Numeraire Portfolio and Weighted Monte Carlo

Bruno Dupire

This piece of research is a tribute to our friend and co-organizer Marco Avellaneda who left us in 2022. It first reviews Marco's work on the Weighted Monte Carlo (WMC) method, which minimally reweights simulated paths such that a set of traded assets is priced correctly. Then it establishes links between the Numeraire Portfolio, a combination of traded assets such that the price of any other non-dividend paying asset expressed in this numeraire is a martingale under the real world probability measure. The Numeraire Portfolio induces weights that reprice the market instruments correctly and provides prices for non traded ones as well.

We apply it to vanilla options of different strikes and analyze the interpolation it provides by pricing options of non traded strikes. We then address the loss of martingality it generates and propose remedies by adding "martingale claims" including signature claims.

From the Quintic model to signature volatility models

Eduardo Abi Jaber

Abstract: We will introduce the Quintic Ornstein-Uhlenbeck model that jointly calibrates SPX-VIX options with a particular focus on its mathematical tractability namely for fast pricing SPX options using Fourier techniques. Then, we will consider the more general class of stochastic volatility models where the dynamics of the volatility are given by a possibly infinite linear combination of the elements of the time extended signature of a Brownian motion. First, we show that the model is remarkably universal, as it includes, but is not limited to, the celebrated Stein-Stein, Bergomi, and Heston models, together with some path-dependent variants. Second, we derive the joint characteristic functional of the log-price and integrated variance provided that some infinite-dimensional extended tensor algebra valued Riccati equation admits a solution. This allows us to price and (quadratically) hedge certain European and path-dependent options using Fourier inversion techniques. We highlight the efficiency and accuracy of these Fourier techniques in a comprehensive numerical study

Markovian approximations to rough volatility models

Christian Bayer, Simon Breneis

The rough Heston model is a very popular recent model in mathematical finance; however, the lack of Markov and semi-martingale properties poses significant challenges in both theory and practice. A way to resolve this problem is to use Markovian approximations of the model. Several previous works have shown that these approximations can be very accurate even when the number of additional factors is very low. Existing error analysis is largely based on the strong error, corresponding to the L^22 distance between the kernels.

Extending earlier results by [Abi Jaber and El Euch, 2019], we show that the weak error of the Markovian approximations can be bounded using the L^11-error in the kernel approximation for general classes of payoff functions for European style options. Moreover, we give specific Markovian approximations which converge super-polynomially in the number of dimensions, and illustrate their numerical superiority in option pricing compared to previously existing approximations. The new approximations also work for the hyper-rough case $H > -1/2$.

Bank Liquidity Management and Payout Policy under Peer Pressure

Diogo Duarte, Ozde Oztekin, Yuri F. Saporito

We present a theoretical model that examines the effects of peer pressure on the trade-off that bank managers face when deciding whether to accumulate reserves or to pay shareholders. We show that under high peer pressure, banks reduce payouts and increase cash reserves, which reduces the probability of default and improves financial stability. Using data from the Federal Reserve's Y-9C report from 1987-2020, we find that a one standard deviation increase in peer pressure corresponds to a 2%-12% rise in cash reserves and a 1%-17% decline in the dividend payout ratio, relative to their respective averages, thereby leading to improved bank risk profiles.

Neural joint S&P 500/VIX smile calibration

Julien Gyon, Scander Mustapha

Abstract : We calibrate neural stochastic differential equations jointly to S&P 500 smiles, VIX futures, and VIX smiles. Drifts and volatilities are modeled as neural networks. Minimizing a suitable loss allows us to fit market data for multiple S&P 500 and VIX maturities. A one-factor Markovian stochastic local volatility model is shown to fit both smiles and VIX futures within bid-ask spreads. The joint calibration actually makes it a pure path-dependent volatility model, confirming the findings in [Guyon, 2022, The VIX Future in Bergomi Models: Fast Approximation Formulas and Joint Calibration with S&P 500 Skew]. Large volatility, fast mean-reversion, and pure path-dependency of the volatility are all learned from scratch.

Regulatory arbitrage and financial stability: how the new output floor compares to other capital requirements

Lakshithe Wagalath

Basel 3 banking regulation has introduced two "new" instruments, the output floor and the leverage ratio, which were designed to complement existing risk-based capital instruments and limit regulatory arbitrage that may occur from the use of banks' internal models when computing risk-weighted assets. In this study, we show how the output floor effectively enables to limit regulatory arbitrage – which we formally define in line with regulators – compared to the leverage ratio. We also quantify how banks react to large shocks on their

assets in order to comply with the output floor and we compare the results with the case of leverage ratio and risk-based capital ratio constraints enabling us to explore the effects of these capital requirements on financial stability and contagion effects after large shocks.

