Anna Aksamit
The University of Sydney

Robust pricing–hedging duality for American options in discrete time financial markets

We investigate pricing–hedging duality for American options in discrete time financial models where some assets are traded dynamically and others, e.g. a family of European options, only statically. We consider an abstract setting, which includes the classical case with a fixed reference probability measure as well as the robust framework with a non-dominated family of probability measures. Our first insight is that by considering an enlargement of the space, we can see American options as European options and recover the pricing–hedging duality, which may fail in the original formulation. This may be seen as a weak formulation of the original problem. Our second insight is that a duality gap arises from the lack of dynamic consistency and hence a different enlargement which reintroduces dynamic consistency is sufficient to recover the pricing–hedging duality: it is enough to consider fictitious extensions of the market in which all the assets are traded dynamically. This is joint work with Shuoqing Deng, Jan Obłój and Xiaolu Tan.

David Aldous
University of California, Berkeley

Markov chains, martingales and weak concentration

My only paper with Larry ("The least variable phase type distribution is Erlang", 1987) relies on a simple martingale-derived identity for variance of a Markov chain hitting time. I will describe surprising recent applications of this identity to bond percolation and first passage percolation on general networks. Maybe further applied probability applications await discovery.

Pieter Allaart
University of North Texas

Some difficult to solve optimal stopping problems

In this talk I will discuss some optimal stopping problems that my current and past Ph.D. students have worked on. These problems are quite natural and easy to state, but appear to be quite difficult to solve.

In the first problem, the objective is to stop at one of the two highest levels reached by a simple random walk with parameter $p < 1/2$ over a finite time interval. The optimal rule, which can be computed by backward induction, involves a sequence (p_n) of critical values for p which seem to follow an oscillating pattern. The problem is to prove that pattern.

The second problem is a random walk version of Robbins' problem of minimizing the expected rank. We consider both discrete- and continuous-time versions of this problem, and give upper and lower estimates for its value. Unfortunately, there remains a large gap between these bounds.

The third problem is to find prophet inequalities for Ferguson and Bruss' problem of selling multiple assets with vector offers and cost for observation. Even the case of two assets is challenging, and we know the exact bound only for time horizon $n = 2$.

This talk is based on joint work with J. Islas, A. Allen and E. Brophy.

Drew Allen
University of North Texas

A random walk version of Robbins' problem

Much progress has been made towards Robbins' problem of finding a stopping rule to minimize the expected rank of i.i.d. random variables. In this talk, we will examine a random walk version of this problem, in which we attempt to instead minimize the expected rank of the partial sums of symmetric, i.i.d. random variables. We will consider both discrete-time and continuous-time versions of this problem, and give upper and lower bounds for the expected rank of an optimal stopping rule. However, these bounds are presently still quite far apart. This talk is based on joint work with P. Allaart.

Luis Alvarez
University of Turku

Ergodic singular control problems of linear diffusions

Singular stochastic control plays an important role in many economic and financial applications ranging from cash flow management problems to the optimal management of renewable resources. Typically the objective is to determine the optimal policy either maximizing the expected cumulative present value of future cash flows or to minimize the expected cumulative present value of future costs. However, especially in harvesting applications an important question associated with the determination of an optimal policy is its long run sustainability. A natural way to tackle this question is to investigate the optimal ergodic singular policy either maximizing the expected long run average returns or minimizing the expected long run average costs. Motivated by these arguments, we will consider a class of stationary singular stochastic control problems of a linear and present a set of conditions under which the optimal policy is of the standard local time reflection type. We also investigate the comparative static properties of the optimal policy and delineate circumstances under which higher volatility expands the continuation region where utilizing the control is suboptimal.

Søren Asmussen
Aarhus University

Longest gap and run problems, with computer science applications

A classical topic in probability is longest runs of ones in i.i.d. Bernoulli sequences. We offer a number of extensions to more general models, including renewal sequences, Markov renewal processes and inhomogeneous Poisson processes. The emphasis is on asymptotic results, both on the first occurrence where the results are of large deviations type and on the typical behaviour in an interval where the study relates to extreme value theory. The problem was motivated from computer science problems such as transmissions of files on a link that may fail, or restart of program execution after a failure.

David Azriel
Technion-Israel Institute of Technology

Optimal sequential designs in multi-arm clinical trials

Consider a multi-arm clinical trial of size n, with N future patients who potentially can benefit from the results of the experiment. Typically, N is much larger than n. The goal of the experimenter is to maximize the number of total successes in the population. In this talk I will present algorithms to compute optimal sequential designs as well as some numerical and theoretical results.

Kerry Back
Rice University

Activism, strategic trading, and liquidity

Activist investors accumulate blocks of shares and then expend costly effort to promote changes in corporate policies that may increase share prices. We analyze equilibrium trading by an activist investor within the framework of Kyle (1985). The innovation in the paper is that the trader's private information is solely about the size of his blockholding, which is important to the market because the activist's willingness to exert effort depends on the blockholding. For example, if the market price is sufficiently high, perhaps because investors expect the trader to act, then it may be optimal for the trader to sell his shares and ultimately take no action. We solve the model by studying a generalization of the Brownian bridge that characterizes equilibrium in the standard Kyle model.

Jay Bartroff
University of Southern California

The bomber problem – Easy as ABC?

A bomber aircraft carrying x units of ammunition is time t from home. Suddenly an enemy appears and the Bomber must decide how much $K(x,t)$ of its ammunition to use to defend itself. Spending more increases the probability of surviving this enemy but leaves less for possible future enemy encounters. The so-called Bomber Problem concerns certain monotonicity conjectures about the optimal K, originally formulated almost 50 years ago, and some of which remain open today. Larry Shepp, with Simons and Yao (Adv. Applied Prob., 1991), studied a version of the Bomber Problem. I will describe some progress on the conjectures and related problems in the last decade, in part inspired by Shepp et al.'s approach, including asymptotic estimates of K and some of the techniques used. Some of this work is joint with Larry Goldstein, Yosi Rinott, and Ester Samuel-Cahn.

Matteo Basei
University of California, Berkeley

Nonzero-sum stochastic differential games with impulse controls

We consider a general class of nonzero-sum impulsive games with N players. By means of a suitable system of quasi-variational inequalities, we provide a verification theorem for the equilibrium strategies and the value functions of the game. In particular, we focus on the regularity conditions required by the theorem. We then present some practical applications. Finally, we consider an extension to mean-field games.

Erik Baurdoux
London School of Economics

Predicting the last zero for spectrally negative Lévy processes

Given a spectrally negative Lévy process drifting to infinity, we consider the last time g the process is below zero. We are interested in finding a stopping time which is as close as possible to g. In the L^1 setting, we show that an optimal stopping time is given by a first passage time above a level based on the convolution with itself of the distribution function of minus the overall supremum of the process. The proof is based on a direct approach without the need to make use of stochastic calculus.

For some more general metrics the problem is more challenging and can be transformed into an optimal stopping problem for a three-dimensional Markov process involving the last passage time. We show that the solution of this optimal stopping problem is given by the first time that the Lévy process crosses a non-increasing, non-negative curve which depends on the time spent above zero.

This talk is based on joint work with José Manuel Pedraza.

Erhan Bayraktar
University of Michigan

Large tournament games

We consider a tournament game in which each player is rewarded based on her rank in terms of the time of reaching a goal. We prove existence, uniqueness and stability of the game with infinitely many players, existence of an approximate equilibrium with finitely many players, and find an explicit characterization when players are homogeneous. In our setup we find that: (i) the welfare may be increasing in cost of effort; (ii) when the total pie is small, the aggregate effort may be increasing in prize inequality, unlike in Fang, Noe and Strack (2018); (iii) the welfare may go up with a higher percentage of unskilled workers, as do the completion rates of the skilled and unskilled sub-populations.

Our results lend support to government subsidies for R&D, assuming the profits to be made are substantial. Joint work with Jaksa Cvitanic and Yuchong Zhang.

Charles Beer
University of Wisconsin-Milwaukee

A stochastic control model for electricity markets

Beginning with some northern European markets in the early 1990's, nations around the world have liberalized the markets for electricity and later other related products. These open markets, however, show many characteristics which do not fit with standard classical commodity pricing models. More complex models including jump processes to account for these unusual behaviors have thus been developed in recent by Fred Espen Benth and others. Electricity producers also face additional challenges acting in these markets due to uncertainties surrounding their fuel supplies. We present a new control model for an electricity producer in a US energy market combining a modern pricing model including compound Poisson processes with a Markov chain model of fuel arrivals at the plant. We then present a first-exit control model for maximizing expected revenue of the plant until the first time fuel runs out. Finally, we briefly discuss the current results on the existence and uniqueness of a solution to this control problem.

Christoph Belak
University of Trier

Utility maximization with constant costs

We study the problem of maximizing expected utility of terminal wealth for an investor facing a mix of constant and proportional transaction costs. While the case of purely proportional transaction costs is by now well understood and existence of optimal strategies is known to hold for a very general class of price processes, the case of constant costs remains a challenge since the existence of optimal strategies is not even known in tractable models (such as, e.g., the Black-Scholes model). In this talk, we present a novel approach which allows us

to construct optimal strategies in a multidimensional diffusion market with price processes driven by a factor process and for general lower-bounded utility functions.

One of the main challenges for the problem under consideration is that the value function turns out to be piecewise but not globally continuous. We establish this result in two steps: (1) We apply the stochastic Perron's method to show that the value function is a discontinuous viscosity solution of the associated dynamic programming PDE (a parabolic non-local free boundary problem). (2) We prove a local comparison principle for viscosity solutions of this PDE, which implies uniqueness of the value function as well as piecewise continuity.

Having established piecewise continuity, we use a characterization of the value function as the pointwise infimum of a suitable set of superharmonic functions to construct optimal trading strategies. The advantage of this approach is that the pointwise infimum (i.e. the value function) inherits the superharmonicity property, which in turn allows us to prove a verification theorem for candidate optimal strategies requiring only piecewise continuity of the value function. An application of the verification theorem entails the existence of optimal strategies.

This talk is based on joint work with Sören Christensen (University of Hamburg).

Tomasz Bielecki
Illinois Institute of Technology

A dynamic model of central counterparty risk

We introduce a dynamic model of the default waterfall of derivatives CCPs and propose a risk sensitive method for sizing the initial margin (IM), and the default fund (DF) and its allocation among clearing members. Using a Markovian structure model of joint credit migrations, our evaluation of DF takes into account the joint credit quality of clearing members as they evolve over time. Another important aspect of the proposed methodology is the use of the time consistent dynamic risk measures for computation of IM and DF. We carry out a comprehensive numerical study, where, in particular, we analyze the advantages of the proposed methodology and its comparison with the currently prevailing methods used in industry.

Nicholas Bingham
Imperial College London

Four themes from the work of Larry Shepp

We deal with four themes (of many!):

(i) Gaussian processes. Here we look at Shepp's work with Mike (M. B.) Marcus, TAMS (1970) and Proc. 6th Berkeley Symposium II (1972), largely on conditions for path-continuity of Gaussian processes, and briefly survey progress in the field since.

(ii) Spurious (nonsense, volatile) correlation. Here we discuss the work of Ernst, Shepp and Wyner, Ann Stat 45 (2017). The standard deviation of the correlation θ of the integrals $\int_0^1 W_i(t)dt$ of two independent Brownian motions W_i is found (by a tour de force of analysis) to be nearly $\frac{1}{2}$. This explains a phenomenon known since Yule in 1926. Links with spurious regression, and its importance in macroeconomics, are considered.

(iii) Stationary Gaussian Markov processes as limits of autoregressive time series. The paper by Ernst, Brown, Shepp and Wolpert, J. Multiv. Anal. (2017), is discussed. We focus particularly on the differentiability (holomorphy) properties of the paths, and the links between discrete and continuous time.

(iv) Occupation times of Brownian motion in quadrants. The paper here is Ernst and Shepp, JAP (2017), on occupation times in the *first and third quadrants* of the plane for planar Brownian motion a problem simple to state but very hard to solve. They obtained partial results using Kontorovich-Lebedev transforms. We discuss related results involving strong (a.s.) forms of the arc-sine law.

Anthony Bloch
University of Michigan

Control and dynamics of nonlinear and quantum systems with dissipation

In this talk we discuss aspects of the physics and mathematics of a quantum control system interacting with its environment as well as some related general ideas in nonlinear control theory. In particular we discuss the control of an finite-dimensional dissipative Lindblad system by considering the geometry of its orbit and interorbit dynamics. This entails considering the geometry of the system, the structure of the Lindblad operator, and the convexity associated with the density equation. Applications are given to constructing pure states. We will also discuss the general structure of controlled quantum and nonlinear systems and the relation of the dissipative dynamics to system noise. This includes recent work with Brockett, Rooney and Rangan.

Nyles Breecher
University of Wisconsin-Milwaukee

Dynamic pricing with variable order sizes for a model with constant demand elasticity

We investigate a dynamic pricing model for constant demand elasticity, generalized to allow customers to have a probability distribution on the number of items they order. We first obtain a closed form expression for the optimal expected revenue and optimal pricing strategy. This expression relates to the average order size, and involves a recursively defined term whose behavior we analyze. We find that comparable models–those with the same average rate of sales, even with different order size distributions–have the same asymptotic pricing behavior in the inventory size. Moreover, the relative difference between comparable models is independent of the magnitude of the sales rate. Numerical results support that these ideas also should apply for models with different types of demand functions.

Edmond Brophy
University of North Texas

Toward prophet inequalities for vector-valued random variables with a cost for observation

I will present the problem of finding prophet inequalities for Ferguson and Bruss' problem of selling multiple assets with vector offers and cost for observation. Restricting the focus to the case of selling two assets, partial results will be given as well as an examination of the difficulty in transforming the usual one-dimensional techniques for solving prophet inequalities to solving problems with multivariate r.v.s.

F. Thomas Bruss
Université Libre de Bruxelles

From the $1/e$-law over time transforms to basket strategies

This talk in honour of Professor Larry Shepp begins with reminiscences of meetings and discussions I was pleased to have with him between 1983 and 2013. Although I have known Prof. Shepp altogether probably less than many other participants at this conference, and have no co-authored paper with him, I will always remember several of these little interactions.

The talk will then quickly turn towards a few elementary results exemplifying these interactions. They mainly concern the theory and practice of optimal stopping. In the talk we tempt a somewhat closer analysis of their "practical" efficiency and discuss possible improvements and open problems. These results are the so-called $1/e$-law of best choice (1984), properties of Pascal processes (1991), the odds-algorithm of optimal stopping (2000), Robbins' Problem of the expected rank (2005), the last arrival problem (2012), and also a result on a related stopping time (Steele 2016). In this short review I will also recall several side comments Prof. Shepp made on some aspects of these, where we agreed in our regret of the lack of applicability in the real world, in particular on the stock market, but also in our feeling "nevertheless very interesting and worth doing."

Finally, we present in our talk an elementary, and possibly new, idea concerning what one could call *basket trading*.

References

Bruss F. T. (1984) *A unified approach to a class of best choice problems with an unknown number of options,* Ann. of Probab., Vol. 12., No 3, 882-889.

Bruss F. T. and L.C.G. Rogers (1991) *Pascal processes and their characterisation,* Stoch. Proc. & Th. Applic., Vol. 37, 331-338.

Bruss F. T. (2000) *Sum the odds to one and stop,* Ann. of Probab., Vol. 28, No 3, 1384-1391.

Bruss F. T. (2005) *What is known about Robbins' problem?,* J. of Appl. Probab. 42 (1), 108-120

Bruss F. T. and M. Yor (2012) *Stochastic processes with proportional increments and the last-arrival problem,* Stoch. Proc. & Th. Applic., Vol. 122 (9), 3239-3261.

Steele J. M. (2016) *The Bruss-Robertson-inequality: Elaborations, Extensions, and Applications*, Math. Applicanda, Vol. 44, No 1, 3-16.

