Logistics and Supply Chain Analytics Group Report

April 2025

Bryan Cheong

Leo Koay

Louis Pedrix

Max Uebele

MSc Business Analytics

Imperial College London

1 Introduction

This report outlines our plan for managing the supply chain for Jacobs Industries' foam technology in the simulation game. Our objectives were to maximise the cash position by the end of the game. We achieved this through rigorous demand forecasting, smart inventory replenishment, strategic capacity decisions, and an effective fulfilment policy. Our approach integrates classical forecasting models (e.g., SARIMA, linear regression, and Croston's method) with simulation-based inventory optimisation heuristics using an adjusted Silver–Meal algorithm. The following sections detail our methodologies.

2 Demand Forecasting Methods

Accurate demand forecasting was central to success. We segmented our approach by region as follows:

2.1 Calopeia

Calopeia exhibited annual seasonality with no trend. To model this, we first split the 730-day dataset into a 365-day training set and a 365-day test set, which allowed us to capture a full seasonal cycle during both model fitting and evaluation. Then, we selected the most suitable model and proceeded to train the model on the whole dataset.

We explored the Exponential Smoothing (ETS), Seasonal ARIMA (SARIMA) and Cosine Fitting forecasting methods:

ETS: We initially explored ETS models but found them unsuitable for Calopeia's data due to R's implementation limitation with seasonal periods greater than 24 observations. With Calopeia's annual (365-day) seasonality, ETS would ignore these critical seasonal patterns, making it significantly less effective than our chosen SARIMA and cosine approaches.

SARIMA: We found that `SARIMA(0,0,0)(0,1,0)[365]` performed better, as it captured the seasonality through seasonal differencing rather than estimating each seasonal effect explicitly.

Cosine Fitting: We implemented a non-linear model fitting a cosine function with a 365-day period to the training data using `nls()`. If this approach encountered convergence issues, we automatically fell back to harmonic regression using both sine and cosine terms: `lm(y~cos(2*pi*t/365) + sin(2*pi*t/365))`. This provided a smooth representation of the seasonal pattern.

Combined Model: After evaluating individual model performance, we created a weighted ensemble combining 20% SARIMA and 80% Cosine forecasts. This hybrid approach offered advantages over both individual models. The SARIMA forecasts captured the natural variability in the data but appeared noisy with possible overfitting to historical patterns. In contrast, the cosine model provided a smooth, no-noise representation of the seasonal pattern but lacked responsiveness to short-term fluctuations. By weighting toward the cosine model (80%) while incorporating some SARIMA influence (20%), our combined approach balanced theoretical seasonal structure with appropriate variability. This integration allowed us to capture the overall seasonal shape while maintaining sensitivity to historical demand patterns.

2.2 Sorange

Sorange demonstrated a clear upward trend in demand with no seasonal variation. We used the first 75 days of demand for training and the last 15 for testing to evaluate model performance. Then we trained the best model using the whole dataset, consistent with our approach for Calopeia.

We explored ARIMA and Linear Regression as our forecasting method:

ARIMA: We found that `ARIMA(5,1,0)` performed well. The first differencing ($d = 1$) removed the linear trend component, while the autoregressive terms ($p = 5$) captured short-term autocorrelations in the differenced series. This structure allowed ARIMA to adapt to the stochastic nature of the trend and provided a good fit.

Linear Regression: We found that a simple linear regression model provided the best performance for Sorange's demand pattern. The model captured the clear linear trend effectively and offered good interpretability with its equation: Demand $= \beta_0 + \beta_1 \times$ Day. This approach avoided potential overfitting while still accommodating the steady growth trend observed in Sorange.

2.3 Entworpe

Entworpe exhibited highly intermittent demand, characterised by infrequent but fixed 250-unit order spikes. To handle this, we applied Croston's method. Croston's method specifically addresses intermittent demand by decomposing the time series into two components: the demand size when non-zero demand occurs (consistently 250 units), and the inter-arrival time between demand occurrences. This makes it more suitable than traditional time series models in this case.

2.4 Tyran and Fardo

For Tyran and Fardo, we assumed stable demand with limited historical data. To replicate the outlier-like behaviour observed in the historical data, we modeled both regions using a simple normal distribution with a mean of 8 units and a standard deviation of 8. Daily demand was generated using:

```
rnorm(days, mean = 8, sd = 8)
```

Negative values were truncated to zero, and all forecasts were rounded to the nearest integer.

