European Journal of Logistics, Purchasing and Supply Chain Management, 14 No.4, pp.29-54, 2026

Print ISSN: 2054-0930 (Print)

Online ISSN: 2054-0949 (Online)

Website: https://www.eajournals.org/

Publication of the European Centre for Research Training and Development-UK

An Integrated Predictive Inventory and Demand-Planning Framework for Improving Operational Reliability and Cost Efficiency in the U.S. Oil and Gas Industry

Esther Chinonye Enyinnia
Business School of Netherlands– Executive MBA (2023)
Nnamdi Azikiwe University-B.SC Applied Microbiology (2008)

doi: https://doi.org/10.37745/ejlpscm.2013/vol14n42954 Published November 22, 2026

Citation: Enyinnia E.C. (2026) An Integrated Predictive Inventory and Demand-Planning Framework for Improving Operational Reliability and Cost Efficiency in the U.S. Oil and Gas Industry, *European Journal of Logistics, Purchasing and Supply Chain Management*, 14 No.4, pp.29-54

Abstract: *Unplanned downtime in U.S. oil and gas operations compels firms to hold expensive critical spares, yet operators still experience stockouts of consequential items while accumulating slow-moving maintenance, repair, and operations (MRO) inventory. This tension between high service-level requirements and working-capital pressure persists because demand forecasting, criticality classification, and working-capital management are typically treated as separate problems, and evidence specific to U.S. oil and gas contexts is fragmented. This paper proposes an integrated predictive inventory and demand-planning framework that segments items by demand pattern, economic value, and operational criticality; assigns forecasting methods to each segment; converts forecast distributions into criticality-differentiated replenishment policies; and evaluates outcomes jointly on service and financial measures. The study adopts a conceptual, integrative design, supported by an illustrative analytical demonstration based on a hypothetical portfolio of 1,000 stock-keeping units. The demonstration indicates that segment-specific forecasting could improve accuracy relative to a uniform baseline, that criticality-differentiated policies could reduce stockout frequency without requiring higher inventory investment, and that selective rationalisation of slow-moving stock could release working capital. These outputs rest on stated assumptions and do not constitute empirical validation. The paper concludes that operators should adopt the framework in sequence: first establishing data integrity and criticality-informed classification, then piloting segment-level forecasting, and finally embedding forecasts in replenishment policy and working-capital governance.*

KEYWORDS: *predictive inventory management; demand forecasting; intermittent demand; critical spare parts; MRO inventory; criticality classification; stockout risk; slow-moving inventory; working-capital optimisation; service-level differentiation; U.S. oil and gas; supply chain analytics*

INTRODUCTION

Unplanned downtime in the U.S. oil and gas industry is operationally disruptive and financially severe. A failed electric submersible pump in a producing well, a compressor outage at a midstream station, or an

unavailable valve in a refinery unit can interrupt production, defer revenue that is seldom recovered in full, and, in some settings, create safety and environmental exposures that extend well beyond the immediate loss of throughput. Because the assets involved are capital-intensive, geographically dispersed, and frequently specified to particular tolerances, the components required to restore them often carry long and uncertain replacement lead times. Operators have therefore historically responded by holding large stocks of maintenance, repair, and operations (MRO) materials, with critical spares positioned as insurance against failure events whose timing is unknown but whose consequences are asymmetric. The resulting inventory systems nevertheless exhibit a paradox familiar to practitioners: organizations that tie up substantial capital in spares still experience stockouts of the items that matter most, while a considerable share of their holdings consists of slow-moving or obsolete items that have not been issued for years.

This paradox reflects a structural tension between two legitimate managerial objectives. On one side, reliability engineers and operations leaders require very high service levels for critical parts, since a single unavailable component can idle an asset whose hourly opportunity cost dwarfs the unit price of the part. On the other, finance functions face sustained pressure to release working capital, a pressure that intensifies in a sector whose capital discipline is shaped by commodity-price cycles, investor expectations for returns, and periodic cost-reduction programs. Excess MRO stock carries not only the financing cost of the capital invested but also costs of storage, handling, deterioration, and eventual write-down. Conversely, aggressive reductions in inventory that are not informed by the criticality of individual items merely transfer cost from the balance sheet to the maintenance backlog and the production ledger. Neither extreme is tenable, and the practical question is how to differentiate stocking decisions item by item in a manner that is both statistically defensible and economically coherent.

The academic and applied literature offers many tools relevant to this problem, but they are typically developed and deployed in isolation. Demand forecasting for spare parts has been advanced considerably by methods designed for intermittent and lumpy demand, beginning with Croston (1972) and refined in subsequent work on bias correction and bootstrapping approaches (for example, Syntetos and Boylan, 2005; Willemain et al., 2004). Criticality classification, by contrast, is usually conducted through ABC, XYZ, and VED-type schemes that rely on managerial judgement and are updated infrequently, and these classifications are seldom linked formally to the probabilistic forecasts that drive replenishment parameters. Working-capital optimization, in turn, is commonly treated as a downstream financial exercise, evaluated through aggregate indicators such as inventory turnover or the cash-to-cash cycle rather than through item-level policy decisions. When these three elements are handled sequentially by different functions, errors and inconsistent assumptions propagate: a forecast optimized for statistical accuracy may be misaligned with the service levels that criticality demands, and a service-level target set without reference to carrying cost may be economically indefensible. Moreover, evidence specific to U.S. oil and gas operating contexts is fragmented. Much of the empirical work on spare-parts management draws on aviation, automotive aftermarket, or generic manufacturing settings, while studies in energy contexts tend to be single-case, proprietary, or focused on one decision at a time. The extent to which findings transfer to environments characterized by commodity-price volatility, drilling-activity cycles, remote field locations, and multi-system ERP landscapes therefore remains insufficiently established.

This paper responds by proposing an integrated predictive inventory and demand-planning framework that treats forecasting, criticality classification, and working-capital optimization as interdependent components of a single decision architecture. The framework first segments MRO items by demand pattern and by consequence of stockout, then assigns forecasting methods suited to each segment, ranging from intermittent-demand estimators to gradient-boosted and probabilistic models that exploit maintenance, operational, and activity-level covariates. The resulting forecast distributions feed criticality-differentiated inventory policies whose service-level targets are set jointly with item-level carrying costs, so that safety stock, reorder parameters, and disposition rules for slow-moving items are derived within one economic calculus. Performance is assessed on forecast accuracy, stockout incidence, slow-moving inventory share, and working-capital measures, thereby linking statistical quality to operational reliability and financial outcomes. Illustrative parameter values used later are labeled as such and are not presented as industry benchmarks.

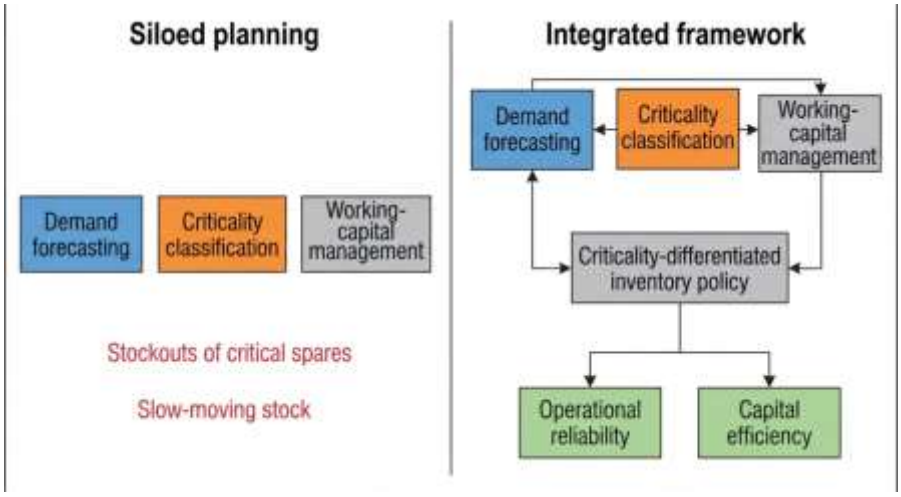


Figure: Siloed planning produces both stockouts and excess stock, whereas an integrated framework links forecasting, criticality, and working capital in one decision process.

