Optimization of Fab Spare Parts Inventory

by

Bratati Das

and

Junyeon Jeon

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

Analog Devices, Inc. (ADI) manages spare parts across four front-end fabrication sites through largely decentralized site-level practices, a strategy that creates limited network-wide visibility and inconsistent inventory logic. To address these challenges, this capstone develops a cluster-based segmentation framework to improve how spare parts are classified and managed across sites. The framework begins by considering six initial indicators—lead time, supplier pool, demand pattern, demand growth, criticality, and price—which are summarized into three composite indicators capturing supply risk, demand risk, and price exposure. It then segments spare parts into six risk-based clusters and applies Data Envelopment Analysis (DEA)-based scoring to prioritize clusters and guide differentiated inventory policies.

The results show that different clusters are driven by different underlying risk factors, supporting the need for differentiated inventory policies rather than a one-size-fits-all approach. For example, some clusters are primarily driven by supply exposure and cost, while others are driven by demand variability. A volume-based ABC validation further shows that high-volume items are distributed across multiple risk profiles rather than concentrated in low-risk segments. These findings suggest that aligning inventory policies with the primary drivers of risk can improve consistency in decision-making and enable more scalable inventory management across sites.

Capstone Advisor: Sreedevi Rajagopalan
Title: Research Scientist
Capstone Co-Advisor: Camilo Andrés Mora Quiñones
Title: Postdoctoral Associate

TABLE OF CONTENTS

1 Introduction

Across the semiconductor industry, spare parts inventory plays a critical role in balancing fabrication uptime, working capital, and supply resilience. This challenge is particularly acute in multi-site environments, where decentralized planning practices, differences in data across sites, and long lead times for critical items can create excess stock in some locations while increasing shortage risk in others. As semiconductor manufacturers expand capacity and face supply uncertainty, inventory harmonization has become both an operational and a strategic priority.

Analog Devices, Inc. (ADI) provides a representative setting for this challenge. ADI is a global semiconductor manufacturer headquartered in Wilmington, Massachusetts, employing approximately 24,000 people worldwide and generating over $11 billion in annual revenue. The company operates four internal front-end fabrication facilities located in Wilmington, Massachusetts; Limerick, Ireland; Camas, Washington; and Beaverton, Oregon, which together form a multi-site manufacturing network requiring coordinated materials and spares management. In the semiconductor industry, which is characterized by cyclicality and long lead times for critical materials, optimizing fab inventory is essential to prevent production interruptions.

1.1 Motivation

ADI's Global Procurement and Front-End Materials Management organization faces significant challenges with duplicated inventory, inconsistent practices, and limited visibility across its four front-end fabrication sites. These issues increase the risk of manufacturing downtime and drive up inventory holding costs, even as the company seeks to support continued growth in capacity.

These challenges are rooted in a decentralized inventory management approach. The four fabs were developed through a mix of organic expansion and acquisitions, leading to fragmented systems, diverse data structures, and varying operational practices. Historically, each site has procured and managed spare parts independently, resulting in inconsistencies in parts naming, procurement processes, and stocking policies. This lack of standardization has made it difficult to gain a consolidated view of inventory and to respond quickly and consistently to supply chain disruptions.

This project addresses these challenges by developing a harmonized methodology that segments spare parts and enables more consistent and scalable inventory decisions across sites.

1.2 Problem Statement and Key Questions

The central problem addressed in this project is *how to optimize ADI's spare-parts inventory across its four fabrication sites while reducing cost and ensuring continuous fab operations*. Because the current site-level management model limits visibility, inter-fab coordination, and scalability, this project develops a harmonized methodology for that begins by considering six initial indicators—lead time, supplier pool, demand pattern, demand growth, criticality, and price. These are summarized into three composite indicators—a Supply Risk Key Performance Indicator (Supply Risk KPI), a Demand Risk Key Performance Indicator (Demand Risk KPI), and Price—which are then used to segmenting spare parts into interpretable risk clusters, evaluating cluster priority through Data Envelopment Analysis (DEA)-based scoring, and linking the results to inventory policies.

The key questions guiding the project are:

1. How should ADI standardize parts data and inventory management processes across four fabs to create a consistent and accurate view of materials?
2. What analytical methods and parameters can be used to determine stocking levels that reduce inventory cost while maintaining reliable fab operations?
3. What inventory management structure best supports consistent and coordinated decision-making across the four fabs?

1.3 Project Goals and Expected Outcomes

The goal of this project is to design an inventory optimization framework that enables ADI to manage spare parts across its front-end fabs through improved inventory visibility, cost efficiency, and operational scalability across the fabs.

Expected outcomes:

- A unified and standardized dataset enhancing clear visibility across sites
- Improved ability to prevent production disruptions by ensuring the timely availability of critical spare parts
- A decision framework that links cluster-level inventory scores to recommendations on where inventory should be held across sites for different clusters, supporting more consistent and coordinated inventory decisions.

This project provides a structured approach for addressing ADI’s key inventory challenges, including excess inventory, supply risk, and limited cross-site coordination. By segmenting spare parts based on supply risk, demand risk, and price exposure, the proposed framework enables more consistent and targeted inventory decisions across sites. This approach supports reducing inventory cost while maintaining reliable fab operations and improving coordination across fabrication facilities. The framework can also be applied to similar multi-site manufacturing environments, providing a scalable basis for improving inventory decision-making beyond ADI.

2 State of the Practice

To design our methodology, we reviewed the literature on multi-criteria inventory classification, clustering and segmentation of spare parts, data envelopment analysis and optimization-bsed inventory scoring, and inventory policies for intermittent and slow-moving demand.

