Program

# 5th Rutgers Applied Probability Conference

## Data Driven Analytics and Related Topics

**Rutgers Business School - Newark and New Brunswick**
1 Washington Park | Newark, NJ 07102
**November 16, 2018**


In Honor of Professor
Eric Denardo

**BREAKFAST** 7:45-8:30 AM Room **609**

**OPENING REMARKS** 8:30 - 8:33 AM, Room **608**

**SESSION 1. 8:33 - 10:05 AM,** Chair Jian Yang, Rutgers University Room **608**

**Opportunity Cost and Marginal Analysis.**
**Eric Denardo**, Yale University.
**Abstract:**
In this (rather elementary) talk, a classic notion in the economics literature is challenged, and very familiar ideas are used to improve on it.

**Making Supply Chain Transparent for a Better World: Information and Analysis.**
**Christopher Tang**, UCLA Anderson School.
**Abstract:** Companies are gaining more supply chain visibility to reduce their supply chain risks, but few are disclosing what they know with the public. Should a firm disclose its supply chain information to the public? What are the risks and opportunities? I plan to present some recent research and case-based studies to illustrate how supply chain transparency can improve our world: planet, people and profit.

**Equicontinuity Conditions for Markov Decision Processes with Application to Inventory Control.**
**Eugene A. Feinberg**, Stony Brook University.
**Abstract:** This talk discusses the equicontinuity condition and its generalization, the lower semi- equicontinuity condition, for Markov decision processes with average costs per unit time. These conditions imply the validity of average-cost optimality equations, convergence of discounted-cost relative value functions to average-cost value functions, and continuity properties of average-cost value functions. Periodic-review stochastic inventory models typically satisfy these conditions. This implies the validity of average-cost optimality equations for inventory models with average cost criteria and provides useful tools to derive structural properties of average-cost optimal policies.

### Connections of Eric Denardo's Papers to Algorithms for Easy Affine MDPs.

**Matthew J. Sobel**, Case Western Reserve University.

**Abstract:** Eric Denardo's papers are enduring gifts, and his beautiful papers with Uriel Rothblum attracted me to MDPs with affine structure. This talk is prompted by a class of affine Markov-modulated MDPs that have embedded "auxiliary MDPs" that are not true MDPs. An auxiliary MDP can have a negative "transition probability," so the justification for the policy improvement algorithm in Eric Denardo's landmark contraction mappings paper cannot be invoked. Nevertheless, Jie Ning and I justify the application of that algorithm to the auxiliary MDP and, thereby, we solve the original MDP. The result is important because the studied class of affine MDPs have vector-valued continuous states and actions, so heretofore they could not have been solved except via a discretized approximation. Yet a solution of the auxiliary MDP can be obtained easily and constitutes an exact solution of the original MDP.

**COFFEE BREAK** 10:05-10:20 AM Room **609**

**SESSION 2. 10:20 - 11:56 AM,** Chair Endre Boros, Rutgers University Room **608**

### Dueling Bandit Problems.

**Sheldon Mark Ross**, University of Southern California.

**Abstract:** There is a set of n bandits and at every stage 2 of the bandits are chosen to play a game, with the result of a game being learned. We suppose there is a "best" bandit that wins each game it plays with probability at least p > 1/2, with the value of p being unknown. In the "weak regret problem" the objective is to choose bandits to maximize the number of times that one of the competitors is the best bandit. In the "strong regret problem" we suppose that bandit i has unknown value $v_i, i = 1, ..., n$, and that i beats j with probability $v_i/(v_i + v_j)$. One version of strong regret is interested in maximizing the number of times that the contest is between the players with the two largest values. Another version supposes that at any stage, rather than choosing 2 arms to play a game, the decision maker can declare that a particular arm is the best, with the objective of maximizing the number of stages in which the arm with largest value is declared to be best. In the weak regret problem we propose a policy and obtain an analytic bound on the expected number of stages over an infinite time frame that the best arm is not one of the competitors when this policy is employed. In the strong regret problem we propose a Thompson sampling type algorithm.

