Mathematical Finance
Finance mathématique
(Org: **Christoph Frei** and/et **Alexander Melnikov** (University of Alberta))

ALEXANDRU BADESCU, University of Calgary
Option Pricing with Recurrent Variance Dependent Stochastic Discount Factors and Realized Volatility

This paper develops an option pricing framework that integrates general Realized EGARCH return dynamics with an exponential linear stochastic discount factor (SDF), in which variance risk aversion is modelled using a recurrent neural network (RNN). Using S&P 500 index options, we show that the RNN-based SDF substantially improves the cross-sectional fit relative to standard autoregressive and constant variance-dependent specifications, with the largest gains for deep out-of-the-money and short-maturity contracts. The results indicate that allowing the pricing kernel to incorporate complex, state-dependent variance risk premia is essential for capturing option market nonlinearities. A GPU-accelerated implementation based on Libtorch (PyTorch's C++ API) and CUDA ensures computational feasibility for large-scale estimation.

FRANÇOIS-MICHEL BOIRE, University of Ottawa
Modeling Systemic House Price Risk

Economists and policy makers have become increasingly aware of the role of house price risk in driving financial fragility. This paper develops a semiparametric framework to model and assess downside risk in the U.S. housing market. First, we use panel quantile regressions to capture heterogeneous effects of supply, demand, and non-fundamental factors across the distribution of state-level house price changes. Second, we estimate the quantile regression jointly with a copula-based structure to capture cross-state dependence. Finally, we construct a measure of systemic housing risk using a weighted composition of state-level Case–Shiller price indices, allowing us to compare tail exposures and quantify cross-state contributions to aggregate risk. This is joint work with S. van Norden.

TAHIR CHOULLI, University of Alberta
Pricing formulas for vulnerable claims and death derivatives

We consider the discrete-time market model described by the triplet $(S, \mathbb{F}, \tau)$. Herein $\mathbb{F}$ is the "public" flow of information which is available to all agents overtime, S is the discounted price process of d-tradable assets, and τ is an arbitrary random time whose occurrence might not be observable via $\mathbb{F}$. This framework covers the credit risk theory where τ represents the default time, the life insurance setting where τ models the death time, and other areas of finance. For various vulnerable claims in credit risk and death derivatives in life insurance, we address the super-hedging pricing valuation problem in many aspects. First of all, we discuss how the Immediate-Profit arbitrage (IP for short), which is the economical assumption that guarantees the existence of the "minimal" super-hedging price $\widehat{P}^{\mathbb{G}}$, is affected by τ. Then we show, as explicit as possible, how the set of all super-hedging prices expands under the stochasticity of τ and its various risks. Afterwards, we elaborate, as explicit as possible, the pricing formulas for vulnerable claims and death derivatives. Finally, we single out explicitly the various informational risks in the dynamics of the price process $\widehat{P}^{\mathbb{G}}$ and quantify them. This latter fact is highly important for the mortality and longevity securitizations.

This talk is based on the following joint work with Emmanuel Lepinette (Paris-Dauphine, France):

T. Choulli and Emmanuel: Super-hedging-pricing formulas and Immediate-Profit arbitrage for market models under random horizon. to appear in Finance and Stochastics. A version of the paper is available at: arXiv:2401.05713.

MATT DAVISON, Western University Canada
A Real Options Approach to Wildfire Evacuations

Joint work with Daniel Guerrero and Doug Woolford

Wildfires pose an increasing threat to human life and property in Canada. Approximately 12% of Canada's population resides in the wildland-urban interface and forest fires are increasing both in frequency and severity. Significant research efforts has been devoted to understanding different components of wildfire risk, including the way in which wildfire moves on the landscape. When fire approaches populated areas, it can be optimal to evacuate the area to reduce danger to life. Evacuating too late will be much more expensive, if possible at all; evacuating too early risks disruptive unnecessary evacuations. In this talk I will examine how financial mathematics approach can help frame this problem in useful ways that not only allow evacuation decisions to be made, but which also provides a way to compare, apples to apples, to financial benefit of measures taken to prevent wildfire spread against their cost.

DENA FIROOZI, University of Toronto

Ranking Quantilized Mean-Field Games and Early-Stage Venture Investments

We study a class of quantilized mean-field game models with a capacity for ranking games, where the performance of each agent is evaluated based on its terminal state relative to the population's α-quantile value, $\alpha \in (0,1)$. This evaluation criterion is designed to select the top $(1-\alpha)\%$ performing agents. We provide two formulations for this competition: a target-based formulation and a threshold-based formulation. In the former and latter formulations, to satisfy the selection condition, each agent aims for its terminal state to be *exactly* equal and *at least* equal to the population's α-quantile value, respectively.

