Applications of Stochastic Control to Finance and Economics

Jakša Cvitanić (California Institute of Technology),
George Georgiadis (Northwestern University),
Dylan Possamaï (ETH Zürich),
Nizar Touzi (École Polytechnique)

April 30–May 5, 2023

1 Overview of the Field

A distinguished feature of the last few decades is the increasing availability of technological innovations to an everyday broader spectrum of society. This phenomenon has rapidly paved the way for the development of new economies such as e-commerce, sharing economies and online advertising. These activities are characterised by, for instance, a higher number of agents from both the supply and demand side, and an unprecedentedly large menu of tailor-made and personalised services. Naturally, this has led to more intricate behaviours and interactions at all levels of the economy. Consequently, there is a pressing need for the creation of economic theories that are able to explain the rationale of the participants, the way their actions are aggregated in society, and which are compatible with documented evidence.

As such, stochastic optimal control theory applied to problems in economics and finance has recently evolved into one of the most active and exciting areas of contemporary applied mathematics. It provides the mathematical community with a rich source of challenges, most of which stem from the deep insights of economists, which, in turn, enrich society's understanding of the complex dynamics which underline its functioning, helping all involved actors to make better and more informed choices. The proximity of these theories to practice, in addition to their purely mathematical appeal, make this field especially attractive to young mathematicians. Moreover, it provides for wider employment opportunities both within the academic world and outside of it.

One of the most challenging aspects has nonetheless remained, in the sense that the communication of recent mathematical results to the economics community has been facing a certain number of hurdles, as the language used is often not the same, making the interactions unfortunately limited. This workshop was aimed at fostering the interactions between researchers from both the theoretical and practical side, and at being one of the limited spaces where a significant number of representatives of these two communities, which do not often meet, can dedicate themselves to the dissemination of new theoretical results and their implementation in the design of new economic models.

2 Presentation Highlights

Beatrice Acciaio

'Equilibria as solutions to Schrödinger problems'

ABSTRACT. I will present a connection between solutions to equilibrium problems in markets with asymmetric information, and solutions of Schrödinger problems. In particular, we will see how in some cases the former ones can be found as limit of the latter ones. Based on joint ongoing work with Umut Çetin.

René Aïd

'A stationary mean-field equilibrium model of irreversible investment'

ABSTRACT. We develop a stationary mean-field model of singular control irreversible investment where production capacity is subject to random gaussian fluctuations while prices are affected by macroeconomic shocks following a two-state Markov chain (boom/burst episodes). We prove existence and uniqueness of the mean-field stationary equilibrium and we characterise it through a system of nonlinear equations. We provide several insights on the effects of crisis on the equilibrium, in particular on the firms size distribution and on the firms value at equilibrium. Joint work with Giorgio Ferrari and Matteo Basei.

Alexander Bloedel

'Persistent private information revisited'

ABSTRACT. This paper revisits [9] (henceforth PPIs) continuous-time principal-agent model of optimal dynamic insurance with persistent private information. We identify three independent issues in PPI that implicate its characterisations of incentive compatible and optimal contracts: (i) the agent cannot over-reportincrements ofhis type, a constraint that does not follow from the common assumption that the agent cannot over-report his type; (ii) the agent's feasible set of reporting strategies does not include standard 'no Ponzi'constraints, without which PPI's main analysis of infinite-horizon incentive compatibility is incomplete; and (iii) most importantly, in PPI's main application, which concerns hidden endowments, the contract identified as optimal isgenerically strictly suboptimal. For this application, we address the first two issues and elucidate the third one by analysing a class of 'self-insurance contracts' that can be implemented as consumption–saving problems for the agent, and which includes the contract derived in PPI as a particular case. We characterise the optimal self-insurance contract and show that, generically, it strictly dominates PPI's. Our analysis does not support PPIs main economic finding that immiseration generally fails or its attribution of this failure to continuous time and persistence. (Joint work with R. Vijay Krishna and Bruno Strulovici).

Umut Çetin

'Order routing and market quality: Who benefits from internalisation?'

