MS-E2177
Seminar on Case Studies in Operations Research

Data-Driven Optimization of Used Car Inventory

Final Report

Team:
Juuso Varho (Project Manager)
Tuomas Mäkelä
Patience Anipa
Linh Duong
Juuso Seuri
Client:
K-Auto

June 5, 2025

Contents

1 Introduction

1.1 Background

Kesko is a leading player in the Finnish retail sector, operating across grocery, building and technical trade, and automotive markets. Within its automotive division, K-Auto imports and sells premium new cars and offers a substantial selection of multi-brand used cars across Finland. In recent years, the used car segment has become increasingly vital to the automotive industry, with nearly 600,000 used passenger cars sold in 2024—significantly outpacing new car sales of around 75,000 (Tiedotuskeskus (2024).

This strong demand highlights the importance of effective inventory management. Dealerships must maintain a well-balanced stock of vehicles to avoid missing sales opportunities or investing in cars that take longer to sell. However, many purchasing decisions are still driven by intuition and personal experience rather than data-driven approaches. This can lead to mismatches between inventory and demand, inefficiencies in operations, and lost revenue.

To address these challenges, this project explores how a data-driven model can optimize inventory composition in the used car market. By analyzing historical sales data and vehicle attributes, the model aims to improve the match between supply and customer demand, reduce excess inventory, and support more informed purchasing decisions. This optimization is very complex in the used car market due to the nature of vehicles with different make (manufacturer), model, age, mileage and condition. These complexities make predictive modeling and optimization both challenging and valuable.

1.2 Objectives

The goal of this project is to develop a data-driven inventory optimization model that enhances the performance of K-Auto's used car operations. The model is intended to support purchasing managers by providing clear, actionable recommendations for which vehicles to acquire or avoid, based on market data and expected returns.

The specific objectives of the project are:

- **Optimizing Inventory Composition:** Build a recommendation model that identifies the ideal mix of used vehicles for a rolling three-month planning horizon, updated monthly.
- **Maximizing Return on Capital Employed (ROCE):** Prioritize vehicle acquisitions based on their potential to generate high returns, improving capital efficiency.

- **Enhancing Decision-Making for Purchases:** Deliver interpretable, data-driven justifications for including or excluding specific vehicles in the inventory, incorporating historical sales performance and demand signals.

By focusing on these goals, the project aims to deliver a robust, practical tool for K-Auto that improves inventory planning accuracy, reduces holding costs, and strengthens the overall responsiveness of the business to evolving market conditions.

2 Literature Review

The used car market in Europe has become increasingly complex and competitive, shaped by shifting consumer preferences, macroeconomic uncertainty, and fluctuating supply-demand dynamics (eCarsTrade (2023)). In this rapidly evolving environment, reliance on traditional decision-making — often grounded in internal expertise, heuristics, or instinct — is proving inadequate for sustaining profitability and operational efficiency. As highlighted by Ellencweig et al. (2023) in a McKinsey & Company analysis, many stakeholders across the automotive value chain—including dealerships, OEMs, and finance providers—have yet to fully capitalize on the potential of data-driven decision-making. One notable area where analytics is generating substantial value is vehicle allocation. Large dealership networks are now leveraging real-time pricing and demand signals to reallocate inventory across regions, optimizing margins by factoring in regional price differentials, local demand variability, and inventory holding periods.

Parallel to these developments, the study of inventory optimization has significantly evolved due to greater supply chain complexity and the proliferation of operational data. In the automotive industry, inventory is both diverse and strategically vital, as it directly influences time-to-market and overall competitiveness. Inefficient stock management ties up significant financial resources and incurs substantial costs, highlighting the importance of optimizing inventory to maintain profitability, operational efficiency and competitive advantage (Saliji (2021)). An inventory model primarily aims to determine the optimal quantity of goods to order and the appropriate timing for those orders, with traditional approaches typically based on deterministic demand models (Babiloni and Guijarro (2020), Stopková et al. (2019), Tamjidzad and Mirmohammadi (2017), Tamjidzad and Mirmohammadi (2015), Choi (2014)). The applicability of such models in real-world scenarios—particularly in the used car market—has become increasingly limited due to numerous constraints and growing uncertainty in demand patterns Maitra (2024). This stochastic and unpredictable business environment underscores the need for more flexible approaches that can effectively incorporate randomness and variability in key decision-making components.

To contextualize the two broad categories of inventory optimization models, it is essential to distinguish between deterministic and stochastic approaches. Deterministic models operate under the assumption of full knowledge of key parameters, such as demand and lead time. In contrast, stochastic models explicitly account for the inherent uncertainties in these variables (Zipkin, 2000). As Turgay (2023) highlights, the stochastic approach is increasingly favored in modern supply chains due to its ability to incorporate variability in demand and supply conditions. These models aim to determine optimal ordering policies by managing uncertainty and making probabilistically informed decisions.

The origins of stochastic inventory theory can be traced back to the seminal work of Arrow et al. (1951), who emphasized the role of demand uncertainty, shortage costs, and the importance of contingency planning. Foundational models such as the (s, S) inventory policy and the newsvendor model have long been used to address uncertainty in both single- and multi-period contexts (Porteus, 1990). However, these classic models often assume static conditions or single-period decisions, which limit their applicability to dynamic, real-world inventory systems. In practice, inventory management frequently involves repeated and adaptive decision-making over time.

To address these limitations, multistage stochastic programming emerges as a more robust framework, enabling decision-makers to revise plans at multiple stages based on newly observed information (Kayacık et al., 2024). This feature makes multistage models particularly well-suited to problems such as used vehicle allocation, where decisions are made on a rolling basis over a planning horizon. For example, in the case examined in this study, inventory and allocation policies must be updated monthly over a three-month horizon in response to evolving market data and operational constraints.

However, the flexibility of multistage models comes at the cost of increased computational complexity. As the number of stages and scenarios grows, traditional solution methods often become intractable. This challenge is especially evident in problems requiring trade-offs between immediate resource utilization and future uncertainty, a context directly relevant to this study. To address these challenges, Stochastic Dual Dynamic Programming (SDDP) has emerged as a prominent algorithmic solution. As described by Füllner and Rebennack (2021), SDDP is designed to efficiently approximate optimal policies in large-scale multistage problems by decomposing them into smaller, tractable subproblems and iteratively refining value function approximations via backward and forward passes (Pereira and Pinto, 1991; Downward et al., 2020).

SDDP operates within the broader framework of stochastic programming, typically modeling uncertainty through a scenario tree that branches at each stage to represent discrete realizations of future events (Füllner and Rebennack, 2021). Since its introduction in 1991, SDDP has gained widespread attention for both its theoretical rigor and practical utility. Today, it is considered one of the state-of-the-art methods for solving complex multistage stochastic problems and has been successfully applied in various domains, including inventory management (Bandarra and Guigues, 2021; Dowson et al., 2020; Guigues, 2017; Guigues et al., 2023). These recent applications demonstrate the practical relevance of SDDP in settings characterized by uncertainty and sequential decision-making. Consequently, adopting SDDP and related stochastic optimization techniques offers a compelling strategy for managing used vehicle inventories in real-world environments.

Based on the literature, data-driven approaches grounded in stochastic optimization and dynamic programming are well-positioned to address the specific challenges of used vehicle inventory management. Traditional deterministic methods, which often rely on static assumptions or expert heuristics, struggle to capture the uncertainty and heterogeneity inherent in the used car market. In contrast, multistage stochastic models—particularly when implemented using advanced algorithms like SDDP—offer the responsiveness, scalability, and adaptability needed for rolling horizon inventory planning. These insights provide a solid theoretical foundation for this study, which aims to develop a customized decision-support tool for K-Auto. By building on these methodologies and incorporating vehicle-specific attributes, market dynamics, and capital constraints, the project aims to bridge the gap between theoretical advances and practical implementation, ultimately enabling data-informed purchasing strategies in the used car market.

