Abstracts

Professor Perry Mehrling

INET & Barnard College, Columbia University

This talk describes progress towards the development of a theory of money that takes as its starting point the inside credit character of modern money, in order to integrate monetary economics with modern finance, and so to contribute to contribute to a money theoretical basis for macroeconomics and policy.

Professor Didier Sornette

Management, Technology & Entrepreneurship, ETH Zurich

The 2007-2008 financial crisis, you might think, was an unpredictable one-time crash. But this talk presents the mathematics and statistics used to develop a set of early warning signs for unstable, growing systems, tracking the moment when any bubble is about to pop. (And it is happening again, right now.)

Professor Mike Giles

Oxford-Mann Institute, Oxford University

The big development in HPC in the last few years has been the emergence of many core GPUs. NVIDIA's current Kepler GPUs have up to 2880 cores on a single chip, and this is complemented by the CUDA software development environment which is based on C/C++ with a few extensions. Leading financial institutions such as Bloomberg, BNP Paribas and JP Morgan have employed GPUs to greatly reduce the time and energy required for pricing , hedging and risk management calculations. This talk discusses GPU efficient use of the underlying Monte Carlo and finite difference algorithms.

Jacques du Toit

Software Developer, Numerical Algorithms Group

The talk describes the use of adjoint algorithmic differentiation running on GPUs to price and hedge FX basket options using a local volatility model.

Professor Antoon Pelsser

Business & Economics, Maastricht University & Kleynen Consultants

The Least Squares Monte Carlo (LSMC) method is widely applied to solve stochastic optimal control problems, such as pricing American-style options. A central part of LSMC is the approximation of conditional expectations across each time-step. Conventional algorithms regress the value function at the end of the time-step on a set of basis functions, which are measurable with respect to the information available at the beginning of the time-step. The corresponding regression error has two sources: an approximation error due to the finite number of basis functions, and a projection error due to the projection onto the coarser filtration at the beginning of the interval. The convergence speed for the conventional algorithms is determined by the projection error component, which converges relatively slowly. Glasserman and Yu (2002) propose the Regress-Later method, wherein the value function at the end of the time-step is regressed on a set of basis functions, which are

measurable with respect to the information available at the end of the time-step. The conditional expectation across the time-step is then computed analytically for each basis function. We show in this paper that by using Regress-Later the projection error component is removed. This implies that the Regress Later method has the potential of converging significantly faster than the conventional algorithms. We provide sufficient conditions for achieving fast convergence on compact and non-compact sets and we give an explicit example.

Professor Elisa Nicolato

Economics & Business, Aarhus University

We consider a tractable affine stochastic volatility model that generalizes the seminal Heston model by augmenting it with jumps in the instantaneous variance. In this framework, we consider options written on the realized variance and we examine the impact of the distribution of jumps on the associated implied volatility smile. The model allows us to isolate the unique impact of the jump distribution. We provide sufficient conditions for the asymptotic behavior of the volatility of variance for small and large strikes. In particular, we show that by selecting alternative jump distributions, one obtains fundamentally different shapes of the implied volatility of variance smile -- some clearly at odds with the upward-sloping volatility skew observed in variance markets.

Professor Juho Kanniainen

Industrial Engineering, Technical University of Tampere

This paper uses a continuous time model to show that under a transversality condition equity dividend yield becomes dependent on stochastic volatility and consequently, squared return volatility can have a negative relationship with the value of deep-in-the-money equity call options. Cloud computing is used to identify the underlying mechanism using forward information on option contracts. The results indicate that the prevailing practice of ignoring time varying dividend yields in option pricing can lead to oversimplification of stock market dynamics.

Professor Teemu Pennanen

Mathematics, Kings College London

This paper studies optimal investment and contingent claim valuation in markets where illiquidity may affect the transfer of wealth over time and between investment classes. In addition to classical frictionless markets and markets with transaction costs, our model covers nonlinear illiquidity effects that arise in limit order markets. We extend basic results on arbitrage bounds and attainable claims to illiquid markets and general swap contracts where both claims and premiums may have multiple payout dates. We establish the existence of optimal trading strategies and the lower semicontinuity of the optimal value of optimal investment under conditions that extend the no-arbitrage condition in the classical linear market model. All results are derived with the \direct method" without resorting to duality arguments.

Dr Rob Coles

Equity Derivatives, Citigroup

The volatility target (sometimes called volatility control) technique aims to construct an underlying with constant target volatility by dynamically adjusting a basket between risky and riskless assets according to the realised volatility of the risky asset. The technique has become increasingly popular and families of 'Risk Control' indices are now available. Despite this there is little information on the risks to an option seller. We derive simple approximations to the greeks, the residual skew and the implied volatility of options on a volatility target underlying. These give some insight into how the

product should be hedged, its sensitivity to jumps and the compromises to be faced in defining the measure of realised volatility.