ABSTRACT. The practice of retail internalisation has been a controversial topic since the late 1990s. The crux of this debate is whether this practice benefits, via the price improvement relative to exchange, or disadvantages, via the reduced liquidity on exchange, retail traders. To answer this question we set two models of market design that differ in their mode of liquidity provision: in the model capturing retail order internalisation the liquidity is provided by market makers (representing wholesalers) competing for the retail order flow in a Bertrand fashion, whereas in the model characterising the open exchange the price-taking competitive agents act as liquidity providers. We discover that, when liquidity providers in both market designs are risk averse, routing of the marketable orders to the wholesalers is preferred by retail traders: informed,

uninformed and noisy. In addition to addressing optimal order routing problem, we identify a universal parameter that allows comparison of market liquidity, profit and value of information across different markets and demonstrate that the risk aversion of liquidity providers fundamentally changes market outcomes. In particular, we observe mean reverting inventories, price reversal, and lower market depth as the result of retail investor (informed or not) absorbing large shocks in their inventory to compensate for the unwillingness of liquidity providers to bear risk.

Tahir Choulli

'Optimal stopping and reflected BSDEs for models under arbitrary random horizon'

ABSTRACT. We consider the setting described by $(\Omega, \mathcal{G}, \mathbb{F}, \tau, \mathbb{P})$. Here $(\Omega, \mathcal{G}, \mathbb{F}, \mathbb{P})$ is a complete filtered probability space, called initial model hereafter, in which $\mathbb{F}$ represents the flow of information available to all agents throughout time. τ is an arbitrary random time, which might not be observable via $\mathbb{F}$, and represents default time of a firm in credit risk theory, or death time of an insured in life insurance, and/or the occurrence time of any event that might impact the market somehow. In this framework of informational model, various valuation approaches and various hedging methods which we are developing point to the following family of reflected backward stochastic differential equations (RBSDE hereafter)

$$\begin{cases} \mathrm{d}Y_t = f(t, Y_t, Z_t)\mathrm{d}(t \wedge \tau) + g(t, Y_t, Z_t)\mathrm{d}U_t^{\tau} + Z_t \mathrm{d}W_t^{\tau} + \mathrm{d}M_t - \mathrm{d}K_t,\ Y_{\tau} = \xi, \\ Y \geq S \text{ on } [0, \tau],\ \displaystyle\int_0^{\tau} \big(Y_{s-} - S_{s-}\big)\mathrm{d}K_s = 0, \mathbb{P}\text{–a.s.} \end{cases}$$

Here W is an $\mathbb{F}$–Brownian motion, and (f, g, ξ, S, U) is the data-quintuplet in which U is an RCLL nondecreasing and $\mathbb{F}$-adapted process, f and g are the drivers functionals, ξ is the terminal condition and the barrier process S is RCLL $\mathbb{F}$-adapted. We allow the initial model $(\Omega, \mathcal{G}, \mathbb{F}, \mathbb{P})$ to be chosen as ""nice enough"" as possible if needed, while τ is let to be as arbitrary general as possible. This will allow to measure adequately the impact of τ on the initial model in various aspects without interferences. In my talk, I will single out the challenges arising from allowing τ to be arbitrary general, and I will detail our approaches for these challenges and the intermediate problems arising from them as well. Among these, we address the impact of τ on the mathematical structures of these RBSDEs, which yields to the mathematical structures in the optimal stopping problem under random horizon τ. Besides being intimately related to RBSDEs and vital in addressing the mathematical structures, the optimal stopping problem is interesting on its own and is useful in solving other financial problems. This talk is based on joint works with Safa Alsheyab.

Théo Durandard

'Under pressure: comparative statics for optimal stopping problems in non-stationary environments'

ABSTRACT. We formulate a general optimal stopping problem that can represent a wide variety of non-stationary environments, *e.g.*, where the decision makers patience, time pressure, and learning speed can change gradually and abruptly over time. We show that, under some mild regularity conditions, this problem has a well-defined solution. Furthermore, we characterise the shape of the stopping region in a large class of monotone environments. As a result, we obtain comparative statics on the timing and quality of decisions for many sequential sampling problems *à la* Wald. For example, we show that accuracy is increasing (decreasing) over time when (i) learning speed increases (decreases) in time, or (ii) the discount rate decreases (increases) over time (i.e., the decision maker values the future more over time), or (iii) time pressure decreases (increases) over time. Since our main comparative static results hold locally, we can also capture non-monotone relations between time and accuracy that consistently arise in both perceptual and cognitive testing.

