Business School
Chair of Supply Chain Management
Prof. Dr. Moritz Fleischmann

UNIVERSITY OF MANNHEIM
Business School

OPM 701 – Research Seminar Supply Chain Management
HWS 2026

General Information:

1. The goal of this seminar is to introduce the participants to conducting scientific research. It thereby prepares the students for the writing of their MSc thesis. The seminar is geared towards students intending to write their thesis at the Chair of Supply Chain Management.

2. Each participant will explore one of the research papers listed below. The task is to review and critically assess the assigned research paper and to relate it to the corresponding stream of scientific literature. Each participant presents their findings in a written report (about 20 pages) as well as in a presentation (20 min + 20 min discussion).

3. A **kick-off meeting** for all participants will be held on **Thursday, 8 June, 15:30** in the room **O226-8.** General guidelines for conducting a scientific literature review will be discussed, and the deliverables of the seminar will be explained in detail.

4. The written reports must be **submitted** electronically and as a hard copy in two-fold by **Wednesday, 4 November**.

5. The **presentations** will be held as a blocked session, most probably on **12 and 13 November** (subject to change; exact times and room to be announced).

6. The final grade for the seminar is composed of the following components: Written report (60%), presentation (30%), contribution to discussion (10%).

7. As the seminar is usually attended by a class of international students, the report and the presentations should be delivered in English.

8. There is a joint application process for all seminars offered by the chairs of the Area Operations Management. In the HWS 2024, this includes the following seminars:

   - OPM 701: Research Seminar Supply Chain Management (Chair of Supply Chain Management), **labeled with 'L'**
   - OPM 760: Project Seminar Operations Analytics (Chair of Production Management), **labeled with 'P'**
   - OPM 761: Research Seminar Production Management (Chair of Production Management), **labeled with 'P'**
   - OPM 781: Research Seminar Service Operations (Chair of Service Operations Management), **labeled with 'S'**
   - OPM 791: Research Seminar Procurement (Chair of Procurement), **labeled with 'B'**

   Detailed information on the seminar topics is available on the home pages of the respective chairs. In their application, students can indicate three to five preferred topics from all seminars.

9. Applications are open within the period from **11 May** to **22 May**. Students have to join the ILIAS group **Seminar Application Area Operations** (link) and complete the **application form** there.

10. Additionally, students applying for a topic of OPM 701 must send an e-mail to scm@uni-mannheim.de, titled "Seminar Application Documents", including a current **CV** and a **grades overview** (the one you can print yourself is enough). If you are applying for topics of the other chairs, please check if you must send documents to them as well.

11. For any questions concerning the seminar, feel free to contact Katrin Waßmuth at mirbeygi@uni.mannheim.de.

Seminar topics

Each participant will be assigned one of the **topics listed below**. The task then is to identify the main issues addressed by the paper, explain its methodology, including potential quantitative models, position it in the corresponding stream of scientific literature, and critically assess the paper's contribution to the literature as well as to practice.

Topic L01: Rojas, B., Larrain, H., Klapp, M., & Banerjee, D. (2026). A Newsvendor Model for Last-Mile Fleet Sizing.

The study addresses the tactical problem of determining a last-mile delivery fleet size while accounting for day-to-day uncertainty in the number and location of customer requests. An optimally sized fleet must balance the cost of contracting vehicles against the penalty costs of unserved customers: a larger fleet reduces the risk of unserved demand, whereas a smaller fleet is less costly. This trade-off resembles the structure of a newsvendor problem, and the fleet-sizing decision problem is modeled accordingly. However, unlike in the classical newsvendor setting, the expected "return," that is, the number of served requests associated with a particular fleet size, is difficult to characterize because of the complexity induced by combinatorial vehicle-routing and customer-selection decisions. Therefore, a key technical challenge lies in estimating how many requests can be served by fleets of different sizes.

As an alternative to solving a hard stochastic team-orienteering problem for each fleet size, the study presents two continuous-approximation approaches that capture the structure of the fleet-sizing problem at an aggregate level. The first treats linehaul time as constant, while the second models linehaul time as variable depending on each vehicle's route location; both approaches rely on the well-known Beardwood-Halton-Hammersley theorem. These approximations require low computational effort while providing structural insights. Crucially, the study shows that the resulting total-cost functions are convex with respect to fleet size, just as the classical newsvendor cost function is convex with respect to inventory quantity. This result allows optimal fleet sizes to be computed efficiently by leveraging first-order optimality conditions. The models are then used to evaluate optimal fleet sizes and associated costs under different linehaul-time formulations, information structures regarding future demand, and depot locations. Finally, the continuous-approximation models are validated through simulation experiments on both synthetic and real road networks, demonstrating their effectiveness and practical usability.