2.1 Multi-Criteria Inventory Classification

Traditional ABC classification groups inventory items based solely on annual dollar value, with a small set of “A” items accounting for most of the value, B items representing medium-value items, and a long tail of “C” items. This simple rule is popular because it is easy to implement. Still, it ignores important drivers of supply risk, such as lead time or criticality. It tends to classify all low-volume items as class C, even when some are indispensable for operations. These shortcomings of traditional ABC classification have been widely discussed in the literature (e.g., Ramanathan, 2006; Teunter et al., 2010). These studies show that when service levels and inventory costs are considered jointly, a richer classification scheme can support differentiated service levels and lower total cost. They propose an ABC policy in which items are ranked by a composite index that combines demand rate, holding cost, order quantity, and shortage cost, and classes are then assigned target service levels. This approach demonstrates the benefit of multi-criteria thinking even within a simple three-class structure.

Multi-criteria ABC methods extend this idea by incorporating several quantitative and qualitative attributes into a single priority score. Ramanathan (2006) argues that relying on a single criterion, such as annual demand value, can lead to suboptimal prioritization, particularly for spare parts whose stock-out cost may far exceed their purchase price. To address this, the study formulates ABC classification as a weighted linear optimization problem. Each item is evaluated on multiple criteria, such as annual dollar spend, criticality, lead time, and unit cost, and the model chooses non-negative weights that maximize each item's score, subject to benchmarking constraints. As a result, items with strong performance on any combination of criteria can be identified without requiring the planner to specify subjective weights in advance.

The literature also emphasizes that criteria like criticality and lead time are particularly important for maintenance and spare parts (e.g., Ramanathan, 2006; Teunter et al., 2010). Items with long replenishment times or high impact on fab uptime may deserve higher service levels even when their usage is small. Multi-criteria classification frameworks therefore often include variables such as demand variability, lead time uncertainty, supplier breadth, equipment criticality, or substitution options alongside monetary value. This body of work motivates our use of six initial indicators—lead time, supplier pool, demand pattern, demand growth, criticality, and price—which are later summarized into composite supply-risk and demand-risk indicators plus a standalone price indicator for ADI's segmentation.

2.2 Clustering and Segmentation for Spare Parts

Moving from item-by-item analysis to cluster-level decisions can substantially reduce complexity. Cohen and Ernst (1988) introduce the concept of "generic" inventory policies, where items with similar demand and lead-time characteristics are grouped into operations-related groups (ORGs). Each ORG shares a common set of control parameters, such as reorder points and order quantities, and the authors show that these generic policies can approximate item-specific optimization with much lower implementation effort. Ernst and Cohen (1990) further develop a clustering procedure that assigns items to ORGs based on their demand and supply attributes, demonstrating that the resulting policies maintain service levels while simplifying policy management.

Recent applied projects also rely heavily on segmentation. Chawla and Miceli (2019) develop a classification-optimization tool for a spare-parts distributor that assigns each SKU to a class based on average unit cost, annual revenue contribution, lead time, out-of-stock cost, and strategic importance, among other inputs. The planner specifies the desired number of groups and cut-off percentages, and the tool outputs a class number for each SKU that can be linked to inventory and service policies. Similarly, Cameron and Efendigil (2021) study slow-moving, high-volatility items for a fastener distributor and segment SKUs into three classes (A, B, and C) based on annual order frequency. They then assign different forecasting and replenishment rules to each class, thereby creating a playbook that can be applied consistently at scale.

These studies illustrate that the main goal of segmentation is not merely statistical clustering; it is to create groups of items that can share inventory policies and be managed with a common logic. For ADI, grouping spare parts into a modest number of clusters with similar composite supply-risk, composite demand-risk, and price profiles will allow us to propose tailored yet operationally manageable stocking guidelines for each cluster rather than for thousands of individual SKUs.

2.3 DEA and Optimization-Based Inventory Scoring

Data Envelopment Analysis (DEA) has been widely used in operations research to evaluate the relative efficiency of decision-making units based on multiple inputs and outputs. Ramanathan (2006) adapts this idea to inventory classification by using multiple criteria as fixed inputs and optimizing for the weights assigned to each. The model maximizes the weighted sum of inputs for each item while constraining the weighted sum for all items, to be less than or equal to one. This formulation yields an "ABC score" that reflects the optimal combination of criteria for each item, without requiring explicit trade-off parameters.

Optimization-based scoring methods are attractive because they provide an internal way to derive weights for criteria such as lead time or criticality (Ramanathan, 2006). In contrast, many multi-criteria ABC schemes require managers to pre-assign weights or points to each criterion, which can be time-consuming and inconsistent across plants. DEA-type models naturally handle criteria measured in different units and scale well when new dimensions or items are added. For our capstone, we adopt the spirit of this approach but apply it at the level of clusters rather than individual SKUs. This allows us to compute a single set of weights across the three composite indicators derived from the six initial indicators and then propagate the resulting inventory score to the SKUs belonging to each cluster.

2.4 Inventory Management for Intermittent and Slow-Moving Demand

Spare parts often exhibit intermittent demand: long periods of zero usage followed by occasional orders of varying size. Managing this kind of demand is challenging because traditional forecasting and replenishment methods designed for fast-moving items tend to underestimate the required safety stock or generate unstable policy parameters (Croston, 1972). Croston's method addresses this issue by separating the estimation of demand size from the estimation of the inter-arrival time between non-zero demands, updating each component independently. This structure improves forecast stability and bias properties for intermittent demand.

