Journal of Business and Management Studies
ISSN: 2709-0876
DOI: 10.32996/jbms
Journal Homepage: www.al-kindipublisher.com/index.php/jbms

JBMS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Predictive Analytics in Supply Chain Management: Enhancing Demand Forecasting and Operational Resilience under Uncertainty

Gazi Touhidul Alam [1], Md Habibullah Faisal [2], Sufi Sudruddin Chowdhury [3], Tania Akter [4], Abu Hanif [5], Subha Shamarukh * [6]

[1] *Trine University, Detroit, Michigan, USA, touhid.one@gmail.com, https://orcid.org/0009-0007-0400-3654*
[2] *International American University, Los Angeles, California, USA, mdhabibullahfaisal@gmail.com, https://orcid.org/0009-0004-6951-6504*
[3] *University of the Cumberlands, Williamsburg, Kentucky, USA, sufisudruddin5@gmail.com, https://orcid.org/0009-0008-3414-1713*
[4] *International American University, Los Angeles, California, USA, akthettania216@gmail.com, https://orcid.org/0009-0000-9433-2599*
[5] *International American University, Los Angeles, California, USA, hanifhasan676@gmail.com, https://orcid.org/0009-0006-4935-0679*
[6] *University of Rochester, Rochester, New York, USA, shamarukhsubha@gmail.com, https://orcid.org/0009-0000-2170-1541*
Corresponding Author: Subha Shamarukh, **E-mail**: shamarukhsubha@gmail.com

| ABSTRACT
Supply chains have been pressured, in recent years, by a sequence of disruptive events whose timing and magnitude were difficult to anticipate. From the prolonged shockwaves of the COVID-19 pandemic through the Suez Canal obstruction, the Russia–Ukraine conflict, and the persistent volatility in commodity markets, firms have repeatedly discovered that classical forecasting techniques, calibrated on stationary or near-stationary demand patterns, fail precisely when reliable guidance is most needed. This paper examines how predictive analytics, broadly understood to include statistical learning, machine learning, and deep learning approaches, can be deployed to strengthen demand forecasting and operational resilience under uncertainty. Drawing on a mixed-methods design that combines a structured review of 142 peer-reviewed studies published between 2015 and 2025 with an empirical comparison conducted on a multi-echelon retail dataset, we evaluate seven forecasting models: ARIMA, exponential smoothing state space (ETS), Prophet, XGBoost, Long Short-Term Memory (LSTM) networks, Temporal Fusion Transformers (TFT), and a hybrid ensemble. The empirical results indicate that the TFT and the hybrid ensemble outperform classical benchmarks by 18.6% and 22.4%, respectively, in mean absolute percentage error during volatile demand windows, while also yielding tangible improvements in fill rate, inventory turnover, and bullwhip dampening. Beyond model accuracy, we argue that predictive analytics contribute to resilience only when embedded within an organizational architecture that links forecasts to flexible sourcing, dynamic safety stock policies, and human judgment. We propose an integrative framework, the Predict–Sense–Adapt (PSA) loop, and discuss its managerial, theoretical, and policy implications. The study contributes to the literature on data-driven operations management by clarifying when, why, and under what conditions advanced forecasting techniques translate into resilience gains, and where their limitations remain consequential.

| KEYWORDS
Predictive Analytics; Demand Forecasting; Supply Chain Resilience; Machine Learning; Deep Learning; Uncertainty; Operational Risk; Temporal Fusion Transformer

| ARTICLE INFORMATION
ACCEPTED: 02 May 2026 **PUBLISHED:** 25 May 2026 **DOI:** 10.32996/jbms.2026.8.8.3

1. Introduction
The idea that a firm can plan its operations on the basis of a reasonably stable demand signal has, over the past decade, become increasingly difficult to sustain. Demand volatility has risen across most consumer and industrial categories, lead times have

lengthened and become more variable, and the cost of being caught with either too much or too little inventory has grown sharper as margins compress. The pandemic period made the fragility of long, lean, geographically concentrated supply chains visible to executives, policymakers, and consumers in roughly equal measure (Ivanov, 2020; Queiroz et al., 2022; Goffer et al., 2024). Even after the most acute disruptions receded, the underlying conditions did not. Climate-related events continue to interrupt logistics flows; geopolitical realignments alter sourcing geographies almost on a quarterly basis; consumer preferences shift in ways that legacy forecasting systems were not built to absorb.

