How Anheuser-Busch InBev is tackling the top challenges of Category Management for modern offline retail channel

Nidhin VC
Rahul Prakash
Growth Analytics Center
Anheuser-Busch InBev, Bangalore, India

Keywords
Offline retail channel, Category Management, assortment optimization, clustering, retailing

Abstract
In this paper we present a practical approach that AB InBev's Global Capabilities Centre ("GCC or ABI or Ab InBev") has developed to solve the challenges of Category Management for Retailers. The approach brings technical rigor from the areas of data science, econometrics, and measurement methodologies very close to business context. This has allowed us to create a solution which is highly contextual and relatable to our business stakeholders. The strength of the presented solution lies in it being a semi-automated framework that allows a wide array of disparate data to be modelled and captures the nuances of different markets - such as socio-demographic profiles, consumption behaviours, local preferences towards beer styles. We also present the ABI created 4C framework to arrive at the optimal assortment recommendation for a Retailer.

Corresponding author: Nidhin VC
Email addresses for the corresponding author: nidhin.chandrasekhar@gmail.com
The first submission received: 15th February 2022
Revised submission received: 27th March 2022
Accepted:7th April 2022

Acknowledgement
We acknowledge the support from AB InBev's Growth Analytics Center (GAC), a unit of GCC, Bangalore. As employees, we had access to data (both owned and syndicated) from our Retailer partners. We have ensured complete confidentiality of our customer data in this paper. The analytical processes proposed in this paper are based on multiple projects that we have executed in different countries. We thank all our business partners from these countries who provided us deep insights into how each Retailer differs from the other, and what would make the most business sense to them from the point of view of Category Management. For extremely helpful feedback and suggestions we would thank the entire CatMan Analytics team.

Introduction

Category Management is a purchasing and retailing concept in which the range of products purchased and sold by a Retailer is broken down into discrete groups of similar or related products; these groups are known as product categories (examples of grocery categories could be beer, washing detergent, toothpastes, etc.). This paper talks about the challenge of Category Management, that in recent years has climbed up the ladder in terms of priority for all the large companies in the Consumer-packaged goods ("CPG"), Fast moving consumer goods ("FMCG") and the retail space due to the opportunity size it brings with it and the gateway it provides to optimize the end consumer's experience today.

Retailers and suppliers have their own Category Managers who lead the Category Management function. Category Managers need to put the shopper and the consumer at the center of their category plan by answering the following questions:

How does the shopper shop?

When and where do they make their purchase decisions? What portion of choices is pre-determined and how much of it is impulsive?

What factors influence their decision? Price, Occasion, Packaging style, Brand Loyalty, etc.?

Who makes the decision to purchase?

Who is the purchase for?

What is the purchase used for? Occasion – is it for outdoor consumption or indoors, for special events or for daily consumption?
How much is bought?
What else did they buy or consider buying?
What type of consumer are they based on demographic, consumption analysis?
How did they feel about the purchase experience?

With the help of the answers to these questions, the focus of Category Management needs to be on a) Assortment Optimization, b) Shelf Segmentation & Navigation c) Omni-channel Optimization, d) Price and Promotion Optimization.

In the scope of this paper, we will focus more on the first three areas. The question of which products to feature in an Assortment comes from the point of view of what the consumer wants.

If the Retailer understands their consumer well, they will not overwhelm them with hundreds of choices for something like soy milk. The presence of umpteen alternatives creates a situation of "brand fatigue" among consumers, wherein decision making of what to buy and finding the right products becomes increasingly tough, and they walk away without making a purchase. Meanwhile, the store also loses out on the opportunity to optimize space usage for products with better sell-through. While on the other hand, limiting the assortment only to those brands which contribute high revenue could lead to losing customers with niche preferences looking for a particular category of product, even though the sales from that category may be less. Hence, it's extremely critical to optimize the assortment that can maximize the revenue and enhance customer experience.

This Paper talks about how the Category Management team at AB InBev is tackling this problem of assortment optimization to enhance consumer experience and achieve a gain in revenue for the retailer. This Paper goes into details of sources of data, significance of collaboration with stakeholders, importance of clustering, pillars of ABI's Assortment Methodology, and an end-to-end description of the framework that we have followed. Even though the presented solution has been created using our experience in the beer market, the parameters that we have used, and the literature referred to could be used for any FMCG industry and could also be extended to other business verticals that require optimization.

The need for Assortment OPTIMIZATION

The problem of what to keep on the shelf, or having an optimized assortment, has for long been a matter of contemplation for retailers and manufacturers alike. Industries with low innovation and seasonality in consumption of products, i.e., industries which sell the same set of products throughout the year (such as hardware, automobile parts, etc.) do not have to undergo a change in assortment every six months, or even a year. But companies in fashion, food and beverage industries are constantly refreshing their assortments based on time of the year and ever-changing consumer preferences.

An optimized assortment goes beyond just optimization of space in stores for revenue maximization. It also helps shoppers in finding their products easier, which can potentially increase their chance of browsing the store and buying more. By enabling an easy decision-making for the shopper, the overall customer satisfaction improves.

Industry Landscape

The way assortment has traditionally been approached was to always keep the top selling SKUs and removing everything else. As per the M L Fisher and R Vaidyanathan (2012) in their HBR article:

Following a survey in which customers said they would like less cluttered stores, Walmart introduced Project Impact, in 2008, removing 15% of the SKUs it carried. Sales declined significantly, and it was forced to roll back most of the changes.

Super Fresh, owned by the grocery retailer A&P, stopped carrying many of its low-selling dry grocery items to allow for an expansion of fresh offerings. But the eliminated products turned out to be essential to customers; when they could not find them, they shifted their business elsewhere, and the retailer entered bankruptcy.

A retailer of home goods used demographic data to localize its assortments to better cater to customers' tastes. It started with fashion bedding and was thrilled to see an 18% revenue lift. But when it applied the data to the fashion bath category, revenues didn't improve. Discouraged, the retailer abandoned the effort.

When the new CEO of a tire retailer shifted its assortment from low-priced tires to more-expensive ones, he learned the hard way that price mattered to his customers. The CEO was replaced after two years, and his successor restored most of the products that had been eliminated.

Fisher also adds "Like so many assortment-strategy shifts, these moves were largely acts of faith. It is easy to spot the dogs in your assortment, of course—sales data will tell you that—but it is far from obvious what slow sellers should be replaced with. And there is always the nagging concern that a slow seller you delete might be an important product to some of your best customers, prompting them to defect to competitors. As all retailers know, picking the best assortment is a balancing act; any one change can have ripple effects". Summarizing the negative impacts of assortment being done the wrong way –

Keeping myriad of SKUs & brands and overwhelming the consumer with choice (a phenomenon called brand fatigue)

Removing the dog-line products but introducing new products purely based on faith

Removing products without measuring their impact or co-dependence on other products

Not accurately measuring the opportunity associated with the instance of "consumer walking away with no purchase"

Retaining top-selling products but losing the richness in diversity of portfolio

The examples mentioned above give us a glimpse of how the problem of assortment and shelf space optimization was handled by organizations, not so long ago. Even today with the influx of analytics and data science tools, and the econometrics models to give a structure to the concepts of consumer decision process, the solutions implemented are far from being perfect primarily due to limitations of data not being available at a granular level across different channels. Some of the other reasons include - models still do not mimic possible human response to the shelf, business acumen needs deeper integration with analytical models, and customization is needed in the outcome as per the nuances of different population clusters.

To understand the kind of limitations posed due to the data, we need to gain some understanding of the different ways in which we can get access to sell-out data and any other data pertinent to store performance and consumer preferences in today's world, e.g., inventory data, promotions data, loyalty, etc.

Talking about sell-out data, it can usually be accessed in two ways:

Licensed data from a third-party market measurement enterprise like Nielsen, IRI, etc.

Retailer (or Key Account as ABI calls them) directly sharing it at a level compliant with government regulations and masking sensitive data about consumers & products

The granularity at which the data can be licensed from the third-party depends on the guidelines laid down by the retailer.

