Data-Driven Stockout Prediction in the Foodservice Supply Chain

by

Mengqi (Lucas) Tang

and

Evgeniya (Jenny) Gerasimov

Submitted on May 8, 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

Large chain restaurant brands operate complex and layered supply chains, where distribution center-level stockouts can spread downstream and cause shortages at the restaurant level. When a stockout is identified at the DC level, Armada Supply Chain Solutions (the sponsor company) responds with aggressive mitigation measures across all downstream locations to recover inventory at any cost. Such uniform action does not consider the likelihood, timing, and severity of the potential downstream stockouts that could occur due to this loss of inventory. This study develops a methodology for quantifying the resiliency of individual restaurants at the SKU level, using weekly inventory snapshots and daily shipment records to estimate consumption rates and project inventory levels for each downstream restaurant. The model classifies each restaurant into a risk tier based on projected days of inventory cover and gives planners a structured framework for prioritizing recovery actions across a network of hundreds of locations. By identifying which restaurants have enough on-hand inventory to absorb a disruption and routing them through standard shipping instead, the tool directly reduces unnecessary emergency freight spend. Even small reductions in the frequency of expedited shipments across a large restaurant network can result in substantial cost savings. Armada is recommended to pilot the tool with one or two brands before scaling it across the network. In the longer term, the company is encouraged to invest in building the data infrastructure needed to support more advanced predictive methods of stockout prediction.

Capstone Advisor: Dr. Milena Janjevic
Title: MIT Supply Chain Design Lab, Director

TABLE OF CONTENTS

1. INTRODUCTION

The foodservice supply chain presents a unique challenge in balancing cost efficiency with service levels. For restaurant businesses, shortages of critical food items cause not only lost sales and operational disruption but also immediate damage to brand reputation. When a customer cannot order a menu item, the entire visit experience is compromised, and the lost sale cannot be recovered. When encountering a stockout, roughly 31% of customers switch to a competitor, while 9% forgo the purchase altogether (Gruen et al., 2002). Unavailable items also prevent customers from purchasing complementary and related products (Anderson et al., 2006), which results in additional revenue losses. Since holding excess safety stock is not practical for perishable goods, the risk of stockouts becomes even more acute.

Supply chain complexity further increases this vulnerability. Large chain restaurant brands often operate extensive, multi-layered networks. Products move from suppliers through several tiers of distribution and forwarding centers before reaching restaurants, where they are finally served to customers. Stockouts at higher echelons of the supply chain can have a ripple effect and cause shortages at downstream nodes. Mitigating such stockouts becomes a costly undertaking, forcing companies to rush freight, execute inter-distribution center (DC) transfers, and expedite orders with suppliers — all to ensure product availability at the restaurant level is maintained.

Our capstone project sponsor, Armada Supply Chain Solutions, provides a case study for this challenge. Armada is a US-based provider of outsourced supply chain services in the foodservice space. Their services include supply chain planning, warehouse and transportation solutions, global logistics, and freight management (Armada, 2025). The company serves large fast-food restaurant brands and positions itself as an extension of their clients' supply chain organizations. They emphasize proactive disruption management, captured in their motto: "Resolving issues before you know they're there" (Armada, 2025).

1.1 Motivation

Armada acts as an integrator and orchestrator within the multi-tier supply chains of its foodservice clients. Armada is not a traditional logistics provider as the company does not own warehouses or control daily operations: its value proposition lies in data management, communication, and decision-making across the network. It functions as a "control tower" that enables visibility and coordination across brands, DCs, and restaurants, as well as proactive disruption and risk management for the entire network. Fast food brands set strategic service and quality standards and contract directly with Armada for orchestration and DCs for warehousing and delivery. Each DC within the network manages forecasting and ordering independently but depends on Armada's analytical oversight and problem-solving both during regular operations and when disruptions occur. In the event of disruptions, Armada uses its network-level visibility to reallocate shipments and mitigate stockouts.

Currently, Armada utilizes various practices to optimize operations at the distribution center level and minimize stockouts. These practices include setting target service levels by item, maintaining safety stock according to the target service level, and implementing tools for inbound pipeline visibility. However, none of these practices can eliminate stockouts completely. When a stockout is identified at the DC level, the sponsor company assumes that the DC-level stockout will lead to a stockout at the restaurant level and responds aggressively to recover inventory at any cost. Likelihood, timing, and severity of the downstream stockouts are not considered in this approach, which leads to expensive and unnecessary mitigation measures. Current approach also overlooks the fact that individual restaurants operate with their own inventory buffering strategies and can often absorb short-term disruptions even without additional resupply. The sponsor company seeks a data-driven way to distinguish critical stockouts from non-critical ones at the DC level, which would allow them to respond selectively rather than uniformly and ultimately reduce unnecessary operating costs while avoiding stockouts at the restaurant level.

1.2 Problem Statement and Key Questions

Our objective is to develop a methodology for quantifying the resiliency of individual restaurants at the stock-keeping unit (SKU) level. One way to measure resiliency is by estimating the probability of stockout for each SKU at each restaurant at any given time. When a stockout is detected at the DC level, this probability can be used to prioritize recovery actions.

Restaurants that have the highest stockout risk would receive intervention immediately, while those with lower risk could be served through more economical options such as Less-Than-Truckload (LTL) or ground shipping, rather than costly expedited shipments.

To examine this approach, we address the following key questions:

- What methodology should the sponsor company follow to calculate the stockout probability for each SKU at each restaurant?
- What are the key inputs, drivers, or variables that determine whether a restaurant will stock out? What assumptions should we make, and what are the limitations of our model?
- How can stockout probability be expressed to reflect both the timing and severity of potential restaurant-level shortages?

1.3 Project Goals and Expected Outcomes

Ultimately, the model will estimate the probability of a restaurant-level stockout given a DC-level stockout for a specific SKU. For example, if SKU A is expected to stock out at the DC today, the model will project the stockout probability for each restaurant that depends on that SKU over the next seven days. With these probabilities, the sponsor company can identify which restaurants are truly at risk and how soon a stockout might occur if they do not receive a new shipment of that SKU. This enables them to determine how much time they have to respond and the latest possible date by which a shipment must arrive to prevent a restaurant-level disruption.

