Service Operations Research Seminar HWS 2026 (OPM 781)

"Current Topics in Service Operations Management Research"

General Information:

1. The goal of this seminar is to introduce participants to conducting applied scientific research in the field of (service) operations management. Also, the seminar aims at practicing presentation skills, such as speaking with clarity, confidence and connection.

2. The master thesis prepares students for writing their M.Sc./Diploma Thesis primarily at the Chair of Service Operations Management, but OPM781 also qualifies you formally for writing a master thesis at any other chair in the Operations Area.

3. The offered topics are presented below and designed to be explored by a single student based on the fundamental literature. Each participant will present his/her findings in a written report (about 20 pages) as well as in an in-class presentation (~20 min), followed by a discussion (~10 min).

4. The **application procedure** for this seminar is combined with those for the seminars of the Chair of Production Management (OPM 761), the Chair of Logistics (OPM 701) and the Chair of Procurement (OPM 791). Students can apply for topics from all chairs by joining the ILIAS application group and completing the online form provided there. Topics labeled with "L" refer to the Chair of Logistics (OPM 701), topics labeled with "P" refer to the Chair of Production Management (OPM 761), topics labeled with "B" refer to the Chair of Procurement (OPM 791) and topics **labeled with "S"** refer to the **Chair of Service Operations Management (OPM 781)**. The **assignment of topics** to students will be preference-based through ILIAS. To better match topic and student background, applicants for OPM 781 may in addition send a CV and official grades overview by post to the chair or by e-mail sekretariat.som@uni-mannheim.de with subject "OPM 781 Seminar Application".[1]

5. The **application period** starts on **May 11th** and ends on **May 22nd, 2026**.

6. **Admission** to the seminar is **binding** and will be confirmed by E-mail by **May 26th, 2026,** at latest.

[1] Data protection: Please note that a breach of confidentiality and the unauthorized access by third parties cannot be excluded when transmitting an unencrypted email. Note on data protection: The submitted documents will be returned only if an envelope with sufficient postage is included. Otherwise they will be destroyed after the application process according to the requirements of the data protection law. Electronic applications will be deleted accordingly.

7. A **kick-off meeting** for all participants will be held on Thursday, **May 28th at 10:15 in SO318**. During this meeting, general guidelines for conducting scientific work will be discussed.

8. The latest **submission** date for the written report incl. appendices is **November 11th (2026),** For submission, please …

    a. **Upload your report** (Word- / Latex-document and PDF) via Task "Upload of final Thesis & Calculations/Software Output" in the OPM781 ILIAS group. If you have multiple files (e.g. a pdf and some Excel analysis), please upload all in a single zip file.

    b. **Submit a hard copy** at our secretary's office (Mon-Thu before noon) or at your thesis supervisor. Please make an appointment for submitting the hard copy.

9. The **final presentations** of the seminar participants will be held by default in the **regular presentation session** on **November 25th**, in room SO 318. A **fast-track presentation track** may be offered to students who desire to start with their master thesis early based on their request. Attendance is mandatory for all presentations on your own presentation date.

10. In addition, we will offer an **optional mock-up presentation** session one week before the regular final presentations, i.e., on **November 18th**. Here, participants can practice their final presentation and get tips on structure, content and presentation style risk-free without being graded – if they want to. In the kick-off meeting, we will provide some guidance and resources on how to train your presentation skills upfront, such as speaking with clarity, confidence and connection.

11. Please **upload your final presentation slides** (ppt and PDF) on Task "Upload of Final Presentation" in the ILIAS group **one day before the presentation**, latest by **18:00 pm**.

12. The final grade for the seminar is composed of the following components: Written report (50%) and presentation (50%).

13. For questions concerning the seminar contact us by email at sekretariat.som@uni-mannheim.de

Seminar topics

Please note:
The amount of recommended literature does NOT indicate more or less workload! Every thesis will contain some literature review, and more recommendations maybe helpful for this. Also, your supervisor may have more recommendations for you, in particular, if the initial list of recommended references for a topic is short.

Table of Contents

Topic S01: Temporal Flexibility in Regulating Environmental emissions

Environmental regulators worldwide have implemented a broad range of policies to mitigate the harmful externalities associated with industrial emissions, including acid rain from sulfur oxides, particulate pollution, and the greenhouse effect driven by carbon dioxide. Among these policy instruments, tradable emission permits have emerged as one of the most widely adopted and economically grounded regulatory tools. Under such systems, regulated entities are required to acquire and surrender permits that correspond to their verified emissions.

Programs such as the European Union Emissions Trading System (EU ETS), the U.S. Acid Rain Program (ARP), the NOx Budget Trading Program (NBP), and the Hunter River Salinity Trading Scheme exemplify the large-scale and durable implementation of permit-based regulation. These systems illustrate that market-based control can produce meaningful abatement incentives while the flexibility of firms is still preserved.