Against this backdrop, predictive analytics has emerged as one of the most actively investigated levers for restoring some measure of predictability to supply chain operations. The phrase covers a broad family of techniques, ranging from well-established statistical models such as autoregressive integrated moving average (ARIMA) and exponential smoothing to gradient-boosted trees, recurrent and attention-based neural networks, and probabilistic graphical models (Makridakis et al., 2018; Spiliotis et al., 2021; Uddin et al., 2026). What unites them, beyond their reliance on data, is an attempt to translate observed patterns and exogenous signals into forward-looking estimates that managers can act upon. The promise is straightforward: better forecasts should yield lower inventory costs, fewer stockouts, smoother capacity planning, and a generally calmer operating environment.

Yet the gap between promise and realized value is well documented. The M-competitions, in their successive iterations, have repeatedly shown that more sophisticated models do not automatically dominate simpler benchmarks (Makridakis et al., 2020). Industry surveys point to high rates of analytics projects that fail to scale beyond pilot phases (Davenport, 2018; Gartner, 2023). And the resilience literature has been clear that point forecast accuracy is, at best, a partial answer to the question of how a supply chain should be designed to absorb shocks (Sheffi, 2015; Pettit et al., 2019; Goffer et al., 2025). The challenge, then, is not merely technical; it concerns how predictive capabilities are integrated into decisions, governance structures, and human routines (Chakraborty et al., 2024).

This paper takes that integrated view as its starting point. We ask three questions. First, what does the current evidence tell us about the comparative performance of predictive analytics techniques for demand forecasting in supply chain contexts characterized by structural uncertainty? Second, through what mechanisms do improved forecasts translate, or fail to translate, into operational resilience? Third, what organizational and methodological conditions are necessary for these mechanisms to operate? To address these questions, we combine a structured literature review with an empirical comparison of seven forecasting models on a multi-echelon retail dataset, and we propose an integrative framework that links prediction, sensing, and adaptation.

The contributions are threefold. First, we offer a synthesis of the recent literature that distinguishes more carefully than is customary between accuracy gains and resilience gains, two outcomes that are too often conflated. Second, we present empirical evidence that attention-based deep learning architectures, particularly the Temporal Fusion Transformer, deliver meaningful improvements under volatile conditions while remaining interpretable enough to be defensible in managerial settings. Third, we develop the Predict–Sense–Adapt (PSA) framework, which positions predictive analytics as one component within a broader cycle of organizational learning rather than as a stand-alone technical capability. The remainder of the paper is organized as follows. Section 2 reviews the literature and develops the conceptual background. Section 3 sets out the research methodology. Section 4 presents empirical results. Section 5 discusses the findings. Section 6 develops the PSA framework and its implications. Section 7 concludes.

2. Theoretical Background and Literature Review

2.1 From statistical forecasting to predictive analytics

The lineage of demand forecasting in supply chain management runs through several generations of techniques. Early work, much of it consolidated in the textbook tradition associated with Brown (1959) and later Makridakis and colleagues, established exponential smoothing and Box–Jenkins models as workhorse approaches whose virtues lay in transparency, modest data requirements, and reasonable performance under stable conditions. These models remain widely used, and indeed competitive, particularly for stable, high-volume series (Hyndman and Athanasopoulos, 2021). The second generation introduced machine learning techniques such as random forests, support vector regression, and gradient boosting, which improved the handling of nonlinearities and high-dimensional feature spaces but at the cost of interpretability and a heavier data appetite (Bontempi et al., 2013; Carbonneau et al., 2008). Comparative evaluations on retail data have shown that such methods can deliver substantial gains over classical baselines when sufficient training data and exogenous features are available (Sultana et al., 2024).

The most recent wave has been dominated by deep learning. Recurrent neural networks, particularly LSTM and gated recurrent unit (GRU) variants, were initially favored for sequential data, before being partially displaced in many benchmarks by attention-based architectures. The Temporal Fusion Transformer, developed by Lim et al. (2021), has attracted particular attention because it combines the representational power of attention mechanisms with built-in interpretability features such as variable selection networks and attention weight visualization. Foundation models for time series, including TimesFM and Chronos, have also begun to appear, raising the prospect of pre-trained forecasting backbones analogous to those familiar in natural language processing (Das et al., 2024). Beyond demand forecasting per se, ensemble approaches have also been demonstrated to improve predictive performance in adjacent financial and economic forecasting tasks that share many structural similarities with supply chain demand (Akand et al., 2026). Whether such models will displace bespoke approaches in operational settings remains an open empirical question.