This visibility will allow the sponsor company to deploy expedited recovery actions only for high-risk locations while slower, lower-cost alternatives can be used for lower-risk locations. This enables the shift from uniform response strategies to a more cost-efficient targeted prioritization.

2. STATE OF THE PRACTICE

In recent years, as data availability and analytical capabilities have grown, data-driven predictive methods have become the prevailing approach to managing stockout risk. Such methods allow planners to anticipate when inventory will run out and intervene before a shortage occurs. Two broad analytical approaches are used to anticipate stockouts in foodservice: forecasting demand and inferring stockouts from projected inventory depletion, and predicting stockout events directly from historical outcomes.

2.1 Demand Forecasting–Based Stockout Prediction

Within the first analytical approach, forecasting demand and inferring stockouts, organizations rely on several techniques. Traditional time-series forecasting is the most commonly used tool in this category. It models historical consumption rates and extrapolates them into the future. Common techniques include moving averages, exponential smoothing, the Holt-Winters approach, and Auto-Regressive Integrated Moving Average (ARIMA) models. These four techniques are well established in the forecasting literature. Moving averages, the baseline method, forecast future demand as the average of a fixed number of recent observations (Hyndman et al., 2021). Exponential smoothing builds on this by assigning greater weight to more recent data. The Holt-Winters model, in turn, extends exponential smoothing and adds components for trend and seasonality (Holt, 1957; Winters, 1960). ARIMA is a more complex statistical method that models demand as a combination of past values and past errors and captures autocorrelation and stochastic noise (Hyndman et al., 2021). Selecting the right forecasting method depends on the availability, granularity, and underlying patterns in historical data, as well as on how stable demand and lead times are.

Regression-based prediction incorporates external variables into the model in addition to historical demand values. Recent studies show that this approach improves forecasting accuracy in perishable food supply chains: for example, a 2025 study of a food-demand forecasting system found that combining time-series models with regressors (such as weather, holidays, and price) improved predictions (Hussain & Prabhu, 2025). This approach relies on the assumption that the relationships between predictors and demand are reasonably stable and are captured by the model.

Machine learning (ML) techniques are also used to improve demand forecasts. Algorithms such as Random Forests, Gradient Boosting, and Neural Networks can capture relationships among both historical demand data and external features such as weather, promotions, and macroeconomic conditions. In some retail environments, ML models have been shown to outperform classical time-series methods, specifically if large SKU-level datasets and complex external drivers are available, as illustrated in a 2020 study on bakery sales predictions (Huber & Stuckenschmidt, 2020). In other studies, traditional approaches still showed strong performance: for example, Lindfors (2021) demonstrates that ARIMA-based models match or exceed multiple ML algorithms for retail sales forecasting. Islam et al. (2024) compare the performance of multiple ML methods, including Random Forest, XGBoost, and Neural Networks, as well as hybrid methods. The study concludes that hybrid machine-learning models can deliver higher forecasting accuracy than individual ML algorithms. To be feasible in practice, however, all of these models require large, clean, granular, and reliable historical data and external feature inputs.

More recently, generative AI methods have begun to emerge. Such methods combine time-series forecasts with Large Language Model (LLM)-derived insights from unstructured data. Abad (2025) introduces an approach that correlates residual forecasting errors with semantic signals found in contextual unstructured data. In a retail-specific example, an LLM analyzes internal emails that describe external events and produces estimates of impact, lag, and confidence. These variables are incorporated as additional inputs to the time-series forecast. This approach requires reliable pipelines to capture, clean, and process large volumes of unstructured text.

Once a demand forecast has been produced, stockouts can be inferred either deterministically or probabilistically. Deterministically, planners estimate when cumulative forecasted demand will exceed on-hand stock plus confirmed replenishments. Some organizations extend their demand forecasting models with probabilistic stockout simulation. Instead of treating demand and lead time as fixed inputs, this method models them with probability distributions, commonly Normal or Poisson (Silver et al., 1998). These models simulate inventory depletion and allow demand planners to move from relying on a single forecast point estimate to calculating stockout probabilities. In many foodservice environments, however, historical data is too limited or noisy to reliably estimate demand distributions.

2.2 Direct Stockout Prediction Approaches

In contrast to methods that first forecast demand and then infer stockouts, the second major approach forecasts stockout events directly. Some organizations apply machine-learning classification to estimate the probability of a future stockout based on past purchase orders (PO), inventory, and stockout data. Such supervised classification models must be trained on labeled historical outcomes, which means that organizations must identify and timestamp each past stockout and reconcile those events with corresponding inventory records. A 2025 study by Liu et al. tested such an approach for a large retail company and found that it performs well in handling large and imbalanced data and significantly improves predictive performance.

There are also non-ML-based direct detection mechanisms that identify stockouts without forecasting demand. For example, Sainsbury's "Shelf Availability Monitor" (SAM) uses point-of-sale-based anomaly rules to flag likely out-of-stock items when expected sales patterns disappear (Corsten & Gruen, 2003). However, these approaches identify stockouts only after they occur and do not allow for forward-looking prediction.

Together, these methods illustrate how stockout prediction in foodservices is largely driven by demand forecasting. While there is evidence that probabilistic, machine-learning, and GenAI-based techniques can improve accuracy, they are still being used selectively only where data and capabilities allow.

