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

UNIVERSITY OF MANNHEIM
Business School

OPM 701 – Research Seminar Supply Chain Management
HWS 2021

General Information:

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

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 his/her 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 **Monday 7 June, 13:45 CEST** in **BWL-ZOOM-02**. During this meeting, general guidelines for conducting a scientific literature review will be discussed and the deliverables of the seminar will be explained in detail.

   BWL-ZOOM-02
   https://uni-mannheim.zoom.us/j/8684542083?pwd=UGt5MFB1dko2cWtmL2c2Q2FYWkFldz09

4. The written reports have to be **submitted** electronically and as a hard copy in two-fold by **Monday 8 November**.

5. The **presentations** will be held as a blocked session, most probably on **18 and 19 November** (exact times and room to be announced, might still be subject to change!).

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 2021, this includes the following seminars:

   - OPM 701: Research Seminar Supply Chain Management (Chair of Logistics and Supply Chain Management), **topics labeled with 'L'**
   - OPM 761: Research Seminar Production Management (Chair of Production Management), **topics labeled with 'P'**
   - OPM 781: Research Seminar Service Operations (Chair of Service Operations Management), **topics labeled with 'S'**
   - OPM 791: Research Seminar Procurement (Chair of Procurement), **topics labeled with 'B'**
   - OPM 7xx: Research Seminar Operations Management (Assistant Professorship of Operations Management), **topics labeled with 'O'**

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 **3 – 16 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 logistics@bwl.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 have to send documents to them as well.

11. For any questions concerning the seminar, feel free to contact Katrin Waßmuth at katrin.wassmuth@bwl.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.

<u>Topic L01:</u> Chen, B. & Chao, X. (2020). Dynamic Inventory Control with Stockout Substitution and Demand Learning. Management Science 66(11), 5108-5127.

We consider an inventory control problem with multiple products and stockout substitution. The firm knows neither the primary demand distribution for each product nor the customers' substitution probabilities between products a priori, and it needs to learn such information from sales data on the fly. One challenge in this problem is that the firm cannot distinguish between primary demand and substitution (overflow) demand from the sales data of any product, and lost sales are not observable. To circumvent these difficulties, we construct learning stages with each stage consisting of a cyclic exploration scheme and a benchmark exploration interval. The benchmark interval allows us to isolate the primary demand information from the sales data, and then this information is used against the sales data from the cyclic exploration intervals to estimate substitution probabilities. Because raising the inventory level helps obtain primary demand information but hinders substitution demand information, inventory decisions have to be carefully balanced to learn them together. We show that our learning algorithm admits a worst-case regret rate that (almost) matches the theoretical lower bound, and numerical experiments demonstrate that the algorithm performs very well.

<u>Topic L02:</u> Oroojlooyjadid, A., Nazari, M., Snyder, L.V., & Takáč M. (2021). A Deep Q-Network for the Beer Game: Deep Reinforcement Learning for Inventory Optimization. Manufacturing & Service Operations Management.

Problem definition: The beer game is widely used in supply chain management classes to demonstrate the bullwhip effect and the importance of supply chain coordination. The game is a decentralized, multiagent, cooperative problem that can be modeled as a serial supply chain network in which agents choose order quantities while cooperatively attempting to minimize the network's total cost, although each agent only observes local information. Academic/practical relevance: Under some conditions, a base-stock replenishment policy is optimal. However, in a decentralized supply chain in which some agents act irrationally, there is no known optimal policy for an agent wishing to act optimally. Methodology: We propose a deep reinforcement learning (RL) algorithm to play the beer game. Our algorithm makes no assumptions about costs or other settings. As with any deep RL algorithm, training is computationally intensive, but once trained, the algorithm executes in real time. We propose a transfer-learning approach so that training performed for one agent can be adapted quickly for other agents and settings. Results: When playing with teammates who follow a base-stock policy, our algorithm obtains near-optimal order quantities. More important, it performs significantly better than a base-stock policy when other agents use a more realistic model of human ordering behavior. We observe similar results using a real-world data set. Sensitivity analysis shows that a trained model is robust to changes in the cost coefficients. Finally, applying transfer learning reduces the training time by one order of magnitude. Managerial implications: This paper shows how artificial intelligence can be applied to inventory optimization. Our approach can be extended to other supply chain optimization problems, especially those in which supply chain partners act in irrational or unpredictable ways. Our RL agent has been integrated into a new online beer game, which has been played more than 17,000 times by more than 4,000 people.

