies. Models analyzing borrower heterogeneity and competition have provided insights into the implications for consumer welfare [6].

Challenges in Climate Finance and Sustainability

- The environmental impact of financial technologies, such as cryptocurrency mining, has sparked concerns about sustainability. Models assessing the economic and environmental trade-offs of these technologies are essential for balancing innovation and climate goals [3].
- Optimal policy mix frameworks for climate change mitigation are emerging as critical tools for governments seeking to balance economic growth and environmental sustainability. These models address the complexities of implementing effective tax and subsidy systems [4].

Open Problems and Future Directions

- Calibrating high-dimensional models, particularly those integrating machine learning, remains a significant challenge. Ensuring model interpretability and robustness in real-time applications is an ongoing research focus.
- Regulatory frameworks must evolve to address the risks posed by FinTech and high-frequency trading. Questions around systemic stability and equitable access to financial services continue to drive policy-oriented research.
- Ethical concerns, including biases in AI-driven financial models and their implications for underrepresented groups, require further exploration to ensure fair and inclusive financial systems.

By addressing these challenges, the field can continue to advance its theoretical foundations while responding to the practical needs of modern financial systems. The workshop facilitated discussions that not only highlighted recent developments but also charted new directions for research and collaboration.

3 Abstracts

3.1 In-person Presentations

The in-person presentations contained a mix of scientific talks by senior, mid-career and junior researchers, as well as a set of "coupled talks" between senior and junior researchers working on related topics. Nearly all sessions featured at least one talk by a member from an under-represented minority (URM), ensuring diverse

representation across the program. The "coupled talks" sessions were each led by a senior researcher, with at least one out of three session leaders being a URM member. Each session included two presentations from early-career researchers (ECRs), among whom at least one-third were URM members.

Ren Ad (Paris-Dauphine University - PSL)

Continuous-Time Persuasion by Filtering

We frame dynamic persuasion in a partial observation stochastic control game with an ergodic criterion. The Receiver controls the dynamics of a multidimensional unobserved state process. Information is provided to the Receiver through a device designed by the Sender that generates the observation process. We develop this approach in the case where all dynamics are linear and the preferences of the Receiver are linear-quadratic. We prove a verification theorem for the existence and uniqueness of the solution of the HJB equation satisfied by the Receiver's value function. An extension to the case of persuasion of a mean field of interacting Receivers is also provided. We illustrate this approach in two applications: the provision of information to electricity consumers with a smart meter designed by an electricity producer; the information provided by carbon footprint accounting rules to companies engaged in a best- in-class emissions reduction effort. In the first application, we link the benefits of information provision to the mispricing of electricity production. In the latter, we show that when firms declare a high level of best-in-class target, the information provided by stringent accounting rules offsets the Nash equilibrium effect that leads firms to increase pollution to make their target easier to achieve. Joint work with Ofelia Bonesini (LSE), Giorgia Callegaro (Dept. Mathematics of Padova University) and Luciano Campi (Milan University)

Carole Bernard (Grenoble Ecole de Management)

Risk Sharing under Ambiguity

The distribution of future losses related e.g., to climate risk is typically not perfectly known. We investigate how to design an optimal sharing scheme among agents (insurers; reinsurers; countries...). We first derive the optimal risk sharing under mean-variance preferences when there is possibly distributional ambiguity on the risk to be shared. It is shown that proportional risk sharing is always optimal and that the presence of ambiguity does not affect the risk sharing. Several generalizations are investigated. Second we will be discussing risk sharing under ambiguity with distortion risk measures.

Ranu Castaneda (University of Alberta)

A Dynamic Model for Open Banking

Open banking allows customers to share their financial data with third-party providers, such as fintech companies, thereby increasing competition and expanding access to financial services. While this shift promises benefits such as improved product offerings and expanded financial inclusion, it also introduces significant technical, social, and economic risks. We develop a continuous-time model using a search-and-match framework to analyze borrower heterogeneity and competition between banks and fintechs. Our preliminary analysis shows that open banking reshapes financial services by leveling the playing field between banks and fintechs. However, while it enhances competition, it could also over-empower fintechs, potentially reversing the benefits to borrowers. The talk is based on joint work with Christoph Frei.

Stphane Crpey (Universit Paris Cit / LPSM)

Statistical Learning of Value-at-Risk and Expected Shortfall

We propose a non-asymptotic convergence analysis of a two-step approach to learn a conditional value-at-risk (VaR) and a conditional expected shortfall (ES) using Rademacher bounds, in a non-parametric setup allowing for heavy-tails on the financial loss. Our approach for the VaR is extended to the problem of learning at once multiple VaRs corresponding to different quantile levels. This results in efficient learning schemes based on neural network quantile and least-squares regressions. An a posteriori Monte Carlo procedure is introduced to estimate distances to the ground-truth VaR and ES. This is illustrated by numerical experiments in a Student-t toy model and a financial case study where the objective is to learn a dynamic initial margin. Based on joint work with D. Barrera, E. Gobet, H.-D. Nguyen and B. Saadeddine