Christoph Frei

'Principal trading arrangements: optimality under temporary and permanent price impact'

ABSTRACT. We study the optimal execution problem in a principal-agent setting. A client (for example, a pension fund, endowment, or other institution) contracts to purchase a large position from a dealer at a future point in time. In the interim, the dealer acquires the position from the market, choosing how to divide his trading across time. Price impact may have temporary and permanent components. There is hidden action in that the client cannot directly dictate the dealer's trades. Rather, she chooses a contract with the goal of minimizing her expected payment, given the price process and an understanding of the dealer's incentives. Many contracts used in practice prescribe a payment equal to some weighted average of the market prices within the execution window. We explicitly characterise the optimal such weights: they are symmetric and generally U-shaped over time. The talk is based on joint work with Markus Baldauf (University of British Columbia) and Joshua Mollner (Northwestern University).

Paolo Guasoni

'Holding stocks, trading bonds'

ABSTRACT. In an economy with random growth, several long-lived agents with heterogeneous risk aversions, time-preferences, and personal income streams make consumption and investment decisions, trading stocks and a consol bond, and borrowing from and lending to each other. We find in closed form equilibrium stock prices, interest rates, consumption, and trading policies. Agents do not trade stocks, although their returns are time-varying and predictable. Agents dynamically trade the consol bond in response to growth shocks, as to hedge their effect on interest rates, dividends, and personal incomes. Static fund separation holds if agents have also access to a linear bond and two additional hedges for dividend and growth shocks. Such additional assets can be dynamically replicated with the stock and the consol bond. No representative agent exists.

Nicolás Hernández Santibáñez

'A continuous-time model of self-protection'

ABSTRACT. We present an optimal linear insurance demand model, where the protection buyer can also exert a time-dynamic costly prevention effort to reduce her risk exposure. This is expressed as a stochastic control problem that consists in maximising an exponential utility of a terminal wealth. We assume that the effort reduces the intensity of the jump arrival process and interpret this as dynamic self-protection. We solve the problem by using a dynamic programming approach, and we provide a representation of the certainty equivalent of the buyer as the solution to a backward stochastic differential equation (BSDE). Using this representation, we prove that an exponential utility maximiser has an incentive to modify her effort dynamically only in the presence of a terminal reimbursement in the contract. Otherwise, the dynamic effort is actually constant, for a class of compound Poisson loss processes. If there is no terminal reimbursement, we solve the problem explicitly and identify the dynamic certainty equivalent of the protection buyer. This shows in particular that the Lvy property of the loss process is preserved under exponential utility maximisation. We also characterise the constant effort as the unique minimiser of an explicit Hamiltonian, from which we can determine the optimal effort in particular cases. Finally, after studying the dependence of the BSDE associated to the insurance buyer on the linear insurance contract parameter, we prove the existence of an optimal linear cover that is not necessarily zero or full insurance.

Camilo Hernández

'Propagation of chaos for Schrödinger problems with interacting particles'

ABSTRACT. In this work, we study the mean field Schrödinger problem from a purely probabilistic point of view by exploiting its connection to stochastic control theory for McKean–Vlasov diffusions. Our main result shows that the mean field Schrödinger problem arises as the limit of 'standard' Schrödinger problems over interacting particles. Due to the stochastic maximum principle and a suitable penalisation procedure, the result follows as a consequence of a novel quantitative propagation of chaos result for forward–backward particle systems. Our stochastic control technique further allows us to solve the mean field Schrödinger problem and characterise its solution, the mean field Schrödinger bridge, by a forward–backward planning equation. The approach described in the paper seems flexible enough to address other questions in the theory. For instance, it allows us to study zero noise limits for the Schrödinger problem.