For the target-based formulation, we obtain an analytic solution and demonstrate the ϵ-Nash property for the asymptotic best-response strategies in the N-player game. Specifically, the quantilized mean-field consistency condition is expressed as a set of forward-backward ordinary differential equations, characterizing the α-quantile value at equilibrium. For the threshold-based formulation, we obtain a semi-explicit solution and numerically solve the resulting quantilized mean-field consistency condition.

Subsequently, we propose a new application in the context of early-stage venture investments, where a venture capital firm financially supports a group of start-up companies engaged in a competition over a finite time horizon, with the goal of selecting a percentage of top-ranking ones to receive the next round of funding at the end of the time horizon. We present the results and interpretations of numerical experiments for both formulations discussed in this context and show that the target-based formulation provides a very good approximation for the threshold-based formulation.

CHRISTOPH FREI, University of Alberta

A Doubly Continuous Model for Equilibrium Trading Dynamics

Analysis of financial markets is usually based on rational expectations, where investors use all available information to trade in order to maximize their expected utility. In equilibrium models, prices are determined so that the market clears, meaning that demand equals supply. Typically, diverging information among homogeneous agents is not enough to generate trade in equilibrium. To address this issue, we introduce and analyze a doubly continuous model with continuous time and continuous agent space. In this setting, each agent is infinitesimally small, contributing zero to trade, while collective trade emerges from the aggregation over non-negligible sets of agents. Our approach leverages tools from Brownian sheets and multiparametric stochastic calculus, providing insights into the interplay of information, behaviour, and equilibrium in financial markets.

This talk is based on joint work with Efstathios Avdis (University of Alberta), Sergei Glebkin (INSEAD), and Raphael Huwyler (University of Alberta).

NIUSHAN GAO, Toronto Metropolitan University

On Continuity and Asymptotic Consistency of Measures of Risk and Variability

Ruszczynski and Shapiro (2006) showed that a convex, real-valued functional on a Banach lattice is continuous whenever it is either increasing or decreasing. This result has played an important role in the development of the theory of risk measures. In this talk, we show that the monotonicity assumption can be relaxed to a much weaker condition: it suffices that the functional be bounded above on every interval. This extension permits new applications, particularly to measures of variability. We

also present an improvement of a result of Kratschmer, Schied and Zahle (2014) concerning the asymptotic consistency of law-invariant risk measures.

This talk is based on joint work with Foivos Xanthos.

GENEVIÈVE GAUTHIER, HEC Montréal
Beyond volatility of volatility: Decomposing the informational content of VVIX

This study investigates the informational content of the VVIX, traditionally viewed as a proxy for the S&P 500 index's volatility of the volatility (VOV). We show that this interpretation is incomplete: the VVIX also embeds a long-run variance (LRV) component. To establish this result, we first demonstrate that regressions of squared VVIX on VOV proxies gain substantial explanatory power once LRV measures are incorporated. We then develop a tractable theoretical framework linking VVIX to three risk drivers—instantaneous variance, LRV, and VOV—and show that the VVIX loads on both VOV and LRV. Our empirical analysis reveals that VVIX dynamics are dominated by LRV in calm markets, but by VOV during financial stress. We further show that these variance components explain option returns in distinct markets: S&P 500 index option straddles load on the instantaneous variance and LRV, while VIX option straddles load on the VOV. Taken together, our results redefine the role of the VVIX, establishing it as a measure of both VOV and LRV uncertainty, with important implications for how it should be read and used by finance practitioners.

FRÉDÉRIC GODIN, Concordia University
Deep Hedging with Options Using the Implied Volatility Surface

We propose a deep hedging framework for index option portfolios, grounded in a realistic market simulator that captures the joint dynamics of S&P 500 returns and the full implied volatility surface. Our approach integrates surface-informed decisions with multiple hedging instruments and explicitly accounts for transaction costs. The hedging strategy also considers the variance risk premium embedded in the hedging instruments, enabling more informed and adaptive risk management. Tested on a historical out-of-sample set of straddles from 2020 to 2023, our method consistently outperforms traditional delta-gamma hedging strategies across a range of market conditions.

