Optimizing box content in try-before-you-buy business models with heterogeneous customer groups

Martin Senftlechner

Dissertation written under the supervision of Professor Joren Gijsbrechts.

Dissertation submitted in partial fulfilment of requirements for the MSc in Business Analytics, at the Universidade Católica Portuguesa on January 4, 2023.

Optimizing box content in try-before-you-buy business models with heterogeneous customer groups

Martin Senftlechner

January 4, 2023

Keywords: try-before-you-buy, customer satisfaction, linear optimization, mixed integer programming, simulation

Abstract

In online retailing, try-before-you-buy business models are emerging. Customers can order items to be shipped home to try and decide which items to keep or send back to the retailer. Usually, this business model is combined with personalized recommendations. In this thesis, an optimization model for supporting the decision of what to put in a box to achieve maximum profit in a sales period is introduced. Multiple periods are simulated and overall profit is analyzed. Profit is evaluated by summing the optimum identified in each period. The optimization model treats customer groups differently based on their purchasing behaviours. Policies and business rules are varied to understand the effects on overall profit and customer groups. From the results, managerial implications are drawn to follow a customer-centric approach. To maximize profit, customers with high lifetime value should be treated as preferred if market demand is high and no other factors limit shipping decisions. Notable limitations include inventory availability and market demand. In market situations with limitations, higher-valued customers should be served first. Once a certain scale of the customer base is reached, it should also be focused on other customer groups. It is shown that the prediction accuracy of input data is a crucial concern for sufficiently optimizing box content. Further works could improve the chosen approach to better understand the effect of actions taking place in future sales periods.

Resumo

Na venda a retalho online, estão a surgir modelos de negócio try-before-you-buy. Os clientes podem encomendar artigos a serem enviados para casa para tentarem decidir que artigos guardar ou enviar de volta para o retalhista. Normalmente, este modelo de negócio é combinado com recomendações personalizadas. Nesta tese, é introduzido um modelo de optimização para apoiar a decisão sobre o que colocar numa caixa para obter o máximo lucro num período de vendas. São simulados vários períodos e o lucro global é analisado. O modelo de optimização trata de forma diferente grupos de clientes com base nos seus comportamentos de compra. As políticas são variadas para compreender os efeitos sobre o lucro global e os grupos de clientes. A partir dos resultados, as implicações de gestão são desenhadas para seguir uma abordagem centrada no cliente. Para maximizar o lucro, os clientes com elevado valor vitalício devem ser tratados como preferidos se a procura do mercado for elevada e nenhum outro factor limitar as decisões de expedição. Em situações de mercado com limitações, os clientes de valor mais elevado devem ser ser servidos em primeiro lugar. Uma vez atingida uma certa escala da base de clientes, esta deve também ser concentrada noutros grupos de clientes. Demonstra-se que a precisão da previsão dos dados de entrada é uma preocupação crucial. Outros trabalhos poderão melhorar a abordagem escolhida para compreender melhor o efeito das acções que têm lugar em períodos de vendas futuras.

Contents

1 Research problem

1.1 Introduction

The retail industry has changed drastically in recent years. Especially in online retailing market players try to meet customers' enhanced expectations by providing sophisticated information, offering personalization and distinctive shopping experiences. This is also important due to the sheer amount of SKUs (stock keeping units) and little differences between many of these products [Grewal et al., 2018, Li et al., 2019]. In addition to the competition of established (online) retailers like Amazon or fashion and apparel-centered companies like Zalando or Asos, new market entrants and start-ups have found their niche with non-traditional business models or exceptional innovations. These innovations are often enabled by cutting-edge technologies, including artificial intelligence, distributed ledger technologies or the internet of things [Kahn et al., 2018, Li et al., 2019, Tao and Xu, 2018].

This thesis focuses on finding the optimal selection of products to pack in boxes and ship to customers at the Belgium-based try-before-you-buy online apparel retailer "CurveCatch". The results are generalized to other companies following similar business models. Try-before-you-buy ("TBYB") companies incorporate the fact that return rates are high in e-commerce into their business model. Retailers determine how many and which products customers may like. A selection of items is put in a box and shipped to a customer. The customer can try the product and decide whether to keep or return it to the retailer [Li et al., 2019, Park et al., 2022]. CurveCatch's operations are centered around a short questionnaire and data from previous orders. Based on these insights, they send lingerie items determined by their experts' recommendation or recommendation system. The company holds data about its products, customers (preferences, persona, etc.), recommendations (virtual box) and the sold or returned items. There are three main sections where decision support models might be applied in try-before-you-buy operations: customer and product classification, customer style-fit prediction and box content prescription (see figure 1). This structure also holds for CurveCatch's downstream recommendation and order fulfillment process.

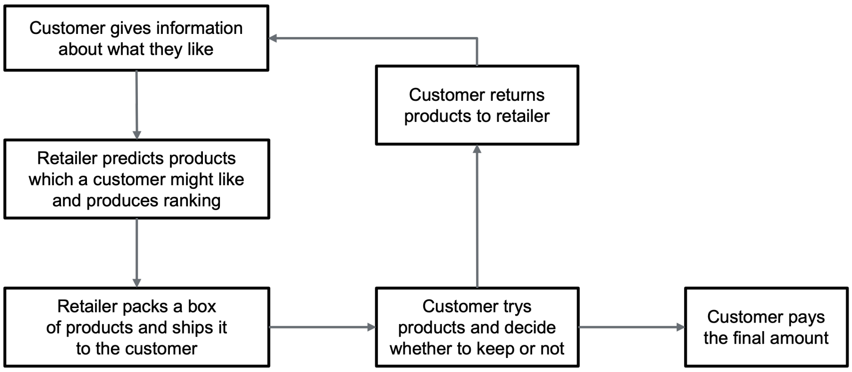

Figure 1: Making box content and quantity decisions in TBYB companies

This thesis will focus purely on the determination of the final box content to be shipped to customers. An optimization model will be introduced and trade-offs between different policies will be analyzed. The general structure of such a multiple-step recommendation system and superiority when optimizing profitability was shown by [Chen et al., 2008]. The prediction of products fitting the customer's style and the initial recommendation of products based on the prediction are out of scope. Also, this work does not cover the classification of items, item categories and customers. Whereas the initial recommendation of fitting products focuses on identifying which products customers might like, the optimization of the box content conducted in this thesis is bottom-line centered. Both aspects need to result in desirable products for customers in order to follow a customer-centric business approach. Key metrics to examine and possibly consider in respective decision support models could include the ratio of purchases to sent items or the average profit from a sent box after handling returns. In this thesis, the main objective will be maximizing the overall profit. Main constraints could include available inventory levels, inventory on the road (e.g. returns), margins associated with products or business rules like a maximum number of shipped items per customer. The results of this thesis are discussed from a managerial perspective, giving insights into possible customer-centric approaches. It is shown how customers related to high purchasing power and lifetime value are especially valuable to satisfy. Because of limited numbers of the most valuable types of customers in market, trade-offs in serving them must be made. Customers with a medium-high purchasing power should also be served well in order to follow a balanced targeting strategy. Sales volume is mainly made up of medium- and lower-purchasing power customers. Although the possible impact on TBYB business operations of different policies will be evaluated, assessing or altering the primary business model of the specific case of CurveCatch is not a goal of this work.

Since the recommendation-based try-before-you-buy business model is novel, only little academic research in this domain has been conducted. This thesis aims to shortly summarize the current state of research and build up on that by setting up an optimization model for supporting quantity and content decisions in TBYB businesses. The model is embedded in a simulation environment to analyze the effects of different policies and business rules on profit over time and to replicate customer purchase behaviour and inventory levels. Thus, input variables and parameters are required and are either based on research, industry-benchmarks or derived from available data. These inputs cover mainly three categories: expected customer behaviour for customer classification, products with their respective product type and inventory levels of items. Customers' profiles vary, particularly regarding the money they are willing to spend on apparel and their probability of purchasing products. Similar profiles are mapped to generalized customer groups, which will subsequently be referred to as classes. Based on that, box content is optimized for each sales period. Results from each period (e.g. business days) are examined in a temporal context to understand how parameters could influence decisions related to the treatment of different customers based on their expected behaviour. Specifically, stock levels and varying customer profiles are affecting the decision-making suggested by the model. General aspects and implications are drawn from the findings and set into an overall business context. This work relies on information derived from data available from CurveCatch and assumptions from business and literature. The thesis gives managerial insights into the treatment of customer classes to maximize period profit over a finite horizon of multiple sales periods, by answering the following the research questions:

1. *RQ1: How can quantity and content decisions of shipped boxes be made for maximizing overall profit in TBYB businesses?*

2. *RQ2: How can components of customer behaviour relevant for simulating the effects of profit-optimized TBYB shipping decisions be replicated?*

3. *RQ3: How can TBYB business models follow a customer-centric approach while optimizing profit?*

RQ1. This question aims to develop an optimization model for making decisions about quantity and content. Specifically, the model supports the decision of which items should be shipped out of a ranked preference list per item type per customer. The model will be based on previous research, which is modified and suited to the business situation of CurveCatch. The base model aims to give support to single-period decisions for multiple customers, but expansions for multiple periods are introduced. As [Dadouchi and Agard, 2018] point out, different policies for boxing low stock level SKUs might be required, depending on the lifetime value of customers. Emerging trade-offs are analyzed and summarized.
RQ2. The introduced model incorporates a variety of variables suggested by existing literature or following a data-driven approach. Simulators help to create a dynamic model test environment, particularly on the demand side and for replicating customer behaviour over multiple sales periods.
RQ3. The focus of this research question is on the business implications of the model results. As different companies have different strategic goals, the policies and approaches they should follow may vary. To stay aligned with their objectives, trade-offs must be made. Businesses have to decide about the customer profiles they are targeting, e.g. a smaller group of customers associated with a high purchasing power or a larger group of customers expected to spend lower amounts of money on average. Additionally, decisions related to customer retention and churn may involve weighing different, potentially conflicting actions [Castéran et al., 2017, Godinho de Matos et al., 2018, Reinartz et al., 2005, Şen and Zhang, 1999].

1.2 Related work and literature review

To reduce product fit uncertainty (subsequently referred to as PFU), which describes the degree to which a customer cannot determine if a product matches their needs and preferences as introduced by [Dimoka et al., 2012] and [Hong and Pavlou, 2014], the concept of "try-before-you-buy" has been emerging [Li et al., 2019]. This can either happen virtually by augmented reality or physically by sending goods to the customer [Marc et al., 2021]. Particularly when sending products for physical inspection and in combination with lenient return policies, TBYB excels in solving PFU [Park et al., 2022, Smink et al., 2019]. TBYB concepts differ among companies. Particularly popular became TBYB business models in combination with recommendation services, sometimes complemented with subscription systems. These systems target customers having insufficient time for shopping or not being interested in doing so, as well as customers curious about items they might not even know. The foundation for the decision making are extensive customer data, prediction and optimization models [Tao and Xu, 2018]. From a retailer's perspective, sending multiple products (bracketing) and offering a free return policy negatively affects profitability. Hence, the accuracy of recommendations is a key enabler for making this business model profitable [Balaram et al., 2022]. The TBYB business model might even outperform the common showroom store's strategy, [Li et al., 2019] argue. The cost imposed by traveling to the physical store and the risk of desired products not being available lead to profit fit uncertainty. The thesis at hand is motivated by CurveCatch[1]. This company focuses on lingerie. Product fit uncertainty is particularly high in this industry sector, especially because of size mismatches between manufacturers and brands. Also, the personal barrier to trying lingerie products in physical stores might be high. The German TBYB company Outfittery[2] or Stitch Fix[3], maintaining operations in the US and UK, address online apparel retailing in general. Stitch Fix can be seen as one of the early movers in the industry and generated a revenue of over 2 billion US$ in 2022.