<u>Topic L03:</u> Alp, O. & Şen, A. (2021). Delegation of Stocking Decisions Under Asymmetric Demand Information. Manufacturing & Service Operations Management 23(1), 55-69.

Problem definition: We consider the incentive design problem of a retailer that delegates stocking decisions to its store managers who are privately informed about local demand. Academic/practical relevance: Shortages are highly costly in retail, but are less of a concern for store managers, as their exact amounts are usually not recorded. In order to align incentives and attain desired service levels, retailers need to design mechanisms in the absence of information on shortage quantities. Methodology: The headquarters knows that the underlying demand process at a store is one of J possible Wiener processes, whereas the store manager knows the specific process. The store manager creates a single order before each period. The headquarters uses an incentive scheme that is based on the end-of-period leftover inventory and on a stock-out occasion at a prespecified inspection time before the end of a period. The problem for the headquarters is to determine the inspection time and the significance of a stock-out relative to leftover inventory in evaluating the performance of the store manager. We formulate the problem as a constrained nonlinear optimization problem in the single period setting and a dynamic program in the multiperiod setting. Results: We show that the proposed "early inspection" scheme leads to perfect alignment when J equals two under mild conditions. In more general cases, we show that the scheme performs strictly better than inspecting stock-outs at the end and achieves near-perfect alignment. Our numerical experiments, using both synthetic and real data, reveal that this scheme can lead to considerable cost reductions. Managerial implications: Stock-out-related measures are typically not included in store managers' performance scorecards in retail. We propose a novel, easy, and practical performance measurement scheme that does not depend on the actual amount of shortages. This new scheme incentivizes the store managers to use their private information in the retailer's best interest and clearly outperforms centralized ordering systems that are common practice.

<u>Topic L04:</u> Gopalakrishnan, S., Granot, D., Granot, F., Sošić, G., & Cui, H. (2020). Incentives and Emission Responsibility Allocation in Supply Chains. Management Science.

Because greenhouse-gas (GHG) emissions from the supply chains of just the 2,500 largest global corporations account for more than 20% of global emissions, rationalizing emissions in supply chains could make an important contribution toward meeting the global CO2 emission- reduction targets agreed upon in the 2015 Paris Climate Agreement. Accordingly, in this paper, we consider supply chains with joint production of GHG emissions, operating under either a carbon-tax regime, wherein a regulator levies a penalty on the emissions generated by the firms in the supply chain, or an internal carbon-pricing scheme. Supply chain leaders, such as Walmart, are assumed to be environmentally motivated to induce their suppliers to abate their emissions. We adopt a cooperative game theory methodology to derive a footprint-balanced scheme for reapportioning the total carbon emissions amongst the firms in the supply chain. This emission responsibility allocation scheme, which is the Shapley value of an associated cooperative game, is shown to have several desirable characteristics. In particular, (i) it is transparent and easy to compute; (ii) when the abatement-cost functions of the firms are private information, it incentivizes suppliers to exert pollution-abatement efforts that, among all footprint-balanced allocation schemes, minimize the maximum deviation from the socially optimal pollution level; and (iii) the Shapley value is the unique allocation mechanism satisfying certain contextually desirable properties.

Topic L05: **Angelelli, E., Archetti, C., Filippi, C., & Vindigni, M. (2017). The probabilistic orienteering problem. Computers & Operations Research 81, 269-281.**

The probabilistic orienteering problem (POP) is defined on a directed graph where a cost is associated with each arc and a prize is associated with each node. Moreover, each node will be available for visit only with a certain probability. A server starts from a fixed origin, has a given budget to visit a subset of nodes, and ends at a fixed destination. In a first stage, a node subset has to be selected and a corresponding a priori path has to be determined such that the server can visit all nodes in the subset and reach the destination without exceeding the budget. The list of available nodes in the subset is then revealed. In a second stage, the server follows the a priori path by skipping the absent nodes. The POP consists in determining a first-stage solution that maximizes the expected profit of the second-stage path, where the expected profit is the difference between the expected total prize and the expected total cost. We discuss the relevance of the problem and formulate it as a linear integer stochastic problem. We develop a branch-and-cut approach for the POP and several matheuristic methods, corresponding to different strategies to reduce the search space of the exact method. Extensive computational tests on instances with up to 100 nodes show the effectiveness of the exact method and the efficiency of the matheuristics in finding high quality solutions in a few minutes. Moreover, we provide an extended analysis on a subset of instances to show the value of explicitly modeling the stochastic information in the problem formulation.

