BOOK OF **ABSTRACTS**

PISA
FINANCE
Week
2026

POLO DIDATTICO
DELLE PIAGGE
MAY 13~15
PISA, ITALY

pisa-finance-week.ec.unipi.it

Participants

Paulin Aubert, Université d'Evry Val d'Essonne, France;
Alessandro Bondi, Luiss, Italy;
Alessia Cafferata, Università di Pisa, Italy;
Giorgia Callegaro, Università di Padova, Italy;
Riccardo Cambini, Università di Pisa, Italy;
Laura Carosi, Università di Pisa, Italy;
Etienne Chevalier, Université d'Evry Val d'Essonne, France;
Cristina Di Girolami, Università di Bologna, Italy;
Giovanna D'Inverno, Università di Pisa, Italy;
Rachele Foschi, Università di Pisa, Italy;
Celine Labart, LAMA - Université Savoie Mont Blanc, France;
Jerome Lelong, Grenoble INP - UGA, France;
Vathana Ly Vath, ENSIIE, France;
Alberto Ohashi, Universidade de Brasilia, Brazil;
Huyen Pham, Ecole Polytechnique, France;
Sergio Pulido, ENSIIE, France;
Anthony Reveillac, INSA Toulouse and Institut de Mathematiques de Toulouse, France;
Cristiano Ricci, Università di Pisa, Italy;
Alexandre Roch, Université du Quebec à Montreal, Canada;
Johannes Ruf, London School of Economics, United Kingdom;
Simone Scotti, Università di Pisa, Italy, and Paris Cité, France;
Carlo Sgarra, Università di Bari, Italy;
Mehdi Talbi, Université Paris Cité, Italy;
Emanuele Vannucci, Università di Pisa, Italy;

Contents

Optimal dividend and capital injection under self-exciting claims

Paulin Aubert

Exiom Partners, France

In this paper, we study an optimal dividend and capital-injection problem in a Cramér-Lundberg model where claim arrivals follow a Hawkes process, capturing clustering effects often observed in insurance portfolios. We establish key analytical properties of the value function and characterise the optimal capital-injection strategy through an explicit threshold. We also show that the value function is the unique viscosity solution of the associated HJB variational inequality. For numerical purposes, we first compute a benchmark solution via a monotone finite-difference scheme with Howard's policy iteration. We then develop a reinforcement learning approach based on policy-gradient and actor-critic methods. The learned strategies closely match the PDE benchmark and remain stable across initial conditions. The results highlight the relevance of policy-gradient techniques for dividend optimisation under self-exciting claim dynamics and point toward scalable methods for higher-dimensional extensions.
This is a joint work with Etienne Chevalier and Vathana Ly Vath.

Osgood and Feller Criteria for Stochastic Volterra Equations

Alessandro Bondi

Luiss, Italy

In this presentation, I will discuss boundary attainment conditions for one-dimensional stochastic Volterra equations (SVEs) of convolution type. In the first part of the talk, I will present an Osgood-type test for explosion to infinity of SVEs driven by additive noise, featuring kernels from a family that includes the fractional kernel. I will also investigate stability results for explosion times with respect to the kernels. In the second part, I will present a Feller-type test that establishes, on a general open interval of the real line, necessary and sufficient conditions for boundary attainment of solutions to SVEs with possibly multiplicative noise. Here, I will consider dynamics governed by nonsingular kernels, which preserve the semi-martingale property of the processes while introducing memory effects through a path-dependent drift.
This is a joint work with Sergio Pulido

Dam Management in the Era of Climate Change

Cristina Di Girolami

Università di Bologna

Climate change has a dramatic impact, particularly by concentrating rainfall into a few short periods, interspersed with long dry spells. In this context, the role of dams is crucial. We consider the optimal control of a dam, where the water level must neither exceed a designated safety threshold nor fall below a minimum level to ensure functionality and sustainability for the downstream river. To model dry spells and intense rainfall events, commonly referred to as water bombs, we introduce a Hawkes process, a well-known example of a self-exciting process characterized by time-correlated intensity, which endogenously reproduces the clustering of events. The problem is formulated as an optimal switching problem with constraints. We establish existence results and propose numerical methods for approximating the solution. Finally, we illustrate the main achievements of this approach through numerical examples focusing in particular on the sensitivity of the self-exciting parameter describing the importance of both water bombs and dry-spells. The main result of our numerical analysis is that the optimal water level inside the dam decreases as the self-exciting parameter increases. This finding shows that, when facing the dilemma of managing the opposing risks of dam overtopping and dry spells, the former proves to be dominant over the latter. In conclusion, dams will increasingly lose their role as water reserves and take on a greater role in flood protection.
This is a joint work with Mhamed Gaigi, Vathana Ly Vath and Simone Scotti.

TBA

Etienne Chevalier

Université d'Evry Paris Saclay, France

TBA
This is a joint work with .....

