

Program Beta Symposium 2025

Click on a speaker to view more info.

07:00 - 09:30	Breakfast			
09:30 - 09:45	Registration			
09:45 - 10:00	Opening			
10:00 - 10:45	Plenary talk by Hein Fleuren (Tilburg University) Fighting Hunger with Analytics: The Zero Hunger Lab Journey			
10:45 - 11:00	Break			
11:00 - 12:30	Parallel sessions PhDs			
11:00 - 11:30	Nele Amiri KU Leuven Non-Stationary Inventory Control with Lead Times	Shahrezad Fahmy UM Joint Optimization of Fixed and On-Demand Public Transportation	Mostafa Bahadornia UHasselt Storage allocation of perishable products in warehouses	Robert Hanusa UGent A production planning tool accounting for uncertainty in renewable energy forecasts
11:30 - 12:00	Casper Bazelmans TU/e The Full-Truckload Pickup and Delivery Problem with Configurations: Theory & Application	Robbert Bosch UT Applying Surrogate Modelling in Large Scale Multi-Objective Infrastructure Renovation Scheduling	Marieke Brederveld WUR TBA	Efe Senyuva TU/e Spoiling a Surprise to Reduce Spoilage? The Impact of Partial Disclosure of Surprise Bag Contents
12:00 - 12:30	Pieter Deleye UAntwerp Locating decisions in space-based applications: a review	Alim Buğra Çınar VU Pricing, bundling, and driver behavior in crowdsourced delivery	Claire Tan TiU The Impact of Cash and Voucher Assistance on Migrants' Food Consumption Along their Journey	Masoud Fazlavi TiU Trustworthy or Not? Authentication as a Credible Signal in Second-Hand Online Platforms
12:30 - 13:30	Lunch - Lounge KonneKt & Atrium			
13:30 - 14:30	Parallel sessions junior staff			
13:30 - 13:50	Daniel Guzman Vargas UGent Towards a Digital Twin Framework for Integrated Production Planning and Scheduling in Reconfigurable Manufacturing Systems	Mirjam Meijer TU/e Channel Selection in Agricultural Supply Chains: Managing Yield Uncertainty, Aesthetic Requirements, and Producer Information	Amin Asadi UT Recharge and dispatch policies for medical drone delivery	Daniela Guericke UT A hybrid solution approach for the Integrated Healthcare Timetabling Competition

Program Beta Symposium 2025

Click on a speaker to view more info.

13:50 - 14:10	Riccardo Lo Bianco TU/e Automated decision-making for dynamic task assignment at scale	Rodrigo Romero Silva WUR The influence of OR/MS research on policymaking	Laura Poreschack KU Leuven Continuity of Care and Tiered Physician Teams	Bryan Galarza Montenegro UAntwerp SMART-NEMT: Structured Models and Algorithms for Real-Time Non-Emergency Medical Transport.
14:10 - 14:30	Xichen Sun TiU Qualitative meta-analysis with system dynamics modeling for process theory development	Damla Yüksel VU Agricultural Subsidies for Risk-Aware Farmers: Bilevel Optimization for Sustainable Crop Planning	Thomas De Lombaert UHasselt Granting autonomy to warehouse workers - Lessons learned from experimental studies	Yueyi Li TU/e Truck Assignment and Scheduling with Mixed-Service Docks: A Q-learning Enhanced ALNS
14:30 - 15:00	Break			
15:00 - 16:00	Parallel sessions PhDs			
15:00 - 15:30	Florentina Hager UT A Dynamic Hypercube Model Approach for Helicopter Dispatching in Wildfire Management Considering Fire Behaviour	Maike van Roijen WUR Suitable branching strategies for SOS1 constraints without unique ordering	Cheshmeh Chamani UGent Production Planning Optimization in a Symbiotic Network Under Demand Uncertainty	Zhongxin Hu TU/e Socially Responsible Operations for Medical Devices
15:30 - 16:00	Brian Overbeek VU Enhancing flexibility and efficiency in Dutch fire services through innovative vehicle deployment strategies	Tugce Canbilen Suticen TU/e Designing Efficient Distribution Networks for Mixed-cropping Systems	Bilgenur Erdogan TU/e Joint Scheduling of AGVs, AMRs, and Machines with A Matheuristic Approach	Niccolò Maccarini TU/e Controlling inventory of spare parts cooperatively
16:00 - 16:15	Break - Lounge KonneKt & Atrium			
16:15 - 16:45	Plenary talk by Heletjé van Staden (University College Dublin), nominee PhD Award 2023 Beyond the Dissertation: Crafting a Career in Academia			
16:45 - 17:15	Pitches nominees PhD Award 2025			
17:15 - 17:30	Ceremony PhD Award 2025 - presentation by jury chair Ton de Kok			
17:30 - 17:45	Closing			
18:00 - 20:00	Dinner			

Non-Stationary Inventory Control with Lead Times

We consider non-stationary inventory control problems under three canonical settings: (i) demand backlogging, (ii) lost sales with zero lead time, and (iii) lost sales with strictly positive lead time. Specifically, the demand distribution may undergo arbitrary changes at unknown time points. We introduce adaptive online algorithms for each setting and establish their dynamic regret guarantees.