3 Demand Forecasting Insights

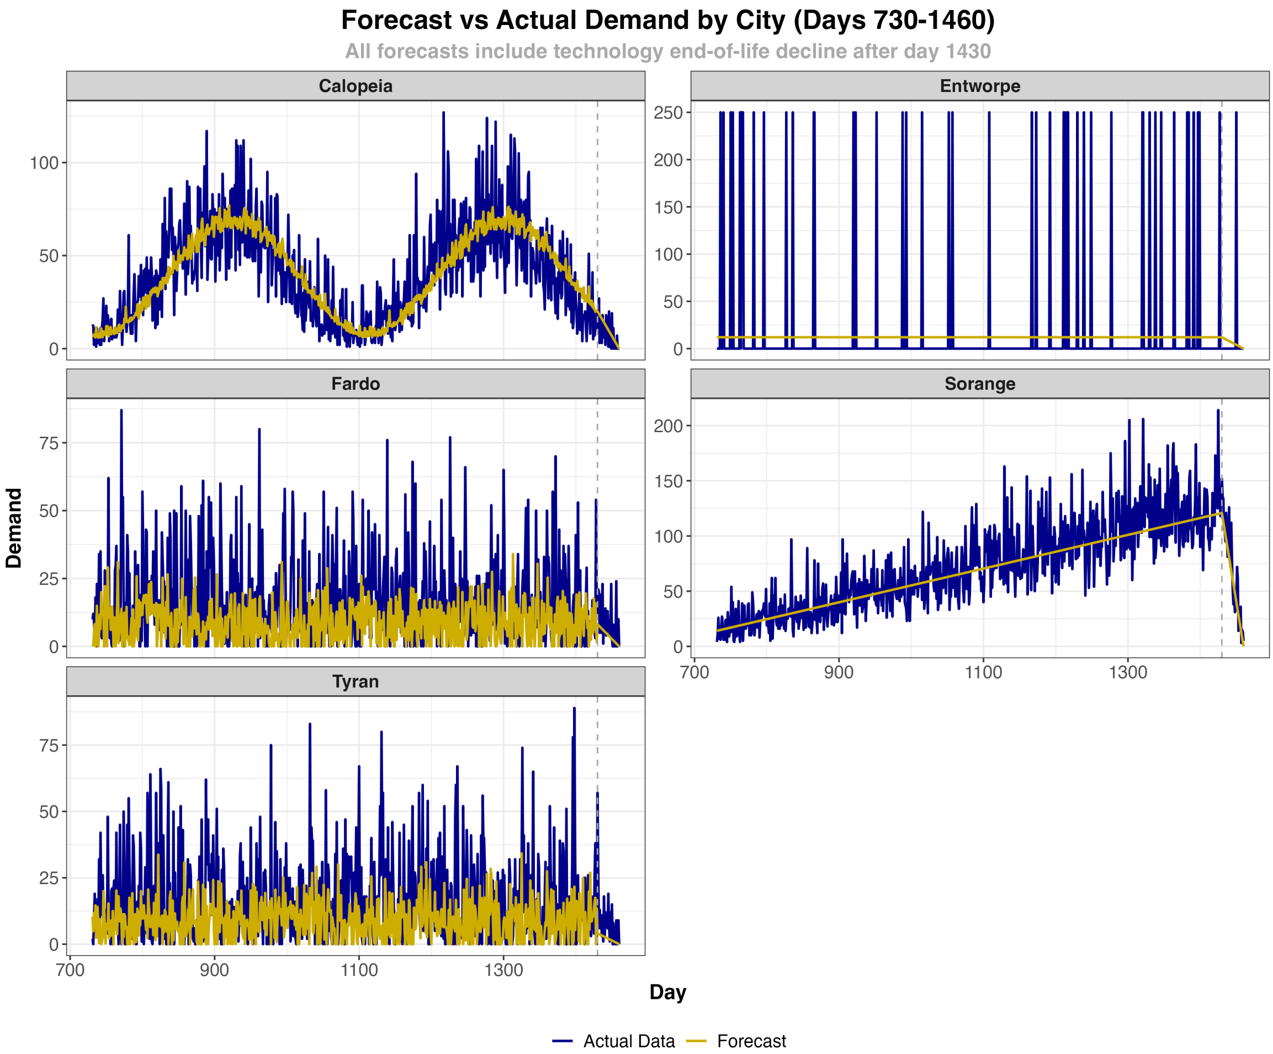

Figure 1: Actual and Forecast Demand from Day 730 to 1460 with Linear Decline from Day 1430

Figure 1 shows the forecasted demand for all cities from Day 730 to 1460. It illustrates the projected demand trends, including a linear decline after Day 1430 to reflect technology end-of-life effects. This figure displays both actual and forecast demand. A combined plot that shows forecasts only is available in Appendix A (Figure 3) for comparisons in the same scale.