3. METHODOLOGY

The methodology in this study is shaped primarily by the limitations of the data available from the sponsor company. Restaurant inventory is reported only at a weekly frequency as end-of-week snapshots, while shipment data is available at a daily level but does not include forward-looking purchase orders or confirmed future deliveries. There is also a reporting lag in inventory reports: weekly ending inventory data for each restaurant is only made available to Armada on the Monday following the reporting week. Therefore, Armada has no visibility into real-time, on-hand inventory at the restaurant level when a disruption occurs. This limitation rules out methods that require continuous or real-time inventory tracking, including the deterministic inventory depletion models, probabilistic stockout simulation, and machine learning and generative AI-based demand forecasting techniques described in Section 2.1. In addition, the sponsor does not maintain a history of restaurant-level stockout events with consistent labels. Since emergency or off-schedule shipments often replenish inventory before a stockout is observed, these stockout events cannot be reliably inferred from the available data. As a result, there is no clean target variable that can be used to train or validate models that predict stockouts directly, eliminating the supervised machine learning classification and point-of-sale anomaly detection approaches described in Section 2.2.

These constraints limit the use of demand forecasting methods that require detailed daily consumption data or full visibility into future supply. As a result, this study adopts a stock-flow–based approach that infers consumption and projects inventory using the information that is available to the sponsor. The resulting framework estimates the probability and timing of restaurant-level stockouts following a known DC-level disruption, enabling the sponsor to identify which restaurant–SKU pairs are most at risk over a seven-day horizon and to prioritize recovery actions accordingly.

This paper uses standard set-theoretic notation for clarity and conciseness. The following sections define the key notations, required datasets, and steps used in the analysis.

3.1 Data Collection and Preparation

Notations and Sets

The following notation is used throughout the methodology:

P: Sets of SKUs

C: Sets of Restaurants

T: Today's date

B: Report date of the latest weekly beginning inventory

S: DC stockout date for a given SKU

E: DC stockout end date for a given SKU

W: Sets of Weeks

Z: Sets of dates

J: Weekday stamp for each date t (convert date to weekday, then covert to J = 0 = Mon,…, $\mathrm{J}=6$ = Sun)

$n_{S-B} = \Delta_{BS} = S - B$:Number of days between stockout date and the latest beginning inventory reported date

$n_{T-B} = \Delta_{BT} = T - B$:Number of days between today's date and the latest beginning inventory reported date

$For\ SKU\ p \in P,\ For\ Restaurant\ c \in C, For\ Weekday\ l\ \in \mathcal{D}_c, For\ date\ t \in Z, For\ week\ w \in W$:

> I_{pc}^{t}: Inventory level for a product restaurant pair at the end of the day on date t
>
> I_{pc}^{w}: Inventory level for a product restaurant pair at week W
>
> d_{pc}: Median daily consumption (calculated)
>
> σ_{pc}: Standard deviation of daily consumption per day
>
> r_{pc}^{t}: Observed shipment quantity received on date t
>
> $\hat{q}_{pcj}$: Projected shipment quantity applied on each schedule delivery day J
>
> $\mathcal{D}_c \subset \{0, \dots, 6\}$: set of weekdays when restaurant C receives delivery (0=Mon, 6=Sun)
>
> So $\mathcal{D}_c = \{0, 2, 4\}$ represents Monday, Wednesday, Friday delivery

Dataset Review and Preparation

The methodology relies on two primary data sources provided by the sponsor company:

Weekly Ending Inventory Report. This report contains weekly ending-inventory levels (I_{pc}^{w}) at each restaurant, for all SKUs that a restaurant is holding. Sunday is considered the end of the week. The current-week ending inventory balance becomes the beginning inventory balance for the following week. This report is also used to derive weekly and daily consumption rates (d_{pc}).

DC-to-Restaurant Sales Order Detail. This report lists shipment quantities, ship dates, and store IDs. The Store ID enables a link to the weekly ending inventory report. Ship dates along with store IDs are used to derive a delivery schedule for each restaurant (Dc). The delivery schedule for a restaurant is fixed. Same-day delivery is always enforced between DCs and restaurants.

Together, using these datasets we can reconstruct restaurant-level inventory behavior, consumption rates, and replenishment timing for each SKU. As part of data preparation, both input files are reformatted into structured formats suitable for simulation. Inventory values are converted to numeric format, SKU identifiers are extracted from descriptive fields, shipment dates are normalized to a daily level, and daily shipment quantities are aggregated by restaurant, SKU, and date.

3.2 Consumption Forecasting

We assume that weekly consumption is a representation of demand. Because we do not have visibility over the daily consumption of each restaurant-product pair, we assume daily consumption is evenly spread throughout the week. The weekly consumption is calculated as the median of weekly consumption using all historical data available, weighted equally. The following formula is used to calculate daily consumption:

$$d_{pc} = \frac{median\,\{I_{pc}^{(w)} + r_{pc}^{(w)} - I_{pc}^{(w+1)} | w = 1, \dots, W\}}{7}$$

Using the median filters out unusual "spikes" or "crashes" in weekly consumption caused by bad inventory counts, atypical shipments, special events, or bad data entries.

3.3 Inventory Projection

Inventory projection is the core of the model. Given the DC stockout date S and stockout SKU P, the goal is to compute projected ending inventory levels for each date t in a forecast horizon ($t = B, \dots, B + n$) for each restaurant c that carries SKU p. The model performs the following steps:

1. Identify relevant restaurants C that have ordered the product P in the past.
2. Identify the most recent beginning inventory reported date; set date $t = B$ and set the value of inventory level on date B as the beginning inventory level I_{pc}^{B}.
3. Extract delivery schedule weekdays for all relevant restaurants, $\mathcal{D}_c$
4. Label each date t with J, starting from $t = B \leq S \leq S + 7.\, S + 7$. The seven-day window after the stockout date allows the company to identify opportunities for cost savings using slower or less aggressive recovery methods.
5. Calculate the ending inventory level I_{pc} for each date t. The ending inventory level for date t becomes the beginning inventory level for the next day $t + 1$:

$$Initial\ Inventory = I_{pc}^{t=B}$$
$$I_{pc}^{t} = I_{pc}^{t-1} - d_{pc} + r_{pc}^{t} + \hat{q}_{pcj}$$