### Distributionally robust and risk averse multistage stochastic programming.

**Alexander Shapiro**, Georgia Institute of Technology.

**Abstract:** We discuss distributionally robust and risk averse approaches to multistage stochastic programming, and the involved concept of time consistency. It turns out that if the respective risk measures are not strictly monotone, then there may exist infinitely many optimal policies which do not satisfy the dynamic programming equations and are not time consistent. This is in a sharp contrast with the risk neutral formulation where all optimal policies are time consistent. As an example we discuss distributionally robust formulations of the inventory models.

### Decision Making under Uncertainty: A Stochastic Optimal Control Problem.

**Mark S. Squillante**, IBM.

**Abstract:** The combination of data proliferation, advances in statistical and machine learning methods, and growth in computational power creates tremendous opportunities for optimization and optimal control under uncertainty to address fundamental decision-making problems arising in the complex large-scale systems of today. We consider a general approach, building on data and on statistical/machine learning and moving from stochastic models of uncertainty to optimization and/or optimal control over these stochastic models, that provides a mathematical foundation for optimal decisions in complex large-scale systems. As a representative example of our general approach, we consider a stochastic optimal control problem and derive an optimal control policy that renders easily and efficiently implementable algorithms for governing dynamic resource allocations over time. Computational experiments further demonstrate and quantify the significant benefits of our approach over previous work. This research includes joint work with X. Gao, Y. Lu, M. Sharma, J.W. Bosman.

### Online Resource Allocation with Applications to Revenue Management.

**David Simchi-Levi**, Massachusetts Institute of Technology .

**Abstract:** Online resource allocation is a fundamental problem in OR and CS with applications such as offering products to customers, distributing jobs to candidates, assigning advertisers to ad slots, and matching drivers to passengers. These problems can be abstracted as follows: there are fixed resources, each of which can be sold at multiple known prices. These resources must be allocated on-the-fly, without assuming anything about future demand. In this talk we cover the CS and OR literature on the problem and in particular focus on two techniques: exploration and exploitation methods, as well as competitive analysis. In the latter case, we review new algorithms that achieve tight competitive ratios under the integral or asymptotic settings. Our algorithms are simple, intuitive and robust and our competitive ratios are provably optimal, for every possible set of prices. In the former case, we discuss an efficient and effective dynamic pricing algorithm, which builds upon the Thompson sampling algorithm used for multi-armed bandit problems by incorporating inventory constraints into the pricing decisions. The algorithm proves to have both strong theoretical performance guarantees as well as promising numerical performance results when compared to other algorithms developed for the same setting. Finally, we compare the performance of both techniques, exploration and exploitation methods and competitive analysis, with real-world and synthetic data from various retail applications.

**LUNCH** 11:56-1:10PM Room **316**

**SESSION 3. 1:10 - 2:46PM,** Chair Adi-Ben Israel, Rutgers University Room **608**

### Synthesis and Generalization of Structural Results in Inventory Management-A Generalized Convexity Property.

**A Federgruen**, Columbia University.

**Abstract:** We address a general periodic review inventory control model with the simultaneous presence of the following complications: (a) bilateral inventory adjustment options, via procurement orders and salvage sales or returns to the supplier; (b) fixed costs associated with procurement orders and downward inventory adjustments (via salvage sales or returns); and (c) capacity limits associated with upward or downward inventory adjustments. We characterize the optimal adjustment strategy, both for finite and infinite horizon periodic review models, by showing that in each period the inventory position line is to be partitioned into (maximally) five regions. Our results are obtained by identifying a novel generalized convexity property for the value functions, which we refer to as (C1K1, C2K2)-convexity. To our knowledge, we recover most existing structural results for models with exogenous demands as special cases of a unified analysis.

### Distributional Robust Optimization in Machine Learning Using the Wasserstein Ball.

**Sanjay Mehrota**, Northwestern University.

**Abstract:** Distributional robust optimization framework allows us to incorporate partial information on the probability distribution followed by the data. In this talk we show that a scenario-wise decomposition is possible when developing algorithms for solving such problems defined using the Wasserstein ball and features with constraints. Computational results show that the framework typically improves out-of-sample performance of the logistic regression model, when it is possible to improve in- sample performance by using the additional information from the Wasserstein-ball.