A numerical scheme for optimal switching problems based on a randomization method

Céline Labart

LaMME Universitè Savoie Mont Blanc, France

In this talk we deal with a numerical method for solving optimal switching problems (OSP) in presence of switching costs. Using the randomization method, we link the value function of the OSP to a non standard BSDE. We propose a numerical method combining a backward discretization and a least square regression method to approximate the solution of the non standard BSDE. We provide a rate of convergence for our algorithm.
This is a joint work with Marie-Amelie Morlais.

How can the dual martingale help solving the primal optimal stopping problem?

Jérôme Lelong

ENSIMAG, France

Motivated by recent results on the dual formulation of optimal stopping problems, we investigate in this short paper how the knowledge of an approximating dual martingale can improve the efficiency of primal methods. In particular, we show on numerical examples that accurate approximations of a dual martingale efficiently reduce the variance for the primal optimal stopping problem.
This is a joint work with Aurélien Alfonsi and Ahmed Kebaier

TBA

Vathana Ly Vath

ENSIIE, France

TBA
This is a joint work with .....

Adaptive Learning via Off-Model Training and Importance Sampling for Fully Non-Markovian Optimal Stochastic Control

Alberto Ohashi

Universidade de Brasilia, Brazil

This paper studies continuous-time stochastic control problems whose controlled states are fully non-Markovian and depend on unknown model parameters. Such problems arise naturally in path-dependent stochastic differential equations, rough-volatility hedging, and systems driven by fractional Brownian motion. Building on the discrete skeleton approach developed in earlier work, we propose a Monte Carlo learning methodology for the associated embedded backward dynamic programming equation. Our main contribution is twofold. First, we construct explicit dominating training laws and Radon–Nikodym weights for several representative classes of non-Markovian controlled systems. This yields an off-model training architecture in which a fixed synthetic dataset is generated under a reference law, while the dynamic programming operators associated with a target model are recovered by importance sampling. Second, we use this structure to design an adaptive update mechanism under parametric model uncertainty, so that repeated recalibration can be performed by reweighting the same training sample rather than regenerating new trajectories. For fixed parameters, we establish non-asymptotic error bounds for the approximation of the embedded dynamic programming equation via deep neural networks. For adaptive learning, we derive quantitative estimates that separate Monte Carlo approximation error from model-risk error. Numerical experiments illustrate both the off-model training mechanism and the adaptive importance-sampling update in structured linear-quadratic examples.
This is a joint work with Dorival Leão, Simone Scotti and Adolfo M. Dias da Silva.

Model-free policy gradient for discrete-time mean-field control

Huyên Pham

Ecole Polytechnique, Institut Polytechnique de Paris, France

We study model-free policy learning for discrete-time mean-field control (MFC) problems with finite state space and compact action space.

In contrast to the extensive literature on value-based methods for MFC, policy-based approaches remain largely unexplored due to the intrinsic dependence of transition kernels and rewards on the evolving population state distribution, which prevents the direct use of likelihood-ratio estimators of policy gradients from classical single-agent reinforcement learning. We introduce a novel perturbation scheme on the state-distribution flow and prove that the gradient of the resulting perturbed value function converges to the true policy gradient as the perturbation magnitude vanishes. This construction yields a fully model-free estimator based solely on simulated trajectories and an auxiliary estimate of the sensitivity of the state distribution. Building on this framework, we develop MF-REINFORCE, a model-free policy gradient algorithm for MFC, and establish explicit quantitative bounds on its bias and mean-squared error. Numerical experiments on representative mean-field control tasks demonstrate the effectiveness of the proposed approach.
This is a joint work with Matthieu Meunier and Christoph Reisinger

---

When do Volterra processes have a finite-dimensional Markovian representation?

Sergio Pulido

ENSIIE, Université Paris Saclay, France

Finite-dimensional Markovian representations of Volterra processes are important for practical applications. In this talk, we provide a characterization of the kernels that enable such representations, based on geometric arguments in an infinite-dimensional lifting framework. We also highlight connections with finite-dimensional realizations in HJM fixed income models.
This is a joint work with Alexandre Pannier.

On self-exciting point processes

Anthony Réveillac

INSA Toulouse and Institut de Mathématiques de Toulouse, France

Based on a recent expansion of general point processes, we will discuss a new class of self-exciting point processes which provide alternatives to Hawkes processes in cyber-insurance risk modelling. We will present some properties of these processes and focus on the particular case of counting processes.
This is a joint work with Caroline Hillairet, Thomas Peyrat and Achille Pommier.