3 Data & Methods

The primary objective of this work is to use historical car sales data to forecast future sales for vehicles and use the forecasted sales to optimize the car dealership used car purchases. The end goal is to improve the purchasing strategy in a way that maximizes the profit and aligns the inventory with consumer preferences. In the following sections, the data and methodology for demand estimation and optimization is presented.

3.1 Data

An extensive dataset of business-to-consumer (B2C) car sales is used, consisting of sales from different vendors across Finland. The dataset covers the period from 2018 to 2025 and contains approximately 2.1 million entries. Each entry corresponds to an individual vehicle sale and includes 53 attributes describing various aspects of vehicle and sale details.

The dataset includes information on a diverse set of vehicle types, such as passenger cars, motorcycles, trucks and caravans. The attributes themselves cover wide range of categories, such as the asking price, vehicle specifications, and sales metadata such as date of sale, time in the inventory, and dealership information.

3.2 Preprocessing of data

The raw dataset is too complex and detailed to be used directly in forecasting or decision-making. Therefore, the data is categorized into larger categories, which we will refer to as elements. Each element represents a set of vehicles that share similar characteristics and are likely to serve the same type of customer need. However, the categorization can not be too broad as the purchasing managers need clear and actionable guidance on the optimal car inventory. In order to find the most meaningful features while keeping the element division actionable, purchasing managers and domain experts were interviewed. This collaboration ensured the elements would reflect both technical and commercial relevance. The elements contain so-called element features, which describe the features used for element division. In addition, each element feature can be divided into smaller categories, referred to as feature classes. Next, the data cleaning and element division are explained step-by-step.

Firstly, the data consists of a wide variety of vehicles. However, this work is only limited to analyzing the sales of passenger cars. Therefore, every other vehicle type is excluded from the dataset. The experts recognized that the most important and actionable features to know would be the make, model, model year, fuel type, engine size, and mileage. This will be used as the basis for deciding the used element features.

The vehicle make (i.e., the manufacturer) is included as one of the element features. However, the dataset contains a large number of different manufacturers. Including all of them would result in an unmanageable number of categories, so we limit the makes used as feature classes. Specifically, we include either the top 30 best-selling makes or all makes with over 5,000 units sold – whichever results in fewer categories. In our dataset, the top 30 makes all exceeded the 5,000-unit threshold, so we use these 30 as our make classes. All other makes are excluded from further analysis.

We add the car model as an element feature. Similarly to the make, not every model can be chosen as a feature class. We choose to only include a limited number of the best-selling models for each make. We only consider models that have sold over 50 units, and then choose at most the top 10 best-selling models. This means the number of considered models is at most $30 \cdot 10 = 300$, but can be less if some models have models that have not sold over the 50 unit limit.

Fuel type is another key factor in car sales and is therefore included as an element feature. While fuel type is already a categorical variable, the dataset includes several less common types. To reduce complexity, we would want to only analyze the most common categories: gasoline, diesel, hybrid, and electric. However, based on expert insight, electric vehicles differ significantly in relevant features from combustion engine vehicles. Therefore, we exclude electric cars from our analysis. Thus, the feature classes for fuel type are gasoline, diesel, and hybrid.

The engine size is an important factor to consider when purchasing used cars and is thus chosen as an element feature. The feature classes are constructed by rounding the engine size, which is given in cubic centimetres, to decilitres with 2 significant figures. This type of feature class construction is warranted as the available engine sizes vary between car models. Therefore, unnecessary feature classes, too big or too small, are not added.

Additionally, we want to filter and divide the elements in terms of mileage and model year. We are mostly focused on used cars. Therefore, we only analyze the sales of cars with model year 2024 or before. However, we don't want too old cars, so we only include those that are at most 10 years old. For the mileage, we only consider cars with less than 200000 km mileage. Overall, we only consider fairly new and low-driven used cars. We would want to add the model year and mileage also as element features, meaning, these features would be divided into categories and then the sale of each of those categories would be done separately. However, due to computational complexity constraints of our model, discussed more in Section 4.3.2, these features were not included as element features.

After the discussed element division and data filtering, we are left with around 800000 sales, which are transformed to represent some specific element. The element features now used are make, model, fuel type (B, D, or H), and engine size (in decilitres). These features now describe the most important features considered in buying a used car, keeping the descriptors broad but still actionable. After converting the sales data to this element format, the future sales of an element can be forecasted with its historical sales.

4 Model

We decided to approach the problem with the stochastic dual dynamic programming (SDDP) technique, which is well suited for multistage stochastic optimization problems where an agent has to make decisions over time under uncertainty. In our case, as we have only the past sales data available, we have to generate the stochasticity and uncertainty ourselves to generate a suitable input for the SDDP model.

First, we need to estimate future sales from the data. From the expert interviews and the initial data exploration we gathered that the car sales are seasonal, so we decided to approach the estimation with a seasonal time series model. As we saw in 3.2, we have split the data into multiple different elements, and we will thus apply the time series model for each element separately. This process will be discussed in Section 4.1 in more detail.

Secondly, the SDDP model requires us to have different states representing different scenarios for all time periods, all of which have their own likelihoods and transition probabilities between the stages. This is done with a Monte Carlo simulation that uses the time series forecasting results. Different states and transition probabilities are then obtained from the simulation results by clustering. This approach is discussed in 4.2.

Finally, we can give the SDDP model the likelihoods for each state and the transition probabilities for each state between time periods, which the model will use to solve the optimization problem. The details of the SDDP model will be discussed in 4.3.2.

Like many optimization problems, those addressed with the SDDP can be combined with an objective function and a set of constraints. Our initial task was to maximize the return on capital employed (ROCE), which we encountered to be a challenge since we do not know the purchasing price of cars. Thus, we estimate a fixed profit of 15 % for each car, and the objective of the SDDP is to maximize total profits.

4.1 Element forecast

When a decision maker wants to maximize their profits from a product, they would ideally like to be aware of the demand that faces that product. Demand for a product is the quantity sold as a function of its price. However, estimating demand is a challenging task, as it requires the decision maker to know how many units of the product they are going to sell at different price points. This requires knowing or estimating price elasticities for the customers.

Due to the complexity of demand estimation, the optimization model will be founded instead on sales forecasting. This approach involves estimating the quantity of different types of cars sold, while assuming that they are sold at the market price. As a result, our solution can be the optimal only under the perfect competition assumption. Any market power that allows the decision maker to make sales above the competitive price would mean that the solution for the profit maximization problem would have to be reached some other way. However, the perfect market assumption will not be far off the truth, as the used car market typically involves a large number of sellers and buyers, which limits the ability for individual sellers to set higher prices.

It is worth noting that there are several different approaches for predicting future sales. Often this problem can be tackled with time series forecasting. However, it has been shown that other techniques can be very efficient and precise in making predictions, e.g., using Long Short-Term Memory (LSTM) method (Ensafi et al., 2022) or XGBoost-based models (Ji et al., 2019) can be quite effective in handling nonlinearities. In order to keep the initial model and the validation phase a bit simpler, we decided to go with the time series route.

In Subsection 4.1.1 we will discuss the chosen time series model and in Subsection 4.1.2 we will validate the chosen model.