Julien Hugonnier

'Asset pricing with costly short sales'

ABSTRACT. We study a dynamic general equilibrium model with costly-to-short stocks and heterogeneous beliefs. The closed-form solution to the model shows that costly short sales drive a wedge between the valuation of assets that promise identical cash flows but are subject to different trading arrangements. Specifically, we show that the price of an asset is given by the risk-adjusted present value of future cash flows which include both dividends and an endogenous lending yield. This formula implies that returns satisfy a modified CAPM and sheds light on recent findings about the explanatory power of lending fees in the cross-section of returns. In particular, we show that once returns are appropriately adjusted for lending fees, stocks with low and high shorting costs offer similar risk-return tradeoffs.

Kostas Kardaras

'Portfolio choice under taxation and expected market time constraint'

ABSTRACT. We consider the problem of choosing an investment strategy that will maximise utility over distributions, under capital gains tax and constraints on the expected liquidation date. We show that the problem can be decomposed in two separate ones. The first involves choosing an optimal target distribution, while the second involves optimally realising this distribution via an investment strategy and stopping time. The latter step may be regarded as a variant of the Skorokhod embedding problem. A solution is given very precisely in terms of the first time that the wealth of the growth optimal portfolio, properly taxed, crosses a moving stochastic (depending on its minimum-to-date) level. The suggested solution has the additional optimality property of stochastically minimising maximal losses over the investment period.

Nabil Kazi-Tani

'The role of correlation in diffusion control ranking games.'

ABSTRACT. In this talk, we will study Nash equilibria in two player continuous time stochastic differential games with diffusion control, and where the Brownian motions driving the state processes are correlated. We consider zero-sum ranking games, in the sense that the criteria to optimise only depends on the difference of the two players' state processes. We explicitly compute the players' equilibrium strategies, depending on the correlation of the Brownian motions driving the two state equations: in particular, if the correlation coefficient is smaller than some explicit threshold, then the equilibrium strategies consist of strong controls, whereas if the correlation exceeds the threshold, then the equilibrium controls are mixed strategies. To characterise these equilibria, we rely on a relaxed formulation of the game based on solutions to martingale problems, allowing the players to randomise their actions. The talk is based on a joint work with Stefan Ankirchner (University of Jena) and Julian Wendt (University of Jena).

Daniel Kršek

'Relaxed principal–agent problem'

ABSTRACT. We study a principal–agent problem with a lump-sum payment on a finite-time horizon. Extending the dynamic programming approach in [3] we consider problems involving constraints on the optimal contract. We introduce a framework, in which the agent is allowed to randomise his actions. This in turn gives more freedom to the principal, when she chooses the contract, and allows us to show existence of an optimal contract in problems with fairly general constraints, for which the standard PDE approach ceased to be tractable.

Ali Lazrak

'Democratic policy decisions with decentralised promises contingent on vote outcome'

ABSTRACT. We study how decentralised utility transfer promises affect collective decision-making by voting. Committee members with varying levels of support and opposition for an efficient reform can make enforceable promises before voting. An equilibrium requires stability and minimal promises. Equilibrium promises exist and are indeterminate, but do share several key characteristics. Equilibria require transfer promises from high to low intensity members and result in enacting the reform. When reform supporters lack sufficient voting power, promises must reach across the aisle. Even if the coalition of reform supporters is decisive, promises must preclude the least enthusiastic supporters of the reform from being enticed to overturn the decision. In that case, equilibrium promises do not need to reach across the aisle. We also discuss a finite sequence of promises that achieve an equilibrium.

Marcel Nutz

'Unwinding stochastic order flow in a central risk book'

ABSTRACT. We study the optimal execution problem for the Central Risk Book (CRB), a centralised trading unit recently established in many large banks and trading companies. The CRB aggregates orders from the other business units within the organisation in real time, netting opposite orders and executing outstanding orders such as to minimise transaction costs. Thus, the in-flow orders of the CRB are a stochastic process. We introduce a tractable model for the price impact and spread cost paid by the out-flow orders and find the optimal execution strategy for a general class of in-flow processes. The strategy highlights how future in-flows are taken into account to determine the optimal trade-off between trading speed and transaction costs. (Joint work with Kevin Webster and Long Zhao.)

