Strategic Optimization in Transforming a Flour Supply Chain Network

by

Nadia Assegaf

and

Shivam Gupta

Submitted on May 8th, 2026, in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management
at the Massachusetts Institute of Technology

ABSTRACT

A leading consumer packaged goods company in Latin America operates a complex supply chain connecting ports, mills, warehouses, and distribution centers to serve over 1,200 clients across a US$3-billion wheat and flour portfolio. With the network running at near-full capacity and modest profit margins, the company must determine how to redesign its network and expand capacity to meet demand growth through 2030, as the risks of data-free investment decisions can be significant. To address this, we propose a multi-period mixed-integer linear programming (MILP) model that simultaneously optimizes commodity flows, facility use, and capacity expansion while maintaining service levels. The optimized network reduces fully loaded cost-to-serve by 66.7%, from $371 per delivered ton in the 2024 historical network to $124 per delivered ton on average across the 72-month planning horizon. Scenario analysis further reveals that central region mills are operating at near-maximum utilization, making proactive capacity expansion a strategic necessity to protect service levels through 2030.

Capstone Advisor: Dr. Christopher Mejia-Argueta

Title: Founder & Director, MIT Emerging Market Economies Logistics Lab

TABLE OF CONTENTS

1. Introduction

1.1 Background

The Sponsor Organization is one of the largest consumer packaged goods companies in Latin America (LATAM) and the Caribbean, operating across seven countries, with the country's headquarters contributing roughly 60% of total sales. Over its 60-year history, the company has expanded primarily through mergers and acquisitions with local firms. Its diversified portfolio spans food and beverage, including pasta, cookies, frozen foods, flour products, home and personal care, and aquaculture feed (Oxford Business Group, 2017). As of December 2024, the Sponsor Organization employed approximately 7,000 people and generated US$3 billion in annual revenue, with US$100 million in profit (Sponsor Organization, 2025).

For its Flour and Mills division, the Sponsor Organization manages a supply chain focused on a single country, encompassing both sales and production. The company imports most of its wheat from Canada, the U.S., Argentina, and the European Union (EU) through three main ports: The main port (Port A) near the capital city handling about 75% of volume, the second port (Port C) in the north of the country with 15%, and the third port (Port B) in the south with the remaining 10%. The Sponsor Organization also has access to two newer ports that are not currently in use due to high costs. The division has five strategically located production plants near these ports, with two facilities near Port A, one near Port B, and two near Port C, helping minimize transport distances. A high-level overview of the Flour and Mills supply chain network is shown in Figure 1.

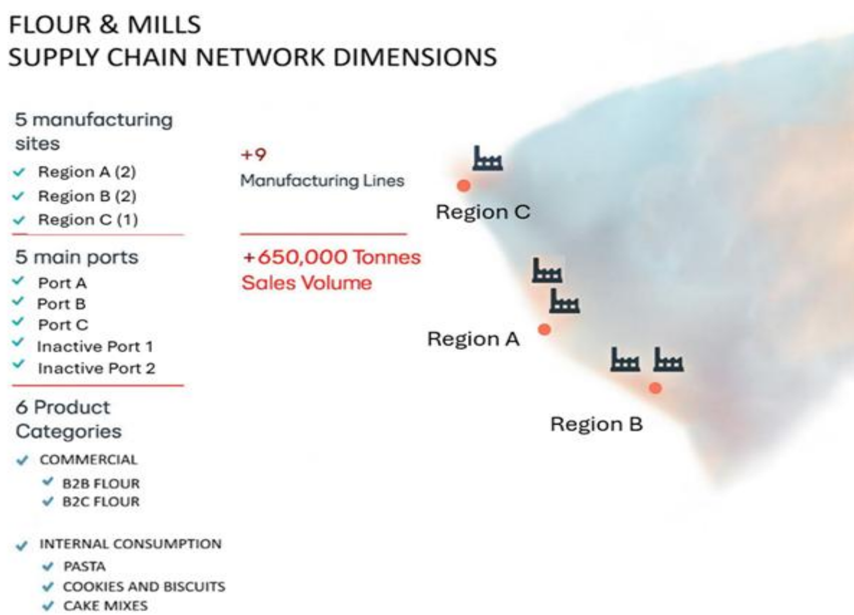

Figure 1: Ports, Manufacturing Plants, and other features (Adapted from Sponsor Organization)

The Sponsor Organization sells flour through two main channels. The business-to-business (B2B) channel supplies over 100,000 bakeries and restaurants nationwide via distributors and wholesalers. The business-to-consumer (B2C) channel offers smaller consumer flour packages sold in supermarkets and local stores. A portion of the flour produced is also used internally by other divisions of the Sponsor Organization to make products such as pasta, cookies, cake mixes, and seasonal bakery items (Sponsor Organization, 2025).

1.2 Motivation and Problem Statement

The Sponsor Organization's Flour and Mills division is currently operating near maximum capacity and facing mounting pressure to sustain competitive service levels amid rising logistics costs. Because this division operates on low margins, it is critical to systematically optimize both capital expenditure (CAPEX) and operating expenditure (OPEX). As the company plans for the next five years, it must determine how to expand existing facilities and reconfigure its supply chain network to accommodate forecasted demand growth, while accounting for non-movable manufacturing sites and a limited CAPEX budget.

Without a data-driven optimization model, the sponsor risks underinvesting in critical locations or unnecessarily overextending capacity, potentially leading to lost customers and disruptions to other internal divisions (such as cookies and pasta). Therefore, this project develops a comprehensive decision-support framework to balance structural efficiency and capital investment. Three core questions drive our research:

1. *Network Design and Configuration*
   - How should the Sponsor Organization redesign its regional supply chain network to efficiently meet projected five-year demand while maintaining required service levels and minimum costs?
2. *Investment and Capacity Planning*
   - Which facilities should receive CAPEX investments, considering brownfield (existing site expansion) opportunities to improve return on investment (ROI) while supporting forecasted demand growth?
3. *Optimization and Resilient Network Design*
   - How can specialized optimization tools (e.g., Gurobi Modeler) be applied to balance cost minimization and capacity needs in the long term in the case of dynamic demand?

1.3 Project Goals and Expected Outcomes

The goal of this project is to design an integrated, data-driven framework that optimizes the Sponsor's regional supply chain network and supports strategic capital investment decisions for the next five years. The study applies advanced analytical tools and optimization modeling to realign production and distribution capacities with projected demand growth while minimizing total logistics costs and maintaining high service-level performance. The model incorporates operational constraints, fixed manufacturing sites, and capacity thresholds, while ensuring recommendations are analytically sound and practically feasible.

The deliverables to the company will include:

1. A network optimization model using mixed-integer linear programming (MILP) to provide expansion recommendations based on financial and operational criteria. This model enables the sponsor to assess various demand and cost scenarios and to evaluate expansions and network configurations under diverse circumstances.
2. Quantifiable results in terms of cost reduction, improved service levels, and resiliency while building a scalable analytical tool for continuous decision-making.

Through this approach, the project determines the optimal configuration of the Sponsor Organization's multi-site manufacturing and distribution network, determining where to expand existing facilities. The optimization framework will provide the management team with clear insights into the trade-offs between cost efficiency, service level, and resiliency, thereby guiding long-term strategic decisions on capital allocation.

2. State of the Practice

The central problem of our capstone is how our Sponsor Organization, operating in a Latin American market, can optimize and redesign its wheat and flour supply chain to achieve minimum cost and resilient capacity. To frame this problem, we reviewed the literature on the structural challenges of food supply chains in emerging markets and the application of mathematical network design to balance cost against resilience.

2.1 Characteristics and Challenges of Food Supply Chains in Emerging Markets

Agri-food supply chains operating in emerging markets face systemic challenges, predominantly driven by underdeveloped infrastructure and fragmented retail sectors (Lorentz et al., 2013). In Latin America (LATAM), poor road infrastructure significantly decreases transportation efficiency and inflates logistics costs (Soysal et al., 2014). Consequently, traditional logistics practices imported from advanced economies often result in bloated, inefficient networks burdened by redundant intermediary warehouses.