### Locks, Bombs and Testing.

**Issac M. Sonin**, UNC Charlotte.

**Abstract:** We present a Defense/Attack resource allocation model, where Defenders have some number of "locks" to protect $n$ vulnerable boxes (sites), and attackers are trying to destroy these boxes, having $m$ "bombs," that can be placed into boxes. Similar models were studied in game theory - (Colonel) Blotto games, but our model has a feature absent in previous literature. Attackers test the vulnerability of all sites before allocating their resources, and these tests are not perfect, i.e., a test can give plus for a box without a lock and minus for a box with a lock. We describe the optimal strategies in two versions of this Locks-Bombs-Testing (LBT) model. The inspiration for this model was the paper by K. Sonin, J. Wilson, A. Wright: Rebel Capacity, Intelligence Gathering, and the Timing of Combat Operations (manuscript, 2018).

### Online Resource Allocation with Limited Flexibility.

**Jiawei Zhang**, New York University.

**Abstract:** We consider a class of online resource allocation problems in which there are several types of resources with limited initial inventory and several demand classes. The resources are flexible in that each type of resources can serve more than one demand class. In this paper, we focus on a special class of structures with limited flexibility, the long chain design, which has been an important concept in the design of sparse flexible processes. We study the long chain design in an online stochastic environment where the requests are drawn repeatedly and independently from a known probability distribution over the different demand classes. Also, the decision on how to address each request must be made immediately upon its arrival. We show the effectiveness of the long chain design in mitigating supply-demand mismatch under a simple

myopic online allocation policy. In particular, we provide an upper bound on the expected total number of lost sales that is irrespective of how large the market size is. This is a joint work with Arash Asadpour and Xuan Wang.

**COFFEE BREAK** 2:46 -3:01 PM Room **609**

**SESSION 4. 3:01 - 4:37PM,** Chair Thomas Lidbetter, Rutgers University Room **608**

### Recent progress on online learning to rank.

**Csaba Szepesvári**, University of Alberta, Canada.

**Abstract:** Online learning to rank is a sequential decision-making problem where in each round the learning agent chooses a list of items and receives feedback in the form of clicks from the user, the goal being to maximize the total expected number of clicks in a given time period. Online learning to rank is challenging because realistic models have many parameters. In this talk I will describe a few of our recent approaches to this problem where a recurring team is to use randomization to eliminate the need to estimate high-dimensional objects. I will conclude by describing a few open problems.

### Surrogate-Based Promising Area Search for Lipschitz Continuous Simulation Optimization.

**Jiaqiao Hu**, Stonybrook University.

**Abstract:** We propose an adaptive search algorithm for solving simulation optimization problems with Lipschitz continuous objective functions. The method combines ideas from several popular strategies in simulation optimization. It employs the shrinking ball method to estimate the performance of sampled solutions and uses the performance estimates to fit a surrogate model that iteratively approximates the response surface of the objective function. The search for improved solutions at each iteration is then based on sampling from a promising region adaptively constructed to contain the point that optimizes the surrogate model. We present the local convergence property of the algorithm and provide numerical examples to illustrate its performance.

### Sample-Based Optimal Pricing.

**Omar Besbes**, Columbia University.

**Abstract:** Pricing is central to many industries and academic disciplines ranging from Operations Research to Computer Science and Economics. In this talk, we study data-driven optimal pricing in low informational environments. We analyze how a decision-maker should price based on a single sample of the willingness-to-pay (WTP) of customers. The decision-maker's objective is to select a general pricing policy with maximum competitive ratio when the WTP distribution is only known to belong to some broad set. We characterize optimal performance across a spectrum of non-parametric families of distributions, lambda-regular distributions, two notable special cases being regular and monotone hazard rate distributions. We develop a general approach to obtain parametric lower bounds on the maximin ratio as well parametric upper bounds. The results have implications on the value of samples in absolute terms but also viz., e.g., increased customer competition in low information environments. (joint work with A. Allouah)