For backlogging systems and zero lead time lost-sales models, our algorithms achieve optimal dynamic regret up to logarithmic factors, without requiring any prior knowledge of the demand distributions or the level of nonstationarity. For lost-sales systems with positive lead time, we derive regret bounds that are near-optimal when the number of distributional changes is known, and sublinear otherwise. Moreover, all proposed algorithms exhibit polynomial dependence on the lead time.

The key driver of the algorithms we present is the deliberate exploitation of convexity and the one-sided feedback structure of the expected asymptotic cost function, which enables counterfactual policy evaluation despite demand censoring under the lost sales model.

Beyond theoretical performance guarantees, our results illustrate how structural insights from inventory theory and online learning can be unified to design adaptive control algorithms for sequential decision-making problems. These advance the development of resilient and efficient supply chain systems capable of adapting to rapidly evolving demand patterns driven by market dynamics and external disruptions.

[Back to program](#)

Joint Optimization of Fixed and On-Demand Public Transportation

What if cities could redesign public transport networks to balance fixed and on-demand bus routes? Traditional fixed-route public transport is efficient in high-demand areas, but becomes less effective when demand is low or variable. Conversely, on-demand services provide flexibility but lack the scalability and cost-efficiency of fixed networks.

This ongoing research addresses that challenge by jointly optimizing both modes within a single, integrated framework. The proposed model integrates the design and operation of fixed-line and on-demand services. The framework combines line planning, frequency setting for fixed lines, hub location selection, and on-demand routing within a unified multilevel heuristic approach.

[Back to program](#)

Storage allocation of perishable products in warehouses

The growing complexity of e-commerce fulfillment has amplified the importance of storage location assignment (SLA) in warehouses, where order-picking is the most resource-intensive activity. While turnover-based, correlated, and scattered storage strategies are well-established, existing approaches typically rely on adopting one strategy in isolation or combine them by using subjective parameter tuning, limiting their adaptability to dynamic operating contexts. For the first part of this PhD, we introduce a data-driven framework that balances these basic SLA strategies according to customer order patterns. First, novel analytical measures are developed to quantify the degree to which each basic SLA principle is realized. Second, a weighting scheme derived from historical order data is proposed to enable context-specific adaptation without decision-maker intervention. These measures and dynamic weights are then integrated into a new multi-objective mathematical model that assigns items to storage locations, taking existing inventory into account. The model is validated in an autonomous mobile robot-assisted order picking system using a factorial experiment across multiple operating contexts. Results demonstrate that the proposed balanced SLA approach significantly outperforms basic strategies by reducing picker travel distance and mitigating order-line splitting.

For the second part, we aim to enrich our study on SLA with a specific business context, that is the SLA of fruits and vegetables in a retail setting. This selection is based on three reasons: (a) retail and consumption have the largest share of food waste; (b) food retailers are the largest commercial users of refrigeration, accounting for 30% of the food sector's electricity consumption; and (c) lower storage temperatures often extend shelf life but at the expense of higher energy consumption. Therefore, our goal is to propose an SLA approach that not only reduces food loss and energy consumption in a retailer's warehouse, but also mitigates food loss further along the supply chain.

[Back to program](#)

A production planning tool accounting for uncertainty in renewable energy forecasts

In academia, we create models that can aid in solving many complex problems that arise in industry. However, it is not always straightforward to convince industrial users to use our models. In this presentation, I first discuss a model that optimizes the energy planning of an electrically-powered chemical reactor. Then, I show how we can present the model in an appealing way to industrial users by making a dashboard that (1) visualizes the model in a way that is interpretable, and (2) permits flexibility by allowing manual adjustments to the output. Presenting complex models in this way can make it easier for a skeptical industrial user to implement them at their facility.

[Back to program](#)

The Full-Truckload Pickup and Delivery Problem with Configurations: Theory & Application

Full-truckload (FTL) transportation is crucial for modern supply chains but is a major contributor to greenhouse gas emissions. Long combination vehicles (LCVs) offer an opportunity to consolidate shipments and reduce costs and emissions, yet their effective use requires complex planning due to the variety of vehicle configurations and the ability to change configurations at equipment yards. We introduce the Full-Truckload Pickup and Delivery Problem with Configurations (FTL-PDPC), a real-world pickup-and-delivery problem that explicitly models vehicle configurations. We present compact and extended mathematical formulations and develop an adaptive large neighbourhood search (ALNS) with an efficient data structure allowing for efficient best-cost insertions. Moreover, the ALNS searches for solutions in a smart parallel framework that explores both the solution space with as well as without yards effectively. Computational experiments on benchmark instances and real-world data from a European FTL carrier demonstrate the applicability and competitiveness of our approach. Our ALNS matches or improves upon state-of-the-art heuristics, and reveals 11.54% emissions can be saved in practice.