Huyên Pham

'Generative modelling for time series via Schrödinger bridge '

ABSTRACT. We propose a novel generative model for time series based on Schrödinger bridge (SB) approach. This consists in the entropic interpolation via optimal transport between a reference probability measure on path space and a target measure consistent with the joint data distribution of the time series. The solution is characterised by a stochastic differential equation on finite horizon with a path-dependent drift function, hence capturing the temporal dynamics of the time series distribution. We can estimate the drift function from data samples either by kernel regression methods or with LSTM neural networks, and the simulation of the SB diffusion yields new synthetic data samples of the time series. The performance of our generative model is evaluated through a series of numerical experiments. First, we test with a toy autoregressive model, a GARCH Model, and the example of fractional Brownian motion, and measure the accuracy of our algorithm with marginal and temporal dependencies metrics. Next, we use our SB generated synthetic samples for the application to deep hedging on real-data sets. Finally, we illustrate the SB approach for generating sequence of images. Based on joint work with M. Hamdouche and P. Henry-Labordère

Zhenjie Ren

'Uniform-in-time propagation of chaos for mean-field Langevin dynamics'

ABSTRACT. In recent times, there has been a growing interest in the study of mean-field Langevin (MFL) dynamics, primarily due to its natural application in training two-layer neural networks. To simulate the MFL dynamics' invariant distribution, one relies on the corresponding N-particle system, hoping that the error between the particle system and the mean-field dynamics remains small over an extended period. Our recent research focuses on the uniform-in-time propagation of chaos for the MFL dynamics with convex mean-field potential, motivated by this observation. We establish that this holds true for the $\mathbb{L}_2$-Wasserstein distance and relative entropy under mild conditions.

Alejandro Rivera

'Contracting with a present-biased agent: Sannikov meets Laibson'

ABSTRACT. This paper develops a methodology to solve dynamic principal–agent problems in which the agent features present-biased time preferences and naive beliefs. There are three insights. First, the problem has a recursive representation using the agent's perceived continuation value as a state variable (i.e., the remaining value the agent (wrongly) anticipates getting from the contract). Second, incentive compatibility corresponds to a volatility constraint on the agent's perceived continuation value. Finally, due to the agent's naiveté, a perceived action constraint needs to be satisfied. This constraint is accommodated by linking the agent's perceived effort policy and the volatility of his perceived continuation value. Novel economic insights regarding optimal time-varying incentives and the term-structure of compensation are also explored.

Marco Rodrigues

'BSDEs and reflected BSDEs driven by general martingales'

ABSTRACT. We study the well-posedness of BSDEs and reflected BSDEs on random time horizons with stochastic Lipschitz constants and driven by general (possibly pure-jump) martingales. We employ a fixed-point theorem and provide various conditions, depending on the complexity of the generator, that ensure the existence and uniqueness within a class of processes.

Chiara Rossato

'Sannikov's principal–agent problem with jumps'

ABSTRACT. Based on recent work by [7] we consider an extension of Sannikov's principal–agent problem by letting the agent control the drift and jump intensity of the output process. We investigate whether the problem exhibits a golden parachute, that is, whether there is a scenario in which the agent retires and receives a continuous stream of payments or the agent receives a lump-sum payment as compensation for the termination of the contract by the principal. With the introduction of the face-lifted utility, we can study the two cases simultaneously in a different principal–agent problem that we reduce to a standard mixed control–stopping problem.

Halil Mete Soner

'Viscosity solutions for the mean-field control'

ABSTRACT. McKean–Vlasov or mean-field control problems are very closely related to the mean-field games. Dynamic programming for these problems results in a nonlinear partial differential equation on the space of probability measures. These equations require the value function to be not only differentiable but also its derivatives (which are in the dual of the set of measures, hence continuous functions) to be twice differentiable. Despite these difficulties, several approaches to characterise the value function as the unique appropriate weak solutions have been developed. We employ the classical viscosity solutions using the intrinsic linear derivative. We obtain a comparison result between the Lipschitz viscosity sub- and super-solutions under a structural assumption on the control set together with the given functions. roof follows the classical variable doubling argument of Crandall & Lions by using the Fourier characterisation of the dual Sobolev norm. The value function is also shown to be Lipschitz continuous with respect to this metric. This is joint work with Qinxin Yan of Princeton University.