Let us try to understand the commonly available levels of data from leading industry data source providers, which have a significant coverage of retail stores across the world-

Syndicated data: This data is useful when we want to gain an insight into the entire market and not just a Key Account or demography or sales channel (online, offline or omni-channel). The real value of data at this level lies in highlighting trends of how the market is evolving.

Point of Sales (POS) data: Sales based on purchases from stores, sourced from retailers' electronic point of sales through checkout scanners. On a lot of occasions, stores do not have an electronic POS, which is where Nielsen uses Field Agents to connect with the store owners and report on sales. The way this is done is by doing a sales audit wherein the difference in the amount of stock on hand for a product in a store at two different points in time becomes the amount of product sold. Based on the guidelines set by the Key Account, Nielsen can then license the data to CPG, FMCG clients at either a store level or an aggregated level such as Region, Channel, ZIP Code, etc. or combinations thereof. In certain cases, the retailers only provide Nielsen with data from a representative sample of stores. Nielsen statistically expands this data to provide a projection of sales in all the retailer's stores.

Panel data: A panel is a sample of shoppers recruited to be representative of a universe. Shopper purchases are recorded by a variety of means such as collecting information from purchased packaging, having the shopper scan their purchases, performing optical character recognition on receipts, etc. The result of panel data capture is an inventory of purchases by store, week, demographics. Like POS data, this must be statistically expanded to represent the universe of stores.

The great strength of Panel data is that it tells the analyst who bought a product, and a lot about the buyer's profile. The drawback is that panel data has a much lower coverage as compared to POS data.

Additionally, Nielsen captures store observations such as Displays, Features, In-store promotions, and Inventory/Stock levels.

In past couple of years, government bodies across the world have also become increasingly conscious and stringent about data sharing and protection policies to maintain the privacy of individuals.

As a result of the way in which licensing and sharing of data takes place, some of the more advanced solutions for assortment and shelf optimization lie with few companies which have access to POS data with an approval from the retailers to use it; and additionally, have a holistic view of other in-store events such as displays and promotions.

AB InBev Landscape

Offline retail chains (Key Accounts) are the most important channel for AB InBev's sales. In many of these retailers, AB InBev holds the position of being the "Category Captain". A Category Captain is a Supplier (manufacturer) nominated by the Retailer and is expected to have the closest and most regular contact with the Retailer. This Supplier will also shoulder the responsibility to invest time, effort, and often financial assets into the strategic development of the category within the Retailer. The Category Captain is often the Supplier with the largest turnover in the category. Being the Category Captain also brings the responsibility of "growing the category", in this case grow beer as a category for the Retailer without any bias for own products. Successful captaincy entails rationalization of SKUs along with efficient shopper solutions that help the Retailer maximize their volume and revenue from the category.

AB InBev, till end of 2019, focused on Assortment in the following ways -

In markets with lesser maturity and complexity, where ABI played the role of Category Captain, proceedings were driven more by business knowledge. SKU rationalization was the primary focus - the lower volume SKUs were axed and the higher were retained. However, the business teams continued to explore to understand what value analytics could bring to the table.

In the more mature markets, ABI would partner with the Retailer's Category Management team and involve either the local analytics team of the Retailer or leverage ABI's Analytics COE to come up with recommendations on what to keep on the shelf. In case ABI was not the Category Captain in some of these markets, the execution of recommendations would solely depend on the Retailer's discretion as they might be working with other manufacturers and vendors providing Assortment solution as well.

For few markets, ABI would also partner with industry leaders to deliver assortment for certain key accounts. In such engagements, they would deliver the finalized set of assortment recommendations.

Gap that our Solution filled

Considering the challenges that each market posed to adequately address the needs of shelf optimization, the Global Category Management team at AB InBev realized it was important to build an in-house assortment optimization capability that could be fully customized, intake business inputs, offer total transparency into the underlying models and win greater trust from the retailers while executing the solution. Through our long engagements with different vendors, we also identified that most solutions offered in the market were more of black boxes with a very limited view of what went on inside and could not be customized for different market needs. None of the available solutions we examined could cluster the stores based on the store consumption patterns and the demographics of the surrounding areas. From our experience we realized that it's important to provide assortment recommendations at the store cluster level as a single assortment for the entire chain of stores may not do justice to many stores. Also having a separate assortment plan for all the individual stores would be impractical as it would be a logistical challenge for the retailer to implement.

Realizing this gap, we embarked on the journey of building an assortment and shelf space optimization capability at AB InBev's Growth Analytics Center (GAC), a unit of GCC – its global COE for analytics and data science programs - with an aim to create a solution that could be used across different markets and retailers.

Literature Review

The business problem of smarter assortment and shelf space planning that we are addressing through this paper can be considered as lying at the intersection of marketing interventions, operations, and economics. So, for our given scenario, we had to adopt a hybrid approach that encompassed a repertoire of techniques and analytical ingenuity to come up with a solution that could work consistently well for different markets. Below are the details of the subject areas that we delved deep into.

Random Utility-Based Discrete Choice Models

Assortment optimization is the problem of deciding which subset of products to offer to customers to maximize the retailer's revenue. Utility-based discrete choice models have been used to understand how customers select from among a group of products that vary in terms of price and quality for a long time. These models assume that every customer associates a utility with each product and chooses the option giving her the highest utility. One of the most popular and widely used discrete choice models is the Multinomial Logit Model ("MNL"). MNL model was presented in McFadden, Train, Tye (1978). The approach of using the multinomial logit model in understanding the customers purchase decision was also studied in detail in Guadagni PM, Little JDC (1983) and Wierenga B (2008). Talluri, K. and van Ryzin, G. (2004), used the MNL to optimize the assortment to maximize the revenue.

The MNL model assumes that the utilities can be decomposed into a deterministic component that represents the average utility derived by the population of customers, and a random component that represents idiosyncrasies across customers. The random component is assumed to be identical and independent Gumbel random variables with mean zero. Under these assumptions, as per ML Fisher and R Vaidyanathan (2009), the utility of each product can be derived from its market share.

Shortcoming of and Alternatives to Multinomial Logit Model (MNL)

A shortcoming of the multinomial logit model is that if a product is added to the offered assortment, then the MNL model predicts that the market share of each product in the new assortment decreases as per the products' proportion in the older assortment. This is because one of the major assumptions of the MNL is that the utilities of products are independent of each other. This phenomenon is referred to as the independence of irrelevant alternatives (IIA), and in real-world conditions might often be violated. To remedy this potential shortcoming of MNL, researchers developed other utility-maximizing models such as the Nested Logit model and Mixed Logit model.

Nested Logit Model

Under the nested logit model which was introduced by Williams (1977), customers first select a nest, and then products within the selected nest. The nested logit model was developed primarily to avoid the independence of irrelevant alternatives property suffered by the multinomial logit model. Davis, Gallego and Topaloglu (2014) discuss how to classify the complexity of the assortment problem for nested attraction models.

Assortment Optimizations at Store Cluster level

Once we calculate the utility of each of the products using MNL, the next step is to maximize an objective function such as revenue. The goal here is to recommend an assortment of products to maximize the expected revenue for the retailer. For coming up with highly contextual recommendations, we relied on first clustering the retailer stores and then proceeding ahead with optimizing the product mix within these clusters.

The optimization of products within each nest to maximize the revenue could be done for individual stores or for a cluster of similar stores. The literature by ML Fisher, R Vaidyanathan (2009) and F Bernstein, S Modaresi, D Sauré (2019) for approaching assortment optimization for the store clusters.

Optimization techniques like linear programming and genetic algorithm are used to identify the optimum assortment for each nest that would maximize the revenue.

Constrained Optimization Techniques

Linear Programming

Gallego and Topaloglu (2014) proposed a linear programming-based method to obtain the optimal solution to both assortment and price optimization problem. We also referred to Feldman, A Paul, H Topaloglu (2019) for the linear programming-based approaches for assortment optimization. In our solution approach, however, we could not use linear programming-based optimization due to difficulties in adding constraints. The unconstrained LP approach was recommending very few products which was not making a lot of business sense.