The study is guided by the following research questions.

1. **RQ1:** To what extent does integrating segment-specific forecasting methods, including probabilistic approaches, improve forecast accuracy for critical, intermittent, and slow-moving MRO items relative to single-method benchmarks?

2. **RQ2:** How does coupling forecast distributions with criticality-differentiated inventory policies affect stockout incidence for critical spares at a given level of inventory investment?

3. **RQ3:** What reductions in slow-moving inventory and working-capital requirements can be achieved without degrading service levels for critical items?

4. **RQ4:** Which demand, criticality, and cost characteristics condition the performance of the integrated framework in U.S. oil and gas operating contexts?

The paper makes the following contributions.

- It formalizes a unified framework that links demand-pattern segmentation, criticality classification, probabilistic forecasting, and working-capital-aware inventory policy within one decision process.
- It provides evidence specific to U.S. oil and gas MRO contexts, addressing a literature that is fragmented and reliant in part on evidence from other industries.
- It evaluates the framework jointly on forecast accuracy, stockout incidence, slow-moving inventory, and working capital, rather than on statistical accuracy alone.
- It offers decision-relevant guidance for practitioners on where predictive sophistication yields economic returns and where simpler policies suffice.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature, Section 3 describes the methodology, Section 4 reports results, Section 5 discusses implications and limitations, and Section 6 concludes.

LITERATURE REVIEW

This review is organized thematically around four bodies of work that bear on the integrated framework: forecasting methods for intermittent and lumpy spare-parts demand, criticality classification and service-differentiated inventory control, the linkage between slow-moving stock and working capital, and the operating characteristics of oil and gas firms. Each stream is assessed for what it establishes, where findings conflict, and how well it transfers to U.S. oil and gas contexts. Claims are limited to what the cited studies support, and no performance figures are asserted that the literature does not report.

Forecasting Intermittent and Lumpy Spare-Parts Demand

Spare-parts demand typically consists of long runs of zero demand interrupted by irregular, highly variable requirements, a pattern that violates the assumptions of conventional exponential smoothing. Croston (1972) addressed this by smoothing non-zero demand sizes and inter-demand intervals separately and forecasting demand per period as their ratio. Syntetos and Boylan (2001) showed that this estimator is biased, and Syntetos and Boylan (2005) proposed a correction, widely termed the Syntetos–Boylan Approximation (SBA), which improved accuracy in tests on industrial data. The same authors also proposed categorizing items by the average inter-demand interval and the squared coefficient of variation of demand sizes, yielding smooth, intermittent, erratic, and lumpy classes (Syntetos, Boylan, and Croston, 2005). Teunter, Syntetos, and Babai (2011) identified a further weakness: Croston-type methods do not update during periods without demand and therefore respond slowly to obsolescence. Their alternative updates the

probability of demand every period, and Babai, Syntetos, and Teunter (2014) examined the resulting trade-off between accuracy and responsiveness to obsolescence empirically.

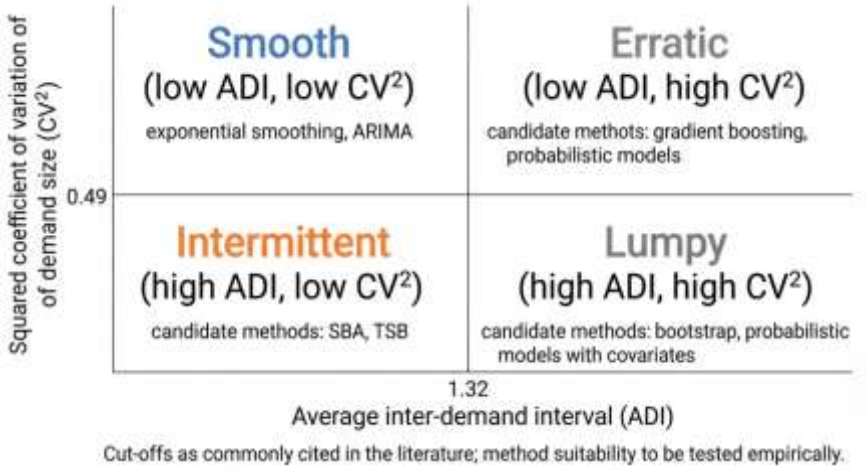


Figure: Demand-pattern classes by inter-demand interval and demand-size variability, with candidate forecasting methods for each.

Other approaches depart from the Croston family. Willemain, Smart, and Schwarz (2004) proposed a bootstrap procedure that targets the distribution of lead-time demand directly and reported improved performance over exponential smoothing and Croston's method on service-parts data. Nikolopoulos et al. (2011) advanced temporal aggregation as a means of reducing intermittency, and Petropoulos and Kourentzes (2015) studied forecast combinations. Zhu et al. (2017) applied extreme value theory to the upper tail of spare-parts demand, which matters because the tail drives safety stock. Collectively, these contributions shift emphasis from point forecasts toward demand distributions, a shift consistent with the general case for probabilistic forecasting made by Gneiting and Katzfuss (2014).

Machine-learning and deep-learning methods have been evaluated more recently, with mixed conclusions. Gutierrez, Solis, and Mukhopadhyay (2008) reported that neural networks compared favorably with several traditional methods on lumpy demand, whereas Kourentzes (2013) found that conclusions about neural networks for intermittent series depend heavily on network specification and on the performance metric adopted. Salinas et al. (2020) introduced an autoregressive recurrent network that produces probabilistic forecasts across many related series. In the M5 competition, which used retail unit sales rather than spare parts, Makridakis, Spiliotis, and Assimakopoulos (2022) reported that gradient-boosting-based approaches featured prominently among the leading methods. Evidence from aviation (Ghobbar and Friend, 2003) and industrial settings is less extensive and less uniform.

Three points of contention follow. First, the ranking of methods is sensitive to the error measure, since scale-dependent and percentage-based measures behave poorly when many observations are zero

(Hyndman and Koehler, 2006; Prestwich et al., 2014). Second, accuracy measures are weakly connected to inventory outcomes, a concern that motivates evaluation of predictive distributions (Kolassa, 2016) and of accuracy-implication metrics. Third, the transferability of retail results to low-volume, high-consequence parts is unestablished. Bacchetti and Saccani (2012) additionally documented a gap between research and practice in spare-parts forecasting, with firms often relying on simple methods despite available alternatives.

Criticality Classification and Multi-Echelon Inventory Optimization

Classification schemes allocate management attention and service targets across large item populations. ABC analysis ranks items by annual usage value, and Flores and Whybark (1987) extended it to multiple criteria, a direction developed further by Ng (2007). VED analysis, which classifies items as vital, essential, or desirable according to the consequence of unavailability, is long established in maintenance practice and is conceptually closer to the needs of asset-intensive operations. Molenaers et al. (2012) presented a structured criticality assessment for spare parts in an industrial case, and Roda et al. (2014) reviewed multi-criteria classification of spare parts, noting considerable heterogeneity in the criteria employed. Bacchetti et al. (2013) proposed hierarchical classification combining demand characteristics and other attributes. A persistent limitation is that criticality scores are usually elicited from engineers, updated infrequently, and not linked formally to the forecasting method or to the stocking parameters.

Teunter, Babai, and Syntetos (2010) made a notable contribution by showing that ranking items by annual usage value does not generally minimize cost for a given aggregate service level, and by proposing criteria that account for demand, unit cost, and the penalty associated with stockouts. Their result supports service-level differentiation, in which higher availability is assigned where the consequence of a stockout is greater. In a petrochemical plant, Gelders and Van Looy (1978) had earlier studied separate inventory policies for slow and fast movers, an early instance of segmenting stock policy in a process-industry setting.