Tolulope Fadina (University of Illinois Urbana-Champaign)

A Framework for Measures of Risk under Uncertainty

A risk analyst assesses potential financial losses based on multiple sources of information. Often, the assessment does not only depend on the specification of the loss random variable but also on various eco-

nomic scenarios. Motivated by this observation, we design a unified axiomatic framework for risk evaluation principles that quantify jointly a loss random variable and a set of plausible probabilities. We call such an evaluation principle a generalised risk measure. We present a series of relevant theoretical results. The worst-case, coherent, and robust generalised risk measures are characterised via different sets of intuitive axioms. We establish the equivalence between a few natural forms of law-invariance in our framework, and the technical subtlety therein reveals a sharp contrast between our framework and the traditional one. Moreover, we provide some characterization results under strong law invariance.

Dena Firoozi (HEC Montral - Universit de Montral)

An Investment Ranking Contest via Quantilized Mean Field Games

We present a model where a principal aims to support a group of homogeneous agents engaged in a competitive ranking contest over a finite time horizon. In this model, only a select proportion of top-ranking agents will advance to the next round at the end of the time horizon. We conceptualize the competition among the agents as a quantilized mean field game where a Nash equilibrium is sought. Furthermore, the principal seeks to achieve a Stackelberg equilibrium with a representative agent. We examine both equilibria and present some illustrative simulation results for an investment contest. This is joint work with Rinel Foguen Tchuendom and Michle Breton.

Geneviève Gauthier (HEC Montréal)

Enhancing Deep Hedging of Options with Implied Volatility Surface Feedback Information

We present a dynamic hedging scheme for S&P 500 options, where rebalancing decisions are enhanced by integrating information about the implied volatility surface dynamics. The optimal hedging strategy is obtained through a deep policy gradient-type reinforcement learning algorithm, with a novel hybrid neural network architecture improving the training performance. The favorable inclusion of forward-looking information embedded in the volatility surface allows our procedure to outperform several conventional benchmarks such as practitioner and smiled-implied delta hedging procedures, both in simulation and backtesting experiments. Joint work with Pascal François, Frédéric Godin, Carlos Octavio Pérez Mendoza.

Matheus Grasselli (McMaster University)

From Debt Crisis to Financial Crashes (and Back)

In this talk I review a model merging two previously proposed models by Steve Keen, namely a monetary model of debt-deflation and a version with Ponzi destabilization, and recall the equilibrium properties and local stability analysis of the merged model. I then add an auxiliary stochastic model of financial markets based on a jump-diffusion process with endogenous jump intensity. This model captures main characteristics of Hyman Minsky's Financial Instability Hypothesis (FIH), and the Quantitative Theory of Credit (QTC) of Richard Werner, with an asset price bubble fueled by pure speculative credit and market crashes impacting the real economy. I then develop and study the fundamental properties of this extended model, its suitability to explain financial crisis and the relationship between growth and private credit. This is joint work with B. Costa Lima and A. Nguyen-Huu.

Daniel Hernndez (CIMAT)

Portfolio Optimization with Path Constraints

In this talk we will consider a utility maximization problem for an agent who has some model beliefs, according to which the agent will try to maximize his utility, coupled with constraints that are based on model-independent considerations. The basic idea is that, assuming that the agent only observes possible paths according to his beliefs, he will pursue a utility maximization objective, but if his losses reach an unacceptable level (for example, due to off-model market behavior), he must be able to meet a budget constraint in all possible models. Under these modeling assumptions, our goal will be to determine the agents optimal investment strategy when he can take (static) positions in certain options, for example, baskets of call options or other simple derivatives. Joint work with A. Cox.

Camilo Hernandez (USC)

The Mean Field Schrdinger Problem: A Mean Field Control Perspective

The mean field Schrdinger problem (MFSP) is the problem of finding the most likely path of a McKean-Vlasov type particle with constrained initial and final configurations. It was first introduced by Backhoff et al. (2020), who studied its existence and long-time behavior. This talk aims to show how ideas from mean field

control theory allow us to derive new interesting results on the MFSP. In particular, we study its existence, characterization, and the so-called convergence problem. The method rests upon studying suitably penalized problems and stochastic control techniques. This talk is based on a joint work with Ludovic Tangpi.

Ulrich Horst (Humboldt University Berlin)

Optimal Trade Execution under Endogenous Order Flow

We consider an optimal liquidation model in which an investor is required to execute meta-orders during intraday trading periods, and his trading activity triggers child orders and endogenously affects future order flow, both instantaneously and permanently. Under the assumptions of risk neutrality and deterministic constants of the impact parameters, we provide closed-form solutions and illustrate the relationship between trading strategies and feedback effects. The optimal trading strategy is of hyperbolic form if the feedback effect of current trading on future order flow is not too strong. If the feedback effect becomes too dominating, a cyclic strategy with possible beneficial round-trips may emerge. Our results extend to risk-averse investors for which semi-closed form solutions involving inverse Laplace transforms are obtained. We set up an estimation framework so that parameter estimates can be made directly from public data and are consistent with the theoretical model. When implementing our model on 110 NASDAQ stocks, the empirical analysis shows that as the level of endogeneity increases, our strategy provides increasingly better performance than the commonly adopted trading strategy. The empirical analysis also shows that too strong feedback effects do not exist in practice, thus ruling out statistical arbitrage. The talk is based on joint work with Ying Chen and Hai Tran.