Amarjit Budhiraja
University of North Carolina at Chapel Hill

Large deviations from the hydrodynamic limit for a system with nearest neighbor interactions

We give a new proof of the large deviation principle from the hydrodynamic limit for the Ginzburg-Landau model studied in Donsker and Varadhan (1989) using techniques from the theory of stochastic control and weak convergence methods. The proof is based on characterizing subsequential hydrodynamic limits of controlled diffusions with nearest neighbor interaction that arise from a variational representation of certain Laplace functionals. The approach taken here does not require superexponential probability estimates, estimation of exponential moments, or an analysis of eigenvalue problems, that are central ingredients in previous proofs. Instead, proof techniques are very similar to those used for the law of large number analysis, namely in the proof of convergence to the hydrodynamic limit (cf. Guo, Papanicolaou, Varadhan (1988)). Specifically, the key step in the proof is establishing suitable bounds on relative entropies and Dirichlet forms associated with certain controlled laws. This general approach has the promise to be applicable to other interacting particle systems as well and to the case of non-equilibrium starting configurations, and to infinite volume systems. Joint work with Sayan Banerjee and Michael Perlmutter.

Bruno Buonaguidi
Università della Svizzera Italiana

Recent developments on optimal variance stopping problems

Optimal variance stopping (O.V.S.) problems are a new class of optimal stopping problems that differ form the classical ones because of their non linear dependence on the expectation operator. In an O.V.S. problem the goal is to determine the stopping time at which the variance of a sequentially observed stochastic process is maximized. In this talk, the solution method proposed by Pedersen (2011) is firstly introduced. Then, using the methodology developed by Pedersen (2016), it will be shown that the solution to an O.V.S. problem can be equivalently obtained by first solving a constrained optimal stopping problem and then by maximizing the obtained solution over the set of the admissible constraints. Some examples will be presented to illustrate this technique. Moreover, it will also be discussed how the solution to the O.V.S. problem for a geometric Brownian motion can be used in trading strategies.

References

Pedersen, J. L. (2011). *Explicit Solutions to Some Optimal Variance Stopping Problems,* Stochastics 83: 505-518.

Pedersen, J. L. and Peskir, G. (2016). *Optimal Mean-Variance Selling Strategies,* Mathematics and Financial Economics 10: 203-220.

Agostino Capponi
Columbia University

The term structure of liquidity: A liquidation game approach

We analyze a dynamic liquidation game where both liquidity demand and supply are endogenous. A large uninformed investor strategically liquidates a position, fully cognizant of the optimal response of competitive market makers. The Stackelberg game solution shows that, if the investor reveals the duration of the trade to the intermediation sector, then he chooses to sell at higher intensity when he has less time to trade. This enables market makers to predict when execution ends, which helps them provide liquidity and thus reduces the liquidity premium they charge. The model explains several empirical facts: order duration and participation rate correlate negatively, and price pressure subsides before execution ends.

Authors: Agostino Capponi (Columbia), Albert Menkveld (VU Amsterdam), and Hongzhong Zhang (Columbia).

Xi Chen
New York University

Statistical inference for model parameters with stochastic gradient descent

In this talk, we investigate the problem of statistical inference of the true model parameters based on stochastic gradient descent (SGD) with Ruppert-Polyak averaging. To this end, we propose a consistent estimator of the asymptotic covariance of the average iterate from SGD — batch-means estimator, which only uses the iterates from SGD. As the SGD process forms a time-inhomogeneous Markov chain, our batch-means estimator with carefully chosen increasing batch sizes generalizes the classical batch-means estimator designed for time-homogeneous Markov chains. The proposed batch-means estimator allows us to construct asymptotically exact confidence intervals and hypothesis tests.

Sören Christensen
University of Hamburg

Finding the optimal threshold for general Markov problems

An approach for solving one-sided optimal stopping problems in discrete and continuous time for general underlying Markov processes on the real line is presented. The main idea is to transform the problem into an auxiliary problem for the ladder height variables. In the case that the original problem has a one-sided solution and the auxiliary problem has a monotone structure, the corresponding myopic stopping time is optimal for the original problem as well. This elementary line of argument directly leads to a characterization of the optimal boundary in the original problem: The optimal threshold is given by the threshold of the myopic stopping time in the auxiliary problem. Supplying also a sufficient condition for our approach to work, we obtain solutions for many prominent examples in the literature, among others

the problems of Novikov-Shiryaev, Shepp-Shiryaev, the American put in option pricing, and the sum-the-odds-theorem. As a further application we show that for underlying random walks (and Lévy processes in continuous time), the reward functions leading to one-sided stopping problems are exactly the monotone and log-concave functions. We furthermore discuss the connection to other approaches in the literature and apply the results to obtain explicit solutions for impulse control problems for general underlying jump processes.

Alexandra Chronopoulou
University of Illinois Urbana-Champaign

Delta-hedging in fractional volatility models

In this talk, we propose a delta-hedging strategy for a long memory stochastic volatility model (LMSV). This is a model in which the volatility is driven by a fractional Ornstein-Uhlenbeck process with long-memory parameter H. We compute the so-called hedging bias, i.e. the difference between the Black-Scholes delta and the LMSV delta as a function of H, and we determine when a European-type option is over-hedged or under-hedged. Finally, we apply our approach to S&P 500 data. This is joint work with Qi Zhao (UIUC).

Asaf Cohen
University of Haifa

Fluctuations in finite state many player games

We consider an n-player symmetric stochastic game with weak interactions between the players. Time is continuous and the horizon and the number of states are finite. We show that the value function of each of the players can be approximated by the solution of a partial differential equation called the master equation. Moreover, we analyze the fluctuations of the empirical measure of the states of the players in the game and show that it is governed by a solution to a stochastic differential equation. Joint work with Erhan Bayraktar.

Alexander Cox
University of Bath

Optimal solutions to the Skorokhod embedding problem via the stochastic maximum principle

We first show the equivalence of a modified version of the Optimal Skorokhod Embedding Problem and a forward-backward stochastic differential equation (FBSDE), the solutions to which can be thought of as randomised stopping times. From the Stochastic Maximum Principle, we are then able to derive conditions which should be satisfied by any optimisers. As an application, we are able to re-derive the form of the dual solution for the Root solution to the SEP. This is joint work with Sam Kinsley and Jiajie Wang.

Dennis Cox
Rice University

Applications of the randomized probability integral transform

The Probability Integral Transform (PIT) applied to a univariate random variable gives a uniformly distributed random variable, but it only applies to continuous distributions. A randomized version of the PIT will produce a uniformly distributed random variable for discrete or mixed discrete/continuous random variables. Based on the randomized PIT, we introduce an empirical measure with discrete components (from continuous random variables) and continuous components (from discrete random variables). We look at several questions that arise from this empirical distribution, including weak convergence of the corresponding empirical process and goodness of fit tests in predictive modeling.

Daryl Daley
University of Melbourne

The lilypond growth protocol for germ–grain models in $\mathbb{R}^d$

This talk describes how I came to work with Larry Shepp on a germ–grain problem posed to me by Dietrich Stoyan and which in turn I posed to Larry when he was visiting Canberra in October/November 1996. He gave a solution to the particular problem, but when he described it in an abstract for a talk after return to US, Colin Mallows gave a slicker solution. The ensuing three-authored paper gave results for even more related problems, all for a 1-dimensional setting of grains. I also describe extensions of the lilypond germ-grain model in $\mathbb{R}^d$ that suggest interest in seeking a limit result for $d \to \infty$ and a percolation property for randomly constructed finite line-segments.

Mark Davis
Imperial College of London

Optimal stopping for piecewise deterministic Markov processes

PDMPs are Borel right processes, so much of general optimal stopping applies to them, but because of their special structure more specific things can be said. The PDMP generator comes with associated boundary conditions, and the talk will focus on the inter-relationship between these conditions and the smooth pasting condition.

Tiziano De Angelis
University of Leeds

On some probabilistic results concerning smoothness of the value function and of the free boundary in optimal stopping

I will present probabilistic proofs of some important regularity properties for the value function of general optimal stopping problems and for the associated optimal boundaries. Most of our arguments rely on fundamental concepts from the theory of Markov processes and bridge the probabilistic and the analytical strands of the literature on free boundary problems. I will also illustrate situations in which our work improves or complements known facts from PDE theory.

One of our results states that, if the underlying process in the optimal stopping problem is strong Feller with differentiable flow, and if the stopping boundary is regular for the stopping set (in the sense of diffusions), then the value V of the optimal stopping problem is C^1 in the whole space (provided also that the gain function is C^1). This improves on the canonical smooth-fit condition, which is instead understood in the sense of continuity for one directional derivative of V. Notice also that the concept of regularity in the sense of diffusions replaces more stringent conditions often imposed in the related PDE literature.

Finally, I will also present probabilistic proofs of continuity and *Lipschitz continuity* for optimal stopping boundaries in multi-dimensional problems. In some instances we succeed in relaxing standard assumptions made in the PDE literature, as for example uniform ellipticity of the underlying diffusion. (Notice that Lipschitz boundaries are very often regular in the sense of diffusions.)

This talk draws from joint work with G. Peskir (University of Manchester) and G. Stabile (Sapienza University of Rome).

Victor de la Pena
Columbia University

On a decoupling approach to boundary crossing

In this talk I introduce a decoupling approach to boundary crossing. This approach involves the decoupling of the process and its associated crossing time. A sharp lower bound for arbitrary processes is provided.

John Dobelman
Rice University

The trader's "2% Rule" for money management

Among proficient traders it is well known that money management is one of the key factors in the success of a trading program. In this paper we examine the so called "2% rule" and observe that it is a modification to the trade exit strategy and is informed by and affects the firm's equity curve (wealth). Using resampled trading data we explore the rule's efficacy. We also examine how this "rule" impacts the optimal strategy for the support and resistance trading system.

Yan Dolinsky
Monash University

Numerical scheme for Dynkin games under model uncertainty

We introduce an efficient numerical scheme for continuous time Dynkin games under model uncertainty. We use the Skorokhod embedding in order to construct recombining tree approximations. This technique allows us to determine convergence rates and to construct numerically optimal stopping strategies. We apply our method to several examples of game option.

Francois Dufour
Université de Bordeaux

Approximation of Markov decision processes

In this talk, we propose an approach for approximating of Markov decision processes with Borel state and action spaces. Under adequate assumptions, which in particular include that the transition probabilities has a density function with respect to a reference measure, together with Lipschitz continuity of the elements of the control model, we approximate the original controlled process by a model with finite state and action spaces. The approximation error is related to the 1-Wasserstein distance between the reference probability measure and a measure with finite support. We study the case when the reference measure is approximated with empirical distributions and we show that convergence of the approximations takes place at an exponential rate in probability. This technique can be applied to several different classes of control problems.

Tyrone Duncan
University of Kansas

Explicitly solvable stochastic differential games

Many physical systems can be modeled by stochastic differential games that evolved from stochastic control and minimax problems. In these games two or more agents choose control strategies for a system to minimize or maximize a payoff functional. R. Isaacs obtained nonlinear partial differential equations for the upper and the lower values of a game and these two nonlinear partial differential equations for a game become one for games that have a value. This latter equation is known as a Hamilton-Jacobi-Isaacs equation which provides the strategies for the agents. Another approach to stochastic differential games is an application of a stochastic maximum principle with the solution of a backward stochastic differential equations. Both of these approaches are difficult to obtain explicit solutions. A direct method has been used to solve a collection of stochastic differential games. This method does not require solving nonlinear partial differential equations or backward stochastic differential equations. A variety of stochastic differential games are described where this direct method provides explicit optimal strategies. These stochastic differential games include

linear stochastic differential games with quadratic or risk sensitive exponential quadratic payoffs, some nonlinear stochastic differential games that evolve in symmetric spaces, and some linear stochastic differential games with quadratic payoffs and state dependent fractional Brownian motion noise. The method can also be used to solve some games described by infinite dimensional systems.

Philip Ernst
Rice University

Yule's "nonsense" correlation solved!

To honor the memory of Larry Shepp, I discuss a problem that interested him for over fifty years. In 2017, we finally closed this longstanding open problem. The problem, formulated by the British statistician Udny Yule in 1926, is to mathematically prove Yule's 1926 empirical finding of "nonsense correlation". We solve the problem by analytically determining the second moment of the empirical correlation coefficient of two independent Wiener processes. Using tools from Fredholm integral equation theory, we calculate the second moment of the empirical correlation to obtain a value for the standard deviation of the empirical correlation of nearly .5. The "nonsense" correlation, which we call "volatile" correlation, is volatile in the sense that its distribution is heavily dispersed and is frequently large in absolute value. It is induced because each Wiener process is "self-correlated" in time. This is because a Wiener process is an integral of pure noise and thus its values at different time points are correlated. In addition to providing an explicit formula for the second moment of the empirical correlation, we offer implicit formulas for higher moments of the empirical correlation. The full paper (joint work with L.A. Shepp and A.J. Wyner) was recently published in *The Annals of Statistics* and can be found at `https://projecteuclid.org/euclid.aos/1498636874`.

References

Ernst, P.A., Shepp, L.A., and Wyner, A.J. (2017). Yule's "nonsense correlation" solved! *The Annals of Statistics* **45**: 1789-1809.

Yule, G.U. (1926). Why do we sometimes get nonsense-correlations between time-series?–a study in sampling and the nature of time-series. *Journal of the Royal Statistical Society* **89**: 1-63.

Eugene Feinberg
Stony Brook University

Lebesgue's convergence theorem and Fatou's lemma for varying probabilities

The talk describes Fatou's lemma, the monotone convergence theorem, and Lebesgue's dominated convergence theorem for varying probabilities. In particular, we discuss the cases of weak convergence, setwise convergence, and convergence in total variation. In general, Fatou's lemma and convergence theorems do not hold in their classic forms for weakly converging probabilities, but they hold under stronger assumptions on convergence of functions. We formulate such assumptions. We discuss applications of these results to Markov Decision Processes as well as applications of the uniform Fatou lemma, which holds for probabilities converging in total variation. The talk is based on joint papers with Pavlo Kasyanov, Yan Liang, Nina Zadoianchuk, and Michael Zgurovsky.

Georgios Fellouris
University of Illinois Urbana-Champaign

Change acceleration and detection

A fundamental problem in e-learning environments and intelligent tutoring systems is to understand when a student has learned a skill, while minimizing the required time for learning. Motivated by these applications, we propose a generalization of the Bayesian sequential change detection problem, where the change is a latent event that should be not only detected, but also accelerated. Specifically, we assume that the sequentially collected observations are responses to treatments selected in real time. The assigned treatments not only determine the distribution of responses before and after the change, but also influence the distribution of the latent change-point. The problem is to find a treatment assignment rule and a stopping rule that minimize the expected total number of observations subject to a bound on the false-detection probability. We will propose an intuitive solution to this problem, which is easy to implement and achieves the optimal performance up to a first-order asymptotic approximation. This asymptotic optimality is established under a large class of change-point models, for which a dynamic programming approach may not be feasible or applicable. Finally, we will illustrate with a simulation study the almost optimality of the proposed approach in the case of a simple (Markovian) change-point model for which the optimal solution is computable. This is joint work with Yanglei Song from the University of Illinois, Urbana-Champaign.

Liming Feng
University of Illinois Urbana-Champaign

American options in Lévy models

Stock options traded on exchanges are usually of the American style. No closed-form solutions are available for American option prices even in the simplest Black-Scholes-Merton model. One must rely on efficient numerical methods. In model calibration and risk management applications, one must constantly evaluate large amounts of options. Being able to compute the option prices quickly and accurately is essential for the success of such applications. We present a Hilbert transform based method for the pricing of discrete time American (Bermudan) options. When the functions being transformed are in a certain analytic class, some simple schemes become very efficient, with exponentially decaying errors. Continuous time American options can be priced through simple extrapolation.

Ricardo Fernholz
Claremont McKenna College

Universality of Zipf's Law for time-dependent rank-based systems

A set of empirical data with positive values follows a power law or Pareto distribution if a log-log plot of the data versus rank is approximately a straight line. A Pareto distribution follows Zipf's law if the slope of the log-log plot is -1. Zipf's law is a form of universality, since many classes of data approximately follow this distribution. Specifically, certain time-dependent rank-based systems follow Zipf's law, and we characterize these systems. We generalize Zipf's law to a less restrictive form in which the log-log plot of the distribution is not necessarily a straight line with slope -1, but rather is concave with a tangent line of slope -1 at some point on the curve. It is this less restrictive version of Zipf's law that usually holds for actual data.