Multi-echelon theory supplies the structural foundation for network-level decisions. Sherbrooke (1968) introduced METRIC for repairable items in a two-echelon system, and Muckstadt (1973) extended the approach to multiple indenture levels; Sherbrooke (2004) provides a comprehensive treatment. Kranenburg and van Houtum (2008) addressed service differentiation across items explicitly, and Basten and van Houtum (2014) reviewed system-oriented spare-parts models. Surveys by Kennedy, Patterson, and Fredendall (2002) and by Driessen et al. (2015) identify the integration of forecasting with inventory control as an open research agenda, and Van Horenbeek et al. (2013) make a parallel observation about joint maintenance and inventory optimization. Most multi-echelon models assume that demand or failure rates are known, which leaves unresolved the estimation error that dominates when demand is intermittent. Trapero, Cardós, and Kourentzes (2019) illustrate the benefit of deriving safety stock empirically from forecast errors, and the prescriptive-analytics literature (Ban and Rudin, 2019; Bertsimas and Kallus, 2020) shows how decisions may be learned directly from covariates. These contributions were developed outside the MRO setting, and their application to critical spares remains limited.

Slow-Moving Inventory and Working-Capital Linkages

Inventory carrying cost comprises capital cost, storage and handling, insurance, taxes, deterioration, and obsolescence risk (Silver, Pyke, and Thomas, 2017). Practitioner rules of thumb for the aggregate rate circulate widely, but they vary across firms and are rarely documented transparently, so this paper does not rely on any single figure. For slow-moving items, the obsolescence component is particularly uncertain. Van Jaarsveld and Dekker (2011) estimated obsolescence risk from demand data in a case study, thereby linking forecasting directly to disposition decisions, while the Teunter et al. (2011) method addresses the same concern within the forecasting step.

At the firm level, the inventory–performance relationship has been studied extensively without producing a uniform conclusion. Eroglu and Hofer (2011) reported that the association between inventory leanness and financial performance is not uniformly positive and varies across firms. Gaur, Fisher, and Raman (2005) related inventory turnover to margins and capital intensity in retail, and Rumyantsev and Netessine (2007) linked inventory levels to demand uncertainty and lead times across industries. Chen, Frank, and Wu (2005) documented long-run trends in inventories of American firms. Working-capital research typically uses the cash conversion cycle introduced by Richards and Laughlin (1980). Shin and Soenen (1998) and Deloof (2003) reported negative associations between cash conversion cycle length and profitability in their samples, while Hofmann and Kotzab (2010) argued for a supply-chain-oriented view of working-capital management.

Two limitations are relevant here. Aggregate indicators such as turnover or the cash conversion cycle generally do not separate MRO supplies from production inventory, so they offer little guidance on item-level policy. Moreover, firm-level studies treat inventory reduction as uniformly desirable until performance deteriorates, whereas for critical spares the cost of excessive reduction materializes as downtime rather than as a financing penalty. This asymmetry implies that working-capital objectives need to be embedded in item-level policy, conditioned on criticality, rather than applied as portfolio-wide targets.

Oil and Gas Operating Characteristics

Huiskonen (2001) argued that maintenance spare parts differ from other inventory in criticality, specificity, demand pattern, and value, all of which characterize oil and gas MRO. Several features, however, are particular to the sector. Equipment populations are heterogeneous across upstream, midstream, and downstream segments, and the same component may appear under different descriptions or manufacturer part numbers across business units and ERP instances. Practitioner accounts frequently cite weak catalog standardization, although systematic published measurement of the problem is sparse. Reliability data standards exist, notably ISO 14224 for the collection and exchange of reliability and maintenance data in the petroleum, petrochemical, and natural gas industries, and the OREDA handbooks compile failure data from offshore equipment. Their adoption at the level of individual operators' spare-parts records appears uneven.

Failure patterns are also site-specific. Operating conditions such as fluid composition, sand production, temperature, pressure cycling, and duty profiles plausibly alter failure rates of nominally identical equipment, which limits the value of generic failure rates and favors models that incorporate local covariates. Demand for MRO items in this sector is further coupled to activity levels, including drilling and completion programs, turnarounds, and shutdowns, which in turn follow commodity-price cycles. Demand is therefore driven by operating events rather than by calendar seasonality, which weakens the assumptions of many time-series methods.

Finally, historical data are limited relative to the requirements of data-hungry models. Failures of critical equipment are rare by design, equipment is replaced or reconfigured, and item histories are fragmented by mergers, asset sales, and system migrations. Condition-monitoring data may mitigate this limitation, but integrating such data with spare-parts forecasts remains an open problem (Van Horenbeek et al., 2013). Published work on oil and gas inventory that was identified for this review tends toward single-case studies or addresses one decision at a time, and the author has not located a multi-operator evaluation integrating forecasting, classification, and working-capital outcomes in U.S. settings.

Synthesis

Table 1 summarizes the state of knowledge and the gap addressed by the proposed framework.

Table 1. Synthesis of the literature and the research gap

Theme	Known	Contested	Gap addressed
Forecasting	SBA corrects Croston bias; bootstrap targets lead-time demand	Machine-learning superiority; error-measure sensitivity	Segment-specific methods judged on distributional criteria
Criticality	Service differentiation lowers cost	Classification criteria; known-parameter assumptions	Criticality linked to forecast distributions
Working capital	Carrying and obsolescence costs matter	Whether leaner inventory always helps	Item-level objectives conditioned on criticality
Oil and gas	Parts are specific, critical, intermittent	Data quality; cross-industry transferability	Integrated U.S. MRO evidence

The literature therefore establishes strong individual components but leaves their integration, and their evaluation in U.S. oil and gas contexts, largely unaddressed. The framework developed in the next section

responds by connecting demand segmentation, probabilistic forecasting, criticality-based service targets, and item-level working-capital objectives, so that each decision inherits the assumptions and uncertainty of the one before it rather than being optimized separately.

METHODOLOGY

Research Design and Methodological Orientation

This study adopts a conceptual, integrative research design to develop a predictive inventory and demand-planning framework for the U.S. oil and gas industry. It does not claim to estimate causal effects or report operational improvements from a proprietary company dataset. Instead, the methodology synthesises established principles from demand forecasting, spare-parts inventory control, asset criticality assessment, and working-capital management into a coherent decision architecture. The design is theory-building and prescriptive: it identifies the information flows, analytical decisions, and performance outcomes through which forecasting can be translated into differentiated replenishment policies and, ultimately, into operational and financial consequences.

The unit of analysis is the stock-keeping unit (SKU) at a defined stocking location over discrete planning periods. The design logic is sequential but iterative: classify each item according to economic value, demand characteristics, and operational criticality; select a forecasting method appropriate to its demand process; convert the forecast and uncertainty into a service-constrained inventory policy; and evaluate the resulting inventory position against reliability and working-capital objectives. Feedback from forecast errors, stockout events, equipment changes, and obsolescence reviews is then used to update the classifications and parameters.

Logical Derivation and Framework Architecture

The proposed framework addresses the limitation of treating inventory planning as a single forecasting problem. Forecast accuracy alone does not determine whether a spare part is available when required. Inventory outcomes also depend on replenishment lead time, demand variability, supplier reliability, repair and substitution options, equipment consequence, and the cost of holding stock. Conversely, a reduction in average inventory is not necessarily an improvement if it increases the probability or duration of an operationally consequential stockout. The proposed framework therefore links four analytical layers rather than optimising each in isolation.

Four linked layers classify SKUs, estimate demand and uncertainty, translate estimates into criticality- and lead-time-sensitive replenishment parameters, and evaluate service and financial consequences. The framework is designed to operate alongside enterprise resource planning (ERP), computerized maintenance management systems (CMMS), and materials-management processes.

