



**CanQueue 2026**  
**October 16-17, 2026**

# Land Acknowledgment

The University of Calgary, located in the heart of Southern Alberta, both acknowledges and pays tribute to the traditional territories of the peoples of Treaty 7, which include the Blackfoot Confederacy (comprised of the Siksika, the Piikani, and the Kainai First Nations), the Tsuut'ina First Nation, and the Stoney Nakoda (including Chiniki, Bearspaw, and Goodstoney First Nations). The city of Calgary is also home to the Métis Nation within Alberta (including Nose Hill Métis District 5 and Elbow Métis District 6)

# Welcome Message

Dear Colleagues,

It is our great pleasure to welcome you to CanQueue 2026 in Calgary, Alberta!

On behalf of the organizing committee, we are delighted to bring the queueing community together in Calgary for the first time. CanQueue has become an important forum for researchers working on queueing theory, stochastic processes, service systems, and their many applications, and we are excited to continue this tradition while welcoming you to our city.

This year's program features an excellent collection of presentations spanning both fundamental methodological developments and applications of queueing and stochastic modeling. We hope the diverse range of topics will spark new ideas, generate lively discussions, and create opportunities for collaboration across research areas and institutions. Just as importantly, we hope CanQueue 2026 provides plenty of opportunities to reconnect with colleagues and to meet researchers who are new to the community.

We are particularly pleased to host the conference in Calgary. Located at the foothills of the Canadian Rockies, Calgary is a vibrant and welcoming city with a strong connection to both innovation and the outdoors. We hope you will have an opportunity to experience some of what Calgary has to offer during your visit, whether by exploring the city, enjoying the local food scene, or simply taking in the spectacular views of the Rockies.

We hope you enjoy the scientific program, the conversations, and the Calgary experience. We look forward to an engaging and memorable CanQueue 2026.

Warm regards,  
Marco Bijvank and Yiwen Jin

# Workshop Schedule

October 16, 2026

| Start    | End      | Topic                                                                                                                                 |
|----------|----------|---------------------------------------------------------------------------------------------------------------------------------------|
| 7:45 am  | 8:30 am  | Registration + Breakfast                                                                                                              |
| 8:30 am  | 8:45 am  | Welcome Message                                                                                                                       |
| 8:45 am  | 9:45 am  | Keynote: <i>Information Design in Queueing Systems: From Fluid Limits to Human Limits</i> , Rouba Ibrahim (University College London) |
| 9:45 am  | 11:00 am | Session 1: Priority Mechanisms in Queueing Systems                                                                                    |
| 11:00 am | 11:30 am | Break                                                                                                                                 |
| 11:30 am | 12:45 pm | Session 2: Queueing Models for Healthcare Capacity and Patient Flow                                                                   |
| 12:45 pm | 2:00 pm  | Lunch                                                                                                                                 |
| 2:00 pm  | 3:45 pm  | Session 3: Learning and Optimization in Congested Systems                                                                             |
| 3:45 pm  | 4:15 pm  | Break                                                                                                                                 |
| 4:15 pm  | 5:30 pm  | Session 4: Pricing and Incentives in Service Systems                                                                                  |
| 6:30 pm  | 8:30 pm  | Dinner (ticketed)                                                                                                                     |

October 17, 2026

| Start    | End      | Topic                                                                                                                  |
|----------|----------|------------------------------------------------------------------------------------------------------------------------|
| 7:45 am  | 8:45 am  | Breakfast                                                                                                              |
| 8:45 am  | 9:45 am  | Keynote: <i>Improving the Judiciary: A Queueing Approach</i> , Ramandeep S. Randhawa (USC Marshall School of Business) |
| 9:45 am  | 11:00 am | Session 5: Queueing and Stochastic Models for Operational Systems                                                      |
| 11:00 am | 11:30 am | Break                                                                                                                  |
| 11:30 am | 12:45 pm | Session 6: Queueing Models for Emergency Healthcare Operations                                                         |
| 12:45 pm | 2:00 pm  | Lunch                                                                                                                  |
| 2:00 pm  | 3:45 pm  | Session 7: Sustainable and Resilient Operations Under Uncertainty                                                      |
| 3:45 pm  | 4:15 pm  | Break + Farewell                                                                                                       |