Giorgio Ferrari
University of Bielefeld

Optimal reduction of public debt under partial observation of the economic growth

We consider a government that aims at reducing the debt-to-GDP ratio of a country. The government can observe the level of the debt-to-GDP ratio, and an indicator of the state of the economy, but does not directly observe the development of the macroeconomic conditions. The government aims at determining a debt reduction policy that minimizes the sum of the total expected costs of holding debt, and of intervening on it. We model this problem as a singular stochastic control problem under partial information. We first provide a general formulation of the model in which the level of debt-to-GDP ratio and the value of the macroeconomic indicator evolve as a diffusion and a jump-diffusion, respectively, whereas the (unobserved) condition of the economy is a continuous-time Markov chain with Q states. We then reduce via filtering the original problem to an equivalent one with full information (the so-called separated problem), and we relate the latter to an optimal stopping problem by means of a verification theorem. Finally, we specialize to a case study in which the economy faces only two regimes, and the macroeconomic indicator is a Brownian motion with regime-dependent drift. In this setting we provide the optimal debt reduction policy in terms of the continuous free boundary arising in an auxiliary truly two-dimensional optimal stopping problem.

Jose Figueroa-Lopez
Washington University in St. Louis

Optimal change-point detection for Lévy processes

Since the work of Page in the 1950s, the problem of detecting an abrupt change in the distribution of stochastic processes has received a great deal of attention. In particular, a deep connection has been established between Lorden's minimax approach to change-point

detection and the widely used CUSUM procedure, first for discrete-time processes, and subsequently for some of their continuous-time counterparts. However, results for processes with jumps are still scarce, while the practical importance of such processes has escalated since the turn of the century. In this work, we show that CUSUM is indeed optimal in Lorden's sense for detecting a change point of Lévy processes. This is joint work with Sveinn Olafsson.

Leopold Flatto
City University of New York

Dixie cup problem and FKG inequality

Let $T_m(n)$ be the number of purchases required to obtain m copies of n given items, each purchase choosing at random one of the n items. $\mathbb{E}[m(n)]$ is the expected value of the random variable $T_m(n)$

The problem of obtaining a formula for $\mathbb{E}[m(n)]$ is known as the dixie cup problem. The problem is easy for $m = 1$, but difficult for $m > 1$. Newman and Shepp (N, S) solve the problem for all m, n. From their formula, they obtain the asymptotics of $\mathbb{E}[m(n)]$, for each fixed m and n tending to infinity.

Later, Erdős and Renyi obtain the limit law for $T_m(n)$, for each fixed m and n tending to infinity. From the limit laws, they also derive and improve on the results of (N, S). The derivation is however incomplete, as they do not address the problem of estimating the tails of the distributions of $T_m(n)$.

In my talk, I provide the estimations, by using the FKG inequality, a correlation inequality which is a fundamental tool in statistical mechanics and probabilistic combinatorics.

Jean-Pierre Fouque
University of California, Santa Barbara

Optimal portfolio under fractional stochastic environment

Rough stochastic volatility models have attracted a lot of attention recently, in particular for the linear option pricing problem. In this talk, starting with power utilities, we propose to use a martingale distortion representation of the optimal value function for the nonlinear asset allocation problem in a (non-Markovian) fractional stochastic environment (for all Hurst indices $H \in (0, 1)$). We rigorously establish a first order approximation of the optimal value, when the return and volatility of the underlying asset are functions of a stationary slowly varying fractional Ornstein-Uhlenbeck process. We prove that this approximation can be also generated by the zeroth order trading strategy providing an explicit strategy which is asymptotically optimal in all admissible controls. Furthermore, we extend the discussion to general utility functions, and obtain the asymptotic optimality of this strategy in a specific family of admissible strategies. If time permits, we will also discuss the problem under fast mean-reverting fractional stochastic environment. Joint work with Ruimeng Hu (UCSB).

Rüdiger Frey
Wirtschafts Universität Wien

Optimal liquidation under partial information with price impact

We study the problem of a trader who wants to maximize the expected revenue from liquidating a given stock position. We model the stock price dynamics as a geometric pure jump process with local characteristics driven by an unobservable finite-state Markov chain and by the liquidation rate. This reflects uncertainty about activity of other traders and feedback effects from trading. Stochastic filtering is used to reduce the optimization problem under partial information to an equivalent one under complete information. This leads to a stochastic control problem for piecewise deterministic Markov processes (PDMPs). We carry out a detailed mathematical analysis of this problem. In particular, we derive the Bellman equation for the value function, we characterize the value function as continuous viscosity solution of the associated dynamic programming equation, and we prove a novel comparison result. The paper concludes with numerical results illustrating the impact of partial information and feedback effects on the optimal liquidation rate. Moreover, we compare the performance of our model to more classical approaches.

Joint work with Katia Colaneri, Zehra Eksi and Michaela Szoelgyenyi.

Eric Friedlander
University of North Carolina at Chapel Hill

Large deviations for load balancing mechanisms in large queueing networks

Large deviation principles (LDP) are an invaluable tool in the estimation of probabilities of rare-events. In this talk, I will explore LDP for several load balancing mechanisms employed in large queueing networks. I will present some of the technical challenges that arise when establishing LDP in such systems and discuss the implications of such results. As one of the goals of load balancing is to improve retrieval speed (i.e. reduce the proportion of long queues), estimating the probability of having long queues under a particular routing scheme is of interest. The LDP studied in our work can be used for estimating such probabilities. This is joint work with Amarjit Budhiraja and Ruoyu Wu.

Ayalvadi Ganesh
University of Bristol

Large deviations for Cox processes and Cox/G/infinity queues, with a biological application

We show that a sequence of Cox processes on a Polish space E satisfy a large deviation principle (LDP), provided their directing measures do so on the space of finite measures on E equipped with the weak topology. Next, we consider a sequence of infinite server queue with general i.i.d. service times, where the arrivals constitute Cox processes with translation invariant directing measures assumed to satisfy an LDP. We show that the corresponding sequence of queue occupancy measures also satisfy an LDP. These problems were motivated by the problem of describing fluctuations of molecule numbers in biochemical reaction networks within cells.

Pavel Gapeev
London School of Economics

On some two-dimensional optimal stopping problems for diffusion processes

We present closed form solutions to certain inherently two-dimensional optimal stopping problems for some two-dimensional continuous Markov diffusion processes which cannot be reduced to the problems for one-dimensional diffusion processes. It is well known that optimal stopping problems for multi-dimensional diffusion processes are analytically more difficult than the appropriate problems for one-dimensional ones and their solutions are very rarely found explicitly. Some necessarily two-dimensional optimal stopping problems for two-dimensional diffusion and jump processes appearing mostly in sequential testing and quickest change-point detection and mathematical finance were studied in the previous literature. We reduce the original inherently two-dimensional problems to the associated families of time-independent optimal stopping problems and derive explicit expressions for the optimal stopping boundaries. Joint work with Goran Peskir (Manchester).

Paul Glasserman
Columbia University

Information-driven price and volatility cycles

We propose and analyze a dynamic model in which stock market prices and volatility shift between higher and lower levels based on changes in the information available to investors about future dividends. The model evolves through overlapping generations. In each generation, investors decide whether to acquire information at a cost. The precision of available information is subject to exogenous shocks; in addition, more information becomes available as more investors choose to become informed. The model gives rise to multiple fixed points for the fraction of informed investors, corresponding to different price and volatility regimes. Investors rationally anticipate potential transitions between regimes. We calibrate the model to historical data and find that the transitions can generate large price drops accompanied by large increases in volatility. The pattern is consistent with the onset of a financial crisis and other scenarios. The model isolates the role of information dynamics in producing this pattern. This is joint work with Harry Mamaysky and Yiwen Shen.

Kris Glover
University of Technology Sydney

A Dynkin game with incomplete and asymmetric information

We study the effects of asymmetry of information for a two-person zero-sum optimal stopping game with linear payoffs. In our set-up, the drift of the underlying diffusion process is unknown to one player, but known to the other one. However, employing a Bayesian approach, the uninformed player is equipped with an initial prior distribution for the drift, and may subsequently learn about the drift from observations of the process and from the actions

(or lack of action) of the informed player. We show that there exists a Nash equilibrium for this game, where the uninformed player uses a stopping time and the informed player uses a randomized stopping time in order to hide their informational advantage. Joint work with Erik Ekström.

David Goldberg
Cornell University

Novel pure-dual algorithm for optimal stopping

We present a new pure-dual algorithm for optimal stopping, inspired by a connection to flows in networks. Our method leads to a hierarchy of approximations, which (in contrast to much past work) are not defined through backwards induction (nor through a restricted family of basis functions or standard dynamic programming techniques). Instead, our approach leads to an expansion in which each term holistically considers how the true process behaves over the entire time horizon, which can enable better approximation with fewer nested conditional expectations. We also prove bounds on the rate of convergence, discuss implications for simulation-based approaches, and (time permitting) explore various applications and generalizations. Joint work with Ph.D. student Yilun Chen.

Ilie Grigorescu
University of Miami

Scaling and optimal strategy in Shepp's urn and the war of attrition

Second order asymptotics are derived for two urn problems related to Larry Shepp's work. The first is Shepp's urn with risk aversion, where we derive an approximate strategy for $V(m,n,k)-k>0$ and critical the k, where $V(m,n,k)$ is the value of the game when the urn has m balls marked $+1$, n marked -1, and k is the initial fortune. The risk aversion consists in stopping when the fortune is zero. The second is the *war of attrition*, where we prove the weak limit of the ruin time, equal to a non-central chi-square distribution. The results are obtained under diffusive scaling to Brownian bridge-like processes. Joint work with Robert Chen, Philip Ernst, and Min Kang.

Paolo Guasoni
Boston University

Asset prices in segmented and integrated markets

This paper evaluates the effect of market integration on prices and welfare, in a model where two Lucas trees grow in separate regions with similar investors. We find equilibrium asset price dynamics and welfare both in segmentation, when each region holds its own asset and consumes its dividend, and in integration, when both regions trade both assets and consume both dividends. Integration always increases welfare. Asset prices may increase or decrease, depending on the time of integration, but decrease on average. Correlation in assets' returns is zero or negative before integration, but significantly positive afterwards, explaining some effects commonly associated with financialization.

Olympia Hadjiliadis
City University of New York

Multi-dimensional quickest detection

We consider the problem of quickest detection in the presence of multiple correlated random sources each driven by distinct sources of noise represented by a Brownian motion.

The first problem we will address is the one of detecting a change in the drift of independent Brownian motions received in parallel at the sensors of decentralized systems. We examine the performance of one shot schemes in decentralized detection in the case of many sensors with respect to appropriate criteria. One shot schemes are schemes in which the sensors communicate with the fusion center only once; when they must signal a detection. The communication is clearly asynchronous and we consider the case that the fusion center employs the minimal strategy, namely an alarm is issued at the fusion center the moment in which the first one of the sensors issues an alarm. We prove asymptotic optimality of the above strategy not only in the case of independent sources of data but also in the presence of across-sensor correlations and specify the optimal threshold selection at the sensors. Moreover, we consider the problem of quickest detection of signals in a coupled system of N sensors, which receive continuous sequential observations from the environment. It is assumed that the signals, which are modeled by a general Itô processes, are coupled across sensors, but that their onset times may differ from sensor to sensor. Two main cases are considered; in the first one signal strengths are the same across sensors while in the second one they differ by a constant. The objective is the optimal detection of the first time at which any sensor in the system receives a signal. The problem is formulated as a stochastic optimization problem in which an extended minimal Kullback-Leibler divergence criterion is used as a measure of detection delay, with a constraint on the mean time to the first false alarm. The case in which the sensors employ cumulative sum (CUSUM) strategies is considered, and it is proved that the minimum of N CUSUMs is asymptotically optimal as the mean time between false alarms increases without bound. In particular, in the case of equal signal strengths across sensors, it is seen that the difference in detection delay of the N-CUSUM stopping rule and the unknown optimal stopping scheme tends to a constant related to the number of sensors as the mean time between false alarms increases without bound. While in the case of unequal signal strengths, it is seen that this difference tends to 0.

Yu-Jui Huang
University of Colorado Boulder

Optimal equilibria for time-inconsistent stopping problems

For time-inconsistent control/stopping problems, it is known that one should employ an equilibrium strategy, formulated in an intertemporal game between current and future selves. Such strategies, however, are not unique. This gives rise to two unsettled problems: (i) How do we find all equilibria? (ii) Among all equilibria, how do we select the appropriate one to use? For stopping problems under non-exponential discounting, we develop a new method,

called the iterative approach, to resolve both (i) and (ii). First, we formulate equilibria as fixed points of an operator, which represents strategic reasoning that takes into account future selves' behavior. Under appropriate regularity conditions, every equilibrium can be found through a fixed-point iteration. When the state process is one-dimensional, we further establish the existence of an optimal equilibrium, which generates larger values than any other equilibrium does at all times. To the best of our knowledge, this is the first time a dominating subgame perfect Nash equilibrium is shown to exist in the literature of time-inconsistency. Our theory is illustrated explicitly in several real options models.

Fred Huffer
Florida State University

The probability of covering the circle by random arcs: qualitative results and numerical explorations

Suppose arcs of fixed lengths ℓ_k, $0 < \ell_k < 1$, $k = 1, 2, \ldots, n$, are thrown independently and uniformly on a circumference $\mathcal{C}$ having unit length. Let $P(\ell_1, \ell_2, \ldots, \ell_n)$ be the probability that $\mathcal{C}$ is completely covered by the n random arcs. Huffer and Shepp (1987) showed that P is a Schur-convex function and that it is convex in each argument when the others are held fixed. We review the results of Huffer and Shepp (1987) and show how they are related to the results in Huffer (1986, 1987) concerning situations where the arc lengths are random. In these papers it is shown that probabilities of coverage and multiple coverage increase when the arc lengths are made more variable. Finally, we show how $P(\ell_1, \ell_2, \ldots, \ell_n)$ may be computed numerically, and use numerical explorations to investigate some conjectures concerning the coverage probability.

Clifford Hurvich
New York University

The slow convergence of OLS estimators of α, β and portfolio weights under long-memory stochastic volatility

We consider inference for the market model coefficients based on simple linear regression under a long memory stochastic volatility generating mechanism for the returns. We obtain limit theorems for the ordinary least squares (OLS) estimators of α and β in this framework. These theorems imply that the convergence rate of the OLS estimators is typically slower than $\sqrt{T}$ if both the regressor and the predictor have long memory in volatility, where T is the sample size. The traditional standard errors of the OLS-estimated intercept ($\hat{\alpha}$) and slope ($\hat{\beta}$), which disregard long memory in volatility, are typically too optimistic, and therefore the traditional t-statistic for testing, say, $\alpha = 0$ or $\beta = 1$, will diverge under the null hypothesis. We also obtain limit theorems (which imply slow convergence) for the estimated weights of the minimum variance portfolio and the optimal portfolio in the same framework. In addition, we propose and study the performance of a subsampling-based approach to hypothesis testing for α and β. Joint work with Jun Liu (Bank of New York Mellon) and Rohit Deo (New York University Stern School).

Tomoyuki Ichiba
University of California, Santa Barbara

An optimal switching problem in rank-based market models

In this talk we discuss a switching problem in the rank-based market models with boundaries where drift and diffusion coefficients depend on the ranking of stocks in terms of their sizes among all the stocks in the market. We start with an optimal switching problem of two competing stocks and then extend to the multiple stocks. This study is motivated by an extension of optimal switching problem studied initially by A. Mandelbaum, L. A. Shepp and R. J. Vanderbei (1990).