Recent literature indicates that LATAM agri-food supply chains can significantly improve operational efficiency and resilience by transitioning to "shorter" supply chains and adopting direct-to-customer logistics (Granillo-Macías et al., 2024). For our capstone, this suggests that an optimized network design must actively evaluate the financial viability of bypassing distribution centers to fulfill demand directly from the production mills.

2.2 Supply Chain Network Design Considering Cost, Service, and Resiliency

To execute these structural shifts, organizations rely on Supply Chain Network Design (SCND) models. Traditional SCND aims to minimize total logistics costs, including transportation, holding, and fixed facility overhead, while satisfying baseline demand. These models are typically formulated as Mixed-Integer Linear Programs (MILP) and solved using commercial optimization engines (Govindan et al., 2017).

However, traditional cost-minimization models often fail to account for operational disruptions (Maharjan & Kato, 2022). Resilient Supply Chain Network Design (RSCND) addresses this by incorporating mechanisms to absorb shocks in highly volatile markets. While some literature proposes complex stochastic modeling, a standard and highly effective proactive strategy is to use scenario analysis to stress-test baseline designs against simulated capacity disruptions (Hosseini-Motlagh et al., 2021). By identifying capacity issues, firms can mathematically justify capital expenditures (CAPEX), such as facility expansions, to ensure service continuity.

Together, these studies provide a clear foundation for our methodology. They highlight the need to use MILP to strip out legacy routing inefficiencies characteristic of LATAM networks, while leveraging scenario-based capacity disruptions to identify where the sponsor must invest in expansions to guarantee long-term resilience.

3. Methodology

This section presents the methodological approach and data used to analyze the Sponsor Organization's supply chain network. We present how the selected methodology is applied to determine the network redesign that minimizes total costs while maintaining service levels and resilience (see Figure 2).

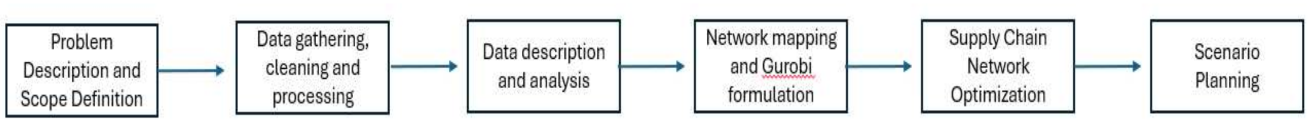

Figure 2: Methodological framework

3.1 Data Description and Analysis and Network Mapping

The data collected from the Sponsor Organization includes the following:

- Sites' master file containing the names, geographic locations, and wheat flows from ports (462 records)
- Customers' master file listing customer names and locations (93,775 records)
- Cost database covering transportation, production, and storage costs in 2024 (1.08 million records)
- Demand dataset providing historical monthly demand from 2024 and future demand forecasts for 2025-2030 (35,903 records)

Descriptive statistics (i.e., averages, variability, minimum and maximum values) were calculated to understand the central tendency and variability of demand and operational data. All datasets were collected directly from the Sponsor through emails and meetings with cross-functional teams. During preprocessing, irrelevant or inconsistent entries were removed using standard data-cleaning and data-quality procedures, including filtering for missing values, correcting formatting discrepancies, detecting outliers, validating entries against operational reports, and cross-verifying to ensure reliability before model integration. The following analysis examines the Sponsor's operational data to identify key demand patterns, import volumes, and transport costs within the wheat supply chain network. The demand analysis, wheat imports, and costs provide the foundation for identifying key drivers of network performance, including capacity utilization, supply concentration, and cost inefficiencies. These insights help parameterize optional configurations and investments. Transportation distances between network nodes were estimated using SCGraph (Makowski et al., 2025), a Python package for road, rail, and maritime shortest-path routing developed at MIT CTL.

3.1.1 Demand Analysis

As part of the demand projection analysis, average monthly demand increases from 37,871 tons in 2025 to 38,946 tons in 2026, reflecting modest annual growth that the network must accommodate through adequate capacity planning. Both years show a 10% coefficient of variation (i.e., the standard deviation of monthly demand divided by the mean monthly demand), indicating stable but meaningful monthly fluctuations that require sufficient buffer storage capacity at ports, mills, and warehouses to avoid bottlenecks during peak months (July – October). The range between the minimum and maximum monthly demand widens slightly in 2026, requiring greater flexibility to manage higher peak volumes, especially in months exceeding 44,000 tons.

3.1.2 Wheat Imports Analysis

The analysis of historical wheat import volumes from 2024 to 2025 provides a deeper understanding of the Sponsor's network structure. Port A is the dominant entry point, accounting for around 75% of wheat imports. Wheat arrivals at Port B remain relatively stable, increasing by 4.2% between 2024 and 2025. This points to consistent utilization of the southern logistics corridor. Port C, in the North, accounts for 9 to 12% of imports and shows the largest decline, dropping by 30% between 2024 and 2025, which may indicate shifting sourcing strategies, changes in vessel routing, or capacity limitations at the northern port.

3.1.3 Logistics Cost Analysis

Wheat imports begin when wheat vessels arrive at designated ports, where the product is unloaded, weighed, and either stored or transported directly to mills for processing. From there, the imported inventory flows through the intermediate warehouse and distribution centers before reaching customers. The Sponsor Organization imports come through from three ports. It possesses five mills, six interims, five distribution centers (DCs), and approximately 1,200 customers (see Figure 3).

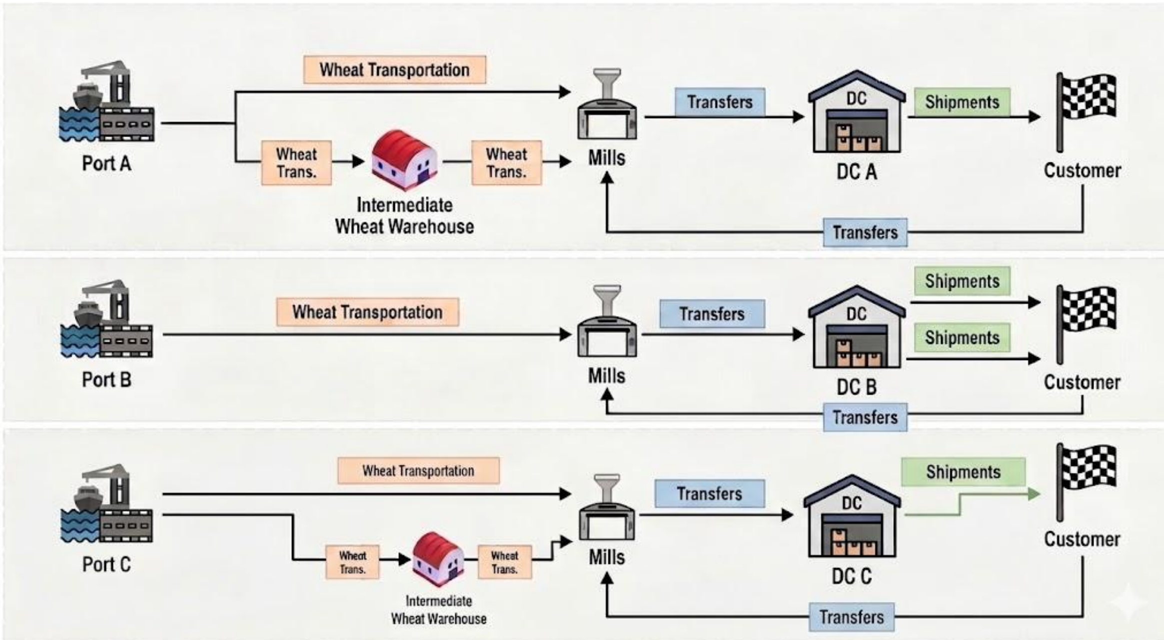

Figure 3: End-to-End Wheat Flow Structure from Ports to Customers (Adapted from Sponsor Organization)

The following presents the average freight cost per ton-kilometer for 2024, based on inbound (ports to mills, ports to interims, interims to mills) and outbound (i.e., mills to DCs, mills to customers, DCs to customers) flows:

- Region A's annual average inbound freight cost is \$0.53 per ton per kilometer, and its outbound freight cost is \$0.67 per ton per kilometer.
- Region B's annual average inbound freight cost is \$0.34 per ton per kilometer, while its outbound freight cost is \$0.86 per ton per kilometer.
- Region C's annual average inbound freight cost is \$0.11 per ton per kilometer, and its outbound freight cost is \$0.33 per ton per kilometer.