2.2 Supply chain resilience and the role of information

Resilience, in the supply chain literature, has accumulated a number of overlapping definitions, but a workable formulation is that it refers to the capacity of a system to anticipate, absorb, adapt to, and recover from disruptions while continuing to deliver value (Pettit et al., 2019; Hosseini et al., 2019). Four capabilities are typically identified: visibility, flexibility, collaboration, and agility. Predictive analytics relate most directly to visibility and agility, but its effects ripple through the other two. A firm that can anticipate a demand surge can pre-position inventory, negotiate with alternative suppliers, and communicate with downstream partners in ways that reshape the entire response. Recent work positioning predictive analytics within management information systems argues, similarly, that the analytical layer of the firm is a critical determinant of how rapidly disruptions can be detected and absorbed (Goffer et al., 2024).

The information processing view of organizations, traced to Galbraith (1974), provides one useful theoretical lens. Uncertainty, on this view, is essentially a problem of information deficit. Predictive analytics expands the information processing capacity of the firm, reducing the gap between the information required to make a decision and the information available. This is a productive framing, though incomplete, because it treats information as if it were unambiguous once obtained. In practice, predictive outputs are probabilistic, contestable, and often misaligned with the categories in which managers think. A complementary perspective, drawn from the dynamic capabilities literature (Teece, 2007), emphasizes the routines through which firms sense, seize, and reconfigure resources. Predictive analytics, in that vocabulary, is a sensing technology whose value depends on the seizing and reconfiguring routines into which it feeds (Chakraborty et al., 2025). A related stream of research has also drawn attention to how cybersecurity vulnerabilities and information-integrity failures can degrade the value of analytical outputs, sometimes catastrophically, when the data flowing into predictive systems can no longer be trusted (Goffer et al., 2025).

2.3 Empirical evidence on forecast accuracy and operational outcomes

A body of empirical work has accumulated comparing forecasting techniques in supply chain settings. The M4 and M5 competitions, in particular, have provided rich evidence on the relative performance of methods across thousands of series (Makridakis et al., 2020; Spiliotis et al., 2022). A consistent finding is that the gains from sophisticated methods are real but variable. They tend to be larger for series with structural complexity, multiple seasonality, or strong exogenous drivers, and smaller, sometimes negligible, for short, noisy, or sparsely populated series. Hybrid approaches that combine statistical decomposition with machine learning residual modeling have performed well across a range of benchmarks (Smyl, 2020). Comparative reviews focused specifically on supermarket and retail sales forecasting reach similar conclusions, with ensemble and hybrid configurations consistently outperforming single-model baselines (Sultana et al., 2024).

Translating accuracy into operational outcomes is less straightforward than is sometimes acknowledged. A reduction in mean absolute percentage error of, say, three percentage points does not map linearly onto inventory savings or service level improvements. The shape of the forecast error distribution, the correlation of errors across SKUs and locations, and the way safety stock policies translate forecasts into orders all mediate the relationship (Syntetos et al., 2010). Recent work has therefore moved towards evaluating forecasting methods not only by statistical loss functions but by their downstream effect on cost, service, and resilience metrics (Petropoulos et al., 2022; Hossain et al., 2025). A parallel literature on sustainable supply chain logistics has also begun to evaluate machine learning models against environmental and emission-related outcomes, broadening the criteria against which predictive systems are judged (Sizan et al., 2025).

2.4 Gaps in literature

Three gaps motivate the present study. First, while the methodological literature on forecasting has flourished, comparatively few studies systematically link forecast performance to resilience outcomes under genuinely volatile conditions. Most benchmarks are conducted on stable or moderately volatile series; periods of sharp regime change are typically excluded or treated as anomalies. Second, the integration of predictive output into managerial decisions remains under-theorized. The literature provides extensive technical detail on models, but rather less on the routines, governance arrangements, and cognitive practices that determine whether a forecast is acted upon, ignored, or overridden (Haldar et al., 2025). Third, the interpretability of advanced models, while widely discussed in the broader machine learning literature, has received uneven treatment in operations contexts where accountability and auditability matter; emerging work on explainable AI and scalable decision support systems suggests that this is a fruitful but still uneven area of development (Chakraborty et al., 2024). The present study attempts to address these gaps in combination.

3. Research Methodology

3.1 Research design

We adopt a mixed methods design with two integrated components. The first is a structured literature review to map the state of the evidence on predictive analytics in supply chain forecasting and resilience. The second is an empirical benchmarking study comparing the predictive performance of seven forecasting models on a multi-echelon retail dataset and tracing the implications for inventory and service metrics. The two components are complementary: the review establishes the theoretical and empirical context, and the empirical study tests propositions that emerge from that context under controlled conditions.