Michael Imerman
Lehigh University

How to bailout if you must or a martingale approach to the question of fiscal stimulus

This talk is based on one of the last problems that I worked on with Larry before his untimely passing. Larry was always very interested in politics and wanted to develop a mathematical model to address the politico-economic debate between Democrats and Republicans. Specially the problem, as it currently stands, examines the debate that Democrats in the US say that government spending can be used to "grease the wheels' of the economy, create wealth, and increase employment; the Republicans say that government spending is wasteful, discourages investment, and so increases unemployment. These arguments cannot both be correct, but both arguments seem meritorious. We address this economic question of fiscal stimulus as a new optimal control problem extending the model of Radner-Shepp (1996). A unique solution is found using traditional martingale methods for stochastic optimization along with a numerical procedure to solve a non-homogeneous ODE as the root of an implicit function. Specifically, we find that there exists an optimal strategy with interesting mathematical properties. When Larry would teach stochastic optimization he always referenced the classical work by Dubins and Savage entitled "How to Gamble if You Must." The original title of our paper pays homage to that.

Yoshiaki Itoh
Institute of Statistical Mathematics

Invariants for interacting particle systems on graphs and rock-paper-scissors interacting particle systems

Suppose initially there are $N_i(0)$ particles at each vertex i of G, and that the particles interact to form a Markov chain: at each instant two particles are chosen at random, and if these are at *adjacent* vertices of G, one particle jumps to the other particle's vertex, each with probability 1/2. The process $\mathbf{N}$ enters a death state after a finite time when all the particles are in some *independent* subset of the vertices of G, i.e., a set of vertices with no edges between any two of them. The problem is to find the distribution of the death state,

$\eta_i = N_i(\infty)$, as a function of $N_i(0)$. We are able to obtain, for some special graphs, the *limiting* distribution of N_i if the total number of particles $N \to \infty$ in such a way that the fraction, $N_i(0)/S = \xi_i$, at each vertex is held fixed as $N \to \infty$. In particular we can obtain the limit law for the graph, $S_2 : \cdot\!\!-\!\!\cdot\!\!-\!\!\cdot$, having 3 vertices and 2 edges. (Itoh, Y., Mallows, C. and Shepp, L.A.: Explicit sufficient invariants for an interacting particle system, Journal of Applied Probability. Vol.35, 633-641 (1998)). For the complete graph G, the model is that of Moran (1958) for the Fisher-Wright random sampling effect in population genetics.

Suppose initially there are $N_i(0)$ particles at each vertices i of 1 (rock) , 2 (paper) and 3 (scissors) and that the particles interact to form a Markov chain: at each instant two particles are chosen at random, and if these are at vertices i and j, one particle jumps to the dominant particle's vertex. The process $\mathbf{N}$ enters a death state after a finite time when all the particles are in one vertex. We are able to obtain the asymptotic probability of coexistence of the three vertices for time t and extend the results to general $2s + 1$ types rock-paper-scissors particle system (Ann. Inst. Statisit. Math. (1973), J. Appl. Prob. (1979)). We find the deterministic approximation of the system is a nonlinear integrable system (Prog. Theor.Phys. (1987)).

Our recent studies on the rock-paper-scissors particle systems will be also considered.

Saul Jacka
University of Warwick

Optimal stopping, smooth pasting and the dual problem

In the Markovian setting we identify sufficient conditions for the value function of the optimal stopping problem to belong to the domain of the extended (martingale) generator of the underlying Markov process. We then show that the dual of the optimal stopping problem is a stochastic control problem for a controlled Markov process, and the optimal control is characterised by a function belonging to the domain of the martingale generator. Finally, we give an application to the smooth pasting condition.

Jean Jacod
Université Paris VI

Statistics for high-frequency observations of a stochastic process

It is often the case that one has to do statistical inference for a stochastic process, based on the observation of a single path of the process, at discrete times and over a finite time interval: in such a framework, estimating the law of the process is usually not feasible, however one can still have reasonable estimators, even consistent ones as the observation frequency increases, for some specific characteristics of the process. We will start with a quick review of those characteristics that can be consistently estimated within this framework, versus those which cannot. Then, restricting our attention to the estimation of the volatility and the degree of activity of jumps in the case of an Itô semimartingale, we will explain some recent developments and new results, including statements about the rate-optimality and in some cases asymptotic efficiency.

Sebastian Jaimungal
University of Toronto

Algorithmic trading with partial information: A mean field game approach

Financial markets are often driven by latent factors which traders cannot observe. Here, we address an algorithmic trading problem with collections of heterogeneous agents who aim to perform statistical arbitrage, where all agents filter the latent states of the world, and their trading actions have permanent and temporary price impact. This leads to a large stochastic game with heterogeneous agents. We solve the stochastic game by investigating its mean-field game (MFG) limit, with sub-populations of heterogeneous agents, and, using a convex analysis approach, we show that the solution is characterized by a vector-valued forward-backward stochastic differential equation (FBSDE). We demonstrate that the FBSDE admits a unique solution, obtain it in closed-form, and characterize the optimal behaviour of the agents in the MFG equilibrium. Moreover, we prove the MFG equilibrium provides an ϵ-Nash equilibrium for the finite player game. We conclude by illustrating the behaviour of agents using the optimal MFG strategy through simulated examples.

Huanqun Jiang
Oregon State University

Optimal dividend strategies under SNLP discounted by the exponential Lévy process.

De Finetti's problem will be extended to a more general model. It could be achieved with the introduction of the eigenfunction of extended generator of Lévy process and recent development on optimal dividend strategies under SNLP(spectrally negative Lévy process)-type model. In the first step, we will see the optimal strategies with Brownian motion-type accumulated interest rate. It is solved by separation method on HJB equation. Then from the perspective of fluctuation theory of SNLP, we see its extension to optimal barrier strategy for SNLP discounted by the exponential Lévy process. Lastly, from the very recent literature review, we will see the extensions of this work to models of doubly-reflected SNLP, refracted SNLP and reflected-refracted SNLP.

Abram Kagan
University of Maryland

Inference on weak signals in presence of an additive noise

The setup under study is

$$X_j = a_j\theta + \xi_j,\ j = 1, 2, \ldots$$

where $\{a_1, a_2, \ldots\}$ is a sequence of known numbers, θ a parameter of interest and $\xi_1, \xi_2, \ldots$ independent random variables with a common cdf $F(x)$. Detecting the presence of a known signal transmitted with an additive noise means testing the null hypothesis $H_0 : \theta = 0$ vs the alternative $H_1 : \theta = 1$.

In Shepp (1965) it was shown that if F is absolutely continuous with $F'(x) = f(x)$ such that the Fisher information is finite,

$$I = \int (f'(x)/f(x))^2 f(x)dx < \infty$$

the condition $\sum_1^\infty a_j^2 = \infty$ is necessary and sufficient for testing with zero probabilities of type 1 and type 2 errors. Shepp noticed "it was unexpected that the Fisher information plays such a central role".
If one turns to estimating the parameter, Shepp's condition becomes necessary and sufficient for consistent estimation of θ. The condition simply means that the Fisher information on θ in the infinite sequence $\{X_1, X_2, \ldots\}$ is infinite.
We study the case of weak signals when $\sum_1^\infty a_j^2 < \infty$ and prove that if $\xi's$ have a uniform distribution, $\xi_j \sim U(-c, c)$, the condition $\sum_1^\infty |a_j| = \infty$ is necessary and sufficient for consistent estimation of θ and also for testing H_0 vs H_1 with zero probabilities of type 1 and type 2 errors.
Denote by $C(\lambda)$ the class of sequences ("signals") $\{a_1, a_2, \ldots\}$ with $\sum_1^\infty |a_j|^\lambda = \infty$. A $\lambda > 0$ is called an exponent of F-detection if the presence of any signal in class $C(\lambda)$ in independent observations $X_1 = a_1 + \xi_1$, $X_2 = a_2 + \xi_2, \ldots$ with $\xi_j \sim F(x)$ can be detected with zero probabilities of type 1 and type 2 error.
We call $\lambda(F) = \inf\{\lambda : \lambda \text{ is an exponent of } F\text{-detection}\}$ the *F-exponent* (of signal detection). Note that $\lambda(F)$ is defined not directly in terms of F but in the terms of the classes of signals which presence is detected when the noise has the cdf $F(x)$. In a sense, the F-exponent characterizes the class of weakest signals which presence can be detected with zero probabilities of errors when $F(x)$ is the cdf of the noise. Here are a few properties of $\lambda(F)$.
(i) For all F, $\lambda(F) \leq 2$. Shepp(1965)
(ii) For $F_c(x) = F(cx)$, $\lambda(F_c) = \lambda(F)$ for any $c > 0$.
(iii) For an absolutely continuous F with finite Fisher information, $\lambda(F) = 2$.
(iv) For the cdf of a uniform distribution, $\lambda(F) = 1$. Some open questions of general interest will be discussed.

Thomas Kaijser
Linkoping University

On random flights and stochastic models for laser bathymetry and water motion

I plan to start my talk by presenting some simple formulas for the mean of the square of the length of a random flight measured from its starting point - both after N changes of direction and at time T.

I will then discuss Monte Carlo simulations of laser bathymetry. Laser bathymetry is a technique to measure water depths using laser pulses emitted from a helicopter or an airplane. In the 1980s the Swedish military had some hope that laser pulses emitted from a helicopter would be useful for detecting objects in the sea below the surface.

The idea behind laser bathymetry is simply to measure the time difference between the epoch of the emitted pulse and the epoch of the peak of the return pulse. The connection

to random flights is that the path of a photon penetrating water can be modelled by a so called anisotropic random flight.

What makes Monte Carlo simulations for laser bathymetry somewhat "tricky" is due to the fact that the energy level of the emitted pulse is many, many times larger then the energy level of the return pulse. Another complication is caused by the backscattering from the sea itself.

In the last part of my talk I plan to present a simple stochastic model for the water motion in a river, a model which is also based on a random flight process.

Some references:

[1] M. Kac, Probability and Related Topics of Physical Sciences, Interscience, NY, 1959, (pages 33-36 in particular),

[2] A.H.Gandjbakhche, R.F. Bonnet and R. Nossal, "Scaling Relationships for Anisotropic Random Walks", J Stat Phys, vol 69, pp 35-53, (1992).

[3] T. Kaijser, "A stochastric model describing the water motion in a river", Nordic Hydrology, vol II, pp 243-265 (1971),

[4] T. Kaijser, "A Monte Carlo simulation model of airborne hydrographic laser systems", FOA report, C 30578-8.1, FOI, Swedish Defence Research Agency, SE-16490 Stockholm.

Jan Kallsen

Christian-Albrechts-Universität zu Kiel

Are American options European after all?

We call a given American option representable if there exists a European claim whose fair price dominates the American payoff at any time and such that the values of the two options coincide in the continuation region of the American option. This concept has interesting implications from a probabilistic, analytic, financial, and numeric point of view. For example, it implies that the American option be hedged statically with European calls/puts of all strikes. Moreover, it allows us to compute the American option price by semi-infinite linear programming. Relying on methods from Jourdain & Martini (2001, 2002), Christensen (2014), and convex duality, we make a step towards verifying representability of American options.

This talk is based on joint work with Matthias Lenga and Sören Christensen.

Ioannis Karatzas

Columbia University

The Harrison-Shepp equation and some of its offspring

In a pioneering article from 1981, Mike Harrison and Larry Shepp provided a stochastic integral equation characterizing the skew Brownian motion of Itô & McKean (1963). We provide similar characterizations for skew-reflected scalar semimartingales, and for a class of planar processes with a roundhouse singularity at the origin which we call "Walsh semimartingales" and which include the Walsh Brownian motion as a special case. Armed with this description, and with an associated calculus we develop, we formulate and solve problems of optimal control with discretionary stopping for such Walsh semimartingales. Joint work with Tomoyuki Ichiba, Vilmos Prokaj and Minghan Yan.

Michael Katehakis
Rutgers University

Reinforcement learning: Connections between MDPs and MAB problems

This talk considers a basic reinforcement learning model dealing with adaptively controlling an unknown Markov Decision Process (MDP), in order to maximize the long term expected average value. In this work, we consider a factored representation of the MDP problem that allows it to be decoupled into a set of individual MAB-problems on a state by state basis. In this way, we show sufficient conditions for efficiently extending classical MAB-type policies to corresponding MDP policies. These constructed MDP policies largely inherit the properties of their MAB-generators, allowing the simple construction of asymptotically optimal MDP policies. We additionally show the construction of a simple UCB-type MDP policy, dramatically simplifying an earlier proof of its optimality. Additional extensions to other MAB policies (e.g., Thompson Sampling) are discussed.

Offer Kella
The Hebrew University of Jerusalem

Some martingales associated with Lévy driven queues

This talk will summarize some continuous time martingales associated with Lévy driven queues. Namely, those that were reported in Kella & Whitt (1992), Asmussen & Kella (2000,2001), Kella & Boxma (2013) and, most recently, Kella & Yor (2017). Various applications will be demonstrated.

Christian Keller
University of Michigan

Path-dependent partial differential equations

Partial differential equations (PDEs) play an important role in mathematical finance and stochastic optimal control. However, non-Markovian problems such as pricing of path-dependent options or control problems involving delays cannot, in general, be treated with standard PDE methods in contrast to the Markovian case. In this talk, I present the theory of path-dependent PDEs as a counterpart to standard PDEs that allows us to deal with those non-Markovian problems in a similar manner. I introduce the notion of viscosity solutions for path-dependent PDEs and point out the difficulties and additional challenges compared to the standard PDE case. In particular, optimal stopping plays a crucial role in our notion and methodology.

This talk is based on joint work with Ibrahim Ekren, Nizar Touzi, and Jianfeng Zhang.

Marek Kimmel
Rice University

Neutral cancer evolution with selective sweep(s): Analysis of site frequency spectra of tumor genomes

This paper concerns extraction of information contained in the genomic sequences of human cancers, which pertains to cancer evolution. This includes history of growth and mutations, among these the effects of genetic drift and selective sweeps. We accomplish this by fitting the genome data available from shotgun sequencing of DNA, by a mathematical model. As it is known, shotgun sequencing scrambles the information contained in haplotypes. Therefore, the information available is limited to the site-specific allele frequency spectra, which provide frequency of reads containing variant (mutated) alleles at the segregating sites. Analysis of such spectra can be obtained, among other, using theoretical expressions that have been developed by mathematicians (Griffiths and Tavaré, 1998), based on the Wright-Fisher coalescent with time-varying population size. In the current analysis, we develop a simple model based on Griffiths-Tavaré theory, with a single selective sweep. Generalization to 2 or more sweeps is straightforward, but it leads to more parameters and therefore to estimation problems. Applying the model to genomic data from The Cancer Genome Atlas (TCGA), we find that for some cancers, the model parameter values are associated with presence or absence of known cancer mutations. We discuss possible ramifications of this finding.

Yerkin Kitapbayev
Boston University

American option pricing under stochastic volatility models via Picard iterations

This paper studies the valuation of American options for a general one- factor stochastic volatility model. Using the local time-space calculus on surfaces we derive an early exercise premium representation for the option price, parametrized by the optimal exercise surface. The exercise surface is the unique solution to an integral equation of Volterra type. The paper proposes a new numerical scheme to solve the integral equation based on the Picard iterations method. The method is flexible and can handle a wide class of non-affine models. Performance is illustrated for the Black-Scholes, Heston and 3/2 models. The approach provides fast convergence, simple implementation and good runtime/RMSE tradeoff and can be extended to other multi-dimensional stopping problems.

Michael Klass
University of California, Berkeley

The S_n/n problem: A paradigm in the development of probability theory

This talk will focus on the ubiquitous and continuing impact of the "S_n/n problem." It has fostered the development of a general theory of optimal stopping and optimal control, extended valued stopping rules, approximation of $\mathbb{E}\left[\Phi\left(||S_n||\right)\right]$, $\mathbb{E}\left[\max_{1\leq n\leq T}\Phi\left(||S_n||\right)\right]$ for stopping times T, and upper quantiles for S_n and $\max_{0\leq i<j\leq n}(S_j - S_i)$, almost sure long run behavior of S_n and more.

Fima Klebaner
Monash University

Persistence of small noise and random initial conditions

The effect of small noise in a smooth dynamical system is negligible on any finite time interval. Here we study situations when it persists on intervals increasing to infinity. Such asymptotic regimes occur when the system starts from initial condition, sufficiently close to an unstable fixed point. In this case, under appropriate scaling, the trajectory converges to solution of the unperturbed system, started from a certain *random* initial condition. The initial condition is of the form $H(W)$, where $H(x)$ is a function obtained from the deterministic equation, and W is a random variable arising in the linear stochastic approximation. This type of result occurs in diffusions, birth-death processes as well as in discrete time dynamics. A special application in discrete time is the model of the Polymerase Chain Reaction (PCR).