4.1.1 Forecasting model

After initial data exploration and the expert discussions, it became clear that there is seasonality in the data. Additionally, the results from Augmented Dickey-Fuller tests (ADF), both on aggregated and sampled elements, indicated that we cannot reject the null hypothesis of a unit root, suggesting non-stationarity in the data. Therefore, to make any successful forecasting models we have to difference the data.

In Figure 1, we have four different Autocorrelation Function (ACF) plots for one of the sampled elements in the data. Each value on the x-axis indicates the number of lags between the time series data point and its lagged version. The y-axis shows the autocorrelation coefficient, which measures the correlation between the time series and its lagged values at each lag.

On the left (A), there is the original non-stationary series. The decaying nature of the different lags on the x-axis indicate that there is some persisting trend with the data. In (B) we have taken the first difference, and in (C) we have taken the 12 month seasonal difference, i.e., for each value in the time series we reduce the value from a year ago. For (B) we still see some large spikes throughout the lags, and for (C) there seems to be some sort of trend, especially in the smaller lags. Finally, by combining both, we seem to get a lot more stationary time series, which makes it more suitable for

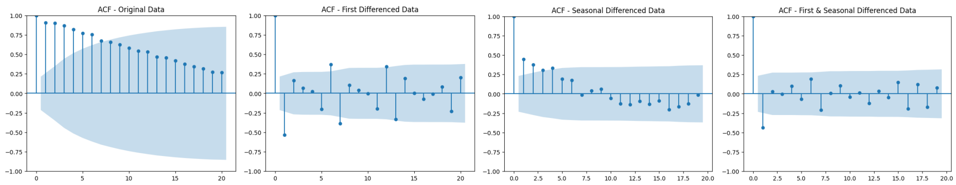

A B C D

Figure 1: Autocorrelation plots for a sampled element (BMW_3-SARJA_H_20.0)

modeling and forecasting with time series techniques.

The same form of the time series data seems to apply for other sampled elements as well. Therefore, it would make sense to use a time series model that takes into account both the first difference and the seasonal difference. A reasonable choice seems to be then the Seasonal Autoregressive Integrated Moving Average (SARIMA) model. The SARIMA model allows us to make predictions using data that has non-stationary and periodic variation.

A SARIMA model can be denoted as

$$\text{SARIMA}(p, d, q)(P, D, Q)_s \tag{1}$$

where p is the number of autoregressive (AR) terms, d is the number of differences, q is the number of moving average (MA) terms, while (P, D, Q) indicate the same respectively for the seasonal part of the model, e.g., P is the number of seasonal AR terms. The parameter s indicates the number of observations per year, for our monthly data this is 12.

The choice of parameters depends on the data. For predicting however, by choosing too many coefficients you run the risk of overfitting your model with the training data. After running tests over multiple sets of samples, we found that having a first difference and one seasonal difference in addition to one non-seasonal AR term yielded reasonable results. Additional terms were often statistically insignificant. In the SARIMA notation this means

$$\text{SARIMA}(1, 1, 0)(0, 1, 0)_{12} \tag{2}$$

The mathematical expression for the chosen model is therefore

$$(1 - \phi_1 B)(1 - B)(1 - B^{12})Y_t = \epsilon_t \tag{3}$$

where Y_t is the observed time series at time t, B is the lag operator ($BY_t = Y_{t-1}$), ϕ_1 is the non-seasonal AR coefficient, and ϵ_t is the white noise error term at time t.

4.1.2 Model validation

Figures 2 and 3 illustrate how the SARIMA model makes predictions for an eight element sample. In the Figure 2, the model predicts the next period's sales, while in the Figure 3, the model predicts the sales three months from the given point. In both Figures, the blue and green lines show the actual sales, the red line depicts the predictions from the model and the gray area around the predictions is the 95 % confidence interval. The confidence interval is always non-negative, as there cannot be a negative number of sales.

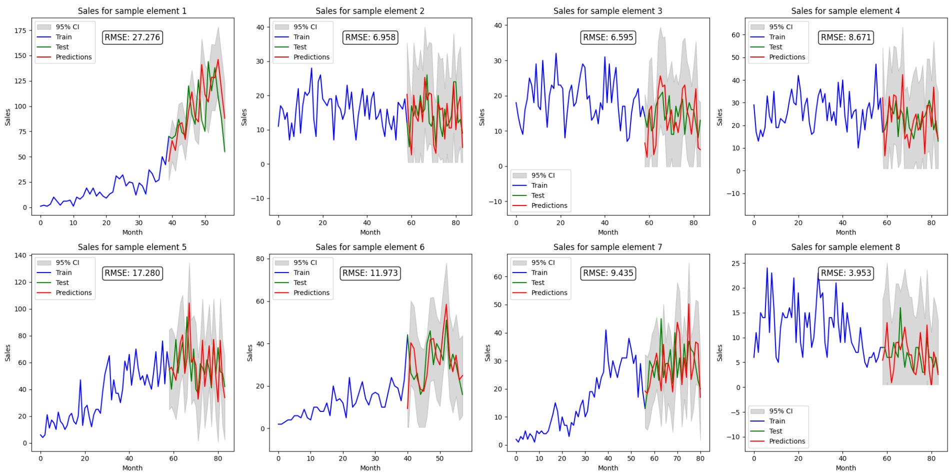

Figure 2: Sales forecasting (t+1) for sampled elements

From the figures, we observe that the prediction model behaves well around the test data, and that the confidence intervals almost always include the actual sales values, suggesting reasonably accurate predictions. However, having elements with fewer sales increases the relative size of the confidence intervals. This can be seen in the case of element 8 in the bottom right corner of Figure 3, where the wider confidence interval represents greater uncertainty in the forecast.

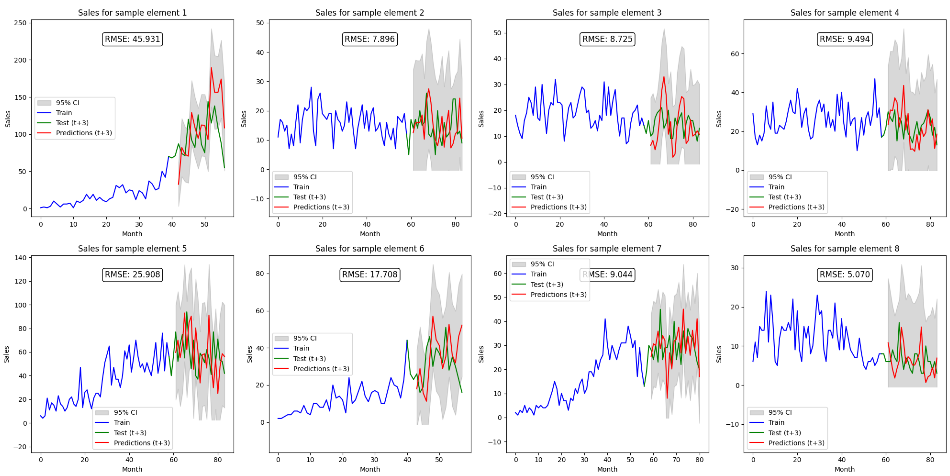

Figure 3: Sales forecasting (t+3) for sampled elements

Furthermore, we see that there is an expected decrease in performance between the two figures, as the model is having a harder time predicting further into the future. However, from the samples we have gathered we can qualitatively say that the model is making predictions to the right direction. The current predictive power of the SARIMA model should not stand in the way of testing the SDDP approach, although better predictions at the element level will inevitably lead to better solutions after optimizing with the SDDP.

For the first iteration of the model we decided to go with a simple SARIMA model. However, we acknowledge that there exists more accurate ways to pick the parameters and other ways to improve the prediction power of the model. These potential future improvements and other ways to approach this problem will be discussed in Section 6.3.