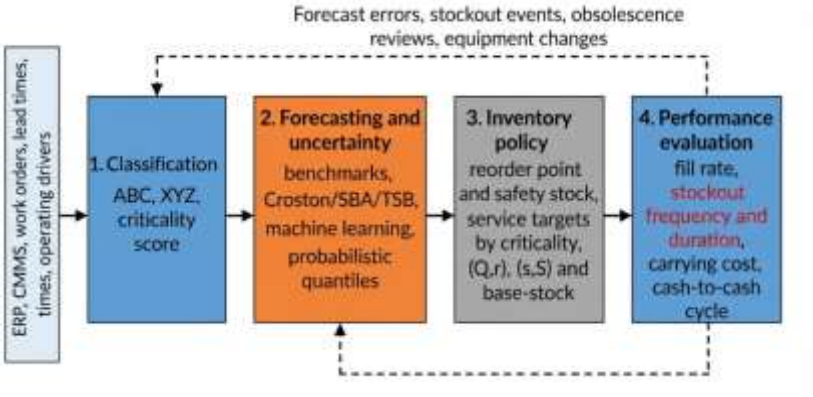


Figure: The framework links classification, forecasting, inventory policy, and performance evaluation in a continuous feedback loop.

Demand Segmentation and Criticality Classification

Segmentation precedes model selection because a single forecasting or stocking rule is unlikely to be suitable for heterogeneous MRO portfolios. The proposed classification combines ABC analysis, XYZ demand-pattern analysis, and a separate operational criticality assessment. ABC classification ranks SKUs by annual consumption value, calculated as annual usage multiplied by unit acquisition cost. Class A items account for the largest cumulative share of portfolio value, B items occupy the intermediate range, and C items comprise the residual. Thresholds should be set from the empirical distribution rather than imposed universally.

XYZ classification describes demand regularity and predictability. It may be based on the coefficient of variation for sufficiently frequent demand, supplemented by measures suited to intermittent series, including average inter-demand interval and the variability of non-zero demand size. X denotes comparatively stable demand, Y denotes moderate variability or identifiable seasonality, and Z denotes irregular or highly intermittent demand. Intermittent-demand diagnostics are used to distinguish sporadic demand from lumpy demand, in which both the interval and the non-zero quantity vary substantially.

ABC-XYZ classification is supplemented by a criticality score that reflects the operational consequence of non-availability. Dimensions include safety and environmental consequences, production impact, regulatory implications, redundancy, repairability, substitutability, and replacement lead time. Criticality is assessed through a documented cross-functional process involving maintenance, operations, procurement, and reliability personnel. A low-cost component with a credible safety consequence may require critical treatment regardless of its ABC class. Accordingly, the framework distinguishes economic priority from consequence-based priority.

Forecasting Layer and Comparative Model Evaluation

Forecasting is conducted at SKU-location level using the finest time interval supported by transaction quality and planning practice. The series should represent actual issues rather than purchase orders, since procurement timing may be decoupled from requirements. Where possible, maintenance work orders, equipment operating hours, planned shutdowns, drilling or completion schedules, and installed-equipment populations are retained as explanatory variables. Stockout periods must be flagged because recorded issues can understate latent demand when requests were unfilled.

Benchmarks include naïve and seasonal-naïve forecasts, moving averages, exponential smoothing, and ARIMA/SARIMAX for sufficiently regular series. SARIMAX is considered where defensible external drivers or seasonal structures exist; explanatory variables should be available at the forecast origin to avoid information leakage. For intermittent demand, Croston's method is included because it separately smooths non-zero demand size and the interval between non-zero events. The Syntetos–Boylan Approximation (SBA) is included as a bias-adjusted alternative, while the Teunter–Syntetos–Babai (TSB) method is relevant where demand may decline towards obsolescence because it updates the estimated probability of demand in zero-demand periods. The literature identifies positive bias in the original Croston estimator and explains the rationale for SBA and TSB adjustments (Syntetos and Boylan, 2001; Teunter, Syntetos, and Babai, 2011). Their suitability is evaluated empirically when data become available.

Machine-learning alternatives comprise regularised regression, random forests, and gradient-boosted tree models. They may capture nonlinear relationships between consumption and covariates such as equipment activity, maintenance plans, asset age, operating conditions, and supplier or project schedules. Probabilistic forecasting is preferred where feasible because inventory decisions require a distribution or quantile of lead-time demand, not only a point estimate. Prediction intervals or empirical quantiles can therefore inform safety-stock calculations and expose uncertainty that a single expected value conceals.

Models are compared through rolling-origin, out-of-sample evaluation. At each origin, only information available at that date is used to estimate forecasts for a prespecified operational horizon. The primary metrics are mean absolute scaled error (MASE), root mean squared error (RMSE), and mean forecast error or bias. MASE permits comparison across series with different scales and avoids some degeneracies of percentage-error measures, although its naïve scaling denominator must be handled where a series is constant (Hyndman and Koehler, 2006). RMSE penalises large errors and is therefore informative where occasional large under-forecasts are consequential, but it is sensitive to outliers. Bias is calculated as the mean of forecast minus actual demand, with the sign convention stated consistently; positive values indicate over-forecasting under this convention. Because aggregate bias can conceal offsetting SKU-level errors, results are reported by criticality and demand segment as well as at portfolio level. Where data permit, inventory simulations assess whether a nominally more accurate model actually improves fill rate, stockout exposure, and total cost.

Print ISSN: 2054-0930 (Print)

Online ISSN: 2054-0949 (Online)

Website: https://www.eajournals.org/

Publication of the European Centre for Research Training and Development-UK

Inventory Policy and Replenishment Logic

The inventory layer converts the selected demand distribution into item-specific control parameters. For continuously reviewed items, the reorder point, (r), comprises expected lead-time demand plus safety stock. In a simplified formulation, (r=\mu_L+SS), where (\mu_L) is expected lead-time demand and (SS) is the buffer required to achieve the chosen service objective. When demand and lead time are independent and demand variability is adequately represented, safety stock may be approximated as a service-factor multiple of the standard deviation of lead-time demand. Where lead time is variable, demand and lead-time uncertainty must both be represented; assuming fixed lead time in such circumstances would understate exposure.

Service objectives are differentiated by criticality. High-criticality parts with severe safety, environmental, or production consequences may warrant high availability targets, expedited-supply arrangements, repair loops, or strategically held insurance stock even when demand is extremely rare. Here, non-availability cost may dominate carrying cost, subject to engineering validation and obsolescence risk. For medium-criticality items, target service levels are balanced against lead time, substitution, repairability, and the financial cost of additional stock. For low-criticality, readily substitutable materials, lower service targets, supplier-held stock, periodic review, or make-to-order arrangements may be appropriate.



Figure: Service targets and replenishment logic are differentiated by operational criticality and demand pattern rather than by item value alone.

For stable, frequently demanded items, conventional continuous-review ((Q,r)) or periodic-review ((s,S)) policies can be considered. Order quantities may be determined using economic order quantity logic, adjusted for minimum order quantities, price breaks, storage constraints, and supplier capacity. For intermittent items, standard EOQ assumptions may be weak because demand arrivals are sparse and lumpy. Replenishment quantities should therefore be tested using empirical demand distributions or discrete-event

simulation where transaction histories support such analysis. Critical spares may also require a separate base-stock or one-for-one replenishment logic, especially where a single unit protects a major asset and repair or exchange cycles determine the effective supply pipeline.