[Back to program](#)

Applying Surrogate Modelling in Large Scale Multi-Objective Infrastructure Renovation Scheduling

This paper addresses the computational challenges of solving large-scale Road Network Maintenance Scheduling Problems with Uncertain Deadlines (RNMSP-UD) by introducing machine-learning based pruning methods integrated within the Non-Dominated Sorting Genetic Algorithm II (NSGA-II). The proposed approach, termed Progressive Lower Bound Evaluation (LBPE), uses machine learning models to quickly estimate traffic congestion impacts and eliminate suboptimal schedules without running expensive simulations, thereby dramatically reducing computational requirements while preserving solution quality. Our experimental evaluation on the Sioux Falls network with 76 renovation projects, significantly larger than the typical 5-33 projects found in existing literature, demonstrates that the proposed PLBE methods achieve a 10x improvement in computational efficiency while maintaining solution quality across multiple performance metrics.

[Back to program](#)

TBA

TBA

Back to program

Spoiling a Surprise to Reduce Spoilage? The Impact of Partial Disclosure of Surprise Bag Contents

Surprise bags, that are opaque bundles of perishable products sold at a fraction of their regular price, are widely used by retailers to generate revenue from near-expired items and reduce spoilage. Traditionally, retailers only disclose the discounted price and the original retail value of these bags, leaving consumers unaware of the specific contents. The composition of surprise bags may significantly influence consumer satisfaction. A well-balanced bag with complementary items may be preferred over a higher-value bag that contains a single product or substitutes. Dissatisfied consumers may ultimately discard undesired items, transferring store waste to households and reducing long-term customer retention. To maintain sustainable revenue streams and effectively minimize waste, retailers must ensure consumer satisfaction. If the consumer does not make a completely blind choice, and has an idea about the contents of the bag, the informed choice will lead to a higher satisfaction and a better match between the consumer and the retailer. This study explores the impact of disclosing partial information about the contents of surprise bags, thereby improving the match between consumer preferences and available products. We compare retailer profitability and waste reduction of the effects of various partial information disclosure models, and quantify the upsides and downsides of the information disclosure.

[**Back to program**](#)

Locating decisions in space-based applications: a review

Over the past two decades, space-based applications have become deeply integrated into our lives, amongst others enabling global navigation, communication, and Earth observation. To provide these capabilities the number of satellites in orbit has increased tremendously causing increased risks of collisions and failures. In response, commercial ventures have recently started offering space-based services, including satellite refuelling or repositioning. With the emergence of On-Orbit satellite Servicing (OOS), fundamental operational questions arise foremost where to locate assets. These assets range from single servicing spacecraft, to depots with shuttles for refuelling or repair. While location theory has long been studied in terrestrial contexts, the space environment has its own, unique challenges: determining a position in space involves six degrees of freedom, compared to just two on Earth.

This literature review examines locating decisions for space-based applications from an Operations Research perspective. A systematic search identified 188 relevant publications, which have been categorized into three domains, reflecting increasing complexity from individual satellites to large-scale networks, and from Earth-based clients to spacebased interactions: single satellite orbit selection, constellation design, and on-orbit services (OOS).

The aim of this study is to unravel how classical Operation Research problems - such as coverage, n-cover, p-median, facility location and location-routing - are adapted to the spatial domain. Moreover, we discuss modelling assumptions, constraints, and solution techniques used to address the high-dimensional nature of space-based location problems. As such, this study provides a comprehensive overview of the current approaches and identifies promising avenues for future research in this rapidly evolving field.

[Back to program](#)

Pricing, bundling, and driver behavior in crowdsourced delivery

Challenges in last-mile delivery have encouraged innovative solutions like crowdsourced delivery, where online platforms leverage the services of drivers who occasionally perform delivery tasks for compensation. A key challenge is that occasional drivers' acceptance behavior towards offered tasks is uncertain and influenced by task properties and compensation. The current literature lacks formulations that fully address this challenge. Hence, we formulate an integrated problem that maximizes total expected cost savings by offering task bundles to occasional drivers. To this end, we simultaneously determine the optimal bundle set, their assignment to occasional drivers, and compensations for each pair while considering acceptance probabilities, which are captured via generic logistic functions. The vast number of potential bundles, combined with incorporating acceptance probabilities leads to a mixed-integer nonlinear program (MINLP) with exponentially many variables. Using mild assumptions, we address these complexities by exploiting properties of the problem, leading to an exact linearization of the MINLP which we solve via a tailored exact column generation algorithm. Our algorithm uses a variant of the elementary shortest path problem with resource constraints (ESPPRC) that features a non-linear and non-additive objective function as its subproblem, for which we develop tailored dominance and pruning strategies. We introduce several heuristic and exact variants and perform an extensive set of experiments evaluating the algorithm performances and solution structures. The results demonstrate the efficiency of the algorithms for instances with up to 120 tasks and 60 drivers and highlight the advantages of integrated decision-making over sequential approaches. The sensitivity analysis indicates that compensation is the most influential factor in shaping the bundle structure.