Region B (South) has the highest overall logistics costs, driven primarily by significantly higher outbound freight costs, suggesting inefficiencies in last-mile distribution. In contrast, Region C (North) is the most cost-efficient across both inbound and outbound flows, while Region A (Center) shows consistently elevated costs, suggesting broader network inefficiencies.

3.2 Methodology Selection

This capstone adopts a Mixed-Integer Linear Programming (MILP) approach to formulate an optimization model to improve the Sponsor's supply chain network. Based on the readings of Caplice (2015), Sadeghi et al. (2020), and Bhowmik and Parvez (2024), we consider a multi-period MILP as the selected approach because it can rigorously determine the optimal allocation of wheat flows across ports, mills, intermediate warehouses, distribution centers, and customers while adhering to constraints related to capacities, service-level requirements, and multiple scenarios over several time periods. The optimization model allows concurrent modeling of continuous variables (e.g., flows, inventories, losses) and binary decisions (e.g., capacity expansions, facility activations/closures), making it ideal for real-world network investment decisions.

Alternative approaches based on simulations or stochastic methods, such as system dynamics (SD) and discrete-event simulation (DES), require substantial data and detail. System Dynamics aggregate structure cannot capture certain detailed interactions, often referred to as detail complexity (Bell, 2015). On the other hand, DES is primarily used for what-if analyses; models are proposed to evaluate policies, but do not directly produce an optimal network configuration (Banks et al., 2005).

3.3 Optimization Model Application

The MILP model contains three essential components. First, the objective function minimizes total logistics cost, incorporating transportation, inventory holding, and CAPEX for storage or processing expansions. Second, the decision variables determine flow quantities across arcs, inventory levels at each node, and selection of capacity expansions. Third, the constraints enforce operational feasibility, including node capacities, flow conservation, demand fulfillment, minimum service levels, and resiliency.

The roadmap to apply the MILP model to the Sponsor Organization is as follows:

1. Strategy definition and problem identification.
2. Determine model set and parameters by collecting and validating:
   - List and types of facilities (e.g., ports, mills, intermediate facilities, DCs, markets/customers)
   - Allowed arcs and transport modes
   - Historical and forecast demand by product and region
   - Storage and milling capacities
   - Transportation, handling, CAPEX, and inventory holding costs
3. Define all decision variables (i.e., flows, inventories, binary expansion).
4. Set the objective function for the minimum cost flow model.
5. Formulate constraints for flow balance, capacity, and service levels.
6. Implement and solve the MILP in Python using the Gurobi solver, then interpret optimal flows, capacity utilization, and recommended expansions.
7. Run scenario analyses.
8. Provide recommendations and analyze trade-offs on the future network design and CAPEX roadmap.

In the scenario analysis, we design a supply chain network that minimizes costs while ensuring capacity alignment, operational efficiency, and reliable customer service. The alternative network corresponds to the MILP-optimized configuration, featuring a streamlined flow structure with fewer transfers. By comparing the current network with this optimized alternative, we demonstrate how strategic reconfiguration and capacity adjustments can deliver resilience and a scalable network to meet the Sponsor's future demand.

To provide managerial insights, we compare optimized outputs with current network performance to derive actionable recommendations for the Sponsor's five-year supply chain master plan. Figure 4 presents a generic supply chain network that considers ports, mills, DCs, and customers for illustrative purposes of our methodology. However, some elements of the MILP model may change following our discussion with the sponsor organization.

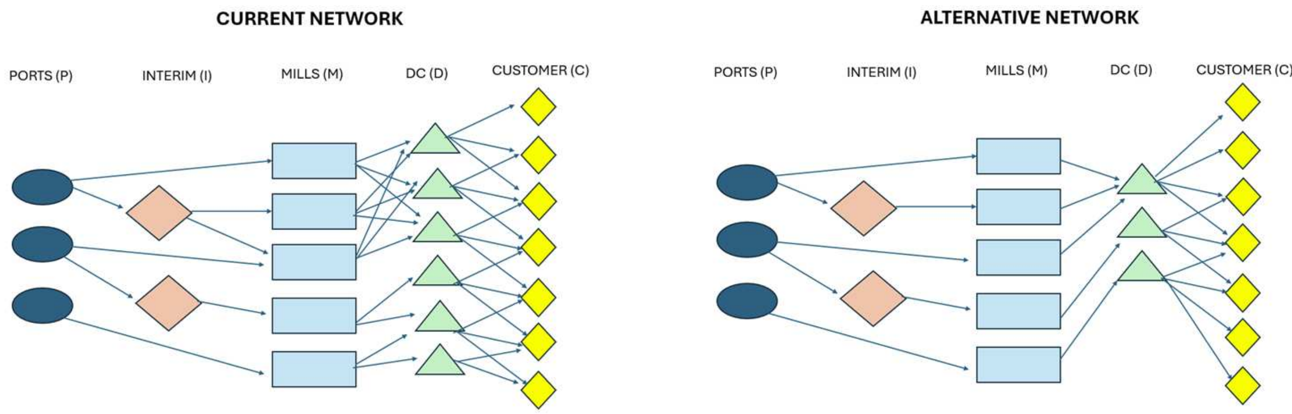

Figure 4. Generic Illustrative Current and Alternative Supply Chain Networks for Sponsor Organization

4. Results and Discussion

This section reports results from the multi-period MILP model applied to the Sponsor Organization's supply chain. Section 4.1 benchmarks the model against 2024 historical operations. Section 4.2 presents the optimized baseline network for the period 2025–2030. Section 4.3 reports results from scenario and sensitivity analyses under four simulated capacity disruptions.

4.1 Model Validation: 2024 Historical Network

The 2024 historical data shows that the network processed 1,878,397 metric tons of wheat across the full supply chain, ultimately delivering 326,602 tons of flour to end customers. The network used the three ports, seven interim facilities, all the mills (five), 10 DCs, and served 1208 unique customers, both internal divisions of Alicorp and external customers. To establish a baseline for comparing the network's operational efficiency, historical cost-to-serve (cost per delivered ton) was calculated across two dimensions: variable logistics cost and fully loaded cost (i.e., variable logistics cost plus fixed facility cost). These figures were derived through a bottom-up, transaction-level aggregation of the 2024 data.

By matching all 93,285 historical shipment records against the provided freight rates, the network incurred $85.7M in total transportation costs across all network legs. Dividing this total system freight by the final delivered volume yields a historical variable cost-to-serve of $262.39 per delivered ton. Furthermore, by identifying the 22 specific facilities actively utilized in 2024 and annualizing their operational overhead, the network incurred $35.6M in fixed facility costs. Incorporating these fixed costs yields a total 2024 network cost of $121.3M, establishing a total cost-to-serve of $371.44 per delivered ton.