Working-Capital and Operational Performance Assessment

The final layer evaluates the financial and operational consequences of alternative decisions. Inventory investment is measured using average inventory at acquisition cost, with additional attention to excess and obsolete stock, slow-moving stock, and inventory ageing. Carrying cost is estimated as the relevant annual carrying-cost rate multiplied by average inventory value. Where possible, the rate is decomposed into capital, storage, handling, insurance, taxes, deterioration, and obsolescence costs. A single assumed rate may be used for scenario analysis only if its provenance and sensitivity are disclosed.

Stockout cost is more difficult to quantify and should not be treated as a uniform penalty. It may include emergency freight, repair or rental premiums, maintenance rescheduling, lost throughput, and contractual or safety consequences. For critical equipment, the financial estimate of downtime should be developed jointly with operations and finance, with assumptions made explicit to avoid overstating avoidable losses. The framework therefore reports physical service measures—such as cycle service level, fill rate, backorders, stockout frequency, and stockout duration—alongside monetised consequences.

Working-capital outcomes include average inventory value, inventory turnover, days inventory outstanding, and the cash-to-cash cycle where receivables and payables data are available. Inventory turnover is interpreted in conjunction with service performance: an increase may reflect better inventory productivity, but it may also reflect understocking if service deteriorates. The overall decision criterion is a total relevant cost and reliability profile, not minimum inventory in isolation.

Propositions and Validation Strategy

As a conceptual study, the framework generates propositions for subsequent empirical testing rather than reporting estimated effects. First, segmentation that combines ABC-XYZ characteristics with operational criticality is expected to produce more appropriate forecast-model and service-policy assignments than value-based classification alone. Second, intermittent-demand methods and machine-learning models are expected to exhibit heterogeneous performance across SKU segments; no universal model advantage is proposed. Third, the use of probabilistic lead-time demand estimates is expected to improve the alignment between safety stock and differentiated service objectives, provided that forecast distributions are sufficiently calibrated. Fourth, criticality-differentiated replenishment policies are expected to reduce consequential stockout exposure without requiring uniform increases in portfolio inventory. Fifth, integrating slow-moving and obsolescence signals into review decisions is expected to reduce avoidable inventory investment, although premature disposal may increase lifecycle risk for legacy assets.

Empirical validation would require a longitudinal, SKU-location dataset containing dated issues and returns, receipts, purchase orders, lead times, supplier performance, unit costs, inventory balances, reservations, stockout and backorder records, work orders, asset criticality, and relevant operating or maintenance drivers. Validation would compare the framework with incumbent planning rules through rolling-origin backtesting and inventory-policy simulation. The backtest would evaluate forecast performance; the simulation would replay replenishment decisions under identical demand and lead-time histories, measuring service and cost outcomes. Parameters would be estimated on training windows and tested on later periods, with sensitivity analysis for lead-time disruption, demand surges, criticality weights, carrying-cost rates, and downtime-cost assumptions. If multiple facilities are available, validation should separate sites or asset classes to assess transferability.

Limitations

The framework's principal limitation is that proprietary oil and gas MRO data are difficult to access and often fragmented across ERP, CMMS, procurement, and operational systems. Transaction records may contain duplicate material codes, inconsistent units of measure, incomplete lead times, censored demand during stockouts, and unreliable links between spare parts and equipment. Criticality scores and downtime costs involve judgement and vary across assets and contexts. Consequently, the proposed architecture establishes a testable logic but does not demonstrate causal effectiveness or a quantified return on investment. Generalisability across U.S. upstream, midstream, and oilfield-services settings cannot be assumed: asset intensity, supplier markets, maintenance regimes, logistics geography, and operational risk differ materially. Future studies should therefore validate the framework across multiple operators and asset classes, disclose data-cleaning and parameter choices, and report both aggregate and segment-level results.

RESULTS

Overview of Analytical Results

The proposed integrated predictive inventory and demand-planning framework generates three principal categories of analytical outcomes: improved demand-forecasting performance across heterogeneous inventory segments, more differentiated inventory and service-level decisions, and potential improvements in working-capital efficiency. As the study is conceptual and does not employ primary company-level data, the results presented in this section constitute illustrative analytical outputs rather than empirically observed improvements. The numerical scenarios are intended to demonstrate how the framework's components could interact under plausible operating assumptions in the U.S. oil and gas industry.

The comparative analysis considers an illustrative portfolio of 1,000 MRO and spare-parts SKUs distributed across three demand categories: stable, intermittent, and highly variable or lumpy. The portfolio is assumed to have an average inventory investment of $50 million, of which $15 million represents slow-moving or potentially excess inventory. These figures are hypothetical and serve exclusively as a consistent reference scenario for examining the framework's financial and operational implications.

Comparative Forecast Accuracy Across Demand Segments

The first analytical outcome concerns the selection of forecasting methods according to demand characteristics. The proposed framework does not assume that machine-learning methods consistently outperform classical forecasting techniques. Rather, it anticipates that forecasting performance will depend on the statistical properties of individual demand series, the availability of explanatory variables, and the quality of historical records.

Table 1 presents an illustrative comparison between conventional planning approaches and forecasting methods selected through the proposed segmentation architecture.

Table 1. Illustrative comparative forecasting performance

Demand segment	**Conventional baseline MASE**	**Selected method MASE**	**RMSE index (baseline = 100)**	**Forecast bias**
Stable (X)	0.82	0.68	82	+3%
Intermittent (Y)	1.25	1.05	84	+5%
Highly variable/lumpy (Z)	1.48	1.27	86	+7%

Source: Illustrative analytical values; not empirical observations or outputs from a validated forecasting experiment.

The figures suggest a potential improvement in forecast accuracy across all three segments when model selection is aligned with demand characteristics. For stable-demand materials, exponential smoothing or ARIMA-based methods may provide relatively reliable forecasts without requiring extensive explanatory data. For intermittent materials, SBA or TSB methods may be more appropriate because they explicitly address irregular demand occurrence. For highly variable materials, gradient-boosting models may offer additional predictive value where equipment activity, maintenance schedules, and operational variables provide meaningful explanatory information.

The reductions in MASE shown in Table 1 represent illustrative improvements of approximately 17%, 16%, and 14% across the respective segments. The RMSE index similarly indicates lower illustrative forecast error relative to the conventional baseline. Nevertheless, these figures should not be interpreted as evidence of universal model superiority. In particular, machine-learning models may perform poorly when historical demand is sparse, explanatory variables are unreliable, or equipment-specific consumption patterns change substantially.

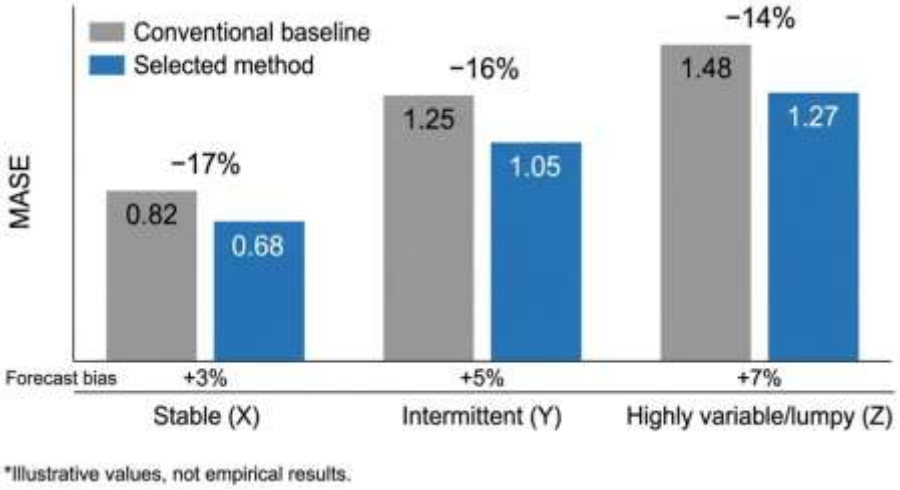


Figure: Illustrative reduction in MASE when forecasting methods are matched to demand segment, with bias shown beneath each group.