Dr Xiao-Jun Zeng

Computer Science, Manchester University

Behavioural economics tells us that emotions can profoundly affect individual behaviour and decision -making. Does this also apply to societies at large, i.e. can societies experience mood dtates that effect their collective decision-making? By extension is the public mood correlated with or even predictive for economic indicators? Here we investigate whether measurements of collective mood states derived from large-scale Twitter feeds are correlated with the value of the Dow-Jones Industrial Average (DJIA) over time. Our results indicate that the accuracy of DJIA predictions can be significantly improved by the inclusion of specific public mood dimensions but not others. We find an accuracy of 86.6% in predicting the daily up and down changes in the closing level of the DJIA and a reduction of the mean average prediction error by 6%.

Professor Damiano Brigo

Mathematics, Imperial College London

We consider nonlinearities in funding costs, margining and gap credit risk and their operational implications in valuation. In particular we discuss

- Asymmetric borrowing and lending rates
- Nonlinear pricing operators
- Deal dependent pricing measures
- Aggregation dependent valuation
- The XVA's practice vs. lack of separability
- Operational implications

Professor Mark Davis

Mathematics, Imperial College London

Recently there has been renewed debate about the relative merits of VaR and CVaR as measures of financial risk. VaR is not coherent and does not quantify the risk beyond VaR, while CVaR shows some computational instabilities and is not "elicitable" (Gneiting 2010, Zwiebel 2013). It is argued in this talk that such questions are best addressed from the point of view of probability forecasting or Dawid's "prequential statistics". We introduce a concept of "consistency" of a risk measure, which is close to Dawid's "strong prequential principle", and show that VaR indeed has special properties not shared by any other risk measure.

Dr Luca Capriotti

Quantitative Strategies, Credit Suisse

Adjoint algorithmic differentiation can be used can be used to implement efficiently the calculation of counterparty credit risk. We demonstrate how this powerful technique can be used to reduce the computational cost by hundreds of times, thus opening the way to real time risk management in Monte Carlo.

Alexander Denev

Senior Analyst, Risk Dynamics

After an introduction to accepted methods for modelling and forecasting extreme events based on historical data, this talk will give an overview of forward looking probabilistic graphical models and their construction. This technique will be illustrated by a range of case studies of applications including scenario analysis, credit risk, tail hedging, optimal portfolios, asset pricing and network theory.

Hicham Lahlou

CEO, Xcelerit

Xcelerit is a leading software provider of cross-platform acceleration tools for financial services, engineering, and research. Xcelerit technology allows Quantitative Analysts to unlock the performance of accelerators (GPUs and multi-core) with minor modifications to their existing source code. In this talk, Hicham Lahlou, Xcelerit CEO & Co-Founder, shows how to design financial software for both performance and productivity on the example of a real-world Credit Value Adjustment implementation.

Dr Stephen Weston

CDO, Maxeler Technologies

Maxeler solutions exploit dataflow computing – a revolutionary way of performing computation, completely different to computing with conventional CPUs. Dataflow computers focus on optimizing the movement of data in an application and utilize massive parallelism between thousands of tiny 'dataflow cores' to provide order of magnitude benefits in performance, space and power consumption. An analogy for moving from control flow to dataflow is the Ford car manufacturing model, where expensive highly-skilled craftsman (control flow CPU cores) are replaced by a factory line, moving cars through a sea of single-skill workers (dataflow cores). We discuss Monte Carlo applications to VaR, CVS and IRS which include hardware, software and services to gain at least an order of magnitude advantage in performance per unit of rack space, computations per Watt (green computing), and top price-performance considering total cost of ownership for monolithic applications.

Erik Vynckier

CIO - Insurance, Alliance Bernstein

Asset management has faced a rapidly changing environment since the 2008-9 crisis. This talk will discuss some of the critical issues. Specifically:

- Capital management and market conduct regulation, Basel III, Solvency II, Dodd-Frank, EMIR
- Yield enhancement and illiquidity
- The future of mortgage finance
- Pensions and insurance for the 21st century

Dr Steven Morrison

Research Director, Moodys Analytics

Stochastic (Monte Carlo simulation) models are used extensively by life insurance companies to value their liabilities to policyholders, and to project these values in the future. The computational demands of these calculations have led to the recent development of 'proxy' methods which approximate the calculation in a fraction of the time. Focusing on the Least Squares Monte Carlo proxy method, this talk will demonstrate how these new techniques are now being successfully applied in areas such as regulatory capital assessment, business planning and appraisal of hedging strategies.