The 2024 transaction data reveals a heavily Region A-dependent network. 72.9% of the total raw wheat shipped to processing facilities in 2024 was processed by Region A, with Region B in the South contributing 14.2% and Region C in the North processing the remaining 12.6%. This structural imbalance creates intense pressure on the central facilities. As shown in Figure 7, the baseline 2025–2030 projection anticipates that Capital City Mill 1 will average 92.8% utilization of its 27,258 tons/month capacity, leaving a realistic spare capacity equivalent to less than two weeks of monthly output. While the remaining mills are projected to operate at safer utilization rates (67-87%), Capital City Mill 1 accounts for 47% of the total milling capacity, establishing it simultaneously as the network's primary production asset and its principal point of vulnerability.

4.2 Optimization Results: Baseline Network Configuration (2025–2030)

The results of the optimized baseline model (i.e., actual facility capacities and demand) produced a 72-month total logistics cost of $608.2M, of which $102M (17%) was fixed cost, $408M (67%) was transportation cost, and $97M (16%) was holding cost, with no unsatisfied demand. The model delivered 4.1 million tons of flour to end customers. Under this scenario, no mill expansion was required.

4.2.1 Cost-to-Serve Improvement

To compare the cost-to-serve improvement, we use the same metric we used to evaluate the 2024 instance: cost-to-serve and total cost-to-serve. The cost-to-serve in the optimized baseline model for year 2025 declined from $262.9/ton in 2024 to $100.02/ton ($63.9M / 0.639M tons). The total cost-to-serve for 2025 declined from $371/ton to $127/ton (($63.9M + $17.1M) / 0.639M tons). Over the entire 72-month planning horizon, the average variable cost-to-serve was $98.96/ton ($408.7M / 4.13M tons), and the average total cost-to-serve was $123.85/ton (($408.7M + $102.8M) / 4.13M tons). As shown in Figure 5, this improvement is sustained across the planning horizon. Cost-to-serve is calculated by summing only fixed and transportation costs, excluding inventory holding costs, to ensure comparability with the 2024 historical benchmark, which does not track holding costs due to a lack of historical inventory-level data. The significant reduction in total cost-to-Serve (i.e., a $247.59 per-ton saving) is driven by two factors: tactical rerouting of flows, which accounted for 66% of the per-ton savings, and structural network rationalization, which accounted for the remaining 34%. These mechanisms are elaborated in the remainder of section 4.2.

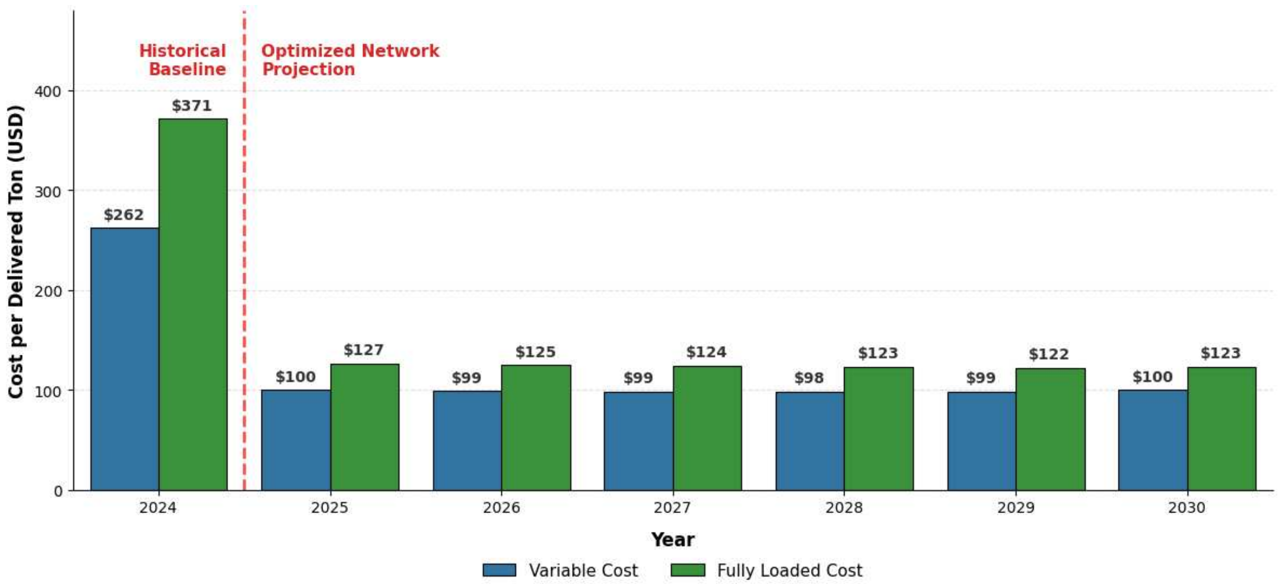

Figure 5. Annual Cost-to-Serve: 2024 Historical vs. Optimized Baseline

4.2.2 Holding Cost Findings

The optimization model also tracked inventory holding costs. Over the 72-month horizon, the model incurred $96.7M in holding costs, representing 15.9% of the total logistics budget of $608.2M. These holding costs are intentionally excluded from the cost-to-serve benchmark and the model's objective function at the Sponsor's direction. This decision is based on the relatively short shelf life of finished flour, which makes prolonged inventory stockpiling operationally unfeasible. Consequently, reducing holding costs was not a strategic priority for the sponsor relative to resolving freight inefficiencies and capacity bottlenecks. We observe that annual holding costs over the planning horizon remain stable at $20–$27 per ton; however, monthly costs fluctuate, ranging from $645,000 in low-demand months (e.g., December) to $2.25M during seasonal peaks (e.g., late Summer). The model uses inventory to substitute physical capacity. It pre-builds safety stock during slow months, incurring temporary holding inventory fees to successfully absorb peak demand surges. The mean inventory held across all periods is 64,600 tons.

4.2.3 Network Rationalization

The model significantly rationalized the facility footprint. As shown in Figure 6, the number of active interim facilities outsourced by the Sponsor was reduced from seven in 2024 to two in the optimized network: Central Interim Hub 2 (Node 6001) and North Interim Hub (Node 6007). Similarly, distribution centers were consolidated from nine down to two: Central DC 1 (Node 1055) and North DC (Node 1001). The consolidation was feasible because the model verified that the nodes had sufficient physical storage capacity to absorb the redirected volume without violating the maximum throughput constraints. All five mills remained active, and no new facilities were opened. The optimization concentrated flow through a select few nodes because the underlying transportation data proved they offer the most cost-efficient freight corridors.

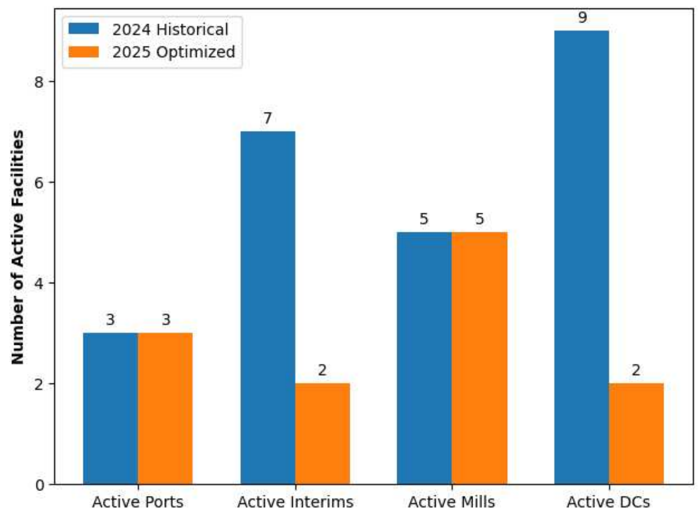

Figure 6. Network Rationalization: Active Facility Count, 2024 Historical vs. 2025 Optimized.