The positive bias values indicate modest over-forecasting under the convention adopted in the methodology. Although over-forecasting may reduce immediate stockout exposure, persistent positive bias can contribute to excess inventory and working-capital absorption. Consequently, forecast accuracy must be evaluated alongside bias and downstream inventory performance rather than treated as an independent measure of planning effectiveness.

Projected Inventory and Service-Level Effects

The second analytical outcome concerns the relationship between differentiated replenishment policies and inventory availability. A conventional siloed planning arrangement may employ relatively uniform reorder parameters, historical consumption averages, and purchasing-driven replenishment decisions. Such arrangements can produce insufficient protection for critical components while maintaining excessive quantities of low-priority materials.

Under the proposed framework, replenishment decisions incorporate forecast uncertainty, lead-time variability, operational criticality, and the economic consequences of non-availability. High-criticality spare parts receive appropriately elevated service targets and protective inventory arrangements, whereas low-criticality and readily substitutable materials are managed through more economical replenishment strategies.

Print ISSN: 2054-0930 (Print)

Online ISSN: 2054-0949 (Online)

Website: https://www.eajournals.org/

Publication of the European Centre for Research Training and Development-UK

Table 2. Illustrative comparison of inventory and operational performance

Performance indicator	Siloed baseline	Integrated framework	Illustrative change
Average inventory investment	$50.0m	$46.0m	−8.0%
Slow-moving inventory	$15.0m	$12.0m	−20.0%
Overall service level	92%	96%	+4 percentage points
Critical-spares service level	95%	98%	+3 percentage points
Stockout frequency index	100	75	−25.0%
Inventory turnover index	100	109	+9.0%

Source: Illustrative scenario outputs based on hypothetical assumptions; no field validation has been undertaken.

The comparison indicates that inventory reduction and service improvement need not be mutually exclusive when inventory is reallocated according to operational requirements. The illustrative reduction in average inventory from $50 million to $46 million represents an 8% decrease, while the overall service level increases from 92% to 96%. The assumed improvement in critical-spares availability is particularly relevant to capital-intensive operations, where the absence of a relatively inexpensive component may delay equipment restoration or interrupt production.

The reduction in the stockout frequency index also illustrates the potential contribution of improved forecasting and differentiated safety-stock policies. However, stockout frequency alone does not capture the severity of operational disruption. A shortage involving a non-critical consumable is fundamentally different from the unavailability of a component required to restore a major production asset. Accordingly, the framework prioritises stockout duration, equipment consequence, and potential production interruption alongside conventional service indicators.

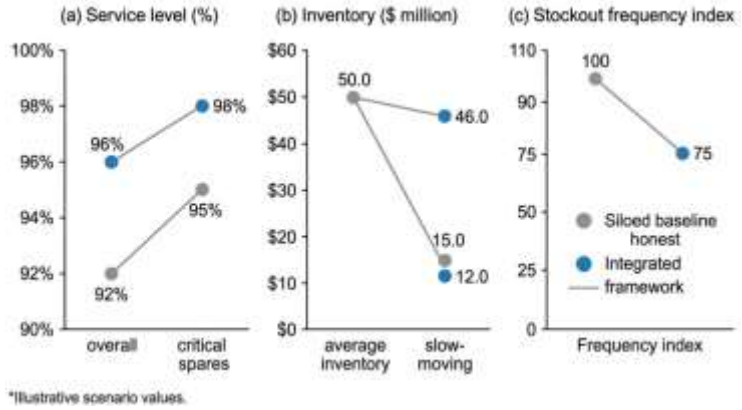


Figure: In the illustrative scenario, service levels rise and stockouts fall while inventory investment declines, because stock is reallocated by criticality.

The figure illustrates the framework's principal analytical sequence. Its central implication is that inventory efficiency emerges from coordinated decisions rather than from forecasting improvements alone. The resulting performance depends on the quality of information transferred between each layer and the extent to which inventory policies reflect actual operational consequences.

Working-Capital Implications and Slow-Moving Inventory

The third analytical outcome concerns the potential financial benefits of integrating predictive planning with inventory rationalisation. Slow-moving and excess spare parts represent a significant working-capital concern in industries where materials are purchased to protect against uncertain equipment failures, extended supplier lead times, and operational interruptions. Nevertheless, reducing slow-moving inventory indiscriminately would be inappropriate because some infrequently demanded components remain essential insurance items.

The proposed framework therefore distinguishes between genuinely excess materials and strategically necessary low-frequency inventory. Items identified for reduction are subject to criticality assessment, obsolescence review, repairability analysis, and substitution checks before disposal, redeployment, or procurement restriction is considered.

Under the illustrative scenario, slow-moving inventory declines from $15 million to $12 million, releasing $3 million in potentially recoverable working capital. This represents a 20% reduction in the designated slow-moving category. The additional reduction in average inventory from $50 million to $46 million includes this $3 million reduction and a further $1 million arising from improved replenishment and inventory positioning elsewhere in the portfolio. The two effects are therefore not additive.

Table 3. Illustrative working-capital and cost implications

Financial indicator	Siloed baseline	Integrated scenario
Average inventory	$50.0m	$46.0m
Slow-moving inventory	$15.0m	$12.0m
Capital released from slow-moving stock	—	$3.0m
Annual carrying-cost rate (assumed)	20%	20%
Annual carrying cost	$10.0m	$9.2m
Annual carrying-cost saving	—	$0.8m
Illustrative annual stockout-cost avoidance	—	$1.2m
Combined illustrative annual benefit	—	$2.0m

Source: Hypothetical scenario calculations; carrying-cost and stockout-cost assumptions are illustrative, not industry benchmarks.

The carrying-cost calculation applies an assumed annual rate of 20% to average inventory value. The resulting annual carrying cost decreases from $10 million to $9.2 million, generating an illustrative annual saving of $800,000. This estimate represents the financial cost associated with maintaining the inventory investment and should not be confused with the $3 million release of working capital, which is a reduction in the inventory asset rather than an annual operating saving.

A further illustrative benefit of $1.2 million is assigned to avoided stockout-related expenditure. This amount represents a hypothetical aggregate estimate incorporating emergency procurement premiums, expedited transportation, maintenance rescheduling, and potentially avoidable operational disruption. It is not derived from observed downtime records or verified operator cost data. In particular, the estimate would require considerable refinement where production losses, contractual penalties, or safety consequences are material.

The combined illustrative annual benefit of $2 million comprises carrying-cost savings and assumed stockout-cost avoidance. It excludes the released $3 million of capital to avoid treating a balance-sheet improvement as an additional annual profit. Similarly, no separate monetary benefit is assigned to the illustrative increase in inventory turnover, since doing so could introduce double counting.

Print ISSN: 2054-0930 (Print)

Online ISSN: 2054-0949 (Online)

Website: https://www.eajournals.org/

Publication of the European Centre for Research Training and Development-UK

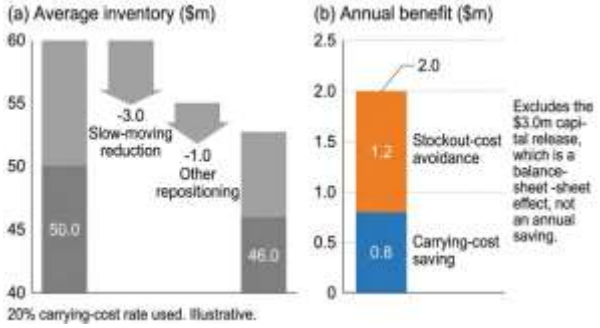


Figure: Illustrative inventory bridge and annual benefit, keeping one-off capital release separate from recurring cost savings.

Integrated Interpretation of Results

Taken together, the analytical outputs suggest that the principal contribution of the framework lies in coordinating forecasting, inventory control, and financial decision-making. Improved forecast accuracy can support more appropriate replenishment decisions, but its value depends on whether forecast information is translated into service-level and inventory parameters that reflect operational criticality.