# Keynote Speakers

**Rouba Ibrahim (University College London)**

## **Information Design in Queueing Systems: From Fluid Limits to Human Limits**

We study the use of delay information in queueing systems from two perspectives: a macro system-level view and a micro behavioral view. Our approach utilizes a diverse methodological toolkit, ranging from heavy-traffic limits to controlled behavioral experiments with human subjects. By bridging these methodologies, we highlight robust insights that emerge across different modeling techniques. Specifically, we show that more information is not always better and that vague or binary signaling can paradoxically enhance performance. Finally, we address the gap between rigorous mathematical modeling and the nuances of human decision-making, offering a way forward to integrate both into a unified, behaviorally-aware framework for operational design.

This talk is based on two papers. The complete list of my co-authors is as follows: Philipp Afeche, Junqi Hu, Vahid Sarhangian, Arturo Estrada, and Mirko Kremer.

**Ramandeep S. Randhawa (USC Marshall School of Business)**

## **Improving the Judiciary: A Queueing Approach**

Justice systems are high-stakes, complex operations that shape societies and lives through the rule of law. In this talk, I will discuss open problems where our community is well suited to tackle issues that can improve access, timeliness, and fairness of justice.

To illustrate the potential of such work, I will present my recent work on improving case adjudication delays in the Supreme Court of India. In this work, we model the court system as a two-stage case-management queueing system, and use data-driven calibrated simulations to surface the underlying system inefficiencies. We then propose policy interventions that can successfully tackle judicial delay, reducing delays by 65%.

# Session 1: Priority Mechanisms in Queueing Systems

October 16 @ 9:45 am – 11:00 am

## Insensitivity in Single-server Non-preemptive Priority Queue Wait Times

*David Stanford (Statistical & Actuarial Sciences, Western University), Peter Taylor (School of Mathematics, Melbourne University), Ilze Ziedins (Department of Statistics, University of Auckland)*

Kesten & Runnenberg's pioneering work on priority queue waiting times assumed that when the server was ready to select a customer from a particular priority class, it did so on a first-come, first served (FCFS) basis among customers of that class. Other authors have since considered models with customer selection rules such as last come first served and random order of service. It has always been assumed that the same customer selection rule applies to all priority classes. In systems dealing with diverse health needs, it might be the case that different service disciplines are warranted within the various priority classes. We consider queues where the selection rules may differ from one class to another. We show that for a work-conserving priority queue, a given class's waiting time distribution does not depend upon the customer selection rules in the other classes, assuming the priority structure among the classes is preserved.

## A Two-class Accumulating Priority Queue with Changing

*David Stanford (Western University), Amir Rastpour (Ontario Tech University), Ehram Safari (Ontario Tech University), Salar Ghamat (Wilfrid Laurier University)*

Motivated by operations of a large hospital's operating room, we consider a two-class accumulating priority queue in which the lower priority class accumulates priority at a lower rate until a time threshold is reached. For waits in excess of the threshold, lower priority customers start accumulating at the same rate as the higher priority class. The rationale for such an approach is to provide some relief for the lower priority class for customers who have already waited a substantial amount of time. We determine the Laplace-Stieltjes transform of the waiting time distributions for each of the classes, and evaluate the waiting-time performance through a series of numerical examples.

### When Waiting Changes Your Priority: Queueing Systems with Stochastic Priority Upgrading

*Sina Taheria (Haskayne School of Business, University of Calgary), Marco Bijvank (Haskayne School of Business, University of Calgary)*

In many service systems, a customer's priority is not fixed at arrival but rises the longer they wait. Patients deteriorate, service-level agreements are breached, and backorders approach their due dates. We study a single-server queue with multiple customer classes in which a low-priority customer is upgraded to high priority after a random, class-dependent waiting threshold. Upgrading creates endogenous high-priority arrivals, so protecting delayed customers also crowds the customers who were urgent to begin with. We obtain the full sojourn time distribution for each class and a tractable approximation that keeps the state space manageable. Two findings stand out. The benefit of faster escalation disappears once the upgrade rate reaches the service rate, so escalating faster than the server can work buys nothing. And when escalation capacity is limited, a fairness goal calls for spreading it equally across classes, while a weighted goal calls for concentrating it on a single class.