Genetic Algorithm

Genetic algorithm is one of the commonly used optimization algorithms along with linear programming in the context of assortment optimization. We referred to TL Urban (1998) for product assortment optimization making use of genetic algorithm. H Hwang, B Choi, and G Lee (2009) also developed an integrated mathematical model for the assortment optimization problem with the objective of maximizing the retailer's profit.

Stakeholders and Decision makers in the Category Management process

Ensuring a successful Category Management process required deep interaction between the AB InBev and the retailers. Since the Category Management and assortment optimization is extremely critical to the revenue of the retailer, it is important to have the senior management of the retailer onboarded to process. Below is the list of people that generally gets involved in the assortment optimization process for a retailer.

Retailer
Head of Category Management or Senior Category Manager
Store Owners or Store Managers
Supplier (AB InBev)
Global Category Management business team
Head of Category Management for the market (country where solution was to be implemented)
Key Account Manager, a person on the supplier side responsible for conversations with the Key Account
Business & Sales teams from the market - who ensure execution of assortment recommendations
Analytics Teams (GAC)
Consumer

Initial steps of Category Management Transformation at AB InBev

When we decided to embark on the journey of Category Management transformation, we wanted to start with one geography, take the learnings from there and apply it to the next and do this a couple of times till we reached a point where we had a good business understanding, knew the right questions to ask to understand a market's complexity, learnt the challenges with implementation; before we could rapidly scale up. The parameters to decide the first geography for pilot were – availability of data, relationship with retailers, market share of ABI (where we were category captains), complexity and maturity of the market.

When we started with our first market, the first step was to understand and know our shoppers.

In a three-month long Shopper Insights project, we mined massive data sets to deliver insights to the business. During the journey, we realized how consumer segments are extremely diverse in terms of their profiles, product preferences and this is when we decided that to have a good understanding of who our consumers are, we need to segment the consumers based on different cuts such as demographic, region, consumption patterns, shopping channels, etc. While we wanted to break down the consumer base, like many real-world analytics challenges, we did not have enough information about the consumers to begin with and hence instead of consumers, we went on to segment the stores so that the impact of neighborhood, occasions, etc. could be captured. We realized that not all these variables added the same kind of value in creating homogenous groups of consumers. After some back and forth, we ended up with 4 segments driven by income and lifestyle index. The number of clusters was decided based on doing a trade-off between the increasing homogeneity of clusters and reducing the complexity of the clusters.

When we profiled the clusters and juxtaposed them with information about sales and product preferences, we realized the actual value the clusters added was a lot more than we anticipated. This exercise helped us to have clustering as a pre-requisite to understand how consumption patterns varied across demographics, regions, stores of different channels and sizes and whether segmenting the store base made sense and added value; and only then we would move on to the next set of tasks such as insights generation, assortment optimization and so on.

How we laid the building blocks for our Assortment and Shelf Optimization

As the first step, we liaised with multiple internal teams from AB InBev that had worked on the problem of assortment optimization or on a similar technical problem in the past. Through this exercise, we brought under one roof, all the learnings, best practices, and existing frameworks to form a knowledge repository. Using this as a foundation, the team researched white papers and journals to understand how we could leverage and mold the existing or proposed solutions for our own use case and bring forth an industry grade solution into action.

While researching on the technical aspects, we were also working closely with the business team and the National Category Manager from our pilot market to gain business context and understand the different aspects and challenges associated with implementation such as -

What role does beer as a category play in the retailer's entire portfolio of products?

How often does re-ranging (change in assortment) occur? How much of a role does seasonality play in this?

What is the current granularity at which assortment optimization is done? That is, is there one assortment for the retailer or is it broken down by clusters, regions or in some cases stores? Where do they want to move from here?

What are the cost and effort implications of making assortment recommendations at a highly granular level? Who creates the planograms – is it a vendor or someone from AB InBev's or retailer's team? Which tool is licensed for this? What is the level of manual intervention needed?

What are the different priorities retailers have? How do these priorities vary depending on the type of retail outlet, i.e., in a grocery vs convenience stores vs liquor stores vs mass merchandiser?

What is the expected outcome from an assortment and shelf optimization exercise? Is it to focus on premiumization, or is it to reduce stock-outs and increase sell-through, etc.?

What is defined as success in this exercise? What are the different KPIs that matter the most for different parties?

How do different markets and retailers track adherence of planograms? What are the complexities associated with tracking adherence and measurement of recommendations?

AB InBev's 4C Framework

Through our work with multiple countries and retailers from different geographies, we have developed what we call as the 4C Framework. The process was built after a lot of back and forth, inputs and suggestions from the business teams, collaboration with other teams that had previously approached or worked on assortment and partnering with MIT. Along the entire journey, the emphasis always was to make the process robust but never complicated. We are immensely proud of the fact that we have simplified the process and made it transparent to an extent where all teams (as mentioned in section 4) come together as owners and contribute to make the solution better every day. The 4C framework is a process that helps us classify the SKUs or Parent SKUs (a level less granular than SKU) in a retailer's portfolio into 4 buckets:

Conserve: This is basically the list of SKUs that should not be discontinued or de-listed at any cost. They are the most significant SKUs and majority volume drivers for the store. The list also includes products which have a relatively lower volume but a sizeable incremental volume. We keep these SKUs under the 'Green List', discussed in detail shortly.

Consult: These are the SKUs that cater to a niche, e.g., craft beer offered in a can, chocolate flavored beer. They do not deliver a lot of revenue or volume but ensure a diverse range of products for the retailer. We keep these SKUs under the 'Amber List'.

Cultivate: These are the SKUs that where we see a potential and want to place our bets on, i.e., we want to increase their distribution (presence) in more stores or increase their shelf space in stores. This is also the list where we add the new SKUs that we want to introduce into the stores.

Cull: These are the SKUs that have a low incremental volume, and in their absence their demand can be substituted by other products. We keep these SKUs under the 'Red List'.

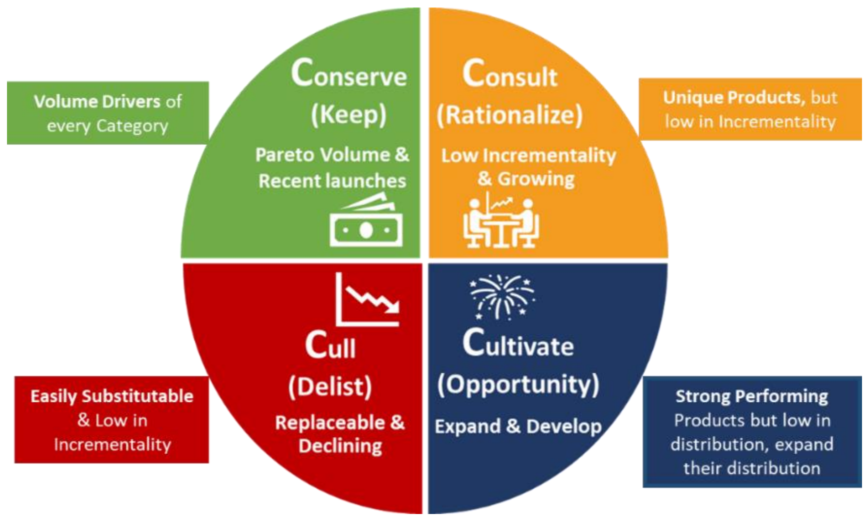

Figure 1: The 4 C Framework

The success of the 4C framework lies in the deep collaboration between business teams (who represent the exact needs and priorities of the retailer) and GAC team (who are responsible for development of the custom-made solution for the retailer.

Overall Process Flow for SKU Optimization

Below (Figure 2) is a high-level overview of the different steps that we follow while performing SKU optimization for retailers. Each of these areas has been in detail in the subsequent sections.