[He et al., 2022] showed that lenient return policies are demanded by customers when profit fit uncertainty is high. These lenient return policies subsequently stimulate demand but increase stock and return handling costs tremendously [Altug and Aydinliyim, 2016, Li et al., 2013, Smink et al., 2019]. Return handling costs (including repackaging, shipping, and mending) significantly lower the profit margins of retailers when offering lenient return policies, particularly when solely covered by the retailer [Altug and Aydinliyim, 2016, Park et al., 2022]. Because of the cost, online retailers should not offer free returns in a monopoly situation. However, because of offers from competition the retailer still has to do so [He et al., 2022]. To stand out in the competitive landscape, making data-driven decisions about the items shipped has become one of the new approaches to efficiently streamline operations of recommendation-based TBYB systems [Altug and Aydinliyim, 2016, Park et al., 2022]. This means that retailers determine based on decision models how

[1] https://www.curvecatch.com/

[2] https://www.outfittery.de/

[3] https://www.stitchfix.com/

many and which exact products are finally boxed and shipped to the customer [Li et al., 2019, Park et al., 2022]. One might assume that sending a larger number of items likely to fulfill customers' preferences would lead to higher profit or sales revenue. However, not only are profits negatively affected by extensive returns but also revenue and sales numbers will not rise necessarily. Customers have preferences and are limited in their choice due to personal reasons. These restrictions could include the product's style (e.g., customers might not tend to buy too many similar items) or the customer's budget. Also, the mix of products sent to the customer might influence the purchase decision [Khouja et al., 2019, Li et al., 2019, Park et al., 2022]. The SKUs designated to be shipped within one order will be referred to as a "box" of products in the following. [Dadouchi and Agard, 2018] showed how stock levels could be included in an optimized recommender system by adding penalties for out-of-stock products. They concluded that SKUs with a low inventory level should only be recommended to high lifetime value customers. For out-of-stock or low-inventory products, alternatives could be recommended [Guzman et al., 2021].
To find the optimal recommended box content, therefore, a trade-off must be made between the value of an immediate sale and a potential sale in the future, also incorporating different likelihoods of customers keeping the sent items [Bernstein et al., 2021]. Thus, the optimal overall profit is subject to the expected future demand, nature of different customers, expected profit of products, probability of returns and stock levels.

Demand classes categorize the source of the demand based on defined characteristics. This can be useful if a trade-off has to be made, which demand should be satisfied (e.g. because of limited stock) and satisfying a specific source of demand seems more desirable than another. Sources of demand also vary in terms of ordering and purchasing behaviour. Reasons given by [Topkis, 1968] include the criticality of a product in a certain geographic region or profit made by selling goods [Dadouchi and Agard, 2018]. In this thesis, the latter is relevant and, is explained more in detail in the following. To identify the demand classes given by a set of customers, the customer lifetime value can be adduced if the objective is profit or revenue maximization. [Berger and Nasr, 1998] presented a mathematical model for calculating the lifetime value based on customer purchasing and spending behaviour. Since spending from CurveCatch customers' initial orders is available, this data-backed approach has been chosen in this thesis. Also, [Dadouchi and Agard, 2018] were successful in splitting their customer data into classes following this methodology. [Castéran et al., 2017] criticize that this approach relies heavily on assumptions, requires customer churn rates to be stable over time and does not account for different customer segments. This proposition might be true and relevant for calculating tangible and specific values. However, in this work, the general behaviour and effect of an optimization model in the context of deciding box content in TBYB will be studied. Thus, the simplified approach shown by [Berger and Nasr, 1998] is perceived to be sufficient. In the following, the terms demand class and customer class will be used likewise, since only one demand class can be applied to a customer in this work.

The decision on which products are eventually boxed and sent to a customer in recommendation-based try-before-you-buy business models can be modeled in a similar manner to assortment planning decision support models. This hypothesis is a foundation of the thesis at hand and will be elaborated further in section 2. Assortment planning is a central task in inventory management, hence, decision support models regarding inventory levels have been researched. [Li, 2007] introduced a single period assortment optimization model and examined the differences between continuous and discrete demand modeling. His work mainly focuses on in-store POS (point-of-sale) assortment optimization, however, there are general components fitting for similar applications. The model incorporates cost parameters, potential cannibalization between SKUs and a specific expected profit. Furthermore, in the paper, a demand model has been set up, and the model is tested with different continuous and discrete generated data with random variance [DeCroix and Arreola-Risa, 1998]. Resource allocation in store retailing and pricing has been studied extensively. Multiple models, mainly based on stochastic or dynamic programming have been introduced and tested against different demand functions. Whereas some approaches involve complex mathematical models, others rely on heuristic procedures to search the most likely solution space due to limitations of computational power [Bhattacharjee and Ramesh, 2000, Hübner et al., 2020, Ma and Fildes, 2017, Maiti and Giri, 2017]. Finally, there is literature assuming homogeneous goods and trying to solve lot sizing problems [Brahimi et al., 2017]. These approaches for solving assortment planning problems can be mapped to the issue at hand. In both situations, it is aimed to optimize the allocation of scarce resources. However, in assortment planning the optimal amount of different items over a certain planning horizon is computed, whereas the allocation problem in this thesis focuses on distributing the available items within a given period.

1.3 Research approach and methodology

The research is being conducted in two phases: building the optimization model (see 2.2) and analyzing its behaviour over multiple periods of time in varying situations via simulation (see 2.3). The model and the simulation environment are meant to conceptually represent the situation at CurveCatch or other similar try-before-you-buy companies. Thus, the approach is based on papers treating similar issues and adjusted to the managerial setting at hand. The model introduced in the following relies on the concept of different demand classes and assumes a certain purchasing behaviour of selected customer types. The number of customers, their demand and the products offered are held limited for computational reasons. For each sales period studied, one optimization problem will be run. The model is embedded in a simulation environment, allowing for analysis over multiple periods. Demand is induced by simulating customers with respective requested items, ranked by buying preference. The optimizer identifies the ideal allocation of items based on the simulated demand and the available inventory in the respective period. Values are drawn from company data, industry insights or are assumptions backed by literature or general business acumen. Details about the model and used parameters and variables are given in section 2. The effects of changes in the various parameters and variables are then studied by conducting sensitivity analysis. This allows to understand how a model and the respective business situation are affected when the environment changes, e.g. because of changes in customer behaviour, competition or technological advances. In this thesis, a baseline case is defined. Subsequently, input variables or parameters are changed. The computed results from the sensitivity cases are compared with the base case and with each other [Borgonovo and Plischke, 2016]. Based on these results, managerial insights are won about what changes in parameters would mean for revenue and profit generation.

2 Model building and simulation

In this thesis, a mixed integer programming model is set up to decide on the box content for an induced demand in a particular sales period to maximize profit. A simulation environment is set up to generate managerial insights into the benefits of applying the optimization model. This allows studying the effects over multiple time periods. The system consists of four main components (see figure 2). The environment includes a demand generator for simulating a number of customers asking for a selection of items ranked by preference. The optimization problem is solved for each period in the "Optimization Model" component. The other components generate input data, run the simulation of the multiple-period business environment and support understanding of the model's results. Once the optimization is done for the respective period, the purchase simulator replicates the buying behaviour of different customers.

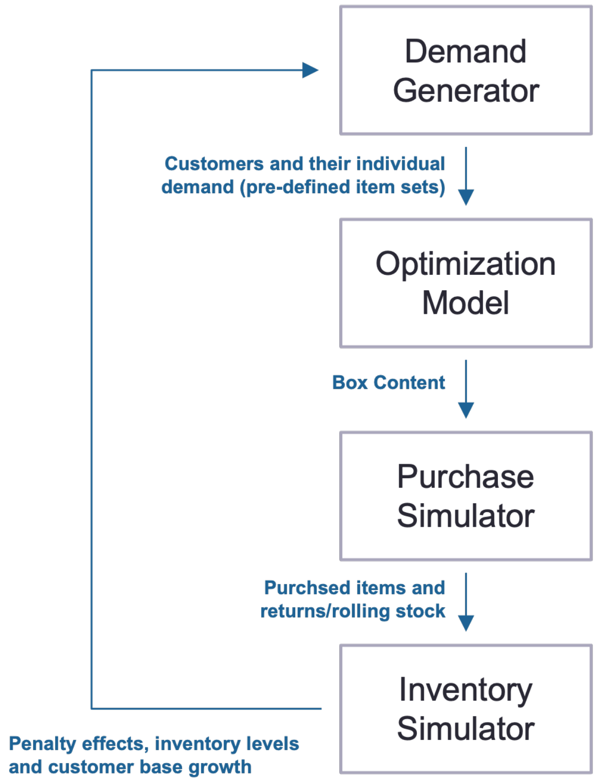

Figure 2: Optimization and simulation components

An inventory simulator keeps track of inventory levels in stock and of products being shipped back from customers. In this work, a sales period is representing a business day. At the end of the business day, all information about demand and supply (inventory levels) are assumed to be available to perform the period's profit maximization by deciding on box contents. Related work focused on integrating (profit) optimization in traditional customer-fit recommendation systems. This, however, could accelerate error rates in the prediction, make it harder to maintain the systems in a productive environment and requires a deep understanding of the products' features [Chen et al., 2008]. As shown before in figure 1, the box optimization in this thesis is treated as a dedicated component within an operational environment. The recommendation system and optimizer are separated. Thus, the optimization requires the results from a preceding style-fit recommendation system as input, which are expected to be available in a usable format. Furthermore, it is assumed that the predictions are of sufficient accuracy since diverging results can distort the model, as shown later. The optimization suggests and chooses the optimal item(s) out of the initial recommendation based on the maximal expected profit for the respective period. In each period, all information about the respective input data, like customers, demand and inventory, must be exhaustively observed before the optimization is performed. This assumption holds in real-world operations if the order information is collected, optimized and shipped at the end of each business day. Subsequently, a short introduction about the components is given and the respective parameters and variables are explained.

2.1 Model components

The first stage, the "Demand Generator", determines a number of customers, their purchase behaviour and their respective demand. Customers are grouped into classes, mainly distinguished by their budget to spend and their average purchasing probabilities for demanded products. So, these classes are associated with a certain behaviour and demand within a given interval. The customer classes used in this thesis are derived from demand classes, which themselves are based on the expected lifetime value a customer has as published by [Berger and Nasr, 1998] and applied by [Dadouchi and Agard, 2018] in a similar work. For example, a customer who regularly orders expensive products and usually keeps most of the shipped items will be associated with a higher lifetime value than a customer ordering only once in a while and returning large proportions of the shipped products. Customer classes follow a particular distribution, which is described in more detail in 2.3. In real-world markets, there might be a limited size of certain customer groups. Among others, reasons could include reaching the absolute total number of customers in a group or strong competitors, making it impossible to grow further in a specific customer group. This market behaviour is replicated in the simulator. The customer base can grow each period. The requirement for that are well served and satisfied customers in the respective class. A maximal number of customers in each group can be defined to simulate a limitation of growth in a specific class. More details about the growth mechanics can be found further down below. Besides the customer profiles for each period, also the demanded products are simulated by the demand generator component. The product demand is structured in item sets, where a customer can request multiple item sets at once. Item sets include items from a preceding customer-style fit prediction, including multiple options for one respective item category. An item set corresponds to a category of items, like a selection of similar or related products, e.g. based on type, size or color. Items within a set are ranked by purchase preferences (see A.2). Each item also has an individual purchase probability assigned, which gets adjusted by a factor for each of the possible customer classes. To simplify the model, the purchase probabilities are taken as averages over an infinite amount of draws and therefore seen as being static and constant for all customers of the same class. A similar approach has been chosen in related work and proved sufficient [Bernstein et al., 2021, Dadouchi and Agard, 2018]. In general, a minimum accuracy of the predicted purchase probabilities per item and the predicted customer class defining the customers' behaviour are key factors for achieving robust results in the up-following optimization [Park et al., 2022].

