

Tapping Into Growth: Strategic Distribution Optimization for a Veteran-Owned Brewery

Yash Patel, Francesca Schaefer, Swayam Singh, Cameron Solomon, Faculty Advisor: Brian Lemay

Operations Research Program, United States Air Force Academy, Colorado Springs, Colorado 80841

Corresponding author's Email: blemay@umich.edu

Author Note: Yash Patel, Francesca Schaefer, Swayam Singh, and Cameron Solomon are cadets in the Operations Research program at the United States Air Force Academy, Colorado Springs, Colorado 80841. Correspondence regarding this paper should be directed to Brian Lemay at blemay@umich.edu. The views expressed herein are those of the authors and do not reflect the position of the United States Air Force Academy, the Department of the Air Force, or the Department of Defense.

Abstract: Red Leg Brewing (RLB), a veteran-owned brewery, seeks to optimize its distribution network by improving delivery efficiency, allowing more time to support high-potential customers. Using 12 months of sales and visitation data, we employed clustering analysis to classify distribution customers based on current performance and growth potential. Analysis reveals that 30% of accounts show significant untapped growth potential, while 25% of current sales visits target accounts with limited opportunities. To capitalize on these insights, we apply the Traveling Salesman Problem (TSP) to streamline delivery routes, reducing travel inefficiencies and reallocating time toward high-value accounts. By optimizing distribution logistics, RLB can increase regional sales and enable more focused engagement with key customers. This data-driven approach provides RLB with actionable recommendations to enhance sales efficiency and achieve strategic growth targets in the competitive craft beer market.

Keywords: Sales Optimization, Route Efficiency, Red Leg Brewing, Operations Research

1. Introduction

Red Leg Brewing (RLB), a veteran-owned brewery located in Colorado Springs, CO, has steadily established itself as a key player in the local craft beer scene. The brewery serves clients through a blend of on-site sales at its taproom and regional distribution. As the brewery looks to expand its footprint, both locally and regionally, the need to refine its delivery and sales operations is becoming increasingly critical. The brewery faces the challenge of identifying high-value clients who require more frequent visits—during which a sales rep checks in on stock levels, promotes new or seasonal beers, offers restock recommendations, and builds relationships to encourage larger or more consistent orders. The next-day delivery model enables RLB to fulfill those orders within 24 hours, making it especially valuable for accounts that run out unexpectedly or need quick turnaround. Conversely, the goal is to reduce the number of trips to low-performing accounts. Moreover, RLB's distribution network has yet to fully leverage data to optimize sales routes and prioritize customer visits, which could improve efficiency and profitability. This includes identifying key accounts that would benefit from increased visits and improving the efficiency of its sales routes to reduce unnecessary travel. By focusing on improving both sales and delivery operations, the brewery aims to achieve greater scalability, cost-efficiency, and customer satisfaction. As a result, this project will address RLB's core operational challenges, providing actionable solutions that align with the company's long-term goals.

1.1 Problem Statement

The current product distribution and sales process lacks efficiency in customer prioritization and route optimization, hindering Red Leg Brewing's ability to expand its distribution reach and increase sales volume. This project aims to develop a data-driven model to optimize salesperson routes and prioritize customer visits, ultimately increasing product sales.

1.2 Related Work

The brewing industry has seen a significant amount of research focused on sales optimization, distribution logistics, and demand forecasting, which provides a foundation for the present project aimed at enhancing the distribution processes of Red Leg Brewing. Jiang et al. (2020), advocate for customer-specific models to enhance accuracy in alcoholic beverage distribution. Their findings indicate that different customers exhibit varying demand patterns, necessitating customized forecasting methods. This approach is particularly relevant to our project, as it supports the need for a data-driven strategy that prioritizes high-value customers for distribution. However, while their study informs the forecasting aspect, it does not delve into the logistics of how these predictions can be operationalized, which is a critical focus of our research. The work of Acharya

(2013) on vehicle routing problems with time window constraints introduces optimization algorithms aimed at improving delivery efficiency. This study is instrumental in providing methodologies for enhancing delivery routes based on customer delivery windows, which directly informs our approach to developing efficient routing strategies for Red Leg Brewing. However, the scope of Acharya's work is limited to theoretical modeling without real-world application in a brewery context, highlighting the necessity of our practical project that aims to implement these concepts in a real-world setting.