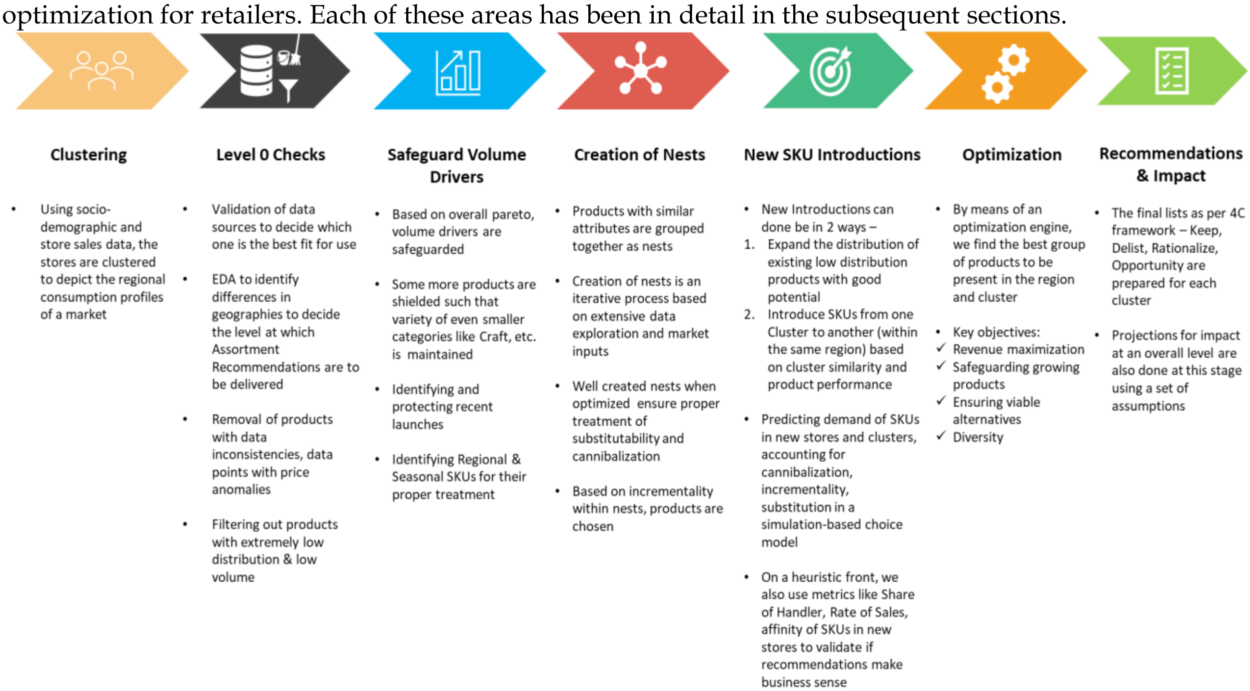

Figure 2: Figure 2: The process flow

Detailed Steps

Let us look at how we approach a new market when we start our Category Management journey. No matter the region, country, type of market, retailer we work with, the first leg of the project always starts with gathering information about the retailer and what they want to achieve from this range review - do they want to focus more on premiumization, do they want to optimize inventory levels and reduce stockouts, do they want to increase profits, do they want to focus on new introductions, etc. This stage of the project makes sure everyone aligns on the objective and shapes the days to come in terms of planning, milestones, etc. Then we get into the nuances of deciding, what should be the level of execution and how frequently a range-reset should take place.

Once we get a certain level of clarity regarding the questions listed above, only then we move on to the analytical steps of the process.

Here is how we go about our process:

Clustering

Our assortment optimization solution is designed to provide optimized SKU assortment for individual stores of the retailer. But since it can be practically difficult to design and roll out separate assortments for individual stores, many retailers that we worked with opted for assortment solutions at a higher granular level - at a Store Group level. For the optimization to be effective, these groups of stores must be similar in terms of their beer consumption and demographics that they serve. We made use of clustering algorithms to come up with these Store Groups.

Data Sources for Clustering

To effectively group the stores into meaningful clusters, we used data related to the stores and the data about the area where the stores were located. Some of the major data sources that we used in the clustering of stores are:

Demographic data- Demographic information of the location where the stores are located, like average age of the residents, average income, car ownership, language spoken, gender ratio, etc. In many geographies, we used census data available at Postal Code/Postal District level for clustering of stores.

Consumption data – Consumption information is one of the major data sources used for clustering stores. Details on the proportion of volume sold by pack size (6 packs, 12 packs, 18 packs, etc.), pack type (bottle, can, keg, etc.), container volume (500 ML, 2L, etc.) help in grouping the stores into meaningful clusters.

Point of Interest data – Points of interest like the number of restaurants, bars, beaches, etc. in the vicinity of each of the retailer outlets can be a useful data point in clustering. We generally retrieve the point of interest data from sources like OpenStreetMap (OSM).

Clustering Methods

To come up with the group of stores which are similar in characteristics, we experimented with both un-supervised clustering techniques like K Means and Hierarchical clustering and semi-supervised clustering techniques like metric learning. K Means clustering turned out to be the most effective algorithm in our case, especially since the business also had an input on the desired number of clusters for their markets. Explaining the clustering techniques in detail is beyond the scope of this paper.

Layers of Clustering

Our experience, working in different markets, suggests that the consumption behaviour of consumers varies drastically across the expanse of any country. For example: in some of the Western European countries, there are regions with a strong preference towards craft beer as compared to a different part of the country which has a strong liking toward stout. To capture these variations existing in the same market, a top-down clustering approach would not have sufficed. This prompted us to perform clustering at multiple levels. While the Demographic Data would help us with market level understanding of the population; Consumption Data from the retailer outlets would tell us which exact brands or styles or price segments were preferred by consumers in different regions of the country. These two levels of consideration would help us build a robust clustering model that would accurately capture the nuances of different regions in any market, without having a uniform model for the entire country.

While markets such as South Africa had distinct demographic patterns and subsequent consumption patterns in different regions which allowed us to create a decently uniform segmentation; in other countries such as the UK, consumption varied drastically across the country. The practical challenge of having two-tier clustering was that it resulted in many clusters specially in countries with huge variations in consumption pattern (such as the UK). This was where we had to sit with the business teams and try to aggregate some of these micro-clusters under a larger umbrella so that the subsequent steps of Category Management analytics such as assortment recommendations, etc. could be carried out in a less complex fashion.

Level Zero Checks

After we had aligned on the clusters with the business, and before subjecting the SKUs to go through the SKU optimization & 4C funnels, it was important to make sure that the data was clean and usable. Below are some of the checks that we do to ensure this:

Erroneous data points are removed or dealt with prior to the analysis

Products that are niche or recently launched or seasonal are retained for 'further consultations with the business' rather than leaving them to the risk of getting removed (as per the 4C framework mentioned above)

Some of the basic checks that we perform are listed below. Many of these are optional and need to be finalized after discussing with the business:

Recency Check – The SKUs which were launched in the last 2- 6 months prior to the assortment exercise could take some time for their sales to pick up and we believe it is not right to compare their sales with the rest of the well-established SKUs. So, we retained these SKUs in the optimized assortment unless asked by the business. Doing this ensures that the recently launched SKUs get more runway to prove their potential.

Price Anomaly Check – Within a category of products, in our case beers, we generally assume that the prices would be in a stipulated range and any product whose price falls outside this range could be flagged as an anomaly. The authenticity of prices of these products would be checked and corrected if necessary.

Seasonality check: For most of the retailers, there could be a considerable proportion of SKUs which would sell predominantly during a particular season. We do not include these SKUs in the assortment optimization as the annual sales revenue or volume of these SKUs could be deceivingly low and could get removed from the assortment. We flag products with exceptionally high sales during some seasons and add these SKUs to the final assortment only during the months when they are in demand.

Absolute Volume Check – We do not include SKUs which sell in extremely low volumes in the optimization step unless specifically asked by the business. This is done under the assumption that there could be erroneous barcodes that might have led to wrong purchase entry, or these SKUs were already in the process of delisting.