4.2 State division

After having a working time series model ready to forecast future sales, we turn to the Monte Carlo simulation phase. From fitting the SARIMA model, we get a set of residuals, ϵ_t, for each observation t in the time series data. From these residuals, we then can formulate a distribution, from which we can pull out random residuals to make predictions. Thus, the prediction for the sales for an element for the next period can be just calculated using the AR coefficient of the SARIMA model and the previous time series value, while ϵ_t gives us some noise for the predictions. Extending this procedure to all elements allows us to forecast scenarios of possible sales trajectories for the whole set of elements.

In the first forecasting stage, we run N Monte Carlo simulations for each element, generating N scenarios of future sales. Each scenario consists of predicted sales values across all elements in the dataset. Then, we cluster the scenarios from the first stage to give us K states. For this we use K-means clustering, so that each scenario belongs to the cluster with the nearest mean. Also, the clustering approach helps us to get transition probabilities as we can very easily get the probability of ending up in one of the clusters by summing up all the elements in that particular cluster and then dividing by N.

We then repeat the same for the second stage, with the exception that we make N predictions at each of the cluster centers obtained in the previous stage. Therefore, in the second stage we make $N \times K$ predictions. We proceed again by clustering and calculating the transition probabilities just like in the first stage.

Unfortunately, the chosen time series approach has its limitations as it is not completely memoryless. This is not an issue with the first and second stages as the transition probabilities can be obtained straightforwardly. However, in the final third stage the issue is that we do not how stochastic process has reached any of the states in the second stage, and therefore obtaining the transition probabilities for the final stage becomes a challenge.

We decided to use deterministic multivariate forecast using SARIMA to acquire the clusters for the second stage, from which we were able simulate the clusters of states and transition probabilities for the last stage. However, this is an approximation of the stochastic process and a limitation that we encountered when using the time series approach. The SDDP model works well with the time series approach when we have only two time periods but after that it becomes a bit harder. On the contrary, we assume that this approximation should not affect the accuracy of the model too much.

After dividing the simulations to states and calculating the transition probabilities, we can move over to using the SDDP model itself.

4.3 Optimization

4.3.1 Problem definition

In this project, the goal is to determine the number of purchased units of each car type over a rolling three-month horizon, updated monthly. The decisions must account for sales uncertainty, varying car margins, and future availability. Unlike one-shot decisions made in a static environment, our challenge involves sequential decision-making under uncertainty, where the outcomes of earlier choices directly influence future decisions. This dynamic and uncertain nature of the problem motivates the use of a multistage stochastic decision problem formulation.

Traditional models such as assortment optimization treat inventory decisions as static or single-period problems. These models may optimize for a single snapshot of time, assuming known demand distributions or ignoring the feedback between periods. While effective for short-term or deterministic decisions, they lack the flexibility needed for rolling planning under uncertainty.

On the other hand, a multistage model allows decisions to evolve over time, incorporating new information as it becomes available. At each stage (monthly decision point), the model accounts for:

- Inventory levels carried over from previous stages
- Realized sales outcomes
- Updated probabilistic knowledge of future uncertainties

This results in a dynamic model that adapts decisions based on the observed system state, unlike static models that commit to decisions upfront without adjustments.

4.3.2 Optimization model

The inventory planning task involves choosing the optimal number of used cars to buy, store, and sell each month over a rolling three-month horizon. The uncertain element is the number of future sales of each element. Given the purchasing cost, inventory holding cost, and selling price, this problem naturally fits a multistage decision framework. The SDDP-related methodological concepts outlined in this section are based on SDDP.jl Developers (2024)

In multistage stochastic optimization problems, an agent makes decisions that affect the state of the system over time. Each decision point is referred to as a *stage*. At each stage, the modelled agent chooses an action, called a *control variable*, which impacts the state of the system – that is, the current level of vehicle inventory.

Control variables are denoted by u. In our context, they include the number of sold and purchased cars at each stage. Though the number of sold cars is referred to as a control variable, in practice the key decision is made on the bought cars variable. We denote the control variables as:

- u_t^{buy}: the number of vehicles purchased at time t,
- u_t^{sell}: the number of vehicles sold at time t.

Similarly, the state of the system is tracked from stage to stage by *state variables*, denoted x. These capture the number of vehicles in inventory. The incoming value of a state variable at a stage is denoted by x, while the outgoing state is x'. In our case:

$$x_t = \text{inventory at the beginning of stage } t, \quad x'_t = x_t + u_t^{\text{buy}} - u_t^{\text{sell}}.$$

At any node i, the transition between stages is governed by a transition function:

$$x' = T_i(x, u, \omega),$$

which determines the next state x' based on the current state x, the chosen controls u, and the realized demand. Stagewise realizations of random variables are denoted by ω. The relation between control variables and the effect of the transition function on the incoming and outgoing state variables is presented in Figure 4.

The control variables are chosen via a decision rule:

$$u = \pi_i(x, \omega),$$

subject to a set of constraints $U(x, \omega)$, including:

- Cannot sell more vehicles than available,
- Cannot sell more vehicles than the realized sales,
- Cannot buy more vehicles than the maximum allowed by supply constraints,
- Total inventory must not exceed maximum capacity $c_{\max}$.

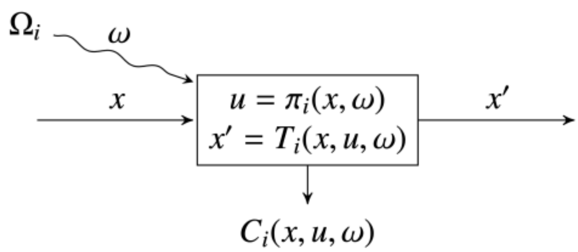

Figure 4: A simplified transition diagram for state, control, and random variables

Each stage is connected to others via a *policy graph*. In the Markovian case, the graph includes multiple nodes per stage, each with probabilistic transitions to nodes in the next stage. An example of this is visualised in Figure 5, where we have 3 stages, each with 2 states. The probabilities of the transitions between the states are defined by the transition matrix. In our

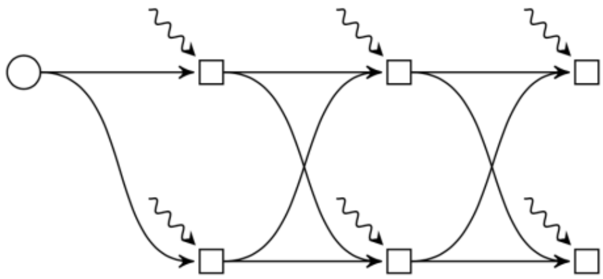

Figure 5: An illustration of a simplified Markovian policy graph with two states and 3 stages

case, the only uncertain element of the model is included via the transition matrix, meaning that the random variable ω

The optimal decision is solved by the SDDP algorithm. For a convex problem, the SDDP algorithm operates in two phases:

- **Forward pass:** Sequentially samples demand scenarios from the initial to final stage. Solves subproblems using current approximations of the cost-to-go function.
- **Backward pass:** Improves the cost-to-go approximation by adding cuts using Kelley's algorithm (Kelley, 1960). This yields a better lower bound for the cost function.