## Session 2: Queueing Models for Healthcare Capacity and Patient Flow

October 16 @ 11:30 am – 12:45 pm

### ICU/SDU Patient Flow Simulation: Exploring Decision Structures Using a Queuing Game Approach

*Yawo Kobara (University of Windsor), Wenshan Deng (University of Windsor)*

Hospitals increasingly rely on step-down units (SDUs) to alleviate intensive care unit (ICU) congestion, yet the operational impact of transfer decision structures between these units remains poorly understood. This study develops a discrete-event simulation model integrating Poisson arrivals, exponential service times, and clinically grounded acuity transitions to evaluate how centralized (Cooperation) and decentralized (Simultaneous) transfer mechanisms influence patient flow in an ICU–SDU network. Results show that the Simultaneous strategy consistently outperforms the Cooperation strategy under low and moderate demand, yielding near-zero ICU holds, shorter waiting times, and reduced length of stay. In contrast, the Cooperation strategy exhibits persistent congestion, with limited benefit gained from expanding SDU capacity. As system demand increases, performance differences narrow, capacity saturation can overwhelm the advantages of improved decision structure. The findings highlight that decision-making architecture, substantially shapes ICU–SDU flow, suggesting that decentralized or protocol-driven transfer processes can enhance efficiency than capacity expansion alone.

### Applying Data-Driven Queueing Methods for NICU/ICU Bed Capacity Planning

*Ali Baloach (Centre for Health Informatics, University of Calgary), Na Li (Cumming School of Medicine, University of Calgary), Maryam Akbari-Moghaddam (Department of Computing and Software, McMaster University), Douglas G. Down (Department of Computing and Software, McMaster University), Catherine Eastwood (Cumming School of Medicine, University of Calgary), Ayman Abou Mehrem (Cumming School of Medicine, University of Calgary)*

Intensive care (ICU/NICU) capacity planning is complicated by time-varying patient demand, heterogeneous lengths of stay, and limited bed availability. We investigate a data-driven infinite server queueing framework that combines time-varying admission rates with site-specific length of stay distributions to estimate expected occupancy and inform bed capacity planning. Patient stay data were analyzed to characterize admissions and lengths of stay and reproduce model-based occupancy and capacity estimates. We compared four capacity planning strategies representing different tolerances for insufficient capacity: planning around average demand, two quantile based approaches designed to limit capacity exceedance, and planning around maximum expected occupancy. Results demonstrate that time-varying demand produces substantial temporal and site-level variation in expected occupancy and required capacity, which fixed average-demand approaches may not capture. We additionally translated the framework into an artificial intelligence assisted interactive web-based planning tool, enabling users to explore how demand, utilization targets, and capacity strategies influence bed requirements.

### Optimal Policies of an Emergency Department with Two Levels of Servers and Seasonal Flu Epidemic Arrivals

*Khushbu Krupal Vekariya (University of British Columbia), Javad Tavakoli (University of British Columbia)*

Emergency departments (EDs) must balance patient congestion, service quality, and the cost of opening temporary surge capacity. We study an ED with regular and surge beds during an influenza season. A seasonally forced susceptible-infected-recovered (SIR) model is used to represent epidemic-related arrivals and is combined with a regular Poisson arrival stream to create a time-varying arrival rate. The ED is modeled as a transient  $M(t)/M/c/c$  queue and linked to a weekly finite-horizon Markov decision process (MDP) that decides when surge beds should be opened or closed. The model includes activation, operating, stretcher occupancy, and overflow costs. Numerical results show a two-threshold pattern: surge capacity opens when occupancy is high and closes when occupancy falls sufficiently. We also develop a robust MDP using Kullback-Leibler ambiguity sets. Simulation results suggest that moderate robustness can reduce average seasonal cost and 95% CVaR while avoiding too many unnecessary surge openings during the influenza season.