Finally, Ritchie and Brindley (2007) present a framework for supply chain risk management that underscores the need for operational resilience amidst potential disruptions. This research complements our project by suggesting that improving the efficiency of distribution logistics can also enhance supply chain resilience. Overall, while existing literature offers valuable insights into various aspects of brewing logistics and sales optimization, there is a notable gap in research that integrates these elements into a cohesive strategy for operational efficiency in breweries. Our project aims to bridge this gap by developing a comprehensive model that not only incorporates advanced forecasting techniques but also addresses the specific logistical challenges faced by Red Leg Brewing. By combining insights from previous studies with practical applications tailored to the brewery's needs, we aim to provide actionable recommendations that will enhance operational efficiency and support the brewery's growth objectives.

1.3 Organization

This paper is organized as follows: Section 2 presents the data collection methods and analytical techniques employed in this study. Section 3 discusses the key insights derived from the analysis. Finally, Section 4 outlines the next steps in the project, including future recommendations for improving Red Leg Brewing's distribution process.

2. Data and Methodology

2.1 Data

The two-pronged approach of this project meant that data collection came from multiple sources. Data on delivery volume and customer potential was compiled using Ekos, a management software for craft beverage producers. This phase of the project draws from information gathered on historical sales figures and customer visits from Red Leg Brewing. To gain accurate insights, we focused specifically on sales data from 2024, as it provides the most up-to-date snapshot of performance for each wholesale customer. One primary goal is to determine which wholesale customers could benefit most from increased visitation to drive sales and which may require fewer visits. The data collection process ensured that only relevant, quality data was analyzed. This involved cleaning and organizing the dataset to eliminate inconsistencies and streamline the focus on essential sales and visitation metrics. For example, symbols such as dollar signs were removed from sales figures, and businesses with zero recorded visits were excluded, as these accounts lacked the interaction data necessary to assess the relationship between visit frequency and sales performance. Through this approach, we ensured that each data point aligned with the goals of the project, allowing us to form a clear and actionable picture of each business's relationship with our brewery.

For phase one of this project (route efficiency), customer addresses and delivery constraints were both supplied by RLB. Delivery constraints included which day of the week customers could not be delivered to due to delivery windows and RLB's capability. Moreover, the delivery volume for each customer was drawn from Ekos. In almost all cases, RLB receives invoices for orders at least one business day before delivery. Thus, accurate delivery volume can almost always be determined before the delivery driver departs from the warehouse on any given day. Furthermore, a specific goal of this project is to create a compatible interface within RLB's current technology infrastructure, wherein they can seamlessly download the delivery volume data from Ekos and determine the most efficient delivery route daily.

Another essential component of the data set is the monthly sales and visit frequency for each selected business. This data will inform the second phase of this project, wherein we will determine "high potential" customers. By examining the patterns within this monthly sales data, particularly in terms of invoices generated, we can gauge each business's engagement level within our distribution network. The number of invoices issued monthly reveals fluctuations in sales volume, providing insight into each customer's active involvement. Additionally, we incorporated columns to represent the number of monthly visits for each business in 2024, facilitating a deeper analysis of how varying visitation frequencies might influence sales performance. After gathering invoice counts per month, we expanded the dataset by creating new columns to calculate the proportion of invoices generated relative to the number of visits for each month. This calculation provides insight into the efficiency and potential sales impact of each visit, helping to identify trends and opportunities. Together, these metrics allow us to identify patterns in customer engagement and sales outcomes, ultimately guiding an optimized distribution strategy that maximizes the effectiveness of each visit.

