

Coverage Adequacy as an Alternative to Demand Forecasting for Aftermarket Spare Parts: The DNEIF Framework

Arjun Pardasani

Independent Researcher
USA

arjunpardasani@gmail.com

Abstract

Aftermarket spare parts inventories in durable goods manufacturing are systematically oversized because planners overstock defensively against uncertain, intermittent demand. Existing remedies either forecast demand at the SKU level, where accuracy is structurally poor, or segment the catalog into policy buckets (ABC-XYZ) without quantifying unit-level excess. This paper proposes a different framing. The Demand-Normalized Excess Identification Framework (DNEIF) treats aftermarket inventory optimization as a coverage adequacy test: rather than predicting what demand will be, it determines whether on-hand stock exceeds what any plausible demand trajectory could consume over a defined service horizon. The framework proceeds in three stages: normalize 36 months of consumption history via IQR filtering, project a 10-year demand ceiling using category-specific geometric decay, and classify each on-hand unit as right-sized, excess, or zero-movement. A simulation calibrated to published industry parameters (11,200 SKUs, \$14.8M in inventory, 3.1% utilization) identifies \$4.2M (28%) as excess and \$820K as zero-movement stock. Six product families concentrate 64% of excess value. Sensitivity analysis confirms classification stability under $\pm 10\%$ perturbation of decay factors. DNEIF requires only standard ERP transactional data and produces transparent, dollar-valued, SKU-level output that neither forecasting models nor segmentation schemes provide.

Keywords

Aftermarket inventory, coverage adequacy, excess stock classification, intermittent demand, and spare parts optimization.

1. Introduction

Aftermarket spare parts occupy an awkward position in manufacturing supply chains. They are operationally essential — a missing replacement component can void service agreements, trigger regulatory penalties, and erode customer loyalty — but economically marginal, generating low and sporadic revenue relative to the capital they immobilize. In multi-product durable goods enterprises, aftermarket parts catalogs routinely grow to thousands of SKUs as each new product generation adds components to the serviceable portfolio without a corresponding retirement of legacy items. The result, documented across industries, is that aftermarket inventory accounts for 12–25% of total inventory value while turning over at rates far below finished goods (Botter and Fortuin 2000; Hu et al. 2018).

The standard response takes one of two forms. The first applies intermittent demand forecasting — Croston's method (Croston 1972), the Syntetos-Boylan Approximation (Syntetos and Boylan 2001), or machine learning models (Lolli et al. 2017) — to predict SKU-level consumption. These methods improve accuracy relative to naive benchmarks but remain unreliable when the majority of SKUs experience fewer than a dozen non-zero demand events over multiple years. The second approach applies ABC-XYZ segmentation to classify SKUs into policy buckets by value and demand variability. This is useful for differentiating replenishment rules but answers the wrong question: it tells the planner which bucket a part belongs to, not how many excess units are sitting on the shelf.

This paper proposes a third approach that sidesteps the accuracy limitations of forecasting and the granularity limitations of segmentation. The Demand-Normalized Excess Identification Framework (DNEIF) asks, for each SKU: does the current on-hand quantity exceed what any plausible demand scenario would consume over the remaining service horizon? This is a coverage adequacy test, not a demand forecast. When on-hand exceeds the ceiling of plausible consumption, the surplus is excess by construction — regardless of how uncertain the point forecast might be.

1.1 Objectives

The paper makes three contributions. First, it formalizes a named, three-stage framework — demand normalization, category-calibrated geometric decay projection, and deterministic excess classification — that operationalizes coverage adequacy testing using only standard ERP data. Second, it provides a simulation-based evaluation calibrated to published industry parameters, producing evidence on excess magnitude, category concentration, Pareto structure, and sensitivity to assumptions. Third, it positions the coverage-adequacy framing against intermittent demand forecasting and ABC-XYZ segmentation, clarifying the conditions under which each approach is appropriate and the specific analytical gap DNEIF fills.