Some relevant highlights per region follow:

- *Region A (Center)* - In 2024, the central region utilized four interim warehouses and at least five distribution centers. The optimized model consolidated this footprint down to one interim and one distribution center. The model selected Central Interim Hub 2 (Node 6001) over competing interims because it provides the lowest total cost for connecting Port A to Mills 1 and 2 in Region A. For example, shipping wheat through competing interim Node 6005 costs up to $20.48/ton. Node 6001, however, feeds Central Mill 1 for just $14.75/ton and Central Mill 2 for $15.57/ton. The model concentrated flow through Central DC 1 (Node 1055) using the same total logic. Calculating the flow from the mill through the warehouse to the customer, Node 1055 has a total downstream cost of roughly $209.67/ton ($15.93 inbound and $193.74 outbound).
- *Region B (South)* - In 2024, this region did not use interim facilities, and it used a single DC (Node 1008). The model consolidated this footprint down to exactly zero active distribution centers. The model shifted the territory to a 100% direct-to-customer fulfillment strategy.
- *Region C (North)* - In 2024, this region utilized one interim [Node 6007] and four active distribution centers [Nodes 1001, 1000, 1002, and 7002]. The optimized model rationalized this footprint down to one interim

hub and one distribution center. It retained only one North DC [Node 1001], whose $100.46/ton outbound cost was the most competitive.

4.2.4 Routing and Sourcing Shifts

As mentioned in 4.2.1, $162.37 of the per-ton savings of the optimization were driven by tactical re-routing shifts. Approximately 70% of this saving is driven by switching to direct deliveries. In 2024, 76% of customer deliveries were routed through a distribution center (DC), while only 24% were served directly from a mill. The optimized model changes this: in 2025, 65% of deliveries are fulfilled directly from mills, while 35% are fulfilled via a DC. This eliminates double-handling, reduces intermediate inventory positions, and shortens the physical supply chain. To illustrate the financial impact of this shift, consider one of the network's highest-volume external customers, Node 30007652, which receives over 90,000 tons across the planning horizon. Shipping this volume through the primary central distribution center, Node 1055, incurs a total downstream cost of $106.84/ton ($15.93 for the inbound leg from Capital City Mill 1 and $90.91 for the outbound leg to the customer). By shifting this customer to a direct-from-mill lane, the total transport cost drops to $83.15/ton, generating logistics savings of $23.68/ton, equating to over $2.1M in cost reduction on this single customer account alone.

Furthermore, for the 35% of the volume that still required DC fulfillment, the MILP algorithm optimally allocated lanes. By re-evaluating the entire transportation matrix, our model systematically removed sub-optimal legacy shipping behaviors from the 2024 baseline, such as fulfilling a customer from a historical default DC when an alternative facility offered a cheaper, localized freight rate.

The tactical savings are achieved through upstream consolidation and raw material handling. Direct port-to-mill shipping fell from 234,747 to 158,075 tons in Year 1 (a reduction of 76,672 tons), with the displaced volume optimally routed through the two retained interims. Central Interim Hub 2 (Node 6001) now functions as a dedicated high-throughput buffer facility for the Capital City mills. This replaces the historical practice of fragmented wheat distribution across five or more equal-weight interim nodes, eliminating the 39,674 tons of redundant interim-to-interim transfers recorded in the 2024 baseline.

4.2.5 Mill Expansion Decision

As mentioned earlier, the company would like to know which of its five milling locations to expand. Thus, the optimization model includes a binary variable to allow expansion to 1 mill, which would add 6,000 tons/month to its production capacity, based on the information provided by the Sponsor. The baseline model did not choose to expand any mill over the 72-month horizon. Instead, it survived demand spikes by rerouting volume and pre-building inventory, rather than paying the high fixed costs to expand a mill. However, capacity at the mills is very tight. As shown in Figure 7 below, Region A averages 92.7% utilization and reaches 99.3% by 2030. The model uses inventory building during slow months to substitute physical capacity. The network successfully carries through the 2025–2030 horizon without expansion, but running the mills this close to their absolute limits leaves the system highly vulnerable to unexpected disruptions.

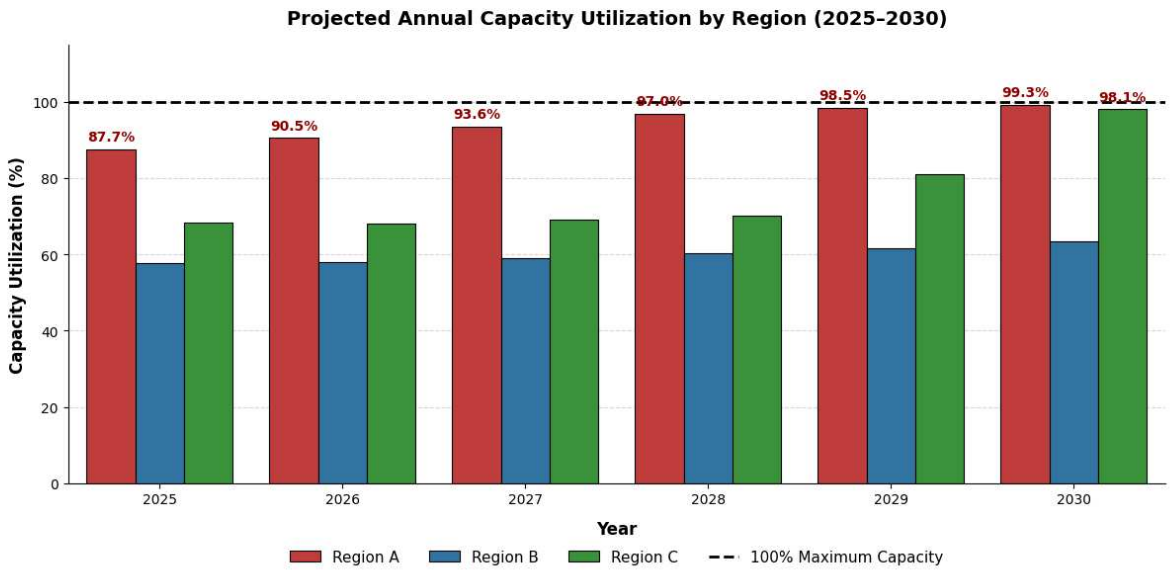

Figure 7. Mill Capacity Utilization – Optimized Baseline (2025–2030)

4.3 Scenario and Sensitivity Analysis

The Sponsor provided feedback that operating mills above 85% capacity is operationally risky, as it leaves insufficient buffer for scheduled maintenance, equipment failures, or new capital installations. To test network resilience under this constraint, four production-stress scenarios were designed and evaluated. Scenarios S1 through S4 reduce system-wide milling production capacity by 10% or 15%, simulating conditions in which simultaneous capacity pressure affects all five mills. For the 15% cases, the model was run three times: once allowing the solver to choose the expansion mill freely (S2), once forcing expansion at Capital City Mill 1 – Node 1015 (S3), and once forcing expansion at Capital City Mill 2 – Node 1016 (S4). This controlled comparison was designed to rigorously determine the lower-cost expansion candidate when the free-choice result was inconclusive. Table 1 below summarizes the key cost metrics and expansion decisions across all scenarios.

Scenario	Description	Total Cost (72-mo)	Variation vs Baseline	Penalty Cost	Penalty % of Total	Expansion
Baseline	Actual capacities	$608.26M	—	$0	0%	None
S1	10% production reduction – ALL mills	$649.90M	+$41.64M	$0	0%	Capital City Mill 2 (Node 1016)
S2	15% production reduction – ALL mills (free expansion choice)	$623.84M	+$15.57M	$23.47M	3.8%	Capital City Mill 1 and Capital City Mill 2

Scenario	Description	Total Cost (72-mo)	Variation vs Baseline	Penalty Cost	Penalty % of Total	Expansion
S3	15% production reduction – ALL mills, expansion fixed at Capital City Mill 1	$623.74M	+$15.48M	$23.47M	3.8%	Capital City Mill 1 (Node 1015)
S4	15% production reduction – ALL mills, expansion fixed at Capital City Mill 2	$624.49M	+$16.22M	$23.47M	3.8%	Capital City Mill 2 (Node 1016)

Table 1. Scenario Analysis Summary.