   Where r_{pc}^{t} is the observed shipment quantity received (applies only for $B \leq t \leq T$), and $\hat{q}_{pcj}$ is the projected shipment quantity (applies only for $t \geq T$ when J matches a scheduled delivery day in $\mathcal{D}_c$.
6. Calculate $\hat{q}_{pcj}$ based on the median value of all historical shipment quantities multiplied by the percentage share in a week that the restaurant is ordering the product on each scheduled delivery weekday J:

$$M_{pc} = median\{r_{pc}^{1}, r_{pc}^{2}, \dots, r_{pc}^{t}\}$$
$$P_{pcj} = \frac{\sum_{t \in Z} r_{pc}^{t} \cdot \beta_{tl}}{\sum_{t \in Z} r_{pc}} \quad with\colon \beta_{tl} = 1\ if\ t\ and\ l\ are\ the\ same\ weekday$$
$$\hat{q}_{pcj} = M_{pc} \times P_{pc}(j)$$

With these inputs in place, the model can project day-by-day inventory levels for each restaurant that carries the impacted SKU from the latest reported snapshot through the end of the disruption window.

3.4 Stockout Risk Classification

Projected inventory levels are the foundation for assessing stockout risk, measured through days of inventory cover. For a given SKU number, a projected inventory level at or below zero indicates a potential stockout event for that restaurant. Days of inventory cover also quantifies how much time a restaurant has before a potential stockout, allowing the company to act accordingly.

Each restaurant is assigned to one of five risk tiers based on its projected days of inventory cover on the DC stockout date, calculated as projected inventory divided by average daily consumption:

$$\text{Days of Inventory on date } t \text{ for restaurant } c \text{ and product } p = \frac{I_{pc}^{t}}{d_{pc}}$$

Risk tiers definition:

- "Most Critical": Projected inventory at or below zero — immediate action required (e.g. expedited shipping)
- "High Risk": Less than one day of inventory cover remaining — urgent action needed
- "Action Needed": Between one and three days of inventory cover — standard replenishment should be initiated
- "Monitor": More than three days of inventory cover — lower immediate priority
- "Unknown": Days of cover cannot be computed due to a missing or zero consumption rate

The thresholds for each tier were determined in consultation with Armada's supply chain team based on their operational experience with typical lead times in the foodservice industry. This classification can be used to determine selective, cost-efficient responses instead of uniform interventions across all restaurants.

3.5 Model Limitations

The model has the following limitations:

- Limited data history reduces the ability to detect trends, seasonality, and the shape of the consumption distribution, as well as the ordering cadence $\hat{q}_{pcj}$
- External factors (promotions, local events, weather) are not explicitly modeled.
- The model assumes accurate Enterprise Resource Planning (ERP) system data synchronization across DCs and restaurants.
- The method does not consider delivery service capacity.

4. RESULTS AND DISCUSSION

The methodology described in Section 3 was applied to real shipment and inventory data provided by the sponsor and implemented as a working prototype. The following results show the tool's practical value in a supply chain disruption situation.

4.1 Model Implementation and Computational Workflow

Prototype Architecture

The model is implemented as a Python-based prototype in a data notebook environment. It is meant to be run manually whenever a distribution center disruption occurs. The workflow is modular: separate functions handle data transformation, parameter estimation, and daily inventory projection. Intermediate outputs are stored in indexed dictionaries to enable efficient lookup during simulation: these outputs include restaurant delivery schedules, expected shipment quantities, and daily consumption rates per restaurant and SKU. The inventory projection module of the tool calculates daily inventory levels for each restaurant-SKU pair and then combines the results into a structured table. These results are shown to the user through a dashboard that displays the following elements:

- A summary of the data processed
- Key metrics, such as how many restaurants have ordered the affected product and how many are at risk during the disruption
- A detailed table that classifies and prioritizes risk for each restaurant

Data Preparation Pipeline

The workflow of the tool begins by ingesting and transforming the two input datasets (Restaurant Inventory and DC-Restaurant Shipments) into structured formats. The weekly inventory file is cleaned by fixing column headers, removing extra non-data rows, and converting all relevant values into numbers. The DC-to-restaurant shipment file is standardized in a similar way: SKU identifiers are extracted from the SKU description field, shipment dates are normalized to a daily level, and entries where the store ID does not match any restaurant in the inventory file are removed. Daily shipment quantities are then aggregated by restaurant, SKU, and date.

Intermediate Parameter Estimation Modules

After the data is prepared, the tool creates several intermediate lookup objects used during projection. First, delivery schedules are created for each restaurant by extracting historical shipment weekdays from the DC-to-restaurant shipment file. These are stored as a dictionary mapping each restaurant to a list of eligible delivery weekdays. Second, two shipment-related components are created for each restaurant: expected shipment volume and a ordering frequency per weekday. These components are combined to estimate an expected inbound quantity for each restaurant-SKU-weekday combination. Next, restaurant consumption rates are estimated based on historical inventory flow. These are stored in dictionaries as weekly values keyed by (restaurant, SKU) and daily values keyed as (restaurant, SKU, weekday).

Inventory Projection Module

The daily inventory projection module is the main “engine” of the tool. It projects inventory forward for each restaurant that carries the selected SKU based on the previously created parameter objects. The simulation moves forward one calendar day at a time and records each day's projected inventory value in an output table. For each day, inventory is updated by subtracting estimated daily consumption and adding any applicable inbound quantity. For all dates up to the analysis date, the model uses actual recorded shipments. For future dates prior to the DC's expected stockout date, expected shipment quantities are used. After the DC stockout date, inbound supply to the restaurant simulated is set to zero. The reliability estimate of both the consumption rate and the projected shipment quantity estimates is based on the number of shipment records available for a given restaurant-SKU pair in the DC-to-restaurant sales order detail file (referred to as “shipment lines” in the tool interface). If a restaurant only has a few historical shipment lines in the provided data, the shipment quantity estimates are less reliable, which affects confidence in the results. The simulation continues for up to 30 days or stops earlier if inventory reaches zero or below.