Cameron and Efendigil (2021) tackle this problem by developing an inventory policy playbook for low-volume, high-volatility items. They classify SKUs by annual order frequency and apply Croston's forecasting method and a hybrid periodic-review (R,s,S) policy for the most frequently ordered class, while using simpler, maximum-based policies for items with very low transaction counts. Their results suggest that this tailored approach can cut inventory cost by roughly half while maintaining a 99% service level.

Kaya et al. (2024) studied a manufacturer facing intermittent demand for a wide range of electrical components. They design an inventory management system that combines appropriate forecasting techniques for intermittent demand, such as Croston's method and the SBA variant, with an integrated Newsvendor and order-up-to policy. The authors implement their solution through a Python algorithm and an Excel-based decision-support system and report that aligning forecasting and inventory decisions in this way reduces holding and backorder costs while improving customer service.

These contributions highlight two themes that are relevant for ADI. First, segmentation is an effective way to manage slow-moving and intermittent-demand items; groups of SKUs with similar demand frequency or volatility can share an appropriate policy. Second, data-driven models that integrate multiple criteria and link forecasting with inventory decisions can unlock substantial savings. Our project concentrates on the segmentation and prioritization step by beginning with six initial indicators—lead time, supplier pool, demand pattern, demand growth, criticality, and price—summarizing them into composite supply-risk and demand-risk indicators plus a standalone price dimension, and then computing an inventory score for each cluster. In later phases of the capstone, these scores can be linked to specific safety-stock or replenishment rules, drawing on policy structures proposed in the slow-moving demand literature.

2.5 Relevance to the Capstone Problem

In summary, the literature suggests that effective inventory management for spare parts and fab materials should (1) use multiple indicators beyond annual dollar usage, (2) summarize those indicators into a manageable number of composite dimensions so that policy design remains operationally feasible, and (3) derive transparent, data-driven weights for those dimensions. Classical ABC analysis provides a useful starting point, but multi-criteria methods and DEA-based scoring offer richer insights into the relative importance of items. Segmentation approaches demonstrate how clusters can be used as the primary unit of decision-making. Finally, research on intermittent demand demonstrates that such segmentation is particularly valuable for slow-moving, high-volatility items.

3 Methodology

This section describes the methodology we propose to translate raw purchase-order data into a small set of inventory clusters, each with an interpretable profile and a quantitative inventory score derived from composite supply-risk, demand-risk, and price dimensions. Specifically, we transform ADI's three-year order history into six initial indicators, three composite indicators, interpretable risk clusters, DEA-based inventory scores, and policy recommendations. These cluster scores provide a foundation for designing differentiated safety-stock and replenishment policies in future work.

3.1 Methodology Selection

Considering the literature, ADI's data availability, and implementation constraints, we propose a two-stage approach. First, to transform item-level operational variables into three composite indicators that summarize supply risk, demand risk, and price. Second, to employ a K-nearest neighbors (KNN)-based clustering approach to group spare parts into a small number of interpretable segments and apply a DEA-style scoring model to rank those clusters and translate them into differentiated policy implications. The overall framework and its key components are illustrated in Figure1.