2.2 Methodology

Phase 1 – TSP:

The initial phase of this project focuses on optimizing RLB's distribution network by building a dynamic routing system that adapts to daily delivery needs while minimizing travel distance, fuel consumption, and delivery time. Unlike traditional fixed-route methods, this solution handles a key operational challenge: the need to visit a mix of customers—some with strict delivery time windows, others with none. Because customer demand varies week to week and visits are determined by recent invoices, routing must remain flexible yet consistent, ensuring reliable coverage of 20 to 25 stops per day without significant manual oversight.

To solve this, we implemented a genetic algorithm—a population-based optimization technique inspired by natural selection. Each route is represented as an individual in the population, and routes evolve over multiple generations through crossover (recombining parts of two routes) and mutation (randomly swapping stop order). A fitness function evaluates each route based on total travel distance and penalties for arriving outside delivery windows. This hybrid approach balances feasibility and efficiency: stops with defined time windows are prioritized appropriately, while flexible stops are optimized purely for geography.

Rather than rely solely on commercial APIs, we custom-built this system using Python and open-source libraries, giving RLB full control and scalability without vendor lock-in or usage fees. A script running in Google Colab ingests daily sales data, evaluates delivery eligibility, and generates an optimized Google Maps link with the correct stop order. To streamline this process, we developed a method to extract relevant delivery data directly from EKOS, RLB's order management system. By integrating a custom script, we automated the transfer of customer and sales data from EKOS into the spreadsheet used to feed the routing algorithm. This eliminated manual data entry, ensured consistency across tools, and allowed the optimization model to reflect real-time demand. The result is a seamless pipeline from order placement to route generation, supporting faster, more accurate decision-making. The system is modular and easily reproducible—RLB can simply update its delivery table to produce new, constraint-aware routes each day.

Phase 2 – Clustering Analysis:

The methodology employed in phase two was a comprehensive, multi-step approach integrating data clustering, visualization, and statistical analysis to classify businesses by performance and optimize visitation strategies for sales growth. The process began with cleaning and organizing the data to ensure its accuracy and consistency. K-means clustering analysis ($k = 3$) was then applied to group businesses based on sales invoice frequency and amount. The first cluster, Low-Performing Businesses (Cluster 0), consisted of businesses showing consistently low sales with modest improvement throughout the year. The second cluster, Moderate-Performing Businesses (Cluster 1), included businesses exhibiting stable but unspectacular sales. The third cluster, High-Performing Outlier (Cluster 2), contained a single business affected by a unique contract, which caused sales spikes during certain months. Because of the irregular nature of this business's performance, it was treated separately. This clustering helped categorize businesses into distinct performance groups, making it easier to tailor strategies specific to the needs of each group.

Further refinement was done by breaking Cluster 0 into two subclusters: one with consistently low revenue across all months and another with moderate revenue improvement in certain months. This granularity allowed for more targeted strategies, such as increasing visits to businesses in the low revenue subcluster or optimizing visits for those showing moderate growth. Following the clustering process, data visualizations were created to illustrate sales trends and differences in visit frequencies across clusters. These visualizations revealed that high-performing businesses received more frequent visits while low-performing businesses received fewer visits, highlighting the need for potential strategic adjustments.

To explore the relationship between visits and sales effectiveness, the analysis compared the number of invoices generated per visit. Invoices were used as a proxy for consistent customer engagement and purchasing behavior, offering a clearer view of how often clients placed orders following a visit. While sales figures can vary widely due to factors like bulk purchases, seasonal trends, or the size and type of customer—such as restaurants, liquor stores, or golf courses—invoice count provided a more stable and comparable measure across accounts. This approach made it possible to identify businesses that were efficient in generating orders per visit and those that might need a reevaluation of their visitation strategy.