A key design dimension of such systems is temporal flexibility. In multi-period settings, compliance obligations may be defined separately for each period or allow for some degree of intertemporal transfer of emission permits, for example, through banking, borrowing, or other transfer mechanisms. Temporal flexibility can improve economic efficiency by enabling firms to shift compliance across periods in response to cost differences and uncertainty. At the same time, it may change the temporal distribution of emissions and thereby affect environmental outcomes. Fu et al. (2025) examine emission-permit designs with varying temporal flexibility across multiple compliance periods and discuss the results of extreme regimes as well as hybrids (e.g., transfer caps/discounts, permit-tax mixes).

The primary objectives of this thesis are to:

- provide an overview of regulatory instruments governing emissions,
- present selected components of the model of Xu et al. (2025), discuss the results, especially the aspect of temporal flexibility and relate them to the literature,
- identify future research opportunities.

Selected Literature Recommendations:

Fu, Xingyu, Ying-Ju Chen, Guillermo Gallego, Pin Gao, Mengqian Lu (2025). Designing Emission Permits Regulation for Multiple Compliance Periods. *Manufacturing & Service Operations Management 27(5):1587-1603*.

Gronwald, M., & Hintermann, B. (Eds.). (2015). *Emissions trading as a policy instrument: evaluation and prospects*. Mit Press.

Yuxuan Cao, Wanyu Ren, Li Yue (2024). Environmental regulation and carbon emissions: New mechanisms in game theory. *Cities*. *149 (104945).*

Topic S02: Global Supply Chains under Different CO_2 Regimes

Global supply chains are a defining feature of modern production systems. Raw materials, components, intermediate goods, and final assembly are often distributed across multiple countries and continents. At the same time, greenhouse-gas emissions generated along these networks have become a material concern for firms, regulators, and investors.

Managing emissions in global supply chains is complicated by the fact that carbon is measured, allocated, and regulated under different regimes and accounting logics around the world. Approaches differ in system boundaries, allocation methods, emissions factors, and regulatory treatment. Regulatory approaches also differ substantially. In Europe, mechanisms such as the EU ETS represent the current system. In the United States, carbon regulation is more fragmented, including systems such as California's cap-and-invest program or the Regional Greenhouse Gas Initiative. These differences can lead to varying emissions estimates, compliance obligations, and supply-chain decisions for the same physical network. Different accounting and regulatory approaches may alter sourcing decisions, production locations, transport choices, and regulatory exposure. For firms operating internationally, these differences raise strategic questions about comparability and optimization.

The primary objectives of this thesis are to:

- motivate the importance of carbon emissions in global supply chains,
- review international approaches for measuring, allocating, and regulating CO_2 emissions and highlight differences,
- use an academic example to illustrate how emissions in a global supply chain can be calculated under different accounting approaches and regulatory regimes,
- discuss how differing regimes may affect supply-chain decisions and identify future research opportunities.

Selected Literature Recommendations:

Caro, F., Corbett, C., Tan, T., & Zuidwijk, R. (2013). Double Counting in Supply Chain Carbon Footprinting. *Manufacturing & Service Operations Management*, 15(4), 545-558.

Dechezleprêtre, A.. Haramboure, A., Kögel, C., Lalanne, G., Yamano, N. (2025). Carbon Border Adjustments: The Potential Effects of the EU CBAM along the Supply Chain. *OECD Science, Technology and Industry Working Papers 2025/02i*.

Dubisz, D., Golinska-Dawson, P., & Koliński, A. (2023). Measuring CO2 emissions level for more sustainable distribution in a supply chain. *Journal of Engineering and Applied Science*. 49. 804-810.

Huang, Y., Li, X., & Wan, F. (2026). How do global value chain positions affect CO2 emissions in supply chains?. *International Review of Economics & Finance*, 104852.

Topic S03: Optimizing Product Return Policies in Omnichannel Retailing

Omnichannel retailing integrates physical stores, online shops, mobile applications, and other customer touchpoints into a unified retail system. While much of the existing omnichannel literature focuses on forward logistics, inventory allocation, pricing, and assortment decisions, product returns represent an equally important operational challenge, as returns are no longer restricted to the channel in which the original purchase was made.

The thesis will introduce the concept of omnichannel retailing and distinguish it from traditional retail settings, with particular emphasis on how product returns become more complex when customers can use multiple return channels. Building on this foundation, the thesis explains the main operational decisions in omnichannel returns management, including return channel selection, routing of returned products, recovery options, inventory implications, and the trade-off between customer convenience, responsiveness, and logistics costs. Following this, the main part of the thesis will review optimization-based models for omnichannel return policy selection. The thesis will then compare the strengths, limitations, and managerial implications of these models. Finally, the thesis will identify open research gaps and outline future research directions.

The objectives of the thesis are to:

- Introduce the concept of omnichannel retailing and explain how returns management differs between traditional and omnichannel retail settings,
- Explain the main operational decisions in omnichannel returns management,
- Review optimization-based models for omnichannel return policy selection,
- Compare strengths, limitations, and managerial implications of existing models for modeling returns in omnichannel retailing,
- Identify open research gaps and outline future research directions

Selected Literature Recommendations:

Hwang, S. O., Üster, H., & Savaskan-Ebert, R. C. (2025). Reverse channel selection for commercial product returns under time-to-market and product value considerations. *Annals of Operations Research*, 349(2), 1403–1440.