2. Literature Review

2.1 Intermittent Demand Forecasting

Croston (1972) decomposed intermittent demand estimation into separate models for demand size and inter-arrival interval, enabling exponential smoothing to be applied to each component independently. Syntetos and Boylan (2001) showed that Croston's estimator carries a positive bias and proposed a corrected approximation (SBA) that has become a standard benchmark. Syntetos et al. (2005) introduced a demand-pattern classification (smooth, erratic, intermittent, lumpy) to guide method selection. Teunter et al. (2011) connected forecast error to obsolescence risk, demonstrating that mechanical safety-stock setting propagates forecast uncertainty directly into excess accumulation. Boylan and Syntetos (2021) provide a comprehensive treatment in their recent monograph. These methods address the prediction problem; DNEIF addresses a different question — given current stock levels, is there more inventory than any plausible demand trajectory could absorb? — which can be answered with lower informational requirements.

2.2 Spare Parts Supply Chain Strategy

Sherbrooke (1968) and Muckstadt (1973) established multi-echelon optimization for military spare parts. Cohen et al. (2006) extended the strategic framing to commercial aftermarket operations. Wagner and Lindemann (2008) modeled end-of-life procurement as a newsvendor problem with uncertain residual demand. Van Jaarsveld et al. (2015) applied Bayesian estimation to repair-shop parts control. These contributions address network design or procurement timing but do not propose methods for identifying which specific SKUs carry excess stock, or by how much.

2.3 Segmentation and Machine Learning

ABC-XYZ classification remains the dominant industrial practice for spare parts policy differentiation (Bacchetti and Saccani 2012). Its principal limitation in the excess-identification context is that it operates at the bucket level: a CZ SKU with 10 excess units and one with 10,000 receive identical treatment. Machine learning approaches (Lolli et al. 2017; Boylan and Syntetos 2021) improve predictive accuracy on structured datasets but face adoption barriers in mid-market manufacturing — interpretability, infrastructure, and the difficulty of explaining black-box outputs to operations teams. DNEIF occupies the space between coarse segmentation and complex prediction: it produces SKU-level, dollar-valued quantification using deterministic, auditable logic.

3. Methods

3.1 Overview

DNEIF operates in three sequential stages applied independently to every active aftermarket SKU (Figure 1). The framework is deterministic: given the same input data and parameter choices, it produces identical output. This property enables auditability and eliminates stochastic variation between runs, a practical advantage for organizations that must justify disposition decisions to finance and operations leadership.




Figure 1. DNEIF three-stage analytical pipeline. Stage 1 normalizes consumption history via IQR filtering; Stage 2 projects a 10-year cumulative demand ceiling using category-calibrated geometric decay; Stage 3 classifies each SKU by comparing on-hand stock to the projection. A composite priority score sequences remediation actions.

3.2 Notation

Table 1 summarizes the principal symbols used throughout the methodology.

Table 1. Principal notation

Symbol	Definition
$s(i,t)$	Observed consumption for SKU i in month t
T	Lookback window length (= 36 months)
$F(i)$	Set of months passing the IQR filter for SKU i
$T^*(i)$	Number of retained months: $ F(i) $
$d_{\sim}(i)$	Normalized monthly demand rate after IQR filtering
$\lambda(c)$	Annual demand retention factor for category c
H	Planning horizon in years (= 10)
$D(i,H)$	Projected cumulative demand over H years
$h(i)$	Current on-hand quantity for SKU i
$c(i)$	Unit standard cost for SKU i
$E(i)$	Excess quantity: $\max(0, h(i) - D(i,H))$
$V_{ex}(i)$	Excess value: $E(i) \times c(i)$
$m(i)$	Count of zero-consumption months in lookback
$r(i)$	Model-relevance ratio for SKU i
$P(i)$	Composite priority score