The second stage holds the mixed integer programming optimization model for suggesting box content for a set of customers. It is formulated as a maximization problem, restricted by customer purchase behaviour and business rules. For each sales period, the model supports the shipping decision for each individual item within item sets of respective customers. In general, favourable customer behaviour, like high purchasing power or high purchase probabilities, is expected to lead to higher profit per shipped box. Thus, these customers should be privileged by the model's shipping decision. To take the behaviour into account in the model, penalties for each customer class can be applied for negative shipping decisions. This means whatever demand remains unsatisfied is facing a penalty cost. A similar approach was proposed by [Dadouchi and Agard, 2018] for recommendation systems or by [Topkis, 1968] with a focus on ordering policies in a less complex business setting, including only a single homogeneous product type. Details about the model formulation are given in 2.2 and the parameterization is described in 2.3.

To reflect the customer purchase behaviour in the profit calculation, a respective purchase simulation component is included. The output of this step is information on whether shipped items are purchased or returned by individual customers. Thus, the purchase simulator draws the actual purchase decision for each item shipped to an individual customer within a period. This step is placed subsequent to the model's shipping decision. It accounts for budget limits, the expected purchase probability per customer class and some added random noise. The simulator is heuristics based. First, the purchase decision is drawn based on a binomial distribution with respect to the purchase probability. After the set of items that the customer wants to keep is determined, the budget limit comes into effect. The heuristics assume that customers prefer the highest ranked item available within the highest ranked (highest rank = first position) item sets. They continue to purchase items in this order until their individual budget is reached. All other items have to be returned and are added to a rolling stock index, which keeps track of items shipped back in each sales period.

The main purpose of the inventory simulator component is removing and adding items shipped or returned to the simulated inventory and keeping track of the items in transit (=rolling stock when items are sent back by customers). Also, in the last step in each period, the effect of not satisfied customers is taken into account. If preferred (=highly ranked) items are not shipped to customers, and they might not order again and churn. If satisfaction is not split evenly over all customer classes, on average, this leads to a reduced portion of customers in the more affected classes. A possible goal for TBYB companies might be increasing the proportion of highly valuable customers who are expected to spend more than their lower-class counterparts. Additionally, for the case of CurveCatch, a constant growth of the customer base is assumed. This growth is only applied to a specific satisfied group of customers if their respective orders are fulfilled. Thus, if customers are served well, the number of customers within their group increases each period gradually.

2.2 Model formulation

The optimization problem for each single sales period is solved by mixed integer programming. The objective is formulated as a maximization function and subject to several constraints, whereby foundations from literature have been leveraged [Bhattacharjee and Ramesh, 2000, Dadouchi and Agard, 2018, Park et al., 2022]. Under the situation introduced before, the unconstrained problem for maximizing total profit $\prod$ of items x with profit π for a single customer in a single period can be formulated as

$$\prod = Max \sum_{n=1}^{m} x_n \pi_n \beta_n - (1 - \beta_n)c \tag{1}$$

with β being the expected average purchase probability and cost c associated with handling returns. The binary decision variable $x \epsilon \{0, 1\}$ determines the shipment of product x_n. Note that the return probability for calculating the expected return handling cost is taken as $1 - \beta$, assuming that the information about the predicted purchase probability is exhaustive. The optimum in each period should be computed over a set of i customers, leading to the following extension of the function

$$\prod = Max \sum_{i=1}^{j} \sum_{n=1}^{m_i} x_{in} \pi_{in} \beta_{in} - x_{ikn}(1 - \beta_{in})c \tag{2}$$

The model is meant to optimize shipping decisions for items within heterogeneous item sets. Thus, these item sets k are added to the objective, which results in

$$\prod = Max \sum_{i=1}^{j} \sum_{k=1}^{l_i} \sum_{n=1}^{m_{ik}} x_{ikn} \pi_{ikn} \beta_{ikn} - x_{ikn}(1 - \beta_{ikn})c \tag{3}$$

In the next step, penalties P are included. The penalty is a fictional cost that is applied to each customer who does not receive their preferred items from each of their demanded item categories. If applied, this cost lowers the profit for each customer affected. Since the penalty is applied for each requested item category, it accumulates to a larger value for unsatisfied customers requesting more different product types. Additionally, the proportion of penalized number customers in each class affects the demand in the upcoming period by a small adjustment factor. This is because unsatisfied customers might churn or spread negative word of mouth within their peer group. To account for different expected customer purchase behaviours, and thus, the importance of the satisfaction of a customer category to a business, the value of the penalty varies with regard to the customer class. Customers with higher expected value for a company are affected by higher imposed penalties. The final objective function for maximizing the profit in one sales period can be formulated as

$$\prod = Max \sum_{i=1}^{j} \sum_{k=1}^{l_i} -P_{jk} + \sum_{n=1}^{m_{ik}} x_{ikn}\pi_{ikn}\beta_{ikn} - x_{ikn}(1-\beta_{ikn})c \tag{4}$$

The maximization function is subject to several constraints. These constraints are driven by business reasons but also due to modelling requirements and limitations. In general, no shipped box contents should exceed a certain number of items nor a pre-defined total value, as formulated in constraint (5) and (6).

$$\sum_{k=1}^{l} \sum_{n=1}^{m_k} x_{ikn} \leq x_{max} \tag{5}$$

$$\sum_{k=1}^{l} \sum_{n=1}^{m_k} p_{ikn}x_{ikn} \leq value_{max} \tag{6}$$

Product profit π, is formulated as

$$\pi_n = p_n\mu_n \tag{7}$$

where $\mu \in [0,1]$ is the profit margin (in %) per product, which is assumed to be known by the business. Note, that (6) uses the actual total sum of prices p of the boxed items.
Besides the total value of items there is another budget constraint in place. In constraint (8) the expected budget a customer i is willing to spend is modelled. It is computed by considering the value of the shipped items multiplied by the respective probability that a customer is actually buying the product. Otherwise, there could be a situation where products of a much higher value are shipped than the respective is able to purchase. So, the value of the shipped items should not exceed this threshold since it would necessarily lead to returns, on average.

$$\sum_{k=1}^{l} \sum_{n=1}^{m_k} x_{ikn}p_{ikn}\beta_{ikn} \leq budget_i \tag{8}$$

As an additional measure to make sure high lifetime value customers are satisfied, a threshold of inventory stock levels can be defined (see 2.3). This threshold allows for a certain number of items to be reserved for customers of the highest-rated classes. Reserved items can only be shipped to these valued customers, ensuring that they have access to the items they demand. By reserving a portion of the inventory, the risk of running out and incurring penalties related to unsatisfied customers can be avoided. Thus, if the inventory level is below this threshold, constraint (9) is applied for the respective item n for a customer i if the customer belongs to a lower-valued class. To avoid using another set of binary variables in the model, the respective constraint (9) is embedded in a conditional statement. Only if the customer class exceeds the defined threshold of customers being classified as "high-value", the constraint for the respective customer will be added to the model.

$$x_{ikn}\pi_{ikn}\beta_{ikn} - x_{ikn}(1 - \beta_{ikn})c = 0 \tag{9}$$

If no item is shipped in a certain demanded category k within a defined threshold, a penalty P is being applied (see 2.3) by constraint (10)

$$P_{ikn} = \rho_i(1 - x_{ikn}) \tag{10}$$

with ρ being the customer-class-specific penalty value.

For simulation purposes, there is one constraint needed to limit the total number of shipped items within a sales period to the stock available in inventory, formulated as constraint (11) with I being the inventory level for product x.

$$\sum_{i=1}^{j}\sum_{k=1}^{l_i}\sum_{n=1}^{m_{ik}} x_{ikn} \leq I_x \tag{11}$$

2.3 Parameterization and Variables

The model behaviour is subject to parameters and initialized variables. Variables change during runtime while parameters mostly remain constant. Both arguments can be used to test and analyze changes in model behaviour. In order to do so, baseline values have been defined and varied. An overview can be found in table A.4 in the appendix. In the following, first, the arguments used for the demand generation and general multiple-period simulation will be explained. Subsequently, it will be focused on the arguments of specific interest for solving the optimization problem.

The number of customers per period, the number of items in the assortment as well as the pre-defined item sets are fixed first. Each period the number of customers is randomly drawn between a defined maximum and a minimum number, following a normal distribution. Items are split into four categories and item sets include items from one of these categories. It is assumed that, on average and over an infinite horizon, certain sets with specific purchase probabilities will become predominant. Thus, for studying the principal effects of the model, no more complex inventory or assortment simulation was chosen. The purchase probabilities are boosted by the customer demand class. The factors used proved to be sufficient but are only there for simulation purposes. When businesses have detailed purchase data available, this factor should be set with respect to the actual data. Each customers is associated with a demand class. The demand class is derived from the customers' initial ordering and purchasing volume under consideration of an assumed churn rate and frequency of future orders. Likewise to the purchase probabilities, these values can be improved by supporting with additional data from business for generating company-specific insights. For the thesis at hand, the calculation for CurveCatch can be found in appendix A.1, which follows an approach explained in 1.2 [Berger and Nasr, 1998]. Customers are split into groups based on their purchased items. For CurveCatch, only little data was available and there was hardly any information available about retention (follow-up orders). Thus, only the purchased products of the first order of each customer were incorporated to calculate the potential customer lifetime value. Other variables, like retention rates, were assumed. For the model and the simulation environment, the actual customer lifetime value does not matter. The lifetime value, as well as many other variables and parameters, is subject to assumptions. So, only the distribution of customers' lifetime value is important for the model, which is sound because it is a ratio rather than a concrete number. For validation, the derived demand classes and the identified distribution among the customer base were checked with what has been found in related research by [Dadouchi and Agard, 2018]. Their results follow a similar distribution why a ratio of 20% highest lifetime value customers has been set as starting point for the classes' distribution. The values of the class penalties have been narrowed down based on the first insights the model generated. Eventually, a certain set of penalties has been taken as a reference point for further analysis. The actual values of the penalties will be different in a real-world setting, depending on all the other variables but mainly because of possibly different expected profits per customer. In a manner similar, other parameters related to the optimization and simulation have been set. These values depend on the respective business setting and the overall objective.

3 Discussion

The issue of customer relationship management to reduce churn and improve retention rates is well known and extensively studied. Whether customers churn or not depends on multiple factors influencing their satisfaction. Retention rates and the composition of the customer base directly affect profits [Braun and Schweidel, 2011, Godinho de Matos et al., 2018, Lemmens and Gupta, 2020]. In the following, several given business scenarios for try-before-you-buy companies are studied. The previously introduced periodic optimization model is executed for a number of periods. The simulation results give insights into the effects of applied policies in a specific situation by varying the parameters introduced in 2.3. The results are analyzed based on insights from inventory levels, cumulative profit and the ratio of high-value customers in each period. The time series plots shown in the following also include confidence intervals on an 80% significance level. Larger samples (= more executions of the simulation with the same parameter set) of observations support validating the observations. Due to resource limitations, for all experiments a maximum sample size between 40 and 120 simulations has been chosen. Subsequently, possible managerial actions to achieve certain strategic goals or objectives are derived. Special attention is paid to the changes imposed by the variation of treatment of different customer groups. In the beginning, a baseline situation is introduced and explained. This baseline serves as a reference point for comparing other simulated parameter combinations. In all business situations, three different customer treatments have been tested: no differentiation in treatment between customer classes, only preferring the highest-lifetime value customer class and preferring all customer classes but not the low-lifetime value customers.

Observation 1: Privileging high-value demand classes leads to the highest overall profit when demand and supply are sufficient.

In the first scenario, the parameters were set to the baseline shown in A.4. The results (see figure 3 and 4) show that if there is unlimited demand from all demand classes, from a profit-focused perspective high-value customers should be served.