4 Demand Forecasting Performance (Post-Game Analysis

City	Method	RMSE	Actual Avg.	Forecast Avg.	% Diff.
Calopeia	20% SARIMA, 80% Cosine	17.24	39.36	**39.04**	-0.81%
Sorange	Linear Regression	25.83	72.23	**67.49**	-6.57%
Entworpe	Croston's Method	NA	16.10	**12.03**	**-25.28%**
Tyran	Mean 8, SD 8	19.24	16.40	**8.99**	**-45.18%**
Fardo	Mean 8, SD 8	18.63	16.58	**8.72**	**-47.40%**

Table 1: Demand forecast performance (RMSE, averages, and % difference) from day 730 to 1460 of our selected models

Table 1 summarises the forecast performance for each city, including RMSE and the percentage difference between forecast and actual average demand post-game. (For the results of all models, see Appendix A's Table 6)

- For **Calopeia**, the Combined model performed slightly better (lowest RMSE) than the SARIMA and Cosine Fitting model but by a small margin, with forecast averages within 2% of the actual.
- **Sorange**'s linear regression, underestimated the actual average, perhaps re-running the regression with new data during the game will improve the results.
- **Entworpe**'s forecast using Croston's Method showed a 25% underestimation, which surprised us and caused us to lose significant demand, especially when it spiked two or even three times in a row.
- The probabilistic models used for **Tyran** and **Fardo** resulted in large underestimations of demand, with differences exceeding **45%**, suggesting that the average is not eight even though it was mentioned in the documents.

5 Long-Term Strategic Decisions

5.1 Cost Benefit Analysis

Since the goal of the simulation is to maximise cash, we evaluated three strategic scenarios for each region and selected the configuration that maximised cash by the end of the game.

Although there are many potential combinations of strategies, we made several simplifying assumptions:

- We used **truck transportation costs** in our forecasts, since trucks have half the inbound logistics cost of mail. While truck delivery takes 7 days and may lead to higher inventory costs, we judged the savings from reduced inbound cost to outweigh these downsides.
- Production cost was calculated using the average unit cost of a **600-unit batch**. We selected this benchmark as it provided a balance between economies of scale and production flexibility. Smaller batches would raise unit costs, while larger batches might delay production for other regions and lead to lost demand.

Using the demand forecasts from Section 2, we analysed whether it would be profitable to build a factory and/or a warehouse in each of the regions not currently served by Jacobs Industries: Sorange, Tyran, Entworpe, and Fardo.

Scenarios Considered:

- **Scenario 1:** Serve the region using Calopeia's warehouse (W) and factory (F)
- **Scenario 2:** Build a warehouse in the region and get inventory from Calopeia's factory
- **Scenario 3:** Build both a warehouse and a factory in the region

5.1.1 Tyran Region Scenario Breakdown

Note that the other breakdowns are in Appendix A, represented by Tables 7, 8 and 9.

Table 2: Tyran: Scenario Cost-Benefit Comparison

Item	Scenario 1	Scenario 2	Scenario 3
Construction cost (Factory)	–	–	1,000,000
Construction cost (Warehouse)	–	100,000	100,000
Total demand	5,760	5,760	5,760
Inbound transportation	432,000	576,000	432,000
Outbound transportation	1,152,000	864,000	864,000
Production cost	5,774,400	5,774,400	5,774,400
Revenue	8,352,000	8,352,000	8,352,000
Profit	993,600	**1,037,600**	181,600

5.1.2 Summary of Most Profitable Scenarios Across Regions

Table 3: Regional Profit Comparison Across Scenarios (Updated)

Region	Avg Demand	Scenario 1	Scenario 2	Scenario 3
Sorange	78	3,611,200	4,759,200	**5,507,200**
Tyran	9	993,600	**1,037,600**	181,600
Entworpe	12	1,324,800	**1,416,800**	608,800
Fardo	9	(154,000)	**306,000**	146,000

Note that in Scenario 1 and 2, the construction cost for Tyran's factory was set to zero. This reflects our assumption that Calopeia's existing excess production capacity could be used to serve Fardo, making a new factory in that region unnecessary. Building a factory in Fardo would simply add to excess capacity and reduce profits.

Final Decisions: The most profitable scenario was to **build both a factory and warehouse in Sorange**, and to **build only a warehouse in Tyran and Entworpe** while serving them from Calopeia's factory. Although Scenario 2 in Fardo offered a marginal profit of $306,000, we opted not to serve this region due to the risk of negative return if demand was overestimated. The upside was too small to justify the working capital requirements and initial investment. Not to mention, there is a 10% annualised daily return rate in the game, which works against tying up capital in low-return opportunities. As a team, we agreed that excluding Fardo was the more prudent and risk-averse decision.

