Christian Bayer

Title: Optimal stopping with signatures

We propose a new method for solving optimal stopping problems (such as American option pricing in finance) under minimal assumptions on the underlying stochastic process X.
We consider classic and randomized stopping times represented by linear and non-linear functionals of the rough path signature $\mathbb{X}<\infty$ associated to X, and prove that maximizing over these classes of signature stopping times, in fact, solves the original optimal stopping problem. Using the algebraic properties of the signature, we can then recast the problem as a (deterministic) optimization problem depending only on the (truncated) expected signature. By applying a deep neural network approach to approximate the non-linear signature functionals, we can efficiently solve the optimal stopping problem numerically.

The only assumption on the process X is that it is a continuous (geometric) random rough path. Hence, the theory encompasses processes such as fractional Brownian motion, which fail to be either semi-martingales or Markov processes, and can be used, in particular, for American-type option pricing in fractional models, e.g. of financial or electricity markets.

(Based on joint work with Paul Hager, Sebastian Riedel, and John Schoenmakers).

Giacomo Bormetti

Title: Deep calibration

We propose a neural network-based approach to calibrating stochastic volatility models, which combines the pioneering grid approach by Horvath et al. (2021) with the pointwise two-stage calibration of Bayer and Stemper (2018). Our methodology inherits robustness from the former while not suffering from the need for interpolation/extrapolation techniques, a clear advantage ensured by the pointwise approach. The crucial point to the entire procedure is the generation of implied volatility surfaces on random grids, which one dispenses to the network in the training phase. We support the validity of our calibration technique with several empirical and Monte Carlo experiments for the rough Bergomi and Heston models. The approach paves the way for useful applications in financial engineering - for instance, pricing under local stochastic volatility models - and extensions to the fast-growing field of path-dependent volatility models.

Jean-Philippe Bouchaud

Title: The Inelastic Market Hypothesis:A Microstructural Interpretation

We will discuss Gabaix and Koijen’s recent Inelastic Market Hypothesis (IMH) from a microstructural point of view. We will review several empirical facts and arguments that give credence to the idea that market price fluctuations are mostly due to order flow, whether informed or non-informed. In particular, the Latent Liquidity Theory of price impact makes a precise prediction for GK’s multiplier M, which measures by how many dollars, on average, the market value of a company goes up if one buys one dollar worth of its stocks. Our central result is that M is of order unity, as found by GK, and increases with the volatility of the stock and decreases with the fraction of the market cap. traded daily.

Luca Capriotti

Title: 15 Years of Adjoint Algorithmic Differentiation (AAD): How to better hedge financial risks, crack some of the puzzles of condensed matter and much more with upside-down derivatives

Adjoint Algorithmic Differentiation (AAD) is a computational technique that, despite being known in its modern form since at least the 1960's, has become mainstream only in the last decade or so when it was "re-discovered" in Finance. In this talk, I will review what makes AAD an important innovation in financial risk management and how the same ideas can be applied in Condensed Matter Physics to speed up ab-initio molecular dynamics simulations or in any context in which computing accurately and efficiently a large number of derivatives is beneficial.

Damien Challet

Title: Filtering the covariance matrix of nonstationary systems with time-independent eigenvalues

We propose a data-driven way to reduce the noise of covariance matrices of time-varying systems. If the true covariance matrix is time-invariant, non-linear shrinkage of the eigenvalues is known to yield the optimal estimator for large matrices. Such a method outputs eigenvalues highly dependent on the inputs, as common sense suggests. When the covariance matrix is time-dependent, we show that it is generally better to use the set of eigenvalues that encode the average influence of the future on present eigenvalues resulting in a set of time-independent average eigenvalues. This situation is widespread in nature, one example being financial markets, where non-linear shrinkage remains the gold-standard filtering method. Our approach outperforms non-linear shrinkage both for the Frobenius norm distance, which is the typical loss function used on covariance filtering, and for financial portfolio variance minimization, which makes our method generically relevant to many problems of multivariate inference. Further analysis of financial data suggests that the expected overlap between past eigenvectors and future ones is systematically overestimated by methods designed for constant covariance matrices. Our method instead takes a simple empirical average of the eigenvector overlap matrix, which is enough to outperform non-linear shrinkage.