MATHEUS GRASSELLI, McMaster University
A Tale of Two Regions: A North and South Macroeconomic-Ecological Model

In this talk, I will describe an extension of the GEMMES climate-economic model proposed in Bovari et al. (2018a) that considers two regions, a Global North and a Global South, interacting through trade. Each region decides on its own carbon pricing policy and abatement subsidies independently, leading to separate paths for industrial emissions, which contribute together to the increase of atmospheric carbon concentration and global average increase in temperature. The two regions are subject to distinct damages caused by climate change, leading to separate paths for economic variables such as output and inflation. I will show a calibration of the model to data available up to 2016 and a test of the predictions up to 2024 with broad agreement for the key variables in the model. I will then investigate three different scenarios for future damages and climate policies and their effects on each of the regions, as well as a case study for financial transfers from the North to the South to help mitigate climate change. I conclude with a sensitivity analysis of the proposed model using similar techniques previously used to analyze the original GEMMES model. This is joint work with B. Badenhorst, K. Baldeo, K. Bopape, E. Kroell, and D. Presta.

CODY HYNDMAN, Concordia University
Optimal annuitization with labor income under age-dependent force of mortality

We consider the problem of optimal annuitization with labour income, where an agent aims to maximize utility from consumption and labor income under age-dependent force of mortality. Using a dynamic programming approach, we derive closed-form solutions for the value function and the optimal consumption, portfolio, and labor supply strategies. Our results show that

before retirement, investment behavior increases with wealth until a threshold set by labor supply. After retirement, agents tend to consume a larger portion of their wealth. Two main factors influence optimal annuitization decisions as people get older. First, the agent's perspective (demand side); the agent's personal discount rate rises with age, reducing their desire to annuitize. Second, the insurer's perspective (supply side); insurers offer higher payout rates (mortality credits). Our model demonstrates that beyond a certain age, sharply declining survival probabilities make annuitization substantially optimal, as the powerful incentive of mortality credits outweighs the agent's high personal discount rate. Finally, post-retirement labor income serves as a direct substitute for annuitization by providing an alternative stable income source. It enhances the financial security of retirees. (Joint work with Criscent Birungi)

ANASTASIS KRATSIOS, McMaster University
A Neural Black–Scholes Formula

Despite its central role in option markets, the implied volatility surface (IVS) remains exceptionally difficult to calibrate to quoted call prices without breaching fundamental economic constraints. We resolve this long-standing problem by deriving a simple, model-free, smooth call option pricing formula describing a (sparse fully-trained) two-layer neural network matching quoted market call prices; both in the strike and maturity. Our formula is *adaptively arbitrage free* (AF) in that is necessrily produces an arbitrage-free call surface whenever the quoted market data is arbitrage-free. The regularity of our data-driven call surface allows us obtain a closed-form reconstruct risk-neutral dynamics for the underlying only using the available market call quotes via the Dupire formula. Moreover, on AF data, our IVS is guaranteed to uniformly approximate call slices at an optimal rate of $\mathcal{O}(1/n^2)$ all all points between any quoted market prices, using n neurons.

We demonstrate state-of-the-art predictive power with virtually no computational overhead, across synthetic data and real-world cryptocurrency markets, by routinely achieving several orders of magnitude greater accuracy than both industry and deep learning benchmarks. Our model-free option pricing formula is subsumes the classical *Black–Scholes* (BS) formula, in that it uses the BS put price as its activation function.

Joint work: Hans Buehler, Blanka Horvath, Yannick Limmer, and Raeid Saiqur

---

ANNE MACKAY, Université de Sherbrooke
Pricing lookback options on quantum computers

Quantum computing promises computational speed up that could have a significant impact across industries. In this presentation, we explore the application of VarQITE, a quantum time evolution algorithm, to option pricing. Extending the work of Fontanela et al. (2021), we consider discretely monitored lookback options and use VarQITE to solve a partial differential equation associated to its price. To address the jump condition in the PDE, which poses a significant challenge in the quantum implementation, we re-write it in terms of multiple continuous equations, thus improving the accuracy of the results. A brief introduction to quantum computing will also be presented.

ROMAN MAKAROV, Wilfrid Laurier University
Spectral Expansions for Structural Credit Risk Models Incorporating Occupation Area and Occupation Time

We develop structural credit risk models with liquidation barriers and hazard rates driven by occupation time, occupation area, and their combinations. Defaults are classified according to Chapter 7 (liquidation) and Chapter 11 (reorganization) of the U.S. Bankruptcy Code. For a firm's value modelled as a diffusion with killing, we obtain a general closed-form representation of the associated Green's function. Using spectral methods, we derive a discrete spectral expansion of the transitional density, which in the geometric Brownian motion (GBM) case can be written in terms of Airy functions. This allows us to derive liquidation probabilities and implied hazard rates through spectral expansions. The framework extends to other solvable processes, including the constant elasticity of variance (CEV) model and state-dependent volatility hypergeometric diffusion models.