Bruno Strulovici

'Smoothness of value functions in general control–stopping diffusion problems'

ABSTRACT. We study the properties of value functions in joint optimal control and stopping problems where (i) the state variable may be multi-dimensional, (ii) the domain may be unbounded, and (iii) the primitives may be time-inhomogeneous. We show that the value function is (i) the unique $\mathbb{L}^p$-solution of the Hamilton–Jacobi–Bellman equation, (ii) twice parabolically differentiable a.e., and (iii) continuously differentiable in the non-time variables, under general conditions relevant for most economic applications. In particular, we show that the smooth-pasting property holds everywhere with respect to all space variables. We also derive sufficient conditions under which smooth pasting also holds with respect to time. Our results imply that numerical solutions obtained by standard methods converge uniformly to the value function.

Mehdi Talbi

'Mean field games of optimal stopping'

ABSTRACT. We are interested in the study of stochastic games for which each player faces an optimal stopping problem. In our setting, the players may interact through the criterion to optimise as well as through their dynamics. After briefly discussing the N-players game, we formulate the corresponding mean field problem. In particular, we introduce a weak formulation of the game for which we are able to prove existence of Nash equilibria for a large class of criteria. We also prove that equilibria for the mean field problem provide approximated Nash equilibria for the N-players game, and we formally derive the master equation associated with our mean field game. Joint work with Dylan Possamaï.

Ludovic Tangpi

'Forward–backward propagation of chaos via displacement monotonicity'

ABSTRACT. In this talk I will present quantitative convergence results for a class of mean field games with common noise and controlled volatility. The basic strategy we employ is the one introduced recently by Laurière and myself—roughly speaking, we use a synchronous coupling argument to prove a 'forward–backward propagation of chaos' result for the FBSDEs which characterise the (open-loop) equilibria of the N-player and mean field games. Unlike in earlier works which have adopted this strategy, we do not require smallness conditions, and instead rely on monotonicity. In particular, (displacement) monotonicity of the Hamiltonian and the terminal cost allow us to establish a (uniform in N) stability estimate for the N-player FBSDEs, which implies the convergence result. The arguments are relatively simple, and flexible enough to yield similar results in the setting of mean field control and infinite horizon (discounted) mean field games.

Stéphane Villeneuve

'Money implements optimal contract'

ABSTRACT. We study a mean-field principal–agent model where interactions take place through the allocation of capital to agents. We analyse the impact of information asymmetry on risk-sharing and investment. The optimal contract is explicitly characterised and implemented with money and taxes.

Hao Xing

'Reward and monitoring in dynamic contracts'

ABSTRACT. We develop a dynamic principal–agent model that examines the intricate balance between monitoring and rewards in executive compensation. In the optimal contract, monitoring and success reward dynamically substitute each other, exhibiting two main characteristics: (i) a reduction in contract sensitivity to monitoring and a rise in the success reward when positive evidence of the agent's effort accumulates, and (ii) a larger success reward when the signal about the agent's effort is more noisy, particularly when negative evidence accumulates. Using the changes in the availability of direct flights for board directors to the firm's headquarters as an exogenous proxy of signal quality, we present empirical evidence that supports our model's predictions. This is a joint work with Kerry Back, Clint Hamilton, and Ali Kakhbod.

Mihail Zervos

'Market equilibrium under proportional transaction costs in a stochastic factor model'

ABSTRACT. We consider an economy with two agents. Each of the two agents receives a random endowment flow. We model this cumulative flow as the stochastic integral of a deterministic function of the economy's state, which we model by means of a general Ito diffusion. Each of the two agents has mean-variance preferences with different risk-aversion coefficients. The two agents can also trade a risky asset. We determine the agents' optimal equilibrium trading strategies in the presence of proportional transaction costs. In particular, we derive a new free-boundary problem that provides the solution to the agents' optimal equilibrium problem. Furthermore, we derive the explicit solution to this free-boundary problem when the problem data is such that the frictionless optimiser is a strictly increasing or a strictly increasing and then strictly decreasing function of the economy's state.