Giorgio Consigli

Benchmarking stochastic optimization approaches for pension fund management

We consider a defined benefit (DB) pension fund management problem formulated as a liability-driven investment problem subject to distributional uncertainty affecting specifically the term structure of interest rates, inflation and mortality intensity. By its nature a DB pension plan has to deal with uncertainties that will only be resolved far into the future and lead naturally to distributionally robust optimization (DRO) formulations carrying a complex and necessarily long-term risk exposure. The ambiguous stochastic optimization approach assumes that the underlying probability distribution is unknown and lies in an ambiguity set of probability distributions.

The DRO approach is analysed in comparison with alternative modelling frameworks for pension fund asset-liability management (ALM) problems and more generally quantitative ALM problems that depending on the underlying model of uncertainty, are formulated as dynamic stochastic control (DSC) or multistage stochastic programming (MSP) problems or as a combination of the two (MSP-DSC).

A set of comparative computational evidence is presented by focusing on the funding ratio as core PF variables.

Masaaki Fukasawa

Title: When to efficiently rebalance a portfolio

A constant weight asset allocation is a popular investment strategy and is optimal under a suitable continuous model. We study the tracking error for the target continuous rebalancing by a feasible finite-time rebalancing under a general Brownian semimartingale model of asset prices. In a high-frequency asymptotic framework, we derive an efficient strategy among simple predictable processes. This is a joint-work with Masayuki Ando.

Julien Guyon

Title: Volatility Is (Mostly) Path-Dependent

We learn from data that volatility is mostly path-dependent: up to 90% of the variance of the implied volatility of equity indexes is explained endogenously by past index returns, and up to 65% for (noisy estimates of) future daily realized volatility. The path-dependency that we uncover is remarkably simple: a linear combination of a weighted sum of past daily returns and the square root of a weighted sum of past daily squared returns with different time-shifted power-law weights capturing both short and long memory. This simple model, which is homogeneous in volatility, is shown to consistently outperform existing models across equity indexes and train/test sets for both implied and realized volatility. It suggests a simple continuous-time path-dependent volatility (PDV) model that may be fed historical or risk-neutral parameters. The weights can be approximated by superpositions of exponential kernels to produce Markovian models. In particular, we propose a 4-factor Markovian PDV model which captures all the important stylized facts of volatility, produces very realistic price and volatility paths, and jointly fits SPX and VIX smiles remarkably well. We thus show, for the first time, that a continuous-time Markovian parametric stochastic volatility (actually, PDV) model can practically solve the joint S&P 500/VIX smile calibration problem.

Blanka Horvath

Title: Robust hedging GANS

The deep hedging framework presented in BuehlerGononTeichmannWood19 has opened new horizons for solving hedging problems under a large variety of models and market conditions. At the same time, any model -- be it a traditional stochastic model or a market generator -- is at best an approximation of market reality, prone to model-misspecification and estimation errors. This raises the question, how to furnish a deep hedging setup with tools that can address the risk of the discrepancy between model and market reality in an automated way. We do so by treating model ambiguity in a ``weighted worst-case'' fashion, that is penalizing deviations (measured in an appropriate metric) from our believed distribution. Automated robustification -- due to its direct implications on risk management -- is a delicate task, which will no-doubt inspire a range of solutions to related challenges.

Antoine Jacquier

Title: Some quantum algorithms in Finance

We introduce several algorithms from Quantum Computing technologies aimed at providing speedup compared to their classical counterparts. We shall highlight in particular how quantum entanglement provides potential expressive explanatory power for neural networks. Time permitting, we will showcase further applications to linear systems and PDEs and to optimization problems.

Sebastian Jaimungal

Title: Risk Budgeting Allocation for Dynamic Risk Measures

We develop an approach for risk budgeting allocation – a risk diversification portfolio strategy – where risk is measured using time-consistent dynamic risk measures. For this, we introduce a notion of dynamic risk contributions that generalise the classical Euler contributions and which allow us to obtain dynamic risk contributions in a recursive manner. Moreover, we show how the risk allocation problem may be recast as a convex optimisation problem and develop an actor-critic approach to solve for risk allocations using deep learning techniques.

(Based on joint work with Silvana Pesenti, Yuri Saporito, and Rodrigo Targino)

Jessica James

Title: Beyond Convexity

Who says fixed income is boring? In 2022 100y bonds were exhibiting price swings larger than in many speculative stocks while yields were still in a relatively confined range. We derive terms beyond duration and convexity and demonstrate how they become essential at times when larger yield moves become more likely. Remarkably, we also illustrate that these instruments can be hedged to a significant extent.