Output Structure

The dashboard presents results to the user in four panels: “Load Summary,” “Summary Metrics,” “Risk Bands,” and “Daily Projections”. The “Load Summary” panel shows the shape of the data processed and displays the number of inventory rows, sales rows, and restaurant–SKU pairs that had been loaded into the tool. The “Summary Metrics” panel shows four headline figures. The first is the total number of restaurants that have ever ordered the affected SKU. The second is the number of restaurants included in the projection model after filters are applied. The third figure is the number of restaurants classified as at-risk during the evaluation window. The fourth is the number of restaurants where projected inventory was already below zero on or before the DC stockout date: the user is able to hide these locations from the main table. The “Risk Bands” panel classifies all restaurants into five risk tiers based on their projected days of inventory cover on the evaluation date. The “Most Critical” restaurant label indicates that projected inventory is at or below zero. “High Risk” indicates less than one day of inventory cover remaining, and “Action Needed” indicates between one and three days. “Monitor” indicates more than three days of cover: restaurants in this tier are lower immediate priority. “Unknown” is assigned when days of cover cannot be calculated due to a missing or zero consumption rate. The user can click on any band to filter the table to only that band. An additional toggle is available that allows the user to hide restaurants whose projected inventory was already negative on or before the DC stockout date. This is done to isolate locations whose risk is directly attributable to the disruption as opposed to other unrelated inventory issues.

The "Daily Projections” panel offers the most granular view of restaurant inventory position. Each row in the table shows the restaurant identifier, daily consumption rate, projected inventory on the evaluation date, days of inventory cover, risk band, the first date within the disruption window on which projected inventory reaches zero or below, and the number of shipment records for the affected SKU found in the DC-to-restaurant sales order detail file. The last measure is a proxy for how actively the restaurant has been ordering the item historically. Clicking on any restaurant row in this table expands a detail panel below. It shows the selected restaurant's delivery schedule, average weekly and daily consumption rates, and shipment history statistics, as well as a day-by-day projection table. Each row of that table contains information on the date, daily consumption, inbound quantity, net change in inventory, and projected ending inventory. This data allows the user to trace exactly how and when a given location is expected to run out of its stock. Illustrative examples of the dashboard output are provided in Section 4.2.

4.2 Operational Application: Illustrative Disruption Scenario

The intended user of the tool is a supply chain planner within Armada's control tower organization. The tool is designed to assist the planner in monitoring inventory risk across restaurants in all DCs' delivery networks and coordinate recovery actions in the event of supply disruptions. When a DC-level stockout of any SKU occurs, the planner must determine which restaurants are truly at risk, how much time is available to intervene, and what the intervention should be.

The tool supports this process end-to-end. From the user's perspective, the tool requires four inputs: the affected item's SKU identifier, the DC's expected stockout start date, the DC's expected stockout end date, and today's date. The DC's expected stockout date and impacted SKU come from the ERP/planning system, which projects when inventory will run out based on current stock and planned incoming shipments. The current date is used to map the analysis to the latest available restaurant inventory value and actual observed shipments. Because the prototype is not connected to live ERP systems, the planner must also upload the most recent restaurant inventory report and DC-to-restaurant shipment detail file before running the analysis. Once these inputs are provided, the tool produces the risk classification and prioritization output described in Section 4.1. It ranks all restaurants that carry the affected SKU by their projected inventory position within the DC stockout window.

To demonstrate how the tool's output translates into planner decisions, consider the following scenario.

Assume that on September 10, 2025, the ERP system indicates that the DC is expected to stock out of Item 16302 on September 20, 2025. The planner uploads the data files, enters Item number 16302, September 20 as the DC expected stockout date, September 25 as the evaluation date, and September 10 as today's date, and runs the tool (see Figure 1).

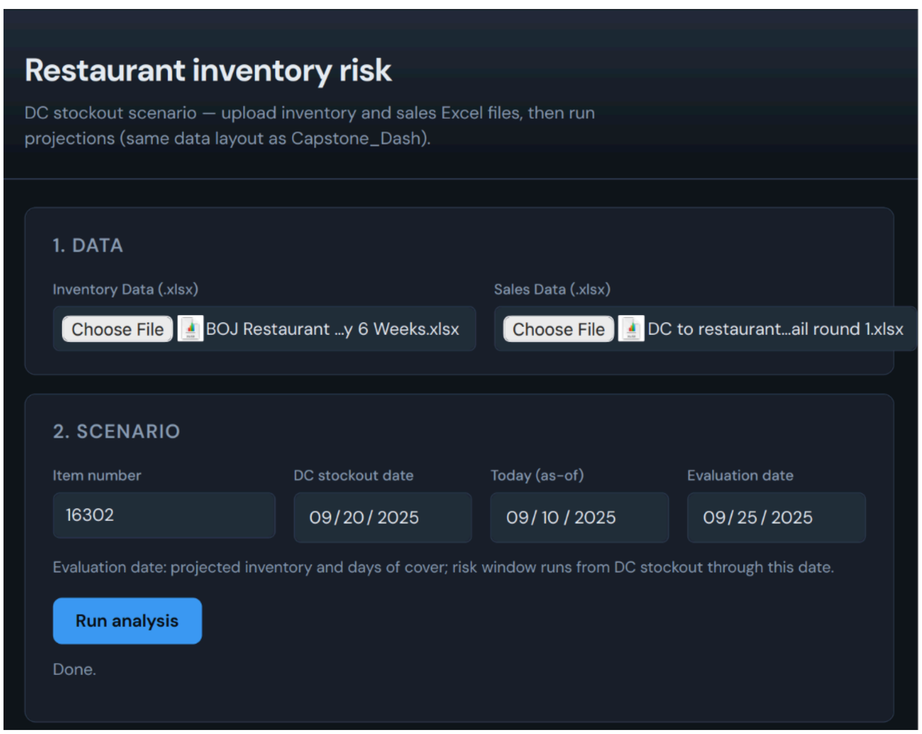

Figure 1: Tool input panel configured for Item 16302 with DC stockout date of September 20, 2025 and evaluation date (i.e. stockout end date) of September 25, 2025.