[Back to program](#)

The Impact of Cash and Voucher Assistance on Migrants' Food Consumption Along their Journey

We use retail sales data combined with individual micro-data from a Cash Voucher Assistance (CVA) for Venezuelan transit migrants throughout Colombia to study the effect of cash transfer timing on calorie and diet diversity purchased. Exploiting two exogenous policy changes throughout the 2.5 year program period, we find that transit migrants' travel trajectories strongly shape food consumption derived from cash transfers. Receiving assistance further away from one's destination is associated with lower calorie and diet diversity purchased, but this effect becomes non-significant once cash transfer value is substantially increased (by 36%), and when a major non-food expenditure such as transportation is provided in-kind. We interpret these findings as evidence that longer budget planning horizons, defined distance left to travel, intensify budget constraints and cause trade-offs between food and other needs. Our analysis illustrates the importance of demand-side considerations when planning cash transfer programs, and suggests that humanitarian agencies may benefit from tailoring transfer amounts or bundling in-kind support to transit migrants' stage of travel.

[**Back to program**](#)

Trustworthy or Not? Authentication as a Credible Signal in Second-Hand Online Platforms

This paper investigates the role of an authentication service in mitigating uncertainty caused by untrustworthy descriptions of second-hand products on online platforms with asymmetric information. In our setting, the platform offers authentication for a fee to both sellers and consumers. If either party purchases the service, the product is inspected and delivered only after successful verification; otherwise, the transaction is canceled. We study two questions: Can consumers infer seller trustworthiness from the seller's authentication decision and pricing? How does offering authentication as an optional service affect seller profit, consumer surplus, and total demand compared with a mandatory policy?

Methodology/Results: We model the interactions as a signaling game in which sellers have private information about product actual condition and consumers update beliefs based on price and authentication inclusion. We show that authentication can serve as a credible signal of trustworthiness for high-type sellers. When the authentication fee is low, high-type sellers signal through lower prices, while at moderate fee ranges they signal by purchasing authentication. Optional authentication benefits high-type sellers more than mandatory authentication, whether separation occurs through authentication or price. Low-type sellers fare worse under optional authentication when they purchase the service, but better when they opt out. Relative to mandatory authentication, the optional mechanism also improves consumer surplus and total demand.

Managerial Implications: Our results suggest that platform managers should consider offering authentication in an optional format. Doing so can enhance high-type seller profits, increase consumer surplus, and strengthen overall platform attractiveness.

[**Back to program**](#)

Towards a Digital Twin Framework for Integrated Production Planning and Scheduling in Reconfigurable Manufacturing Systems

Modern manufacturing is increasingly challenged by global competition, market volatility, and the demand for customized products. Reconfigurable Manufacturing Systems (RMSs) offer the flexibility needed to address these pressures, but their complexity requires integrated and responsive planning—something traditional methods often fail to deliver within the fast-paced context of Industry 4.0. This work presents a Digital Twin (DT) framework designed to support real-time decision-making in RMS environments. At its core is a novel Responsive Decision-Making Support (RDMS) system that efficiently solves the Integrated Production Planning and Scheduling (IPPS) problem in industrial RMS settings. Instead of relying on slow conventional optimization, the RDMS uses fast-to-evaluate surrogate models to predict the performance of candidate plans and schedules. This enables rapid scenario analysis and the identification of high-quality, integrated solutions that enhance system agility and responsiveness. The proposed framework bridges the gap between advanced optimization theory and the operational demands of smart factories, offering a scalable pathway toward more agile, efficient, and resilient manufacturing operations in the Industry 4.0 era

[Back to program](#)

Channel Selection in Agricultural Supply Chains: Managing Yield Uncertainty, Aesthetic Requirements, and Producer Information

Direct-to-consumer (D2C) channels, where farmers sell directly to end-consumers, are gaining attention in agri-food supply chains due to increasing consumer interest in product origin and sustainability. However, farmers face complex decisions when choosing between D2C and traditional retail channels, especially considering yield uncertainty and esthetical requirements of retailers. This study analyzes farmers' channel choices and pricing decisions, considering demand uncertainty and yield uncertainty, regarding both quantity and esthetics. We explore the viability of a dual-channel strategy, selling through both retail and D2C channels, and its impact on contractual agreements with retailers. We develop a game-theoretical modeling framework to examine the optimal channel choice of a farmer and, if applicable, the associated wholesale-price contract between farmer and retailer. We develop conditions under which a dual channel choice is viable, highlighting the importance of using the opportunity in the D2C channel to share additional information on the product and its background. Finally, we investigate the effect of yield uncertainty on supply chain profitability and food waste.