2.3 Assumptions

Several assumptions can be outlined to form the basis for our analysis and model development. First, we assume that any benefits gained from uprooting RLB's current region-based delivery tactic will not outweigh the costs for the business. Moreover, given the relatively standardized road infrastructure in the region, Euclidean distance serves as a reasonable proxy for driving distances, as the road network allows for relatively direct travel between points. While both 2023 and 2024 data were reviewed to identify and confirm seasonal trends, the 2024 dataset was used as the primary basis for modeling visitation

and sales strategies. This decision was made because 2024 reflects substantial growth—nearly 150 more customers than the previous year—making it a more accurate representation of RLB’s current scale and operations. Seasonal fluctuations observed in 2023 were consistent with those in 2024, supporting their continued relevance for future planning. Furthermore, we assume that the customer segments identified (e.g., low performers, moderate performers, and special contracts) will remain consistent in behavior within each cluster across the analysis period. This stability allows us to tailor visit frequencies and forecast sales for each group, assuming that customers in a given cluster are likely to exhibit similar purchasing and engagement patterns in the near term. For simplicity, the analysis does not account for certain external factors, such as economic shifts, competitor actions, or sudden changes in customer preferences, which could impact customer demand or operational efficiency.

3. Modeling And Analysis

3.1 Modeling Strategy

The modeling for this project occurred in two distinct phases. The first phase consists of using an open-source TSP model to optimize delivery efficiency for RLB. Phase two employed a multi-step methodology combining k-means clustering and visualization to optimize visitation strategies for sales growth.

3.2 Analysis

TSP (Phase 1)

The TSP-based routing system offers a meaningful advancement in RLB’s ability to deliver efficiently while adapting to fluctuating customer needs. Its dynamic structure allows for smarter allocation of delivery resources, ensuring that routes remain both efficient and responsive without requiring constant manual input. This has operational value not just in terms of time or mileage saved, but also in preserving RLB’s ability to fulfill next-day delivery promises—a core aspect of its customer service model.

From a strategic perspective, the model enables better planning by aligning delivery routes with actual sales activity, rather than relying on fixed schedules. It prioritizes deliveries that support stronger account relationships and makes room to shift away from low-activity accounts when necessary. As customer volume continues to increase, this adaptability becomes even more critical.

In addition, the modular design of the routing tool positions it as a long-term asset. It creates a sustainable, repeatable framework that scales with business growth and can be expanded with additional constraints in the future. Most importantly, the system sets a foundation for consistent, data-driven decision-making around route planning—reducing inefficiencies and freeing up capacity to better serve high-value accounts.

Clustering (Phase 2)

The data is broken down into three clusters based on performance: Cluster 0 represents low-performing businesses, Cluster 1 consists of high-performing businesses, and Cluster 2 includes a single, unique, high-value contract which significantly impacts revenue but is not typical of the general customer base. Figure 1, which illustrates the average revenue per cluster across each month in 2024, reveals a stark contrast in revenue generation. Cluster 2, representing these special contracts, shows significant revenue spikes, particularly in certain months, while Cluster 0 consistently generates much lower revenue. These outlier cases in Cluster 2 highlight the need for RLB to treat these accounts differently while also ensuring that resources are not disproportionately allocated to them at the expense of more typical customers in Cluster 1.




Figure 1. Average Revenue per Customer




Figure 2. Average Visits per Customer

Figure 2, showing the average visit count per cluster, complements this view by demonstrating that Cluster 2 receives the highest number of visits, followed by Cluster 1, and Cluster 0 receives the fewest. Despite its high revenue, Cluster 2 may be receiving more visits than necessary, given that the contract-based account is stable and requires less frequent attention once the relationship is established. This insight suggests that RLB could potentially optimize its distribution by reducing visits to Cluster 2 and reallocating those resources to Cluster 1, where high-performing businesses could benefit from more frequent visits to maintain or boost their revenue generation. In contrast, Cluster 0 businesses, which are low performers, might need more targeted interventions to increase their visit count and sales.