Topic L02: Silva, W. A., Carvalho, M., & Jena, S. D. (2026). Dynamic facility location under cumulative customer demand. *Transportation Science*, *60*(1), 46-63.

Dynamic facility-location problems address the placement of one or more valuable resources over a planning horizon in order to meet customer demand. The existing literature commonly assumes that customer demand quantities are defined independently for each time period. In many planning contexts, however, unmet demand carries over to future periods. As a result, unmet demand in some periods may affect decisions in subsequent periods.
This work studies a novel location problem in which the decision maker places facilities over time to capture cumulative customer demand. Two mixed-integer programming formulations are proposed for this problem. The study shows that one formulation has a tighter continuous relaxation and allows for the representation of more general customer-demand behavior. It also characterizes the computational complexity of the problem and analyzes which problem characteristics lead to NP-hardness. The study then proposes an exact branch-and-Benders-cut method and shows, through computational experiments, that this method is approximately five times faster on average than solving the tighter formulation directly. The results also quantify the benefit of accounting for cumulative customer demand within the optimization framework, showing that the corresponding planning solutions perform substantially better than those obtained by ignoring cumulative demand or by employing myopic heuristics. Finally, the study provides managerial insights into the quality of service perceived by customers when a provider places facilities under cumulative customer demand.

<u>Topic L03:</u> Mao, W., Zhang, K., Zhu, R., Simchi-Levi, D., & Başar, T. (2025). Model-free nonstationary reinforcement learning: Near-optimal regret and applications in multiagent reinforcement learning and inventory control. *Management Science*, *71*(2), 1564-1580.

The study considers model-free reinforcement learning in nonstationary Markov decision processes. Both the reward functions and the state-transition functions are allowed to vary arbitrarily over time, as long as their cumulative variations do not exceed certain variation budgets. The work proposes Restarted Q-Learning with Upper Confidence Bounds, referred to as RestartQ-UCB, as the first model-free algorithm for nonstationary reinforcement learning. The study shows that this algorithm outperforms existing solutions in terms of dynamic regret. Specifically, RestartQ-UCB with Freedman-type bonus terms achieves a dynamic-regret bound, where $S$and A denote the numbers of states and actions, respectively, $\Delta > 0$ denotes the variation budget, Hdenotes the number of time steps per episode, and Tdenotes the total number of time steps.

The study further presents a parameter-free algorithm, called Double-Restart Q-UCB, which does not require prior knowledge of the variation budget. It shows that the proposed algorithms are nearly optimal by establishing an information-theoretic lower bound, which is the first lower bound for nonstationary reinforcement learning. Numerical experiments validate the advantages of RestartQ-UCB in terms of both cumulative rewards and computational efficiency. The study also demonstrates the applicability of its results through examples in multi-agent reinforcement learning and inventory control across related products.

<u>Topic L04:</u> Lei, D., Qi, Y., Liu, S., Geng, D., Zhang, J., Hu, H., & Shen, Z. J. M. (2025). Pooling and boosting for demand prediction in retail: A transfer learning approach. *Manufacturing & Service Operations Management*, *27*(6), 1779-1794.

The study examines how retailers should leverage aggregate category-level sales information to improve individual product demand prediction. Motivated by the idea of inventory risk pooling, the work develops a new prediction framework that integrates category-level and product-level sales information in order to exploit the benefits of pooling. The proposed approach combines data from different aggregation levels within a transfer-learning framework. Specifically, it treats top-level sales information as a form of regularization when fitting the bottom-level product demand prediction model. The study characterizes the error performance of the model in linear settings and demonstrates the benefit of pooling. Moreover, the approach exploits a natural connection to regularized gradient-boosting trees, enabling scalable implementation in large-scale retail applications.

Using an internal study with JD.com based on more than 6,000 weekly observations from 2020 to 2021, the study evaluates the out-of-sample forecasting performance of the proposed approach against state-of-the-art benchmarks. The results show that the approach consistently delivers superior forecasting performance, achieving more than a 9% improvement over JD.com's benchmark method. The study further validates the generalizability of the framework using a Walmart retail dataset and through alternative pooling and prediction methods. From a managerial perspective, the study shows that using aggregate sales information directly may not necessarily improve product-level demand prediction. Instead, the results highlight the value of transfer learning for retail demand prediction, supported by both theoretical and empirical evidence. Based on a conservative estimate from JD.com, the improved forecasts can reduce

operating costs by 0.01–0.29 renminbi per sold unit on the retail platform, implying significant cost savings for a low-margin e-retail business.