4.3.3 Mathematical Formulation

Let N denote the number of vehicle types. For each vehicle type $i \in \{1, \ldots, N\}$ and time t, we define:

$$\begin{aligned} x_i &: \text{Inventory level (state variable)}, \\ u_i^{\text{buy}} &: \text{Vehicles purchased (control variable)}, \\ u_i^{\text{sell}} &: \text{Vehicles sold (control variable)}, \\ \omega_i &: \text{Sales realization (random variable)}. \end{aligned}$$

The objective of the model is to maximize the generated profit across all stages. Thus, the general formulation of the model can be presented as:

$$
\begin{aligned}
\max_{u_i^{\text{buy}}, u_i^{\text{sell}}, x_i'} \quad & \sum_{i=1}^{N} \left(p_i^{\text{sell}} u_i^{\text{sell}} - p_i^{\text{buy}} u_i^{\text{buy}} - c_i^{\text{inv}} x_i' \right) \\
\text{s.t.} \quad & x_i' = x_i + u_i^{\text{buy}} - u_i^{\text{sell}} && \text{(inventory update)} \\
& u_i^{\text{sell}} \leq \omega_i && \text{(cannot exceed demand)} \\
& u_i^{\text{sell}} \leq x_i && \text{(cannot sell more than in stock)} \\
& u_i^{\text{buy}} \leq s_i^{\max} && \text{(supply constraint)} \\
& \sum_{i=1}^{T} x_i' \leq c_{\max} && \text{(total inventory constraint)} \\
& u_i^{\text{buy}},\, u_i^{\text{sell}} \in \mathbb{Z}_+, && \text{(non-negative integers)}
\end{aligned}
$$

where p_i^{sell} and p_i^{buy} represent the selling and purchase price of each car type, c_i^{inv} represents the inventory cost, c_{max} represents the maximum storage capacity, and ω_i represents the demand for element i.

4.3.4 Model Dimensions and Complexity

- **Number of decision variables:** $2N$ per stage (buy/sell per vehicle type)
- **Number of state variables:** N per stage
- **Number of constraints:** $4N + 1$ per stage

The complexity of the model grows linearly with the number of stages and vehicle types. However, the Markov transition matrix introduces exponential growth in the number of possible paths. This makes the SDDP algorithm particularly suitable, as it approximates the value function through sampled scenarios and cutting planes.

Implementation

This model is implemented using `SDDP.jl`, a Julia package for solving multistage stochastic optimization problems using dual dynamic programming. The Markovian policy graph structure allows for stage-specific demand distributions, enhancing realism in the vehicle inventory context.

5 Results and model validation

In this section we discuss the results of the inventory optimization model. We compare the recommendations for bought vehicles of our inventory optimization model. Furthermore, we include two heuristic models in the comparison, to validate the performance of our model against a potential real-life simplistic strategy.

We ran our model with 5 different setups. These contained different parameters for the inventory cost. Furthermore, we included two additional runs: one with a continuous set of decision variables and another with a dynamically constrained maximum purchased vehicles. The aim of the continuous run was to examine the performance for a potentially less computationally expensive setup. For the dynamic model, the objective was to guide the model towards feasible values while allowing the model to overbuy if it sees a potentially profitable situation. The setups are summarized below:

- Integer model: Highest inventory cost (1200 €/car)
- Integer model: High inventory cost (900 €/car)
- Integer model: Medium inventory cost (600 €/car)
- Integer model: Low inventory cost (500 €/car)
- Integer model: Dynamic maximum (calculated by $max(10, 1.5 \cdot \text{maximum monthly historical sales})$
- Continuous model: Medium inventory cost (600 €/car)

The optimal buying strategy was evaluated by extracting the optimal policy at the node for time T. The model was solved in Triton, using the Gurobi solver set to use the Dual Simplex method. The maximum number of iterations was restrained to 7000 iterations, to keep the solution times and computational resource usage at acceptable limits. Each solution terminated at the maximum amount of iterations. Thus, the solution is not necessarily a global optimum but instead a "good-enough" approximation.

For the comparison, we construct two alternative strategies based on a simple heuristic. These strategies purchase the 50 most sold car elements in proportion to their sold amount across the entire history. One of the heuristic models is set to purchase the approximately the same number of units as the SDDP model. This is the realistic variant, as we have a fixed inventory capacity, which the SDDP models also adhere to. The second heuristic model is set to use the same amount of approximate capital. Thus, it can buy as many units as possible, even though it might result in a solution higher than the maximum inventory capacity. These models will be referred to as

- Heuristic model: same units
- Heuristic model: same capital

5.1 SDDP and heuristic model Comparison

To compare the performance of the SDDP and heuristic models, we purchase the number of each element at time T. We then evaluate the cumulative profit for each strategy using historical sales data. Furthermore, we compare the ROCE and number of vehicles purchased, as well as discuss some reasons behind the differences in performance.

The SDDP model, with the dynamic maximum number of purchased vehicles, achieves the highest cumulative potential profit over a 3-month horizon (Fig. 6), outperforming all the other models by over 1.2 M€. In a one month period, the heuristic strategy outperforms the SDDP models. This is expected as the turnover of these cars is higher and likely more constant. However, the SDDP models all outperform the heuristic model with the same units across the entire time period. The performance of each of these runs is similar for all variants, converging to a value around approximately 12 M€.

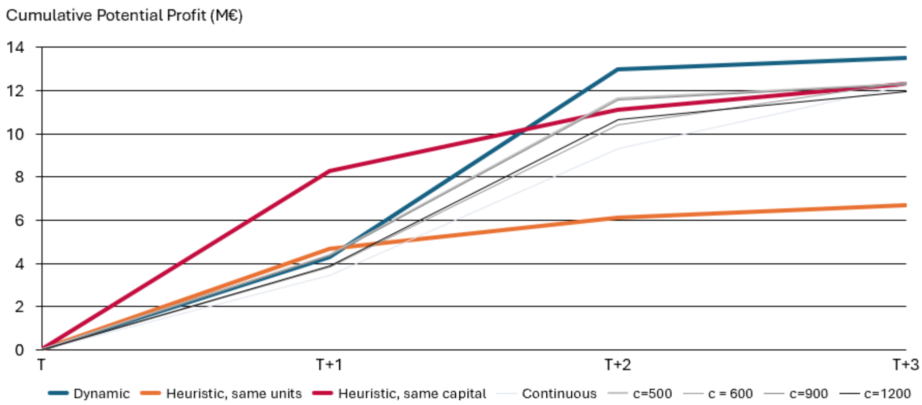

Figure 6: The cumulative potential profit of the purchased vehicles for the 6 SDDP and 2 heuristic models across 3 time stages. The best performing SDDP model and the heuristic models have been highlighted

The same phenomenon is seen when comparing the ROCE, where the heuristic models see a high ROCE in the first month, however the dynamic and other SDDP models see a notable improvement across the longer time period (Fig. 7). The ROCE is similar for both heuristic strategies, while the dynamic maximum SDDP model once again outperforms the other SDDP models across the time period. The ROCE for all strategies converges to approximately 17.6% at the end of the time period. This is because we have set the profit margin for a purchased car to 15%. Thus, as most of the purchased cars are likely sold at the end of the time period, the ROCE converges to the ratio 0.15/0.85, which is the ratio between the profit and the cost.