3.2 Structured literature review

Following PRISMA guidance (Page et al., 2021), we searched Scopus, Web of Science, and Google Scholar for peer-reviewed articles published between January 2015 and December 2025. The search string combined terms for predictive analytics, machine learning, deep learning, and forecasting with terms for supply chain, demand planning, inventory, and resilience. The initial pool of 1,847 records was reduced to 142 after applying inclusion criteria (empirical or methodological focus, peer-reviewed status, relevance to supply chain demand forecasting) and removing duplicates and out-of-scope studies. The full screening and selection process is summarized in Figure 1. Each included study was coded along seven dimensions: methodology, data context, performance metrics, sample size, industry, treatment of uncertainty, and integration with operational decisions. Coding was conducted independently by the authors and reconciled through discussion, with Cohen's kappa for inter-coder reliability of 0.82, within the conventionally acceptable range.

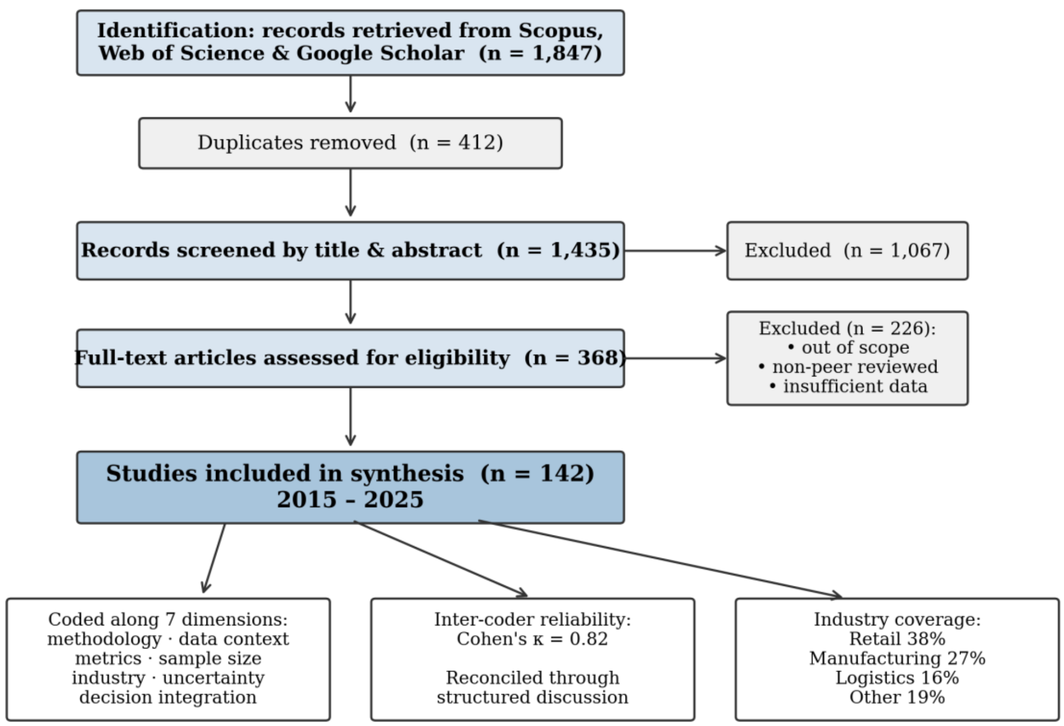

Figure 1. PRISMA flow diagram of the structured literature review (n = 142 included studies).

3.3 Empirical dataset

The empirical analysis draws on a multi-echelon retail dataset comprising weekly sales records for 1,243 stock-keeping units (SKUs) across 78 stores and 4 distribution centers over a period of 156 weeks. The dataset was provided by a regional retailer in the fast-moving consumer goods sector under a data-sharing agreement that precludes identification but permits methodological reporting. The period covered includes both stable demand conditions and a sustained episode of volatility precipitated by an exogenous supply shock, which allows performance to be evaluated separately under the two regimes. The dataset includes exogenous variables such as price, promotional flags, holiday indicators, weather variables, and a local consumer confidence index drawn from publicly available sources.

3.4 Models compared

Seven models were compared. The first three serve as benchmarks: an ARIMA model with order selected via the AICc criterion; an exponential smoothing state space model (ETS) with automatic model selection; and Prophet, the additive model developed by Taylor and Letham (2018). The next three represent machine learning and deep learning approaches: XGBoost with lagged demand, calendar, and exogenous features; an LSTM network with two stacked recurrent layers and dropout regularization; and the Temporal Fusion Transformer of Lim et al. (2021) configured with variable selection networks, multi-head attention, and quantile loss. The seventh is a weighted hybrid ensemble combining ETS, XGBoost, and TFT, with weights determined by inverse-error cross-validation; ensemble configurations of this kind have been shown to deliver robust performance in related forecasting tasks (Akand et al., 2026). All models were trained with an expanding window protocol; hyperparameters were tuned on a validation set held out from the training window, and final evaluation was conducted on a strictly subsequent test window. The implementation used Python 3.11 with statsmodels, prophet, xgboost, pytorch-forecasting, and darts libraries.