Jianfeng Zhang

'Viscosity solutions for fully nonlinear path dependent HJB equations on the Wasserstein space'

ABSTRACT. In this talk we investigate path dependent mean field optimal control problems with both drift and volatility controls. The value function is characterised by a fully nonlinear path dependent HJB equation on the Wasserstein space of probability measures. By lifting to the space of processes, we introduce a new notion of viscosity solutions and establish both existence and comparison principle. The main feature of our notion is that the test function consists of an extra component, besides the standard smooth part, which helps to get around of some major difficulty arising from the volatility control. We shall use the doubling variable arguments, combined with the Borwein–Preiss variational principle in order to overcome the non-compactness of the state space. A smooth gauge-type function on the path space is crucial for our estimates. The talk is based on an ongoing joint work with Nizar Touzi and Jianjun Zhou.

Participants

Acciaio Beatrice, ETH Zürich
Aïd René, Université Paris Dauphine–PSL
Bernard Carole, Grenoble École de Management
Bloedel Alexander, University of California Los Angeles
Bouchard Bruno, Université Paris Dauphine–PSL
Cadenillas Abel, University of Alberta
Çetin Umut, London School of Economics
Choulli Tahir, University of Alberta
Cvitanić, Jakša, California Institute of Technology
Durandard Théo, Northwestern University
Frei Christoph, University of Alberta
Georgiadis George, Northwestern University
Guasoni Paolo, Dublin City University
Hernández Santibáñez Nicolás, Universidad de Chile
Hernández Camilo, Princeton University
Hou Songyan, ETH Zürich
Hubert Emma, Princeton University
Hugonnier Julien, École Polytechnique Fédérale de Lausanne
Jofré Alejandro, Universidad de Chile
Kardaras Constantinos, London School of Economics

Kazi-Tani Nabil, Université de Lorraine
Ke Rongzhu, Zhejiang University
Kršek Daniel, ETH Zürich
Lazrak Ali, University of British Columbia
Mastrolia Thibaut, University of California, Berkeley
Nutz Marcel, Columbia University
Pham Huyên, Université Paris Cité
Possamaï Dylan, ETH Zürich
Ren Zhenjie, Université Paris Dauphine–PSL
Rivera Alejandro, University of Texas at Dallas
Rodrigues Marco, ETH Zürich
Rossato Chiara, ETH Zürich
Soner Halil Mete, Princeton University
Strulovici Bruno, Northwestern University
Talbi Mehdi, ETH Zürich
Tan Xiaolu, The Chinese University of Hong Kong
Tangpi Ludovic, Princeton University
Touzi Nizar, École Polytechnique
Villeneuve Stéphane, Toulouse School of Economics
Xing Hao, Boston University
Zariphopoulou Thaleia, The University of Texas at Austin
Zervos Mihail, London School of Economics
Zhang Jianfeng, University of Southern California
Zhou Chao, National University of Singapore

3 Scientific Progress Made

Contract theory. A large number of researchers shared their recent contributions on moral hazard problems in frameworks unavailable until very recent developments, which started from the seminal contribution of Sannikov [8], and culminated more recently in a series of papers by Possamaï, Cvitanić and Touzi [2, 3]. Hence Alexander Bloedel discussed deep new insights on a continuous-time principal–agent model of optimal dynamic insurance with persistent private information, while Christoph Frei and Nicolás Hernández Santibáñez respectively addressed an optimal execution problem, and a model of self-protection in an insurance context. Daniel Kršek presented completely new results on the general problem of existence of randomised optimal contracts, and Alejandro Rivera tackeled a new class of moral hazard problems with present-bias. In the same vein, Chiara Rossato offered an analysis of extension of the model of [8] to a setting incorporating potential sudden accidents, Stéphane Villeneuve considered a mean-field principal–agent model where interactions take place through the allocation of capital to agents, and Hao Xing examined the intricate balance between monitoring and rewards in executive compensation.