Anran Hu (Columbia University)

Optimization and Learning for Mean-Field Games via Occupation Measure

Mean-field games (MFGs) and multi-agent reinforcement learning (MARL) have become essential frameworks for analyzing interactions in large-scale systems. This talk presents recent advancements at the intersection of MFGs and MARL. We begin with a new framework MF-OMO (Mean-Field Occupation Measure Optimization), which reformulates Nash equilibria for discrete-time MFGs as a single optimization problem over occupation measures, offering a fresh characterization that enables the use of standard optimization algorithms to identify multiple equilibria without relying on restrictive assumptions. We also extend these results to continuous-time finite state MFGs. Building on the concept of occupation measures, we then introduce MF-OML (Mean-Field Occupation Measure Learning), the first fully polynomial online RL algorithm capable of finding approximate Nash equilibria in large population games beyond zero-sum and potential games. We establish regret bounds for the N-player games that can be approximated by MFGs under monotonicity conditions. Together, these advancements provide a comprehensive approach to characterizing and solving Nash equilibria in complex multi-agent environments.

Sebastian Jaimungal (U. Toronto)

Kullback-Leibler Barycentre of Stochastic Processes

We consider the problem where an agent aims to combine the views and insights of different experts models. Specifically, each expert proposes a diffusion process over a finite time horizon. The agent then combines the experts models by minimising the weighted Kullback-Leibler divergence to each of the experts models. We show existence and uniqueness of the barycentre model and proof an explicit representation of the Radon-Nikodym derivative relative to the average drift model. We further allow the agent to include their own constraints, which results in an optimal model that can be seen as a distortion of the experts barycentre model to incorporate the agents constraints. Two deep learning algorithms are proposed to find the optimal drift of the combined model, allowing for efficient simulations. The first algorithm aims at learning the optimal drift by matching the change of measure, whereas the second algorithm leverages the notion of elicitability to directly estimate the value function. The paper concludes with a extended application to combine implied volatility smiles models that were estimated on different datasets. Joint work with Silvana Pesenti.

Steven Kou (Boston University)

Bitcoin Mining and Climate Damage

We propose a tractable dynamic equilibrium model to study the impact of Bitcoin mining on climate damage. To incorporate miners endogenous exit and entry with technology innovation, a critical modeling step is to draw on the idea of Dammon, Spatt, and Zhang (2001) and Ben Tahar, Soner, and Touzi (2010) to track the average operating costs rather than the exact operating costs, thus overcoming the difficulty of strong

path-dependency incurred by the interaction among endogenous exit, entry, and technology innovation. The model can capture empirical co-movements of miners computing power and mining revenue. The model predicts that Bitcoin mining is not economically sustainable in terms of long-term climate damage. This is joint work with Min Dai. Shuaijie Qian, and Ling Qin.

Anastasis Kratsios (McMaster & Vector Institute)

Exponential Expression Rates for Neural Operator Approximations to the Solution Operator of Certain FBSDEs

The numerical solution to forward-backwards stochastic differential equations (FBSDEs) plays a central role in optimal control and its applications to game theory, economics, finance, and insurance. Most classical numerical and modern deep-learning schemes, however, have the disadvantage that they must be re-run every time the user specifies a new set of parameters and/or terminal conditions for an FBSDE, meaning that these methods cannot feasibly solve large families of FBSDEs. One possible solution is to consider a neural operator (NO) which "learns to solve FBSDEs"; the NO outputs the solution to an FBSDE given inputs: a terminal condition and a generator of the backward process. Though the existence of such NOs is not surprising, it is unclear if they can be implemented using a few parameters. We establish exponential rates for NO approximations of the solution operator to a broad class of fully coupled FBSDEs with random terminal time. Our result is based on new exponential approximation rates for a class of convolutional NOs, which can efficiently encode Green's function to the Elliptic boundary problems associated with our FBSDEs. Joint work with: Takashi Furuya.

Anne MacKay (Universit de Sherbrooke)

Continuous-Time Markov Chain Approximations for an Optimal Stopping Problem with Discontinuous Reward Function

We consider an optimal stopping problem with an unbounded, time-dependent and discontinuous reward function. Using an alternative representation for the value function of the optimization problem, we study its analytical properties and the resulting exercise region, and we obtain different representations for the value function. From these results, we develop fast and efficient pricing algorithms for equity-linked insurance contracts and convertible bonds under a wide range of market models. We do so by approximating the diffusion processes describing the market by a two-layer continuous-time Markov chain. Numerical examples show the efficiency of our pricing algorithms. This is joint work with Marie-Claude Vachon (UQAM).