Ma, B., Mao, B., Liu, S., Meng, F., & Liu, J. (2024). Managing physical inventory and return policies for omnichannel retailing. *Computers & Industrial Engineering*, 190, 1–25

Yang, L., Li, X., Xia, Y. & Aneja, Y. P. (2023). Returns operations in omnichannel retailing with buy-online-and-return-to-store. *Omega*, 119

Topic S04: Operations Research and Machine Learning for Personalized Medical Decision-Making

Personalized medicine aims to tailor medical treatments to individual patients, with the potential to significantly improve outcomes for a wide range of diseases such as cancer, diabetes, and cardiovascular conditions. In this context, operations research can contribute to better, more transparent, and data-driven decision-making, ultimately supporting improved patient outcomes. At the same time, designing personalized treatments is challenging due to several factors, including patient heterogeneity, limited and noisy data, and the need to account for dynamic and long-term treatment effects. Recent advances increasingly combine machine learning with optimization techniques to address these challenges, enabling more flexible, adaptive, and scalable approaches to personalized decision-making.

This thesis will explore how operations research and machine learning methods can be applied to personalized medicine. It will explain the key challenges in this domain. Building on this foundation, the thesis will review methodological literature on personalized treatment design, with a particular focus on optimization-based approaches and their integration with machine learning methods. The core part of the thesis will provide a structured overview of prescriptive approaches to personalized medicine. The thesis will also discuss the strengths and limitations of existing approaches, particularly with respect to their practical applicability in healthcare settings. Finally, it will identify open research gaps and will outline promising directions for future research at the intersection of operations research, machine learning, and healthcare.

The objectives of the thesis are to:

- Introduce the concept of personalized medicine,
- Explain key challenges in personalized treatment design,
- Review and classify methodological literature on personalized medicine, with a focus on optimization-based and machine learning approaches,
- Discuss strengths and limitations of approaches from the reviewed literature,
- Identify open research gaps and outline future research directions,

Selected Literature Recommendations:

Bertsimas, D., Dunn, J., & Mundru, N. (2019). Optimal Prescriptive Trees. *INFORMS Journal on Optimization, 1(2)*, 164–183.

Du, J., Gao, S., & Chen, C.-H. (2024). A Contextual Ranking and Selection Method for Personalized Medicine. *Manufacturing & Service Operations Management, 26(1)*, 167–181.

Keyvanshokooh, E., Zhalechian, M., Shi, C., Van Oyen, M. P., & Kazemian, P. (2025). Contextual Learning with Online Convex Optimization: Theory and Application to Medical Decision-Making. *Management Science, 71(12)*, 10442–10464.

Lee, E. K., Wei, X., Baker-Witt, F., Wright, M. D., & Quarshie, A. (2018). Outcome-Driven Personalized Treatment Design for Managing Diabetes. *Interfaces, 48(5)*, 422–435.

<u>Topic S05: Queuing Analysis – Review and Application to Baria Planning Solutions</u>

Queueing analysis is a fundamental tool in operations management for studying systems in which demand arrives over time and must be processed by limited resources. It provides a structured framework to analyze congestion, delays, and resource utilization in service and production systems such as call centers, healthcare operations, and supply chains. By modeling key elements such as arrival processes, service-time distributions, routing, and system capacity, queueing analysis helps firms understand the trade-offs between efficiency and service quality, as well as the role of variability in shaping system performance. In relatively simple settings, analytical queueing formulas can be used to derive performance measures such as waiting times, queue lengths, utilization, and service levels. In more complex, multi-stage systems, discrete-event simulation provides a flexible complementary approach by explicitly tracking stochastic arrivals, service completions, routing, and resource constraints. Both methods are useful for studying variability propagation, that is, how variability generated at upstream stages amplifies delays downstream, and for identifying bottlenecks and evaluating process improvements.

The objective of the thesis is to review queueing analysis with variability propagation in terms of its modeling primitives (e.g., arrival processes, service-time distributions, routing, and resource capacity), key performance measures (e.g., waiting time, queue length/WIP, and service level), and methodological tools for capturing and diagnosing variability propagation in multi-stage systems. Furthermore, in the second part of the thesis, queueing analysis shall be applied to the Baria Planning Solutions case.

In the Baria Planning Solutions case, the Sales Support process operates as a multi-stage workflow in which requests arrive over time and must pass through several teams with finite capacity. The system exhibits strong demand seasonality and heterogeneous processing times, creating congestion and long turnaround times. The proposed thesis will conduct a simulation-based (or analytical formula-based) queueing analysis to evaluate the current Sales Support process and illustrate how variability drives delays and how variability propagates across stages. It will also propose suggestions for improving the Sales Support process and evaluate the performance of these suggestions.

Selected Literature Recommendations:

Wheelwright S. C., Schmidt W. (2011). Baria Planning Solutions, Inc.: Fixing the sales process. https://hbsp.harvard.edu/product/4568-PDF-ENG (case study will be provided by the chair)

Cachon, G. & Terwiesch, C. (2011): *Matching supply with demand, 3rd edition*, McGraw-Hill.