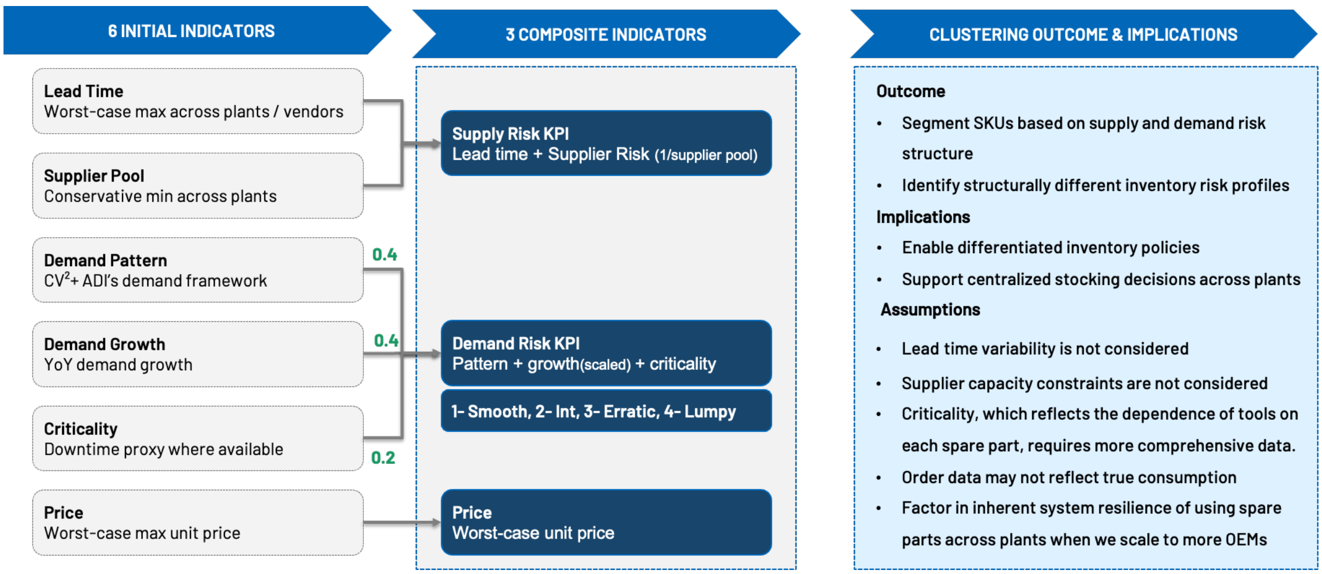

Figure 1. SKU segmentation framework

The Integrated Risk-Based Inventory Segmentation (IRIS) Framework begins with six initial indicators that are directly observable or derivable from ADI-provided data: lead time, supplier pool, demand pattern, demand growth, criticality, and price. These are then summarized into three composite indicators. Supply Risk KPI combines lead time and supplier risk, where supplier risk is represented as the inverse of supplier pool size. Demand Risk KPI combines demand pattern, scaled demand growth, and criticality, with working weights of 0.4, 0.4, and 0.2, respectively, in the current slide framework. Price remains a standalone indicator because of its direct relevance to capital tied up in inventory and to centralization decisions.

This structure reflects three considerations. First, ADI's fab inventory consists of thousands of spare parts with heterogeneous supply and demand behaviors, so grouping items into a limited number of segments is operationally more realistic than setting a fully optimized policy for each SKU. Second, the six-indicator structure captures a broad range of operational variation while remaining interpretable through composite risk dimensions. Third, the ADI managers needs a method that can be refreshed as new order history becomes available and that can classify new items consistently. A K-nearest neighbors (KNN)-based clustering approach supports this requirement by assigning items to clusters based on similarity in a way that planners can understand and explain.

3.2 Data Collection and Preparation

We use ADI's "Past 3 Yr Order History" dataset, which contains purchase order line items for spare parts over the last three years. Each row represents a purchase order (PO) line and includes, among other fields, vendor alias, vendor country, plant, purchasing document and line item, PO issue date, material number and description, material group, scheduled quantity, quantity delivered, lead time in days, existing minimum and maximum levels, lot size parameters, unit price, order unit, and net order value in USD.

From this transactional data, we built an aggregated item level dataset and, where relevant, material–supplier combinations. For each item, we derived the six initial indicators used in the latest framework:

- Lead time: worst-case or representative supplier-to-fab procurement lead time, using conservative maximum values across relevant plant/vendor combinations.
- Supplier pool: conservative minimum count of feasible suppliers across plants; supplier risk is later represented as (1 / supplier pool).
- Demand pattern: categorical pattern classification based on variability metrics, coefficient of variation squared (CV^2), and average demand frequency, mapped to four categories of demand: Smooth (low variability, high frequency), Intermittent (low frequency, low variability), Erratic (high variability, high frequency), and Lumpy (high variability, low frequency).
- Demand growth: year-over-year delivered-quantity weighted average growth, scaled for comparability across items. Growth is calculated year-over-year, with first-time purchases set to zero growth, revival purchases set to 100% growth, and greater weight given to more recent years.
- Criticality: proxy for the importance of a spare part, measured by the risk of production disruption if the part is unavailable.
- Price: worst-case maximum unit price or representative high-end unit-cost exposure.

We also computed supporting measures such as annualized quantity and order frequency for validation and interpretation, including the later volume-based ABC check. Before indicator construction, we removed cancelled or zero-quantity transactions, filtered noninventory adjustments, reconciled units of measure, and reviewed outliers with ADI planners. The six initial indicators were then normalized or scaled as needed so that composite supply- and demand-risk calculations were comparable across items.

3.3 Model Application

The IRIS framework consists of five main steps, as illustrated in Figure 2. First, we clean purchase-order history and transform it into SKU-level attributes. Second, we derive the six initial indicators and summarize them into three composite indicators: Supply Risk KPI, Demand Risk KPI, and Price. Third, we apply a K-nearest neighbors (KNN)-based clustering approach in the composite-indicator space to generate interpretable clusters. Fourth, we score the resulting clusters through a DEA-style linear program. Fifth, we map the cluster profiles and scores to candidate inventory policies and alternative stocking structures.

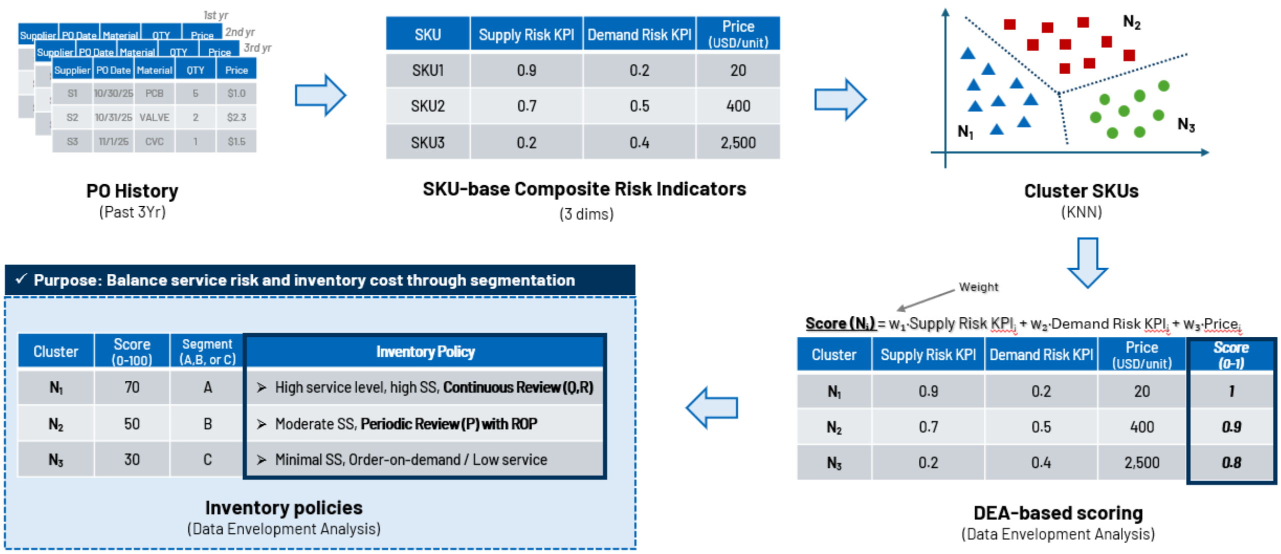

Figure 2. A DEA-Based Framework for Risk Segmentation and Inventory Policy Design