Figure 3. Subcluster Analysis (Cluster 0)




Figure 4. Average Visits per Customer (Cluster 0)

Figure 3, which breaks down average revenue per subcluster within Cluster 0, provides a look into the performance of the two subclusters within this group. Subcluster 1, although still part of the low-performing cluster, generates significantly more revenue than Subcluster 0. This suggests that there may be a subset of businesses within Cluster 0 that have potential for increased sales, making them a good candidate for more focused efforts. Subcluster 0, with its consistently low revenue, might require a different approach. Comparing the size of Subclusters 0 and 1, we note that 30% of all accounts show significant untapped growth potential, while 25% of current sales visits target accounts with limited opportunities. This insight shows that RLB could consider redistributing resources, focusing more on Subcluster 1, while potentially reducing visits to Subcluster 0, which may not respond well to increased frequency. Figure 5 illustrates the average visit count per subcluster within Cluster 0, showing that both Subcluster 0 and Subcluster 1 receive similar numbers of visits, despite the large disparity in revenue between them. This suggests a misalignment between effort and performance—Subcluster 1, which generates more revenue, could likely benefit from more visits to capitalize on its potential, while Subcluster 0 may not warrant the same level of attention.

By reducing visits to high-revenue but low-need accounts like those in Cluster 2, reallocating efforts toward high-potential businesses in Cluster 1, and reassessing the approach for low-performing businesses in Cluster 0, RLB can improve operational efficiency, maximize revenue growth, and better balance its sales efforts between taphouse and distribution. The insights derived from these figures will help RLB refine its strategy, ultimately contributing to more effective resource allocation and a stronger customer base.

4. Next Steps and Recommendations

The next steps for optimizing Red Leg Brewing's distribution strategy involve focusing on two main areas: implementing optimized delivery routes and determining the optimal frequency of customer visits. First, routing optimization will be crucial for improving the efficiency of the brewery's delivery network. By the TSP algorithm, RLB can reduce travel distances and time, minimize fuel consumption, and streamline the delivery process. These optimization methods will allow RLB to prioritize visits to high-performing businesses in Cluster 1 while still managing deliveries to lower-performing businesses in Cluster 0 efficiently. Furthermore, it is important to take into account traffic patterns, and customer-specific needs to create a flexible and effective routing strategy.

In addition to optimizing routes, determining the frequency of customer visits is another critical next step. Currently, RLB follows a fixed sales schedule that may not be tailored to the actual performance or needs of each customer. Moving forward, a dynamic visit scheduling system for sales should be implemented, driven by data analysis. The frequency of visits should be adjusted based on customer performance, with high-performing customers in Cluster 1 receiving more frequent visits to maintain or increase their revenue, while low-performing customers in Cluster 0 may benefit from fewer visits or focused on specific performance issues. As market conditions change and new data becomes available, RLB can further refine its approach to maximize its profitability. By implementing these strategies, Red Leg Brewing will be better positioned to grow its distribution business while maintaining a balanced revenue split between its taphouse and distribution channels.

5. References

- Acharya, S. (2013). *Vehicle routing and scheduling problems with time window constraints – Optimization based models*. International Journal of Mathematical Sciences & Applications, 3(1), 169-172.
- Jiang, L., Rollins, K. M., Ludlow, M., & Sadler, B. (2020). *Demand forecasting for alcoholic beverage distribution*. SMU Data Science Review, 3(1), 14-29.
- Lerro, M., Marotta, G., & Nazzaro, C. (2020). *Measuring consumers' preferences for craft beer attributes through Best-Worst Scaling*. Agricultural Economics, 8(1), 1-10.
- Ritchie, B., & Brindley, C. (2007). *An emergent framework for supply chain risk management and performance measurement*. The Journal of the Operational Research Society, 58(11), 1398-1411.