## Session 3: Learning and Optimization in Congested Systems

October 16 @ 2:00 pm – 3:45 pm

### Sample-Augmented Reinforcement Learning for Event-Driven Queueing Control

*Parvin Malekzadeh (Rotman School of Management, University of Toronto), Opher Baron (Rotman School of Management), Dmitry Krass (Rotman School of Management), Simai He (Antai College of Economics and Management, Shanghai Jiao Tong University), Hongsong Yuan (School of Information Management and Engineering, Shanghai University of Finance and Economics)*

The control of jobs in queueing systems is central to applications in healthcare, manufacturing, transportation, and communication networks. Reinforcement learning is promising for such problems but is challenged by stochastic dynamics, unbounded queues, and continuous-time evolution at random event epochs. We develop an event-driven reinforcement learning framework for queueing-control problems with general event-time distributions. Exploiting the conditional structure of queueing dynamics, each observed transition is augmented with synthetic transitions from feasible states, improving state-space coverage and reducing sample complexity. We provide theoretical guarantees quantifying the resulting sample-efficiency gains. Experiments on a single queue with controllable service speed and a two-station tandem queue with shared service capacity demonstrate faster convergence and performance close to dynamic-programming benchmarks when available. The framework requires only samples of the stochastic primitives and applies broadly to event-driven queueing-control problems.

### Learning-Based Approximation in Congested Service Systems

*Borzou Rostami (University of Alberta, Alberta School of Business), Parang Zadtootaghaj (Wilfrid Laurier University)*

We study the assignment of heterogeneous job classes to fixed congested facilities under first-come-first-served (FCFS) and non-preemptive priority-queue (PQ) disciplines. Assignment decisions jointly determine direct assignment costs and facility congestion, making queueing delays endogenous to the optimization problem. We develop a learning-based approximation framework in which parametric models learn queueing behavior from feasible assignment configurations and are embedded directly within the optimization model. Consequently, each candidate assignment generates its own predicted congestion state. We establish performance guarantees linking approximation error to the true optimality loss of the resulting assignment decisions and identify conditions under which such errors become most consequential. Extensive computational experiments separately evaluate predictive accuracy, realized decision quality, and computational efficiency. The results show that the learning-based approach substantially reduces solution times while preserving high-quality decisions, and can outperform time-limited exact optimization on computationally difficult priority-queue instances.

### Predict-then-Optimize for Livestream Shopping under Uncertainty

*Mingyao Song (McMaster University), Kai Huang (McMaster University), Manish Verma (McMaster University)*

Livestream shopping demand is operational decision-sensitive, and interactions across products and employees create a large stochastic combinatorial optimization problem with decision-dependent uncertainty. We propose a scenario-based predict-then-optimize approach integrating machine learning and a two-stage stochastic model for event planning optimization. A scalable and generalizable solution method was developed and demonstrated through case study.

### A Simulation Study of Boost: A Strongly Tail-Optimal Scheduling Policy

*Maryam Elahi (Mount Royal University)*

Scheduling policies determine the order jobs are served in a queue, shaping response time—the time a job spends in the system. For light-tailed job-size distributions, First-Come-First-Served (FCFS) achieves the optimal asymptotic decay rate for the response-time tail, but not the optimal tail constant. Yu and Scully (2024) proposed Boost, a policy that closes this gap and achieves the best possible tail constant. Boost prioritizes jobs using a boosted arrival time that combines arrival time with a size-dependent boost. We use simulation to evaluate Boost beyond its theoretical tail guarantees, comparing its average response time and fairness with FCFS, Processor Sharing (PS), Shortest-Remaining-Processing-Time (SRPT), Fair Sojourn Protocol (FSP), and our proposed FEST policy. We find that Boost improves average response time and fairness over FCFS while retaining its tail advantage. However, size-based policies achieve lower average response times, and Boost remains unfair to small jobs despite its improved fairness over FCFS.

## Session 4: Pricing and Incentives in Service Systems

October 16 @ 4:15 pm – 5:30 pm

### Pricing in Dynamic Matching Platforms with Strategic Agents

*Chiwei Yan (University of California, Berkeley), Tonghun Lee (University of California, Berkeley),  
Matthew Tsao (Stanford University)*

We study pricing policies for dynamic matching platforms with strategic agents. State-dependent pricing can induce agents to delay service requests in anticipation of more favorable prices. Despite this complication, we establish two positive results. First, a pricing policy optimized under the assumption that agents are myopic—requesting service immediately—never backfires when agents behave strategically: it generates weakly greater platform profit than the optimal profit attainable with myopic agents. Second, a simple policy that offers a discount and clears the market once two matchable agents are waiting to request is optimal under strategic behavior. Large-scale ridesharing simulations using Chicago data further confirm that strategic delays can increase platform profits even when they are not explicitly accounted for. Interestingly, platforms may benefit from restricting matching to higher-quality matches than would be optimal under myopic behavior, encouraging agents to wait for more efficient opportunities and thereby increasing market thickness.