This is a continuation of the joint work with Giuseppe Campolieti and Hiromichi Kato.

ALEXANDER MELNIKOV, University of Alberta
On Market Completions Approach to Option Pricing and Related Questions

We consider a financial market with a reducible incompleteness. It means that the market can be embedded to a complete market by adding new risky assets. We call such embedding as a market completion. In the framework of such a market one can give a dual characterization of upper and lower option prices via maximization/minimization of expectations of discounted payoffs over market completions instead of martingale measures. Moreover, the method also works for the so-called indifference option pricing. To improve option price approximations, we explore a combination of the market completion method and machine learning technique in an incomplete jump-diffusion market model. Finally, we show how this approach work in life insurance applications.

ADAM METZLER, Wilfrid Laurier University
Comparing Life-Cycle and Contrarian Investment Strategies

Conventional wisdom holds that, when saving for retirement, individual investors should reduce their exposure to equities as retirement approaches. A recent strand of (high profile) literature criticizes this approach and purports to empirically demonstrate the superiority of so-called contrarian strategies, where exposure to equities is increased as retirement approaches. In this talk we formally (i.e. theoretically and rigorously) demonstrate that the underlying analysis is flawed and misleading, and prove that, within a certain parametric class, decreasing allocations (those that reduce their exposure to equities over time) strictly dominate increasing allocations in the mean-variance sense.

JINNIAO QIU, University of Calgary
Some recent progress on stochastic HJB equations

In this talk, we shall present some recent progress in the study of stochastic Hamilton-Jacobi-Bellman (HJB) equations, which arise naturally in the context of non-Markovian control problems, particularly within the field of mathematical finance. The non-Markovian nature of these problems may also involve path dependence or mean-field interactions, in addition to general randomness in the coefficients. The discussion will cover various aspects, including the well-posedness of such stochastic HJB equations, numerical approximations, and their applications.

MARK REESOR, Wilfrid Laurier University
Approximating the Money-Weighted Rate of Return

We develop a closed-form approximation to the so-called money-weighted rate of return (MWRR). The approximation is general in the sense that (i) it allows for contributions of varying sizes made at irregularly-spaced times (including both discrete and continuous contributions), (ii) it allows the composition of the underlying portfolio (as manifested through the mean and standard deviation of its instantaneous return) to vary through time and (iii) it does not make any specific assumptions on the stochastic dynamics of the underlying portfolio return. The approximation facilitates insights into a complicated object, which in turn allows us to explain and/or resolve findings elsewhere in the literature.

This is joint work with A. Metzler, M. Lau and D. Polegato.

ALEXANDRE ROCH, ESG UQAM
Optimal Green Transition for a Firm

I present a stochastic singular control problem that models a firm's optimal transition from Brown to Green technologies. Remaining in the Brown regime generates ongoing costs, while switching entails a proportional investment cost. The firm may distribute dividends but must maintain solvency through capital injections. Using viscosity-solution methods and comparison principles, I characterize the optimal transition policy and show that it is governed by endogenous threshold rules. Numerical experiments illustrate how parameters impact the viability and timing of the transition.

DAVID SAUNDERS, University of Waterloo
Exploratory Investment-Consumption with Non-Exponential Discounting

We extend the classic Merton optimal investment-consumption problem to the reinforcement learning (RL) framework. Additionally, we incorporate a general non-exponential discounting function to capture an investor's risk preferences, which leads to time inconsistency in the exploratory control problem. Under entropy regularization and logarithmic utility, we obtain closed-form equilibrium investment-consumption policies. Specifically, the optimal investment policy follows a Gaussian distribution, while the optimal consumption policy follows a Gamma distribution. Our results show that uncertainty about the discount rate leads the investor to adopt more conservative policies, with the Gaussian-distributed investment policy retaining the same mean but lower variance, and the Gamma-distributed consumption policy having both a lower mean and variance. We further develop and implement two RL algorithms- one based on the policy evaluation approach and the other on the q-learning approach- demonstrating their effectiveness through simulation studies. This is joint work with Y. Chen and Y. Li from the University of Waterloo.