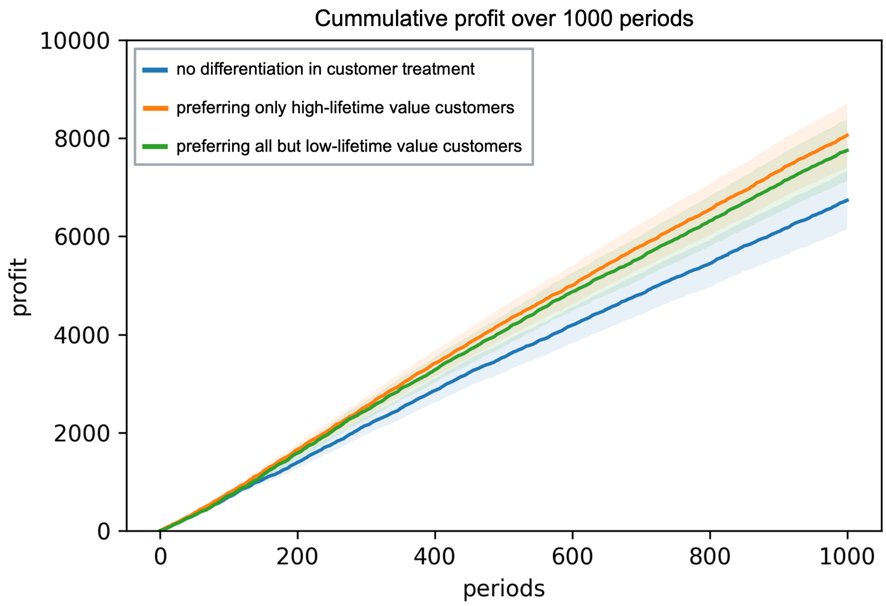

Figure 3: Baseline model profit

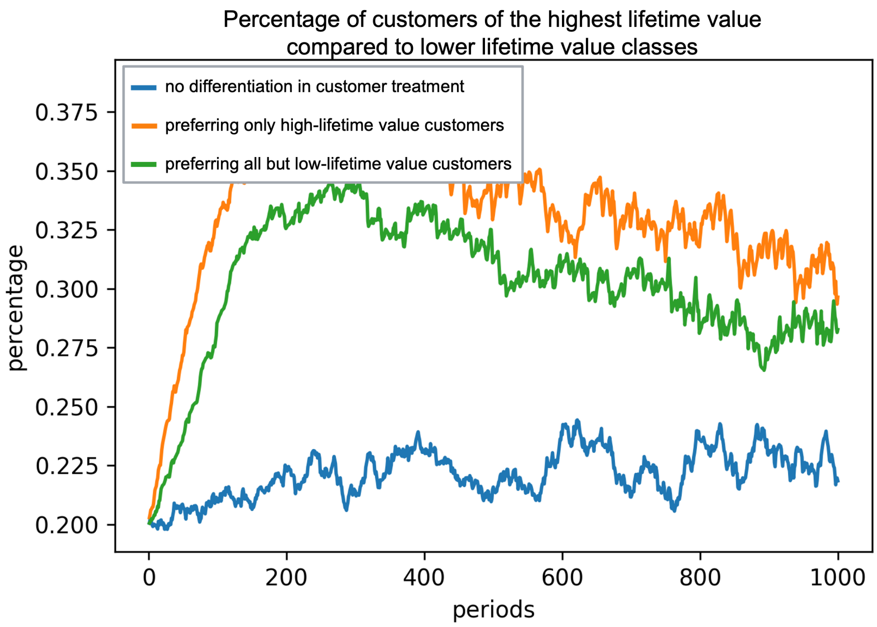

Figure 4: Baseline model share of highest class customers

By serving the highest customer class, their share increases, unfolding a self-accelerating effect up to a certain level. Once this level is reached, the ratio of high-value customers compared to lower classes starts decreasing. This is due to excess demand and because of the fact that high-lifetime value customers are expected to purchase more products. Since a growing number of customers require a larger number of products, not all of the high-class customers can be satisfied anymore, given the maximum inventory levels remain static and do not increase over time. The ratio decreases because, at that point, a share of premium customers can't be satisfied anymore, e.g. because of a lack of resources like inventory levels. However, in the baseline scenario the advantage over satisfying also customers of medium lifetime value is little and cannot be completely proven at the 80% significance level with the number of simulations (=samples).

If there is an unlimited supply for all items in inventory, in addition to the unlimited demand, the self-accelerating effect increases and drives profits further. More insights into the results of this scenario are shown in the appendix section A.5 in figure 12 and 13.

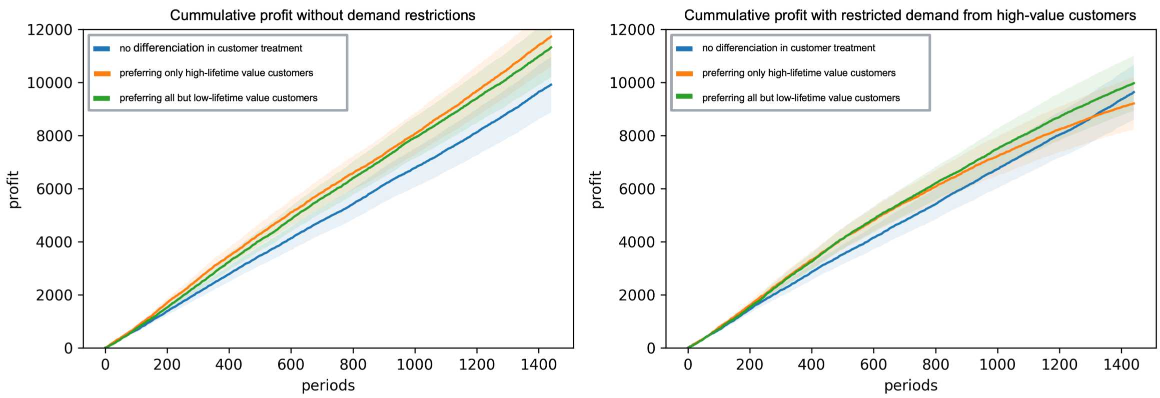

Figure 5: Baseline model (left) compared to a scenario where demand is limited (right)

In a market where multiple customer classes, e.g. defined by purchasing power and probability, like in this work, the total number of each class is limited (see 2.1). It might be the case that a company reaches this natural limitation, either because of total numbers if the market size is small or because of strong competitors where it is hard to alienate customers. With a limitation of the demand of the highest-value customer class up to a certain point, purely focusing on high-value customers might not lead to the highest overall profits in the long term, as shown in figure 5. Although the results are not highly significant on the chosen 80% level, the trend is clearly visible and might become significant with higher computational power to simulate a larger number of samples. In the periodic optimization model given, mainly two effects come into place. First, the customer base grows when customers are satisfied, e.g. because of word of mouth and recommendations. This behaviour is simulated and results in lower-class sales overtaking the higher-classes by the sheer number of sales. Second, reserved for higher customer classes cannot be used to fulfill the larger share of lower-class demand. Thus, because the ratio of demand changes, the negative effect of this reservation on the profit growth is over-proportionally.

To decide which customers to target, try-before-you-buy businesses should outline customer target groups aligned with their business strategy. Based on the results given, it could be a possibility to privilege higher-class customers first and later focus on all customers equally.

Observation 2: Customers' budgets in target groups are the most significant differentiation feature.

A key feature in customer class differentiation is the available budget of a customer, which reflects the possible expected spending [Berger and Nasr, 1998]. The gap between budgets can be substantial [Dadouchi and Agard, 2018]. In this scenario, therefore, the parameters defining the average budget of the different customer classes were set closer to each other.

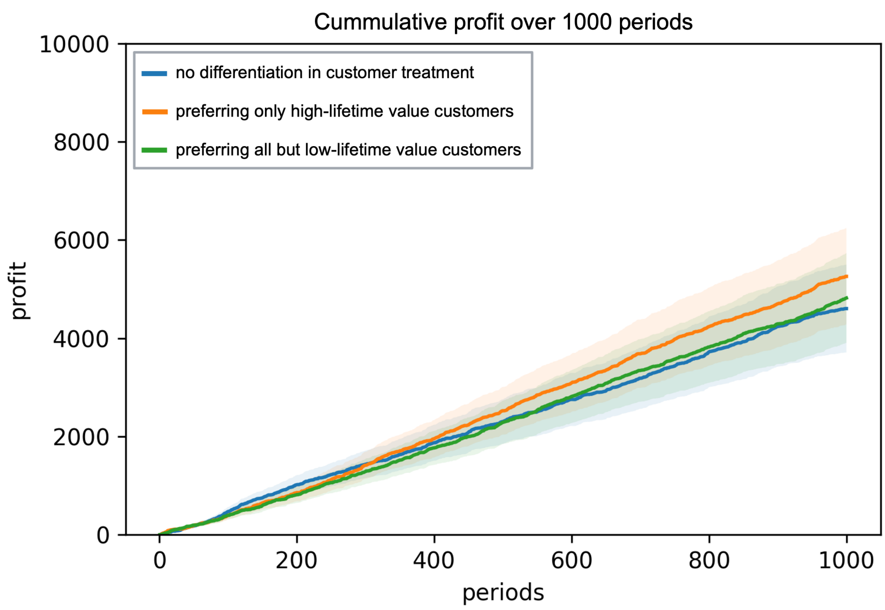

Figure 6: Profit in a scenario with only little differences in customers' purchasing power

In figure 6, it is shown that serving customers differently is less successful when budgets are similar among groups. However, a deep understanding of a company's customers is needed since other factors might affect the purchase behaviour too. Additionally, since budgets are more evenly distributed in this simulated business situation, there are no customers with exceptionally high purchasing power who would rapidly drive profits with a low number of shipped boxes. Also, the relevant factors for making shipping decisions in try-before-you-buy business models depend on factors like purchase probabilities or expected return duration per customer too. In the model at hand, these factors were averaged among customers within a class or among all classes in general. In a real-world market situation, there might be other crucial differentiation features, even if budgets are not. A detailed analysis has to be made on a company level, assessing the respective relevant features for customer segmentation. For example, higher cost of return handling for the retailer makes it more relevant to target the right customer than the ones with the highest budget (see more details in figure 14 in appendix section A.5). Differentiating between all classes rather than just high- and low-lifetime value becomes more important since return handling cost will significantly impact the relevance of customers' budgets compared to their purchase probabilities. A higher cost of handling returns means a larger negative impact on profit if products are sent back by customers. When the cost of return handling is high, customer differentiation based on purchase probabilities becomes more important than based on expected average spending.

Observation 3: Differentiating between customer classes cannot be justified when inventory levels are limited or restock periods are large.

When available inventory is very restricted, so that only a small share of demand can be satisfied, there are several drivers affecting the results. In this situation, period's profit is largely dependent on the demand composition. In this experiment, the number of available entities per product has been set to only a fraction of what was available in the baseline scenario. Available inventory is limited, and therefore, only if the requested products are in stock, profits can be generated.

As shown in figure 7 (see more details in the zoomed-in figure 15 and the distribution among classes in figure 16 in the appendix section A.5), there is not much room for optimizing profits based on customer classes and their respective purchasing behaviour, if demand from high-value customers is out of stock anyways.

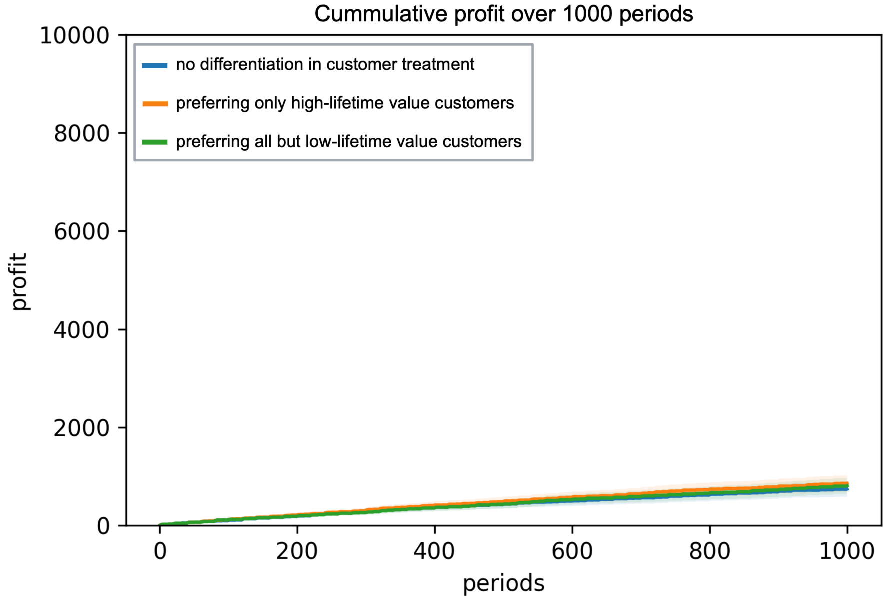

Figure 7: Model profits when inventory levels are limited

Lower-class customers might return more items, which get added to the limited or even empty inventory again and can be sold later, whereas higher-class customers are more likely to keep products instead of sending back. Also, if only a few items can be shipped to each customer because of limitations, customers with lower purchasing power do not run out of budget and do not have to send back items. Overall, when limited stock levels are the active constraint, all these effects average out and do not justify different treatment of customer classes based on profit generation. However, there might be other reasons why focusing on a particular group of customers could be the preferred strategy. Maybe inventory is limited only for a specific time period and there will be large supply in the future. Also, higher number of returns mean higher administrative effort, requiring more staff members and possibly larger facilities.