Thibaut Mastrolia (UC Berkeley)

Transaction Fees and Auction Market Design

Flaws of a continuous limit order book mechanism raise the question of whether a continuous trading session and a periodic auction session would bring better efficiency. This talk wants to go further in designing an auction when both a continuous market and a periodic auction market are available to traders. In a periodic auction, we discover that a strategic trader could take advantage of the accumulated information available along the auction duration by arriving at the latest moment before the auction closes, increasing the price impact on the market. Such price impact moves the clearing price away from the efficient price and may disturb the efficiency of a periodic auction market. We thus propose and quantify the effect of two remedies to mitigate these flaws: randomizing the auction's closing time and optimally designing a transaction fees policy. Our results show that these policies encourage a strategic trader to send their orders earlier to enhance the efficiency of the auction market, illustrated by data extracted from Alphabet and Apple stocks. Joint work with Tianrui Xu (UC Berkeley).

Felix Matthys (ITAM)

Optimal Climate Policy Mix

In this paper, we derive a general equilibrium framework to study an optimal policy mix problem to tackle the negative impact of climate change. Our model economy contains a clean and a polluting sector, a representative agent that chooses his or her optimal consumption and investment in each sector, plus a benevolent government that sets optimal climate taxes, subsidies to the clean sector, adaptation and mitigation policies to navigate the transition to a cleaner economy without excessively hampering economic growth. For each of the four RCP Scenarios 26, 45, 60 and 85, we disciple our model by matching both, the atmospheric concentration of CO2 and surface temperatures, while simultaneously calibrating the optimal policies to current projected values implied by each RCP scenario. Our novel framework is flexible enough to match

both the climate moments, as well as the economic moments, provided we use the full optimal policy mix. We then discuss in detail why such a rich policy mix is indeed necessary to meet both targets and discuss why simpler policies fail. Finally, we document several unintended consequences of policymaking to highlight their complex relationship.

Silvana Pesenti (University of Toronto)

Portfolio Choice with α-Bregman Wasserstein Penalisation

We consider the problem of active portfolio management, where an agent aims at finding the portfolio with maximal expected utility of terminal wealth subject to deviation constraints from a benchmark portfolio. As the agent values gains and losses differently, they utilise an asymmetric divergence on the space of distribution. Moreover, the agent aims at outperforming the benchmark, thus penalises outcomes where the portfolio wealth is below that of the benchmarks. This is achieved by the recently introduced α-Bregman Wasserstein divergence, generalising the Bregman Wasserstein and the popular Wasserstein divergence. We prove existence and uniqueness of the optimal portfolio strategy and discuss when the strategy coincides with the Merton strategy. We further give explicit criteria when the divergence constraints and the budget constraints are binding. This is joint work with Thai Nguyen

Dylan Possama (ETH Zrich)

Randomness and Early Termination: What Makes a Game Exciting?

In this talk, we revisit an open problem posed by Aldous on the max-entropy win-probability martingale: given two players of equal strength, such that the win-probability is a martingale diffusion, which of these processes has maximum entropy and hence gives the most excitement for the spectators? From a stochastic control perspective, the corresponding value function can be characterised by a nonlinear parabolic PDE, for which we show existence and uniqueness of classical solutions. We establish key qualitative properties of the solution including concavity, monotonicity, and convergence to a steady state. Moreover, we construct convergent numerical approximations, which allow us to highlight the behaviour of the win-probability process in the present case where the match may end early, in contrast to recent work by Backhoff-Veraguas and Beiglbck where the match always runs the full length. This is a joint work with Gaoyue Guo, Sam Howison, and ChristophReisinger.

Jinniao Qiu (University of Calgary)

A Particle Consensus Approach to Solving Nonconvex-Nonconcave Min-Max Problems

A zero-order optimization method is introduced for sequential min-max problems based on two populations of interacting particles. The systems are coupled so that one population aims to solve the inner maximization problem, while the other aims to solve the outer minimization problem. The dynamics are characterized by a consensus-type interaction with additional stochasticity to promote exploration of the objective landscape. Without relying on convexity or concavity assumptions, theoretical convergence guarantees of the algorithm are established via a suitable mean-field approximation of the particle systems. Numerical experiments illustrate the validity of the proposed approach. In particular, the algorithm is able to identify a global min-max solution, in contrast to gradient-based methods, which typically converge to possibly suboptimal stationary points. This talk is based on joint work with Giacomo Borghi and Hui Huang.

Alexandros Saplaouras (NTUA)

Stability of Backward Propagation of Chaos

It will initially be considered the asymptotic behavior of the solution of a mean-field system of Backward Stochastic Differential Equations with Jumps (BSDEs), as the multitude of the system equations grows to infinity, to independent and identically distributed (IID) solutions of McKeanVlasov BSDEs. This property is known in the literature as backward propagation of chaos. Afterwards, it will be provided the suitable framework for the stability of the aforementioned property to hold. In other words, assuming a sequence of mean-field systems of BSDEs which propagate chaos, then their solutions, as the multitude of the system equations grows to infinity, approximates an IID sequence of solutions of the limiting McKeanVlasov BSDE. The generality of the framework allows to incorporate either discrete-time or continuous-time approximating mean-field BSDE systems.