Maike Klein
Friedrich-Schiller-Universität Jena

Optimal stopping with expectation constraints

We consider the problem of optimally stopping a continuous-time process with a stopping time satisfying an expectation constraint. By introducing a new state variable, we show that one can transform the problem into an unconstrained control problem. Moreover, we characterize the value function in terms of the dynamic programming equation, which turns out to be an elliptic, fully non-linear partial differential equation of second order. In a second approach we use results on Skorokhod embeddings in order to reduce the stopping problem to a linear optimization problem over a convex set of probability measures.

This talk is based on joint work with Stefan Ankirchner, Nabil Kazi-Tani and Thomas Kruse.

Thomas Kruse
University of Duisburg-Essen

An inverse optimal stopping problem for diffusion processes

Let X be a one-dimensional diffusion and let $g\colon [0,T]\times\mathbb{R}\to\mathbb{R}$ be a payoff function depending on time and the value of X. In this talk we analyze the inverse optimal stopping problem of finding a time-dependent function $\pi : [0,T]\to\mathbb{R}$ such that a given stopping time $\tau^\star$ is a solution of the stopping problem $\sup_\tau \mathbb{E}\left[g(\tau,X_\tau)+\pi(\tau)\right]$. Under regularity and monotonicity conditions, there exists a solution π if and only if $\tau^\star$ is the first time when X exceeds a time-dependent barrier b, i.e., $\tau^\star=\inf\left\{t\geq 0\,|\,X_t\geq b(t)\right\}$. We prove uniqueness of the solution π and derive a closed form representation. The representation is based on an auxiliary process which is a version of the original diffusion X reflected at b towards the continuation region. The results lead to a new integral equation characterizing the stopping boundary b of the stopping problem $\sup_\tau \mathbb{E}\left[g(\tau,X_\tau)\right]$.

Jeffrey Kuan
Columbia University

Markov duality and the KPZ equation

The asymmetric simple exclusion process (ASEP) can be viewed as a discretization of the Kardar-Parisi-Zhang (KPZ) stochastic partial differential equation. The limit from ASEP to KPZ can be viewed through Markov duality. We explain this connection to duality as well as its generalizations.

Hui-Hsiung Kuo
Louisiana State University

Itô formula and near-martingale property for anticipating stochastic integrals

We briefly review a stochastic integral for adapted and instantly independent stochastic processes. Then we show an extension of the Itô formula to this new stochastic integral. We describe an observation leading to the concept of near-martingale property, which is the analogue of the martingale property for the Itô integral. Doob–Meyer decomposition is extended to near-submartingales. We also study exponential processes and stochastic differential equations regarding to this new stochastic integral.

Andreas Kyprianou
University of Bath

Entrance and exit from infinity of stable jump diffusions

In his seminal work from the 1950s, William Feller classified all one-dimensional diffusions on $-\infty \leq a < b \leq \infty$ in terms of their ability to access the boundary (Feller's test for explosions) and to enter the interior from the boundary. Feller's technique is restricted to diffusion processes as the corresponding differential generators allow explicit computations and the use of Hille-Yosida theory. In this talk we study exit and entrance from infinity for the most natural generalization, that is, jump diffusions of the form

$$dZ_t = \sigma(Z_{t-})\, dX_t,$$

driven by stable Lévy processes for $\alpha \in (0, 2)$. Many results have been proved for jump diffusions, employing a variety of techniques developed after Feller's work, but exit and entrance from infinite boundaries has long remained open. We show that the presence of jumps implies features not seen in the diffusive setting without drift. Finite time explosion is possible for $\alpha \in (0, 1)$, whereas entrance from different kinds of infinity is possible for $\alpha \in [1, 2)$. We derive necessary and sufficient conditions on σ so that (i) non-exploding solutions exist and (ii) the corresponding transition semigroup extends to an entrance point at 'infinity'.

Our proofs are based on very recent developments for path transformations of stable processes via the Lamperti–Kiu representation and new Wiener–Hopf factorisations for Lévy processes that lie therein. The arguments draw together original and intricate applications of results using the Riesz–Bogdan–Żak transformation, entrance laws for self-similar Markov processes, perpetual integrals of Lévy processes and fluctuation theory, which have not been used before in the SDE setting, thereby allowing us to employ classical theory such as Hunt–Nagasawa duality and Getoor's characterisation of transience and recurrence.

This is joint work with Leif Döring.

Damien Lamberton
Universite Paris-Est

On American option prices in the Heston model

In this talk, based on joint work with Giulia Terenzi, I will discuss some properties on American option prices in the Heston model. The analysis of the associated variational inequality will be based on the approach used in recent papers by Daskalopoulos and Feehan (2011) and Feehan and Pop (2015). The link between the optimal stopping problem and the variational inequality will be clarified by using the affine property of the model. We will also present some qualitative properties of the value function and of the exercise boundary.

Hans Rudolf Lerche
University of Freiburg

From sequential statistics to optimal stopping and back

This presentation surveys an approach to optimal stopping which originates in sequential statistics. The Bayes risks of several testing problems have a certain structure, which has led to the following main idea. Find a representation of the expected gain (or risk) as

$$R(\tau) = \int g(X_\tau) M_\tau I_{\{\tau<\infty\}} dP,$$

where M_t, $t \geq 0$, is a positive martingale under P of the underlying process X_t, $t \geq 0$, g is a function with unique maximum or minimum and τ is a stopping time. An optimal stopping time stops at argmax (or argmin). In the case of exponentially discounted diffusions the optimal stopping set can be characterized by this approach.

The method can also be extended to stopping games to find the Nash-equilibria.

Finally, for discrete observations the interplay of optimality and overshoot is discussed, mainly in the setting of sequential testing.

This is joint work with M. Beibel, D. Stich, and M. Urusov.

Michel Mandjes
University of Amsterdam

Overdispersion and multi-timescale models: exact asymptotics

Poisson processes are frequently used, e.g. to model the customer arrival process in service systems, or the claim arrival process in insurance models. In various situations, however, the fluctuations in the arrival rate are so severe that the Poisson assumption ceases to hold. In a commonly followed approach to remedy this, the deterministic parameter λ is replaced by a stochastic process $\Lambda(t)$. In this way the arrival process becomes overdispersed.

The first part of this talk considers the case that the Poisson rate is sampled periodically, with a focus on an infinite-server queue fed by the resulting overdispersed arrival process. After having presented a functional central limit theorem, we concentrate on tail probabilities under a particular scaling of the arrival process and the sampling frequency. We derive logarithmic tail asymptotics, and in specific cases even exact tail asymptotics.

In the second part of the talk we embed our overdispersion setting in a more general framework. The probability of interest is expressed in terms of the composition of two Lévy processes, which can alternatively be seen as a Lévy process with random time change. For this two-timescale model we present exact tail asymptotics. The proof relies on an adaptation of classical techniques developed by Bahadur and Rao, in combination with delicate Edgeworth expansion arguments. The resulting asymptotics have a remarkable form, with finitely many sublinear terms in the exponent. Joint work with Mariska Heemskerk and Julia Kuhn.

Randall Martyr
Queen Mary University of London

Nonzero-sum optimal stopping games and generalised Nash equilibrium problems

In this talk I will discuss how questions related to Nash equilibria in nonzero-sum games of optimal stopping (Dynkin games) can be studied using the theory of (deterministic) generalised Nash equilibrium problems (GNEP). The talk focuses on novel equilibria of threshold type in the Dynkin game for Brownian motion. I will also show how the analysis of an equilibrium's stability in the GNEP gives insight into the convergence of a best-response policy iteration scheme commonly used to prove the existence of an equilibrium point in the Dynkin game.

Isaac Meilijson
Tel Aviv University

Variations on a theme of Skorokhod embeddings

Skorokhod (1965) showed that for every distribution with mean zero and finite variance there exist integrable stopping times τ in SBM $B(\cdot)$ such that $B(\tau)$ has the given distribution. The consecutive application of this idea embeds random walks into SBM. Skorokhod then inferred central limit theorems and invariance principles from the SLLN applied to the incremental stopping times.

History: The history of Skorokhod embedding has been wide and productive. Doeblin (1940) and Dubins & Schwarz (1965) extended embeddability to continuous martingales and Monroe (1972) to cádlàg martingales. After briefly describing Skorokhod's original randomized stopping time construction, attention will be focused on the Chacon & Walsh (1976) family of stopping times, of which Dubins (1968) and Azéma & Yor (1978) are special cases.

Billingsley & Ibragimov: Holewijn & Meilijson (1981) extended Skorokhod random walk embedding to martingales with stationary ergodic increments to infer the Billingsley & Ibragimov CLT from the Birkhoff ergodic theorem applied to the incremental stopping times.

Selective risk aversion: Risk averters, with concave utility functions, prefer the mean to the random variable itself. Preference by all risk averters of $X_1 \sim F_1$ to (equal mean) $X_2 \sim F_2$ occurs iff there exist Skorokhod embeddings of F_1 and F_2 in SBM such that $\tau_1 \leq \tau_2$ a.s. A weaker notion of risk aversion (Landsberger & Meilijson 1990) will be shown to be characterized by Skorokhod embedability in the Azéma martingale, to be described.

How variable can a mean-zero martingale with given final variance σ^2 be?: Dubins & Schwarz (1988) identified the maximal expected maximum as σ and the maximal expected maximal absolute value as $\sigma\sqrt{2}$. Dubins, Gilat & Meilijson (2009) identified the maximal expected diameter as $\sigma\sqrt{3}$. Gilat, Meilijson & Sacerdote (2017) identified the maximal expected number of up-crossings of intervals. These martingale maximands are bounded from above by the corresponding variables defined on the background SBM, following Monroe's result that martingales are optional sampling of SBM. Hence, all of these problems have as worst case martingale SBM stopped at suitable stopping times with mean σ^2. These optimal stopping times will be briefly outlined.

The spider problem: This is an offshot of the diameter problem that has eluded attempts by Dubins, Ernst, Gilat, Meilijson, and Shepp to solve explicitly.

Martingale exponential transforms: For a regular enough random walk S_n with positive drift there is (the adjustment coefficient or Aumann-Serrano index) $a > 0$ such that $\exp\{-aS_n\}$ is a martingale. Meilijson (2008) Skorokhod embedded this martingale in the corresponding exponential transform of a suitable Brownian Motion to infer on random walk Lundberg-type approximations and inequalities related to global minimum and drawdown behavior, from the corresponding answers for Brownian Motion.

Finance applications: Various other applications of the adjustment coefficient via Skorokhod embeddings to Finance (under study with Amir and Tuvik Beker, and with M.Sc. student Nitay Alon) will be described.

Ely Merzbach
Bar-Ilan University

A martingale characterization of the multi-parameter fractional Poisson process

We present new properties for the Fractional Poisson process and the multi-parameter fractional Poisson field. A martingale characterization for fractional Poisson processes is given. We extend this result to fractional Poisson fields, obtaining some other characterizations. The fractional differential equations are studied.

Joint work with G. Aletti and N. Leonenko.

Masakiyo Miyazawa
Tokyo University of Science

An alternative approach for diffusion approximation of queueing networks in heavy traffic

Diffusion approximation is a versatile technique in stochastic analysis. We are interested in its application to queueing networks in heavy traffic. When such a network is described by a multidimensional reflecting process, a typical limit under time and state scaling is a semimartingale reflecting Brownian motion on the nonnegative orthant, SRBM for short. To obtain this limiting process, it is usually used a mapping from a net flow process to joint queue length process and joint cumulative idle times of servers, which is called a Skorohod map. When the net flow process is well approximated by a function of Brownian motion, the Skorohod map nicely produces SRBM as the limiting process.

This approach is a bit inconvenient when studying the limit of a sequence of the stationary distributions under the scaling because a Markov process to describe the network is suppressed in its analysis and one needs to verify the interchange of the limits in time and a sequence of models. It also hides the intermediate steps to the limit behind the curtain of the Skorohod map.

We here look at this diffusion approximation problem in a different way. We directly work on the Markov process which may have complicated jumps due to arrivals and service completions. A key idea is to represent its time evolution by a semimartingale with an absolutely continuous bounded variation, suitably choosing an exponential test function. Namely, the complicated jump terms are taken care by a martingale. This enables us to nicely handle the stationary distribution. It also allows the intermediate steps for the approximation to be more visible. For example, in each step, this semimartingale can be used to obtain the tail asymptotics of the stationary distribution through change of measure. Hence, we can see how the tail decay rates change under the limiting operation. It may tell us the quality of the diffusion approximation.

In this talk, we first outline this semimartingale approach using a simple example, then discuss its application to a generalized Jackson network and its multiclass customer extensions. These applications are based on the joint work with Jim Dai and Anton Braverman.

Ilya Molchanov
University of Bern

Non-linear expectations of random sets

We consider sublinear and superlinear expectations defined on random closed sets. These expectations are set-valued and, unlike expectations of random variables, cannot be transformed to each other by altering the sign. Several general construction methods for non-linear expectations are presented. It is explained that sublinear expectations are naturally related to detection of outliers in multivariate samples, while superlinear ones can be used to assess utilities and risks of random vectors.

Joint work with Ignacio Cascos (Madrid) and Anja Muehlemann (Bern)

Ernesto Mordecki
Universidad de la República, Uruguay

On optimal stopping of multidimensional diffusions

We develop an approach for solving perpetual discounted optimal stopping problems for multidimensional diffusions, with special emphasis on the d-dimensional Wiener process. We first obtain some verification theorems for diffusions, based on the Green kernel representation of the value function. Specializing to the multidimensional Wiener process, we apply the Martin boundary theory to obtain a set of tractable integral equations involving only harmonic functions that characterize the stopping region of the problem in the bounded case. The approach is illustrated through the optimal stopping problem of a d-dimensional Wiener process with a positive definite quadratic form reward function.

Joint work with Sören Christensen, Fabián Crocce and Paavo Salminen.

John Moriarty
Queen Mary University of London

A geometric answer to an open question of singular control with stopping

We solve a problem of singular stochastic control with discretionary stopping, suggested as an interesting open problem by Karatzas, Ocone, Wang and Zervos (2000), by providing suitable candidates for the moving boundaries in an unsolved parameter range. We proceed by identifying an optimal stopping problem with similar variational inequalities and inspecting its parameter-dependent geometry, which reveals a discontinuity not previously exploited.

Marcel Nutz
Columbia University

Convergence to the mean field game limit: A case study

We study the convergence of Nash equilibria in a game of optimal stopping. If the mean field game has a unique equilibrium, any sequence of n-player equilibria converges to it as $n \to \infty$. However, both the finite and infinite player versions of the game often admit multiple equilibria. We show that mean field equilibria satisfying a transversality condition are limits of n-player equilibria. But we also find mean field equilibria that are not limits, thus questioning their interpretation as "large n" equilibria.

Richard Olshen
Stanford University

A story, Larry Shepp's diverse interests, and IgG heavy chain V(D)J rearrangements

This talk traces some of the interactions of the speaker with the late Larry Shepp that took place over 50 years. As attendees know, Larry was a person of enormous mathematical/scientific gifts and diverse interests that go well beyond the topics of this meeting. Some of these will be reported, as will be some of Larry's varied research. There will also be a report on quantifying numbers of IgG heavy chain genes of subsets of B cells and T cells of the adaptive human immune system. The latter involves many collaborations, especially with Lu Tian and Yi Liu, but also Andrew Fire, Scott Boyd, and Jorg Goronzy.

Mariana Olvera-Cravioto
University of California, Berkeley

The giant strongly connected component on random digraphs

I will introduce a directed random graph model where the presence of arcs between vertices are independent of each other and are influenced by vertex attributes. Our model includes as special cases the directed versions of the well-known Erdős-Renyi graph, the Chung-Lu model, the Norros-Reittu model, and the generalized random graph. The talk will focus on the phase transition for the existence of a giant strongly connected component, and how the vertex attributes can be chosen to produce graphs with inhomogeneous degrees and highly dependent in-degree and out-degree for each vertex. We also compare the phase transition result with its undirected counterpart, and explain how the direction of the arcs introduces some new interesting phenomena.