Discontinuity check: Discontinued SKUs are identified by checking the quarter-over-quarter change in revenue. SKUs showing decline above a cut-off are assumed discontinued and hence de-listed in our process. We do this under the assumption that these SKUs are being sold to sell off the existing stock and would not be resupplied. This list gets vetted by the business.

Identifying and Safeguarding Top Performing SKUs

While doing the SKU rationalization, we do not want to include the SKUs which are major volume drivers in the optimization process. The reason for this is, historically when we ran optimization iterations with the top-performing SKUs, a few of them got delisted. In such a scenario where a top-selling SKU gets de-listed, it would take a while before the lower selling SKUs would organically capture the lost volume. To avoid this, we keep the top performing SKUs out of the optimization process. The fact that all the major volume drivers are retained in the new assortment also helps in building the confidence of the retailers in the solution delivered.

Generally, in each category, there are around 20% SKUs which contribute to 80% of the sales volume. So, we protect those ~20% SKUs from optimization exercise and we call them Pareto SKUs.

Segmenting Products into Consideration Sets (Nests)

Nests are logical groupings of products with similar attributes – like beer style, price segment, pack size, etc. The underlying assumption of our approach is that a consumer would consider the products in a nest as alternatives while making a purchase. Nests are created after the above-mentioned Level Zero Checks and safeguarding the top performing SKUs. We run our SKU optimization codes on these Nests. This ensures that the optimization funnel is applied on comparable products.

To create nests, we group products that are similar in their attributes. Each nest carries products that belong to the same beer categories (lager, craft or domestic, etc.), same pack type (can, bottle, etc.), same price range (say $5-$10), and the same range of Rate of Sale (ROS) (say 20-30L per week). We perform binning using attributes such as price or ROS, and then different bins become part of various nests. The entire nest creation process goes through discussions with business teams to agree on which parameters should be fixed and where we can have some flexibility.

Below are few illustrative nests. As can be seen, the objective is to ensure full coverage of all attributes so that the end consumers' choices do not get limited.

Figure 3: Pictorial representation of a Nest

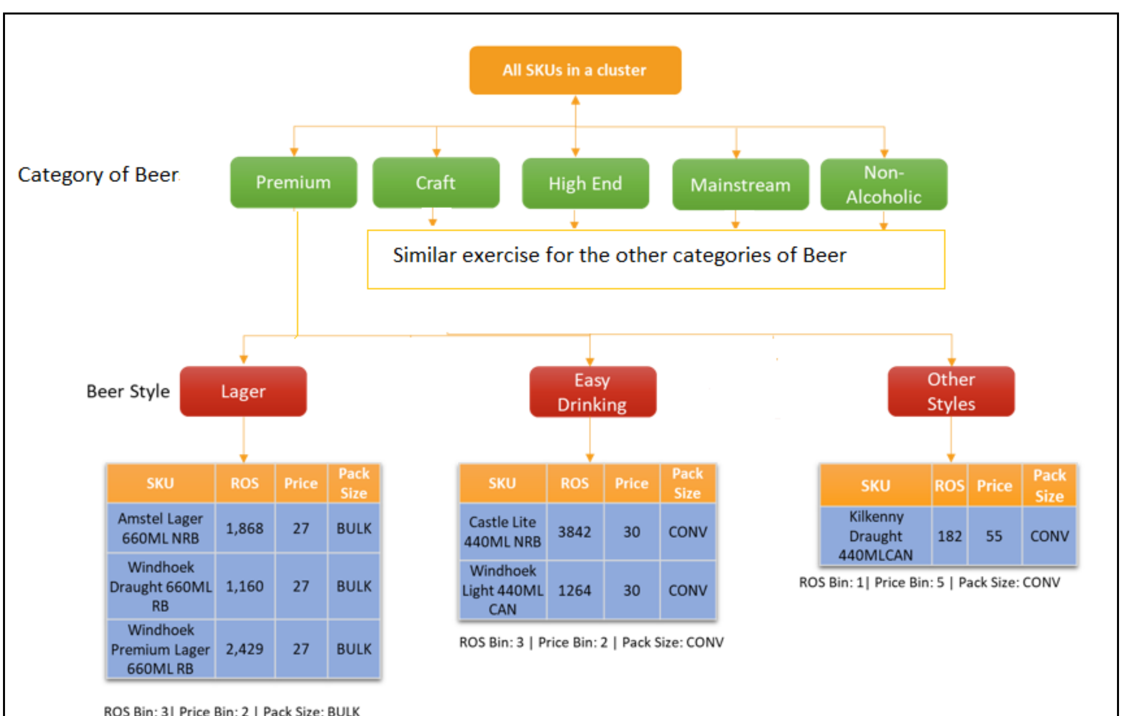

Figure 4: Example of a nest creation

Finding the Optimal Assortment by Maximizing Utility
Utility

The word Utility comes from Economics, and it means "usefulness" or "value", the value a product or service holds for the end consumer. Based on this we can see that the perceived value of a product could be different for different types of consumers. It is important to note that this utility of a product is determined by the attributes of the product and whether those attributes are relevant to the consumer or not (Fishburn, Peter C. (1970), McFadden, Train, Tye (1978).).

Random Utility Theory

Our solution for finding the optimum assortment comes from the source of Random Utility Models and these models aim at modeling the choices of individuals among discrete sets of alternatives. It is assumed in these models that the preferences of an individual among the available alternatives can be described by a utility function. The individual chooses the alternative with the highest utility (McFadden, Train, Tye (1978)).

According to Random Utility Theory, the utility (*U*) for an individual *i* making a choice *j* is a function of one or more observed features of the choice (*Vij*), and an error term representing unobserved attributes (*εij*)

$$\mathbf{U_{ij} = V_{ij} + \varepsilon_{ij}}$$

Here *Vij* is the Systematic utility which is a function of observable variables and *εij* is the Random utility.

Utility is being looked at as a product's importance for an individual; but for our use case where most of the times the data is not available at a consumer level (i.e., transaction-level data or loyalty data) and we do not know the attributes of the end consumer and the choices they make, we must settle with syndicated data at store-product-week/month level. We have come up with a way to work around this shortcoming. Since we cluster stores based on demographics, we go with the assumption that most of the consumers walking into a store in this demographic region would have similar attributes (ex: income, age, occupation, etc.). Based on this approach, we try to arrive at the list of products that would optimize the value/utility delivered by the entire store at an aggregate level.

Option of No Purchase and Purchasing Outside the Category

Our choice model is based on the Random Utility theory which says that people generally choose what they prefer, and where they do not, this can be explained by random factors. There is criticism associated with random utility theory which says that people are not always rational and most of the decisions are impulsive and made on the shelf. The way we look at this is that while the stated preference might differ from the final action of the consumer on the shelf, nevertheless sales truly reflect the revealed preference.

We have leveraged Multinomial Logit (MNL) which comes from the class of Generalized Extreme Value models. For example, a person may choose her preferred ice cream 9 out of 10 times and on the 10th occasion she chooses something else due to some random factor. To model the decision-making process mathematically, we need to make sure that the choices are exhaustive in nature. For example, if a store offers 10 products such that their selection is mutually exclusive, the consumer will have a probability associated with the purchase of each of these products and the sum of probabilities must add up to 1. Further, we would need to have a sample outcome set that shows all the possible choices, one of which being the option to not purchase anything. The utility associated with this no-purchase option governs the volume that would be lost because of changes made such as de-listing or a price change.

In our case, we consider the no purchase utility as the least useful choice for a consumer, because the assumption is that when a consumer walks into the store, she walks with the intent of purchasing something and only if she does not find any value in purchasing the offered items, she walks away with a no-purchase.

Thinking of this mathematically, we use sales as a proxy to identify the utility of a product. We call this a proxy because actual utility would come from a conjoint survey exercise where people provide their stated preferences of what they would choose when provided with n alternatives. Any product that got sold in a store, has a volume greater than zero, and hence will have some utility associated to it. Another assumption here is that no-purchase option/outcome is the least popular outcome in the universe of possibilities, i.e., the probabilities of products getting bought are a function of their utilities, higher the utility, higher the probability of a consumer buying it. Keeping this in mind, we assign the no-purchase option a nominal utility value which is lower than the least popular choice in the nest, i.e., the product with the lowest utility.