Hopp, W. J., & Spearman, M. L. (2011). *Factory physics. Waveland Press*. Chapter on variability basics.

Topic S06: Assortment and Price Optimization under Context-Dependent Choice Models

Discrete choice models are used to describe and predict individuals' decision-making among a set of alternatives. It captures how consumers trade off attributes of different options based on their preferences. Common models include the multinomial logit (MNL) and nested logit models. They are widely applied in marketing, economics, and operations to optimize pricing, product design, and assortment planning. These models were mostly derived from the assumption that the perceived utility that customers get from individual products is merely dependent on the product's own features. However, numerous studies in the marketing and cognitive science literature have shown that the assumption of utility being independent of the presence or absence of other options often fails to align with observed consumer choices and is in contrast with the so-called "context effects". Context effects refer to that consumers evaluate each option by considering both its absolute utility and its relative standing in the choice set (Trueblood et al. 2013). Recently, researchers incorporated context effects into discrete choice models to study its predictive power and address assortment and price optimization problems within the framework of context-dependent preferences.

The objectives of the seminar thesis are to

- review the literature on discrete choice models with context effects, published in leading academic journals, such as INFORMS journals, European Journal of Operational Research, Production and Operations Management, Journal of Operations Management, Journal of Revenue and Pricing Management,
- delve into a specific paper that addresses assortment or price optimization problems based on a discrete choice model with context effects, examining its model, advantages, limitations, and applications,
- identify and discuss open research gaps and future trends in the field of discrete choice models with context effects.

Selected Literature Recommendations:

Bai, Y., El Housni, O., Rusmevichientong, P., & Topaloglu, H. (2023). Assortment and price optimization under an endogenous context-dependent multinomial logit model. *Available at SSRN 4534984*.

Bai, Y., El Housni, O., Rusmevichientong, P., & Topaloglu, H. (2025). An Endogenous Multinomial Logit Model with Population-Based Purchase Feedback: Assortment Optimization, Pricing and Estimation.

Guo, L. (2016). Contextual deliberation and preference construction. *Management Science*, 62(10), 2977-2993.

Topic S07: Capacity Control and Pricing in Air Cargo: A Literature Review

Although the air cargo industry has become increasingly important for global supply chains, research in revenue management (RM) for cargo shipments still lags behind passenger RM. One reason is that the selling process is more complex: cargo capacity is sold through both long-term contracts and short-term spot market bookings. While spot-market cargo RM shares some similarities with passenger RM, several characteristics prevent a direct transfer of passenger methods. In particular, cargo capacity is typically two-dimensional (weight and volume) rather than one-dimensional (seats). In addition, capacity is often uncertain before departure (e.g., because passenger baggage and payload are prioritized for combination carriers or due to weather-related take-off weight limits), and shipment weight and volume requirements may not be known precisely at booking. Finally, cargo customers typically care less about routing complexity than passengers, as long as the shipment is delivered reliably and on time at an attractive price. As a result, much of the literature focuses on single-leg and network capacity allocation under independent demand assumptions (with and without overbooking), while choice-based models and dynamic pricing have received comparatively less attention

The objectives of this seminar thesis are to:

- Review and classify the literature on air cargo RM and discuss differences to airline passenger RM,
- Review relevant empirical studies from the literature that model and estimate demand response or customer choice as a function of different determinant attributes of cargo offerings,
- Discuss whether capacity allocation or dynamic pricing is the more suitable RM approach to cargo RM, considering opportunities, challenges, and limitations of moving from capacity allocation to dynamic pricing in the air cargo business

Selected Literature Recommendations:

Amaruchkul, K., Cooper, W. L., & Gupta, D. (2007). Single-leg air-cargo revenue management. *Transportation science, 41(4),* 457-469.

Barz, C., & Gartner, D. (2016). Air cargo network revenue management. *Transportation Science*, *50*(4), 1206-1222.

Becker, B., & Dill, N. (2007). Managing the complexity of air cargo revenue management. *Journal of Revenue and Pricing Management*, *6*(3), 175-187.

Feng, B., Li, Y., & Shen, Z. J. M. (2015). Air cargo operations: Literature review and comparison with practices. *Transportation Research Part C: Emerging Technologies*, *56*, 263-280.

Topic S08: Sustainable Airline Planning and Scheduling with Choice-Based Demand

With new aircraft frequently costing over $100 million, it is critical for airlines to maximize their utilization. For this, there are two main decisions to be made: 1. Which aircraft to assign to each flight, 2. How many tickets to sell for each itinerary? These decisions are interdependent, as the size of the assigned aircraft determines how many tickets can be sold. Furthermore, if there is little demand for a route, it should not be assigned a large aircraft, as its seating capacity can be used more profitably on other routes. Therefore, different optimization models have been developed to integrate schedule design and fleet assignment with demand modeling. One such approach is the Sustainable Airline Planning and Scheduling (SAPS) model by Krömer et al., which integrates the five most strategic stages of airline planning and scheduling while explicitly accounting for passenger choice, fuel burn, and CO2-related sustainability considerations. This makes the paper especially relevant as airlines increasingly face pressure to balance profitability with environmental targets and climate policy constraints.