Woo Chang Kim

Title: Deep Financial Planning

Professor Michael Dempster has pioneered and made a significant advancement in solving large-scale stochastic programming problems including the use of decomposition techniques, sample average approximation, and scenario reduction methods. He has also developed software for solving stochastic programming problems and has applied these methods to a wide range of practical problems, such as portfolio optimization, energy systems, and telecommunications. His work is currently serving as the foundations for related research. Based on his seminal works, we propose a framework that utilizes machine learning in stochastic programming. Specifically, we establish theoretical foundations for approximating value functions and optimal policies through Neural Networks by providing conditions under which the Universal Approximation Theorem can be applied to parametric optimization problems. This study addresses a gap in the analysis of when Neural Networks can be effectively used as approximators for parametric optimization and demonstrates that the approximation error decreases as the amount of training data increases. Based on this analysis, we introduce the 'Deep Financial Planning (DFP)' model, which combines stochastic optimization and Deep Neural Networks (DNNs) to solve financial planning problems, and show that it performs well in a variety of settings. Additionally, we propose the use of transfer learning in DFP, resulting in improved performance with less data. These results suggest that DFP has the potential to provide real-time, personalized lifetime financial planning in a scalable manner.

(Based on joint work with Hyunglip Bae, Jinkyu Lee, Sanghyun Bae, and Yongjae Lee).

Natalie Packham

Title: Correlation scenarios and correlation stress testing

We develop a general approach for stress testing correlations of financial asset portfolios. Based on methods from interest rate modelling, the correlation matrix of asset returns is specified in a parametric form, where correlations are represented as a function of risk factors, such as country and industry factors. A sparse factor structure linking assets and risk factors is built using Bayesian variable selection methods. Regular calibration yields a joint distribution of economically meaningful stress scenarios of the factors. As such, the method also lends itself as a reverse stress testing framework: using the Mahalanobis distance or Highest Density Regions (HDR) on the joint risk factor distribution allows to infer worst-case correlation scenarios. We give examples of stress tests on a large portfolio of European and North American stocks. In an outlook, I will present some ideas on how to aggregate existing risk factors giving rise to further risk scenarios.

Wim Schoutens

Title: Sustainable Finance and ESG Investing: Sense or nonsense

We take a critical look at the current trends of sustainable finance and ESG investing. We elaborate on green bonds, sustainable linked bonds and other alternatives. We assess whether statements concerning lower risks and higher returns are true in a constrained setting of sustainable portfolio investments. We elaborate on the extension of the classical portfolio theory and the optimal frontiers from a two dimensional setting (risk and return) into a third ESG dimension. We also touch upon issues of greenwashing. We finally comment on potential systemic risks arising from an enforcement of sustainable investment into the regulation.

Sasha Stoikov

Title: Where Market Making Meets Market Microstructure

How can a market making algorithm use information in the order book when computing bid and ask quotes? Market making models, such as Avellaneda and Stoikov (2008), compute bids and asks around the midprice, to minimize inventory risk. In practice, the midprice may be a poor estimate of the fair value, particularly for cryptocurrencies, where the tick size is relatively small. Using Bitcoin data, I backtest market-making strategies around the midprice, as well as other microstructure adjusted prices. In particular, a new definition of the fair price, which we call the Volume Adjusted Mid Price (VAMP) consistently outperforms the mid price, from the perspective of a market maker.

Claudio Tebaldi

Title: The origins of scaling and power law fluctuations in a competitive equilibrium

First-order optimality conditions in general equilibrium models establish an empirically testable relationship between agents' beliefs, preferences, and security return moment dynamics. These restrictions are hard to reconcile with empirically well-documented scale invariant security returns' statistics at short- and long-term horizons if preferences and beliefs are assumed to be time consistent. On the contrary, we show that scale invariant properties arise naturally introducing an alternative assumption that we name scale consistency. Remarkably, the scale consistency property has a micro foundation: at an individual level it rationalizes the present-bias and the myopic behavior that has been documented by field experiments. We illustrate the implications and highlight the relevance of scale consistent first-order conditions in a number of financial applications.