Specifically speaking, the outside options are assigned a utility which is lesser than the utility of the smallest selling SKU in the nest, which then makes the no purchase scenario the least taken option or least popular choice.

Outside option is pretty similar to no-purchase situation except that here the consumer, in the event of not finding a product on the shelf, shifts to a different category (ex: shifting away from Lager beer to Flavored Alcoholic Beverage or Soda).

Optimizing the Assortment using Genetic Algorithm

As mentioned before, we run the optimization within each of the nests to maximize the value of the objective function.

Objective Function

The objective function calculates the revenue from the exponential of utility and the market share of the product and the original volume. Please see the formula below for the objective function

$$\text{Revenue} = \sum_{i=1}^{n} \frac{p_i * \exp v_i}{\sum_{i=1}^{n} \exp(v_i) + \exp v_{no\ purchase}} * original\ volume$$

Here p_i denotes the market share of the product,

$\exp v_i$ denotes the exponential of utility of the product,

$\sum_{i=1}^{n} \exp(v_i)$ denotes the sum of the exponential of utilities of all the products in the nest, and

$\exp\ v_{no\ purchase}$ denotes the exponential of utility of the outside option

The objective is to arrive at the list of SKUs within each nest that would maximize the retailer's revenue at the cluster level. To identify the top-performing SKUs which have a meaningful contribution to the cluster revenue, we experimented with different optimization techniques like differential evolution optimization (using the DEoptim package), linear programming (using lpsolve) and genetic algorithm (using genoud package). We did not get the best results with either DEoptim (as it was taking extremely long to converge and also had output in decimal values between 0 and 1 making it difficult to identify the cut off values for recommending the SKUs) or lpsolve (this linear programming based algorithm was recommending very few products per nest in an unconstrained approach and it was practically very difficult to add constraints to make sure more products which would maximize the revenue got recommended). We therefore decided to use genetic algorithm using the genoud package for identifying the optimized assortment. The convergence time for the genetic algorithm was much lower than DEoptim and the algorithm also provided the option of receiving the outputs as integers (1 or 0 which implies whether we are recommending a product or not).

Genetic Algorithm

Genetic algorithm is a search-based optimization algorithm based on Darwin's theory of natural evolution (TL Urban (1998)). It works on the basic principle of Darwin's theory that the fittest would survive, and they would be selected for reproducing the next generation. It works on a random selection process and can solve complex and non-linear problems. Details of genetic algorithms are beyond the scope of this paper.

We run genetic algorithm for each cluster at the level of different constituent nests. Here we maximize the revenue and from the algorithm, we identify the best combination of SKUs that get selected and directly go to "Green" list, while the remaining SKUs go to "Amber" list for further investigation.

Identifying High Growth SKUs and SKUs with Close Substitutes

After completing the optimization, the SKUs recommended to be delisted would further go through a few more rounds of checks to make sure that the SKUs without close substitutes and SKUs with high growth are preserved.

Identifying High Growth SKUs

There could be SKUs which may not be selling in great volumes currently, but their market share shows a consistent quarter on quarter increase. In this step, we protect those SKUs with a QoQ growth above a certain threshold. SKUs above a particular cut off are marked as growing SKUs and they would be removed from the list of SKUs marked for delisting. Further, we consult with the business team on how the retailer would like to deal with these high growth SKUs. The threshold growth rate for determining the 'growing SKUs' depend on the overall growth of the SKUs across segments and generally falls between 5% to 10%.

Checking for Close Substitutes

Before delisting an SKU from the shelf, we need to make sure that there are substitutable products available so that a customer walking in for a particular type of product would not be disappointed. This is necessary because there could be customers with niche preferences looking for a particular category of product and even though the sales from that category may be less, delisting the category completely could have negative halo effect on the sales of other brands which are often picked within this category.

Checking for substitutes is done both within and outside of nests. This is because SKUs from different bins could be alternatives and they could have got separated into different nests due to marginal differences in attributes such as price or Rate of Sale.

Identifying Potential Super Performers with Low Distribution – for 'Cultivate' list

The Cultivate List in the 4C Framework focuses and brings forth products that have the potential to grow and generate more revenue for the retailer but got overshadowed and neglected in our analysis due to their low utility which in turn is dependent on their volume. This could be due to reasons such as low numeric distribution (not being present in enough number of stores) or low weighted distribution (not being present in the right set of stores, i.e., the product not being present in stores where that particular beer type is in demand), etc. We try to look at these products more deeply to eliminate the inherent disadvantages associated with them.

We do this by identifying SKUs which have performed well with respect to their distribution by making use of a metric ***Over performance index*** which is calculated by dividing the volume percentile of a SKU by the distribution percentile.

Over performance index = **Volume percentile/ Distribution percentile**

An ***over performance index of greater than one*** indicates that the SKU is performing better than the similarly distributed SKUs and may have a potential to do even better if the distribution is improved. These SKUs could be considered as candidates for cultivate.

Demand Model – Introducing SKUs Across Clusters

There could be products that are performing well in some of the clusters or regions and can potentially add value if introduced to other clusters as well. To identify SKUs that could do well when introduced to a new cluster, we have developed a machine learning model which uses algorithms like Random Forest and XGBoost to predict the volumes that an SKU can generate when added to a new cluster.

To predict the volumes of potential SKUs, the model uses three sets of variables

SKU related variables like product description, brand, pack size, etc.

Store related variables like Regional Sales Division, Strategic Business Unit, number of bars in the area, number of hotels in the area, store parking space, etc.

Sales related variables like net revenue, base sales, price, number of orders for the SKU in the week, etc.

The predicted volume of these new SKUs would be compared with the volumes of existing SKUs in each segment and the SKUs that do well compared to the existing SKUs can be considered for recommendation. Generally, for introduction, we consider those SKUs whose predicted volumes exceed the top ten percentile of each segment.

Financial Impact Calculation

Once we complete the assortment optimization process and make the SKU recommendations, it is important to understand the financial implications of the new assortment. We calculate the financial impact at the store level, and they are later rolled up to the retailer level to understand the overall impact. To calculate financial impact, we assume that:

The space emptied by the ‘delisted SKUs’ is distributed among the SKUs that we are retaining (Keep) and to the new SKU introductions, in proportions of their expected Rate of Sale (ROS)

Share of the shelf cleared due to SKUs getting removed (Delist SKUs) is equivalent to the volume share of those SKUs

The ROS of SKUs which are being newly introduced in a store is approximated from the ‘similar’ stores in the neighborhood. The definition of ‘similarity’ can vary across retailers and needs to be finalized after discussions with the retailer and business teams. We generally consider stores with same store format (Supermarkets, Convenience stores, etc.) Cluster, Region, etc. as similar. A store can be considered to be in the neighborhood’ of another store if the geographically distance between them is within a threshold. This threshold could be 0-3 kilometers or more and to be decided after understanding the geographical realities of the market. Once the ‘similar’ stores in the neighborhood are identified, the ROS for the new SKUs are calculated as the mean ROS, and the number of weeks the SKU was sold is calculated as the median of the number of weeks it was sold in the similar stores in the neighborhood. Once we have the ROS and the number of weeks the SKU was sold, the revenue from the product is calculated by multiplying the ROS, price, with the number of weeks the SKU was sold.