Xiaofei Shi (University of Toronto)

Reinforced-GANs for Financial-Market Equilibria

We incorporate trading costs into a dynamic multiagent general-equilibrium model, in which participating agents optimally decide their trade and their aggregated demands for the stocks matches with the total shares outstanding in the market. Although global equilibrium is achieved under specific market dynamics, the nonlinear nature of the system makes it falls outside the scope of any known well-posedness results. In this work, we show how to leverage deep-learning techniques to obtain numerical solutions with calibrated parameters to market prices and trading volumes. In particular, we propose the architecture of reinforced generative adversarial networks (Reinforced-GANs) as a numerical algorithm for equilibrium models, where GANs not only overcome the curse-of-dimensionality and but also show their great scalability.

Costas Smaragdakis (University of the Aegean)

A Deep Implicit-Explicit Minimizing Movement Method for Option Pricing in Lévy Models

Solving high-dimensional differential and integro-differential equations remains a significant challenge in mathematical finance, particularly in the context of option pricing. Recently, innovative approaches have emerged that approximate solutions by training neural networks with loss functions tailored to the differential operator of the equation, incorporating initial/terminal and boundary conditions. In this talk, we will present advanced machine-learning procedures for pricing European basket options, where the underlying assets are subject to correlated dynamics with random discontinuities. The neural network architecture we propose is designed to ensure the solution's known asymptotic behaviour for extreme values of the underlying assets. Furthermore, the architecture is designed to align the network outputs with the solution's known qualitative properties, enhancing their consistency and reliability. We will present results corresponding to various Lévy models to demonstrate the merits of our model in solving high-dimension option pricing problems.

Yang Yang (University of Calgary)

Stochastic Path-Dependent Volatility Models for Price-Storage Dynamics in Natural Gas Markets and Discrete-Time Swing Option Pricing

This talk is about the price-storage dynamics in natural gas markets. A novel stochastic path-dependent volatility model is introduced with path-dependence in both price volatility and storage increments. Model calibrations are conducted for both the price and storage dynamics. Further, we discuss the pricing problem of discrete-time swing options using the dynamic programming principle, and a deep learning-based method is proposed for numerical approximations. A numerical algorithm is provided, followed by a convergence analysis result for the deep-learning approach.

Fenghui Yu (TU Delft)

Signal-Adaptive Optimal Execution Quotes

This talk focuses on optimal execution strategies for sequentially placing limit orders in a limit order book at specific quote prices. Unlike previous studies that primarily derive optimal trading speed within the Almgren-Chriss framework, we address the problem by detailing how limit orders should be placed over time. Our model considers that price quotes in the limit order book may be influenced by signals and incorporates both execution probabilities and price impact. We frame the optimal execution problem with four objectives: execution without risk aversion (Case I), execution with running inventory risk (Case II), execution with exponential utility (Case III), and execution with both running inventory risk and exponential utility (Case IV). By deriving the corresponding HJB equations for each case, we demonstrate that these problems converge to a specific type and can be solved explicitly, leading to fully explicit solutions for optimal execution across all cases.

Jiacheng Zhang (University of Berkeley)

Ito's Formula for Flows of Conditional Measures on Semimartingales

Motivated by recent development of mean-field systems with common noise, this paper establishes Ito's formula for flows of conditional probability measures under a common filtration associated with general semimartingales. This generalizes existing works on flows of conditional measures on Ito processes and flows of deterministic measure on general semimartingales. The key technical components involve constructing conditional independent copies and establishing the equivalence between stochastic integrals with respect to the conditional law of semimartingales and the conditional expectation of stochastic integrals with respect to copies of semimartingales. Ito's formula is then established for cylindrical functions through conditional independent copies, and extended to the general case through function approximations.

Yufei Zhang (Imperial College London)

α-Potential Games: A New Paradigm for N-Player Games

Static potential games, pioneered by Monderer and Shapley (1996), are non-cooperative games in which there exists an auxiliary function called static potential function, so that any player's change in utility function upon unilaterally deviating from her policy can be evaluated through the change in the value of this potential function. The introduction of the potential function is powerful as it simplifies the otherwise challenging task of finding Nash equilibria for non-cooperative games: maximizers of potential functions lead to the game's Nash equilibria. In this talk, we propose an analogous and new framework called α-potential game for dynamic N-player games, with the potential function in the static setting replaced by an α-potential function. We present an analytical characterization of α-potential functions for any dynamic game. For stochastic differential games in which the state dynamic is a controlled diffusion, α is explicitly identified in terms of the number of players, the choice of admissible strategies, and the intensity of interactions and the level of heterogeneity among players. We provide detailed analysis for games with mean-field interactions, distributed games, and crowd aversion games, for which α is shown to decay to zero as the number of players goes to infinity, even with heterogeneity in state dynamics, cost functions, and admissible strategy classes. We also show α is capable of capturing the subtle difference between the open-loop and closed-loop strategies. The talk is based on joint work with Xin Guo and Xinyu Li: https://arxiv.org/abs/2403.16962

Haosheng Zhou (University of California, Santa Barbara)

Stochastic Differential Games on Graphs

In this talk, we present a new model for stochastic differential games on graphs, aiming to bridge game theory with network structures to capture the influence of graph structures on strategic interactions. Our framework supports heterogeneous player interactions across general graph structures, extending current models to encompass more complex, network-driven dynamics. We establish two main results: firstly, we demonstrate the convergence of fictitious play, along with numerical estimates of convergence rates that reflect key aspects of the graph structure. Secondly, we provide a semi-explicit construction of the Nash equilibrium, validated through numerical simulations and offering a reliable computational baseline for future applications in deep learning. This is joint work with Ruimeng Hu and Jihao Long.