4.3.1 Scenario 1: 10% Production Reduction – Zero Penalty, Capital City Mill 2 Expansion

Scenario 1 applies a uniform 10% reduction to the production capacity of all five mills. Under this scenario, the model achieves zero unmet demand across all 72 months. Total network cost rises to $649.90M, an increase of $41.64M (+6.8%) over the baseline. The cost increase is driven entirely by fixed costs: the model triggers an expansion at Capital City Mill 2 (Node 1016), adding 6,000 t/mo of production capacity starting in 2028. This expansion adds $52.75M to the fixed-cost component, bringing it to $155.57M (vs. $102.82M in the baseline). In comparison, transport costs fall slightly ($405.14M vs $408.70M) due to distribution optimization enabled by the additional capacity. Holding costs also decline to $89.19M when the model prepositions less safety stock, as capacity buffers are tighter. The zero-penalty result confirms that the network can absorb a 10% system-wide capacity reduction without service degradation. See Figure 8 for further details.

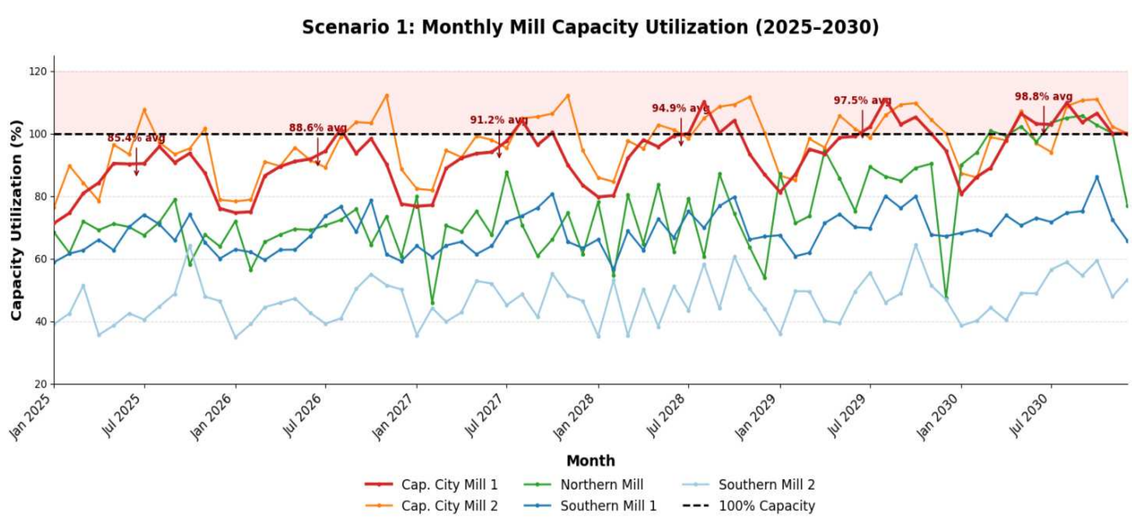

Figure 8: Mill Utilization Rates for Scenario 1

4.3.2 Scenario 2-4: 15% Production Reduction – Structural Penalty and the Expansion Choice

Scenario 2 (S2) simulates a more severe 15% capacity reduction across all mills. Due to the computational necessity of solving the MILP in 12-month sequential batches, Scenario 2 yielded a fragmented, operationally infeasible capital expenditure plan. Thus, the model recommended expanding Capital City Mill 2 in 2028 and 2030, but shifting the 2029 expansion to Capital City Mill 1. Recognizing that the Sponsor must commit to a single, unified facility for a major CAPEX investment, we designed two targeted variations to evaluate the optimal Capital City mill. Scenario 3 (S3) limits the expansion to Capital City Mill 1, while Scenario 4 (S4) limits it to Capital City Mill 2.

At this stage of capacity shortage, the network cannot fully satisfy demand regardless of which site receives the investment. Consequently, all three 15% reduction variants (S2, S3, and S4) incur an identical penalty cost of $23.47M, indicating a structural capacity shortfall that a single 6,000 t/mo expansion cannot overcome. Total network costs across these variants fall within a very narrow range, from $623.74M (S3) to $624.49M (S4). This represents a $15.48M to $16.22M increase (+2.5% to +2.7%) over the optimized baseline.

Since the penalty cost is structurally fixed, the decision between expansion candidates reduces to a comparison of logistics costs (i.e., transportation and holding costs). On this basis, Scenario 3 (i.e., Capital City Mill 1 expansion) is the marginally lower-cost solution, cheaper than Scenario 4 (i.e., Capital City Mill 2 expansion) by $746,133 (0.12%). The choice between Capital City Mill 1 and Capital City Mill 2 should therefore be made on operational and strategic grounds rather than cost, as neither expansion candidate produces a meaningfully different total network cost under the 15% stress scenario.

A notable inventory effect accompanies the 15% stress scenarios. Holding costs collapse from $96.75M in the baseline to approximately $37M across Scenarios 2, 3, and 4, a reduction of nearly $60M. This is not a network efficiency gain: it reflects the model entering a failure regime in which production shortfalls prevent the pre-building of safety stock entirely. The model no longer has sufficient output to buffer seasonal demand peaks, so inventory accumulation becomes physically impossible.

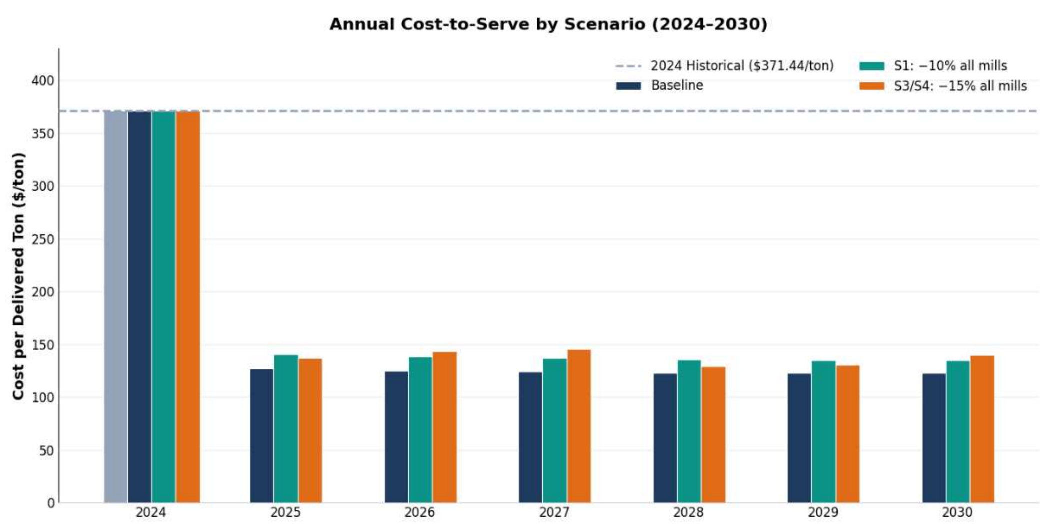

Figure 9: Annual Cost- to-Serve for Scenarios 1-4

Figure 9 extends the cost-to-serve analysis from Section 4.2.1 to the scenario runs. As shown, all scenarios remain below the 2024 historical benchmark of $371.44/ton throughout the planning horizon, confirming that even under production stress, the optimized network is structurally more efficient than the pre-optimization baseline. Scenario 1 (i.e., 10% reduction in capacity across mills) tracks $10–15/ton above the optimized baseline, reflecting the fixed expansion costs spread across delivered volume. Scenarios 3 and 4 (15% reduction in capacity) start at $10/ton above baseline in 2025–2026, spike slightly in 2027 as capacity shortfalls begin to materialize before the 2028 expansion activates, then converge back toward baseline levels in 2028–2030 once expansion capacity comes online. Cost-to-serve is calculated consistently with Section 4.2.1: transport plus fixed costs divided by delivered tons, excluding holding costs. Penalty costs of $23.47M (Scenarios 3 and 4 only) are excluded from this calculation to maintain comparability with the 2024 historical benchmark; including penalty would add approximately $18/ton to the cost-to-serve of Scenarios 3 and 4 in affected years.