Once the analysis is complete, the tool classifies all modeled restaurants into risk tiers based on their projected days of inventory cover on the DC stockout end date, as shown in Figure 2. Of the 129 restaurants included in the projection model, 33 are classified as Most Critical, 9 as High Risk, 12 as Action Needed, and 75 as Monitor.

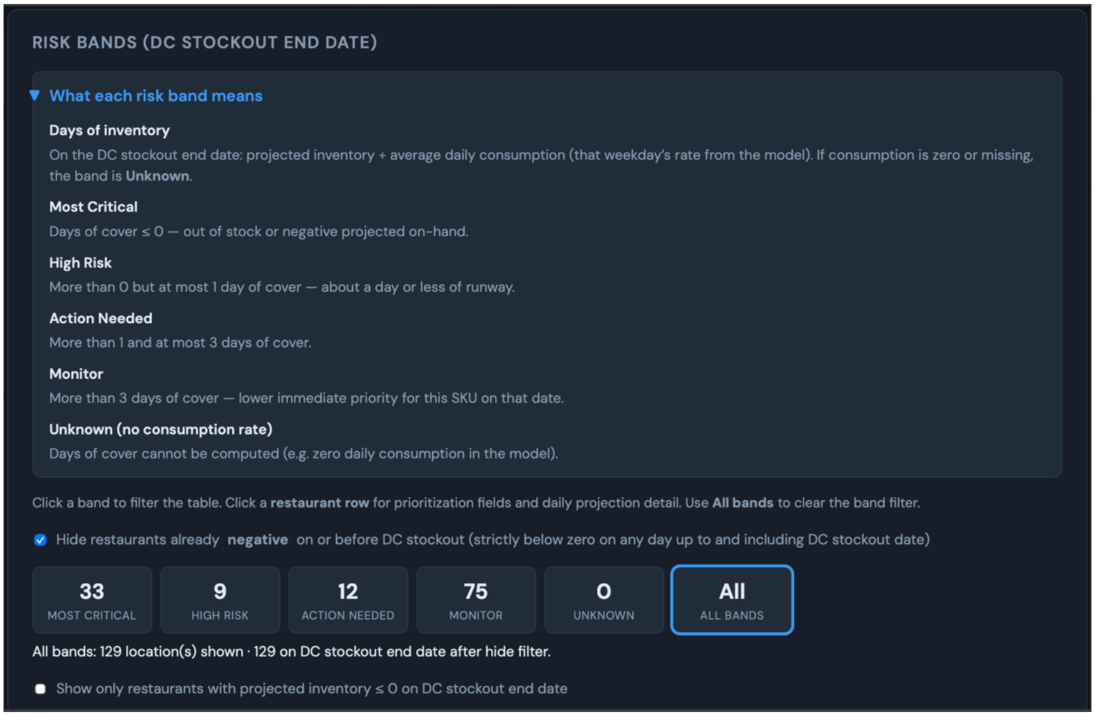

Figure 2: Risk band summary and prioritization table for Item 16302.

To illustrate the example in more detail, consider two restaurants. As shown in Figure 3, Restaurant 38 is classified as "Most Critical" with a first stockout date of September 21, while Restaurant 993 is classified as "Monitor" with no projected stockout in the window.

RESTAURANT NUMBER	PROJECTED INVENTORY	DAILY CONSUMPTION	DAYS OF INVENTORY	RISK BAND	FIRST DAY ≤0 IN WINDOW	SHIPMENT LINES (ITEM)
38	-189.4643	41.9286	-4.5187	Most Critical	2025-09-21	13
993	167.3532	21.0714	7.9422	Monitor	—	11

Figure 3: Prioritization table for Item 16302; it shows that Restaurant 38 is classified as "Most Critical" with a projected stockout on September 21, 2025, and Restaurant 993 is classified as "Monitor" with no projected stockout in the stockout window.

Restaurant 993 starts with 335 cases of item 16302 on September 17 and consumes about 21 cases per day. The model assumes that regular scheduled deliveries continue until September 20, after which no new supply arrives. Even without additional shipments, inventory declines relatively slowly and remains positive throughout the evaluation period: there are still 167 cases on hand as of September 25, which is more than a week of buffer past the DC stockout date. With 11 shipment records in the sales data, the consumption and replenishment estimates for this restaurant are based on a reliable order history. For this restaurant, expedited freight is unnecessary. Figure 4 shows the day-by-day inventory projection for Restaurant 993, with rows through September 20 displayed.