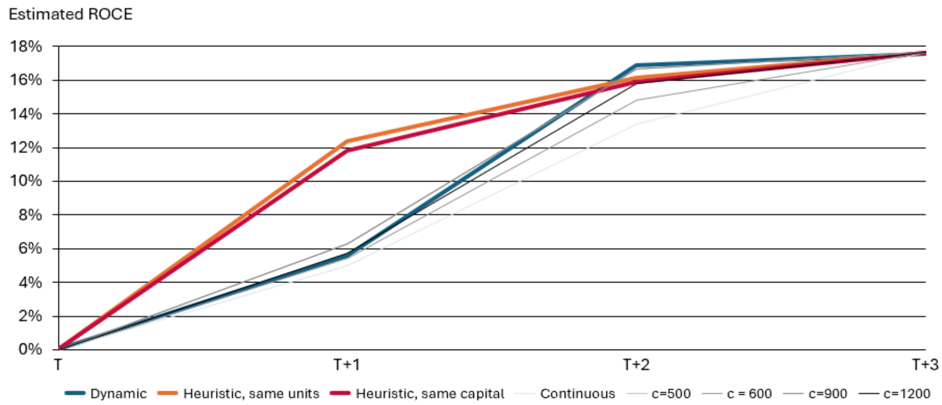

Figure 7: The estimated ROCE of the purchased vehicles for the 6 SDDP and 2 heuristic models across 3 time stages. The best performing SDDP model and the heuristic models have been highlighted

The same capital heuristic model performs similarly to a majority of the SDDP models, only clearly outperformed by the dynamic maximum variant. However, When comparing the number of vehicles purchased, it is clear that the same capital heuristic is unfeasible when examining the inventory capacity (Fig. 8). The same capital heuristic model purchases over 1000 vehicles more than the strategy that purchases the second-most cars.

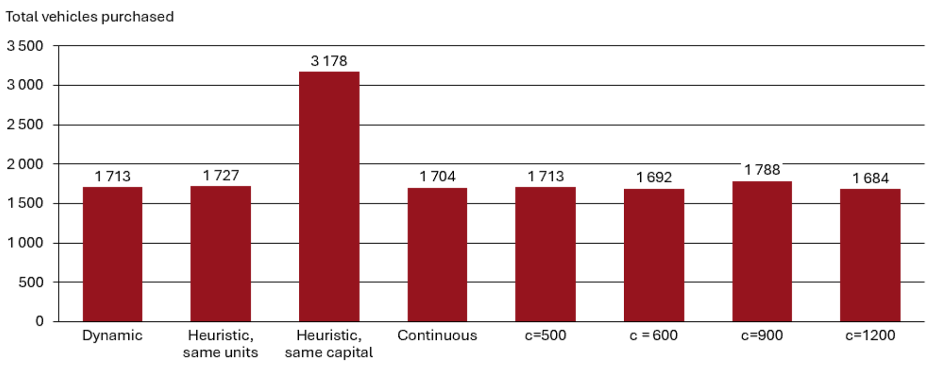

Figure 8: The number of purchased vehicles at time T for all 6 SDDP and 2 heuristic models

The dynamic maximum model has a clear focus on buying vehicles with the largest possible spread between the purchase and selling price. On the

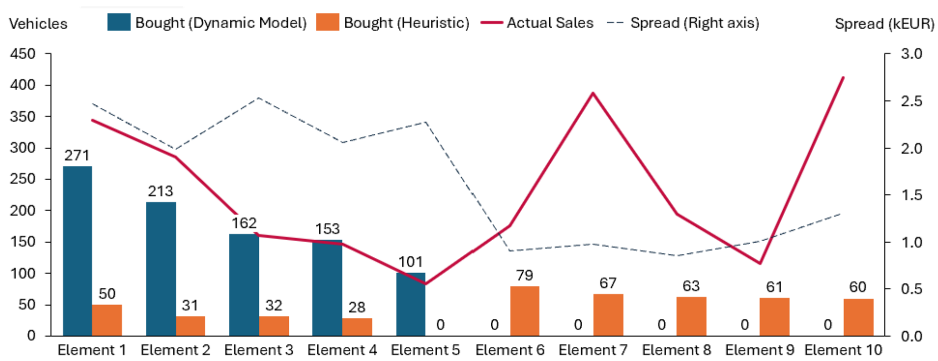

Figure 9: A comparison of the bought vehicles for the dynamic maximum SDDP and same vehicles heuristic models. The colored bars and line plot relate to the left axis, and the dashed line representing the profit margin to the right axis.

other hand, the simple heuristic typically captures a smaller spread between sales and purchase price (Fig. 9). Furthermore, the dynamic model clearly adheres to a Min/Max type of strategy, where it overbuys vehicles with a high spread and a large potential number of sales. On the other hand, it purchases no units of the top 5 sold car elements, which the simple heuristic buys the most of. Simplistically stated, the key driver between the differences in cumulative profit and ROCE is that the bought dynamic model attempts to maximize the total profit buy purchasing a lot of cars with a high profit margin, while ensuring these cars will most likely be sold. On the other hand, the simpler heuristic model is indifferent to the spread between purchase and sales price, instead prioritizing sales volume.

6 Discussion and future research

6.1 Model validity and limitations

The model consists of three parts: the element division, the state division, and the SDDP model. The validity of the results is dependent on the validity of all these intermediate stages. Each of the modeling choices made are justified based on the problem requirements. However, the limitations of the project simplify the complex dynamics involved in the used car trade, which affects the validity of the results.

The element division was done accordingly to expert opinion. The requirements were to create an actionable element division, which is as precise as possible. Therefore, broad descriptions of the vehicles were not feasible. The previous sales data is analyzed with regard to these elements, and the future sales prediction is used as an estimator for the demand. However, with increasingly precise element division, the demand estimation becomes more inaccurate. The customers typically have some requirements for the vehicle they wish to purchase, such as the number of seats or the body type, such as an SUV. Describing the elements with broad categories, such as *SUV, 2-3 years old and automatic transmission*, would more accurately estimate the customer demands. However, as this definition is so broad that it provides little value for the purchasing managers, the elements need to be defined more accurately. The current element division is make, model, fuel type, and engine displacement. This more accurate element division does not take into account the overlap for the element demands, and can then produce results with too little variety. On the other hand, due to the already complex nature of our model, the total number of elements needed to be limited. Therefore, the suggested element division by purchasing experts could not be fully implemented, meaning that even more precise element division could have produced more insight for the purchasing managers.

The chosen approach of the multistage decision problem required the sale of the elements to be in discrete states. This state division, discussed in Section 4.2, was done by simulating possible scenarios using element-wise sales forecasting, and then clustering these scenarios into finite states. As was already discussed in Section 4.1, the chosen approach to use SARIMA for sales forecasting is not necessarily the optimal way, as other models may perform better if nonlinear dependencies are present in the data. Also, the memoryless nature of the time series model forced us to make an approximation in the third stage of the state division, as was discussed in 4.2. To avoid the approximation, one would have to use some other approach for the sales forecasting. Moreover, some simplifications were made to keep the initial forecasting model straightforward. For instance, we used a single set of SARIMA parameters across all different elements, even though a different set of parameters might have been more suited for some individual elements.

As discussed earlier, the element division is a necessary step to make the model actionable. However, once the data is split into narrowly defined elements, the number of historical data points per element becomes very limited. This also affects the reliability of the forecast and the estimated transition probabilities in the Markov model, which can ultimately affect the optimization outcome. On the other hand, the model operates on a fairly limited number of features. Key attributes such as color, mileage, transmission type, drivetrain, vehicle history, and condition are not included. While this reduction is partly due to data availability and partly due to the desire for computational simplicity, it reduces the descriptive power of the model. As a result, the model cannot fully account for the complexity of customer preferences or the nuances involved in vehicle valuation and demand.

