This event is jointly organised by Professor Johannes Ruf, Professor of Mathematics and Deputy Head (Teaching) in the Department of Mathematics and at the London School of Economics, and Professor Carmine Ventre, Professor of Computer Science, Head of the Department of Informatics and Director of the King's Institute for Artificial Intelligence, King's College London.

Confirmed speakers:

Dr Silvia Bartolucci is Associate Professor of Computer Science, in the Department of Computer Science, University College London, where she is part of the Financial Computing and Analytics Group.

Abstract: *Deep Learning for Limit Order Book Forecasting*

Limit Order Book (LOB) markets drive modern electronic rading, where buy and sell orders are dynamically matched to determine asset prices. Forecasting price movements in this high-frequency setting is inherently complex due to the high-dimensional and noisy nature of LOB data. This presents both a challenge and an opportunity for deep learning models, which can extract hidden patterns but must be carefully evaluated for real-world applicability.

By analysing a heterogeneous set of stocks traded on the NASDAQ exchange in this talk we will I) assess stocks' predictability in relation to their microstructural features, such as their so-called tick size, ii) extend model's evaluation beyond standard machine learning metrics to consider the feasibility of actual trading strategies carried out using the forecast.

Finally, we will discuss the integration of deep learning architectures with network-theoretic representations that capture complex and non-trivial dependency structures among volume levels. This approach offers new insights into the spatial distribution and temporal degradation of information in LOBs, bridging the gap between microstructural modelling and deep learning-based forecasting in high-frequency financial markets.

Dr James Hamp is Director and Head of FX Data Strategy and Analytics, Citigroup London, where he leads the development of data-driven analytics and automation for the FX options trading business.

Abstract: To follow.

Dr Christian Julliard is Associate Professor of Finance in the Department of Finance, at the London School of Economics.

Abstract: *The Co-Pricing Factor Zoo.*

We analyze 18 quadrillion models for the joint pricing of corporate bond and stock returns. Only a handful of factors, behavioural and nontradable, are robust sources of priced risk. Yet, the true latent stochastic discount factor is dense in the space of observable factors. A Bayesian Model Averaging Stochastic Discount Factor (BMA-SDF), combining the corporate bond and stock factor zoos, explains risk premia better than all existing models, both in- and out-of-sample. We show that multiple factors are noisy proxies for common underlying sources of risk, and the BMA-SDF aggregates them optimally. The SDF, as well as its conditional mean and volatility, are persistent, track the business cycle and times of heightened economic uncertainty, and predict future asset returns. Finally, we show that stock factors price the credit component of corporate bond excess returns well, while the Treasury component is priced almost exclusively by the bond factors.

Dr Bart De Keijzer is Senior Lecturer in Computer Science, in the Department of Informatics, King's College London, where he is Deputy Director of the Distributed AI Research Group and Informatics Finance Hub Champion.

Abstract: *Clearing in Financial Networks with Credit Default Swaps: Hardness and Algorithms.*

I will give an overview of our research on the clearing problem for the extension of the Eisenberg-Noe model of financial networks where credit default swaps (CDSes) may be present in the network. The process of clearing a financial network identifies insolvent banks in the network and can be used as a tool to evaluate a network's exposure to systemic risk. We prove that clearing is a computationally very difficult task when CDSes are present, as the problem requires the ability to compute fixed points of arbitrary functions defined by any composition of standard arithmetic operations.

In this talk, I will explain the main ideas required to prove this fact, and will discuss algorithms that can clear a network efficiently when certain natural structural conditions on the network are imposed.

Lastly, I will present a mixed linear integer program for the problem, for which we show experimentally that modern solvers (such as Gurobi) can find a solution quickly, for reasonably large financial networks.

Joint work with Stavros Ioannidis, Jinyun Tong, and Carmine Ventre.