Jan Palczewski
University of Leeds

Impulse control of non-uniformly ergodic processes with average cost criterion

We study a problem of impulse control of a general Feller-Markov process that maximises the average cost per unit time criterion:

$$\liminf_{T\to\infty} \frac{1}{T} E^x \Big\{ \int_0^T f(X_s)ds - \sum_{i=1}^{\infty} 1_{\tau_i \le T} c(X_{\tau_i -}, \xi_i) \Big\},$$

where f is a running reward and $c \ge 0$ is an impulse cost. We characterise optimal strategies via a solution to an auxiliary Bellman equation. The novelty of the paper is a general treatment of models in which (X_t) is a process supported on an unbounded space and not uniformly ergodic. We also allow for the cost of an impulse to be unbounded, e.g., proportional to the distance the process is shifted. Our results have applications in balancing of energy systems and in managing inventories. The talk is based on a joint work with Łukasz Stettner.

Zbigniew Palmowski
Wroclaw University of Science and Technology

Double continuation regions for American and swing options with negative discount rate in Lévy models

In this talk we analyze perpetual American call and put options in an exponential Lévy model. We consider a negative effective discount rate which arises in a number of financial applications including stock loans and real options, where the strike price can potentially grow at a higher rate than the original discount factor. We show that in this case a double continuation region arises and we identify the two critical prices. We also generalize this result to multiple stopping problems of swing type, that is, when successive exercise opportunities are separated by i.i.d. random refraction times. We conduct numerical analysis for the Black-Scholes model and the jump-diffusion model with exponentially distributed jumps.

This talk is based on joint work with Marzia De Donno and Joanna Tumilewicz.

Guodong Pang
Pennsylvania State University

Ergodicity properties of Lévy-driven SDEs arising from multiclass many-server queueing networks

We study the ergodic properties of multidimensional piecewise Ornstein-Uhlenbeck processes with jumps, arising in multiclass many-server queues with bursty arrivals and/or asymptotically negligible service interruptions in the Halfin-Whitt regime. The Itô equations have a piecewise linear drift, and are driven by either (1) a Brownian motion and a pure-jump Lévy

process, or (2) an anisotropic Lévy process with independent one-dimensional symmetric alpha-stable components, or (3) an anisotropic Lévy process as in (2) and a pure-jump Lévy process. We identify conditions on the parameters in the drift and the Lévy measure which result in polynomial and/or exponential ergodicity.

Juan Carlos Pardo
Centro de Investigación en Matemáticas A.C.

The excursion measure away from zero for spectrally negative Lévy processes

In this talk, we provide a description of the excursion measure from a point for a spectrally negative Lévy process. The description is based in two main ingredients. The first is building a spectrally negative Lévy process conditioned to avoid zero and the study of its entrance law at zero. The latter is connected with both the excursion measure from zero of the process reflected in its infimum and reflected in its supremum. This leads us to establish a connection between the excursion measure from the state zero and the excursion measure from zero for the process reflected at the infimum and reflected at the supremum, respectively, which is the second main ingredient of our description. Joint work with V. Rivero and J.L. Perez (CIMAT).

Jesper Pedersen
University of Copenhagen

Finite horizon mean-variance optimal stopping problems

Assuming that the underlying process follows a geometric Brownian motion, we study finite horizon nonlinear mean-variance optimal stopping problems. By employing first-order conditions for convex functions we then show that the nonlinear problem can be reduced to a family of linear problems. Solving the latter using a free-boundary approach we find the optimal stopping time. The dynamic formulation of mean-variance problems are also considered.

Shige Peng
Shandong University

Optimal stopping and reflected BSDE with uncertainty of probabilities

Reflected Backward Stochastic Differential Equations (BSDE) is a sharp and powerful tool in various situations to analysis and calculate optimal stopping problems of a nonlinear high dimensional system. Often the environments are so complex such that the probability model uncertainty cannot be negligible. A direct approach is to introduce a G-Brownian motion to robustly model the dynamic model uncertainty of the system.

In this talk we present our recent research results of reflected BSDE driven by G-Brownian motion with which we can treat the optimal stopping problem. We have found a very

interesting formulation to describe such new type of reflected BSDE. Some fundamental results such as existence, uniqueness and comparison theorems, as well as a new type of probabilistic interpretation of the free-boundary problem for fully nonlinear parabolic PDE have been obtained under this framework. Joint work with Hanwu Li.

Goran Peskir
The University of Manchester

Optimal real-time detection of a drifting Brownian coordinate

Consider a three-dimensional Brownian motion X whose two coordinate processes X^i and X^j are standard Brownian motions with zero drift, and the third (unknown) coordinate process X^k is a standard Brownian motion with a non-zero drift μ. Given that only the position of the three-dimensional Brownian motion X is being observed, the problem is to detect, as soon as possible and with minimal probabilities of the wrong terminal decisions, which coordinate process has the non-zero drift. We solve this problem in the Bayesian formulation under any prior probabilities of the non-zero drift being in any of the three coordinates when the passage of time is penalised linearly.

Based on joint work with P.A. Ernst and Q. Zhou.

Huyen Pham
University Paris Diderot

Portfolio diversification and model uncertainty: a robust dynamic mean-variance approach

This talk is concerned with multi-asset mean-variance portfolio selection problem under model uncertainty. We develop a continuous time framework for taking into account ambiguity aversion about combined expected rate of return and correlation matrix of stocks, and for studying the effects on portfolio diversification. We prove a separation principle for the associated robust control problem, which allows to reduce the determination of the optimal dynamic strategy to the parametric computation of the minimal risk premium function. Our results provide a justification for under-diversification, as documented in empirical studies, and that we explicitly quantify in terms of correlation and Sharpe ratio ambiguity parameters. In particular, we show that an investor with a poor confidence in the expected return estimation does not hold any risky asset, and on the other hand, trades only one risky asset when the level of ambiguity on correlation matrix is large. This extends to the continuous-time setting the results obtained by Garlappi, Uppal and Wang [RFS 07], and Liu and Zeng (2017) in a one-period model.

Based on joint work with X. Wei (Paris Diderot) and Chao Zhou (NUS).

Eckhard Platen
University of Technology Sydney

Less expensive investing for the long run

The presentation suggests less expensive production of long term investments that maximize either expected utility from terminal wealth or from consumption. It emphasizes the importance of a tradeable stochastic discount factor (SDF) and the related minimum price to attain target payouts. In this case the SDF equals the inverse of the growth optimal portfolio or numeraire portfolio. An advanced model for the dynamics of well-diversified equity indexes allows accurate production of targeted payouts. Optimal wealth evolution is then closely linked to dynamic versions of Kelly strategies.

Camelia Pop
University of Minnesota

Obstacle problems for nonlocal operators

We prove existence, uniqueness, and regularity of viscosity solutions to the stationary and evolution obstacle problems defined by a class of nonlocal operators that are not stable-like and may have supercritical drift. We give sufficient conditions on the coefficients of the operator to obtain Hölder and Lipschitz continuous solutions. The class of nonlocal operators that we consider include non-Gaussian asset price models widely used in mathematical finance, such as Variance Gamma Processes and Regular Lévy Processes of Exponential type. In this context, the viscosity solutions that we analyze coincide with the prices of perpetual and finite expiry American options. This is joint work with Donatella Danielli and Arshak Petrosyan.

Lea Popovic
Concordia University

Large deviations for two time-scale jump-diffusions and Markov chain models

For a number of processes in biology the appropriate stochastic modelling is done in terms of multi-scale Markov processes with fully dependent slow and fast fluctuating variables. The most common examples of such multi-scale processes are deterministic evolutions, jump-diffusions, and state dependent Markov chains. The law of large numbers limit, central limit theorem, and the corresponding large deviations behaviour of these models are all of interest in applications. In this talk I will give a proof of the large deviations principle for such multi-scale systems, and give an example of an intracellular reaction model on two time-scales for which these results apply.

Ernst Presman
Central Economics and Mathematics Institute

General one-dimensional diffusion and its optimal stopping. Revisited.

The talk is devoted to the general one-dimensional diffusion. We discuss the definition and characterization of such processes: scale, speed measure and killing measure. The generating operator is considered on an extended space of functions (as compared with a classical approach). We give a local characterization of potential functions and excessive functions. For the general one-dimensional diffusion we give necessary and sufficient conditions that the optimal strategy in the optimal stopping problem has a threshold or an island character.

Philip Protter
Columbia University

Semimartingale decompositions under a continuous expansion of the filtration

We are concerned with the role of information plays in various settings, including the financial markets. A few years ago, Younes Kchia and the speaker developed a model for expanding a filtration continuously with new information, for example coming from an external stochastic process. Sufficient conditions were obtained for a semimartingale to remain a semimartingale under the expansion, but the techniques did not supply access to the new resulting decomposition. Within a Brownian paradigm examples showed that the new resultant finite variation terms could or could not remain absolutely continuous with respect to dt.

In new work with Léo Neufcourt we show how one can access the new semimartingale decompositions and determine when the finite variation terms retain the absolute continuity of their paths. Such properties are key to a use of Girsanov's theorem and the absence of arbitrage opportunities in financial markets. Various possible applications are discussed. In honor of Larry Shepp and his strong opinions about fair play, we will indicate how these new techniques give a way to model insider trading and the resulting advantage the insider can gain.

This talk is based on joint work with Léo Neufcourt of Michigan State University.

Kavita Ramanan
Brown University

Scaling limits and local characterization of stochastic dynamics on sparse graphs

Given a sequence of regular graphs G_n whose size goes to infinity, and dynamics that are suitably symmetric, a key question is to understand the limiting dynamics of a typical particle in the system. When each G_n is a clique or its vertices have degrees that tend to infinity, it is known that (under suitable assumptions) the scaling limit is the so-called mean-field limit and the dynamics of a typical particle is governed by a nonlinear Markov process. In this talk, we consider the complementary sparse case when G_n converges in a suitable sense to a countably infinite locally finite graph G, and we describe various limit results, both in the setting of diffusions and Markov chains.

Neofytos Rodosthenous
Queen Mary University of London

Optimal stopping of one-dimensional diffusions with generalised drift

We consider the problem of optimally stopping a one-dimensional diffusion with generalised drift over an infinite horizon. We develop a complete characterisation of the problem's value function and optimal stopping strategy in terms of a variational inequality. We then solve the special case that arises when the state process is a skew geometric Brownian motion and the reward function is the one of financial call option. We show that the optimal stopping strategy can take several qualitatively different forms, depending on parameter values. This is joint work with Mihail Zervos.

Tomasz Rolski
University of Wroclaw

Fluctuation theory for level-dependent Lévy risk processes

A level-dependent Lévy process $U(t)$ is the solution of the sde $dU(t) = dX(t) - \phi(U(t))\,dt$, where X is a spectrally negative Lévy process. We consider the following cases for rate function ϕ, which is always non-negative and non-decreasing. Thus

(i) $\phi(t) = \delta \mathbf{1}_{\{t \geq b\}}$ (then by Loeffen and Kyprianou is said to be a refracted process),

(ii) $\phi_k(t) = \sum_{j=1}^{k} \delta_j \mathbf{1}_{\{t \geq b_j\}}$ (then U is said to be a multi-refracted process),

(iii) ϕ is continuously differentiable.

For all the cases we prove the existence and uniqueness of solutions.

The main work deals with the so called scale functions for process $U(t)$, which are counterparts of the scale functions from the theory of Lévy processes. We show how fluctuation identities for U can be expressed via these scale functions. We also demonstrate that scale functions fulfill an integral equation, in particular the derivatives of the scale functions are solutions of Volterra equations.

References

I. Czarna, José-Luis Pérez, T. Rolski and K. Yamazaki. Fluctuation theory for Level-dependent Lévy risk processes. MS.

A. Kyprianou and R.L. Loeffen, Refracted Lévy processes , *Ann. Inst. Henri PoincarLé Probab. Stat.*, 2010, 24-44.

Alet Roux
University of York

Game options under proportional transaction costs with instantaneous and gradual exercise and cancellation

A game (or Israeli) option is a contract between an option buyer and seller, which allows the buyer the right to exercise the option, and the seller the right to cancel the option at any time up to expiry. The payoff associated with this option is due at the earliest of the exercise and cancellation times. If the option is cancelled before it is exercised, then the buyer also receives additional compensation from the seller.

In this talk, we will consider both the case of instantaneous exercise and cancellation (at ordinary stopping times) and gradual exercise and cancellation (at randomised or mixed stopping times) in a financial model with proportional transaction costs. Efficient constructions for optimal hedging, cancellation and exercise strategies are presented, as well as probabilistic dual representations for bid and ask prices will be presented for both cases.

The probabilistic dual representations involve randomised stopping times, even in the case of instantaneous exercise and cancellation. Moreover, allowing gradual exercise and cancellation leads to increased flexibility in hedging, and hence tighter bounds on the option price as compared to the classic case of instantaneous exercise and cancellation.

Some of the work presented in this talk is joint with Prof. Tomasz Zastawniak.

Marek Rutkowski
The University of Sydney

Nonlinear optimal stopping problem and valuation of American options

Our goal is to reexamine and generalise the findings from recent works by Dumitrescu et al. (2017), Grigorova et al. (2018) and Kim et al. (2018) who have studied optimal stopping problems and their applications to American options using the nonlinear arbitrage-free pricing theory within the framework of financial markets driven by BSDEs. We provide a detailed study of pricing and exercising problems for contracts of American style in a general nonlinear trading model. We first establish a suitable version of the Doob-Meyer-Mertens decomposition extending some results from Peng (1999, 2004). Next, we show that solutions to unilateral valuation problems for the issuer and holder of an American option can be obtained by solving optimal stopping problems associated with nonlinear evaluations corresponding to trading strategies. Joint work with Tianyang Nie and Shige Peng (Shandong University).

Paavo Salminen
Åbo Akademi University

Optimal stopping of strong Markov processes via expected suprema

In this talk we discuss the approach to optimal stopping based on the representation of excessive functions via expected suprema. Let $X = (X_t)_{t\geq 0}$ be a real valued strong Markov process, e.g. a Hunt process, and T an exponential random time independent of X. Our starting point is the result which says that given non-negative and upper semicontinuous function f then

$$u(x) := \mathbb{E}_x\big(\sup_{0\leq t\leq T} f(X_t)\big), \qquad (*)$$

is β-excessive for X. Using this we derive a verification theorem for the OSP

$$V(x) := \sup_{\tau\in\mathcal{M}} \mathbb{E}_x\big(e^{-\beta\tau}G(X_\tau)\big) = \mathbb{E}_x\big(e^{-\beta\tau^*}G(X_{\tau^*})\big), \qquad (**)$$

where $\mathcal{M}$ is the set of stopping times and $G \geq 0$ is the reward function satisfying some regularity conditions. The verification theorem holds, in particular, for diffusions, Lévy processes and continuous time Markov chains. We also discuss another view in this framework which allows us to formulate problem $(**)$ as a problem with monotone structure.

Morover, we compare the present approach with the one based on the Riesz representation of excessive functions. Notice that the representation in $(*)$ is non-linear whereas the classical Riesz representation is linear. Finally, it is seen that the approach is also applicable in optimal control problems for X. We focus, in particular, on impulse control, and – if time allows – on singular control of diffusions.

David Saunders
University of Waterloo

Optimal withdrawal from shared-loss fee structures for hedge funds

Traditional fee structures for hedge funds involve a flat fee expressed as a percentage of assets under management, together with a performance fee that has the structure of a call option. This structure has disadvantages for investors both in terms of expenses, as well as the incentives it provides for hedge fund managers. We will discuss a new fee structure that has been adopted by some funds in the industry, referred to as the shared-loss fee structure. In this framework, in return for receiving upside participation, the fund manager provides some downside protection against losses to the investors. We show that this fee structure can be formulated as a portfolio of options, and discuss the resulting optimal stopping problem of determining the optimal time for an investor to withdraw from the structure.