The illustrative scenarios also indicate that slow-moving inventory reduction should be selective rather than universal. In oil and gas operations, a component with infrequent demand may nevertheless have substantial operational importance because of its relationship to equipment redundancy, replacement lead time, or production continuity. Consequently, working-capital optimisation requires distinguishing avoidable excess from strategically justified insurance inventory.

The framework also suggests that inventory turnover should be interpreted jointly with service performance. Higher turnover accompanied by deteriorating critical-spares availability would not constitute an unambiguous improvement. Conversely, a reduction in inventory investment accompanied by improved service performance would provide stronger evidence of enhanced inventory productivity, subject to validation of the underlying cost assumptions.

Evidential Status and Results Limitations

The numerical comparisons presented in Tables 1–3 are illustrative analytical outputs constructed to demonstrate the potential operation of the proposed framework. They are neither empirical findings nor results generated through a completed simulation using actual company transaction histories. The forecasting comparisons represent hypothetical performance differences, while the inventory and financial calculations are scenario-based estimates derived from explicitly stated assumptions. The conceptual relationships between demand segmentation, forecasting, replenishment, and working-capital management

are grounded in established operations-management principles, but their combined effectiveness remains to be empirically established.

Accordingly, no claim of field validation, statistically significant improvement, or demonstrated financial return is made. The projected benefits may vary considerably across U.S. upstream, midstream, and oilfield-services operations because of differences in equipment criticality, procurement lead times, supplier concentration, maintenance practices, and operational exposure. Future empirical research should test the framework using longitudinal SKU-level data from multiple operators, comparing integrated policies against existing planning arrangements through out-of-sample forecasting evaluation and inventory-policy simulation. Such validation would establish whether the projected improvements are achievable, economically sustainable, and generalisable across different operating environments.

DISCUSSION

Interpretation in Relation to the Literature

The proposed framework's principal contribution is not the substitution of one forecasting algorithm for another, but the integration of demand characterisation, forecasting, inventory control, and working-capital evaluation. This distinction is important in spare-parts management, where the demand process is frequently intermittent, lumpy, and driven by maintenance events rather than smooth consumption. Reviews of the field emphasise both the diversity of forecasting approaches and the need to assess methods in relation to inventory consequences, rather than relying exclusively on statistical error measures (Pinçe, Turrini, and Meissner, 2021). The framework extends this logic by making operational criticality an explicit decision variable alongside demand regularity and economic value.

The illustrative results suggest that segmentation-based model selection can improve forecast performance relative to a uniform baseline. This is consistent with research showing that intermittent-demand methods are not interchangeable and that the information contained in demand intervals, installed equipment populations, and maintenance activity can improve spare-parts forecasts (Pennings, van Dalen, and van der Laan, 2017; Van der Auweraer, Boute, and Syntetos, 2019). Machine learning is not inherently superior: intermittent methods may be preferable for sparse histories, while tree-based methods require reliable explanatory variables. The contribution is contingent model selection, not algorithmic novelty.

The integrated approach is expected to outperform single-function planning when the latter optimises a local objective while ignoring downstream consequences. Forecasting may minimise MASE or RMSE without considering lead-time exposure; procurement may reduce prices while increasing stockout risk; and inventory functions may apply uniform reductions to obsolete materials and essential insurance spares. By linking these decisions, the framework evaluates forecast usefulness through service and cost outcomes. This position is supported by research distinguishing forecast accuracy from forecast utility in inventory settings; improvements in statistical accuracy do not necessarily produce proportionate inventory benefits

(Ali, Boylan, and Syntetos, 2012). The literature also identifies a persistent gap between forecasting and inventory control, reinforcing the need to evaluate them jointly (Syntetos et al., 2021).

Why Integration Matters for Operational Reliability and Cost

The illustrative scenario indicates a possible reduction in average inventory alongside higher overall and critical-spares service levels. The mechanism is reallocation rather than indiscriminate reduction: inventory is protected where shortage consequences are high and challenged where stock is low-value, substitutable, or genuinely obsolete. In oil and gas, acquisition cost may bear little relationship to shortage consequence: an inexpensive component can constrain restoration when substitutes are unavailable, while a costly item may be managed through repair, exchange, or supplier support where redundancy exists.

The projected working-capital benefit is similarly conditional. The illustrative $3 million reduction in slow-moving stock represents potential capital release, whereas the $0.8 million annual carrying-cost saving is a recurring cost effect under the assumed carrying-cost rate. Neither should be interpreted as a realised financial return. Moreover, forecast accuracy will not translate linearly into working-capital savings: critical parts may retain high minimum protection levels, so better forecasts primarily reduce emergency procurement or stockout exposure. For low-criticality items, supplier minimums, long procurement cycles, and low residual values may constrain buffer reductions. Thus, the framework should be judged by a joint service-cost frontier, not by inventory reduction or forecast accuracy in isolation.

Managerial Implications and Adoption Sequence

Implementation should proceed in stages. First, operators should establish material-master integrity and a cross-functional classification process. ABC-XYZ segmentation should be combined with engineering criticality, substitutability, repairability, and lead-time exposure. This stage creates a defensible basis for deciding which items require advanced forecasting and which require policy protection regardless of forecast sophistication. Classification requires periodic review as assets, contracts, and decommissioning plans change.

Second, organisations should improve the demand signal before upgrading algorithms. Issues, returns, reservations, stockout requests, work-order consumption, purchase receipts, and lead-time records should be reconciled across ERP and CMMS environments. Stockout-censored demand must be distinguished from genuine zero demand; reliable maintenance plans, installed-base data, equipment utilisation, and project schedules should be incorporated. Forecasting upgrades should then be piloted by segment, comparing simple benchmarks, intermittent-demand methods, and machine-learning alternatives through rolling-origin testing. Forecast value added should be monitored so that manual overrides are retained only when they improve decisions; recent evidence finds that judgemental adjustments do not reliably improve accuracy across all SKUs (Fildes, Goodwin, and De Baets, 2025).

Third, forecasts should be embedded in replenishment policies rather than delivered as stand-alone dashboards. Inventory planners, maintenance engineers, procurement teams, and finance personnel should jointly set service targets, reorder points, review frequencies, and exception rules. High-criticality parts may require base-stock protection, repair-loop visibility, or contingency sourcing; lower-criticality items may be candidates for periodic review, supplier-held inventory, or rationalisation. Finally, working-capital governance should review inventory ageing and turnover together with fill rate, stockout duration, and critical-asset exposure. This sequence avoids automating unreliable data or forecasts disconnected from purchasing and maintenance decisions.

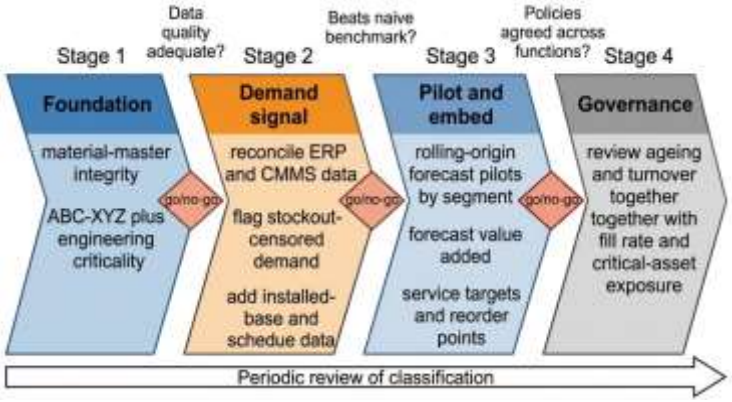


Figure: Operators should adopt the framework in stages, each gated by data quality and demonstrated decision value.