3.3 Stage 1: Demand Normalization

Let $s(i,t)$ denote observed consumption for SKU i in month t , with $t = 1, \dots, T$ and $T = 36$. Define the first and third quartiles of the monthly series as $q_1(i)$ and $q_3(i)$, with $IQR(i) = q_3(i) - q_1(i)$. A month t is retained if its observation falls within the Tukey fence:

$$F(i) = \{ t : q_1(i) - 1.5 \cdot IQR(i) \leq s(i,t) \leq q_3(i) + 1.5 \cdot IQR(i) \} \quad (1)$$

The normalized monthly demand rate is the mean of retained observations:

$$d_{\sim}(i) = [1 / T^*(i)] \cdot \sum \{ s(i,t) : t \in F(i) \} \quad (2)$$

where $T^*(i) = |F(i)|$ is the count of retained months. SKUs for which $d_{\sim}(i) = 0$ are flagged as *zero-movement* inventory — items with no qualifying consumption across the entire lookback period.

3.4 Stage 2: Forward Projection with Category Decay

Aftermarket demand declines as the installed product base shrinks through retirement and technology substitution. DNEIF models this decline with a geometric decay function parameterized by product category:

$$D(i, H) = \sum_{y=1..H} 12 \cdot d\sim(i) \cdot \lambda(c)^{y-1} \quad (3)$$

The closed-form geometric-series expression is:

$$D(i, H) = 12 \cdot d\sim(i) \cdot [1 - \lambda(c)^H] / [1 - \lambda(c)], \lambda(c) < 1 \quad (4)$$

where H is the planning horizon (set to 10 years) and $\lambda(c)$ is the annual demand retention factor for category c. Lower λ applies to categories with faster technology turnover (e.g., electronic control modules, 0.82–0.88); higher λ applies to structural or mechanical categories serving multiple product generations (0.93–0.97).

3.5 Stage 3: Excess Classification

Let $h(i)$ and $c(i)$ denote on-hand quantity and unit standard cost. The excess quantity and excess value are:

$$E(i) = \max(0, h(i) - D(i, H)), \quad V_ex(i) = E(i) \cdot c(i) \quad (5)$$

Classification is mutually exclusive and exhaustive. Right-sized: $E(i) = 0$. Excess: $E(i) > 0$ and $d\sim(i) > 0$. Zero-movement: $d\sim(i) = 0$ and $h(i) > 0$ — no qualifying consumption in 36 months, the tier carrying the strongest removal signal.

3.6 Prioritization Scoring

Each excess SKU receives a composite priority score:

$$P(i) = w_1 \cdot V_ex(i) + w_2 \cdot [m(i)/36] + w_3 \cdot [1 - r(i)] \quad (6)$$

where $V_ex(i)$ is the min-max normalized excess value across excess SKUs, $m(i)$ counts zero-consumption months, $r(i)$ captures model relevance, and $w_1 = 0.50$, $w_2 = 0.30$, $w_3 = 0.20$ are stakeholder-calibrated weights. Higher $P(i)$ values indicate stronger candidates for early remediation.

4. Data Collection

The framework was developed in an applied context whose data cannot be disclosed. To enable evaluation, we constructed a synthetic portfolio calibrated to published industry parameters (Hu et al. 2018; Bacchetti and Saccani 2012); all reported results are simulation-derived. The simulated catalog comprises 11,200 active SKUs across seven product families (28 categories), with total on-hand value of \$14.8M. Unit costs follow a Pareto Type I distribution ($\alpha = 1.5$). Thirty-six months of demand were generated per SKU: occurrence follows a Bernoulli process with category-specific probability in [0.05, 0.45], and magnitude is log-normal. Approximately 19% of SKUs were assigned zero demand to generate zero-movement profiles. Retention factors were assigned by category (0.82–0.88 electronic, 0.93–0.97 mechanical), and on-hand quantities were set to a log-normal multiple of cumulative consumption to reproduce defensive overstocking. Table 2 summarizes the parameters.