Xunyu Zhou (Columbia University)

Generative AI for Diffusion Models by q-Learning

We study a generative AI problem in which one trains continuous-time score-based diffusion models to generate samples that maximize some reward function while keeping the generated distribution close to the true (yet unknown) distribution of a given set of samples. We formulate the problem as reinforcement learning and solve it by the (little)q-learning theory. Joint work with Xuefeng Gao and Jiale Zha.

3.2 Online Presentations

ECRs attending online have been offered the opportunity to give short talks during an online session. The shorter format of online presentations allowed ECRs to hone their concise presentation skills and receive valuable initial feedback from experts, with opportunities for follow-up discussions after the online session. This approach also helped mitigate the fatigue often associated with longer online sessions, while still providing significant professional development opportunities to the ECRs.

Daniel Bartl (University of Vienna)

Statistical Estimation of Stochastic Optimization Problems and Risk Measures

We develop a novel procedure for estimating the optimizer of general convex stochastic optimization problems from an iid sample. This procedure is the first one that exhibits the optimal statistical performance in heavy tailed situations and also applies in highdimensional settings. We discuss the portfolio optimization problem and the estimation of risk measures. Joint works with Stephan Eckstein and Shahar Mendelson.

Gökçe Dayanıklı (University of Illinois at Urbana-Champaign)

Cooperation, Competition, and Common Pool Resources in Mean Field Games

The tragedy of the commons (TOTC, introduced by Hardin, 1968) states that the individual incentives will result in overusing common pool resources which in turn may have detrimental future consequences that

affect everyone negatively. However, in many real-life situations this does not happen and researchers such as the Nobel Prize winner Elinor Ostrom suggested mutual restraint by individuals can be the preventing factor. In mean field games (MFGs), since individuals are insignificant and fully non-cooperative, the TOTC is inevitable. This shows that MFG models should incorporate a mixture of selfishness and altruism to capture real-life situations that include common pool resources. Motivated by this, we will discuss different equilibrium notions to capture the mixture of cooperative and non-cooperative behavior in the population. First, we will introduce mixed individual MFGs and mixed population MFGs where we also include the common pool resources. The former captures altruistic tendencies at the individual level and the latter models a population that is a mixture of fully cooperative and non-cooperative individuals. For both cases, we will briefly discuss definitions and characterization of equilibrium with the forward backward stochastic differential equations. Later, we will discuss a real-life inspired example of fishers where the fish stock is the common pool resource. We will analyze the existence and uniqueness results and discuss the experimental results. (This is a joint work with Mathieu Lauriere.)

Mathieu Lauriere (NYU Shanghai)

Deep Learning for Stackelberg Mean Field Games via Single-Level Reformulation

We propose a single-level numerical approach to solve Stackelberg mean field game (MFG) problems. In Stackelberg MFG, an infinite population of agents play a non-cooperative game and choose their controls to optimize their individual objectives while interacting with the principal and other agents through the population distribution. The principal can influence the mean field Nash equilibrium at the population level through policies, and she optimizes her own objective, which depends on the population distribution. This leads to a bi-level problem between the principal and mean field of agents that cannot be solved using traditional methods for MFGs. We propose a reformulation of this problem as a single-level mean field optimal control problem through a penalization approach, and we prove convergence of the reformulated problem to the original problem. We propose a machine learning method based on (feed-forward and recurrent) neural networks and illustrate it on several examples from the literature. Joint work with Gokce Dayanikli.

Julian Sester (National University of Singapore)

Uncertainty-Aware Calibration of Affine Models

Robust modeling approaches often start with a given parameter set. In this talk we study the inverse engineering question how to obtain this parameter set through calibrating stochastic models to historical data. We focus on the class of affine term structure models under parameter uncertainty, while the approach can be applied similarly to other model classes or other financial markets. Our approach introduces a dynamic calibration to an observed time series of market prices by determining parameter sets in such a way that bid/ask price intervals are reproduced by the non-linear model. Through numerical illustrations employing US bond market data and SOFR rates, we show that our dynamic calibration approach outperforms non-robust approaches in typical applications such as hedging of derivatives. (Joint work with Thorsten Schmidt and Eva Ltkebohmert)

Marko Hans Weber (National University of Singapore)

General Equilibrium with Unhedgeable Fundamentals and Heterogeneous Agents

We examine the implications of unhedgeable fundamental risk, combined with agents' heterogeneous preferences and wealth allocations, on dynamic asset pricing and portfolio choice. We solve in closed form a continuous-time general equilibrium model in which unhedgeable fundamental risk affects aggregate consumption dynamics, rendering the market incomplete. Several long-lived agents with heterogeneous risk-aversion and time-preference make consumption and investment decisions, trading risky assets and borrowing from and lending to each other. We find that a representative agent does not exist. Agents trade assets dynamically. Their consumption rates depend on the history of unhedgeable shocks. Consumption volatility is higher for agents with preferences and wealth allocations deviating more from the average. Unhedgeable risk reduces the equilibrium interest rate only through agents' heterogeneity and proportionally to the cross-sectional variance of agents' preferences and allocations.