Boundary Conditions Across Operating Contexts

The framework's effects are unlikely to be uniform across operating environments. Offshore facilities generally face more constrained logistics, longer or weather-sensitive replenishment routes, limited storage capacity, and potentially severe consequences from equipment unavailability. These conditions strengthen the case for criticality-led protection and lead-time uncertainty modelling, although they may limit the scope for reducing insurance stock. Onshore operations may benefit from denser supplier networks, road access, and opportunities for lateral transfers between sites, making pooling and network-level inventory positioning more feasible. Remote onshore basins may face similar constraints, so geography should be modelled explicitly.

Operators and oilfield-services firms also face different demand structures. Operators typically plan around asset maintenance, production continuity, and long-lived installed equipment. Oilfield-services firms may experience demand linked to customer activity, contract awards, equipment fleets, project cycles, and geographically shifting deployments. Their forecasting models may therefore require customer- and activity-related variables, while inventory policies must account for redeployment and ownership arrangements. Capital-constrained firms may favour pooling, repair, consignment, and supplier

collaboration; capital-abundant firms may tolerate larger buffers, but remain exposed to obsolescence and lifecycle risk.

Implementation Barriers and Research Implications

The principal implementation barriers are data quality, system integration, and organisational fragmentation. ERP material masters may contain duplicate or obsolete codes, inconsistent units, and incomplete supplier lead times. CMMS work orders may not reliably identify the parts consumed or distinguish planned from corrective maintenance. Integration requires common identifiers, data ownership, interface governance, and historical-data controls. Organisational silos create a further obstacle: maintenance may prioritise availability, procurement may prioritise price and supplier compliance, and finance may prioritise inventory reduction. Without shared service-cost measures and clear decision rights, the framework may remain an analytical layer disconnected from operational action.

These limitations also define the research agenda. Future studies should validate the framework using longitudinal, multi-site SKU data and compare integrated policies with incumbent practices under identical demand and lead-time histories. Results should be disaggregated by criticality, demand segment, asset type, and operating context. The illustrative outputs in this paper demonstrate analytical plausibility only; they do not establish field effectiveness, causal impact, or generalisability. The framework's value will ultimately depend on whether organisations can sustain the required data discipline and cross-functional governance, and whether measured reliability gains justify the capital and implementation costs.

CONCLUSION

Unplanned downtime in capital-intensive U.S. oil and gas operations compels firms to hold costly critical spares, yet stockouts of consequential items and the accumulation of slow-moving stock persist because demand forecasting, criticality classification, and working-capital management are usually treated as separate problems. This paper has proposed an integrated predictive inventory and demand-planning framework that links demand segmentation, criticality assessment, probabilistic forecasting, and item-level replenishment policy to service and financial outcomes within a single decision architecture.

Four conclusions are defensible on the basis of the analysis presented. First, forecast accuracy is necessary but insufficient for sound inventory decisions: its value depends on whether forecast distributions are translated into lead-time- and criticality-sensitive stocking parameters, so forecasting methods should be judged by their service and cost consequences as well as by error measures. Second, no forecasting approach should be presumed superior across a heterogeneous MRO portfolio. Intermittent-demand estimators, classical time-series methods, and machine-learning models are each likely to suit different segments, and method selection should therefore be contingent on demand characteristics and data quality. Third, operational criticality, rather than acquisition cost alone, should govern service targets, because the consequence of an inexpensive component being unavailable can far exceed the cost of protecting it. Fourth, reductions in slow-moving inventory should be selective rather than uniform. The framework distinguishes

avoidable excess from strategically justified insurance stock, and inventory turnover should be interpreted jointly with critical-spares availability rather than in isolation. The illustrative scenarios in Section 4 show how these principles might interact, with lower inventory investment accompanying higher availability through reallocation rather than blanket reduction, but they rest on stated assumptions and should not be read as empirical effects.

The limitations of the study are material. The design is conceptual, the numerical outputs are hypothetical, the carrying-cost and stockout-cost assumptions are not industry benchmarks, and the framework has not been tested on operator transaction data. Any real implementation would also contend with data fragmented across ERP and CMMS systems, demand censored during stockouts, and judgement-based criticality scores. Generalisability across upstream, midstream, and oilfield-services settings cannot be assumed.

Future research should proceed along three lines. The first is empirical validation using longitudinal, multi-operator SKU-location data, comparing the framework with incumbent planning rules through rolling-origin backtesting and inventory-policy simulation, with results disaggregated by criticality, demand segment, and asset class. The second is extension to multi-echelon settings in which pooling, lateral transfers, repair loops, and network-level positioning across sites are modelled explicitly alongside forecast uncertainty. The third is integration with real-time condition-monitoring and predictive-maintenance data, which may offer earlier and more asset-specific signals of impending demand for critical spares than historical consumption alone, while partly alleviating the data scarcity that constrains conventional models. Pursuing these lines would establish whether the benefits identified here are achievable, economically sustainable, and transferable across U.S. operating environments.

REFERENCES

Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. International Journal of Production Economics, 71(1–3), 457–466.

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688.

Teunter, R. H., Syntetos, A. A., & Babai, M. Z. (2011). Intermittent demand: Linking forecasting to inventory obsolescence. European Journal of Operational Research, 214(3), 606–615.

Syntetos, A. A., & Boylan, J. E. (2006). On the stock control performance of intermittent demand estimators. International Journal of Production Economics, 103(1), 36–47.

Pinçe, Ç., Turrini, L., & Meissner, J. (2021). Intermittent demand forecasting for spare parts: A critical review. Omega, 105, 102513.

Pennings, C. L. P., van Dalen, J., & van der Laan, E. A. (2017). Exploiting elapsed time for managing intermittent demand for spare parts. European Journal of Operational Research, 258(3), 958–969.

Van der Auweraer, S., Boute, R. N., & Syntetos, A. A. (2019). Forecasting spare part demand with installed base information: A review. International Journal of Forecasting, 35(1), 181–196.

Ali, M. M., Boylan, J. E., & Syntetos, A. A. (2012). Forecast errors and inventory performance under forecast information sharing. International Journal of Forecasting, 28(4), 830–841.

Print ISSN: 2054-0930 (Print)

Online ISSN: 2054-0949 (Online)

Website: https://www.eajournals.org/

Publication of the European Centre for Research Training and Development-UK

Fildes, R., Goodwin, P., & De Baets, S. (2025). Forecast value added in demand planning. International Journal of Forecasting, 41(2), 649–669.

Syntetos et al. (2021). Inventory – forecasting: Mind the gap. European Journal of Operational Research.

Driessen, M. A., Arts, J. J., van Houtum, G.-J., Rustenburg, W. D., & Huisman, B. (2015). Maintenance spare parts planning and control: A framework for control and agenda for future research. Production Planning & Control, 26(5), 407–426.

Basten, R. J. I., & van Houtum, G.-J. (2014). System-oriented inventory models for spare parts. Surveys in Operations Research and Management Science, 19(1), 34–55.

Zhang, S., Huang, K., & Yuan, Y. (2021). Spare parts inventory management: A literature review. Sustainability, 13(5), Article 2460.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. Operational Research Quarterly, 23(3), 289–303.

Teunter, R. H., & Duncan, L. (2009). Forecasting intermittent demand: A comparative study. Journal of the Operational Research Society, 60(3), 321–329.

Hasni, M., Aguir, M. S., Babai, M. Z., & Jemai, Z. (2019). Spare parts demand forecasting: A review on bootstrapping methods. International Journal of Production Research, 57(15–16), 4791–4804.

Feizabadi, J. (2022). Machine learning demand forecasting and supply chain performance. International Journal of Logistics Research and Applications, 25(2), 119–142.

Kiymaz, H., Haque, S., & Choudhury, A. A. (2024). Working capital management and firm performance: A comparative analysis of developed and emerging economies. Borsa Istanbul Review, 24(3), 634–642.