A similar situation can be observed when there is a larger number of periods between restocking inventory, which means it takes longer to refill up stock levels. Items run out of stock, leading to a similar result where differentiation is less important or even possible anymore, as can be seen in figure 8.

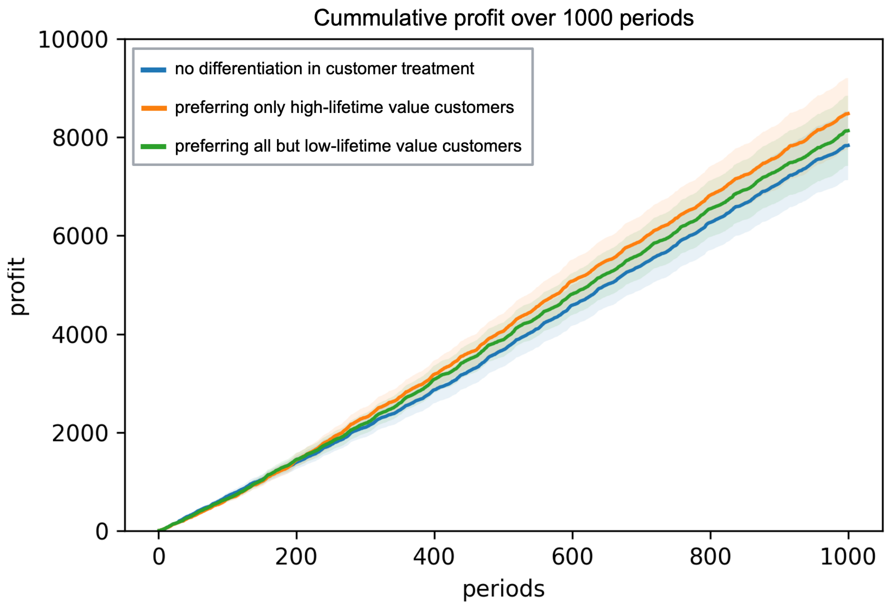

Figure 8: Longer periods until restock deplete inventory

Reserving a certain number of items for higher-class customers does not necessarily significantly affect this result (see figure 17 in appendix section A.5). This is mainly due to the fact that optimization is performed for each single period. The reservation level might be too low for higher-class customers in multiple up-following periods, so the inventory runs out of stock even with this given threshold. On the other hand, the reservation level might be too high, preventing from shipping to lower-class customers to generate additional profits without loosing higher-class customers. Based on the simulation results, effects are averaging out when only considering a single-period optimization. Thus, the decision whether to keep a certain amount of items in stock for premium customers or not, must be seen in alignment with the overall business strategy and in the context of other parameters like brand perception.

Observation 4: Solely implemented stricter return policies can lower cost imposed by working capital.

In this scenario, it is assumed that a company implements stricter return policies. In the simulation environment, the number of average days it takes for an item to get shipped back to the company is set to half the baseline value. In figure 9, it can be seen that the overall profit is similar to what can be expected in the baseline scenario.

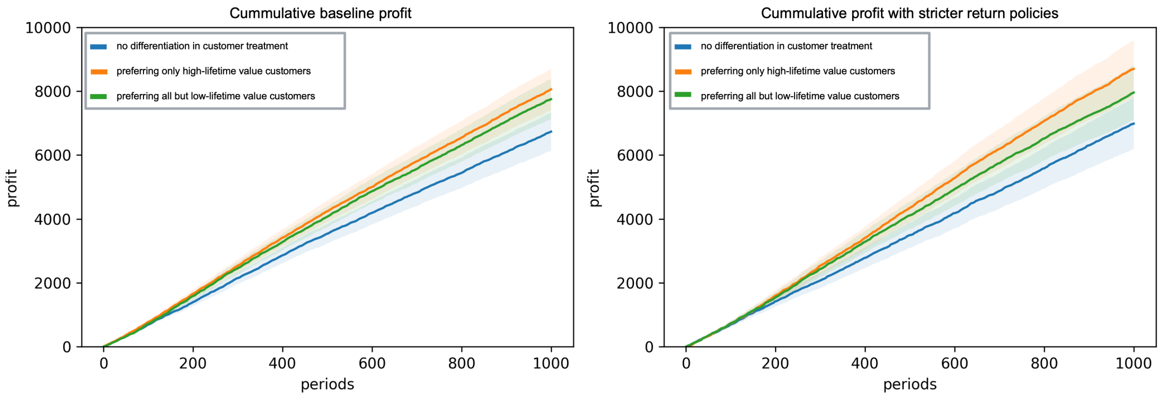

Figure 9: Baseline model (left) compared to a scenario with stricter return policies (right)

A shorter return duration means that more inventory is in stock at a certain point in time and can be shipped directly again instead of being bound on the road. So, consequently, more lenient return policies (=longer periods until items are back in inventory) require larger inventory levels to serve demand, even when products are in transit. If returns take longer, a larger growing rolling inventory imposes higher investments for working capital. In a monopoly situation, businesses should urge customers to send items back as soon as possible. However, because of market pressure and fierce competition, this might be difficult to assert [He et al., 2022]. Thus, companies have to make a trade-off on which return policy they should follow in their respective business environment in order to align with specific strategic goals. Quantification could be performed by comparing the differences in profit of lenient return policies compared to the loss imposed by churning customers because of the stricter return policy. More lenient return policies might be suited when mainly premium customers are targeted since potentially higher working capital costs can be compensated by higher average profit per buyer. Following stricter return policies with medium-class customers might be preferred to reduce costs. However, because of competition, demand could be affected negatively, reducing overall profit in the long run.

Observation 5: Accurate customer behaviour prediction is required for effective optimization.

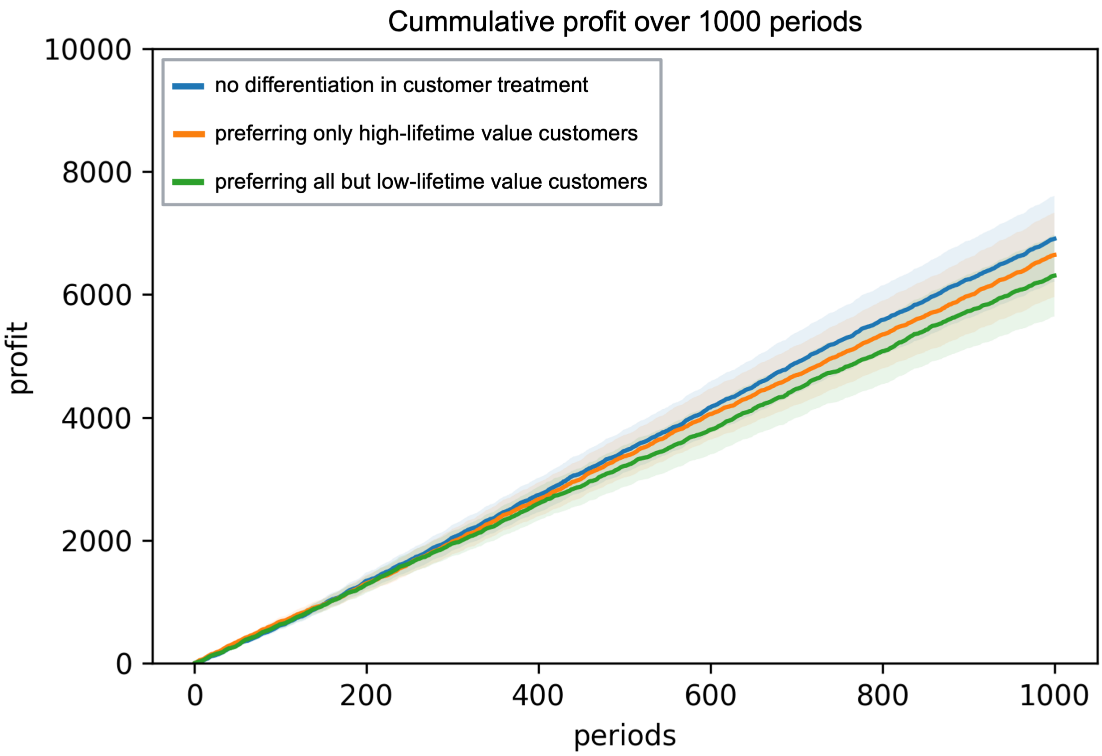

Figure 10: Profit in a scenario where the prediction of customer preferences and budgets is of insufficient accuracy

For decision-making based on customer behaviour, achieving high accuracy rates when predicting the respecting attributes is essential. Only if the preceding customer-style-fit prediction and the customer classification are sufficient, sensible optimization can be performed subsequently [Balaram et al., 2022, Park et al., 2022]. For the model introduced, a robustness check has been performed to understand the effects of misalignment between predicted values and actual customer behaviour. Particularly, the budget, purchase probability and a combination of both have been studied since these factors have been used to assign customers to their respective classes in the simulation environment. In this scenario, during the simulation the budgets and/or the purchase probabilities have been assigned randomly during the simulation of the purchase decision. Thus the values in the last component (see simulation and model structure in 2.1) differ from the initially determined values in the demand simulator. The misalignment between predictions and actual values results in lower total profits, as can be seen in figure 10. Additional insights about separated effects of wrongly predicted purchase probabilities and budgets can be found in the appendix section A.5 in figure 18, 19 and 20. As outlined in observation 2, the customers' budgets have been found to be the most significant attribute in segmentation. If the predicted budget does not correspond to the actual one, customers are served in an insufficient way. Customers actually belonging to a higher class are expected to spend less than they would do, leading to reduced products in their ordered boxes. These customers remain unsatisfied since the items demanded are not shipped. On the other hand, actual lower-class customers might receive more products than they are able to purchase, leading to excessive return rates. Similar effects can be studied when the assumed purchase probability predictions per class differ from the actual ones. Here it can be observed that the effect is not that significant if the overall probabilities are still high. If the actual probabilities end up higher than expected, profits also are higher than predicted by the model. Thus, the lower the actual probabilities are, the more problematic the misalignment gets. Consequently, having a robust prediction model is essential for using its results in an optimization model.

Summing up observations, the research questions introduced in section 1.1 can be answered. Models supporting box content and quantity decisions in companies following a customer-centric try-before-you-buy business model can be set up similarly to assortment optimizations. To make sure the most important customers are satisfied, penalty costs can be applied. Multiple up-following sales periods are hard to reflect in a single optimization problem, however, a simulation environment can help examine the performance of the decision model.

To answer the second research question, the basic model components in 2.1 and parameters in A.4 should be considered. In principle, customer behaviour is needed as input for the decision support model and subsequently for simulating the actual customer decisions. Most important for replicating the purchasing behaviour of customers are simulated purchasing probabilities of shipped products and the available budget customers are willing to spend. Additionally, it is important to reflect market dynamics, like the growth of the customer base in certain groups when customers are satisfied.