5.2 Factory Capacity Decision

We carefully considered our factory capacity planning, especially given that any further capacity adjustment would result in a 90-day operational downtime, and capacity reductions are non-refundable. This made it essential to provision sufficient capacity upfront to avoid operational constraints or missed demand later in the game.

Based on our latest demand forecasts, the average daily demand across regions is as follows:

Table 4: Average Daily Demand by Region

Region	**Avg Demand**
Calopeia	42
Tyran	9
Entworpe	12
Total (served by Calopeia)	**62**
Sorange	78
Overall Total	**141**

Calopeia has an existing factory capacity of 70 units per day. Given that it needs to serve itself, Tyran, and Entworpe (total average demand of 63), this leaves a buffer of 7 units per day for flexibility and safety stock. Sorange, with an average daily demand of 78 units, requires a dedicated factory due to its high volume.

To ensure operational robustness, we opted to build an 80-unit capacity factory in Sorange. This gives us a total system-wide capacity of **150 units per day**, which provides a small buffer over the total projected demand of 141 units per day.

Strategic Rationale:

- **Buffer for volatility:** Although demand forecasts are based on historical averages, real-time fluctuations are inevitable. A 9-unit buffer (150 vs. 141) allows us to respond to demand spikes without sacrificing service levels.
- **Inventory repositioning flexibility:** Higher capacity improves our ability to quickly produce and redistribute inventory across regions. This is particularly important when switching focus between regions during periods of uneven demand, allowing us to adjust to fulfillment needs without delay.
- **Risk mitigation against production delays:** Any unplanned disruptions, such as temporary under-utilisation, batch transitions, or replenishment lag, can be absorbed more easily with a modest overcapacity.
- **Avoiding future adjustment penalties:** Overbuilding slightly is a safer long-term decision than underbuilding, since capacity increases come with a 90-day production halt, which would likely lead to severe lost revenue.

Given these considerations, our choice of **70 units of capacity at Calopeia and 80 units at Sorange** reflects a conservative, risk-managed approach to long-term production planning while maintaining the agility to support dynamic regional demand.

5.3 Calopeia Inventory Optimisation with Adjusted ROPs and Silver-Meal

To address the seasonal demand patterns in Calopeia, we employ a two-step approach: first, we compute an adjusted reorder point; then, we integrate this into a modified Silver-Meal heuristic for determining order quantities.

5.3.1 Adjusted Reorder Point

Given that demand varies seasonally, a static reorder point is suboptimal. We capture the seasonal pattern using a cosine function similar to the cosine fitting model used for forecasting:

$$D(t) = a + b \cos\left(\frac{2\pi t}{365} + \phi\right),$$

where:

- a is the baseline (average) demand,
- b determines the amplitude of seasonal fluctuations, and
- ϕ is the phase shift.

We then adjust a base reorder point, ROP_{base}, to account for expected demand during lead time:

$$ROP_{\text{adjusted}}(t) = ROP_{\text{base}} \cdot \frac{D(t + LT)}{a},$$

where LT is the total lead time (production plus shipping). Here, $\frac{D(t+LT)}{a}$ acts as a seasonal multiplier, scaling ROP_{base} (typically explored over a range, e.g., 600–1000 units) to reflect anticipated demand at the time of order arrival.