Dr Mark Cathcart

Quantitative Analytics, Standard Life

Hedging methods to mitigate the exposure of variable annuity products to market risks require the calculation of market risk sensitivities (or "Greeks"). The complex, path-dependent nature of these products means these sensitivities typically must be estimated by Monte Carlo simulation. Standard market practice is to measure such sensitivities using a "bump and revalue" method. As well as requiring multiple valuations, such approaches can be unreliable for higher order Greeks, e.g., gamma. In this article we investigate alternative estimators implemented within an advanced economic scenario generator model, incorporating stochastic interest-rates and stochastic equity volatility. The estimators can also be easily generalized to work with the addition of equity jumps in this model.

Mario Skoric & Dr Angelo Uristano

Investment Management, Allianz, Munich

The financial environment for global property and casualty (P/C) insurers has changed dramatically in recent years mainly due to: record P/C claims, historically low fixed income and real estate returns, ongoing introduction of regulatory, risk-based, capital requirements (Solvency II) and increasing market competition. Jointly or individually those pressures have determined increasing liquidity deficits in core P/C activity, stressed technical and financial scenarios and stimulated integration between insurance, investment and risk management divisions of insurance corporations. Using dynamic stochastic optimization methods we present an analysis of the specific problem of an investment manager subject to technical and maximum risk exposure constraints employing a dynamic portfolio strategy over a 10 year horizon. Key to the practical solution are: long-term investment targets, the enlargement of the investment universe to include alternative investments, the control of risk capital exposure over time and the sensitivity of short, medium and long term goals to stressed insurance and financial scenarios.

Dr Francesco Sandrini

Pioneer Investments, Munich

This talk presents the fiduciary solution recently provided to the German pension fund Pensionkasse using dynamic stochastic optimization to manage the asset liability problem facing the fund to remain 100% funded through difficult market conditions. We find that going forward the probability to have funding gap remain positive (representing a surplus) at a 1 year horizon is 98%, which is due to the total asset portfolio wealth decreasing because of the combined effect of lower cash flows from the fixed income direct investments and the return of the remaining asset portfolio that is very weak and for some assets negative. However, the funding gap is positive with around 100% probability at the other intermediate and at final horizons.

Dr Elena Medova

Pure Mathematics & Statistics, Cambridge University & MD, Cambridge Systems Associates

Recent times have seen the transfer of risks to individual households as the unintended consequences of pension reforms and new rules for regulation of financial advisory companies. We present new theory and technology for individual Asset Liability Management (iALM) which addresses these critical issues for individuals. The iALM tool uses a cutting edge dynamic stochastic programming formulation based on ideas from behavioural finance and decision theory. The resulting system will be Illustrated in terms of solutions for representative households. The technological challenges of practical mass implementation requiring current HPC solutions such as the cloud and the progress in this direction made to date will also be discussed.

Dr Sylvain Corlay

Quantitative Research, Bloomberg, New York

This talk concerns the theory and some applications of finite dimensional approximation of infinite dimensional function spaces. In particular it will cover:

- Background on functional quantization
- Cubature for pricing with a stochastic volatility application to long-memory stochastic volatility models
- Functional quantization-based stratified sampling and generalized bridges.

Professor Andrew Rau-Chaplin

Computer Science, Dalhousie University & Willis Research Network

Risk hedging strategies are at the heart of financial risk management. As with many financial institutions, insurance companies try to hedge their risk against potentially large losses, such as those associated with natural catastrophes. Much of this hedging is facilitated by engaging in risk transfer contracts with the global reinsurance market. Devising an effective hedging strategy depends on careful data analysis and optimization. Such an insurance company faces a reinsurance contract optimization problem in which a given reinsurance contract consists of a fixed number of contractual layers and requires a simulated set of expected loss distributions (one per layer) and a model of reinsurance market costs. The task is to identify optimal combinations of placements such that for a given expected return the associated risk value is minimized. The solution to this high-dimensional multi-objective data analysis and optimization problem is a Pareto frontier that quantifies the best available trade-offs between expected risk and returns. One approach to this reinsurance contract optimization problem is to adapt the evolutionary heuristic search method called Population Based Incremental Learning (PBIL), to work with discretized solution spaces. Our multi-threaded Discretized PBIL (DiPBIL) method is able to solve larger “real world” problem instances than previous methods. For example, using clusters problems with a 5% discretization and 7 or less contractual layers can be solved in about an hour, while previously infeasible problems that would have taken weeks or even months to run with as many as 15 layers can be solved in less than a day. Cloud and GPU implementations yield even faster solutions.