Topic L05: Esenduran, G., Letizia, P., & Ovchinnikov, A. (2022). Customization and returns. *Management Science*, *68*(6), 4517-4526.

Recent advances in information technology, advanced manufacturing technologies such as robotics and 3D printing, and logistics have enabled firms to customize products according to the specifications of individual consumers. Consumers often prefer such customized products to standard alternatives. However, when customized products do not meet expectations, consumers may still feel entitled to return them. This raises the question of whether firms should offer returns on customized products.

The study examines this question using a Stackelberg game model. In the model, the firm acts as the leader and decides the prices and return policies for both customized and standard products. Consumers act as followers and decide which product to buy based on their initial noisy valuations. After experiencing the product, they then decide whether to return it. Both the firm and consumers behave strategically: forward-looking consumers incorporate the real-option value of possible returns into their initial purchasing decisions, while the firm incorporates consumers' optimal purchase and return responses into its pricing and return-policy decisions.

The model generates three main insights. First, firms can use customized products to induce some consumers who would otherwise buy and return a standard product to switch to customized products with lower return rates. Second, it may be optimal for firms to offer returns on customized products, even though these products have lower salvage value. Third, firms can increase profits and reduce total returns by offering returnable customized products.

Topic L06: Jalili, M., Pangburn, M. S., & Yazdani, A. (2024). Trend-Chasing Versus Minimalism: Selling Fewer, Better Products to Fashion-Sensitive Customers. *Production and Operations Management*, *33*(4), 922-942.

The study examines the criticism that fashion sellers contribute to waste by selling low-durability products, as well as the criticism directed at consumers who continue purchasing new fashion items despite already having accumulated many prior fashion products. In this context, the slow-fashion movement encourages sellers to produce more durable products, thereby supporting less frequent consumer purchases.

Using an infinite-horizon model with strategic consumer behavior, the study analyzes the seller's profit-maximizing pricing and product-durability decisions. The model explicitly accounts for consumers' ability to accumulate a "closet" of product varieties over time. The study initially assumes static pricing and later examines the potential profit gains from dynamic pricing. In the heterogeneous consumer market, consumers are first allowed to differ, according to a uniform distribution, in their sensitivity to fashion. The study then explores alternative distributions of fashion sensitivity, as well as the correlation between consumers' fashion sensitivity and their product valuations. Within this framework, the study shows how the seller's optimal pricing and durability decisions generate distinct shopping segments, referred to as minimalists and trend-chasers. The results show that when the degree of fashion uncertainty is moderate, the seller's optimal product-durability choice supports the emergence of both types of behavior. As variety uncertainty increases, if the seller's costs are sufficiently low, the seller supports a throwaway

culture by offering disposable products. Otherwise, when costs are high, the seller optimally targets a slow-fashion-type outcome, in which consumers focus on reuse through durable products rather than variety. Overall, the findings shed light on consumers' optimal purchasing behavior as a function of market parameters and the firm's pricing and durability decisions.

Topic L07: Ürkmez, S., & Kayış, E. (2026). Can carbon border adjustment mechanisms reduce carbon leakage? Revisiting the dual sourcing problem for a manufacturer. *Journal of Cleaner Production*, *554*, 148055.

The study examines the Carbon Border Adjustment Mechanism, which is designed to complement emissions trading systems by addressing carbon leakage through carbon pricing on selected imports. It fills a gap in the literature by providing a firm-level perspective on the effects of CBAM on supplier selection, dual-sourcing decisions, and green-investment incentives. The study also examines the effectiveness of CBAM in reducing carbon leakage and promoting green-technology investment under decentralized supply-chain decision-making. A game-theoretical model is developed to analyze supplier selection under a split-sourcing strategy. In this setting, a manufacturer chooses between a domestic supplier located in an ETS region and an overseas supplier located in a non-ETS region that is subject to CBAM. The study derives the equilibrium decisions of each firm and presents a case study using real-world supplier data from the steel industry and other hard-to-abate industries.