Andrew Schaefer
Rice University

A mixed-integer programming approach to consensus stopping games

We consider consensus stopping games, a class of stochastic games that requires the consent of all players to terminate the game. We show that each consensus stopping game may have many pure stationary equilibria, which in turn raises the question of equilibrium selection. We develop an efficient algorithm to find a best pure stationary equilibrium of a consensus stopping game. We discuss an application of this class of games in the context of kidney exchanges.

David Scott
Rice University

A minimum distance partial mixture approach to finding clusters

Clustering remains a challenging problem. The goal of finding K clusters and appropriately partitioning the data is both exploratory (graphical) and confirmatory (hypothesis on K). Even if the clusters are known to be multivariate normal and K is assumed known (oracle-style assumption), finding initializations is an almost impossible task as K and p increase.

In this talk, we discuss a minimum distance approach (L2E) for fitting Gaussian mixtures that allows useful fitting to sequentially build up to all K components in some situations. Minimum distance methods are robust, which facilitates its ability to ignore (downweight) groups of data points that the algorithm views as "outliers." L2E has the additional feature that the fitted mixture density does not have to have the weight 1 (unlike MLE). If the normal components are all spherical, then classical algorithms such as K-means can be useful.

These ideas are discussed and illustrated in several dimensions.

Sanjay Shakkottai
The University of Texas at Austin

Regret of queueing bandits

We consider a variant of the multiarmed bandit (MAB) problem where jobs or tasks queue for service, and service rates of different servers (agents) may be unknown. Such (queueing+learning) problems are motivated by a vast range of service systems, including supply and demand in online platforms (e.g., Uber, Lyft, Airbnb, Upwork, etc.), order flow in financial markets (e.g., limit order books), communication systems, and supply chains.

We study algorithms that minimize queue-regret: the expected difference between the queue-lengths (backlogs) obtained by the algorithm, and those obtained by a genie-aided matching algorithm that knows exact service rates. A naive view of this problem would suggest that queue-regret could grow logarithmically: since queue-regret cannot be larger than classical regret, results for the standard MAB problem give algorithms that ensure queue-regret increases no more than logarithmically in time. Our work shows surprisingly

more complex behavior specifically, the optimal queue-regret decreases with time and scales as $O(1/t)$. We next consider holding-cost regret in multi-class (multiple types of tasks) multi-server (servers/agents have task-type dependent service rate) systems. Holding costs correspond to a system where a linear cost (with respect to time spent in the queue) is incurred for each incomplete task. We consider learning-based variants of the c-mu rule a classic and well-studied scheduling policy that is used when server/agent service rates are known. We develop algorithms that result in constant expected holding-cost regret (independent of time). The key insight that allows such a regret bound is that service systems we consider exhibit explore-free learning, where no penalty is (eventually) incurred for exploring and learning server/agent rates. We finally discuss the implications of our results on building platforms for matching tasks to servers/agents. Base on joint work with Subhashini Krishnasamy, Rajat Sen, Ari Arapostathis and Ramesh Johari.

Paul Shaman
The Wharton School, University of Pennsylvania

Properties of the bias mapping of the Yule–Walker estimator

The Yule–Walker estimator of the coefficients of a causal autoregressive process can exhibit substantial bias. This can be understood by noting that the bias mapping produced by the estimator is a contraction. For each order of an autoregressive process and each degree of polynomial time trend for the process mean, iteration of the bias mapping gives convergence to a unique Yule–Walker fixed-point process. The fixed-point processes essentially provide maximal separation of spectral peaks, and there is no estimation bias for them. Yule–Walker estimated coefficients are pulled toward those of a fixed-point process. Ernst and Shaman (2018) have proved that the bias mappings for a second-order autoregressive process are contractions. An induction argument is being explored to provide a proof of this result for orders greater than two. Recursions satisfied by the autoregressive covariances and the inverse of the $p \times p$ covariance matrix of the process, where p is the autoregressive order, are being studied to attempt to complete the induction argument.

Qi-Man Shao
Chinese University of Hong Kong

From limit distributions of self-normalized sums to Cramér type moderate deviations

Let $X_1, X_2, \cdots$ be independent and identically distributed random variables, and let $S_n = \sum_{i=1}^n X_i$ and $V_n^2 = \sum_{i=1}^n X_i^2$ be the partial sum of $\{X_i, 1 \le i \le n\}$ and $\{X_i^2, 1 \le i \le n\}$, respectively. Consider the self-normalized sum S_n/V_n. Logan, Mallows, Rice and Shepp (1973) obtained the limit distribution of self-normalized sums when X_1 is in the domain of attraction of a stable or normal law. Significant progress on the limit theorems for self-normalized sums has been made since then. In this talk we shall give a brief survey on recent developments on limiting distributions and Cramér type moderate deviations for self-normalized sums of independent random variables as well as for general self-normalized processes.

Arnav Sheth
Saint Mary's College of California

Explaining 'irrationality' using 'rational' techniques

The theme of my work with Larry had to do with explaining behavioral ideas using a mathematical framework. For example, why do people take risks when near bankruptcy? What is the impact of politicking on firm value? Why do people spend sub-optimally when they are cash-rich? We solved for the optimal solutions using a numerical technique that included applying Itô calculus, the use of martingale theory, and linear programming. With sparse matrix techniques, a manageable number of constraints and using the MOSEK solver embedded into Matlab, our methodology provides for a quick and easy way by which this class of stochastic control problems can be solved. We provide explicit solutions where it is possible to do so.

David Siegmund
Stanford University

Detection and estimation of local signals

I discuss a general approach to segmentation of independent normal observations according to changes in their mean. The changes may occur continuously, e.g., a change in the slope of a regression line, or discontinuously, e.g., a jump in the level of a process. Theoretical results will be illustrated by simulations and applications to copy number changes and to historical weather patterns. Confidence regions for the change-points and difficulties associated with dependent observations will also be discussed.

Aspects of this research involve collaboration with Fang Xiao, Li Jian, Liu Yi, Nancy Zhang, Benjamin Yakir and Li (Charlie) Xia.

Arno Siri-Jégousse
Universidad Nacional Autónoma de Mexico

Refracted continuous state branching processes

In this talk, we use a Lamperti-type transform to define and study an example of size-dependent continuous state branching process. This process is built from a refracted Lévy process, leaing to a population model with two regimes according to if the size of the population is above or below a certain level. Results on fluctuations of Lévy processes can be adapted to obtain exit times, probability of extinction,... of the branching process.

Michael Sklar
Stanford University

Singular stochastic control via optimal stopping with applications to finance and queueing networks

In a seminal paper, Benes, Shepp, and Witsenhausen (1980) used the connection between singular stochastic control and optimal stopping to obtain explicit solutions for certain stochastic control problems. We review subsequent work that establishes the equivalence between the bounded variation follower problem and optimal stopping associated with a Dynkin game and uses this equivalence to develop backward induction algorithms based solely on optimal stopping. We then extend this idea to develop coupled algorithms when such equivalence does not hold, and describe applications to (a) control of queueing and communications networks, and (b) option pricing in the presence of transaction costs. This is joint work with Tze Leung Lai (Stanford University) and Tiong-Wee Lim (National University of Singapore).

Isaac Sonin
University of North Carolina at Charlotte

Elimination and insertion operations in Markov chains and their application to optimal stopping and other problems

It is well known that a Markov chain (MC) with values in the state space S, observed only when it is outside of a subset $D \subset S$, is again a MC, called the Embedded (Censored) MC. The transition matrix of the embedded MC can be obtained in $|D|$ iterations, each requiring $O(m^2)$ operations, $m = |S|$, when the states from D are "eliminated" one at a time. Recently, it was found that this Elimination operation has an inverse operation, called Insertion, which allows any state, previously eliminated, to be "reinserted" into the state space in one iteration. We will discuss the relationship between these operations and optimal stopping of MC. We also, briefly, describe the applications of these operations to several probability models, and even to some models outside of probability theory.

References

[1] I. Sonin (1999). The state reduction and related algorithms and their applications to the study of Markov chains, graph theory and the optimal stopping problem. *Advances in Mathematics*, **145**, 159-188.

[2] E. Presman, I. Sonin (2010). On optimal stopping of random sequences modulated by Markov chain. *Theory of Probability and Its Applications,* **54**, 3, 534-542.

[3] I. Sonin, C. Steinberg (2016). Continue, quit, restart probability model. *Annals of Operations Research*, **241**, 1-2, 295318.

Siamak Sorooshyari
Johnson & Johnson

A continuous linear-nonlinear model of retinal ganglion cells, and some optimal stopping pitches from neuroscience

The linear-nonlinear (LN) system with a continuous input and output is a common model representing sensory systems including the retina and early visual system, auditory system, and psychophysical behavior. An analysis and expression is presented for mutual information in continuous LN systems with signal-dependent noise with the hope to facilitate the understanding of how the parameters of the LN system influence information transmission. In the second half of the talk several topical problems from behavioral neuroscience that will benefit from an application of optimal stopping are discussed. Background will be provided on three problems that pertain to decision making, inhibitory control, and brain-machine interface (BMI) design.

Lukasz Stettner
Institute of Mathematics, Polish Academy of Sciences

Zero-sum Markov games with impulse controls

In the talk we consider a zero-sum Markov stopping game on a general state space with impulse strategies and infinite time horizon discounted payoff. Under two sets of quite general (weak) assumptions, we show the existence of the value of the game and the form of saddle-point (optimal) strategies. We provide sufficient conditions for some of these assumptions and prove, by a counterexample, the necessity of these assumptions for the uniqueness of the value function. As is also shown, the game can be practically restricted to a sequence of Dynkin's stopping games. Our methodology is somewhat different from the previous techniques used in the existing literature and is based on purely probabilistic arguments. As a consequence, certain interesting properties of Feller-Markov processes are proposed and proved in the process. The talk is based a joint paper with Arnab Basu from Indian Institute of Management Bangalore.

Richard Stockbridge
University of Wisconsin-Milwaukee

A weak convergence approach to inventory control using a long-term average criterion

This talk revisits a stochastic inventory problem in which the inventory is modelled by a diffusion process and a long-term average cost criterion is used to make decisions. The class of such models under consideration have general drift and diffusion coefficients and boundary points that are consistent with the notion that demand should tend to reduce the inventory level. The conditions on the cost functions are greatly relaxed from those in Helmes et al. (2017). Characterization of the cost of a general (s, S) policy as a function of two

variables naturally leads to a nonlinear optimization problem over the ordering levels s and S. Existence of an optimizing pair (s_*, S_*) is established for these models under very weak conditions; non-existence of an optimizing pair is also discussed. Using average expected occupation and ordering measures and weak convergence arguments, weak conditions are given for the optimality of the (s_*, S_*) ordering policy in the general class of admissible policies. The analysis involves an auxiliary function that is globally C^2 and which, together with the infimal cost, solves a particular system of linear equations and inequalities related to but different from the long-term average Hamilton-Jacobi-Bellman equation. This approach provides an analytical solution to the problem rather than a solution involving intricate analysis of the stochastic processes. The range of applicability of these results is illustrated on a drifted Brownian motion inventory model, both unconstrained and reflected, and on a geometric Brownian motion inventory model under two different cost structures.

Philipp Strack
University of California, Berkeley

Optimal stopping under prospect theory preferences

For many decisions in economics and finance, timing is crucial. People decide when to sell an asset, when to exercise an option, when to invest in a project, or when to stop gambling in a casino. We provide a general result on the dynamic decision behavior of a prospect theory agent who is either sophisticated or nave, i.e., unaware of being time-inconsistent.

Dan Stroock
Massachusetts Institute of Technology

A Gaussian version of Cauchy's equation in infinite dimensions

The classical Cauchy equation for a function $f : \mathbb{R}^N \longrightarrow \mathbb{R}$ is

$$f(x+y) = f(x) + f(y) \tag{1}$$

for $x, y \in \mathbb{R}^N$. When f is Borel measurable, it is well known that f satisfies (1) if and only if $f(x) = (x, \xi)_{\mathbb{R}^N}$ for some $\xi \in \mathbb{R}^N$. In fact, if f satisfies (1) for Lebesgue a.e. $(x, y) \in \mathbb{R}^N \times \mathbb{R}^N$, then one can show that there is a $\xi \in \mathbb{R}^N$ such that $f(x) = (x, \xi)$ for a.e. $x \in \mathbb{R}^N$. Because there is no translation invariant measure on an infinite dimensional, real Banach space E, the corresponding result has to be formulated somewhat differently. Namely, assume that $\mathcal{W}$ is a centered Gaussian measure on E and that $f : E \longrightarrow \mathbb{R}$ is a Borel measurable function. Then f satisfies

$$f\left(\frac{x+y}{2^{\frac{1}{2}}}\right) = \frac{f(x) + f(y)}{2^{\frac{1}{2}}} \tag{2}$$

for $\mathcal{W}^2$-a.e. $(x, y) \in E^2$ if and only if there exists an h in the associated Cameron-Martin space such that $f = \mathcal{I}(h)$ $\mathcal{W}$-a.e., where $\mathcal{I}(h)$ is the Paley–Wiener integral h. In particular, if $\mathcal{W}$ is non-degenerate and f satisfies (2), then f is continuous if and only if $f(x) = \langle x, \xi \rangle$ for some $\xi \in E^*$.

Krzysztof Szajowski
Wroclaw University of Science and Technology

On multilateral incomplete information stopping models

The subject of this talk is a model of stopping a multivariate process by the agents having their own payoffs dependent on the state of stopped process. Each are able to declare his suggestion for when the process should be stopped. The players' decision to stop has various effect which depends on the type of the decision makers (players). The ultimate stopping is defined by the simple game over the set of players. Such aggregation of individual decisions or expertise is well known in decision theory. This model was implemented for stopping problems of i.i.d vectors by Yasuda et al. (1982) and extended for the homogeneous Markov sequences by Szajowski & Yasuda (1996). The aggregated individual stopping decisions define the ultimate payoffs. The rational players action should form equilibrium. Another setting can be found when the simultaneous stopping is not allowed (see e.g. Mashiah-Yaakovi (2014)) or the simultaneous stopping is allowed but it does not close the game for all players (Neumann et al. (2002)). Information about the future is revealed by the players stage by stage. Adapting to the new information the players revise their strategies accordingly. The objective is to look for a Nash equilibrium point for the game (cf. Yeung & Petrosyan(2017)). The various setting to multilateral stopping games with information updating and their application (cf. Szajowski(2015)) is compared.

References

Mashiah-Yaakovi, A.: Subgame perfect equilibria in stopping games. *Int. J. Game Theory*, 43(1):89–135, 2014.

Neumann, P., Ramsey, D., Szajowski, K.: Randomized stopping times in Dynkin games. *ZAMM. Z.Angew. Math. Mech.*, 82(11–12):811–819, 2002.

Szajowski, K.: On some distributed disorder detection. In: Steland, Ansgar et al. (eds.) Stochastic Models, Statistics and their Applications. Springer Proc. in Math. & Stat. 122, 187-195 (2015).

Szajowski, K., Yasuda, M.: Voting procedure on stopping games of Markov chain. In: Anthony H. Christer, Shunji Osaki and L.C. Thomas (eds.) UK-Japanese Research Workshop on Stochastic Modelling in Innovative Manufecuring, July 21-22, 1995. Lecture Notes in Economics and Mathematical Systems, vol. 445, pp. 68–80 (1996).

Yasuda, M., Nakagami, J., Kurano, M.: Multi-variate stopping problem with a monotone rule. J. Oper. Res. Soc. Jap. 25, 334–350 (1982).

Yeung, D., Petrosian, O.: Infinite horizon dynamic games: a new approach via information updating. Int. Game Theory Rev. 19(4), 1750026, 23 pp. (2017).

Gábor Székely
National Science Foundation

Independence, uncertainty, and correlation

During our long friendship I discussed many interesting probabilistic topics with Larry Shepp. In this talk I plan to focus on three on them: Independence and atoms (our first manuscript with this suspicious title was confiscated at the Iron curtain), The uncertainty principle of game theory, and the problem of dependence measures.

Topic 1: If the range of an atomic probability measure is the whole $[0, 1]$ interval then does this condition itself guarantees the existence of infinitely many independent events in the probability space?

Topic 2: What is the sharp lower bound of the entropies of minimax strategies of zero sum games in terms of the min–max commutator? How can this lower bound of uncertainty be applied to find simple and approximately optimal strategies?