Table 2. Summary of simulation parameters and calibration basis

Parameter	Value	Source / rationale
Active SKUs	11,200	Industry median, multi-brand mfg.
Product families / categories	7 / 28	Representative hierarchy
Total inventory value	\$14.8M	Industry range \$5M–\$50M
Unit cost distribution	Pareto ($\alpha=1.5$)	Right-skewed catalog
Demand probability $p(c)$	[0.05, 0.45]	Syntetos–Boylan ranges
Demand size occurrence	Log-normal	Bacchetti and Saccani (2012)
Zero-movement fraction	19.2%	Industry range 15–25%
Retention factors $\lambda(c)$	[0.82, 0.97]	Technology-turnover calibration
On-hand multiplier	LogN(2.5, 1.8)	Defensive overstocking
Utilization rate	3.1%	Industry range 2–5%

5. Results and Discussion

5.1 Numerical Results

Table 3 summarizes the three-tier classification. DNEIF classifies 72% of inventory value (\$10.6M) as right-sized and 28% (\$4.2M) as excess, with a zero-movement subset comprising 5.5% of value and 19% of SKU count. The 28% excess ratio falls within the range reported for unoptimized aftermarket portfolios (Botter and Fortuin 2000).

Table 3. Portfolio-level classification output

Tier	SKUs	Units (K)	Value (\$M)	% Value
Right-sized	7,480	1,340	10.60	71.6%
Excess (non-zero demand)	1,570	418	3.38	22.8%
Zero-movement	2,150	272	0.82	5.5%
Total	11,200	2,030	14.80	100.0%

Excess is concentrated rather than diffuse. Six of seven named product families account for 64% of total excess value (Table 4, Figure 4). This Pareto-like concentration means a remediation campaign focused on six families captures nearly two-thirds of recoverable value without a portfolio-wide review of all 11,200 SKUs.

Table 4. Six product families concentrating 64% of excess value

Product Family	Total (\$K)	Excess (\$K)	% of Excess
Power Modules	3,200	582	13.9%
Structural Assemblies	2,100	518	12.3%
Flow Regulators	1,300	476	11.3%
Thermal Components	1,100	421	10.0%
Drive Systems	900	364	8.7%
Interface Panels	800	318	7.6%
Subtotal	9,400	2,679	63.8%

The top 10 SKUs by excess value account for approximately 14% of total excess (\$588K); Table 5 shows the top five. The leading item — a digital control module — contributes 6.3% alone, illustrating the leverage available through targeted action on a small number of high-value items.

Table 5. Top five SKUs by excess value (anonymized, simulated)

SKU	Description	Inv. (\$K)	Excess (\$K)	% Excess
PM-4012	Digital control module, Gen-4	612	263	6.3%
SA-7085	Universal mounting kit	82	80	1.9%
IP-2094	Handle assy., top-mount	75	73	1.7%
SA-3061	Shelf assy., full platform, SS	96	68	1.6%
SA-1047	Multi-config chassis, Ser-4	229	64	1.5%

The 2,150 zero-movement SKUs (\$820K) carry the strongest analytical signal for removal. Their classification is independent of decay-factor assumptions — it depends only on the empirical absence of consumption over 36 months — making them the lowest-risk entry point for remediation (Table 6).

Table 6. Top five zero-movement SKUs by on-hand value (simulated)

SKU	Description	Units	Value (\$K)
ZM-001	Thermal enclosure, sand-coat	700	43.8
ZM-002	Ignition switch module assy.	11,400	22.5
ZM-003	Drive motor assy., 12VDC	605	14.9
ZM-004	Integrated basin housing	182	14.3
ZM-005	Combustion chamber, built-in	168	12.3

5.2 Graphical Results

Figure 2 shows the projected portfolio value under a phased remediation plan; the right-sized base remains while excess declines through three-phase removal. Figure 3 decomposes the classification by dollar value and by unit count, showing that excess constitutes 28% of value but 34% of units, reflecting lower average cost among excess items. Figure 4 displays total versus excess inventory by product family, with Power Modules and Structural Assemblies together accounting for 26% of total excess value.




Figure 2. Projected portfolio value under phased remediation (2025–2035). Green: right-sized base inventory; red: excess declining through three-phase removal. All values simulated.