Topic L06: **Ulmer, M.W. (2020). Dynamic pricing and routing for same-day delivery. Transportation Science 54(4), 1016-1033.**

An increasing number of e-commerce retailers offers same-day delivery. To deliver the ordered goods, providers dynamically dispatch a fleet of vehicles transporting the goods from the warehouse to the customers. In many cases, retailers offer different delivery deadline options, from four-hour delivery up to next-hour delivery. Due to the deadlines, vehicles often only deliver a few orders per trip. The overall number of served orders within the delivery horizon is small and the revenue low. As a result, many companies currently struggle to conduct same-day delivery cost-efficiently. In this paper, we show how dynamic pricing is able to substantially increase both revenue and the number of customers we are able to serve the same day. To this end, we present an anticipatory pricing and routing policy (APRP) method that incentivizes customers to select delivery deadline options efficiently for the fleet to fulfill. This maintains the fleet's flexibility to serve more future orders. We model the respective pricing and routing problem as a Markov decision process (MDP). To apply APRP, the state-dependent opportunity costs per customer and option are required. To this end, we use a guided offline value function approximation (VFA) based on state space aggregation. The VFA approximates the opportunity cost for every state and delivery option with respect to the fleet's flexibility. As an offline method, APRP is able to determine suitable prices instantly when a customer orders. In an extensive computational study, we compare APRP with a policy based on fixed prices and with conventional temporal and geographical pricing policies. APRP outperforms the benchmark policies significantly, leading to both a higher revenue and more customers served the same day.

Topic L07: **Li, F., Fan, Z.-P., Cao, B.-B., & Li, X. (2021). Logistics Service Mode Selection for Last Mile Delivery: An Analysis Method Considering Customer Utility and Delivery Service Cost. Sustainability 13(1), 284.**

Last mile delivery is an important part in the logistics service process of express enterprises since it can directly contact with the customer and affect customer satisfaction. How to select a suitable logistics service mode for last mile delivery with the objectives of lower delivery service cost and higher customer satisfaction is a noteworthy research topic. In this paper, we focus on the analysis method for logistics service mode selection for last mile delivery considering customer utility and delivery service cost. First, we conduct the market survey of customer needs and discuss the market segmentation, and then we propose a customer utility value calculation model based on utility theory. Next, we propose a delivery quantity prediction method based on the time series prediction and customer selection probability calculation. Furthermore, we construct a cost accounting model to determine the delivery service cost. On this basis, we show the selection of the suitable logistics service mode for last mile delivery according to the analysis results of customer utility and delivery service cost. Finally, we show the feasibility and effectiveness of the proposed method by a case analysis.

Topic L08: **Zhou, L., Lin, Y., Wang, X., & Zhou, F. (2019). Model and algorithm for bilevel multisized terminal location-routing problem for the last mile delivery. International Transactions in Operational Research 26(1), 131-156.**

The last mile delivery is regarded as one of the most expensive but least efficient stretches in the business-to-customer supply chain. Designing the last mile delivery system in a lean way is crucial to serve customers efficiently and economically. To address this issue, we propose a bilevel multisized terminal location-routing problem (BL-MSTLRP) with simultaneous home delivery and customer's pickup services. The solution method is proposed by combining genetic algorithm (GA) and simulated annealing (SA), called self-adaptive SGA. Studies for designing the last mile delivery system in a real-world environment indicate the validity of the proposed model based on the comparison of different scenarios. Numerical experiments are also conducted to evaluate the performance of the presented SGA. Computational results show that the hybrid approach efficiently solves the BL-MSTLRP.