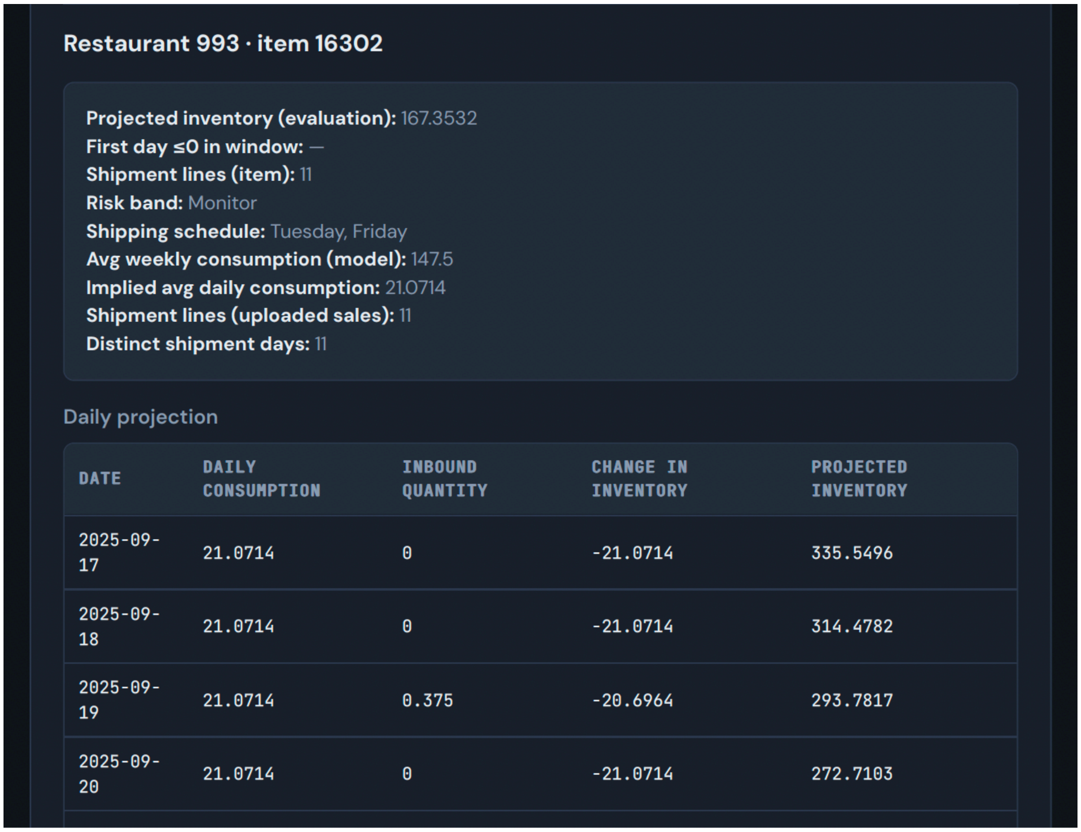

Restaurant 993 · item 16302

Projected inventory (evaluation): 167.3532
First day ≤0 in window: —
Shipment lines (item): 11
Risk band: Monitor
Shipping schedule: Tuesday, Friday
Avg weekly consumption (model): 147.5
Implied avg daily consumption: 21.0714
Shipment lines (uploaded sales): 11
Distinct shipment days: 11

Daily projection

DATE	DAILY CONSUMPTION	INBOUND QUANTITY	CHANGE IN INVENTORY	PROJECTED INVENTORY
2025-09-17	21.0714	0	-21.0714	335.5496
2025-09-18	21.0714	0	-21.0714	314.4782
2025-09-19	21.0714	0.375	-20.6964	293.7817
2025-09-20	21.0714	0	-21.0714	272.7103

Figure 4: Day-by-day inventory projection for Restaurant 993, Item 16302; it displays projected on-hand inventory remaining positive through the evaluation date of September 25, 2025 (rows through September 20 shown)

As demonstrated in Figure 5, Restaurant 38 has a consumption rate of approximately 42 cases per day, which is higher than Restaurant 993. The model shows that Restaurant 38 is projected to stock out on September 21 (only one day after the DC stockout date). By the evaluation date its projected inventory position reaches negative 189 cases. With 13 shipment records available for the restaurant, the user can have a solid level of confidence that the estimated consumption rate reflects its actual ordering behavior.

Restaurant 38 · item 16302

Projected inventory (DC stockout end date): -189.4643
First day ≤0 in window: 2025-09-21
Shipment lines (item): 13
Risk band: Most Critical
Shipping schedule: Monday, Tuesday, Friday
Avg weekly consumption (model): 293.5
Implied avg daily consumption: 41.9286
Shipment lines (uploaded sales): 13
Distinct shipment days: 13

Daily projection

DATE	DAILY CONSUMPTION	INBOUND QUANTITY	CHANGE IN INVENTORY	PROJECTED INVENTORY
2025-09-13	41.9286	0	-41.9286	311.1786
2025-09-14	41.9286	0	-41.9286	269.25
2025-09-15	41.9286	0	-41.9286	227.3214
2025-09-16	41.9286	0.75	-41.1786	186.1429
2025-09-17	41.9286	0	-41.9286	144.2143
2025-09-18	41.9286	0	-41.9286	102.2857
2025-09-19	41.9286	1.75	-40.1786	62.1071
2025-09-20	41.9286	0	-41.9286	20.1786
2025-09-21	41.9286	0	-41.9286	-21.75
2025-09-22	41.9286	0	-41.9286	-63.6786
2025-09-23	41.9286	0	-41.9286	-105.6071

Figure 5: Day-by-day inventory projection for Restaurant 38, Item 16302. Projected on-hand inventory turns negative on September 21, 2025, one day after the DC stockout date (rows through September 23 shown).

Without the tool, both restaurants might be treated as equally urgent. However, based on the output of the tool, Restaurant 993 can use its existing buffer stock and wait until the next scheduled replenishment or be served through lower-cost recovery options, while Restaurant 38 requires immediate intervention. The tool, therefore, enables cost-efficient, targeted prioritization.

4.3 Limitations

The prototype has the following limitations related to implementation:

- As noted in section 4.2, the tool is not integrated with Armada's automated ERP databases, and therefore the supply planner must execute the tool manually. The user is required to load updated data files and specify all required parameters.
- The projection module uses a sequential method to process restaurants for a selected SKU as opposed to being optimized for large-scale parallel computation. The processing time of the tool may be higher if larger datasets with more historical data are supplied. Longer forecast horizons would also affect the tool runtime.
- The forecast horizon is defined within the prototype and does not adjust automatically based on restaurant consumption patterns or ordering frequency.
- The input validation performed by the tool is limited: it assumes that all the data files provided are properly formatted and pre-cleaned.

These constraints are expected at the prototype stage and may be resolved according to the future work recommendations in Section 5.2.

5. RECOMMENDATIONS

This study developed a methodology for estimating restaurant-level stockout risk in the event of a DC-level supply disruption. The methodology was also implemented as a working prototype for Armada's planning team. Our tool is designed to be applied to real brand data used by Armada; it differentiates restaurants by risk level and gives planners a clear basis for deploying targeted recovery actions for every disruption instead of responding uniformly. The following recommendations outline how Armada can take advantage of these findings and build toward a more powerful version of this tool over time.