Finally, the third research question aims to give general insights into the managerial challenges in try-before-you-buy companies. Differentiation between customer groups will lead to improved profit in certain market circumstances, especially when there is a high demand from customers belonging to classes associated with high purchasing power and lifetime value. However, in a real-world situation, there is only a limited number of the most valuable customers available. It depends on the exact scenario which strategy pays off the most for TBYB companies. Examples of potential scenarios are given above. Besides profit, companies might also want to reach other targets, like sustainability goals or a specific brand image.

4 Conclusion and further research

4.1 Conclusion

In this master thesis, an optimization model for supporting quantity and content decisions when packing boxes in try-before-you-buy has been set up. For multiple up-following sales periods, the problem of optimizing the allocation of items for shipping to customers has been solved. To account for customer behaviour and market development, a simulation environment has been created. The results have been studied by conducting sensitivity analysis and examining the implications when varying different parameters and variables. The results help in understanding under which circumstances customer differentiation drives profit further and how a customer-centric approach could be followed while optimizing TBYB box content.

It has been shown that customers belonging to classes related to high purchasing power and lifetime value are especially valuable to satisfy. Satisfied customers are less likely to churn and more likely to order again. Also, they spread positive word of mouth and can help increase the customer base. Since the most valuable types of customers are limited in the market, trade-offs in serving them must be made. Dedicating resources to satisfy them is an important measure, but middle-class customers should not be discriminated against too much if all classes should be targeted. Middle- and lower-class customers make up for the largest share of customers and also sales volume. On the contrary, by satisfying only high-value customers, a premium-brand strategy could be followed if the market supplies enough high-value demand. Applying penalty cost can be a suitable tool to account for different classes in an optimization model. The exact values, however, are arbitrary and hard to determine.

If inventory levels are low and product supply is restricted, differentiating between customer types does not necessarily result in higher profits. Stock levels are the limiting factor if demand is high and every sale contributes to the total profit. Differentiation could be important if different customer classes would react differently if they remained unsatisfied. In this case, reserving a certain amount of items for high-value or sensitive customers could be beneficial. Otherwise, there is no significant effect of holding a threshold of products for a certain customer group.

For sufficient optimization and simulation along multiple periods, accurate input data is needed. Thus, the prediction of customer data, purchase behaviour and demand estimation is a crucial success factor. The identified most important feature is customer budget. If the prediction of the budget is much different from the actual value, the box content decisions supported by the model are far off the optimal ones. This leads to significantly reduced profits and contradicts any customer-centric approach, where customer profiles determine their individual treatment.

4.2 Limitations and further perspectives

When simulating customer behaviour and market patterns, a number of limitations are given. Also, the optimization model is subject to some restrictions and assumptions, mainly due to the complexity of the business model replicated and analyzed. In the following, a short introduction to relevant limitations is given.

This thesis provides an optimization model considering one sales period at a time. The model described evaluates the optimum for the current period but does not include past shipments or forecasts of upcoming periods. Adding these additional periods would make computation much harder, requiring tremendous amounts of computing power [Armstrong and Jacobson, 2003]. [Tulabandhula et al., 2022] showed an approach for solving assortment optimization at scale by using binary search and approximating the near-optimal solution. [Ayoub and Poss, 2016] decomposed the problem set to simplify the computational problem. For large-scale applications, calculating optimal solutions for a large number of periods, [Glomb et al., 2022] suggest using a rolling-horizon approach.

Managerial insights can vary when decisions regarding the future must be made in the current sales period. Thus, further research can expand the optimization model and simulation environment by including and comparing these different approaches for approximating the global optimum over a certain number of periods and the imposed effects on possible business strategies.

Instead of separating the recommendation system and optimizer, both systems could also be combined. With an approach like shown by [Chen et al., 2008], the recommendation system not only predicts customer preferences but also includes profitability for the retailer. The advantage of the combined approach lies in its ability to account for cross-selling opportunities and possible relationships between purchase probabilities of multiple products. For example, if multiple items of the same item category are shipped to customers, they might be less likely to buy all of them and purchase probabilities of additional products might decrease. So, relationships between items can be included more easily, but there are also downsides to the combined approach. To generate such a detailed prediction and get the relationship between products right, much data is needed. In a real-world business situation, where several thousand SKUs are kept in the assortment, it might be difficult to gather the required amount of information. Additionally, with different products and product types, it might be difficult to identify the relevant features for an accurate prediction. By splitting the problem in two systems, uncertainty in each system can be tackled separately. [Chen et al., 2008] showed that the combined approach is superior to recommendation systems merely focusing on customer-style fit, however, no comparison to a system consisting of a separated recommender and optimizer has been made. Further work could compare the combined approach with separated systems like those used in this thesis.

Computational power is not only a limiting factor with respect to the model complexity but also limits the robustness of the model and simulation results. For computing reasonably robust confidence intervals, an adequate number of repeated simulations is needed. With a larger number of simulated sales periods, the need for narrow confidence intervals becomes even more important. Margins add up and might corrupt results in later periods. The model optimizes box content and quantity based on profit. It might be of interest to not only focus on profit but also include additional business objectives like a target growth of the customer base or a percentage of a specific type of items (e.g. "organic" products). For adding these objectives, performance metrics must be found and quantitatively added to the model. When multiple objective metrics are included, the trade-offs between them must be represented in the model, e.g. by adding weights to each of them. From a managerial perspective, such an expanded objective function leads to a result that is more aligned with a holistic business strategy. By optimizing for a varying set of objectives, the results of differently weighted objective functions on key metrics like profit could be analyzed.

Little data on customer behaviour and relationships between product categories were available for conducting this thesis. Thus, there are underlying assumptions. Customer samples and product set sizes used for simulation are small and behaviour is assumed to average out among customer groups. For some variables, industry benchmarks and reference values suggested by preceding research have been used, e.g. for the return handling cost and purchase probabilities. The distribution of customer classes has been derived from data and validated with an example given by [Dadouchi and Agard, 2018], showing a similar pattern. Thus, before applying the shown model to a particular real-world business situation, proper data sources have to be identified and more data has to be gathered first. Additionally, variables like cost penalties for different customer groups or the maximal inventory level have been held constant over all simulated sales periods. Further projects could focus on deriving managerial suggestions on how and under which circumstances these variables and parameters could be adjusted and which long-term effects could be expected.

Finally, the thesis at hand requires an accurate customer-style fit prediction. A recommendation system must provide a list of suitable items, sorted by preference rank and complemented by purchase probabilities, for each product demanded by customers. Also, customer classification must precede the optimization process. As shown in this thesis and argued by [Dadouchi and Agard, 2018, Balaram et al., 2022] and [Park et al., 2022], prediction results are only useful for optimization if they yield high recommendation and classification accuracy. Before box content and quantity can be optimized, therefore, try-before-you-buy companies should focus on maintaining well performing recommendation systems and customer classification.

References

[Altug and Aydinliyim, 2016] Altug, M. S. and Aydinliyim, T. (2016). Counteracting strategic purchase deferrals: The impact of online retailers' return policy decisions. *Manufacturing and Service Operations Management*, 18:376–392.

[Armstrong and Jacobson, 2003] Armstrong, D. E. and Jacobson, S. H. (2003). Studying the complexity of global verification for np-hard discrete optimization problems. *Journal of Global Optimization*, 27:83–96.

[Ayoub and Poss, 2016] Ayoub, J. and Poss, M. (2016). Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Computational Management Science*, 13:219–239.

[Balaram et al., 2022] Balaram, A., Perdikaki, O., and Galbreth, M. R. (2022). Bracketing of purchases to manage size uncertainty: Should online retailers be worried? *Naval Research Logistics*, 69:783–800.

[Berger and Nasr, 1998] Berger, P. and Nasr, N. (1998). Customer lifetime value: Marketing models and applications. *Journal of Interactive Marketing*, 12.

[Bernstein et al., 2021] Bernstein, F., Chakraborty, S., and Swinney, R. (2021). Intertemporal content variation with customer learning.

[Bhattacharjee and Ramesh, 2000] Bhattacharjee, S. and Ramesh, R. (2000). A multiperiod profit maximizing model for retail supply chain management: An integration of demand and supply-side mechanisms. *European Journal of Operational Research*, pages 584–601.

[Borgonovo and Plischke, 2016] Borgonovo, E. and Plischke, E. (2016). Sensitivity analysis: A review of recent advances. *European Journal of Operational Research*, 248:869–887.

[Brahimi et al., 2017] Brahimi, N., Absi, N., Dauzère-Pérès, S., and Nordli, A. (2017). Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research*, 263:838–863.

[Braun and Schweidel, 2011] Braun, M. and Schweidel, D. A. (2011). Modeling customer lifetimes with multiple causes of churn. *Marketing Science*, 30(5):881–902.

[Castéran et al., 2017] Castéran, H., Meyer-Waarden, L., and Reinartz, W. (2017). *Modeling Customer Lifetime Value, Retention, and Churn*. Springer International Publishing.

[Chen et al., 2008] Chen, L. S., Hsu, F. H., Chen, M. C., and Hsu, Y. C. (2008). Developing recommender systems with the consideration of product profitability for sellers. *Information Sciences*, 178:1032–1048.

[Dadouchi and Agard, 2018] Dadouchi, C. and Agard, B. (2018). Lowering penalties related to stock-outs by shifting demand in product recommendation systems. *Decision Support Systems*, 114:61–69.

[DeCroix and Arreola-Risa, 1998] DeCroix, G. A. and Arreola-Risa, A. (1998). Optimal production and inventory policy for multiple products under resource constraints. *Management Science*, 44:950–961.

[Dimoka et al., 2012] Dimoka, A., Hong, Y., and Pavlou, P. A. (2012). On product uncertainty in online markets: Theory and evidence theory and evidence1. *Source*, 36:395.

[Glomb et al., 2022] Glomb, L., Liers, F., and Rösel, F. (2022). A rolling-horizon approach for multi-period optimization. *European Journal of Operational Research*, 300:189–206.

[Godinho de Matos et al., 2018] Godinho de Matos, M., Pedro, F., and Rodrigo, B. (2018). Target the ego or target the group: Evidence from a randomized experiment in proactive churn management. *Marketing Science*, 37(5):793–811.

[Grewal et al., 2018] Grewal, D., Motyka, S., and Levy, M. (2018). The evolution and future of retailing and retailing education. *Journal of Marketing Education*, 40:85–93.

[Guzman et al., 2021] Guzman, S. P., Gil-Costa, V., and Marin, M. (2021). A simulation tool to provide alternative products in out-of-stock situations for b2b companies. *Proceedings - Winter Simulation Conference*, 2021-December.

[He et al., 2022] He, J., Liu, J., Yang, F., and Zhao, X. (2022). Returns freight insurance policy and the impact from a bops retailer. *Operations Research Forum*, 3.

[Hong and Pavlou, 2014] Hong, Y. and Pavlou, P. A. (2014). Product fit uncertainty in online markets: Nature, effects, and antecedents. *Information Systems Research*, 25:328–344.

[Hübner et al., 2020] Hübner, A., Schäfer, F., and Schaal, K. N. (2020). Maximizing profit via assortment and shelf-space optimization for two-dimensional shelves. *Production and Operations Management*, 29:547–570.

[Kahn et al., 2018] Kahn, B. E., Inman, J. J., and Verhoef, P. C. (2018). Introduction to special issue: Consumer response to the evolving retailing landscape. *Journal of the Association for Consumer Research*, 3:255–259.

[Khouja et al., 2019] Khouja, M., Ajjan, H., and Liu, X. (2019). The effect of return and price adjustment policies on a retailer's performance. *European Journal of Operational Research*, 276:466–482.

[Lemmens and Gupta, 2020] Lemmens, A. and Gupta, S. (2020). Managing churn to maximize profits. *Marketing Science*, 39(5):956–973.

[Li et al., 2019] Li, Y., Li, G., and Tayi, G. K. (2019). Try-before-you-buy: Online retailing strategy with customer self-mending. *Electronic Commerce Research and Applications*, 36.

[Li et al., 2013] Li, Y., Xu, L., and Li, D. (2013). Examining relationships between the return policy, product quality, and pricing strategy in online direct selling. *International Journal of Production Economics*, 144:451–460.