[Back to program](#)

Recharge and dispatch policies for medical drone delivery

Drones are increasingly used in healthcare logistics to deliver medical supplies to remote or congested areas that are difficult to access through traditional transportation. However, managing a drone fleet under stochastic demand for urgent items (e.g., blood units) is challenging due to trade-offs between charging mode choices, drone availability, and battery lifespan. Fast charging enables rapid redeployment but accelerates battery degradation and increases replacement costs, whereas slow charging prolongs battery life but limits responsiveness. To address these challenges, we propose a decentralized stochastic multi-hospital allocation and recharging model for blood transportation that captures realistic charging-discharging dynamics, spatiotemporal stochastic demand, and charging trade-offs. The system is formulated as a Markov Decision Process (MDP) that tracks drone charge states, mission durations, and resource availability. A multi-agent deep reinforcement learning (DRL) framework is then developed to derive adaptive dispatching and recharging policies. Experiments show that our decentralized approach improves cost efficiency, fleet availability, and demand responsiveness relative to benchmark methods. Overall, the results highlight the importance of integrated recharging and dispatching management for enhancing both operational performance (via higher demand satisfaction and equitable resource allocation) and sustainability (via smarter charging strategies) in medical drone delivery.

[Back to program](#)

A hybrid solution approach for the Integrated Healthcare Timetabling Competition

The Integrated Healthcare Timetabling Competition (IHTC) 2024 posed a planning problem relevant to the hospital sector. The IHTC considers four coupled scheduling problems in an integrated problem setting, since patient flow through the hospital creates dependencies between them. The scheduling problems are: 1) deciding the patients' admission day, 2) scheduling the patients' surgeries in operating theatres, 3) assigning patients to rooms, and 4) assigning nurses to rooms to care for the patients. These decisions must satisfy several hard and soft constraints. The objective function of the IHTC consists of weighted penalties for the eight soft constraints. Since mixed-integer programming models for the entire problem are prohibitively large, we propose a hybrid solution approach utilizing mixed-integer programming, constraint programming, and simulated annealing. Our approach decomposes the planning into different phases and explores feasible solutions from previous phases in parallel to find solutions for the overall planning problem. During the presentation, we will discuss the relevant models and design decisions, as well as results from the competition instances and the insights we gained into future research directions while solving them. Our approach, which proved robust across different instance settings, ultimately placed third among the finalists of the competition. This is joint work with Rolf van der Hulst, Asal Karimpour, Ieke Schrader and Matthias Walter.

[Back to program](#)

Automated decision-making for dynamic task assignment at scale

The Dynamic Task Assignment Problem (DTAP) concerns matching resources to tasks in real time while minimizing some objectives, like resource costs or task cycle time. In this talk, we consider a DTAP variant where every task is a case composed of a stochastic sequence of activities. The DTAP, in this case, involves the decision of which employee to assign to which activity to process requests as quickly as possible. In recent years, Deep Reinforcement Learning (DRL) has emerged as a promising tool for tackling this DTAP variant, but most research is limited to solving small-scale, synthetic problems, neglecting the challenges posed by real-world use cases. To bridge this gap, we propose a DRL-based Decision Support System (DSS) for real-world scale DTAPS.

The proposed DRL agent has two novel elements: a graph structure for observations and actions that can effectively represent any DTAP and a reward function that is provably equivalent to the objective of minimizing the average cycle time of tasks. The combination of these two novelties allows the agent to learn effective and generalizable assignment policies for real-world scale DTAPs.

The proposed DSS is evaluated on five DTAP instances whose parameters are extracted from real-world logs through process mining, matching or outperforming the best alternative in all DTAP instances and generalizing on different time horizons and across instances.

[Back to program](#)

The influence of OR/MS research on policymaking

Operations Research and Management Science (OR/MS) has long debated the existence of a theory-practice gap, where academic research is perceived as disconnected from the needs of practitioners. While this gap is difficult to trace in industrial contexts, the public policy domain offers a unique opportunity to measure research impact through formal citations. This study provides the first large-scale empirical investigation of the knowledge pathways connecting core OR/MS research to policymaking, by analyzing a comprehensive dataset of articles from core OR/MS journals linked to citing policy documents. We conceptualize these pathways through three key elements: the topic of the research, the academic institutions that produce it, and the policy institutions that cite it. Our results identify specific research topics and approaches (e.g., data envelopment analysis, innovation, organizational studies), journals (e.g., Management Science), and universities that serve as the most effective bridges between academia and policy, demonstrating that OR/MS research successfully bridges the gap when it aligns with pressing societal needs. Finally, our work also identifies significant opportunities for future policy impact in domains such as crime, housing, and immigration.