The 15%-capacity-reduction stress scenarios also trigger a change in the facility footprint. As shown in Table 3, Scenarios 3 and 4 reopen one additional interim warehouse (Node 6005, 6k tons) and one additional distribution center (Node 1008, 60k tons) compared to the baseline and Scenario 1. These are marginal facilities — Node 6005 handles less than 0.2% of total interim volume — but their activation shows the proposed model seeking additional shipping flexibility when the tighter capacity constraint forces flow redistribution across the network. Scenario 1, by contrast, preserves the baseline's lean two-interim, two-DC footprint, confirming that a 10% reduction with a Capital City Mill 2 expansion is sufficient to serve customers.

Facility Type	Baseline	S1 (−10%)	S3/S4 (−15%)	Change S1 → S3/S4
Active Interims	2 (6001, 6007)	2 (6001, 6007)	**3 (6001, 6005, 6007)**	**+1 (Node 6005, 6k tons)**
Active DCs	2 (1001, 1055)	2 (1001, 1055)	**3 (1001, 1008, 1055)**	**+1 (Node 1008, 60k tons)**
Active Mills	5	5	5	No change

Table 2: Network Facility Footprint by Scenario

4.3.3 Expansion Site Selection: Why Region A Mills?

Across all production-stress scenarios, the optimization model consistently selects a Region A-region mill for capacity expansion. This section documents the two main drivers behind that decision: demand concentration and freight economics.

i. *Region A's demand is the largest and most capacity-constrained in the network.* From 4.32M tons of flour delivered over the 72-month planning horizon, 60.7% are consumed at Region A, compared to 20.1% for Region C (North) and 19.2% for Region B (South). This concentrated demand is already being served by mills operating at near-maximum utilization, as shown in Figure 7. When system-wide production capacity is reduced by 10–15%, Region A mills breach their limits first and most severely - the Southern and Northern mills absorb the same proportional cut but from a position of relative slack. At the same time, Region A has no buffer to draw on. Expanding capacity at Region A-region mills, therefore, directly addresses the demand segment and the geographic node that is simultaneously the largest and the most vulnerable.

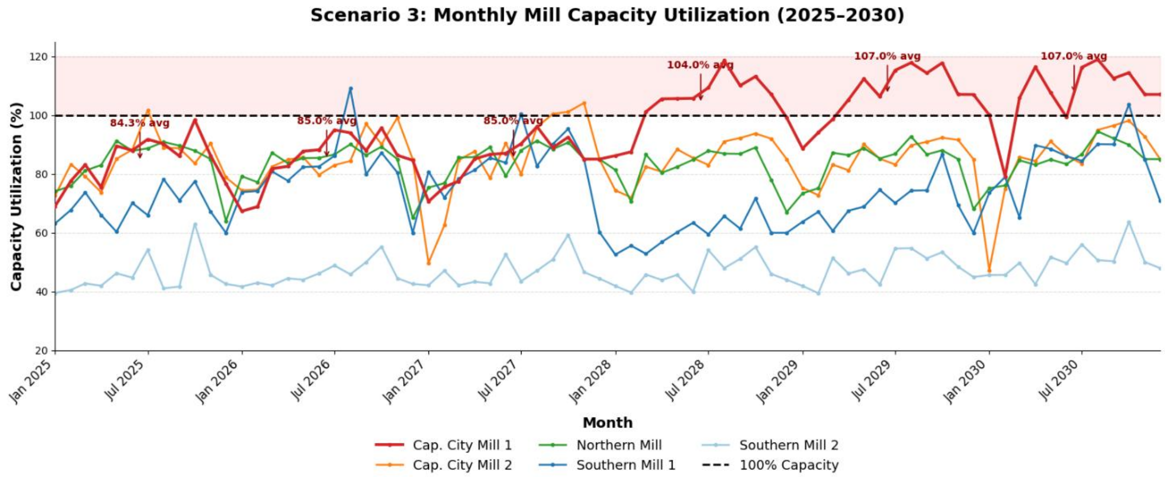

Figure 10: Mill Utilization Rates for Scenario 3

ii. *Cross-regional freight makes expansions in other regions prohibitively expensive.* Table 4 below quantifies the total cost of expanding each mill to serve 6,000 t/mo of demand in Region A over the 36-month active expansion window (2028–2030). The expansion fixed cost is identical across all mills ($7.06M), so the differentiating factor is the freight rate for customers in Region A. Both Region A mills are low-cost candidates: Capital City Mill 2 at $11.67M and Capital City Mill 1 at $16.95M. The difference between them reflects baseline distribution patterns rather than a structural constraint. The Southern and Northern mills are categorically ruled out on cost grounds: their estimated freight rates to Region A (≈$375–437/ton, derived from straight-line distance to the Region A customer centroid, 766 km from Region B in the South, 892 km from Region C in the North, multiplied by the system median long-haul freight rate of $0.490/ton/km from all routes exceeding 200 km in the transportation cost matrix) produce total expansion costs of $88–101M, 7.5–8.7 times the cost of a Region A'S expansion. Even if new direct lanes from Southern or Northern mills to Region A were established, no realistic trucking arrangement over 766–892 km could match the freight rates achievable within Region A.

Mill	Region	Baseline Cap (t/mo)	Freight to Region A (per ton)	Freight Cost (36 mo)	Expansion Fixed (36 mo)	TOTAL EXPANSION COST
Cap. City Mill 2 (Node 1016)	Region A	18,266	$21.35	$4.61M	$7.06M	$11.67M
Cap. City Mill 1 (Node 1015)	Region A	27,258	$45.81	$9.90M	$7.06M	$16.95M
Southern Mill 1 (Node 1007)	Region B	7,840	$375.50*	$81.11M	$7.06M	$88.16M
Southern Mill 2 (Node 1023)	Region B	5,000	$375.89*	$81.19M	$7.06M	$88.25M

Mill	Region	Baseline Cap (t/mo)	Freight to Region A (per ton)	Freight Cost (36 mo)	Expansion Fixed (36 mo)	TOTAL EXPANSION COST
Northern Mill (Node 1605)	Region C	8,699	$436.81*	$94.35M	$7.06M	$101.41M

Table 3. Total cost of expanding each mill to serve 6,000 t/mo of Region A demand over 36 active months (2028–2030).

4.4 Model Limitations

Despite the robust strategic findings, the proposed methodology relies on certain mathematical and operational assumptions that introduce limitations. First, the model utilizes a deterministic MILP framework. It relies on fixed monthly demand forecasts and static cost parameters, inherently excluding the impacts of intra-month demand stochasticity, macroeconomic inflation, and escalating fuel surcharges over the 2025–2030 planning horizon. Furthermore, the model's temporal resolution aggregates flows into monthly periods, assuming that all inter-facility transit times are resolved within the same period as their initiation.

Operationally, the model relies on historical proxies where granular data was unavailable. For example, the silo utilization portion of the model estimates raw material drawdowns using historical wheat-type flow proportions as a proxy for true SKU-level bill of materials (BOM) data. Incorporating exact BOM formulations in future iterations would significantly improve the accuracy of upstream inventory tracking. Operationally, the model currently lacks a constraint to cap direct-to-customer shipment volumes originating from the mills. While the model aggressively routed 65% of volume directly to customers to bypass distribution center costs, mills typically lack the outbound dock infrastructure required to load hundreds of small customer trucks simultaneously. This unconstrained direct-shipping assumption must be treated carefully, and future iterations of the model should introduce outbound dock-door capacity constraints to ensure loading feasibility.

