UNIVERSITY OF MANNHEIM
Business School

CHAIR OF SERVICE
OPERATIONS MANAGEMENT
Prof. Dr. Cornelia Schön

Master Thesis Proposal

Smart Predict-Then-Optimize for the Newsvendor Problem: A Direct Loss Minimization Approach

In operational decision-making under uncertainty, a common approach is to first build a predictive model for unknown inputs—such as future demand—and then optimize decisions using the predicted values. This sequential method, known as Predict-Then-Optimize (PTO), is widely used in inventory management, pricing, logistics, and workforce planning. However, PTO models are typically trained using statistical loss functions like Mean Squared Error (MSE) or Mean Absolute Error (MAE), which are not aligned with the actual costs of the downstream decision. As a result, such models may yield accurate forecasts while still leading to poor or costly decisions. This issue is particularly evident in the single-product newsvendor problem, where the goal is to decide how much inventory to order in advance of observing uncertain demand. Due to the asymmetric costs of understocking (lost sales) and overstocking (excess inventory), the optimal decision is based not on average demand but on a specific quantile of the demand distribution. Traditional PTO methods often ignore this asymmetry, resulting in suboptimal inventory levels. To address this, the Smart Predict-Then-Optimize (SPO) framework—introduced by Elmachtoub and Grigas (2022)—proposes aligning the training loss of the prediction model with the actual decision objective. In the context of the newsvendor problem, this leads to minimizing the quantile loss, which directly reflects inventory costs.

This thesis investigates how minimizing the SPO loss affects inventory performance compared to traditional PTO models. The evaluation considers decision quality, predictive accuracy, and robustness under varying cost structures.

The objectives of this thesis are to:

- Conduct a literature review on the PTO framework and its limitations, with a focus on integrated prediction–optimization methods and the SPO approach.
- Formulate and solve the SPO-based training problem for the single-product newsvendor, using a linear demand model and quantile loss minimization.
- Implement a baseline PTO approach by training a linear regression model using MSE and applying the critical fractile formula to determine optimal order quantities.
- Evaluate both approaches using synthetic or real-world datasets, comparing decision cost, prediction accuracy (e.g., RMSE), and sensitivity to cost asymmetry between underage and overage costs.

Requirements

- OPM 781
- Profound knowledge in Operations Research

- Excellent analytical skills and an ability to transform real-world business problems into Operations Research models

Administrative information for writing a master thesis at the Chair of Service Operations Management can be found here.

Selected Literature Recommendations

Ban, G. Y., & Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. Operations Research, 67(1), 90-108.

Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize". Management Science, 68(1), 9-26.

Liu, C., Letchford, A. N., & Svetunkov, I. (2022). Newsvendor problems: An integrated method for estimation and optimisation. European Journal of Operational Research, 300(2), 590-601.

Qi, M., & Shen, Z. J. (2022). Integrating prediction/estimation and optimization with applications in operations management. In Tutorials in operations research: emerging and impactful topics in operations (pp. 36-58). INFORMS.