This thesis introduces the concept of discrete choice modeling and aims to explore recent advancements in integrated demand modeling, fleet assignment, and sustainable airline planning models.

The objectives of the seminar thesis are to:

- provide an introduction to discrete choice models,
- Identify and analyze key challenges in integrated demand, schedule design, and fleet assignment models,
- explain the SAPS model by Krömer et al. in detail, including limitations and improvement opportunities (you do not need to discuss the solution method),
- Identify gaps and directions for future research

Selected Literature Recommendations:

Krömer, M. M., Topchishvili, D., & Schön, C. (2024). Sustainable airline planning and scheduling. *Journal of Cleaner Production*, *434*, 139986.

Yan, C., Barnhart, C., & Vaze, V. (2022). Choice-based airline schedule design and fleet assignment: A decomposition approach. *Transportation Science*, *56*(6), 1410-1431.

Aktürk, M. S., Atamtürk, A., & Gürel, S. (2014). Aircraft rescheduling with cruise speed control. *Operations Research*, *62*(4), 829-845.

Jalalian, M., Gholami, S., & Ramezanian, R. (2019). Analyzing the trade-off between CO2 emissions and passenger service level in the airline industry: Mathematical modeling and constructive heuristic. *Journal of cleaner production*, *206*, 251-266

Topic S09: Product-Line Design with New and Remanufactured Products

Firms across industries such as consumer electronics and automotive components increasingly offer both newly manufactured and remanufactured products alongside each other, driven by regulatory requirements, rising raw material costs, and growing consumer acceptance of remanufactured alternatives. Offering both product types simultaneously raises fundamental challenges. On the demand side, the market is segmented between consumers committed to new products and those open to remanufactured alternatives, and price discounts do not always increase attractiveness since very low prices may signal poor quality. On the supply side, the quantity and quality of cores returned for remanufacturing are not fixed — they depend on the firm's own take-back incentives and directly shape remanufacturing cost and feasibility. These demand- and supply-side considerations are deeply intertwined, as current sales create future remanufacturing opportunities.

Against this background, a growing stream of research asks how firms should jointly design and price a product line that includes both new and remanufactured products. Aydin et al., (2015) provide an integrated treatment, developing a multiobjective product-line design model that jointly determines product specifications, prices, and launch timing using conjoint-based demand estimation. Kwak & Kim (2017) extend this direction through a mixed-integer programming model that simultaneously optimizes buyback pricing, selling prices, and production planning. Both models rest on simplifying assumptions whose realism warrants careful examination.

The primary objectives of this thesis are to:

- provide a concise overview of the challenges that arise when new and remanufactured products are offered together, covering demand segmentation, cannibalization, and reverse-flow management,
- review the literature on product-line design models that incorporate remanufacturing,
- present Aydin et al., (2015) or Kwak & Kim (2017) as the focal model in depth, covering its assumptions, formulation, solution approach, and managerial insights,
- critically assess the focal model's assumptions against real-world conditions and the broader literature, and discuss how they could be relaxed or improved.

Selected Literature Recommendations:

Atasu, A., Sarvary, M., & Van Wassenhove, L. N. (2008). Remanufacturing as a marketing strategy. *Management Science*, 54(10), 1731–1746.

Aydin, R., Kwong, C. K., & Ji, P. (2015). A novel methodology for simultaneous consideration of remanufactured and new products in product line design. *International journal of production economics, 169*, 127-140.

Debo, L. G., Toktay, L. B., & Van Wassenhove, L. N. (2005). Market segmentation and product technology selection for remanufacturable products. *Management Science, 51(8)*, 1193–1205.

Kwak, M., & Kim, H. (2017). Green profit maximization through integrated pricing and production planning for a line of new and remanufactured products. *Journal of cleaner production, 142*, 3454-3470.

Topic S10: Artificial Intelligence in Service Operations and Revenue Management: LLMs as Decision Interfaces

Service operations and revenue management increasingly rely on optimization tools for pricing, capacity allocation, demand forecasting, and resource scheduling. However, practitioners often struggle to interpret solver outputs, perform what-if analyses, or adjust models when business conditions change. These tasks usually require support from data scientists or engineers, slowing decision-making and creating a gap between analytical capability and operational execution. Large language models (LLMs) can help reduce this gap by serving as natural-language interfaces to existing optimization systems. Rather than replacing mathematical optimization, they can help managers query, interpret, and modify decision tools in plain language.

Recent research explores how LLMs can support optimization-based decision-making while maintaining correctness and control. Li et al. (2023) introduce OptiGuide, a framework that uses LLMs to translate plain-language questions into code, interact with optimization solvers, and explain the resulting outputs — a methodology whose core design is transferable to service operations and revenue management settings. Simchi-Levi et al. (2025) extend this perspective to broader operational decision environments. Menache et al. (2025) provide a managerial view of how generative AI can reduce the time required for explanation, scenario analysis, and decision support. Cohen et al. (2026) position these developments within a broader AI-enabled operations framework, emphasizing the integration of AI with operations management principles, human judgment, and governance. Together, these works highlight both the promise and the risks of LLM-based decision interfaces, including hallucination, code verification, query ambiguity, data privacy, and cost-performance tradeoffs.