Financially, the model's penalty cost for unmet demand is set strictly at the median product cost per ton. This serves as a highly conservative proxy; it does not capture the cascading contractual penalties, loss of market share, or reputational damage associated with service failures. Consequently, the true financial impact to the Sponsor Organization during high-stress capacity disruptions is likely understated.

Finally, due to processing constraints, the 72-month planning horizon was solved sequentially in 12-month batches. While this rolling-horizon approach successfully passed closing inventory to the subsequent year, it occasionally forced the solver to make "myopic" capital expenditure decisions, such as in scenario two, which recommended two different mills for expansion, because the algorithm could not simultaneously optimize across the entire six-year demand curve. To resolve this, future strategic reviews could migrate the underlying mathematical logic into commercial enterprise network design software, enabling fully simultaneous, end-to-end horizon optimization.

5. Recommendations and Conclusions

The primary objective of this capstone project was to design a five-year supply chain master plan (2025–2030) for the Sponsor Organization's Flour and Mills division. Using a multi-period Mixed-Integer Linear Programming (MILP) optimization model, we evaluated facility footprints, capacity allocations, and product flows across the national network. The analysis revealed that redundant intermediary facilities and sub-optimal routing logic burdened the 2024 historical network. By mathematically streamlining the network, the baseline optimization achieved a 66.7% reduction in fully loaded cost-to-serve, bringing the system-wide average down to $123.85 per delivered ton. Furthermore, scenario analysis revealed that the central region mills are currently operating at or near their absolute maximum capacity. According to the Sponsor Organization's internal policy of not exceeding 85% facility utilization, targeted expansion is needed to maintain a comfortable 10–15% buffer to absorb demand surges or supply disruptions. These findings show that targeted capital investments, structural realignment, and dynamic planning are needed to protect future service levels and profitability.

5.1 Management Recommendations

Based on the MILP model's quantitative outputs, we recommend that the Sponsor execute three primary strategic initiatives. First, the company must carry out a structural network rationalization and transition to direct fulfillment. We recommend consolidating the secondary facility footprint by reducing the network from 16 intermediary locations to four. Specifically, active interim warehouses should be reduced from seven to two (retaining only Node 6001 and Node 6007). At the same time, distribution centers (DCs) should be consolidated from nine to two (retaining only Node 1001 and Node 1055). Simultaneously, the downstream fulfillment strategy should shift to prioritize direct-from-mill deliveries, increasing these direct shipments from 24% of total volume in 2024 to 65%. This structural consolidation eliminates double-handling, redundant fixed-facility overhead, and suboptimal lane assignments. For example, correcting routing inefficiencies on low-volume accounts saved upwards of $247 per ton. Ultimately, streamlining the physical footprint directly supports corporate profit-growth objectives by drastically improving operating margins, and shortening the supply chain enhances operational agility to respond to market fluctuations more quickly.

Second, the Sponsor should approve proactive capital expenditure for the central region milling expansion. Scenario analysis (specifically Scenario 3) indicates that expanding Capital City Mill 1 (Node 1015) yields the lowest total network cost. However, because the cost differential between expanding Capital City Mill 1 and Capital City Mill 2 is extremely small, the company should evaluate both sites and select the final location based on local business knowledge, operational readiness, and zoning feasibility. The optimization model proved that relying on the North and South mills to cover central capacity shortfalls is not viable; cross-shipping from the South incurs prohibitive freight premiums and risks unmet demand penalties. Expanding the central region yields a $271.5M net benefit over 36 months by avoiding those penalties. Because the capital region represents over 60% of national demand, protecting this territory is paramount.

Finally, we recommend institutionalizing the optimization engine for strategic Sales and Operations Planning (S&OP). The developed MILP optimization model should be integrated into the Sponsor’s quarterly S&OP meetings and annual CAPEX review cycles. Human planners operating without mathematical optimization naturally accumulate legacy inefficiencies, which cost the 2024 baseline network millions in redundant freight. By using the MILP algorithm to continuously govern lane allocations, the company locks in the $123.85/ton optimized cost-to-serve and prevents sub-optimal routing from creeping back into the network.

5.2 Future Work

To build upon the foundation established in this capstone, future research should transition the model from deterministic to stochastic parameters. Incorporating probabilistic models for external disruptions such as macroeconomic volatility, sudden demand spikes, or fuel-driven fluctuations in inter-regional freight rates would allow the Sponsor to stress-test the network against highly unpredictable global events. Additionally, future iterations of the model could expand beyond the Flour and Mills division to include cross-business-unit logistics, identifying opportunities for shared warehousing and consolidated truckloads across the Sponsor's broader consumer packaged goods portfolio, as well as inventory policies.

In a broader context, the methodological framework developed in this study provides a scalable blueprint for Latin American CPG firms operating in emerging markets by demonstrating how to mathematically isolate the most resilient network design by expanding capacity at the center of gravity rather than relying on cross-regional redundancy.

References

Banks, J., Carson, J. S., Nelson, B. L., & Nicol, D. M. (2005). *Discrete-event system simulation (4th ed.)*. Pearson Prentice Hall.

Bell, M. A. (2015). Methods for enhancing system dynamics modeling [Doctoral dissertation, Lancaster University].

Bhowmik, O., & Parvez, S. (2024). Supply chain network design: An MILP and Monte Carlo simulation approach. *Brazilian Journal of Operations & Production Management*, 21(1), 2-3.

Caplice, C. (2015). CTL.SC2x-supply chain & logistics design – intro to network design. MITx.

Govindan, K., Fattahi, M., & Keyvanshokooh, E. (2017). Supply chain network design under uncertainty: A comprehensive review and future research directions. *European Journal of Operational Research*, 263(1), 108–141. https://doi.org/10.1016/j.ejor.2017.04.009

Granillo-Macías, R., Rivera-Gómez, H., González-Hernández, I. J., & Santana-Robles, F. (2024). Reconfiguration of agrifood supply chain management in Latin America during COVID-19: A brief literature review. *Sustainability*, 16(9), 3743. https://doi.org/10.3390/su16093743

Hosseini-Motlagh, S.-M., Samani, M. R. G., & Saadi, F. A. (2021). Strategic optimization of wheat supply chain network under uncertainty: A real case study. *Operational Research*, 21, 1487–1527. https://doi.org/10.1007/s12351-019-00515-y

Lorentz, H., Kittipanya-ngam, P., & Srai, J. S. (2013). Emerging-market characteristics and supply network adjustments in internationalizing food supply chains. *International Journal of Production Economics*, 145(1), 220–232. https://doi.org/10.1016/j.ijpe.2013.04.033

Maharjan, R., & Kato, H. (2022). Resilient supply chain network design: A systematic literature review. *Transport Reviews*, 42(6), 739–761. https://doi.org/10.1080/01441647.2022.2080773

Makowski, C., Saragih, A., Guter, W., Russell, T., Heinold, A., & Lekkakos, S. (2025). SCGraph: *A dependency-free Python package for road, rail, and maritime shortest path routing generation and distance estimation (MIT CTL Research Paper No. 2025-028)*. MIT Center for Transportation & Logistics.Oxford Business Group. (2017). Country report: Consumer goods and manufacturing sector overview. https://www.oxfordbusinessgroup.com

Sadeghi, A. H., Sun, Z., & Handfield, R. (2020). A mixed-integer linear formulation for a dynamic modified stochastic p-median problem in a competitive supply chain network design. *Processes*, 8(12), 1586, 1-2. https://doi.org/10.3390/pr8121586

Soysal, M., Bloemhof-Ruwaard, J. M., & van der Vorst, J. G. A. J. (2014). Modeling food logistics networks with emission considerations: The case of an international beef supply chain. *International Journal of Production Economics*, *152*, 57–70. https://doi.org/10.1016/j.ijpe.2013.05.017

The Sponsor Organization. (2025). Internal data and interviews with company representatives [Unpublished source].