Mean-field games and mean-field Schrödinger problems. A strong emphasis was also put in studying state-of-the-art problems in mean-field game theory and their applications, which evolved from the seminal contributions by Lasry and Lions [6] on the one hand, and Huang, Caines and Malhamé [5] on the other hand. Hence, René Aïd talked about a stationary mean-field equilibrium model of irreversible investment, Zhenjie Ren presented a study of uniform-in-time propagation of chaos for mean-field Langevin dynamics, Mehdi Talbi concentrated on mean-field games of optimal stopping, while Ludovic Tangpi commented on very recent results for forward–backward propagation of chaos via displacement monotonicity. In a related fashion, several researchers presented progress on both the Schrödinger problem in the framework of [4], as well as the very recent mean-field Shrödinger problem, coined by Backhoff-Veraguas, Conforti, Gentil and Léonard [1]. In particular, Beatrice Acciaio explored how solutions to equilibrium problems in markets with asymmetric information could be connected to solutions of Schrödinger problems, while Huyên Pham showed us how one could design generative modelling for time series via Schrödinger bridges, and Camilo Hernández explored the limit theory for for Schrödinger problems with interacting particles.

Optimal control theory. Several very interesting contributions related to theoretical issues stemming from control problems, and how these could find applications for a deeper understanding of the solutions to certain economics problems were also presented. As such, both Tahir Choulli and Marco Rodrigues were concerned with different classes of backward stochastic differential equations and their applications for optimal control problems and optimal stopping problems. We also had the pleasure to listen to Théo Durandard and Bruno Strulovici who presented respectively deep results in optimal control problem allowing to obtain comparative statics on the timing and quality of decisions for many sequential sampling problems *à la* Wald, and new regularity results for Hamilton–Jacobi–Bellman equations. There were as well two talks, by Halil Mete Soner and Jianfeng Zhang, presenting two approaches to notions of viscosity solutions on the Wasserstein space of probability measures, which is an extremely recent and active research area, finding applications both for mean-field games and contract theory as described above. Let us also mention the talk by Nabil Kazi-Tani exploring the role of correlation in diffusion control ranking games.

Equilibrium theory and mathematical finance. The topic of competitive equilibria—traditionally housed in economics—has found its way to the mainstream of mathematical research in quantitative finance thanks to the versatility and mathematical richness of the stochastic models it supports. We thus had the pleasure to listen on this topic to Paolo Guasoni, Julien Hugonnier and Mihail Zervos. Other talks on recent takes on timely problems in mathematical finance were presented by Umut Çetin, on order routing and market quality, Kostas Kardaras on portfolio choice under taxation and expected market time constraint, as well as Marcel Nutz on the optimal execution problem for the central risk book.

Finally, let us mention the talk by Ali Lazrak, which outlined a whole new domain of research at the interface of financial mathematics and political economics, studying how decentralised utility transfer promises affect collective decision-making by voting.

References

[1] J. Backhoff-Veraguas, G. Conforti, I. Gentil, and C. Léonard. The mean field schrödinger problem: ergodic behavior, entropy estimates and functional inequalities. *Probability Theory and Related Fields*, 178:475–530, 2020.

[2] J. Cvitanić, D. Possamaï, and N. Touzi. Moral hazard in dynamic risk management. *Management Science*, 63(10):3328–3346, 2017.

[3] J. Cvitanić, D. Possamaï, and N. Touzi. Dynamic programming approach to principal–agent problems. *Finance and Stochastics*, 22(1):1–37, 2018.

[4] H. Föllmer. Random fields and diffusion processes. In P.L. Hennequin, editor, *École d'été de Probabilités de Saint-Flour XV–XVII, 1985–87*, volume 1362 of *Lecture notes in mathematics*, pages 101–203, 1988.

[5] M. Huang, P.E. Caines, and R.P. Malhamé. Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions. In C. Abdallah and F. Lewis, editors, *Proceedings of the 42nd IEEE conference on decision and control, 2003*, pages 98–103. IEEE, 2003.

[6] J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. I–Le cas stationnaire. *Comptes Rendus Mathématique*, 343(9):619–625, 2006.

[7] D. Possamaï and N. Touzi. Is there a golden parachute in Sannikov's principal–agent problem? *ArXiv preprint arXiv:2007.05529*, 2020.

[8] Y. Sannikov. A continuous-time version of the principal–agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.

[9] N. Williams. Persistent private information. *Econometrica*, 79(4):1233–1275, 2011.