ALEXANDER SCHIED, University of Waterloo
Exploring Roughness in Stochastic Processes: From Weierstrass Bridges to Volatility Estimation

Motivated by the recent success of rough volatility models, we introduce the notion of a roughness exponent to quantify the roughness of trajectories. It can be computed in a straightforward manner for many stochastic processes and fractal functions and also inspired the introduction of a new class of stochastic processes, the so-called Weierstrass bridges. After taking a look at Weierstrass bridges and their sample path properties, we discuss the relations between the roughness exponent and other roughness measures. We show furthermore that the roughness exponent can be statistically estimated in a model-free manner from direct observations of a trajectory but also from discrete observations of an antiderivative—a situation that corresponds to estimating the roughness of volatility from observations of the realized variance. This is joint work with Xiyue Han and Zhenyuan Zhang.

XIAOFEI SHI, University of Toronto
The Price of Information

When an investor is faced with the option to purchase additional information regarding an asset price, how much should she pay? To address this question, we solve for the indifference price of information in a setting where a trader maximizes her expected utility of terminal wealth over a finite time horizon. If she does not purchase the information, then she solves a partial information stochastic control problem, while, if she does purchase the information, then she pays a cost and receives partial information about the asset's trajectory. We further demonstrate that when the investor can purchase the information at any stopping time prior to the end of the trading horizon, she chooses to do so at deterministic time(s).

KRISTINA STANKOVA, University of Western Ontario
Applying ruin theory to retirement savings: A case study

In this talk, we will discuss how an advanced ruin theory model can be applied to retirement savings of individuals aiming at evaluating long-term risks related to their portfolios. To illustrate our approach, we fit the model to transactional data provided by a registered investment provider to the Financial Wellness Lab at Western University. We split the clients by gender and risk tolerance and examine how investment portfolios evolve over time in each group of clients.

LARS STENTOFT, University of Western Ontario
In estimation, the key is the volatility index, not the returns

This paper proposes a new methodology to estimate a GARCH model using only returns and volatility index (vli) data. The approach is centered on applying likelihood to the vli with an approximation to returns, denoted A-C-VIX-Ret, rather than the

standard approach of applying likelihood to returns with an approximation to the vli, denoted A-Ret-VIX. The new approach overcomes the ill-posed problem of an infinite likelihood from the vli, proposing the well-posed A-C-VIX for working with vli data only. We apply and compare the methodologies on three GARCH models, several volatility indexes, and stock data sets, with the main focus on the GARCH model of Heston & Nandi (2000), i.e., the HN-GARCH model, and the time series of S&P 500 and VIX from CBOE. Our analyses demonstrate that the volatility index holds more information on the parameters of a GARCH model than the returns, leading to A-C-VIX-Ret improving the quality of estimation compared to A-Ret-VIX, as seen by RMSE reductions of up to 90%; with significant improvements also in predicting the variance process. As part of the novelty, our methodology overcomes the problem of infinity-likelihood arising from Chi-squares with 1 degree of freedom, delivering a robust numerical procedure.

ANTONY WARE, University of Calgary

Generative Pricing of Basket Options via Signature-Conditioned Mixture Density Networks

We present a generative framework for pricing European-style basket options by learning the conditional terminal distribution of the log arithmetic-weighted basket return. A Mixture Density Network (MDN) maps time-varying market inputs—encoded via truncated path signatures—to the full terminal density in a single forward pass. Traditional approaches either impose restrictive assumptions or require costly re-simulation whenever inputs change. Trained on Monte Carlo (MC) under GBM with time-varying volatility or local volatility, the MDN acts as a reusable surrogate distribution: once trained, it prices new scenarios by integrating the learned density. Across maturities, correlations, and basket weights, the learned densities closely match MC (low KL) and produce small pricing errors, while enabling train-once, price-anywhere reuse at inference-time latency.

This is joint work with MD Hasib Uddin Molla, Ilnaz Asadzadeh and Nelson Fernandes

---

TING-KAM LEONARD WONG, University of Toronto

Excess growth rate and axiomatic characterizations

The excess growth rate is a fundamental logarithmic functional in portfolio theory. After reviewing its financial definition and properties, we present three axiomatic characterization theorems in terms of (i) the relative entropy, (ii) the gap in Jensen's inequality, and (iii) the logarithmic divergence that generalizes the Bregman divergence. We also consider maximization of expected excess growth rate and compare its solution with the growth optimal portfolio. Joint work with Steven Campbell.

FOIVOS XANTHOS, Toronto Metropolitan University

Star-Shaped Risk Measures: Representations and Cash-Additive Hulls

In this talk, we present representation results for star-shaped risk measures defined on general model spaces. We further investigate the cash-additive hulls of star-shaped risk measures and establish conditions under which these hulls preserve key continuity properties. The results provide new insights into the structure of Optimized Certainty Equivalents and Haezendonck–Goovaerts risk measures.

The talk is based on joint work with Denny Leung and Niushan Gao