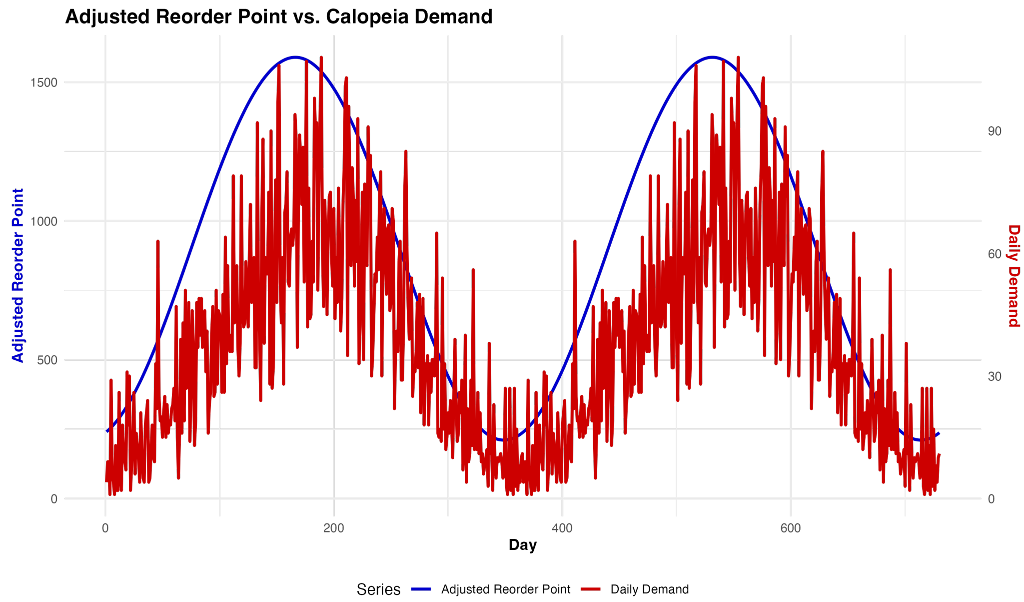

Figure 2: Adjusted Reorder Points Guideline through Cosine Fitting, Note That Blue Curve is Shifted Left Slightly Which Shows Lead-Time Adjustment

5.3.2 Integration into Modified Silver-Meal

The traditional Silver-Meal heuristic calculates an optimal order quantity by minimising the average cost, assuming constant demand. However, our environment is dynamic. We modify the approach as follows:

- **Triggering Orders:** The dynamically computed $ROP_{\text{adjusted}}(t)$ serves as a trigger. When the current inventory level falls below this threshold, it signals the need to replenish stock.
- **Determining Order Quantity:** Upon triggering, the Silver-Meal heuristic is applied to compute the economical order quantity, balancing fixed ordering costs and holding costs over a given planning horizon.

Thus, while the adjusted reorder point indicates *when* to order based on seasonal demand, the Silver-Meal heuristic determines *how much* to order. Together, they provide an inventory policy that adapts to both the seasonal variations in demand and the underlying cost structure.

The output of the Silver-Meal algorithm is as follows:

Table 5: Silver-Meal Ordering Schedule

Order #	Start Day	T (Days)	Quantity (drums)
1	45	48	1,402
2	93	39	1,764
3	132	33	1,979
4	165	38	2,623
5	203	37	2,384
6	240	49	2,169
7	289	46	906
8	335	48	480
9	383	43	766
10	426	43	1,445
11	469	33	1,628
12	502	35	2,236
13	537	36	2,389
14	573	42	2,662
15	615	39	1,591
16	654	46	906
17	700	30	321

Due to the inherent uncertainty in the game and the many assumptions required, the Silver-Meal output and the adjusted ROP are utilised primarily as benchmarks and guidelines rather than as definitive inventory policies. Moreover, given that there are four other regions to serve, high order quantities are not ideal as they can result in longer production times, potentially affecting the production schedule and service levels for the other regions.

5.4 Aggressive Production in Sorange

According to our demand forecast, Sorange's production capacity would be fully utilised by approximately Day 1160. From that point onward, we would have to rely entirely on both ongoing production and pre-built inventory to meet the continuing excess demand. Our calculations indicated that a total of approximately 5,000 units would be required to bridge the gap until Day 1440, when demand would eventually fall back below our production threshold.

Given this insight, we made a strategic decision to begin operating Sorange's factory at full capacity starting from Day 860—well before peak demand—while demand was still relatively low. The objective was to accumulate sufficient inventory in advance to meet future shortfalls.

Although this approach resulted in a high inventory holding cost and left our cash reserves minimal for a significant portion of the game, the team ultimately agreed that the future profits from fulfilling peak demand would outweigh both the holding costs and the opportunity cost of capital, including the 10% annualised daily return available elsewhere in the game. In hindsight, this aggressive production strategy was essential to unlocking the full revenue potential of Sorange during its high-demand phase.

5.5 Low Quantity Was Assumed to Be Sufficient for Entworpe

Our forecasted average demand indicated that Entworpe as a region, will on average consume 12 drums a day. This means that orders would arrive approximately once every 20 days. Based on this assumption, we believed that maintaining an order quantity of around 600 units, with a reorder point set at 400 units, would be sufficient to meet demand.

This inventory strategy was intended to ensure that we could fulfill up to two consecutive demand spikes without stockouts, while also allowing time for replenishment. At the time, our supply lead time was approximately 14 days, which appeared safe given that it was shorter than the expected

20-day interval between spikes. However, this assumption was quickly proven incorrect, as detailed in Section 7.3.

5.6 Reliable Demand in Tyran

Tyran's demand was forecasted to average around 9 units per day. Based on this, we adopted a replenishment strategy with an order quantity of 400 units and a reorder point set at 300 units. This configuration was expected to provide approximately 30 days of coverage before requiring replenishment.

The rationale behind this decision was grounded in the observed stability of Tyran's demand. The region did not exhibit apparent seasonality or trend over time, which was also indicated by the information provided in the game briefing. This consistent demand pattern gave us the confidence to maintain a modest level of inventory that would balance reliability and cost-efficiency.