Manufacturer	Recommendation	SKU	Original ROS (Volume/ week in L)	Price	Total Revenue (Volume*Price* Weeks)	MS (Pre-Assortment)	Weights (based on original ROS)	Distribution in New Portfolio	New Shares (Post Assortment)	New Volumes	New Revenue (Volume * Price)	Total New Revenue (Revenue*Weeks)	Overall Impact
ABI	Keep	Budweiser	500	4	130,000	42%	0.4	10%	52%	624	2,496	129,792	3%
HK	Keep	Amstel	400	5	64,000	33%	0.32	8%	41%	492	2,460	98,400	
HK	Delist	Kronenbourg	300	3	42,000	25%	-	-	-	-	-	-	
ABI	Introduce	Budlight	350	5			0.28	7%	7%	584	420	14,700	
								Total 25%					

For SKUs which are getting introduced in the store, their ROS & Price get approximated from the nearby “Similar” stores in the same Cluster Region & belonging to the same store format.

We distribute the delisted SKU’s share into the existing and new introductions based on their weights.

Table 1: Illustration the Revenue Impact of De-listing & Introducing SKUs

Collating Recommendations

Classifying Products under the 4 Cs

After performing all the above-mentioned steps, we finally classify the SKUs under the 4C Lists (Conserve, Consult, Cultivate and Cull). These lists help us in highlighting the reason due to which a certain recommendation was made for a SKU.

Conserve list captures SKUs that are volume drivers i.e., Pareto SKUs or SKUs that have a high incremental volume.

Consult list captures the SKUs which are either not easily substitutable or grew in market share in the last quarter. Even though these SKUs do not add a high incremental volume to the category, they nevertheless cater to a niche and hence add to the diversity of the product portfolio. **Cull list** highlights the SKUs that are low on incremental volume, do not have growth in market share and can be easily substituted by the other products retained in the portfolio of SKUs. **Cultivate list** includes SKUs that have shown promising sales in the limited avenues they were present in. The reach of these SKUs can hence be expanded to a larger number of stores.

Here's a flowchart which shows the process in a nutshell –

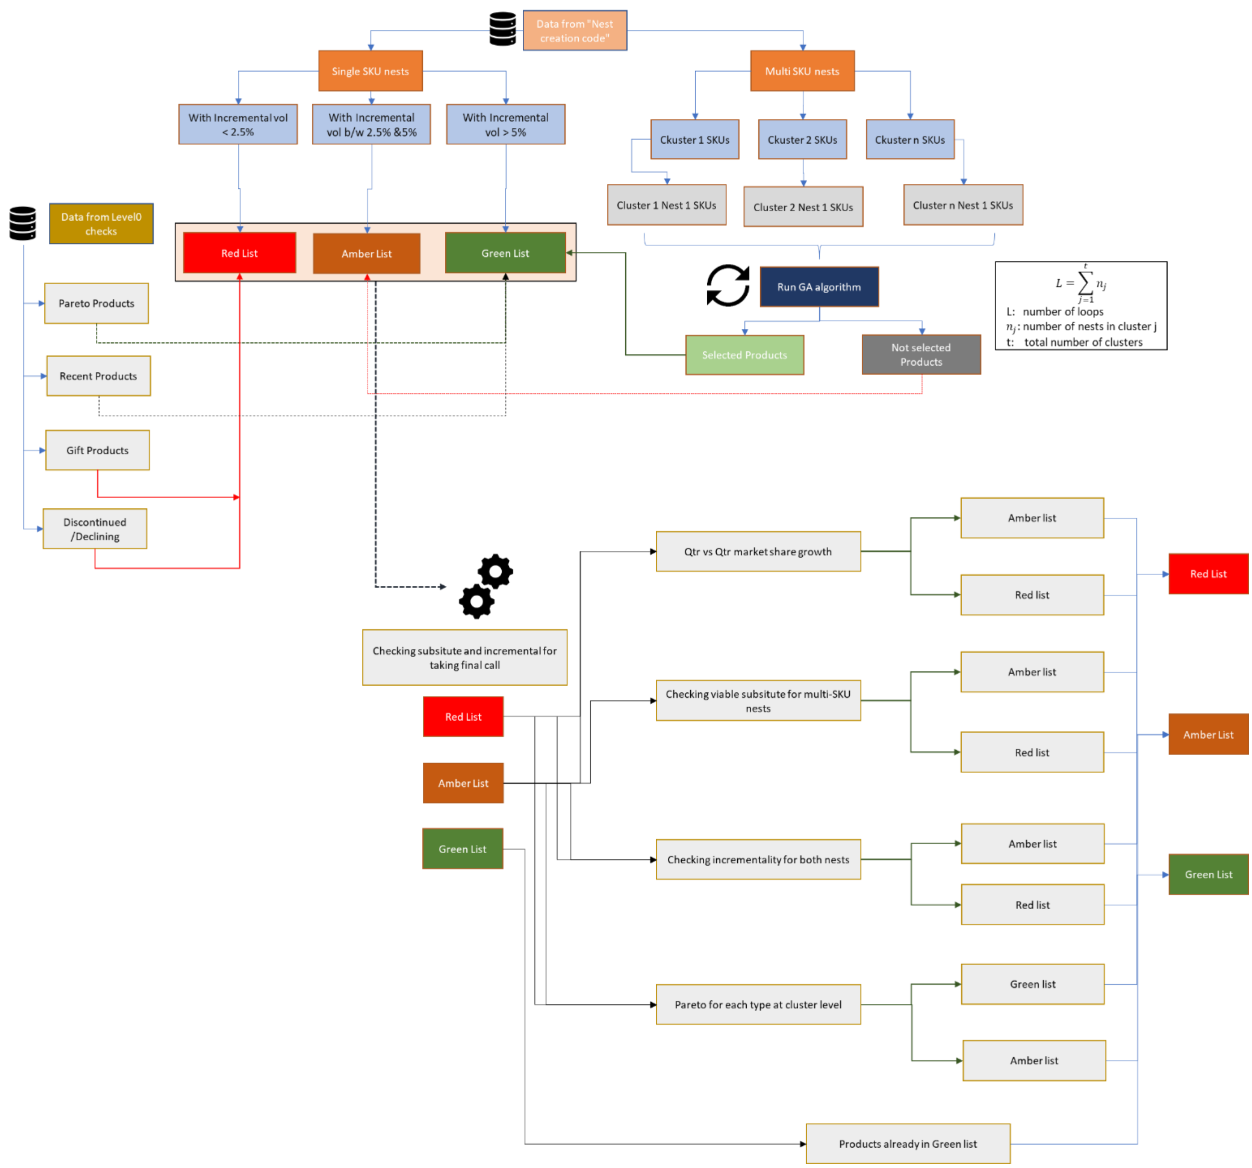

Figure 5: The Process Flowchart

Recommending Quantities of Products

Along with recommending SKUs for a store or a cluster, we also keep in consideration their quantities because that is how we would ultimately convert these recommendations into shelf space and put them out on planograms.

While we perform the SKU optimization exercise using genoud, we run simulations to find the predicted volume for each SKU and the volume that would be lost to another store due to a no-purchase situation. Using these predicted volumes, we can estimate the recommended shelf space for products.

But often volume is not the only parameter to consider when assigning space to a product. Hence, we create a rank of SKUs based on a composite score of predicted volume, ROS, and recency of a product.

Once we finalize on the shelf space to be allocated to each product, we then use this as an input to create planograms using any planograming tool.

The final leg of implementation requires deep collaboration with the retailer and the business teams to ensure the strategic priorities for the retailer are fully taken care of.

Measurement of results

The last step of the process is to measure the effectiveness of the recommendations. For this we need to measure the lift in sales which can be statistically attributed to the change of assortment and planogram.

Below we are elaborating the results from a retailer from Mexico that we will be referring to as Retailer X

Identifying the test and control stores

For measuring the lift in implemented stores, we follow a test-control approach. In a design of experiment, a test group is a sample or a group that receives an experimental procedure or a treatment. A control group is a group separated from the rest of the experiment such they are shielded from the independent variable that is being tested (new assortment).