### Pricing and Information Design for Nonstationary Congested Service Systems with Two Facilities

*Hamid Arzani (Rotman School of Management, University of Toronto), Philipp Afèche (Rotman School of Management, University of Toronto), Shreyas Sekar (Rotman School of Management, University of Toronto)*

We study how a revenue-maximizing provider should choose static prices and a queue information design for two congestion-prone facilities under time-varying demand that alternates between peak and off-peak. Customers are price- and time-sensitive; they choose a facility or balk to maximize their expected utility. We first study and compare three static, practically relevant regimes: fully observable (FO, both queues revealed), fully unobservable (FU, neither revealed), and partially observable (PO, one queue revealed, the other hidden). We identify conditions for PO to strictly outperform both FO and FU, by enabling the provider to monetize information about the observable queue while diverting overflow demand to the unobservable queue at an average-delay price. We then characterize the optimal (time-dependent) information design. Specifically, we show that the optimal policy couples peak- and off-peak-demand periods through time-dependent join/balk recommendations, thereby shifting congestion toward off-peak-demand periods to reduce expected delays and sustain higher prices.

### Blended Payment in Healthcare: Incentive Issues and Optimal Reforms

*Sara Mohammadi (UCL School of Management), Kenan Arifoglu (UCL School of Management), Bilal Gokpinar (UCL School of Management)*

Blended payment is a mixed prospective reimbursement scheme combining fixed lump-sum payments with variable activity-based payments. Hospitals receive an upfront payment based on planned activity or readiness, followed by a year-end adjustment reflecting deviations from planned activity. We develop a game-theoretic model of hospitals' treatment-cost and planned-activity decisions under stochastic demand and information asymmetry between hospitals and the regulator. We compare blended payment with block contracts and Payment by Results (PbR), and support the theoretical predictions using detailed NHS empirical analysis. Blended payment improves on these alternatives but does not fully align incentives for cost reduction and readiness because the regulator cannot identify the optimal variable rate, break-glass provisions may weaken readiness incentives, and hospitals may face long-run deficits. We therefore propose a modified scheme with a 100% variable rate, peer-based activity benchmarking, and no break-glass provisions, improving incentive alignment, budget balance, and system performance across hospitals facing demand uncertainty.

# Session 5: Queueing and Stochastic Models for Operational Systems

October 17 @ 9:45 am – 11:00 am

## Analysis of Multiclass Queues with Load-Dependent Service Rates

*Berk Gorgulu (DeGroote School of Business, McMaster University), Jing Dong (Columbia Business School, Columbia University), Vahid Sarhangian (Mechanical and Industrial Engineering, University of Toronto)*

We study scheduling in multiclass multi-server queues in which each class's service rate decreases in queue length. In the fluid model, slowdown produces bad equilibria at strictly positive queue lengths: even at loads that would be stable without slowdown, whether the system empties depends on the initial state. We define the good stability region of a policy as the set of initial states whose fluid trajectory converges to the empty state, and the maximum good stability region as its union over admissible policies. We show that emptying the system reduces to reaching a single point in the space of required capacities, where a class's required capacity is the server allocation that holds its fluid queue stationary. For two classes with linear slowdown, we characterize the maximum region explicitly and show it is attained by a myopic policy serving the class with the larger urgency index.

## Transit Delays in a Simple System with Parallel Routes

*Brian Fralix (Clemson University), Mark Holmes (University of Melbourne), Andreas Lopker (Dresden University of Applied Sciences)*

We study a model for delays experienced by traffic along  $L$  lanes of a highway segment, where there is no opportunity to switch between lanes, and where arriving cars are routed to a lane based on a policy that does not depend on the congestion at the time of arrival. The single lane model corresponds to the so-called tollbooth tandem queue introduced in the work of Chao et al. (Journal of Applied Probability, 2015).

We establish structural properties associated with the delay rate in the case of a single lane. These properties allow us to compare the long run delay rates arising from different routing policies in the multi-lane case. In particular we prove that optimal partition routing (where cars are assigned to lanes based entirely on their velocities) yields a smaller long run delay rate than optimal simple random routing.