[**Back to program**](#)

Continuity of Care and Tiered Physician Teams

Context: Continuity of care - sustained therapeutic relationships between patients and physicians - is fundamental to effective primary care but increasingly difficult to maintain due to workforce fragmentation and part-time employment. In this talk, I will introduce a tiered physician team framework where each patient is assigned both a primary physician (highest consultation frequency) and secondary physician (second-highest frequency), creating accountability hierarchies that preserve continuity benefits while accommodating workforce constraints. Our theoretical contribution develops a multi-level familiarity framework distinguishing between patient-physician familiarity (consistently beneficial) and physician-physician familiarity (potentially counterproductive through accountability erosion).

Main Findings: Using 11 years of consultation-level data from 381 English primary care practices, we operationalize clinical productivity as inter-arrival times between consultations. Secondary physicians preserve substantial continuity benefits: approximately 30-40% of primary physician effects. Time until patients' next appointments extends by 12.9% when seen by primary physicians and 4.5% by secondary physicians, relative to other practice physicians. Both effects increase with patient-physician familiarity and are amplified for older patients and those with chronic conditions. Physician-physician familiarity exhibits an inverted U-shaped relationship with team productivity, with optimal performance when physicians share 10-20% of consultations, followed by decline as excessive familiarity breeds accountability erosion. A counterfactual analysis shows that implementing tiered teams increases consultations with familiar physicians from 59% to 78% and eliminates 1.84% of follow-up appointments, representing annual NHS savings exceeding £65 million.

[Back to program](#)

SMART-NEMT: Structured Models and Algorithms for Real-Time Non-Emergency Medical Transport.

Operations Research and Management Science (OR/MS) has long debated the existence of a theory-practice gap, where academic research is perceived as disconnected from the needs of practitioners. While this gap is difficult to trace in industrial contexts, the public policy domain offers a unique opportunity to measure research impact through formal citations. This study provides the first large-scale empirical investigation of the knowledge pathways connecting core OR/MS research to policymaking, by analyzing a comprehensive dataset of articles from core OR/MS journals linked to citing policy documents. We conceptualize these pathways through three key elements: the topic of the research, the academic institutions that produce it, and the policy institutions that cite it. Our results identify specific research topics and approaches (e.g., data envelopment analysis, innovation, organizational studies), journals (e.g., Management Science), and universities that serve as the most effective bridges between academia and policy, demonstrating that OR/MS research successfully bridges the gap when it aligns with pressing societal needs. Finally, our work also identifies significant opportunities for future policy impact in domains such as crime, housing, and immigration.

[Back to program](#)

Qualitative Meta-Analysis with System Dynamics Modeling for Process Theory Development

Management research has been dominated by variance theories that explain outcomes through statistical associations. While process theory complements variance theory by focusing on how and why events, activities, and choices interact over time, it remains challenging to acquire rich, longitudinal process data and translate them into novel, abstract process theories. We propose integrating qualitative meta-analysis (QMA) with system dynamics (SD) modeling to address these challenges. QMA systematically synthesizes temporally rich evidence across multiple qualitative case studies, while SD provides a formal language for articulating causal mechanisms, accommodating temporal complexity, and testing dynamic hypotheses. Using a recent application to servitization, we illustrate how QMA-SD can be implemented in a step-by-step manner. Beyond its substantive contribution to the servitization literature, we demonstrate the broader potential of QMA-SD for generating empirically grounded, generalizable, and testable process theories in domains characterized by abundant qualitative evidence but limited theoretical development.

[Back to program](#)

Agricultural Subsidies for Risk-Aware Farmers: Bilevel Optimization for Sustainable Crop Planning

Regulatory incentives and subsidies play a crucial role in encouraging sustainable food production practices such as crop rotation, reduced pesticide and fertilizer use, or crop diversification. However, such practices may lead to very high variation in yield and revenue. Current subsidy schemes may not be sufficient to incentivize risk-averse farmers to adopt these practices. To address this issue, we develop a bilevel optimization model that integrates a Target-MOTAD (Minimization of Total Absolute Deviation) approach into subsidy allocation, providing new insights on how regulatory incentives can simultaneously support farmer income stability and promote sustainable crop planning. The proposed framework aims to tackle two key challenges: (i) helping policymakers design effective mechanisms that encourage environmentally sustainable practices and minimize negative environmental impact of crop production, and (ii) supporting farmers in adopting these practices by reducing the financial risk they face.