Optimization of capital injections and dividend payments in a general diffusion model

Alexandre Roch

Université du Quebec à Montreal

We study an optimal dividends control problem with capital injection for a firm's surplus modeled as a general diffusion. Dividend rates are constrained by a level-dependent upper bound given by a concave function of the current surplus. We first analyze the auxiliary problem in which capital injections are forced to keep the surplus nonnegative at all times. The crucial step is the concavity of the value function which we prove using concave envelopes and viscosity-type comparison principles. This allows us to characterize the optimal policy as a double-barrier strategy: the surplus is minimally reflected at zero by injections, and dividends are paid at the maximal admissible rate whenever the surplus exceeds the upper bound. Our main result establishes a dichotomy for the general optimization problem: the optimal value coincides with that of one of two simpler problems, namely the no-injection problem (dividends until ruin) or the forced-injection problem (dividends with reflection at zero). The optimal injection decision is obtained by a simple boundary test comparing the two value functions at the origin.
This is joint work with Hélène Guérin, Dante Mata Lopez and Jean-François Renaud

---

The numeraire e-variable and reverse information projection

Johannes Ruf
London School of Economics, United Kingdom

We consider testing a composite null hypothesis $\mathcal{P}$ against a point alternative Q using e-variables, which are nonnegative random variables X such that $\mathbb{E}_{\mathsf{P}}[X] \leq 1$ for every $\mathsf{P} \in \mathcal{P}$. This paper establishes a fundamental result: under no conditions whatsoever on $\mathcal{P}$ or Q, there exists a special e-variable X^* that we call the numeraire, which is strictly positive and satisfies $\mathbb{E}_{\mathsf{Q}}[X/X^*] \leq 1$ for every other e-variable X. In particular, X^* is log-optimal in the sense that $\mathbb{E}_{\mathsf{Q}}[\log(X/X^*)] \leq 0$. Moreover, X^* identifies a particular sub-probability measure P^* via the density $d\mathsf{P}^*/d\mathsf{Q} = 1/X^*$. As a result, X^* can be seen as a generalized likelihood ratio of Q against $\mathcal{P}$. We show that P^* coincides with the reverse information projection (RIPr) when additional assumptions are made that are required for the latter to exist. Thus P^* is a natural definition of the RIPr in the absence of any assumptions on $\mathcal{P}$ or Q. In addition to the abstract theory, we provide several tools for finding the numeraire and RIPr in concrete cases. We discuss several nonparametric examples where we can indeed identify the numeraire and RIPr, despite not having a reference measure. Our results have interpretations outside of testing in that they yield the optimal Kelly bet against $\mathcal{P}$ if we believe reality follows Q. We end with a more general optimality theory that goes beyond the ubiquitous logarithmic utility. We focus on certain power utilities, leading to reverse Rényi projections in place of the RIPr, which also always exist.
This is a joint work with Martin Larsson and Aaditya Ramdas.

Semi-static variance-optimal hedging with self-exciting jumps.

Carlo Sgarra

Universitá di Bari, Italia

The aim of this talk is to investigate a quadratic, i.e. variance-optimal, semi-static hedging problem in an incomplete market model where the underlying log-asset price is driven by a diffusion process with stochastic volatility and a self-exciting jump process of Hawkes type. More precisely, we aim at hedging a claim at time $T > 0$ by using a portfolio of available contingent claims, so to minimize the variance of the residual hedging error at time T. In order to improve the replication of the claim, we look for a hybrid hedging strategy of semi-static type, in which some assets are continuously rebalanced (the dynamic hedging component) and for some other assets a buy-and-hold strategy (the static component) is performed. We discuss in detail a specific example in which the approach proposed is applied, i.e. a variance swap hedged by means of European options, and we provide a numerical illustration of the results obtained.
This is a joint work with Giorgia Callegaro, Paolo Di Tella and Beatrice Ongarato.

Deep Learning for the Multiple Optimal Stopping Problem

Mehdi Talbi

Université Paris Cité, France

This talk presents a novel deep learning framework for solving multiple optimal stopping problems in high dimensions. While deep learning has recently shown promise for single stopping problems, the multiple exercise case involves complex recursive dependencies that remain challenging. We address this by combining the Dynamic Programming Principle with neural network approximation of the value function. Unlike policy-search methods, our algorithm explicitly learns the value surface. We first consider the discrete-time problem and analyze neural network training error. We then turn to continuous problems and analyze the additional error due to the discretization of the underlying stochastic processes. Numerical experiments on high-dimensional American basket options and nonlinear utility maximization demonstrate that our method provides an efficient and scalable method for the multiple optimal stopping problem.
This is a joint work with Mathieu Laurière.

Author Index

SPONSORSHIP AND PATRONAGE:

Università di Pisa – Dipartimento di Economia e Management

université evry val-d'essonne | université PARIS-SACLAY

ensiie PARIS-ÉVRY

LaMME – Laboratoire de Mathématiques et Modélisation d'Évry

Université Franco Italienne | Università Italo Francese

Grenoble INP UGA – Ensimag

A.M.A.S.E.S. – Association for Mathematic Applied to Social Economic Sciences

CO MAT – Comitato dei Matematici Applicati in Toscana

pisa-finance-week.ec.unipi.it