Topic L09: **Trochu, J., Chaabane, A., & Ouhimmou, M. (2018). Reverse logistics network redesign under uncertainty for wood waste in the CRD industry. Resources, Conservation & Recycling 128, 32-47.**

This paper addresses the reverse logistics network (RLN) design problem under environmental policies targeting recycled wood materials from the construction, renovation and demolition (CRD) industry. The main objective is to determine the location and the capacities of the sorting facilities to ensure compliance with the new regulation and prevent the wood from being massively landfilled. We formulated the problem as a mixed-integer linear programming model (MILP) to minimise the total cost of the wood recycling process collected from CRD sites. The main contribution lies in the consideration of important uncertain factors such as supply sources locations, the available quantity of recycled wood at the collection sites, and the various quality grades of the collected wood. A scenario-based analysis is conducted to evaluate the impact of uncertainties on the RLN design. In addition, the proposed MILP model has been applied for a case study in the CRD industry within the province of Quebec, Canada. The results of this study show the adjustment of the reverse logistics network leads to the reduction of wood recycling cost due to the improved efficiency of sorting facilities and the economy of scale achieved under the new policy. Moreover, sorting facilities are now located near the CRD collection points and not close to landfilling site as for the actual situation. Finally, the study demonstrates that efforts to obtain accurate information about the supply sources locations and the expected wood quantity recovered from sorting facilities will guarantee a more efficient RLN redesign.

Topic L10: **Wang, L., Cai, G., Tsay, A.A., & Vakharia, A.J. (2017). Design of the Reverse Channel for Remanufacturing: Must Profit-Maximization Harm the Environment? Production and Operations Management 26(8), 1585-1603.**

A key attribute of a remanufacturing strategy is the division of labor in the reverse channel, especially whether remanufacturing is performed in-house or outsourced. We investigate this decision for a retailer who accepts returns of a remanufacturable product. Our formulation considers the cost structures of the two strategies, uncertainty in the input quality of the collected/returned used products, consumer willingness-to-pay for remanufactured product, the extent to which the remanufactured product cannibalizes demand for a new product, and the power structure in the channel. For the profit-maximizing retailer, the differentials in variable remanufacturing costs drive strategy choice, and higher fixed costs of in-house remanufacturing favors outsourcing. The variable remanufacturing costs and the balance of power in the prospective outsourced reverse channel are the key drivers of environmental impact, as measured by the retailer's propensity to remanufacture. While profitability and environmental goals often conflict, they align under certain conditions. These include (a) the third party has less bargaining power; or (b) the fixed cost for in-house remanufacturing is relatively high. All else equal, when remanufacturing is outsourced, the environment fares better if the third party has leadership power. We generalize to the cases when remanufacturing achieves a quality level less than "good-as-new" and when used items have non-zero salvage value. Analysis of these extensions illuminates how relative power in the reverse channel drives the firms' preferences, as well as the end customers' consumption experience.

Topic L11: **Mak, H.-Y., Dai, T., & Tang, C.S. (2021). Managing Two-Dose COVID-19 Vaccine Rollouts with Limited Supply. Available at SSRN.**

We consider the problem of rolling out an approved COVID-19 vaccine (e.g., the one developed by Pfizer) that requires two doses to be spaced apart by a fixed time interval. Facing severely limited supply and mounting pressure to rapidly achieve herd immunity through vaccination, the U.K. and U.S. governments, along with the E.U., face the decision of whether to reserve stocks for the required second doses for returning recipients. We model the vaccine rollout process and complement it with an SIR (susceptible-infected-recovered) model that captures the disease transmission process. For a two-dose vaccine, we model the inventory dynamics under three stocking policies: (1) holding back second doses, (2) releasing second doses, and (3) stretching the lead time between doses. Counterintuitively, we show that even if one intends to release all second doses, only less than half of the available doses can be allocated to first-dose appointments in order to avoid delays in administering the required second doses for returning recipients, no matter how quickly supply grows over time. We show analytically that under a (weakly) increasing vaccine production rate, releasing second doses reduces the number of infections but creates uneven vaccination patterns. Contrary to popular belief that releasing the required second doses will reduce infections significantly, our numerical results reveal the reduction in infections is quite modest. Compared with the policy of releasing second doses, stretching the between-dose lead time flattens the infection curve but leads to a higher total case count. Finally, we consider an alternative single-dose vaccine with a lower overall efficacy (e.g., the Johnson & Johnson vaccine) and show it can be more effective than its two-dose counterparts in slowing down infections and ending the pandemic. As more than one hundred countries have yet to start their vaccination efforts as of February 2021, our research has important policy implications for various national and local governments to develop their vaccine rollout strategies especially when their supply of vaccine will be likely to be limited initially.