### A Spatial Queueing and Stochastic Dynamic Programming Model for Single Airtanker Fire Suppression

*Yiming Zhou (University of Windsor), Kevin Granville (University of Windsor)*

We develop a spatial queueing model for dispatching a single airtanker over a directed graph of cells. Fire ignition and growth events in each cell follow Poisson processes with time-varying rates subject to finite buffers, while suppression and travel times are independently exponentially distributed. The system is formulated as a continuous-time Markov chain whose state records fire levels, airtanker's location, and the cell that is the target of the airtanker's current action. A nested block construction defines the infinitesimal generator matrix, and uniformization produces transition probabilities for applying finite-horizon stochastic dynamic programming (SDP) to optimize the service policy. After each observed event (fire ignition or growth, fire suppression, or travel completion), a target matrix identifies feasible post-event states, and the SDP's backward recursion minimizes expected cumulative fire costs and yields time-dependent dispatch decisions for each possible state and event.

## Session 6: Queueing Models for Emergency Healthcare Operations

October 17 @ 11:30 am – 12:45 pm

### Designing Specialist-Response Policies in Hospital Emergency Departments

*Emily Fainman (Texas State University), Beste Kucukyazici (Smith Business School, Queen's University), Zhankun Sun (College of Business, City University of Hong Kong)*

This study develops a systematic specialist-response strategy to reduce delays in specialist consultation (SC) and emergency department (ED) length of stay (LOS) without increasing specialist workload. Using data from over one million ED visits, we model SC requests as a queueing system with time-varying cyclical Poisson arrivals and specialists available at specified times. We analytically determine optimal specialist visit times using the Martingale Representation Theorem and compare this fixed-time policy with alternative response policies. Based on these analyses, we develop guidelines for selecting appropriate response strategies across specialist types and hospital settings. Comprehensive simulation models calibrated with data from two EDs validate the analytical findings. Results show that the proposed strategy can reduce SC waiting times by approximately 25–42% and ED LOS by more than 10%, without increasing the frequency of specialist visits, providing a practical approach to improving ED efficiency and patient flow.

### Brownian Motion Approximation for Developing Benchmark Emergency Department Standing Order Policies

*Masoum Talebi (University of Alberta, Alberta School of Business), Armann Ingolfsson (Alberta School of Business), Saied Samiedaluie (Alberta School of Business)*

Standing orders (SOs) allow triage nurses to initiate tests for selected patients before physician evaluation. Although SO guidelines specify eligible medical conditions, they typically do not indicate when test initiation is operationally desirable. The congestion-based threshold policy proposed by Samiedaluie et al. (2025) addresses this limitation, but its performance under time-varying demand and capacity remains unclear. We develop a rigorous benchmark for evaluating this policy in a non-stationary setting. We model the ED as a multiclass queueing network and derive a heavy-traffic Brownian control problem that incorporates time-dependent system conditions. We then use a neural-network-based computational method to approximate the optimal value function and its gradient, yielding a near-optimal dynamic control policy. Comparing the threshold policy with this benchmark quantifies its optimality gap and determines whether its simple structure remains effective under non-stationary conditions or requires refinement.

### Looking Past the Emergency Department: Ambulance Routing in Hospital Networks with Boarding Considerations

*Amirhossein Shabani (University of Calgary, Haskayne School of Business), Marco Bijvank (University of Calgary, Haskayne School of Business), Mohammad Delasay (Stony Brook University)*

Ambulance routing decisions are often based on emergency department (ED) congestion, but downstream inpatient congestion can reduce the capacity available to new arrivals. We study a two-hospital network in which ambulances are dynamically routed between hospitals, while walk-in patients arrive directly to each ED. Each hospital consists of an ED followed by an inpatient ward, with blocking after service: patients requiring admission remain in ED treatment spaces when no inpatient bed is available. We evaluate a family of shortest-expected-delay routing policies that differ in whether they use nominal or boarding-adjusted ED service capacity and whether they explicitly incorporate boarding pressure. To evaluate these policies, we develop decomposition and coupled continuous-time Markov chain approximations and compare them with existing blocking heuristics and discrete-event simulation. Numerical experiments show that accounting for boarding-induced capacity loss improves routing performance, while the additional value of explicit boarding information depends on congestion and hospital asymmetry.