[Li, 2007] Li, Z. (2007). A single-period assortment optimization model. *Production and Operations Management*, 16:369–380.

[Ma and Fildes, 2017] Ma, S. and Fildes, R. (2017). A retail store sku promotions optimization model for category multi-period profit maximization. *European Journal of Operational Research*, 260:680–692.

[Maiti and Giri, 2017] Maiti, T. and Giri, B. C. (2017). Two-period pricing and decision strategies in a two-echelon supply chain under price-dependent demand. *Applied Mathematical Modelling*, 42:655–674.

[Marc et al., 2021] Marc, R., Jakob, K., Nannan, X., Rüdiger, Z., and Juho, H. (2021). The use of augmented reality in retail: A review of literature. *53rd Hawaii International Conference on System Sciences (HICCS).*

[Park et al., 2022] Park, Y. S., Sim, J., and Kim, B. (2022). Online retail operations with "try-before-you-buy". *European Journal of Operational Research*, 299:987–1002.

[Reinartz et al., 2005] Reinartz, W., Thomas, J. S., and Kumar, V. (2005). Balancing acquisition and retention resources to maximize customer profitability. *Journal of marketing*, 69(1):63–79.

[Smink et al., 2019] Smink, A. R., Frowijn, S., van Reijmersdal, E. A., van Noort, G., and Neijens, P. C. (2019). Try online before you buy: How does shopping with augmented reality affect brand responses and personal data disclosure. *Electronic Commerce Research and Applications*, 35.

[Tao and Xu, 2018] Tao, Q. and Xu, Y. (2018). Fashion subscription retailing: an exploratory study of consumer perceptions. *Journal of Fashion Marketing and Management*, 22:494–508.

[Topkis, 1968] Topkis, D. M. (1968). Optimal ordering and rationing policies in a non-stationary dynamic inventory model with n demand classes. *Management Science*, 15(3):160–176.

[Tulabandhula et al., 2022] Tulabandhula, T., Sinha, D., and Karra, S. (2022). Optimizing revenue while showing relevant assortments at scale. *European Journal of Operational Research*, 300:561–570.

[Şen and Zhang, 1999] Şen, A. and Zhang, A. X. (1999). The newsboy problem with multiple demand classes. *IIE Transactions (Institute of Industrial Engineers)*, 31:431–444.

List of Figures

A Appendix

A.1 Calculating customer lifetime value and demand classes

The purchases of the initial order of customers at CurveCatch for the data sample analyzed follows the distribution shown in figure 11. Also, the relevant percentiles are shown below.

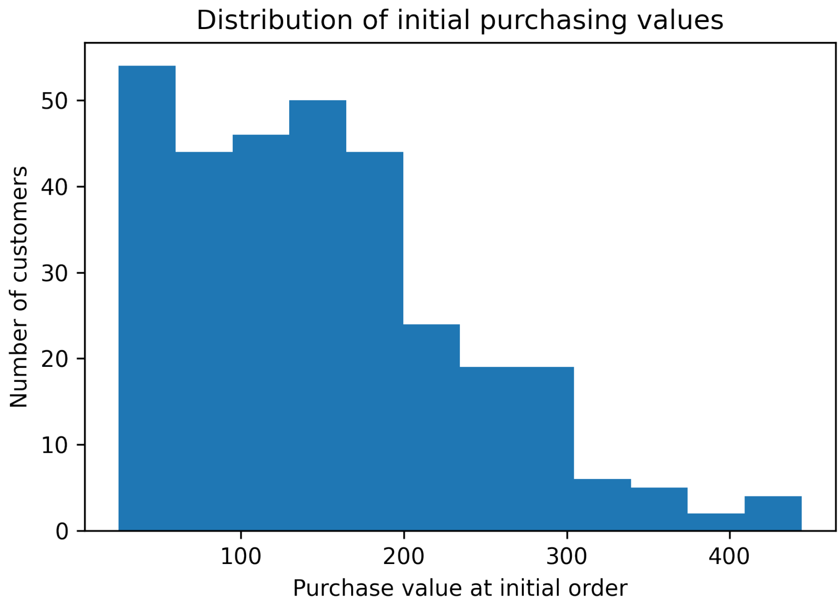

Figure 11: Purchase value of initial orders at CurveCatch

Percentile	Purchased value of initial order
40%	120.31
80%	223.57
99%	431.92

Table 1: Percentiles of initial order values

The customer lifetime value has been calculated by applying formula (12) as introduced by [Berger and Nasr, 1998] with ov being the order value, t the number of periods a customer will retain, r the average retention rate and λ a discount factor to discount future incomes compared to earlier ones.

$$CLV = ov * \sum_{i=0}^{t} \frac{r^i}{(1+\lambda)^i} \tag{12}$$

For the available data the resulting lifetime values are shown in table below. Note, that these numbers are subject to many assumptions and might not hold in a real-world situation. However, the ratio of the customer groups' lifetime values are of interest in this thesis.

Customer class	Customer lifetime value
Class 1	352.38
Class 2	654.81
Class 3	1265.04

A.2 Example data structures for customer and demand data

```
{
    "itemsets": [[item1, item14, item8], [item6, item11, item23]]
    "budget" : 521,
    "class" : 1
}
```

Listing 1: Example customer data

```
{
    "ID": 3
    "price" : 90,
    "profit_margin": 0.15,
    "category": "B"
}
```

Listing 2: Example item data

A.3 Python code

The MIP optimization model has been implemented in Python 3.8.8 using the Gurobi Optimiser version 9.5.1 build v9.5.1rc2 under macOS 12.6. Also, all simulation components have been implemented in Python.

The code for the baseline scenario is available on GitHub: `https://github.com/MartinS-AT/MasterThesis.git`

A.4 List of parameters and variables

Argument ID	Name	Baseline Value(s)	Description	Type	Component
1	min_customers	20	start value	variable	Demand Generator
2	max_customers	30	start value	variable	Demand Generator
3	budget_per_class	[500, 250, 100]	actual budget per customer drawn randomly distributed from parameter	parameter list	Demand Generator
4	purchase_probability_class1	1.25	purchase probability factor for class 1 customers	parameter	Demand Generator
5	purchase_probability_class3	0.75	purchase probability factor for class 2 customers	parameter	Demand Generator
6	class1_to_class23_ratio	0.2	initial split between class 1 and class 2 and 3 customers; start value	variable	Demand Generator
7	penalty_per_class	[10, 0, 0]	penalties for each customer class with the first being applied to class 1 customers	parameter list	Optimization Model
8	cost_return	1	assumption: cost account for average varible and fixed cost per item	parameter	Optimization Model
9	stock_level_threshold	20	entities reserved for higher class customers	parameter	Optimization Model
10	customer_class_reservation	1	which customer classes reserved entities can be shipped to	parameter	Optimization Model
11	penalty_rank_threshold	1	defines after how many ranks penalties are applied	parameter	Optimization Model
12	max_box_value	500	max value of all items combined	parameter	Optimization Model
13	max_number_items	10	max number of items per box	parameter	Optimization Model
14	avg_days_return	10	average duration it takes for returns to get back into the inventory	parameter	Inventory Simulator
15	restock_policy	20	days until stock levels of inventory are filled up	parameter	Inventory Simulator
16	customer_growth	0.00025	linear growth for satisfied customers	parameter	Inventory Simulator

A.5 Additional optimization and simulation results

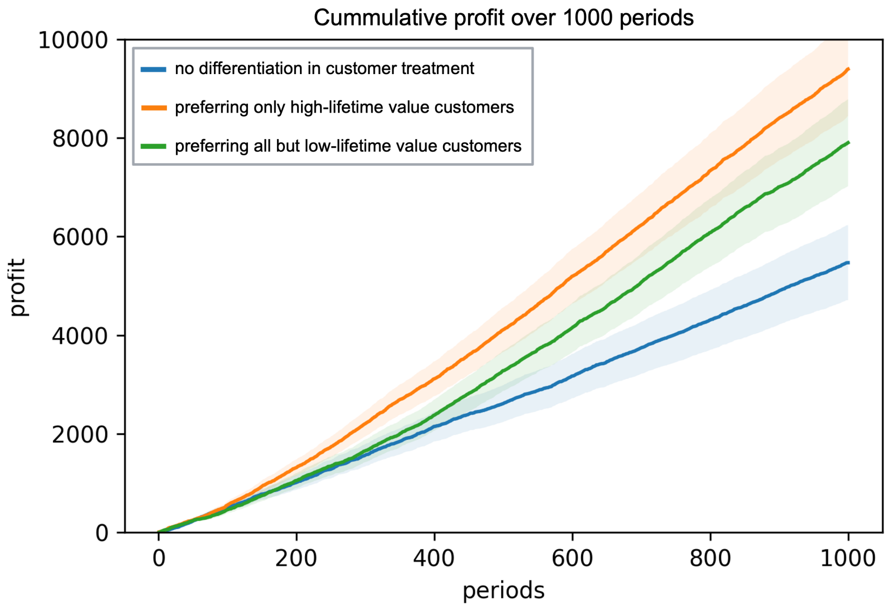

Figure 12: Profit development over time in a scenario where demand (customer classes) and supply (inventory levels) are unlimited.

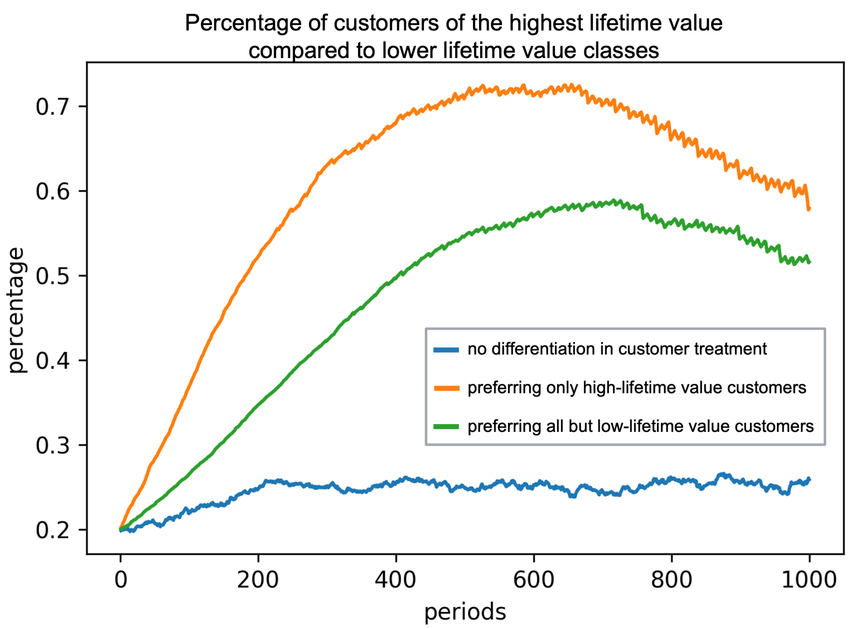

Figure 13: Percentage of highest-value customers in a scenario where demand (customer classes) and supply (inventory levels) are unlimited.