5.7 Summary of Reorder Points and Order Quantities

Calopeia During non-peak periods, Calopeia's reorder point was set at 800 units with an order quantity of 600 units. This setup provided a balance between holding cost and production efficiency. During peak demand periods, we increased the reorder point to 1,500 units and reduced the order quantity to 400 units. The lower order quantity allowed for more flexible and responsive replenishment, helping us quickly replace depleted stock and avoid missed demand during consecutive demand spikes.

Sorange For most of the game, Sorange operated at full production capacity. We set its reorder point at 5,000 units with an order quantity of 1,000 units. The higher order quantity allowed us to benefit from economies of scale while maintaining sufficient cash reserves to ensure operational flexibility. This was especially important during inventory accumulation phases when Sorange's factory would be running at a net negative cash flow.

Entworpe As discussed earlier, we initially set Entworpe's reorder point at 400 units with an order quantity of 600 units. However, this proved inadequate due to the region's spiky demand pattern. We eventually revised the reorder point to 750 units while keeping the order quantity at 600 units. Entworpe was designated as the second priority for Calopeia's factory, following Tyran.

Tyran Tyran's initial settings included a reorder point of 300 units and an order quantity of 400 units. However, actual demand in the region turned out to be nearly double our forecast. As a result, we increased the reorder point to 400 units and maintained the order quantity at 400 units. Given its higher-than-expected and consistent demand, Tyran was assigned the highest priority for Calopeia's factory capacity.

6 Short-Term Replenishment Decisions

6.1 Mail or Truck?

Throughout the simulation, we predominantly used **truck shipping** for inbound logistics. This decision was driven by two main factors: (1) truck costs are exactly half the cost of mail, with a substantial difference of $75 per unit, and (2) truck shipments enforce a minimum batch size of 200 units, which helped us reduce the per-unit production and transportation cost by leveraging economies of scale.

For the majority of the game—up until around Day 1200—demand remained below our total production capacity. As a result, truck shipments were both cost-effective and operationally feasible. The only exceptions were rare and specific cases, such as when Entworpe experienced three consecutive demand spikes and inventory was completely depleted. In those cases, we used **mail shipments** to quickly replenish stock and prevent further lost demand.

One critical insight we gained during the game was that switching from truck to mail and then back to truck did not offer much strategic benefit. In effect, using mail simply pulled forward production volume and did not create net value unless we had confidence that demand would soon return to levels below our production capacity. Otherwise, once mail shipping began, we were often forced to continue relying on it—leading to sustained high costs—rather than accepting temporary stockouts and preserving long-term margin via truck shipments.

Therefore, our policy was to use mail **only in two specific scenarios**:

- When inventory levels were critically low and we were unable to meet short-term demand spikes, particularly in intermittent-demand regions like Entworpe.
- During the final days of the game (after Day 1430), when demand was declining and we needed the flexibility to bring in small batches quickly to match tapering demand without overstocking.

6.2 Supply Chain Configuration

As outlined earlier in the report, we structured our supply chain so that the Calopeia factory would serve three regions: Calopeia itself, Tyran, and Entworpe. Based on our forecasts, Calopeia's existing production capacity exceeded the combined average daily demand of these three regions, giving us confidence that it could reliably fulfill all orders under normal operating conditions.

In situations where the regional warehouses in Tyran or Entworpe experienced stockouts, we planned to temporarily fulfill their demand using Calopeia's warehouse, provided sufficient inventory was available. However, during high-demand periods—particularly in the first year when Calopeia's own demand peaked—we made a strategic decision to prioritise Calopeia. This was due to both profitability and capacity constraints because fulfilling Tyran and Entworpe's orders from alternative warehouses would have incurred an extra $50 per unit in outbound logistics costs, further reducing their contribution margins and making them less economically viable under constrained conditions.

Meanwhile, the Sorange factory was dedicated exclusively to serving Sorange. As explained previously, its production capacity was just sufficient to meet Sorange's projected demand, assuming it ran near full capacity for most of the game. As such, redirecting Sorange's output to support other regions was neither feasible nor economically justified.

7 Things We Could Have Done Better

7.1 Misinterpretation of Initial Demand Description

In the Round 2 briefing, we relied on the note that stated the average order size for Fardo, Tyran, and Entworpe was 8 units. From this, we assumed the actual mean demand would hover around that value. Despite observing that the stabilised demand data (Days 670–730) was significantly higher than the stated average of 8, we dismissed it, assuming the sample size was still too small and likely influenced by outliers. This led to a significant underestimation of demand in those regions. With a more realistic assessment, we likely would have opted to serve Fardo and potentially generated approximately $1 million in additional cash by the end of the game. Unfortunately, by the time we recognised this underestimation, it was too late to act, as building a warehouse or factory involves significant lead time.