As discussed in Section 4.3.2, the multistage decision problem approach is valid based on the given project objectives. The decisions are made for a three-month planning horizon with 1-month intervals. The individual cars can only sell an integer amount, and the sold amounts can be combined into states describing the realized sales of a stage. Additionally, the transition from state to state between stages can be presented as a Markovian transition matrix. Then, SDDP modeling is used to solve the formulated problem. The validity of the claim that the transition between stages can be described by a Markov process can be questioned. In a Markov process, the probability of transitioning into the next state is only dependent on the current state and does not consider the previous states. However, in the real world, the demand for vehicles might not be best described this way. It is reasonable to assume that the current state of sold vehicles alone does not fully capture the probability of sales in the next period. The state transition matrices can be dependent on the overall economic situation, which could be estimated from previous state transitions. The model also gives an estimate of the number of units bought of some element. However, the supply of cars to be bought is also an unknown variable. The model results give the number of units one should buy of each element. However, if the number of units to be bought is not available, the results become suboptimal.

In the problem setup, we also analyze the type of cars to buy for the next period, such that we would buy all the cars immediately in to the inventory. However, the solution from this might not reflect the real-life situation where you could buy additional cars to the inventory once you have sold some units. This means the total number of cars bought in one month is higher than the inventory capacity. It is possible that the optimal solution would be one where the dealership buys cars with a lower profit margin but a very high turnover rate. However, due to the problem setup, dynamic allocation is not possible.

Additionally, the model only gives recommendations on elements to be bought if we assume a static profit margin in the average selling price. However, the decision for the dealership to buy a used car is dependent on many factors not described by the element division features, such as the actual asking price and the car's condition. Therefore, purchasing managers are still needed as decision-makers to analyze the profitability of potential procurements.

6.2 Project achievements

Despite the limitations of our approach, this project offers a valid basis for the inventory management process, which can be improved upon in later implementations. Our model shows a proof of concept that data-driven inventory management is feasible with the existing data. Our approach is constructed in a modular way, meaning parts of the process can be modified easily. Therefore, this project offers a basis that can be improved upon or modified for a different purpose while keeping the overall structure the same. The constructed approach also matches fairly well with the objectives presented in Section 1.2, and the limitations arise from the available data and the complex nature of the task.

6.3 Future research

The provided model offers a basis for future developments. The modular structure of the approach means every part of the approach can be easily changed or improved without changing the overall structure.

The element division can be improved upon with more detailed discussion with experts and deeper data analysis. Currently, the element division is static, as every car type is divided similarly. However, a more dynamic division can be implemented, meaning, more element categories can be introduced or removed depending on the model. For example, the demand for a truck can be dependent on whether it has a towing hitch. Therefore, this should be one of the element division features. However, this does not provide any value if the car is a convertible sports car. Additionally, with dynamic element features, features can be added or removed based on their usefulness, which keeps the total number of elements negligible. Therefore, dynamic element division could provide additional value for the purchasing managers without increasing the computational complexity too much.

The current forecast of the element sales is done with a SARIMA model. This is due to the need to easily provide justifications for buying some element. However, this time series approach is limited to only considering past sales. With a more sophisticated method, the forecasting model could take into account the more complex nature of the demand. For example, the XGBoost model, which is a supervised machine learning model that can be used for various tasks, such as regression. The model is able to handle large datasets and would also be able to handle external factors such as economic situation, internet search trends, and other possibly affecting factors.

The SARIMA model itself could also be improved and tuned. First, the chosen parameters may not be the optimal combination. For the first iteration we emphasized simplicity, which could mean that in some cases we did not utilize the full predictive power of the SARIMA. Secondly, since we are attempting to forecast the sales of multiple different car elements at the same time, choosing one set of parameters is not necessarily the optimal choice. Ideally, one would like to have optimized parameter choices for each element separately.

Future implementation of the model could also optimize inventory composition based on what cars are currently available to be purchased. The current implementation only gives what is optimal if every possible element is available for purchase. However, if some elements have a limited supply and the optimal amount can not be purchased, the solution becomes suboptimal. To guarantee optimality, a constraint for the purchased elements needs to be introduced. However, this type of data was not available for us during this project.

Additionally, sensitivity analysis should be done before this process is implemented. Due to the scope of this project, only simple validation of the methods could be executed. Sensitivity of the optimal solution should be tested, how it behaves with slightly different beginning inventories, and how much the optimal inventory changes with slightly different assumptions, prices, and inventory costs, etc.

The multistage decision problem approach is not the only option, and other possible approaches, such as assortment models, could also be used for the optimization. After constructing alternative models, the optimal inventories should be compared, and the best approach for the problem can then be decided accordingly.

7 Conclusions

This project presents a data-driven framework for optimizing the used car inventory at K-Auto, addressing the core objectives of improving purchasing decisions, maximizing ROCE, and enhancing overall inventory composition. By leveraging SDDP, seasonal time series forecasting, and Monte Carlo simulation, we constructed a modular optimization pipeline capable of generating actionable purchase recommendations under sales uncertainty.

When comparing the performance of the SDDP model against the heuristic strategies, the dynamic maximum variant of the SDDP model consistently delivered the best overall results. It achieved the highest cumulative profit over a three-month period, significantly outperforming both heuristic baselines by more than 1.2 million euros. While the heuristics showed stronger short-term returns due to high-turnover vehicles, they lacked efficiency over time. The SDDP model prioritized vehicles with higher profit margins and managed inventory constraints better, making it both more profitable and more feasible in a real-world setting.

Despite data limitations and the complexity of real-world used car operations, our model demonstrates the feasibility of applying advanced analytics to support strategic inventory decisions. Key challenges such as element division granularity, the lack of true demand data, and simplified assumptions regarding profit margins and supply availability limit the model's precision. Nevertheless, the modularity of our approach enables iterative improvements and adaptation to real operational constraints.

The developed model lays a strong foundation for future enhancements, including dynamic element definitions, more advanced forecasting techniques (e.g., XGBoost), and the incorporation of supply constraints and economic indicators. While the current results are indicative rather than prescriptive, the project validates the potential of data-driven optimization in improving operational efficiency and responsiveness to market conditions in the used car business.

Ultimately, this project delivers a valuable proof of concept and a robust baseline for K-Auto's ongoing efforts toward smarter, more adaptive inventory management.

8 Self Assessment

8.1 Implementation of the project with regard to original plan

The project implementation largely followed the structure and purpose of the initial plan, but there were several notable deviations and adjustments along the way. Our original proposal included a clear focus on refining the inventory model for K-Auto through a structured process: identifying key features, dividing elements meaningfully, and validating the model with client data. In practice, time constraints, the complexity of the task, and challenges with scheduling led to significant changes in scope and execution.

One of the key differences was the limitation in feature engineering and element division. Initially, we intended to carry out a more detailed segmentation of inventory elements based on variables such as car type, fuel, and make. However, due to the capacity of the model and the ability to generalize the results, we had to scale back these ambitions. For example, many modeling refinements were delayed because they depended on decisions made during earlier stages, such as element division or feature aggregation.

Another deviation was in the validation and verification phase. Although we planned to engage in an active feedback loop with K-Auto to fine-tune our models, the time frame and complexity of the model limited this collaboration. Nevertheless, we remained in contact with both K-Auto and course personnel to ensure that our model remained aligned with business logic. In cases where initial models underperformed or lacked interpretability, we adapted by trying alternative approaches and refining assumptions.

Despite some of these unanticipated challenges, our implementation still met the core objectives of the original project plan.

8.2 In what regard was the project successful

The project was successful in demonstrating that a data-driven approach to inventory optimization at K-Auto is both feasible and valuable. Even though the final solution did not reach full deployment readiness, our work clearly shows that existing data can be used to support more informed decisions around used car inventory. The model, while basic in its current form, lays a robust foundation that can be refined and extended in the future.