In the first step, the raw purchase-order data are transformed into an analysis-ready dataset. Each row represents a material–plant or material–supplier observation, and the relevant fields are used to compute lead time, supplier pool, demand pattern, demand growth, criticality, and price. These six initial indicators are then standardized or scaled as required to support consistent comparison across items. Demand pattern is ordinally scaled using Smooth -1, Intermittent -2, Erratic -3, and Lumpy -4. Lead time, price, and demand growth are min-max scaled. Supplier pool and criticality have very few variations in data, so they are not scaled.

In the second step, we apply a K-nearest neighbors (KNN)-based clustering approach to the composite-indicator dataset. Candidate solutions are compared using clustering cost, stability of cluster memberships, and business interpretability for ADI's planners. For a given number of clusters, the algorithm assigns each item to the cluster that is most common among its nearest neighbors in the composite-indicator space. The result is a set of clusters that differ meaningfully in supply risk, demand risk, and price exposure.

In the third step, we apply a DEA-style optimization at the cluster level. For each cluster c, we compute representative values of the three composite indicators, denoted y_c^S for Supply Risk KPI, y_c^D for Demand Risk KPI, and y_c^P for Price. Following Ramanathan (2006), we treat each cluster as a decision-making unit with a single dummy input equal to 1 and three outputs corresponding to these composite indicators. The linear program selects non-negative weights that maximize the weighted sum of outputs for the focal cluster, subject to benchmarking constraints on all other clusters. Solving this model produces a score between 0 and 1 for each cluster and reveals whether its relative policy priority is driven primarily by supply risk, demand risk, or price.

Because the optimization is run at the cluster level, the resulting weights and scores capture ADI's implicit preferences across groups of items rather than individual SKUs. This reduces sensitivity to noise in the data for very low-volume items and aligns with our objective of designing cluster-level policies. In the final step, we interpret cluster profiles, DEA scores, and validation checks together with ADI to develop a policy map. Because ADI currently operates a hybrid min–max logic through weekly Material Requirements Planning (MRP), the policy implications are interpreted in terms of minimum levels, order-up-to levels, centralization candidates, and review intensity rather than as purely item-by-item optimal formulas. This also allows the same framework to be refreshed as new order history becomes available.

4 Results and Discussion

This section presents the results of the clustering and scoring analysis, and discusses the implications for inventory policy design. Specifically, we analyze the structure of the resulting clusters, evaluate their relative priority using DEA-based scores, and interpret how these results translate into differentiated inventory policies across clusters.

4.1 Cluster Structure and Interpretation

The proposed methodology and the given data yield six interpretable clusters. Figure 3 plots the clusters on a supply-risk versus demand-risk plane, where the horizontal axis represents the Supply Risk KPI and the vertical axis represents the Demand Risk KPI, both normalized to enable comparison across items. Bubble size represents unit price.

The chart is directionally consistent with expectations: clusters separate high-demand-risk items from high-supply-risk items, while also highlighting a small number of very high-price items.

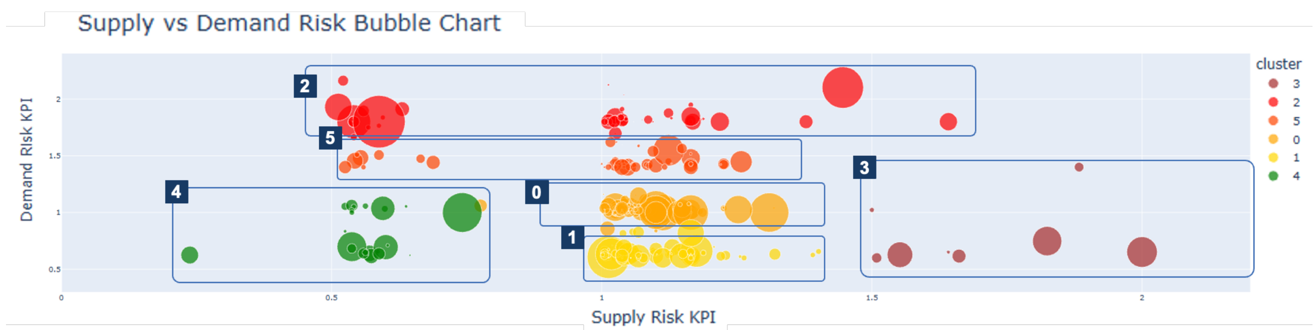

Figure 3. Supply-Risk versus Demand-Risk Cluster Solution

Cluster 2 sits high on the demand-risk axis with moderate supply risk and corresponds to a lumpy-demand profile. Cluster 5 also shows high demand risk but with a different pattern structure, consistent with erratic items. Cluster 3 is the clearest supply-risk outlier; it lies farthest to the right and contains the most price-intensive items in the current solution. Clusters 0 and 1 occupy the middle of the chart and represent intermittent and smooth demand regimes with moderate-to-higher supply risk. Cluster 4 occupies the lower-left region and represents the lowest combined risk profile in the current six-cluster draft.