7.2 Fulfillment Policy Selection

At the start of the game, we were unfamiliar with the impact of the fulfillment policy setting. As a result, we did not change the default setting from "soonest" to "nearest." In hindsight, choosing the "nearest" fulfillment policy from the beginning would have allowed us to use nearby warehouses more effectively, enabling regions to assist each other in case of stockouts. The "soonest" policy prioritised immediate fulfillment, but it led to suboptimal logistics, particularly when a local warehouse could not meet demand. This limited flexibility ultimately contributed to missed demand.

7.3 Missed Demand from Inventory Mismanagement

Across the four regions we served, we estimated that approximately 7,000 units of demand were lost due to stockouts and poor inventory positioning. Several contributing factors include:

- **Calopeia Inventory Accumulation:** We failed to accumulate sufficient inventory in Calopeia during its low-demand phase in Year 1. When demand in Calopeia surged in the middle of Year 1, we were unable to simultaneously support Entworpe and Tyran, resulting in unmet demand in those regions.
- **Suboptimal Reorder Point for Entworpe:** Our inventory management in Entworpe was overly conservative, largely due to concerns about incurring holding costs. However, in retrospect, the holding cost was minimal compared to the profit lost from failing to fulfill 250-unit batch orders. We also failed to account for the operational delay caused by simultaneous production scheduling. Specifically, when Entworpe experienced a sudden spike in demand, the factory was often already engaged in producing batches for other regions. As a result, switching to Entworpe's batch required waiting for the current production to finish, followed by truck shipping—which introduced an additional 7-day delay. In the worst-case scenario, the total lead time was excessive, and we underestimated the risk and impact of this bottleneck.
- **Sorange Inventory Buffer:** Although we initially planned to accumulate 5,000 units of inventory in Sorange before production reached parity with demand, we later reduced the buffer to 4,000 units due to concerns over high inventory holding costs. This decision was also influenced by our poor standing on the leaderboard at the time, which led us to second-guess our original strategy. In hindsight, our initial analysis was sound, and we would have needed approximately 6,000 units of inventory to effectively bridge the gap.

7.4 Lack of Adaptive Forecasting

While we did build forecasting models and made a few operational adjustments throughout the game—like updating Entworpe and Tyran's reorder points and order quantities once we realised demand was higher than expected—our forecasting process overall was quite static.

Looking back, we should have been more proactive in incorporating real-time data from the simulation and regularly updating our model parameters. Reforecasting with the additional data points we gathered during the game would have likely led to more accurate predictions and helped keep our operations better aligned with the actual demand trends as they evolved.

8 Conclusion

This simulation gave us a valuable hands-on experience in navigating the complexities of a dynamic supply chain. Our careful selection of forecasting models led to relatively small errors in key regions like Calopeia and Sorange, allowing us to make proactive decisions around capacity planning, inventory build-up, and replenishment strategies. These strong foundations contributed to our team's ability to maximise cash flow and ultimately secure 2nd place in the game.

However, one major lesson we learned was the importance of interpreting briefing information more carefully. Our early assumptions about the average demand in Entworpe, Tyran, and Fardo turned out to be inaccurate, leading to underestimations that could have been avoided by clarifying ambiguities or verifying assumptions with the professor before the game began.

Despite that, our strategy overall was logically sound. We built up inventory during low-demand periods, adjusted reorder points and quantities based on real-time observations, and maintained a hands-on approach by checking the game frequently to react quickly. Moving forward, we aim to enhance our forecasting by fitting models more adaptively and using new data as it becomes available. Additionally, we will place greater emphasis on stress-testing our decisions against worst-case scenarios to fully understand the risks and trade-offs of each strategy before implementation.

A Appendix:

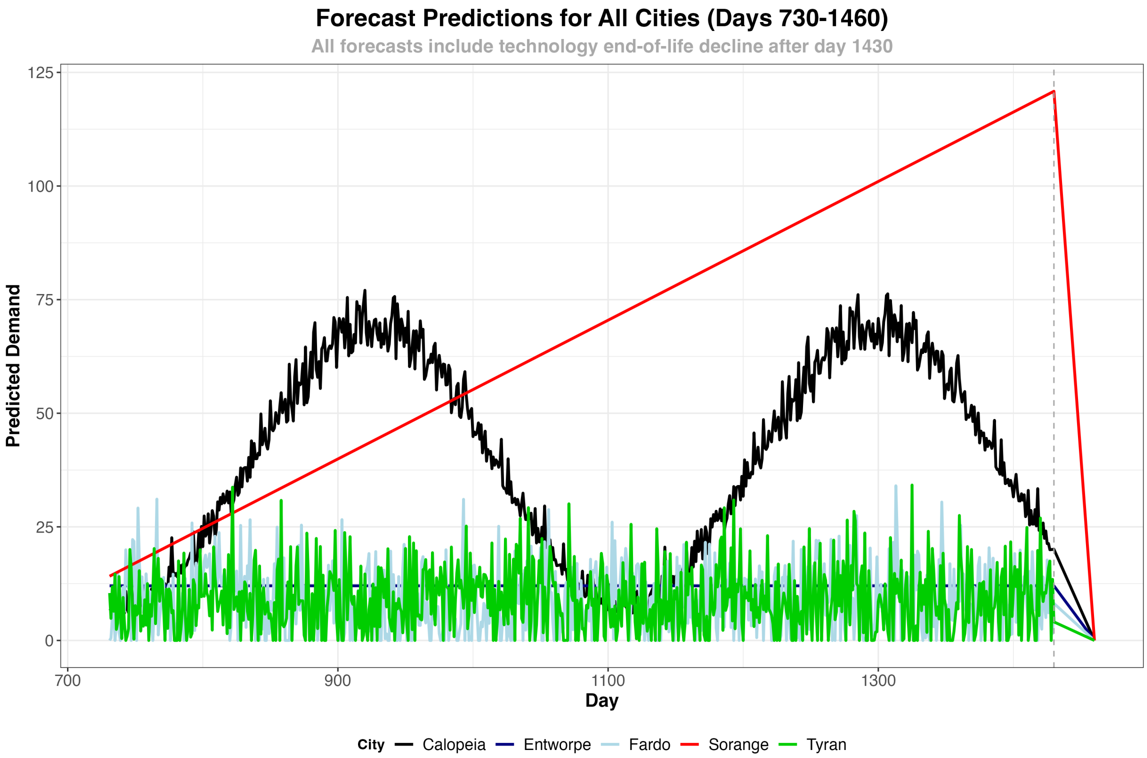

Figure 3: Demand Forecast from Day 730 to 1460 with Linear Decline from Day 1430

City	Method	RMSE	Actual Avg.	Forecast Avg.	% Diff.
Calopeia	SARIMA(0,0,0)(0,1,0)[365]	21.54	39.36	38.62	-1.88%
Calopeia	Cosine Fitting	17.58	39.36	39.14	-0.56%
Calopeia	20% SARIMA, 80% Cosine	17.24	39.36	39.04	-0.81%
Sorange	ARIMA(5,1,0)	27.23	72.23	66.22	-8.32%
Sorange	Linear Regression	25.83	72.23	67.49	-6.57%
Entworpe	Croston's Method	NA	16.10	12.03	**-25.28%**
Tyran	Mean 8, SD 8	19.24	16.40	8.99	**-45.18%**
Fardo	Mean 8, SD 8	18.63	16.58	8.72	**-47.40%**

Table 6: Demand forecast performance (RMSE, averages, and % difference) from day 730 to 1460 for all models

Table 7: Fardo: Scenario Cost-Benefit Comparison

Item	Scenario 1	Scenario 2	Scenario 3
Construction cost (Factory)	–	–	1,000,000
Construction cost (Warehouse)	–	100,000	100,000
Total demand	5,600	5,600	5,600
Inbound transportation	420,000	1,260,000	420,000
Outbound transportation	2,240,000	840,000	840,000
Production cost	5,614,000	5,614,000	5,614,000
Revenue	8,120,000	8,120,000	8,120,000
Profit	(154,000)	**306,000**	146,000

Table 8: Entworpe: Scenario Cost-Benefit Comparison

Item	**Scenario 1**	**Scenario 2**	**Scenario 3**
Construction cost (Factory)	–	–	1,000,000
Construction cost (Warehouse)	–	100,000	100,000
Total demand	7,680	7,680	7,680
Inbound transportation	576,000	768,000	576,000
Outbound transportation	1,536,000	1,152,000	1,152,000
Production cost	7,699,200	7,699,200	7,699,200
Revenue	11,136,000	11,136,000	11,136,000
Profit	1,324,800	**1,416,800**	608,800

Table 9: Sorange: Scenario Cost-Benefit Comparison

Item	**Scenario 1**	**Scenario 2**	**Scenario 3**
Construction cost (Factory)	5,000,000	5,000,000	5,500,000
Construction cost (Warehouse)	–	100,000	100,000
Total demand	49,920	49,920	49,920
Inbound transportation	3,744,000	4,992,000	3,744,000
Outbound transportation	9,984,000	7,488,000	7,488,000
Production cost	50,044,800	50,044,800	50,044,800
Revenue	72,384,000	72,384,000	72,384,000
Profit	3,611,200	4,759,200	**5,507,200**