The model reveals nontrivial managerial insights for the relevant stakeholders, including supply-chain participants and regulatory bodies. The study highlights the complex interaction between carbon-pricing mechanisms and supply-chain dynamics. It emphasizes the need for a balanced approach to pricing strategies, green-technology investments, and environmental objectives. The results suggest that policymakers should account for these interdependencies when designing ETS and CBAM, since their effectiveness depends on well-calibrated policy parameters. The study further shows that CBAM and ETS should not be considered in isolation from the broader economic context, particularly demand and supply dynamics, because doing so may lead to unintended consequences.

Topic L08: Zhao, S., Gui, L., & Cui, S. (2026). Emission Reduction Through Regulating Indirect Sources. *Manufacturing & Service Operations Management*, *28*(1), 193-211.

The study examines emissions from diesel semitrucks, such as nitrogen oxides, which have contributed significantly to air pollution and prompted government intervention. However, directly regulating trucking companies' diesel-truck usage often falls outside the jurisdiction of local governments. As an alternative, governments may regulate other sectors within the local region that indirectly drive diesel-truck usage, referred to as indirect emission sources.

The study is motivated by Southern California's Rule 2305, the warehouse indirect source rule, which is the first regulation of this type. Under this rule, warehouses are held accountable for diesel-truck visits to their facilities through a mitigation fee. The policy aims to incentivize the adoption of electric semitrucks and reduce air pollution. The study examines the effectiveness of the indirect source rule by focusing on its environmental impact and the industry burden it imposes, compared with a hypothetical direct source rule that would regulate trucking companies directly.

The study develops game-theoretical models to analyze the decision-making processes of warehouses and trucking companies under both the indirect source rule and the direct source rule. The findings suggest that the indirect source rule has the potential to encourage greater adoption of electric semitrucks while reducing the industry's cost burden, particularly when mitigation fees for diesel-truck trips are kept low. However, the current practice of using mitigation fees to subsidize electric-semitruck investments for trucking companies may backfire and potentially hinder electric-truck adoption.

The study further shows that factors such as competition in the trucking sector and regional characteristics, including truck-trip distance distributions, can either exacerbate or alleviate these effects. From a managerial and policy perspective, the results indicate that indirect source rules can be effective tools for emissions control. However, governments must carefully consider potential unintended consequences when designing and implementing such policies. The study also contextualizes its findings using real warehouse data from Southern California.

Topic L09: Wang, T., Atasu, A., & Kurtuluş, M. (2012). A multiordering newsvendor model with dynamic forecast evolution. Manufacturing & Service Operations Management, 14(3), 472-484.

The study considers a newsvendor who dynamically updates the forecast of market demand over a finite planning horizon. The forecast evolves according to the martingale model of forecast evolution. The newsvendor can place multiple orders over time, with ordering costs increasing as the demand realization approaches the end of the planning horizon.

In this setting, the study examines the trade-off between improving demand information and incurring higher ordering costs. It shows that the optimal ordering policy is a state-dependent base-stock policy. The study further analytically characterizes how the base-stock level depends on the information state: the relationship is linear under additive MMFE and log-linear under multiplicative MMFE. The study also analyzes a benchmark model in which the newsvendor is restricted to ordering only once. By comparing the multi-ordering and single-ordering models, it quantifies the impact of allowing multiple orders on the newsvendor's expected profit and risk exposure.

Topic L10: Forel, A., & Grunow, M. (2023). Dynamic stochastic lot sizing with forecast evolution in rolling-horizon planning. Production and Operations Management, 32(2), 449-468.

The study addresses the gap between academic stochastic lot-sizing approaches and industry practice. Although academic models often explicitly account for demand uncertainty, such approaches are seldom used in practice. Instead, firms typically implement deterministic lot-sizing models and handle uncertainty through rolling-horizon planning with frequent forecast updates.

To bridge this gap, the study proposes a stochastic lot-sizing methodology adapted to rolling-horizon planning processes. Using the martingale model of forecast evolution, the approach anticipates future forecast updates within the stochastic lot-sizing framework. The formulation is further extended with production recourse in order to reflect the replanning flexibility available in rolling-horizon planning. Extensive simulations using both synthetic and real-world data demonstrate the value of forecast-evolution models. The results show that such models reduce actual costs by 14% on average compared with traditional deterministic planning. The benefit of the extended model with production recourse depends on several factors, including capacity,

correlation, and uncertainty. Sensitivity analyses show that recourse can reduce costs by an additional 3% on average and by up to 10% in specific settings. Using both real-world and synthetic data, the study provides the first analysis of the value of additive and multiplicative MMFE-based planning models when the true forecast-evolution process is unknown. Contrary to the existing consensus, the results show that the additive model performs more robustly than the multiplicative model across a wide range of problem settings.