These results matter because they show that a single inventory rule would mask important differences. High-demand-risk clusters should not be controlled the same way as high-supply-risk clusters, and high-price items should not automatically receive the same stocking logic as lower-value items. The clustering therefore provides the missing bridge between item-level heterogeneity and operationally manageable policy families.

4.2 DEA-Based Cluster Scoring Results

Table 1 summarizes the cluster-level profile values and the resulting DEA scores. Table 2 reports the corresponding DEA weights. Together, these tables show which clusters define the current efficient frontier and which criteria drive their relative rankings.

Cluster	Demand Risk KPI	Supply Risk KPI	Price (scaled)	Total Qty	DEA score
3	0.665	1.621	0.138	214	1.000
2	1.828	0.862	0.020	2,822	1.000
5	1.436	1.058	0.005	18,538	0.971
0	1.038	1.076	0.017	2,299	0.853
1	0.633	1.055	0.010	11,640	0.715
4	0.646	0.580	0.007	4,017	0.487

Table 1. Cluster Profile and DEA Score

Cluster	w_Demand Risk	w_Supply Risk	w_Price
3	0.000	0.000	7.268
2	0.547	0.000	0.000
5	0.318	0.487	0.000
0	0.318	0.487	0.000
1	0.318	0.487	0.000
4	0.318	0.487	0.000

Table 2. DEA Weights by Cluster

Two clusters—Cluster 2 and Cluster 3—receive DEA scores of 1.000 and therefore define the efficient frontier in the updated working solution. Cluster 3's weight vector remains entirely price-driven, with a price weight of 7.268, indicating that its policy priority comes from a very expensive set of items with high supply exposure. Cluster 2 is also efficient, but in the updated results, it is driven primarily by demand risk, with a demand-weight coefficient of 0.547 and zero supply- and price-weight coefficients. Cluster 5 remains close to the frontier with a score of 0.971, while Clusters 0, 1, and 4 share a common weighting structure in which both demand and supply risk contribute positively to policy priority.

By contrast, Cluster 4 receives the lowest score (0.487), which is consistent with its position as the lowest-risk cluster in the bubble chart (Figure 3). Clusters 0 and 1 occupy the middle of the ranking and are driven by both demand and supply risk rather than by price. This pattern reinforces the logic that not all clusters should be managed through a single value lens: for some groups, cost exposure dominates, while for others, demand-pattern and supply-risk effects matter more.

4.3 Volume-Based ABC Validation

To stress-test the six-cluster solution from a quantity perspective, we overlaid a volume-based ABC segmentation on the same supply-risk / demand-risk space. A, B, and C were defined by cumulative delivered quantity with cutoffs of 80%, 15%, and 5%, respectively. Table 3 reports the distribution by demand pattern, and Table 4 reports the distribution by cluster. This additional analysis is intended to validate that the clustering captures meaningful differences beyond volume and to ensure that high-volume items are not concentrated only in specific clusters, which would indicate that the clustering is primarily driven by quantity rather than risk characteristics.

Demand Pattern	A	B	C	Total
Smooth	38	57	81	176
Intermittent	5	25	97	127
Lumpy	7	10	29	46
Erratic	44	41	26	111
Total	**94**	**133**	**233**	**460**

Table 3. Volume-Based ABC Counts by Demand Pattern

Cluster	A	B	C	Total
0	5	22	91	118
1	35	49	69	153
2	7	14	31	52
3	0	3	5	8
4	3	8	14	25
5	44	37	23	104
Total	**94**	**133**	**233**	**460**

Table 4. Volume-Based ABC Counts by Cluster

The main result is that A, B, and C items are spread across the risk space rather than concentrating only in the lowest-risk regions. This is reassuring because it suggests that the clustering is not merely proxying for volume; even high-volume items appear across multiple demand-risk and supply-risk profiles.

The main exception is the erratic demand pattern, which contains 44 A items and 41 B items. In parallel, 79 of the 94 A items fall into Clusters 1 and 5, showing that high-volume items are concentrated in a small number of moderate-to-high risk segments rather than in the leanest cluster. A plausible explanation is that some of these items are ramp-down parts or are tied to decommissioned tools or replacement transitions; three-year cumulative quantity remains high, but recent ordering has become more variable.

4.4 Inventory-Policy Implications by Cluster

We developed a working policy map by combining cluster results, updated DEA scores, and volume-based ABC validation, as summarized in Table 5. These were informed by discussions with ADI's material management team on their current min–max / weekly MRP operating logic.