The primary objectives of this thesis are to:

- provide a concise overview of how AI, particularly machine learning and LLMs, is transforming service operations and revenue management,
- review literature on LLM-based decision support in operations, classified by use case, methodology, and reported impact,
- present Li et al. (2023) and Simchi-Levi et al. (2025) as focal contributions, covering the OptiGuide framework, evaluation approaches, and implementation insights, with emphasis on the transferability of these methods to service operations and revenue management contexts,
- critically assess key limitations, including hallucination, correctness verification, ambiguous queries, data quality, privacy, and cost-performance tradeoffs,
- optionally conduct a small empirical test using a well-defined service operations or revenue management problem, such as a newsvendor pricing model or a single-leg seat allocation problem, to evaluate whether an LLM can translate plain-language problem descriptions and what-if modifications into valid mathematical formulations and working solver code.

Selected Literature Recommendations

Cohen, M. C., Dai, T., Perakis, G., Agrawal, N., Allon, G., Boute, R. N., ... & Zhang, D. (2026). OM Forum—Supply Chain Management in the AI Era: A Vision Statement from the Operations Management Community. *Manufacturing & Service Operations Management*.

Li, B., Mellou, K., Zhang, B., Pathuri, J., & Menache, I. (2023). Large language models for supply chain optimization. *arXiv preprint arXiv:2307.03875*.

Menache, I., Pathuri, J., Simchi-Levi, D., & Linton, T. (2025). How generative AI improves supply chain management. *Harvard Business Review, 104(1-2)*, 86-95.

Simchi-Levi, D., Mellou, K., Menache, I., & Pathuri, J. (2025). Large language models for supply chain decisions. *arXiv preprint arXiv:2507.21502*.

<u>Topic S11: Regulating Environmental Emissions</u>

Environmental regulators worldwide have implemented a broad range of policies to mitigate the harmful externalities associated with industrial emissions, including acid rain from sulfur oxides, particulate pollution, and the greenhouse effect driven by carbon dioxide. Among these policy instruments, tradable emission permits have emerged as one of the most widely adopted and economically grounded regulatory tools. Under such systems, regulated entities are required to acquire and surrender permits that correspond to their verified emissions.

Programs such as the European Union Emissions Trading System (EU ETS), the U.S. Acid Rain Program (ARP), the NOx Budget Trading Program (NBP), and the Hunter River Salinity Trading Scheme exemplify the large-scale and durable implementation of permit-based regulation. These systems illustrate that market-based control can produce meaningful abatement incentives while preserving flexibility for firms.

Xu et al. (2025) examine emission-permit designs with varying temporal flexibility across multiple compliance periods and discuss the results of extreme regimes as well as hybrids (e.g., transfer caps/discounts, permit-tax mixes).

The primary objectives of this thesis are to:

- provide an overview of regulatory instruments governing emissions,
- present selected components of the model of Xu et al. (2025), discuss the results and relate them to the literature,
- identify future research opportunities.

Selected Literature Recommendations:

Fu, Xingyu, Ying-Ju Chen, Guillermo Gallego, Pin Gao, Mengqian Lu (2025). Designing Emission Permits Regulation for Multiple Compliance Periods*. Manufacturing & Service Operations Management 27(5):1587-1603*.

Gronwald, M., & Hintermann, B. (2016). Emissions trading as a policy instrument: evaluation and prospects. Cambridge, Massachusetts.

Yuxuan Cao, Wanyu Ren, Li Yue (2024). Environmental regulation and carbon emissions: New mechanisms in game theory. *Cities*. 149 (104945).

<u>Topic S12: Evaluating the Performance of MNL and MMNL Choice Models in Assortment Optimization and Product Line Selection</u>

Assortment optimization and product line selection are crucial strategic decisions for firms aiming to align product offerings with consumer preferences. These decisions are typically modeled using discrete choice frameworks such as the Multinomial Logit (MNL) and Mixed Multinomial Logit (MMNL) models, each capturing customer behavior with varying degrees of complexity and realism.

Under the MNL model, optimal assortments can often be computed efficiently using convex optimization techniques. A notable example is provided by Chen and Hausman (2000), who examined a joint pricing and assortment selection problem. Their formulation as a nonlinear binary integer program could be relaxed without loss of optimality due to the problem's favorable structural properties, allowing the relaxed solution to remain integral and globally optimal. However, the MNL model assumes homogeneous consumer preferences, which may not accurately reflect real-world behavior. The MMNL model addresses this by incorporating preference heterogeneity, most probably offering a better empirical fit. Despite its advantages, the MMNL-based assortment problem is computationally intractable and can only be solved exactly for small problem sizes using mixed-integer linear programming (MILP) formulations. This seminar examine how much solution quality is lost when an MNL model is used as an approximation for an MMNL model?

The primary objectives of this thesis are to:

- explore recent work on assortment optimization and product line selection under MNL and MMNL models,
- utilize parameter estimates from a relevant empirical study (e.g., Keane & Wasi, 2013) to formulate and solve the problem under both MNL and MMNL frameworks,
- quantify the performance bias introduced by using a single-segment MNL model instead of MMNL.