5.1 Management Recommendations

We recommend that Armada pilots the tool with one or two brands before scaling it across the full network. A pilot program would allow the team to validate risk classifications against planner judgment and assessment in real disruption scenarios: over time, this may enable the users to adjust the risk band thresholds based on operational experience. Brands with reliable, clean, and consistent inventory and shipment data and a history of DC-level disruptions would be strong candidates for the pilot program. Once confidence in the outputs is built, the tool may be used to make recovery decisions at scale.

The financial opportunity is meaningful because Armada's current uniform response to DC stockouts carries significant cost. The tool directly reduces unnecessary emergency freight spend by identifying which restaurants have enough on-hand inventory to absorb a disruption and routing them through standard shipping instead. Even small reductions in expedited shipment frequency across numerous restaurant locations can create substantial cost savings. In addition, ensuring that truly at-risk restaurants are identified promptly and receive immediate intervention can help the company protect service levels and customer experience. This supports Armada's value proposition to its brand clients and, ultimately, their ability to retain customers.

During the pilot stage, Armada is recommended to track and record recovery actions taken in response to the tool's recommendations, as well as their outcomes. This will generate labeled historical data which may be then used to validate and improve the model over time. In addition, such historical data will allow Armada to quantify cost savings from avoided expedited shipments. Sharing this information with brand clients may serve as a demonstration of Armada's analytical capabilities.

5.2 Future Work

There are multiple areas of future development for this tool. First, connecting the prototype to live DC ERP systems would eliminate the need for manual file uploads. This would allow the tool to be used in real time as disruptions occur. Second, once the company accumulates a sufficient record of labeled stockout events, the model should be validated against historical stockout outcomes. That way, risk band thresholds can be adjusted empirically. That same record would function as labeled historical data and unlock more advanced predictive approaches, such as the machine learning classification methods described in Section 2. Third, the tool currently analyzes one SKU at a time, but DC-level disruptions often affect multiple items simultaneously. Extending the tool to support multi-SKU processing would allow planners to assess network-wide risk across all affected items in a single run of the analysis. Fourth, as the tool is scaled to larger datasets, the projection module should be optimized for parallel computation to keep runtime manageable across the full restaurant network.

The accuracy of the model's estimates could also be improved through better input data. Currently, restaurant inventory snapshots are only available at a weekly frequency and are not shared with Armada until the following Monday. Access to daily inventory data reported in real time would enable the model to start projections from a much more current point in time. In addition, daily consumption is currently assumed to be evenly distributed across the week because daily inventory or point-of-sale data is not available at the restaurant level. With access to daily usage data, the model could capture day-of-week consumption patterns, which would improve both the inventory projection and the stockout probability estimate. In addition, incorporating forward-looking PO data would help estimate future replenishment more accurately, particularly for restaurants with limited order histories. In the longer term, as Armada builds richer historical data infrastructure, more sophisticated methods may become viable, including machine learning classification and probabilistic inventory simulation, as described in Section 2. Building this data infrastructure should be a strategic priority.

This work is relevant not only for Armada, but also for other supply chains. The challenge of distinguishing critical from non-critical disruptions is common in foodservice and other perishable supply chains. The framework developed in this study offers an approach for companies that face similar data constraints without requiring real-time inventory visibility or extensive historical records.

REFERENCES

Abad, A. G. (2025, November 18). When your model missed the memo: Combining time-series + LLM signals. *LinkedIn*. https://www.linkedin.com/pulse/when-your-model-missed-memo-combining-time-series-llm-abad-ph-d--mjqke

Anderson, E. T., Fitzsimons, G. J., & Simester, D. (2006). Measuring and mitigating the costs of stockouts. *Management Science.* DOI: 10.1287/mnsc.1060.0577

Armada. (2025). *Company website.* Retrieved September 30, 2025, from https://www.armada.net/

Corsten, D., & Gruen, T. (2003). Desperately seeking shelf availability: An examination of retail out-of-stocks. *International Journal of Retail & Distribution Management.* DOI: 10.1108/09590550310507731

Gruen, T. W., Corsten, D. S., & Bharadwaj, S. (2002). *Retail out-of-stocks: A worldwide examination of extent, causes and consumer responses.* Grocery Manufacturers of America.

Holt, C. C. (1957). Forecasting seasonals and trends by exponentially weighted averages (ONR Memorandum No. 52). *Carnegie Institute of Technology, Pittsburgh USA.* Reprinted in the *International Journal of Forecasting, 2004.* DOI: 10.1016/j.ijforecast.2003.09.015

Huber, J., & Stuckenschmidt, H. (2020). Daily retail demand forecasting using machine learning with emphasis on calendric special days. *International Journal of Forecasting, 36*(4). DOI: 10.1016/j.ijforecast.2020.02.005

Hyndman, R.J., & Athanasopoulos, G. (2021). *Forecasting: principles and practice, 3rd edition*, OTexts: Melbourne, Australia*.* Retrieved December 6, 2025, from OTexts.com/fpp3.

Islam, M. T., Ayon, E. H., Ghosh, B. P., Chowdhury, S., Shahid, R., Rahman, S., Bhuiyan, M. S., & Nguyen, T. N. (2024). Revolutionizing retail: A hybrid machine learning approach for precision demand forecasting. *Journal of Computer Science and Technology Studies.* DOI: 10.32996/jcsts.2024.6.1.4

Lindfors, A. (2021). Demand forecasting in retail: A comparison of time series analysis and machine learning models (Master's thesis). *Åbo Akademi University.* https://www.doria.fi/handle/10024/181510

Liu, Y., Kalaitzi, D., Wang, M., & Papanagnou, C. (2025). A machine learning approach to inventory stockout prediction. *Journal of Digital Economy.* DOI: 10.1016/j.jdec.2025.06.002

Silver, E. A., Pyke, D. F., & Peterson, R. (1998). Inventory management and production planning and scheduling (3rd ed.). *John Wiley & Sons.*

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management Science, 6(3).* DOI: 10.1287/mnsc.6.3.324