## Session 7: Sustainable and Resilient Operations Under Uncertainty

October 17 @ 2:00 pm – 3:45 pm

### Sustainable Access to Sanitation Facilities

*Setareh Farajollahzadeh (McGill University)*

Access to sanitation remains limited in developing economies, where constrained public budgets make universal free restroom provision difficult. Private and nonprofit providers can expand capacity, but charging for access may disproportionately benefit users with greater ability to pay. We develop a spatial service-operations model in which users differ in location and price sensitivity and choose among congestible facilities based on travel distance, waiting time, quality, and price. We show that unrestricted private entry can increase total access while leaving the least-affording users no better off and widening disparities across affordability types. In contrast, regulating the location of external capacity can induce interaction with the public facility, diverting users from the free queue, reducing congestion, and expanding access for users who depend on provision. We further characterize how a budget-constrained municipality should allocate investment among expanding existing facilities, constructing new public facilities, and supporting external capacity to improve efficiency and fairness.

### Dynamic Control of Battery Swapping Networks

*Arash Asgari (Rotman School of Management, University of Toronto), Opher Baron (Rotman School of Management, University of Toronto), Nasser Barjesteh (Rotman School of Management, University of Toronto)*

Battery swapping has emerged as a promising alternative to plug-in charging for electric vehicles, offering refueling within minutes and enabling battery-as-a-service business models. However, operating battery swapping networks efficiently is challenging due to stochastic demand, limited charger capacity, fluctuating electricity prices, and the need to maintain high battery availability across locations. We develop a model for a network of battery swapping stations, in which the system manager makes dynamic scheduling (i.e., charging) and pricing decisions to maximize total profit over a finite horizon. Demand is time-varying and depends on station-level prices, inducing substitution across locations; each station has limited charging capacity and a finite inventory of batteries. To obtain tractable policies, we construct a deterministic relaxation that serves as an upper bound on the stochastic problem and use it to design scalable control policies. The results provide operational insights for capacity planning, pricing design, and reliable deployment of swapping infrastructure.

### **A Stochastic MultiCriteria Approach to Incompatible Uncertainty in Canada–Brazil SME Data**

*Renata Pelissari (Telfer School of Management, University of Ottawa), Sharfuddin Ahmed Khan*

This study proposes a Stochastic Multicriteria Acceptability Analysis (SMAA) approach for cross-national comparisons based on statistical surveys that represent uncertainty in incompatible forms. The method is developed in the context of digital technology adoption by SMEs in Canada and Brazil, where estimates are affected by sampling uncertainty, differences in uncertainty reporting, and non-aligned enterprise-size definitions. Each survey estimate is represented probabilistically: Brazilian data use published margins of sampling error, whereas Canadian estimates require modelling coefficient-of-variation bands, introducing nested uncertainty. Differences in enterprise-size definitions are incorporated as an additional source of uncertainty rather than eliminated through deterministic harmonisation. These sources are jointly propagated through a SMAA-based model to classify or rank sectors and size bands according to technology adoption and reported barriers. Instead of deterministic comparisons, the approach produces acceptability indices, allowing identification of robust cross-national conclusions and conclusions that depend on particular uncertainty or alignment assumptions.

### **Warehouse Congestion and Order Fulfillment Performance: An Empirical Analysis of Supply Chain Operations**

*Muhammad Aamir (Virtual University of Pakistan), Naveed Ahmad (Riphah Institute of Hospitality and Culinary Arts, Riphah International University)*

This study examines how warehouse congestion influences order fulfillment performance through order processing delay and whether warehouse capacity flexibility mitigates this effect. Drawing on queueing and supply chain operations perspectives, the study conceptualizes warehouse congestion as accumulated operational pressure arising from high order volumes, limited processing capacity, and workflow interference. Data will be collected from warehouse and supply chain professionals and analyzed using structural equation modeling. The study expects warehouse congestion to increase order processing delays and reduce fulfillment performance, while greater warehouse capacity flexibility is expected to weaken these adverse effects. The findings will offer practical implications for congestion and capacity management.

# Organizing Committee



Dr. Marco Bijvank



Dr. Yiwen Jin