Selected Literature Recommendations:

Chen, K. D., & Hausman, W. H. (2000). Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. *Management Science, 46(2)*, 327–332.

Keane, M., & Wasi, N. (2013). Comparing alternative models of heterogeneity in consumer choice behavior. *Journal of Applied Econometrics, 28(6),* 1018–1045.

Schön, C. (2010). On the optimal product line selection problem with price discrimination. *Management Science, 56(5*), 896–902.

Topic S13: Aligning User Preferences and Profits: Exploring the Intersection of Recommender Systems and Choice-Based Optimization

With digital services becoming ever more embedded in daily life, users face an overwhelming abundance of options — whether selecting a restaurant on Lieferando.de, booking a hotel on Booking.com, or finding a new mobile phone contract on Check24.de. Recommender Systems (RS) have become vital tools for helping users navigate this abundance by filtering and ranking options based on inferred preferences.

Econometric choice models, while generally simpler and more theory-driven than data-driven approaches, such as RS, have long served as powerful tools for predicting consumer behaviour. Naturally, they are regularly used to solve pricing, product design, and assortment planning problems. Both recommender systems (RS) and choice-based optimization address a selection problem: identifying the optimal subset of items. In choice-based optimization, this process is typically guided by explicit business objectives (e.g., maximizing expected revenue) and constrained by operational limits such as inventory, budget, or shelf space. Recommender Systems focus on user-oriented objectives such as engagement, satisfaction, and retention, which may not always align directly with a company's strategic or financial goals. In addition, RS frequently integrate design considerations like diversity, novelty, and catalogue coverage to improve the overall user experience. These aspects, however, are typically embedded implicitly within the model training process rather than being formulated as explicit optimization constraints. Despite these conceptual parallels between RS and choice-based optimization, a systematic investigation of their relationship is largely missing from the existing literature. This presents a valuable opportunity to integrate theoretical insights from choice modelling with data-driven approaches from recommender systems, ultimately leading to more effective, user-aware, and business-aligned decision-making frameworks.

This thesis will critically examine the relationship between RS and choice-based optimization, with a particular focus on how these fields can complement and inform one another. To this end, the thesis will begin by introducing the fundamentals of discrete choice models and choice-based optimization, alongside a comprehensive overview of RS, including collaborative filtering, content-based filtering, and hybrid approaches. Literature at the intersection of RS and choice-based optimization will be critically reviewed to assess the conceptual and methodological relationships between the two fields, briefly evaluating their respective strengths and weaknesses, and areas of divergence and convergence. Finally, the thesis will identify open research gaps at the intersection of these domains and outline promising directions for future work.

The primary objectives of this thesis are to:

- introduce the theoretical foundations of RS and choice-based optimization,
- critically review literature at the intersection of RS and choice-based optimization,

- evaluate the strengths and weaknesses of RS and choice-based optimization in the context of managing company operations,
- identify open research gaps and outline future research directions.

Selected Literature Recommendations:

Feldman, J., Zhang, D. J., Liu, X., & Zhang, N. (2022). Customer Choice Models vs. Machine Learning: Finding Optimal Product Displays on Alibaba. *Operations Research, 70(1)*, 309–328.

Kallus, N., & Udell, M. (2020). Dynamic Assortment Personalization in High Dimensions. *Operations Research, 68(4),* 1020–1037.

Jiang, H., Qi, X., & Sun, H. (2014). Choice-Based Recommender Systems: A Unified Approach to Achieving Relevancy and Diversity*. Operations Research, 62(5),* 973–993.

Ricci F, Rokach L, Shapira B (2011). Introduction to recommender systems handbook. Ricci F, Rokach L, Shapira B, Kantor PB, eds*. Recommender Systems Handbook (Springer, Boston),* 1–35.

Roy, D., & Dutta, M. (2022). A systematic review and research perspective on recommender systems*. Journal of Big Data, 9(1)*, 59.

Topic S14: Learning to Optimize: Machine Learning for Combinatorial Optimization

Combinatorial optimization (CO) problems are ubiquitous in real-world decision-making and industrial planning. Companies such as Amazon, UPS, and FedEx must solve large-scale vehicle routing and delivery scheduling problems daily to minimize transportation costs and delivery times. Likewise, airlines face complex crew scheduling and aircraft assignment problems, while retail and e-commerce, firms like Walmart and Zara tackle assortment optimization and product line design problems, where the goal is to select an optimal subset of products or configurations that maximize expected profit or customer satisfaction subject to production, space, or cannibalization constraints. Each of these settings can be formalized as a combinatorial optimization problem, where decisions are discrete and interdependent, and the feasible set grows exponentially with the problem size.

Traditionally, such problems have been addressed using operations research (OR) techniques such as branch-and-bound algorithms and various meta heuristics. These approaches rely heavily on handcrafted rules and expert knowledge to guide the search process. However, recent research has shown that machine learning (ML) methods can learn optimal or near-optimal decision policies for combinatorial optimization problems by leveraging data from past instances, for example, historical delivery routes or previous product configurations. Nevertheless, other important classes of combinatorial optimization problems, unrelated to transportation, remain relatively underexplored, offering promising avenues for advancing learning-based optimization schemes.