Our model consists of three main components: element division, sales forecasting, and optimization using SDDP. Each of these components can be independently modified or replaced. For example, different methods can be applied to define elements, forecast sales, or construct the optimization model. This modular structure allows the model to adapt to varying business requirements, data availability, or methodological preferences.

Additionally, the project added value in terms of raising awareness of how data can enhance operational efficiency. We demonstrated that, even with limited time and resources, a data science process can offer tangible benefits. Our analysis also reinforced the importance of setting practical constraints and simplifying decision-making processes for end-users.

On the human side, the collaboration with K-Auto and the learning-by-doing environment offered by the course gave us a strong understanding of real-world project work. We developed skills in client communication, data wrangling, modeling under uncertainty, and team-based problem-solving—skills that go beyond technical modeling and are highly relevant in professional practice.

8.3 In what regard was the project not successful

Despite these achievements, there were several areas where the project fell short, particularly regarding scheduling and workload management. The project plan could not be followed strictly during the final phase, especially after mid-April, when we began writing the report. At this point, some technical elements were still in progress, which made coordination between documentation and model completion more difficult. The initial delay at the start of the course also contributed to a compressed timeline later on.

Moreover, the client collaboration, while useful, was not so active or continuous. Limited meetings and feedback sessions reduced our ability to iterate based on real business constraints or insights from K-Auto. As a result, the model remained more theoretical in some aspects than we would have liked.

Another challenge was the uneven division of tasks within the team. Some responsibilities, particularly around modeling and technical execution, ended up falling on individual members, creating an imbalance in workload. Since many steps of the project were sequential, this also meant that if one part was delayed, the rest of the work could not proceed in parallel, which compounded the delays. While everyone contributed meaningfully to the final product, the lack of task parallelization and uneven workload distribution limited our overall efficiency.

8.4 What could have been done better

There are several key lessons we take from this project in terms of process improvement and project management. First, more regular and structured internal meetings would have helped keep the project on track and improved collaboration. Weekly check-ins could have ensured more accountability, quicker problem resolution, and earlier feedback on ongoing work.

Second, communication with the client could have been improved. While K-Auto was responsive when approached, more proactive scheduling of meetings would have created a more continuous feedback loop. This would have allowed us to adapt to their needs more effectively and validate our assumptions earlier.

Third, the timeline could have been better structured. The initial month of the course saw limited progress due to a lack of concrete deliverables, which delayed our momentum. An earlier project kickoff, coupled with clearer task division and progress tracking, would have enabled better use of the available time. Similarly, the interim report came at a point when there was not enough substantial content to reflect on, making it less effective. Modifying the course's schedule structure or aligning reporting deadlines with the actual project pace would be helpful in the future.

Lastly, improving task division by assigning overlapping responsibilities or collaborative pairs on key components would have reduced the dependency on single individuals and created a more resilient workflow. Building in more parallel workstreams would also mitigate the impact of delays in one part of the project on the entire timeline.

In conclusion, while the project had its limitations, it provided a strong learning experience and delivered real value to the client. With a few adjustments in project planning, communication, and team coordination, the same framework could yield even more impactful results in future implementations.

References

Arrow, K., Harris, T., and Marschak, J. (1951). Optimal inventory policy. *Econometrica*, 19(3):250–272.

Babiloni, E. and Guijarro, E. (2020). Fill rate: from its definition to its calculation for the continuous (s, q) inventory system with discrete demands and lost sales. *Central European Journal of Operations Research*, 28(1):35–43.

Bandarra, M. and Guigues, V. (2021). Single cut and multicut stochastic dual dynamic programming with cut selection for multistage stochastic linear programs: convergence proof and numerical experiments. *Computational Management Science*, 18:125–141.

Choi, T. M., editor (2014). *Handbook of EOQ Inventory Problems*, volume 197. Springer US, Boston, MA.

Downward, A., Dowson, O., and Baucke, R. (2020). Stochastic dual dynamic programming with stagewise-dependent objective uncertainty. *Operations Research Letters*, 48(1):33–39.

Dowson, O., Morton, D. P., and Pagnoncelli, B. K. (2020). Partially observable multistage stochastic programming. *Operations Research Letters*, 48(4):505–512.

eCarsTrade (2023). Will used car sales change in europe in 2023? a complete analysis. `https://ecarstrade.com/blog/used-car-sales-europe`. Accessed: 2025-05-14.

Ellencweig, A. et al. (2023). Data and analytics in the driver's seat of the used-car market. `https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/data-and-analytics-in-the-drivers-seat-of-the-used-car-market#/`. Accessed: 2025-03-01.

Ensafi, Y., Amin, S. H., Zhang, G., and Shah, B. (2022). Time-series forecasting of seasonal items sales using machine learning–a comparative analysis. *International Journal of Information Management Data Insights*, 2(1):100058.

Füllner, C. and Rebennack, S. (2021). Stochastic dual dynamic programming and its variants – a review. *Optimization Online*. Published: 2021/01/19, Updated: 2024/10/16.

Guigues, V. (2017). Dual dynamic programming with cut selection: convergence proof and numerical results. *European Journal of Operational Research*, 258(1):47–57.

Guigues, V., Shapiro, A., and Cheng, Y. (2023). Duality and sensitivity analysis of multistage linear stochastic programs. *European Journal of Operational Research*, 308(2):752–767.

Ji, S., Wang, X., Zhao, W., and Guo, D. (2019). An application of a three-stage xgboost-based model to sales forecasting of a cross-border e-commerce enterprise. *Mathematical Problems in Engineering*, 2019(1):8503252.

Kayacık, S. E., Basciftci, B., Schrotenboer, A. H., and Ursavas, E. (2024). Adaptive multistage stochastic programming. *arXiv preprint arXiv:2401.07701.*

Kelley, Jr, J. E. (1960). The cutting-plane method for solving convex programs. *Journal of the society for Industrial and Applied Mathematics*, 8(4):703–712.

Maitra, S. (2024). Inventory management under stochastic demand: A simulation-optimization approach. *arXiv preprint.* arXiv:2406.19425.

Pereira, M. and Pinto, L. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52:359–375.

Porteus, E. L. (1990). Stochastic inventory theory. In Heyman, D. and Sobel, M., editors, *Handbooks in Operations Research and Management Science*, volume 2, chapter 12, pages 605–652. Elsevier Science Publishers B.V. (North-Holland), Amsterdam.

Saliji, M. (2021). Effective inventory management in the automotive industry, a literature study. Dissertation. https://urn.kb.se/resolve?urn=urn:nbn:se:mdh:diva-54644.

SDDP.jl Developers (2024). Sddp.jl documentation.

Stopková, M., Stopka, O., and Ľupták, V. (2019). Inventory model design by implementing new parameters into the deterministic model objective function to streamline effectiveness indicators of the inventory management. *Sustainability*, 11(15):4175.

Tamjidzad, S. and Mirmohammadi, S. H. (2015). An optimal (r, q) policy in a stochastic inventory system with all-units quantity discount and limited sharable resource. *European Journal of Operational Research*, 247(1):93–100.

Tamjidzad, S. and Mirmohammadi, S. H. (2017). Optimal (r, q) policy in a stochastic inventory system with limited resource under incremental quantity discount. *Computers & Industrial Engineering*, 103:59–69.

Tiedotuskeskus, A. (2024). Uusien ja käytettyjen autojen kauppa maakunnittain. Accessed 7 May 2025.

Turgay, S., E. K. (2023). Optimizing inventory allocation for fast moving consumer goods: A stochastic model approach.

Zipkin, P. (2000). *Foundations of Inventory Management.* McGraw-Hill Companies, Incorporated, New York, illustrated edition.