<table>
<tr><th>Cluster</th><th>Working interpretation</th><th>Recommended control logic</th><th>Suggested review mode</th></tr>
<tr><td>3</td><td>Highest supply risk and highest price intensity</td><td>Centralized stocking to minimize duplicate holding cost</td><td rowspan="2">Continuous review with a reorder point</td></tr>
<tr><td>2</td><td>Lumpy; highest demand risk; centralization candidate</td><td>Centralized stocking with high safety stock</td></tr>
<tr><td>5</td><td>Erratic; high demand risk; moderate supply risk</td><td rowspan="3">Higher safety stock and higher reorder levels</td><td rowspan="3">Periodic review with an order-up-to level</td></tr>
<tr><td>0</td><td>Intermittent; moderate combined risk</td></tr>
<tr><td>1</td><td>Smooth; higher supply exposure</td></tr>
<tr><td>4</td><td>Lowest combined risk</td><td>Keep current logic / lean local replenishment</td><td>Hybrid review with min- max levels</td></tr>
</table>

Table 5. Inventory Policy Mapping by Cluster

The updated policy map aligns with the cluster characteristics and ADI's current operating logic. The implications for each cluster are as follows:

- Cluster 3: Prioritize centralization due to high price and supply risk. Consolidating inventory across sites reduces duplicate holdings and improves cost efficiency.
- Cluster 2: Also a strong candidate for centralization, but with an emphasis on higher safety stock to manage high demand variability.
- Clusters 5, 0, and 1: Require higher safety stock and higher order-up-to levels to cover demand during both supplier lead time and the review period, reflecting greater demand uncertainty and supply exposure.
- Cluster 4: Can be managed with leaner local replenishment due to its relatively low overall risk.

The policy implications follow directly from the updated cluster profiles and ADI's current operating logic. To align the recommendations with ADI's weekly Material Requirements Planning (MRP) process, the near-term recommendation is to use a periodic review order-up-to policy for each cluster. Under this policy, inventory is reviewed at fixed intervals, and replenishment raises the inventory position to an order-up-to level (S). Following Hadley and Whitin (1963), the order-up-to level can be expressed as:

$$S = \mu(T + L) + z\,\sigma\,\sqrt{T + L} \quad (1)$$

where for each SKU:

S = Order-up-to level

μ = Mean demand

T = Review period

L = Lead Time

σ = Standard deviation of demand

$$z = \phi^{-1}(\alpha) \quad (2)$$

where α = target cycle service level. and

φ^{-1} = inverse cumulative distribution function of the standard normal distribution

Given that our sponsor company, may not be able to calculate exactly their stock-out costs per spare part, due to the complexity across tools, locations, and process dependencies, we recommend specifying service levels rather than explicitly estimating shortage costs. Muckstadt (2005) formulates service-parts inventory problems as optimization models that maximize fill rate or impose service-level constraints, highlighting the practical approach to keeping service levels as exogenous variables, as shown in Table 6 below.

Cluster	DEA score	Service Levels (α)	z-score (Safety Stock)
3	1.000	99%	2.33
2	1.000	99%	2.33
5	0.971	98%	2.05
0	0.853	95%	1.64
1	0.715	90%	1.28
4	0.487	85%	1.04

Table 6. Recommended Service Levels by Cluster

We also recommend keeping a cluster-based review period as opposed to the current weekly MRP run. The review period of each cluster is given by:

$$T = \frac{Q^*}{D}$$

Where D = annual demand (units/year)

Q* = Economic Order Quantity, which is given by,

$$Q^* = \sqrt{\frac{2 * D * Order\ Cost}{Holding\ Cost}}$$

The sponsor has data on ordering cost per purchase order, estimated to be 25 USD. We can estimate the holding cost to be 20% of price, as recommended by sponsor, based on estimates for semiconductor industry. Thus, the review period will vary depending on the cluster's annual demand and the average holding cost, derived from its average price.

This helps us set cluster-differentiated policies on both order-up-to level and review period, factoring in demand risk, supply risk, and unit price of SKUs comprising each cluster.

4.5 Limitations and Considerations

This section outlines key limitations of the analysis that have implications for the interpretation of the results.

First, lead-time variability is not modeled directly. The current framework uses conservative or representative lead-time values to reflect replenishment exposure, but it does not estimate the full distribution of lead times. As a result, the model captures relative differences in lead-time risk across items, but it does not explicitly account for cases in which the same supplier or route may deliver much earlier or much later than expected.

Second, purchase-order history may not fully reflect true consumption. The analysis relies primarily on historical order records, but order quantity does not always equal actual usage quantity at the time of purchase. Some spare parts may be purchased in large batches and consumed gradually over time. In those cases, the data can make a part appear more erratic or lumpy than its actual consumption pattern. This limitation is especially relevant for parts ordered infrequently but in large quantities, because the observed order pattern may reflect procurement behavior, minimum order quantities, or planning decisions rather than actual operational demand.

Beyond the analytical limitations, practical implementation considerations must also be taken into account.

Centralized stocking can reduce duplicate inventory across sites, but it requires operational validation before implementation. Supplier qualification, supplier capacity, inter-plant transfer feasibility, and warehouse-space availability all affect whether a centralized or regional stocking strategy can be executed successfully. These factors do not change the analytical logic of the framework, but they influence how the recommended policy can be translated into practice.

In addition, spare parts used at only one fabrication site may require a different policy logic than parts used across multiple sites. Multi-site parts can benefit more directly from pooling because demand and inventory can be coordinated across fabs. Single-site parts, by contrast, may have fewer pooling opportunities and may require more local stocking logic. Therefore, the recommended framework is most directly applicable to parts with cross-site relevance, while future implementation may require additional rules for single-site items.

Overall, these limitations and considerations do not undermine the value of the segmentation framework. Rather, they define the boundary between the current analytical model and a fully deployable inventory-policy system. The current framework provides a structured way to classify spare parts, identify the main drivers of inventory risk, and connect those risk profiles to differentiated stocking policies.

5 Recommendations

The findings of this capstone indicate that ADI can improve spare-parts inventory performance by moving from decentralized, fab-level decision-making to a unified, cluster-based policy framework. By segmenting spare parts based on supply risk, demand risk, and price exposure, the proposed approach enables more targeted and consistent inventory decisions across sites.