Figure 3. Classification split by dollar value (a) and unit count (b). Excess constitutes 28% of value but 34% of units.




Figure 4. Total (dark) versus excess (red) inventory by product family. Power Modules and Structural Assemblies together account for 26% of total excess value.

5.3 Statistical Analysis (Sensitivity)

Uniform perturbation of $\lambda(c)$ by $\pm 5\%$ and $\pm 10\%$ yields excess estimates ranging from \$3.60M (-14%) to \$4.70M ($+12\%$). The asymmetry reflects the convexity of cumulative demand in λ (Eq. 4). Importantly, zero-movement classification is invariant to λ perturbation, since it depends only on the presence of historical consumption (Table 7, Figure 5).

Table 7. Sensitivity of total identified excess to λ perturbation.

Scenario	λ shift	Excess (\$M)	Change vs. baseline
Conservative	+10%	3.60	-14.3%
Mildly conservative	+5%	3.87	-7.9%
Baseline	0%	4.20	—
Mildly aggressive	-5%	4.55	+8.3%
Aggressive	-10%	4.70	+11.9%

nce of historical consumption, not by the forward projection.




Figure 5. Sensitivity of excess identification to λ perturbation. The dashed line marks the baseline scenario; classification direction is stable across the full $\pm 10\%$ range.

5.4 Proposed Improvements (Remediation Roadmap)

The composite priority score (Eq. 6) sequences remediation into three phases (Figure 6). Phase 1 targets zero-movement stock (\$820K) for immediate scrap validation or supplier return. Phase 2 addresses the top-decile high-priority excess SKUs (\$580K) through negotiated disposition. Phase 3 implements category-wide stocking policy adjustments across the six concentrated product families (\$2.8M). Estimated total working-capital recovery is \$4.2M over a 24-month execution horizon.



Figure 6. Three-phase remediation roadmap, ordered by analytical confidence (Phase 1 highest) and execution complexity (Phase 1 lowest).

5.5 Validation and Positioning

DNEIF vs. ABC-XYZ. ABC-XYZ assigns SKUs to behavioral buckets useful for policy differentiation but does not answer how many units of a specific SKU are excess. In the simulation, the top-10 excess SKUs span three distinct ABC-XYZ segments, confirming that excess is not confined to a single behavioral class. DNEIF complements ABC-XYZ by adding a dollar-valued, unit-level quantification layer.

DNEIF vs. intermittent demand forecasting. Croston/SBA methods produce point forecasts converted to stocking decisions via safety-stock formulas — a pipeline that compounds forecast uncertainty with inventory-model assumptions. DNEIF bypasses this by posing a ceiling question rather than a point-estimate question. The trade-off is explicit: DNEIF is a classification tool, not a replenishment engine. In practice the two are complementary.

DNEIF vs. ML forecasting. Machine learning models achieve strong accuracy on datasets with sufficient training signal, but for aftermarket parts where most SKUs have fewer than 12 non-zero observations over 36 months, model complexity provides diminishing returns (Boylan and Syntetos 2021), and the required infrastructure and interpretability burden exceed what many mid-market manufacturers can support. DNEIF is infrastructure-light by design.

Three features make DNEIF operationally practical: the 10-year coverage horizon provides a credibility anchor across finance, operations, and legal; the category concentration finding reduces the coordination problem from portfolio-wide to manageable scope; and the zero-movement tier provides a politically safe starting point, since 36 months of zero activity is difficult to dispute. Limitations are symmetrically clear: results derive from a calibrated simulation rather than longitudinal field deployment; the framework assumes normalized historical consumption approximates future demand; the $\lambda(c)$ parameters are the most assumption-sensitive component; and a supplier-feasibility overlay is needed before executing removal.