Estimation risk in conditional expectiles

Marcelo Fernandes, Victor Henriques, Duda Mendes

We establish the consistency and asymptotic normality of a two-step estimator of conditional expectiles in the context of location-scale models. We first estimate the parameters of the conditional mean and variance by quasi-maximum likelihood and then compute the unconditional expectile of the innovations using the empirical quantiles of the standardized residuals. We show how replacing true innovations with standardized residuals affects the asymptotic variance of the expectile estimator. In addition, we also obtain asymptotic-valid bootstrap-based confidence intervals. Finally, our empirical analysis reveals that conditional expectiles are very interesting alternatives to assess tail risk in cryptomarkets, relative to traditional quantile-based risk measures, such as value at risk and expected shortfall.

Optimal Execution among N Traders with Transient Price Impact

Marcel Nutz, Steven Campbell

We study an N-player game of optimal execution in the Obizhaeva—Wang model of transient price impact. A unique equilibrium exists when the problem is regularized by an additional instantaneous cost, but existence fails without regularization. We show how a modification of the model restores existence. The newly found equilibrium is consistent with the limit for small regularization, and has a simpler form.

Causal Discovery in Financial Markets: What is it?

Mohammad Fesanghary

A deeper understanding of financial markets requires understanding not only the statistical dependencies among various entities but also the causal dependencies. In this talk, we introduce the (basic) concepts of causal discovery which is how we can learn causal dependencies directly from data without randomized control trials (RCT). We will also cover a new algorithm we developed called CD-NOTS (Constraint-based Causal Discovery from Nonstationary Time-Series) which is able to handle non-stationarity, non-linearity, non-Gaussianity, and lagged and contemporaneous effects, all of which are typical in financial data. Finally, we show the results of CD-NOTS on Synthetic and real data and show that our algorithm is a more effective alternative to other causal discovery algorithms.

Insider Trading with Temporary Price Impact

Ryan Donnelly

We model an informed agent with information about the future value of an asset trying to maximize profits when the agent's trades are subjected to a transaction cost as well as a market maker tasked with setting fair transaction prices. In a single auction model, equilibrium is characterized by the unique root of a particular polynomial. Analysis of this polynomial with small levels of risk-aversion and transaction costs reveal a dimensionless parameter which captures several orders of asymptotic accuracy of the equilibrium behaviour. In a continuous time analogue of the single auction model, incorporation of a transaction costs allows the informed agent's optimal trading strategy to be obtained in feedback form. Linear equilibrium is characterized by the unique solution to a system of two ordinary differential equations, of which one is forward in time and one is backward. When transaction costs are in effect, the price set by the market maker in equilibrium is not fully revealing of the informed agent's private signal, leaving an information gap at the end of the trading interval. When considering vanishing transaction costs, the equilibrium trading strategy and pricing rules converge to their frictionless counterparts.

Partial Information Nash Equilibria between Broker and Traders

Sebastian Jaimungal, Álvaro Cartea, Leandro Sanchez-Betancourt, Xuchen Wu

We study the partial information Nash equilibrium between a broker and an informed trader. In our model, the broker trades in the lit exchange where trades have instantaneous and transient price impact with exponential resilience. The informed trader trades solely with the broker but has an additional trading signal that the broker lacks. We characterise the Nash equilibrium of the trading strategies as the solution to a coupled system of FBSDEs and prove uniqueness and existence. Further, we provide insights on the trading strategies when compared to the Stackleberg equilibria.

Portfolio choice with α-Bregman Wasserstein penalisation

Silvana Pesenti

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.

Duality in Convex Stochastic Optimization

Teemu Pennanen

We study a general class of convex stochastic optimization
problems that unifies many common problem formulations from operations
research, financial mathematics and stochastic optimal control. We develop a general duality framework that allows for a unified and simplified treatment of various special problem classes found in the literature. In particular, we establish the existence of primal solutions and the absence of a duality gap without compactness or boundedness assumptions. In the context of financial mathematics, the relaxed assumptions are satisfied under the well-known no-arbitrage condition and the reasonable asymptotic elasticity condition of the utility function. The existence of dual solutions is established in the general format under conditions that extend those given by Rockafellar and Wets for stochastic problems of Bolza. The extended theory allows also for significant generalizations to existing problem formulations in financial mathematics and elsewhere.

The Price of Information

Xiaofei Shi, Sebastian Jaimungal

Abstract: 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).