From a business perspective, this framework supports key operational priorities, including reducing excess inventory, improving service reliability, and enhancing coordination across fabrication sites. By aligning inventory policies with the underlying drivers of risk, ADI can better balance working capital efficiency with production continuity, which is critical in semiconductor manufacturing environments.

Beyond operational improvements, the proposed framework also provides a scalable decision-support structure that can be extended as additional data becomes available. This positions ADI to continuously refine its inventory strategy in line with broader supply chain and corporate objectives, including cost optimization, resilience, and efficient capacity utilization.

5.1 Management Recommendations

This section presents the recommended actions for ADI to implement the proposed inventory framework and improve spare-parts management across its fabrication sites.

Adopt cluster-based policy governance. ADI can use the six-cluster structure as the primary decision layer for spare-parts policy design. This will allow planners to manage a small number of interpretable risk profiles instead of applying one rule to all parts or treating every SKU as a special case.

Align policy logic with ADI’s actual weekly MRP process. For near-term implementation, all clusters can be managed through a periodic review policy, with cluster-specific service levels and review parameters. Higher-risk clusters require higher service levels and more conservative order-up-to targets, while lower-risk clusters can be managed with leaner periodic settings.

Use value and scarcity to guide centralization decisions. The current evidence suggests that Clusters 2 and 3 are the strongest candidates for centralized or regional stocking, especially when price intensity and supply exposure are both high. Cluster 4 is the clearest candidate for lean local replenishment.

Institutionalize data standardization and ongoing refresh. A common clustering and policy layer will only be sustainable if ADI standardizes part identifiers, lead-time logic, and criticality fields across fabs and refreshes the clustering as order history evolves.

Although continuous review would be the preferred long-term policy for the highest-risk spare-parts clusters, the recommended near-term implementation is periodic review because it aligns with ADI’s current weekly MRP operating baseline.

5.2 Future Work

An important advantage of the IRIS framework is that it can be extended using ADI’s existing data without requiring new systems or additional data collection. The current three-year purchase-order history already captures the core drivers of supply risk, demand risk, and price exposure, enabling the model to be refreshed periodically as new transactions are recorded. Using the data as-is, future work can focus on recalculating indicators on a rolling basis, updating cluster memberships, and tracking movement of parts across clusters to flag emerging risks or stabilization trends.

The same dataset can also be used to validate and tune cluster-level policies. Future analysis can compare historical stock-out events, expedited orders, inventory turns, and service performance against assigned risk segments. This would allow ADI to test whether high-risk clusters are receiving appropriate service levels and whether low-risk clusters can be managed with leaner inventory settings.

Future work can also refine the policy parameters by incorporating ordering costs, holding costs, and service-level targets more explicitly. These inputs would support more precise selection of service factors and review intervals for each cluster. If ADI later develops system capabilities for automated reorder triggers, continuous-review or EOQ-based policies can be evaluated for selected high-risk spare parts. Until then, periodic review remains the more practical operating baseline.

Finally, the framework can be extended to spare parts used at only one plant and to broader original equipment manufacturer (OEM) part-number coverage. These items may require different pooling logic because they do not always benefit from cross-site inventory sharing. Expanding the framework in this direction would help ADI create a recurring decision-support layer that evolves with the business while maintaining consistency and minimizing implementation overhead.

References

Cameron, K. K., & Efendigil, E. (2021). *Inventory Management for Slow-Moving and High-Volatility Items* (Capstone project). Massachusetts Institute of Technology, Program in Supply Chain Management.

Chawla, G., & Miceli, V. (2019). *Demand Forecasting and Inventory Management for Spare Parts* (Capstone project). Massachusetts Institute of Technology, Program in Supply Chain Management.

Cohen, M. A., & Ernst, R. (1988). Multi-item classification and generic inventory stock control policies. *Production and Inventory Management Journal, 29*(3), 6–8.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, *23*, 289–303. http://www.jstor.org/stable/3007885

Ernst, R., & Cohen, M. A. (1990). Operations related groups (ORGs): A clustering procedure for production/inventory systems. *Journal of Operations Management*, *9*(4), 574–598. https://doi.org/10.1016/0272-6963(90)90010-B

Hadley, G. and Whitin, T. (1963) Analysis of Inventory Systems. Prentice-Hall, Englewood Cliffs. https://www.scirp.org/reference/referencespapers?referenceid=1157422

Kaya, B., Karabağ, O., Çekiç, F. R., Torun, B. C., Başay, A. Ö., Işıklı, Z. E., & Çakır, Ç. (2024). Inventory Management Optimization for Intermittent Demand. In N. M. Durakbasa & M. G. Gençyılmaz (Eds.), *ISPR 2023, Lecture Notes in Mechanical Engineering* (pp. 768–782). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-53991-6_59

Muckstadt, J. A. (2005). *Analysis and Algorithms for Service Parts Supply Chains*. Springer, New York. https://doi.org/10.1007/b138879

Ramanathan, R. (2006). ABC inventory classification with multiple-criteria using weighted linear optimization. *Computers & Operations Research*, *33*(3), 695–700. https://doi.org/10.1016/j.cor.2004.07.014

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains* (4th ed.). CRC Press. https://doi.org/10.1201/9781315374406

Teunter, R. H., Babai, M. Z., & Syntetos, A. A. (2010). ABC Classification: Service Levels and Inventory Costs. *Production and Operations Management*, *19*(3), 343–352. https://doi.org/10.1111/j.1937-5956.2009.01058.x