Our experiments highlight how our subsidy design affects crop allocation and its environmental impact for varying levels of risk aversion and yield uncertainty. The findings demonstrate that carefully designed subsidies can stabilize farmer incomes while promoting sustainability, providing both policymakers and farmers with guidelines for decision-making when balancing economic and environmental concerns.

[**Back to program**](#)

Granting autonomy to warehouse workers - Lessons learned from experimental studies

Warehouses play a vital role in the supply chain of a company and contribute to its failure or success. Many activities are performed within a warehouse, but it has been shown that order picking (OP) is by far the costliest. Therefore, managers aim for high efficiency levels in the OP system. There are many planning decisions involved to devise an efficient OP system. One of them is the job assignment planning problem which manages the allocation of orders to pickers. Very often, a central planning system has full agency over this planning problem and sets out directives for human workers, potentially harming their (perceived) autonomy, which is found to affect worker well-being. In response, I developed a working system that grants more autonomy to order pickers. The efficacy of such a system was tested on several occasions and in different environments. I will present the overarching insights from this system's implementation, as well as its long-term implications. In several studies, I have investigated the impact of an order assignment mechanism in which order pickers get the opportunity to choose their next order at the depot. This assignment mechanism was tested on several occasions, for instance in a lab experiment. In total, 165 students with an education background in logistics took part in this study. I have also conducted a similar study in a real-world field experiment to test the generalisability of our findings. I bundled the results of these studies to have a better understanding of the effects of involving workers in operational decision-making. Using a holistic evaluation approach, I will show the beneficial impacts of an autonomy-increasing intervention on different organisational outcome measures.

[**Back to program**](#)

Truck Assignment and Scheduling with Mixed-Service Docks: A Q-learning Enhanced ALNS

Mixed service mode (MSM) docks enhance efficiency by flexibly handling both loading and unloading trucks in warehouses. However, existing research often predefines the number and location of MSM docks prior to planning truck assignment and sequencing. This predefined approach becomes less effective in high-demand systems and increases operational complexity, as warehouse operators must manually test various configurations. This paper addresses this challenge by proposing a new model that integrates dock mode decision, truck assignment, and scheduling, enabling more flexible dock mode arrangements. To solve the complex problem, we introduce a Q-learning-based adaptive large neighborhood search (Q-ALNS) algorithm, which adaptively adjusts dock modes through perturbation operators while simultaneously solving truck assignment and scheduling with destroy and repair operators. The Q-learning mechanism selects these operators based on their performance history and future gains, employing the epsilon-greedy strategy. Comprehensive experimental results and statistical analysis indicate that the Q-ALNS outperforms the benchmark in terms of optimality gap, with an average drop of 12.1%, while maintaining competitive computation efficiency. Compared to the predefined approach, our proposed adaptive strategy reduces tardiness by 22.5% and makespan by 7.6% on average, demonstrating its superiority in improving operational efficiency and supporting demand-driven assignment of MSM docks.

[Back to program](#)

A Dynamic Hypercube Model Approach for Helicopter Dispatching in Wildfire Management Considering Fire Behaviour

Wildfires pose a growing threat to ecosystems and human communities, necessitating rapid and effective suppression strategies to mitigate widespread damage. One such strategy is the strategic placement of helitack resources in high-risk areas. The effectiveness of this approach, however, depends critically on the positioning of these resources.

To address this challenge, we introduce a framework for determining the optimal positioning of helitack resources while incorporating fire behavior dynamics. We model the interaction between origin-destination dependent service times and subsequent resource availability using a hypercube model with dynamic service times, optimized using the Grey Wolf Optimizer. We validate the model through a case study in Québec, Canada.

[Back to program](#)

Suitable branching strategies for SOS1 constraints without unique ordering

In many mixed-integer programming problems, special ordered set (SOS) constraints are used to model mutually exclusive choices among decision variables. Although widely applicable in areas such as scheduling, assignment, and diet optimization, these constraints often lead to large branch-and-bound trees due to the high number of associated binary variables. When there is a natural ordering present in these SOS constraints, there are SOS branching strategies that can be exploited to decrease the number of branch-and-bound iterations required to solve the problem. However, when there is no natural ordering, these techniques often do not help to increase performance compared to standard branching rules. Therefore, this paper aims to propose and analyze new branching strategies by exploiting less obvious structures of matrices containing SOS constraints. In addition, these ideas are implemented and tested along a model taken from meal planning.

[Back to program](#)

Production Planning Optimization in a Symbiotic Network Under Demand Uncertainty

Industrial Symbiosis (IS) promotes eco-efficient collaboration between factories by repurposing one plant's byproducts as alternative raw materials for another. This research addresses the optimization of production planning in a symbiotic network under uncertain demand. To manage this uncertainty, both stochastic programming and distributionally robust optimization (DRO) models are developed. Computational results demonstrate that the DRO approach provides more robust and stable solutions than the stochastic model, particularly under complex and unpredictable demand patterns. Practically, the DRO model reduces cost variability and limits dependence on costly outsourcing, while the symbiotic collaboration itself lowers overall network costs by eliminating waste disposal and reducing raw material purchases. The findings highlight the potential of robust planning to enhance cost efficiency and resilience for industries engaged in IS.