### Stochastic dynamic optimization of service stacking .

**Boris Defourny**, Lehigh University.

**Abstract:** This talk is motivated by the management of electric storage resources performing several functions in different markets. A challenge is to avoid control conflicts when more than one decision maker is given access to the storage capacity. Several questions also arise when the functions have different time scales. In this talk we adopt a continuous-time framework and describe random events by point processes. We discuss the analytics of real-time data in order to calibrate model parameters. The problem of optimizing operations is then formulated as a continuous-time stochastic optimization problem. In our application, the problem has structural properties that allow us to derive an effective numerical solution method and discuss the impact of problem parameters.

**COFFEE BREAK** 4:37 - 4:52 PM Room **609**

**SESSION 5. 4:52 - 6:04PM,** Chair Mert Gurbuzbaalaban, Rutgers University Room **608**

### Pricing Variance Derivatives Using Trees.

**Ionut Florescu**, Stevens Institute of Technology.

**Abstract:** In this talk we will show how to calculate the fair price of variance swaps as well as options written on variance swaps. Due to increased market volatility, these instruments have gained importance over the last 10-15 years. We believe this work is the first to show how to use any approximating tree structure to price these path dependent financial instruments. The approach presented makes it possible to use any underlying model as long as it may be approximated by a tree. We further show how to price VIX derivatives (the most liquid variance related instruments) within our framework. We finally introduce and price a new instrument: the Bermudan variance swaption. Similar with the Bermudan interest swaption we believe this later instrument has the potential to be useful for the financial markets.

### Distributed Inexact Newton-type Pursuit for Non-Convex Sparse Learning.

**Bo Liu**, SJD Finance AI Lab.

**Abstract:** Sparse model learning refers to the methods that look for a trade-off between model fitting to the training data and model sparsity. It has been demonstrated to be an effective way of improving model interpretability, alleviating model overfitting, reducing model space and computational complexity in extensive machine learning problems. With the model sparsity degree as prior knowledge, several greedy algorithms have been recently designed to approximately solve the cardinality constrained minimization problem for sparse model learning. We propose a communication-efficient algorithm that solves the cardinality-constrained sparse model learning problem when training samples are distributed on multiple-machines. The proposed method alternates between inexact optimization of a local training loss and centralized global results aggregation. Theoretical analysis shows that for a general class of convex functions with Lipschitz Hessian, the method converges linearly and the contraction factor scales inversely with the data size. Moreover, the communication complexity required to reach desirable statistical accuracy scales logarithmically with the number of machines for some popular statistical learning models. Numerical results demonstrate the efficiency and accuracy of our method when applied to large-scale sparse learning tasks including deep neural nets pruning.

### Online Balancing of Bias-Variance Tradeoff in Stochastic Estimation.

**Henry Lam**, Columbia University.

**Abstract:** We consider biased stochastic estimators, such as gradient estimators with only noisy black-box simulation oracles. Ensuring good performances of these estimators typically requires selecting procedural parameters to balance bias and variance, which in turn depends on a priori knowledge on the simulation budget and other unknown characteristic quantities of the model. We investigate a scheme that matches, in an online manner, the optimal order of error without specifying the budget in advance. Moreover, we explain and explicitly calibrate the online scheme such that its performance is provably more competitive than the conventional scheme, regardless of the model characteristics, even if the budget is correctly planned. We show how the latter is achieved via a refined bias-variance balancing revealed from the minimax relative performance between the online and the conventional schemes.

**DINNER** 6:10 - 8:30 PM Room **316**

**Program Committee:** Farid Alizadeh, Adi Ben-Israel, Endre Boros Mert Gürbüzbalaban, Douglas Jones, Odysseas A. Kanavetas, Michael N. Katehakis, Thomas Lidbetter, Spiros Papadimitriou, Karti Puranam, Andrzej Ruszczyński, Jim Shi, Flora Spieksma, Jian Yang

**Conference Coordinator:** Ms. Luz Kosar kosar@andromeda.rutgers.edu