Topic 3: Ceterum censeo correlationem... as Senator Cato the Elder (234-149 BC) used to say long time before Galton and Pearson.

Denis Talay
Inria

A Wasserstein type distance between diffusion processes and related stochastic control problems

In this joint work with Jocelyne Bion-Nadal (CMAP) we consider the set of the solutions to stochastic differential equations with smooth coefficients. We introduce a Wasserstein type distance which can be represented by means of a stochastic control problem. This representation allows one to get mathematical and numerical estimates for the distance under consideration, and to handle calibration problems in an innovative way.

Peter Tankov
ENSAE ParisTech

An optimal stopping mean-field game of resource sharing

We consider a group of producers sharing a common resource reservoir (such as coal-fired power plants using fresh water from the same river for cooling). Each producer faces a random demand for her product and may withdraw the required quantity of resource if the reservoir level is sufficient. If the reservoir does not allow meeting the total demand of all producers, resource is shared among producers proportionally to their demand levels, and the producers pay a penalty for not meeting their demand. Each producer has the opportunity to switch, at a cost, to a technology not requiring the resource (e.g., build a desalination plant), after which she will always be able to meet the demand. Each producer therefore solves an optimal stopping problem, and the problems of different producers are coupled through the resource reservoir level. Assuming that the number of producers is large, we formulate the

problem of finding a Nash equilibrium as a mean-field game of optimal stopping. Such games have only recently been considered in the literature [1, 2, 3]: building on the results of [1] we prove existence and uniqueness of equilibrium in our setting. We then develop a convergent numerical algorithm for computing the equilibrium and present a numerical illustration for the problem of coal power plants using river water.

This is joint work with Géraldine Bouveret (Oxford) and Roxana Dumitrescu (King's College London)

References: [1] Bertucci, C. Optimal stopping in mean-field games, an obstacle problem approach, J. Math. Pures Appl., in press [2] Carmona, R., Delarue, F., and Lacker, D. (2017). Mean field games of timing and models for bank runs. Applied Mathematics & Optimization, 76(1), 217-260. [3] Nutz, M. A mean field game of optimal stopping. arXiv preprint arXiv:1605.09112, 2016.

Murad Taqqu
Boston University

Lévy driven Ornstein-Uhlenbeck type processes and intermittency

We consider finite variance Lévy driven Ornstein-Uhlenbeck processes. Although these processes provide a rich class of stationary models, their dependence structure is rather limited from a modeling perspective. But superpositions of such Ornstein-Uhlenbeck processes introduced by Barndorff-Nielsen, provide far more flexibility and can exhibit both short-range dependence and long-range dependence. They have found many applications, especially in finance where they are used in models for stochastic volatility. The asymptotic behavior of such processes, integrated and normalized, can be, however, unusual. The cumulants and moments turn out to have an unexpected rate of growth, akin to a physical phenomenon called intermittency. Self-similarity of the limit and intermittency, however, are typically not compatible. We will try to shed some light on this unusual situation. This is joint work with Danijel Grahovac, Nikolai Leonenko and Alla Sikorskii.

Nizar Touzi
École Polytechnique

Infinite horizon Principal-Agent optimal stopping

The Principal-Agent problem is a Stackelberg non-zero sum stochastic differential game which lies at the heart of contract theory in economics. The infinite horizon formulation involves an optimal stopping problem representing the Principal's choice of the Agent's retirement time. We provide a reduction of this problem to a standard stochastic control problem which can be addressed by standard tools. The key argument is a nonlinear representation result by means of a second order backward stochastic differential equation.

Robert Vanderbei
Princeton University

The complex zeros of random sums

In this talk, I will discuss an extension of earlier work with Larry Shepp on the distribution in the complex plane of the roots of random polynomials. In this extension, the random polynomials are generalized to random finite sums of given basis functions. The basis functions are assumed to be entire functions that are real-valued on the real line. The coefficients are assumed to be independent identically distributed Normal (0, 1) random variables. An explicit formula for the density function is given in terms of the set of basis functions. We also consider some practical examples including Fourier series. In some cases, we derive an explicit formula for the limiting density as the number of terms in the sum tends to infinity.

Frederi Viens
Michigan State University

Robust dynamic portfolio selection with model ambiguity, in finance and insurance

Optimal portfolio selection problems with mispricing and model ambiguity can be applied to financial decisions and to the insurance industry. For instance, an investor can hope to exploit temporary mispricings in an immature market with adequate liquidity but where there is little hope of estimating model parameters in a credible way. By adopting an ambiguity-averse framework, one in which the agent's aversion to ill-estimated model parameters can be specified, a strategy can be defined which is highly robust to mis-specifications, but for which acceptable levels of expected future utility can be established, which are optimal in a moderated worst-case-scenario sense. We will present one or more works in which an explicit solution for such a robust optimal strategy, and its value function, are derived. In an insurance context, which is more complex than a quantitative finance context, a strategy includes deciding how much of a company's reserves to invest in financially risky and risk-free assets, and how much insurance risk to seek out or to pass on to a reinsurer. Time permitting, we will address what happens when some of the stochastic model factors, such as financial volatility or stochastic risk premium, are unobserved. And we will discuss the distinction, in the context of quant finance, between a robust optimization problem, which is appropriate for fund managers, and an inventory and impulse control optimization problem, which is appropriate for market makers working at high-frequency with limit orders and market orders. This talk will cover several joint papers with Bo Yi from Shenzhen Venture Capital, Ailing Gu from Guangdong University of Technology, and Baron Law from Agam Capital Management.

Richard Voepel
Rutgers University

Leveraging experimental mathematics on certain stopping problems

In the context of optimal stopping problems, the "goal" of any given problem is typically to maximize some expected payout function associated with the observations at hand. But focusing purely on the expectation of a stopping rule loses a great deal of information concerning the random process involved – what if we desire or require knowledge of the variance of the reward? Or perhaps the skewness, kurtosis, or any higher moment? In this talk we will apply an experimental methodology to the study of the stopping problem of Shepp's Urn in order to demonstrate how to gain access to the full probability distribution associated to play under the expectation maximizing strategy, and then bootstrap our analysis to handle a wider variety of stopping rules for this problem.

Moe Win
Massachusetts Institute of Technology

Network localization and navigation (based on what I learned from Larry A. Shepp)

The availability of positional information is of extreme importance in numerous wireless applications including autonomous driving, assisted living, Internet-of-Things, crowdsensing, medical services, as well as search-and-rescue operations. The coming years will see the emergence of location-aware networks with sub-meter localization accuracy, minimal infrastructure, and high robustness in harsh (GPS challenged) environments. To reach this goal we advocate network localization and navigation, a new paradigm that exploits a combination of wideband transmission and spatiotemporal cooperation. Our work relies on statistical communication theory and has addressed this problem from three perspectives: theoretical framework, cooperative algorithms, and network experimentation. This talk will provide an overview of our recent research results, based on lessons learned from Larry A. Shepp, in this exciting field.

Abraham Wyner
The Wharton School, University of Pennsylvania

When might it be optimal to not stop at all? Explaining the success of AdaBoost and random forests

There have been enormous advances in statistical learning (i.e. "Machine Learning") in the last 20 years. Some of the most successful of these algorithms have been interpreted as statistical optimizations of loss functions on training data, which necessarily leads to a challenging problem of "optimal stoppage" to prevent overfitting. Nevertheless, many of these algorithms work best when allowed to interpolate all the data, which is to say that they fit without error even noisy training data. In this talk, I will explore the behavior of two popular ensemble methods of classification, Boosting and random forests. I will explain, in contrast to the statistical view, why it may be possible to achieve optimal classification for the Adaptive Boosting algorithm "AdaBoost" without stopping at all.

Yizhou Xia
Rice University

Longest consecutive patterns in Markov chains

Consider a discrete-time homogeneous Markov chain with initial state i. We study the distribution of $L(j,n)$, the length of the longest consecutive visits of this chain to state j until time n. We provide an explicit formula for $\mathbb{P}\big(L(j,n) < k\big)$ and offer two limiting theorems for $L(j,n)$. Further, we establish asymptotics for the moment generating function of $L(j,n)$ and provide two large deviation principles for $L(j,n)$.

Minge Xie
Rutgers University

A 'probabilistic' development and interpretation of p-value as strength of evidence measured by confidence distribution

We develop a new framework for hypothesis testing based on confidence distribution, which provides a 'probabilistic' interpretation of p-value as strength of evidence for null hypothesis. Unlike the typical three-step testing approach under the Neyman-Pearson framework (i.e., construct an explicit test statistic, establish or approximate its sampling distribution under the null hypothesis, and then obtain a p-value or a rejection region for making inference), we use the confidence distribution approach to obtain a p-value directly, bypassing the first two steps. The development provides a coherent interpretation of p-value as evidence in support of the null hypothesis, as well as a meaningful measure of degree of such support. Therefore, it places an interpretable meaning for a large p-value, e.g. a p-value of 0.9 has more support than 0.5. It may also provide direct connections and meaningful comparisons between testing results obtained from frequentist and Bayesian paradigms, and a coherent inference outcomes for general scientific pursuits. Numerical examples are used to illustrate the wide applicability and computational feasibility of our approach.

This is joint work with Sifan Liu and Regina Liu at Rutgers University.

Kazutoshi Yamazaki
Kansai University

American options under periodic exercise opportunities

We study a version of the perpetual American call/put option where exercise opportunities arrive only periodically. Focusing on the exponential Lévy models with i.i.d. exponentially-distributed exercise intervals, we show the optimality of a barrier strategy that exercises at the first exercise opportunity at which the asset price is above/below a given barrier. Explicit solutions are obtained for the cases the underlying Lévy process has only one-sided jumps. This is joint work with Jose Luis Perez (CIMAT, Mexico).

Yi-Ching Yao
Institute of Statistical Science Academia Sinica

One-sided solutions for optimal stopping problems with logconcave reward functions

In the literature, the problem of maximizing the expected discounted reward over all stopping rules has been explicitly solved for a number of reward functions under general random walks in discrete time and Lévy processes in continuous time. All such reward functions are continuous, increasing and logconcave while the corresponding optimal stopping times are of threshold type (i.e. the solutions are one-sided). In this paper, we show that all optimal stopping problems with increasing, logconcave and right-continuous reward functions admit one-sided solutions for general random walks and Lévy processes. We also investigate the principle of smooth fit for Lévy processes when the reward function is increasing and logconcave. (This talk is based on joint work with Yi-Shen Lin.)

George Yin
Wayne State University

Dynamic systems with random switching and applications

Many problems arising in applications involve random noise influence as well as interactions of continuous and discrete events. In this talk, we study dynamic systems with random switching. A number of examples in applications are given. Then we take a closer look at switching diffusions. In addition to the state space of the switching being a finite set, we allow the switching process taking values in a countable set and the associate operator being past dependent. We examine recurrence, ergodicity, and stability of the system.

Jiongmin Yong
University of Central Florida

Feynman-Kac formula and beyond

The classical Feynman-Kac formula gives a representation of the solution to the linear parabolic equation in terms of the solution to a corresponding (forward) stochastic differential equation. Backward stochastic differential equation can be used to represent solutions to semilinear/quasilinear parabolic equations. On the other hand, quasilinear parabolic systems can be used to represent adapted solutions to (coupled) forward-backward stochastic differential equations. Further, for backward stochastic Volterra integral equations, the so-called adapted solutions and adapted M-solutions can be represented by the solutions to a new type system of parabolic equations. In this talk, we will survey some relevant results that we have obtained in recent years.

Thaleia Zariphopoulou
The University of Texas at Austin

Mean field and n-agent games for optimal investment under relative performance criteria

In this talk, I will present a family of portfolio management problems under relative performance criteria, for fund managers having CARA or CRRA utilities and trading in a common investment horizon in log-normal markets. Explicit constant equilibrium strategies for both the finite-population games and the corresponding mean field games will be constructed, and it will be shown that are unique in the class of constant equilibria. In the CARA case, competition drives agents to invest more in the risky asset than they would otherwise, while in the CRRA case competitive agents may over- or under-invest, depending on the levels of their risk tolerance.

Cun-Hui Zhang
Rutgers University

A sequential probability ratio test for higher criticism

Sequential probability ratio test (SPRT) is a well known optimal stopping rule for testing two simple statistical hypotheses. We develop a one-sided SPRT for testing multiple null hypotheses with nearly optimal power in detecting the presence of signals which are rare and weak. This makes an interesting connection between test of power one and higher criticism. A connection of one-sided SPRT to controlling the false discovery rate is also discussed. This is joint work with Wenhua Jiang.

Jianfeng Zhang
University of Southern California

Some thoughts about time inconsistent problems

Time inconsistency appears in many practical problems and there are typically two approaches in the literature: the strategy of consistent planning and the strategy of precommitment. The former one is the quite popular game approach, the latter one is static, and the two approaches lead to different value functions and thus mathematically are two different problems. Inspired by the notion of forward utility, we shall propose a dynamic approach which (i) is time consistent in certain sense, and (ii) coincides with the strategy of precommitment at initial time. We shall justify through several examples that our approach is appropriate in many practical situations.

Yiqiang Zhao
Carleton University

Asymptotic independence and tail asymptotics for two-dimensional sticky Brownian motions

In this talk, we consider a two-dimensional time-changed semimartingale reflecting Brownian motion (SRBM), or a sticky Brownian motion.

This type of time-changed SRBM finds applications in many areas including mathematical finance and queueing theory. For example, such a process can be used as a model for a market that experiences a slowdown due to a major event, or as an approximation of queueing systems with exceptional services. The main results reported here include asymptotic independence for the joint extreme value distribution and tail asymptotic properties in the joint distribution for the sticky Brownian motion through applications of the kernel method, extreme value theory and the concept of copula.

This talk is based on the joint work with Dr. Hongshuai Dai.

Mikhail Zhitlukhin
Steklov Mathematical Institute

A Bayesian sequential test for the drift of a fractional Brownian motion

We construct a sequential test for the sign of the drift of a fractional Brownian motion. We work in the Bayesian setting and assume the drift has a prior normal distribution. The problem reduces to an optimal stopping problem for a standard Brownian motion, obtained by a transformation of the observable process. The solution is described as the first exit time from some set, whose boundaries satisfy certain integral equations, which are solved numerically. This is a joint work with Alexey Muravlev.

Quan Zhou
Rice University

When is it best to follow the leader?

I discuss a problem that first interested Professor Larry Shepp during his early days at Bell Labs. An object is hidden in one of N boxes. Initially, the probability that it is in box i is $\pi_i(0)$. You then search in continuous time, observing box J_t at time t, and receiving a signal as you observe: if the box you are observing does not contain the object, your signal is a Brownian motion, but if it does contain the object your signal is a Brownian motion with positive drift μ. It is straightforward to derive the evolution of the posterior distribution $\pi(t)$ for the location of the object. If T denotes the first time that one of the $\pi_j(t)$ reaches a desired threshold $1-\varepsilon$, then the goal is to find a search policy $(J_t)_{t\geq 0}$ which minimizes the mean of T. This problem was studied by Posner and Rumsey (1966) and by Zigangirov (1966), who derive an expression for the mean time of a conjectured optimal policy, which we call *follow the leader* (FTL); at all times, observe the box with the highest posterior

probability. Posner & Rumsey assert without proof that this is optimal, and Zigangirov offers a proof that if the prior distribution is uniform then FTL is optimal. In this paper, we show that if the prior is not uniform, then FTL is *not* always optimal; for uniform prior, the question remains open. Joint work with Philip Ernst and L.C.G. Rogers.

Chao Zhu
University of Wisconsin-Milwaukee

Regime-switching jump-diffusion processes with countable regimes

We consider a class of regime-switching jump-diffusion processes in which the switching component has countably many states or regimes. The corresponding stochastic differential equation in each regime has super linear and non-Lipschitz coefficients. After establishing non-explosion and pathwise uniqueness criteria for such SDEs, we obtain the existence and uniqueness of the regime-switching jump-diffusion process by an interlacing procedure. Then we use the coupling method to study the Feller and strong Feller properties of such processes. Irreducibility and ergodicity will also be investigated.

This is a joint work with Fubao Xi (Beijing Institute of Technology).