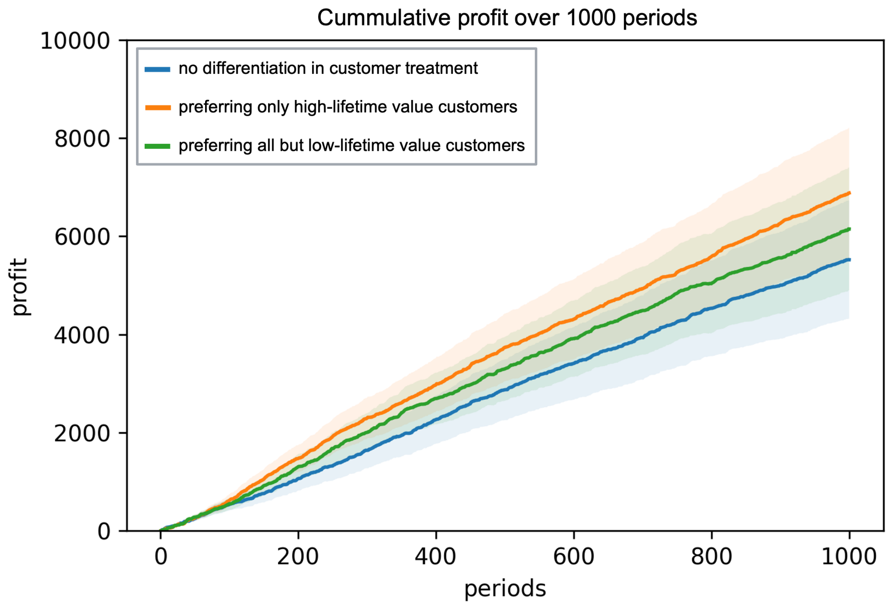

Figure 14: In a scenario where cost of returns are high for the retailer, differentiation between the particular customer groups becomes more important.

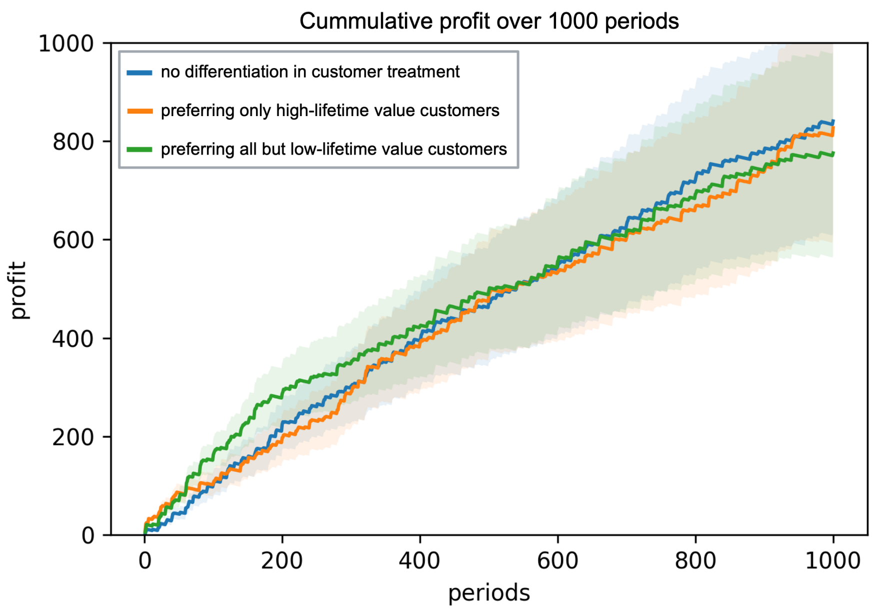

Figure 15: In low-inventory scenarios profits are low in general and the effect of differentiation is limited.

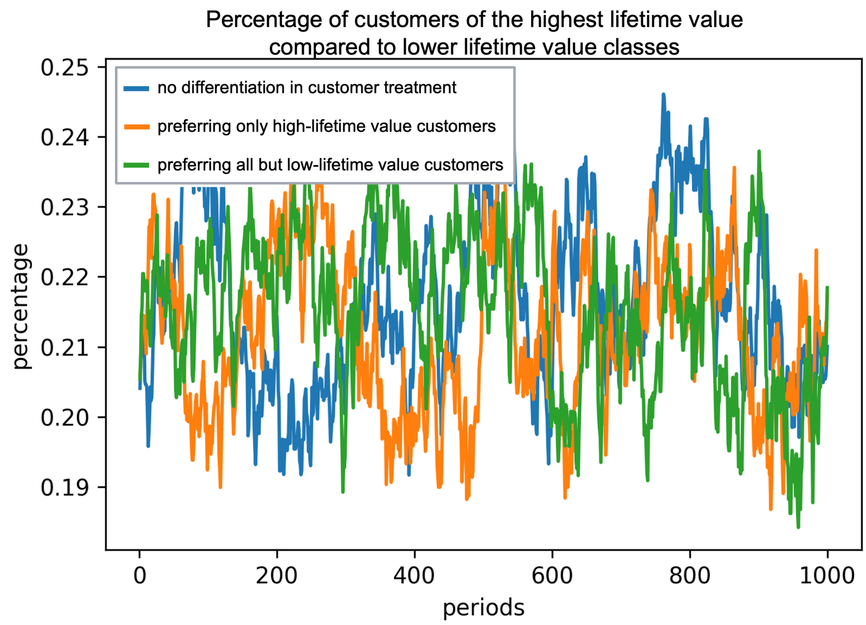

Figure 16: The share of highest-value customers does not change much in low-inventory situations since differentiation is hardly possible.

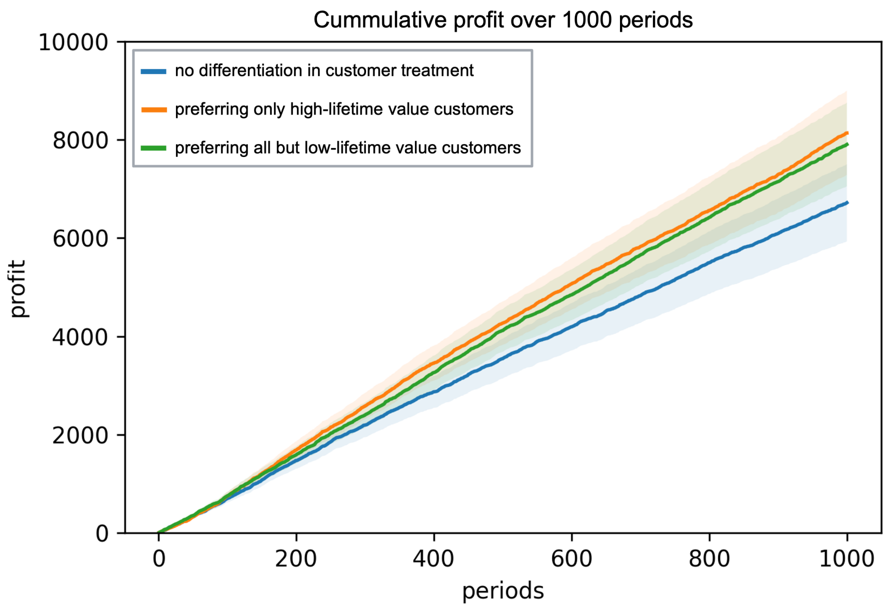

Figure 17: In a scenario where less items are reserved for high-value customers the profit still remains similar to the baseline scenario.

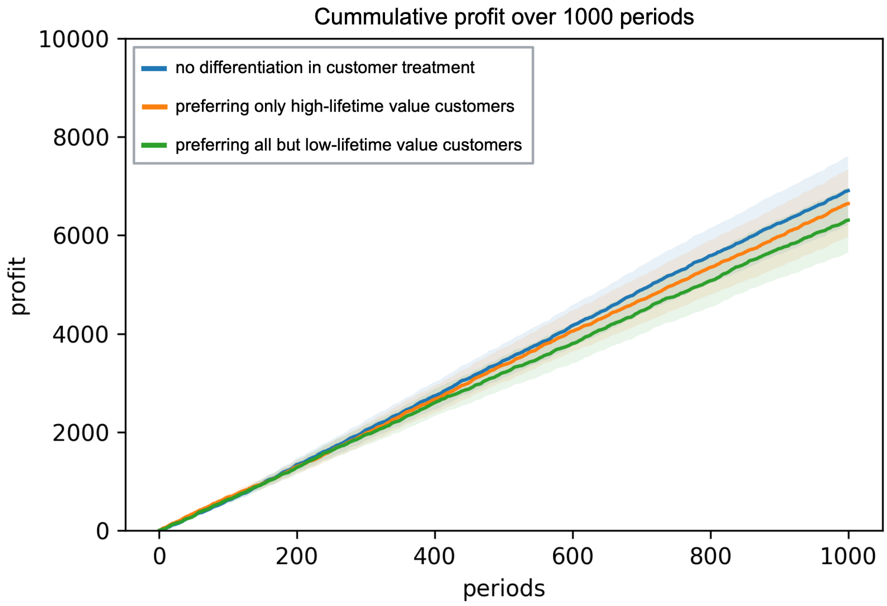

Figure 18: If expected customer budget (=spending) is misaligned with the actual numbers, profit is significantly reduced.

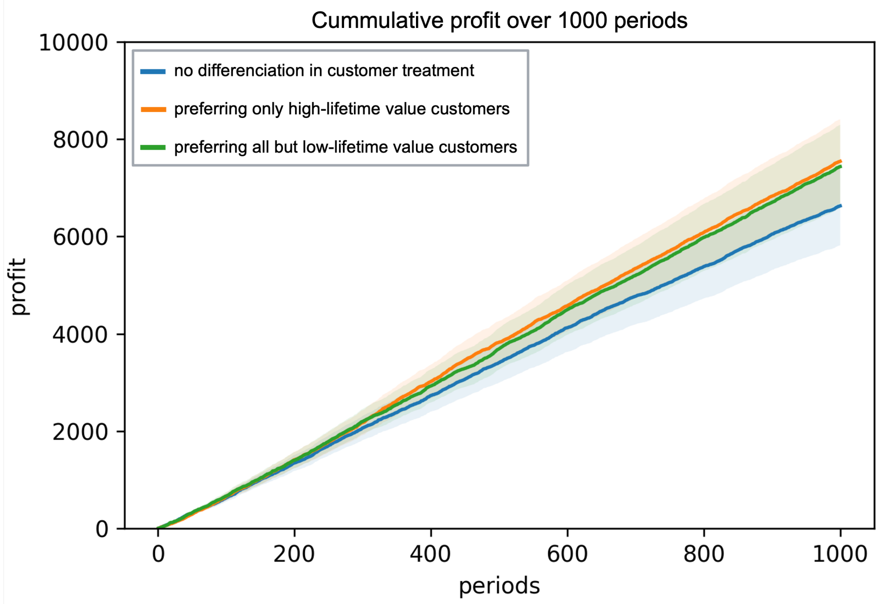

Figure 19: Misaligned expected purchase probabilities can lead to unexpected actual profits.

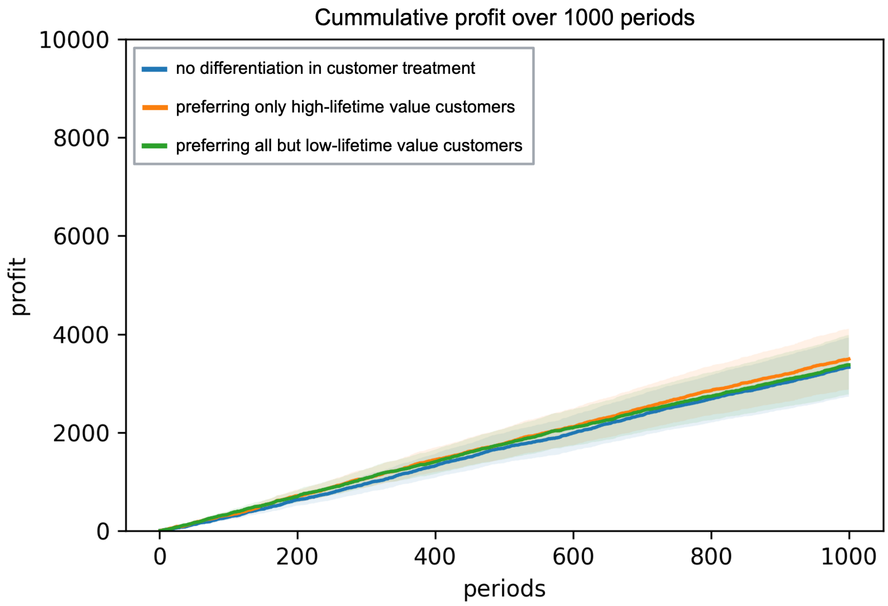

Figure 20: If the expected purchase expected purchase probabilities are inaccurate and estimated too high, targeted resource allocation is difficult and profit drastically reduced.