6. Conclusion

This paper introduced DNEIF, a three-stage framework that reframes aftermarket spare parts inventory optimization from demand forecasting to coverage adequacy testing. Rather than asking what demand will be, the framework asks whether there is more stock than any plausible scenario would consume — a question that can be answered deterministically using standard transactional data. A simulation calibrated to industry parameters identified \$4.2M (28%) of a \$14.8M portfolio as excess, with 64% concentrated in six product families and \$820K in zero-movement stock offering the lowest-risk remediation entry point. The framework requires no specialized software, no machine learning infrastructure, and no distributional assumptions beyond a category-level decay rate. Its limitation is that it is a classification tool, not a replenishment engine, and it has not yet been validated through longitudinal field deployment; closing that gap through field validation, integration with statistical forecasting platforms such as SAP IBP, and dynamic product-to-part compatibility matrices is the most important next step.

References

- Aberdeen Group, Spare parts planning and optimization, Aberdeen Research Report, 2019.
- Bacchetti, A. and Saccani, N., Spare parts classification and demand forecasting for stock control, *Omega*, vol. 40, no. 6, pp. 722-737, 2012.
- Botter, R. and Fortuin, L., Stocking strategy for service parts, *International Journal of Operations & Production Management*, vol. 20, no. 6, pp. 656-674, 2000.
- Boylan, J.E. and Syntetos, A.A., *Intermittent Demand Forecasting*, Wiley, 2021.
- Cohen, M.A., Agrawal, N. and Agrawal, V., Winning in the aftermarket, *Harvard Business Review*, vol. 84, no. 5, pp. 129-138, 2006.
- Croston, J.D., Forecasting and stock control for intermittent demands, *Journal of the Operational Research Society*, vol. 23, no. 3, pp. 289-303, 1972.
- Hu, Q., Boylan, J.E., Chen, H. and Labib, A., OR in spare parts management: A review, *European Journal of Operational Research*, vol. 266, no. 2, pp. 395-414, 2018.
- Lolli, F., Ishizaka, A., Gamberini, R., Rimini, B. and Messori, M., FlowSort-GDSS for sorting problems with application to FMEA, *Expert Systems with Applications*, vol. 82, pp. 217-225, 2017.
- Muckstadt, J.A., A model for a multi-item, multi-echelon, multi-indenture inventory system, *Management Science*, vol. 20, no. 4, pp. 472-481, 1973.
- Sherbrooke, C.C., METRIC: A multi-echelon technique for recoverable item control, *Operations Research*, vol. 16, no. 1, pp. 122-141, 1968.
- Syntetos, A.A. and Boylan, J.E., On the bias of intermittent demand estimates, *International Journal of Production Economics*, vol. 71, no. 1-3, pp. 457-466, 2001.
- Syntetos, A.A., Boylan, J.E. and Croston, J.D., On the categorization of demand patterns, *Journal of the Operational Research Society*, vol. 56, no. 5, pp. 495-503, 2005.
- Teunter, R.H., Syntetos, A.A. and Babai, M.Z., Intermittent demand: Linking forecasting to inventory obsolescence, *European Journal of Operational Research*, vol. 214, no. 3, pp. 606-615, 2011.
- Van Jaarsveld, W., Dollevoet, T. and Dekker, R., Improving spare parts inventory control at a repair shop, *Omega*, vol. 57, pp. 217-229, 2015.
- Wagner, S.M. and Lindemann, E., Spare parts management in the engineering industry, *Production Planning & Control*, vol. 19, no. 4, pp. 397-407, 2008.

Biography

Arjun Pardasani is a business and demand analyst working in the consumer goods sector, where he builds production-grade demand forecasting and supply chain analytics systems using SAP Integrated Business Planning, Power BI, and Python across large multi-SKU portfolios. He holds a Master of Science in Engineering Management from Northeastern University and a Bachelor of Technology from VIT Vellore. His research interests span demand forecasting, multi-echelon inventory optimization, aftermarket spare parts management, and the application of statistical and AI-driven methods to enterprise planning systems. He has authored peer-reviewed publications in IEEE venues and has served as a reviewer and session judge for several international academic conferences. His current work focuses on bridging rigorous quantitative methods with auditable, practitioner-deployable frameworks for inventory and supply chain decision-making.