Topic L12: **Savaskan, R.C., Bhattacharya, S., & Van Wassenhove, L.N. (2004). Closed-Loop Supply Chain Models with Product Remanufacturing. Management Science 50(2), 133-279.**

The importance of remanufacturing used products into new ones has been widely recognized in the literature and in practice. In this paper, we address the problem of choosing the appropriate reverse channel structure for the collection of used products from customers. Specifically, we consider a manufacturer who has three options for collecting such products: (1) she can collect them herself directly from the customers, (2) she can provide suitable incentives to an existing retailer (who already has a distribution channel) to induce the collection, or (3) she can subcontract the collection activity to a third party. Based on our observations in the industry, we model the three options described above as decentralized decision-making systems with the manufacturer being the Stackelberg leader. When considering decentralized channels, we find that ceteris paribus, the agent, who is closer to the customer (i.e., the retailer), is the most effective undertaker of product collection activity for the manufacturer. In addition, we show that simple coordination mechanisms can be designed such that the collection effort of the retailer and the supply chain profits are attained at the same level as in a centrally coordinated system.

Topic L13: **Amiri-Aref, M., Klibi, W., & Babai, M.Z. (2018). The multi-sourcing location inventory problem with stochastic demand. European Journal of Operational Research 266(1), 72-87.**
This paper deals with a multi-period location-inventory optimization problem in a multi-echelon supply chain network characterized by an uncertain demand and a multi-sourcing feature. The aim of the paper is to propose a generic modeling approach to integrate key features of the inventory planning decisions, made under a reorder point order-up-to-level (s, S) policy, with the location-allocation design decisions to cope with demand uncertainty. Given the hierarchical structure of the problem, a two-stage stochastic mathematical model that maximizes the total expected supply chain network profit is proposed. This optimization model is intractable due to its non-linearity. Therefore, a linear approximation is proposed and a sample average approximation approach is used to produce near-optimal solutions. Numerical experiments are conducted to validate the proposed modeling and solution approaches. The results show the efficiency of the linear approximation of the (s, S) policy at the strategic level to produce robust design solutions under uncertainty. They underline the sensitivity of the design solution to the demand type and the impact of the inventory holding costs and backorder costs, especially under non-stationary processes.

Topic L14: **Thevenin, S., Adulyasak, Y., & Cordeau, J.F. (2021). Material Requirements Planning Under Demand Uncertainty Using Stochastic Optimization. Production and Operations Management 30(2), 475–493.**
Material Requirements Planning (MRP), a core component of enterprise resource planning (ERP) systems, is widely used by manufacturers to determine the production lot sizes of components. These lot sizes are typically computed based on deterministic and dynamic demand assumptions, while safety stocks, which hedge against demand uncertainty, are determined independently based on different assumptions. As the lot sizes and safety stocks are not determined simultaneously, sub-optimal decisions are often used in practice. The critical impact of inventories and service levels in manufacturing motivates the study of stochastic optimization methods for MRP. In this study, we investigate stochastic optimization methods for MRP systems under demand uncertainty. A two-stage and a multi-stage model are proposed to deal with the static-static and static-dynamic decision frameworks, respectively. We first derive structural properties of the two-stage and multi-stage models to provide insights on the differences between the plans created with these two models. As multi-stage stochastic programs are not convenient in real-world applications, several practical enhancements are proposed. First, to address scalability issues, we employ heuristics in combination with advanced sampling methods. Second, to allow real-time static-dynamic decisions, we derive a policy from the solution of the multi-stage model. Third, to deal with the dynamic-dynamic decision framework, we employ a rolling horizon implementation. The effectiveness and performance of stochastic optimization for MRP are validated by numerical experiments, which demonstrate that the stochastic optimization approaches have the potential to generate significant cost savings compared to traditional methods for production planning and safety stocks determination.