This thesis will critically review the existing literature on the application of ML to CO problems. It will provide an overview of the different types of combinatorial problems and their defining characteristics, examine the range of ML schemes employed in the literature, and analyse how these learning-based approaches operate and perform in comparison to traditional optimization methods. Furthermore, the thesis will briefly discuss the strengths and limitations of current approaches, highlighting their practical implications and theoretical underpinnings. Finally, it will identify open research gaps at the intersection of machine learning and combinatorial optimization and outline promising directions for future research in this rapidly evolving field.

The primary objectives of this thesis are to:

- introduce the field of ML for CO,
- explain the principal classes of combinatorial optimization problems along with their defining characteristics and elaborate on the theoretical foundations that underpin the various optimization schemes analyzed in this thesis,
- review the range of ML schemes employed in the literature to solve CO problems and analyze how these learning-based approaches operate and perform in comparison to traditional optimization methods,
- discuss the strengths and limitations of ML for solving CO problems,
- identify open research gaps and outline future research directions,

Selected Literature Recommendations:

Bello, I., Pham, H., Le, Q. V., Norouzi, M., & Bengio, S. (2017). Neural Combinatorial Optimization with Reinforcement Learning (No. arXiv:1611.09940). *arXiv.*

Bengio, Y., Lodi, A., & Prouvost A. (2021). Machine learning for combinatorial optimization: A methodological tour d'horizon. *European Journal of Operational Research, 290(2),* 405-421.

He, H., Daumé, H., & Eisner, J. (2014). Learning to Search in Branch and Bound Algorithms. *Advances in Neural Information Processing Systems,* 27.

Zhou, F., Lischka, A., Kulcsar, B., Wu, J., Chehreghani, M. H., & Laporte, G. (2025). Learning for routing: A guided review of recent developments and future directions. *Transportation Research Part E: Logistics and Transportation Review*, 202, 104278.

Topic S15: Reinforcement Learning Approaches for Dynamic Revenue Management

Revenue management (RM) plays a critical role in industries such as airlines, hospitality, and retail, where firms must optimally set prices and allocate limited capacity to uncertain and time-dependent demand. Traditional RM systems rely on dynamic programming or stochastic optimization models that assume known demand distributions and customer behaviors. However, in practice, these assumptions are often violated: demand functions are unknown, customer preferences evolve, and competitive environments are highly dynamic. Consequently, the ability to learn optimal pricing and allocation policies directly from interaction with the environment – without requiring precise model specifications – has motivated growing interest in reinforcement learning (RL) as a powerful alternative.

In recent years, RL has shown promise in addressing several limitations of classical RM models. By treating pricing and inventory control as sequential decision-making problems under uncertainty, RL can adaptively balance exploration (learning demand patterns) and exploitation (maximizing revenue based on current knowledge). Moreover, RL frameworks naturally handle high-dimensional state spaces, dynamic customer segments, and complex competitive settings where closed-form solutions are infeasible. Despite these advantages, applying RL to RM remains challenging due to issues such as sparse feedback, non-stationary environments, and the need for explainability in decision-making systems that affect pricing fairness and customer trust.

This seminar thesis focuses on the intersection of reinforcement learning and revenue management, exploring how modern RL methods can enhance dynamic pricing and capacity allocation under uncertainty. In particular, it builds upon recent work that applies RL to model dynamic customer behavior, learn pricing strategies from data, and optimize revenues across multiple time horizons and market conditions.

The objectives of this thesis are to:

- provide a comprehensive literature review of reinforcement learning approaches for revenue management,
- explain the underlying key methods and assumptions of RL models for dynamic pricing and capacity control theoretically and with academic examples,
- discuss a selected approach in detail, including contribution, model and method, limitations and improvement opportunities,
- theoretically compare RL-based RM approaches with classical stochastic and optimization-based methods in terms of scalability, adaptability, and robustness,
- identify future research opportunities at the interface of machine learning and revenue management.

Selected Literature Recommendations

Bondoux, N., Nguyen, A.Q., Fiig, T. et al. Reinforcement learning applied to airline revenue management. *J Revenue Pricing Manag 19*, 332–348 (2020). https://doi.org/10.1057/s41272-020-00228-4

Lange, F., & Schlosser, R. (2025). Dynamic pricing with waiting and price-anticipating customers. *Operations Research Perspectives*, 14, 1-20.

Ilan Lobel (2020) Revenue Management and the Rise of the Algorithmic Economy. *Management Science 67(9):5389-5398*. https://doi.org/10.1287/mnsc.2020.3712

Meng, H., Chen, N., & Gao, X. (2024). Reinforcement learning for intensity control: An application to choice-based network revenue management. *arXiv preprint arXiv:2406.05358.*

Shihab, S.A.M., Wei, P. A deep reinforcement learning approach to seat inventory control for airline revenue management. *J Revenue Pricing Manag 21*, 183–199 (2022). https://doi.org/10.1057/s41272-021-00281-7