4 Scientific Progress Made

The workshop showcased significant advancements across various themes in financial mathematics, financial technology, and financial economics. Researchers presented innovative methodologies and applied insights, leading to valuable progress in understanding and addressing complex challenges in these fields. Key areas of progress include:

Dynamic Models and Systemic Risk

- René Aïd introduced a framework for dynamic persuasion with applications to energy markets, highlighting the role of information provision in reducing mispricing and pollution.
- Stéphane Crépey proposed advanced statistical learning methods for conditional value-at-risk and expected shortfall, offering robust tools for risk assessment under heavy-tailed financial losses.

Market Microstructure and Optimization

- Ulrich Horst explored optimal trade execution under endogenous order flow, providing insights into hyperbolic and cyclic trading strategies that improve market performance.
- Thibaut Mastrolia presented innovative auction market designs aimed at mitigating inefficiencies in periodic trading systems, with practical implications for improving equity market operations.

Advances in Financial Technology

- Geneviève Gauthier introduced deep learning approaches for dynamic hedging of S&P 500 options, incorporating volatility surface dynamics to outperform conventional benchmarks.
- Ranu Castaneda developed a dynamic model for open banking, analyzing competition between traditional banks and FinTech firms, which reshapes financial inclusion and borrower benefits.

Mean-Field Games and Multi-Agent Systems

- Anran Hu presented mean-field optimization methods integrated with reinforcement learning, enabling efficient solutions to Nash equilibria in large-scale systems.
- Haosheng Zhou extended stochastic differential games to graph structures, bridging network theory with game dynamics to capture heterogeneous interactions.

Sustainability and Climate Finance

- Steven Kou analyzed the economic and environmental implications of Bitcoin mining, presenting a dynamic equilibrium model to assess long-term climate damage.
- Felix Matthys proposed a comprehensive policy mix for climate change mitigation, balancing economic growth and environmental sustainability through optimal taxation and subsidies.

Machine Learning in Financial Modeling

- Anastasis Kratsios demonstrated exponential approximation rates for neural operators in solving forward-backward stochastic differential equations, with implications for optimal control problems.
- Costas Smaragdakis applied machine-learning techniques to high-dimensional differential equations, advancing the pricing of options under Lévy models.

The workshop not only advanced theoretical foundations but also fostered interdisciplinary collaborations, linking mathematical finance with cutting-edge computational methods and real-world applications. It underscored the critical role of integrating traditional modeling with innovative technologies to address emerging challenges in financial markets.

5 Outcome of the Meeting

In addition to the scientific presentations, the workshop featured two 90-minute discussion sessions: one focused on unconscious biases in academia and another in the form of a career panel discussion.

5.1 Discussion on Unconscious Biases

The session on unconscious biases was structured into three segments: (i) viewing the Canadian Tri-Agency's unconscious bias training module[1], (ii) engaging in group discussions to explore diverse perspectives and experiences, and (iii) reconvening as a whole to share key insights from the groups.

The training module provided an introductory framework for understanding unconscious bias and set the stage for deeper dialogue. During the group discussions, participants examined various questions such as:

- Can you share a personal experience where you encountered or witnessed unconscious bias in your academic career?
- What are some common indicators of unconscious bias that faculty, staff, and students should be aware of in academic settings?
- In your opinion, what are the most effective strategies for individuals and institutions to address unconscious bias in hiring and promotion processes?
- How effective do you find bias training programs in academia? What improvements would you suggest to make them more impactful?
- How do unconscious biases affect student experiences in the classroom, particularly for underrepresented groups?
- What evidence exists regarding unconscious bias in research funding allocations, and how can these biases be mitigated?
- How can mentors and advisors play a role in identifying and mitigating unconscious biases among their mentees?
- What advice do you have for individuals who want to raise awareness about unconscious bias within their academic communities?

The group discussions yielded valuable insights on identifying and addressing unconscious bias. Key highlights included strategies to increase awareness of unconscious biases, both personally and within institutional processes, such as hiring committees. Notable proposals included ensuring a sufficiently large and diverse pool of candidates while approaching the "excellence regime", as well as deferring the evaluation of reference letters until later stages of hiring. Participants also explored ways to empower individuals affected by bias to voice concerns effectively and proposed institutional-level changes to create more equitable environments. This session underscored the importance of continual dialogue and action to mitigate unconscious bias, fostering a culture of inclusivity and fairness in academia.