The implemented stores in our case become the test stores and the new assortment is the treatment. To find a set of control stores, we performed few high-level checks:

Control stores should correspond to each test store and should be from the same cluster

Control stores should be from a similar area such that the socio-demographic and consumption profile of the customers served is consistent. This is done by taking the zip of a test store, then using cosine similarity on zip level demographic data to find the top 2 most similar zip codes to get our control stores. We identify our control stores use the following store characteristics to identify the similarity-

Average sales of the stores should be similar (within a defined variance limit)

Test and control stores should have similar SKU lists (within a defined variance limit)

For this pilot Retailer X, we identified 14 Test and 17 Control stores from the same cluster.

Checks on data consistency and selection bias

Once matchable test-control pairs of stores are arrived at, we first check the robustness of test and control groups to make sure there is no selection bias present in the selection of control stores. We generally use Propensity Score Matching (using K Nearest Neighbor algorithm) to make sure the test and control stores were similar before the experiment. The results we have from the Retailer X is shown below.

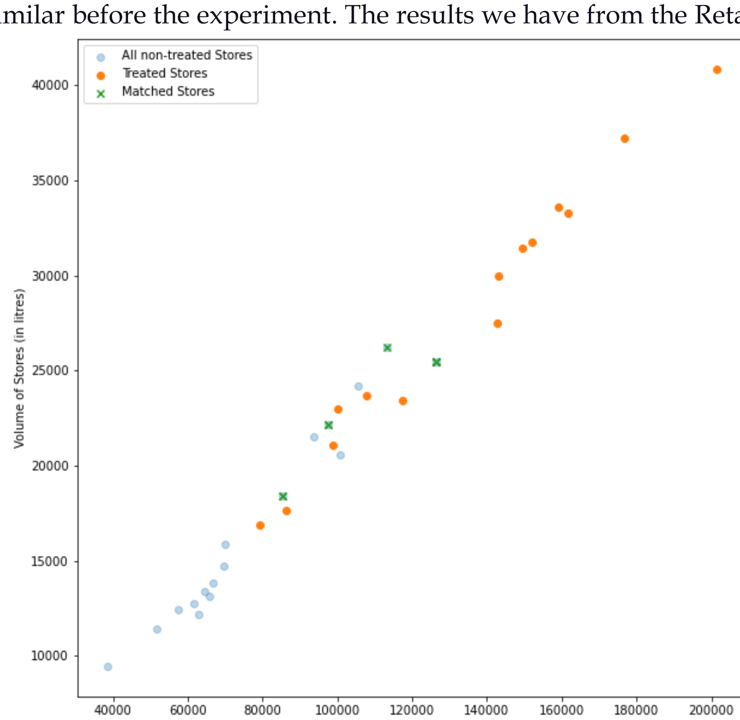

Figure 6: Propensity Score Matching for the stores

Here the treated stores are the ones where our assortment optimization solution got implemented and the matched stores are the ones identified as control stores. From the graph it’s identifiable that they are similar.

Check for statistically significant difference in revenue between Test and Control stores

Once we have the test and control sets of stores finalized, we conduct hypothesis testing to identify if there is any statistically significant difference between the average sales of Test and Control stores in the treatment period. We use Difference in Difference (DiD) method to check and validate if the impact of treatment exists. In the below chart we have the results of the DiD estimate between the test and control stores of the pilot Retailer. And a positive sign of the time*treatment variable(did), indicates that treatment has a positive impact on the stores.

Coefficients:					
	Estimate	Std Error	t value	Pr(>|t|)	Significance level
(Intercept)	8022	432.4	18.552	< 2e-16	***
treated	4521.4	643.4	7.027	6.23E-12	***
time	677.7	611.5	1.108	0.2682	
did	1662.5	910	1.827	0.0682	.

Table 2: Results of Difference in Difference Estimation

Further details of the measurement strategy are beyond the scope of this paper.

Financial Impact Assessment

The financial impact attributed to the impact of assortment optimization is calculated by comparing the revenues of test stores vs. the control stores as below.

	Test	Control
No. of stores	14	17
Average Sales Revenue 2021	133,953	78,298
Average Sales Revenue 2020	112,891	72,198
Aggregated Sales Revenue 2021	1,875,339	1,331,060
Aggregated Sales Revenue 2020	1,580,468	1,227,372
% Revenue Lift in Test (14) Vs Control (17)*		10.2%
% Revenue Lift=Change in revenue % in test data-Change in revenue % in control data		

Table 3: Comparison of test and control store revenues

As shown in Table 3, there was a revenue lift of 10% in the stores where the assortment optimization was applied compared to the store where it was not implemented.

Conclusion

Ensuring efficient Category Management is important to any Retailer. A Retailer cannot carry the same assortment forever and needs to refresh it based on the market trends and changing customer tastes, to maximize their revenue and customer satisfaction.

As discussed in detail in our paper, we referred to the leading approaches being adopted in the industry to address the challenge of assortment optimization. We spent a great amount of time researching the best-in-class academic approaches to solve this very important and long existing industry challenge. Not only did we incorporate the most suitable econometric models together with cutting edge genetic algorithms, but we also gave an equal amount of importance to heuristics. This hybrid approach enabled us to come up with a robust mathematical approach which is overlaid by a strong layer of business acumen as well.

We believe one of the most significant enhancements that our approach brings forward is the flexibility of its framework which very well captures the nuances of different markets. It also provides a great deal of control to business practitioners who can tweak the knobs for certain group of products which they feel are important to be continued in the recommended assortment. Although it may be difficult to take care of these borderline products which would otherwise not find a place in any standard assortment optimization exercise, our approach provides just enough space to accommodate these products as well so that at the end of an assortment refresh exercise, the final consumer is able to find what she was interested in.

For an organization like AB InBev which has markets across the globe and caters to a wide range of consumer tastes and preferences, it is extremely important to empower its Retailer partners with a highly customised approach to serve the end consumers in a holistic way.

We have formulated a process for finding optimal assortments, comprised of an optimization model and heuristics-based interventions for choosing the right assortment.

Limitations and directions for future research

Although the study results create a comprehensive framework for finding the right assortment, there are some limitations that should be mentioned and noted.

Impact of promotional pricing - The purchase decision of a consumer can get heavily impacted by the promotions being run by the retail at various points in time. However, in this study, we have assumed that there are no promotions at any point in time as the promotional data was not available.

Use of Transactional and loyalty card data – Transactional data combined with the demographic information of the consumer can provide useful insights about shopping behavior. But this study was limited to the store level sales data and demographics data.

To further enhance this approach, a similar study should be carried out with more data sources like transactional data, loyalty card data, and data on promotional pricing.

References

McFadden, Train, Tye (1978). An Application of Diagnostic Tests for the Independence from Irrelevant Alternatives Property of the Multinomial Logit Model, Transportation Research Record Journal of the Transportation Research Board

Guadagni PM, Little JDC (1983). A logit model of brand choice calibrated on scanner data. Marketing Science

Wierenga B (2008). Handbook of Marketing Decision Models, International Series in Operations Research and Management Science, Vol. 121 (Springer, New York)

Talluri, K. and van Ryzin, G. (2004). 'Revenue management under a general discrete choice model of consumer behavior', Management Science

ML Fisher, R Vaidyanathan (2009). An algorithm and demand estimation procedure for retail assortment optimization

JM Davis, G Gallego, H Topaloglu (2014). Assortment optimization under variants of the nested logit model- Operations Research

TL Urban (1998). An inventory-theoretic approach to product assortment and shelf-space allocation - Journal of Retailing

H Hwang, B Choi, G Lee (2009). A genetic algorithm approach to an integrated problem of shelf space design and item allocation   Computers & Industrial Engineering

R Chen, H Jiang (2017). Capacitated assortment and price optimization for customers with disjoint consideration sets - Operations Research Letters

F Bernstein, S Modaresi, D Sauré (2019). A dynamic clustering approach to data-driven assortment personalization- Management Science

Feldman, A Paul, H Topaloglu (2019). Assortment optimization with small consideration sets - Operations Research

ML Fisher, R Vaidyanathan (2012). Which Products Should You Stock? Harvard Business Review, November 2012

Fishburn, Peter C. (1970). Utility Theory for Decision Making. Huntington, NY: Robert E. Krieger