[Back to program](#)

Socially Responsible Operations for Medical Devices

To address the limited access to medical equipment in underserved regions, we propose a refurbish-to-lease model where manufacturers refurbish equipment from well-resourced regions and lease it to resource-limited hospitals using a pay-per-use payment. Modeling the problem as a Markov Decision Process (MDP), our analysis shows that optimal pay-per-use prices follow a threshold policy. Numerical results demonstrate that this business model can increase manufacturer profits and expand healthcare access in resource-limited settings.

[Back to program](#)

Enhancing flexibility and efficiency in Dutch fire services through innovative vehicle deployment strategies

Dutch fire departments are increasingly faced with ever-growing demand, financial resource constraints, tight labor markets, and an increasing variety of incident types. These trends have led to the recent introduction of several innovative vehicle deployment strategies to enhance flexibility and resource efficiency. These include the deployment of swift intervention vehicles - relatively small vehicles operated by a two-person crew that are specifically used to provide a prompt initial response - and the possibility of composite deployments where multiple smaller units are dispatched and combined on scene to function as a single, larger operational unit. We analyze the impact of these new deployment strategies on key performance metrics and how to re-optimize logistical decision-making following their widespread implementation, specifically focusing on fleet composition and vehicle locations. We utilize a new coverage-based optimization approach that aims to minimize the resources required to deliver adequate fire services, where we allow for adequate fire services to be provided in multiple ways. This optimization approach is applied to a case study on the Netherlands to analyze how the new deployment strategies allow Dutch fire departments to deliver adequate fire services more efficiently. Results from this case study and corresponding insights will be presented.

[Back to program](#)

Designing Efficient Distribution Networks for Mixed-cropping Systems

Given the substantial environmental impact of current agricultural practices, rethinking agri-food systems is imperative. Mixed-cropping systems, growing multiple crops side by side, can increase productivity, improve product quality, enhance biodiversity, and promote soil health. However, they introduce significant supply chain challenges due to the increased variety of crops and smaller batch sizes, which complicate transportation and storage activities and increase farm production costs. Exploring alternative sales channels, such as short food supply chains (SFSCs), is therefore necessary, though the costs and benefits of this supply chain structure must be carefully evaluated.

We present a multi-commodity, multi-period, multi-echelon network design model to optimize key supply chain decisions, including sales channel selection, facility location, processing, transportation, and inventory management, while accounting for crop perishability. The model is used to design and assess SFSC configurations for mixed-cropping systems. Our contribution includes a holistic mixed-integer linear programming formulation and the introduction of novel constraints inspired by real-world storage incompatibilities between crops. Computational experiments on randomly generated instances include detailed sensitivity and structural analyses to evaluate the impact of varying model parameters and problem features, thereby demonstrating the value of addressing these interdependent decisions within an integrated framework.

[Back to program](#)

Joint Scheduling of AGVs, AMRs, and Machines with A Matheuristic Approach

With Industry 4.0 advancements, production environments are increasingly automated and intelligent, enhancing flexibility and efficiency. However, the integration of smart robots, including automated guided vehicles (AGVs) and autonomous mobile robots (AMRs), introduces complex scheduling and intralogistics challenges. This study addresses the simultaneous scheduling of assembly, product transport, and material supply tasks in a pull production system, ensuring the efficient coordination of automated and manual workstations. Inspired by an industry case study on electrical battery pack production, we develop a mixed-integer linear programming (MILP) model to minimize makespan while accounting for production, material supply, and transportation interdependencies. Due to the computational complexity of larger instances, we further propose a heuristic approach. Using real-world data, our approach is validated through simulation across varying fleet sizes, job numbers, and product types, demonstrating its applicability to high-mix-low-volume production and autonomous manufacturing systems.

[Back to program](#)

Controlling inventory of spare parts cooperatively

In recent years, there has been a rise in pooling and sharing projects among European Ministries of Defense to improve collaborative defense capabilities. With initiatives such that of Spare Parts Pooling, different armed forces commit to supplying each other with replacements for worn or broken components, when these are unavailable from national stocks. Motivated by this setting, we study a model in which a set of buyers collaborate in purchasing products jointly, pooling their inventories and their demands. The objective of this research is to investigate the stability of such cooperation. We achieve this with the tools of cooperative game theory and inventory theory, formalizing the problem as a cooperative game. We show that cooperation can be convenient, but not always stable, and present a possible solution in the form of a subsidy from an external organization. Furthermore, we give sufficient conditions for which an upper bound of the size of the subsidy holds.

[**Back to program**](#)