5.2 Career Panel Discussion

The workshop featured a career panel discussion led by Genevive Gauthier and Matheus Grasselli, and moderated by Antonis Papapantoleon. The session aimed to engage participants in a dialogue about the difficulties encountered in academic careers by scientists with an underrepresented background, the multifaceted responsibilities of modern academics, and how to balance these demands with personal life and individual choices. While the session was primarily targeted at early-career academics, it drew significant interest from participants across all career stages.

[1] Available at https://www.chairs-chaires.gc.ca/program-programme/equity-equite/bias/en/

The discussion began cautiously, as attendees seemed reluctant to raise topics or ask questions. Genevive Gauthier broke the ice by mentioning that she was the second woman appointed in a Business School, discussing the challenges this brought, but also highlighting her trust in the system in order to deliver a positive outcome. This relatable anecdote not only eased the initial tension but also set a tone of openness and candor that encouraged participants to share their own concerns. Matheus Grasselli stated that institutional measures are required in order to ensure a sense of identity and belonging for scientists from underrepresented minorities, while Ruimeng Hu discussed the "mentoring initiave" in financial mathematics[2].

Another key focus of the discussion was the broad range of tasks academics must navigate. Beyond teaching and research–often perceived as the only responsibilities of academics–the discussion revealed a far more complex reality. Many participants noted they had been unaware of the breadth of tasks involved until they experienced them firsthand. These responsibilities include:

- Securing funding: writing grant proposals and seeking financial support.
- Supervising students: mentoring undergraduate, graduate, doctoral, and postdoctoral researchers.
- Service to the institution: participating in PhD defense committees, admissions processes, and departmental leadership roles.
- Service to the community: reviewing papers, contributing to journal editorial boards, and organizing conferences.
- Larger outreach: developing collaborations with industry, filing patents, and engaging in outreach activities.

This multiplicity of roles requires a wide array of skills. As Matheus Grasselli aptly put it: "*You should not expect to be A+ in everything, and must choose in which areas its ok to be B+. Otherwise life will choose it for you, and it will be an F.*"

The panelists and participants acknowledged that most academics tend to prioritize research, often at the expense of service or community engagement. This imbalance is reinforced by institutional metrics that overemphasize research output, while undervaluing contributions like mentoring or peer review. Without reviewers, for instance, papers would not be published, yet this critical task is rarely rewarded. Genevive and Matheus emphasized the need for academia to recognize and reward excellence across all dimensions of academic life, not just research.

Academics were encouraged to reflect on their strengths, interests, and institutional expectations to focus their efforts effectively. Institutions, in turn, could benefit from intentionally building balanced teams that reflect diverse skill sets. This diversity strengthens collective performance and reduces the burden on individuals by ensuring no single person is expected to excel in every area. Recognizing diverse skill sets should also align with fostering broader diversity in hiring practices, avoiding clichs or assumptions about roles based on gender or background.

The discussion delved into frustrations with the prevailing criteria for academic excellence. A well-established senior professor expressed concerns about traditional metrics of success, such as publication counts and citation indices, which often incentivize quantity over quality. Their candid critique highlighted that even those who have thrived within the current system recognize the need for cultural and structural change. Participants also reflected on the accelerating pace of academic life–from rapid publishing cycles to an overwhelming volume of papers–and its detrimental effects on research quality and personal well-being. The session concluded with this thoughtful reminder of the importance of continuously re-evaluating the trajectory of academia and the values it strives to uphold.

A short essay [7] resulted from this panel discussion, that elaborates on the various tasks of modern academics and relates them to career and personal choices. Let us finally mention that, following the panel discussion, several participants stated that they were initially reluctant to participate in this activity, but they felt very happy that they did attend in the end.

References

[1] Y. Armenti, S. Crépey, S. Drapeau, and A. Papapantoleon, Multivariate shortfall risk allocation and systemic risk, *SIAM Journal on Financial Mathematics* **9** (2018), 90–126

[2] https://sites.google.com/view/finmathmentor/home

[2] M. Avellaneda and S. Stoikov, High frequency trading in a limit order book, *Quantitative Finance* **8** (2008), 217–224

[3] M. Dai, S. Kou, S. Qian, and L. Qin, Bitcoin mining and electricity consumption, available at https://ssrn.com/abstract=3994797 (2024)

[4] E. Dimanchev and C. Knittel, Designing climate policy mixes: Analytical and energy system modeling approaches, *Energy Economics* **122** (2023), 106697

[5] P. François, G. Gauthier, F. Godin, and C. Pérez Mendoza, Is the difference between deep hedging and delta hedging a statistical arbitrage?, *Finance Research Letters* **73** (2025), 106590

[6] Z. He, J. Huang and J. Zhou, Open banking: Credit market competition when borrowers own the data, *Journal of Financial Economics* **147** (2023), 449–474

[7] A. Papapantoleon, What makes a good academic?, Unpublished Essay, 2025.

[8] S. Pesenti and S. Jaimungal, Portfolio optimization within a Wasserstein ball, *SIAM Journal on Financial Mathematics* **14** (2023), 22M1496803