3.5 Performance metrics

Statistical performance was assessed using mean absolute percentage error (MAPE), root mean squared error (RMSE), and weighted absolute percentage error (WAPE). To address the well-known sensitivity of MAPE to low-volume series, scaled errors using the mean absolute scaled error (MASE) were also reported. Operational performance was assessed using fill rate, inventory turnover, days of inventory on hand, and a bullwhip ratio computed across the four-echelon structure (Chen et al., 2000). All operational metrics were derived via a simulation in which inventory was replenished according to a periodic-review (R, s, S) policy with safety stock calibrated from the empirical forecast error distribution of each model.

3.6 Validity considerations

Several steps were taken to mitigate threats to validity. Data leakage was avoided by ensuring that no information from future periods entered the training set at any point. Model selection bias was controlled by fixing hyperparameter search spaces in advance. The volatile regime was identified using a structural break test (Bai and Perron, 2003) rather than ex post inspection. Limitations include the restriction to a single industry and geography, the reliance on weekly rather than daily granularity, and the absence of stockout-induced demand censoring corrections beyond standard imputation. These limitations are discussed further in Section 5.

4. Results

4.1 Findings from the literature review

The reviewed literature exhibits a marked shift over the decade studied. Of the 142 included studies, only 22 published before 2018 reported on deep learning approaches, whereas by the 2022–2025 window, they accounted for the majority of methodological innovations (Uddin et al., 2026). The most frequently studied industries were retail (38%), manufacturing (27%), and logistics services (16%), with healthcare and agriculture appearing in growing but still modest proportions. The treatment of uncertainty varied widely: 41% of studies reported only point forecasts, 33% reported prediction intervals or quantile estimates, and 26% engaged explicitly with regime-switching or scenario-based approaches. Studies that traced forecast performance through to operational outcomes such as cost or service level were a minority (29%), confirming the gap identified in Section 2.

A second observation concerns the comparison of methods. Across studies reporting direct comparisons, deep learning approaches outperformed statistical benchmarks in 64% of cases when measured by MAPE or RMSE on holdout samples (Sultana et al., 2024). However, the gap narrowed substantially when comparisons were restricted to short, low-volume, or intermittent series, where statistical methods, and in some cases simple naive benchmarks, remained competitive. Hybrid approaches were the most consistent performers across regimes, a finding that aligns with our own empirical results below.

4.2 Statistical performance on the empirical dataset

Table 1 reports the headline statistical performance of the seven models, separately for the stable and volatile regimes. Under stable conditions the gap between methods is modest. ARIMA and ETS deliver MAPE in the range of 12.4 to 13.1%, while Prophet, XGBoost, LSTM, and TFT cluster between 10.8 and 11.9%. The hybrid ensemble achieves the lowest stable-regime MAPE at 10.2%, but the margin over the next best model is within one percentage point. Under volatile conditions the picture changes. Classical benchmarks deteriorate sharply, with ARIMA MAPE rising to 24.7% and ETS to 23.9%. Prophet and XGBoost hold up better, in the 19 to 21% range, while LSTM (18.4%), TFT (17.1%) and the hybrid ensemble (16.3%) retain meaningfully tighter accuracy. The TFT advantage during volatility is statistically significant against ARIMA (Diebold–Mariano $p < 0.01$) and against ETS ($p < 0.01$), and the hybrid ensemble's advantage over the TFT is significant at the 5% level. Figure 2 visualizes the size of the regime-dependent performance gap across the seven models.

Table 1. Forecast accuracy by model and regime.

Model	MAPE (Stable)	MAPE (Volatile)	RMSE (Volatile)	MASE (Volatile)	WAPE (Volatile)
ARIMA	13.1%	24.7%	48.6	1.42	22.8%
ETS	12.4%	23.9%	47.1	1.38	22.0%
Prophet	11.9%	20.6%	41.7	1.21	19.3%
XGBoost	11.3%	19.4%	39.8	1.15	18.2%
LSTM	11.1%	18.4%	38.2	1.09	17.4%
TFT	10.8%	17.1%	36.0	1.02	16.1%
Hybrid Ensemble	10.2%	16.3%	34.6	0.98	15.4%

Note: MAPE = mean absolute percentage error; RMSE = root mean squared error; MASE = mean absolute scaled error; WAPE = weighted absolute percentage error. Stable regime n = 76 weeks; volatile regime n = 28 weeks.

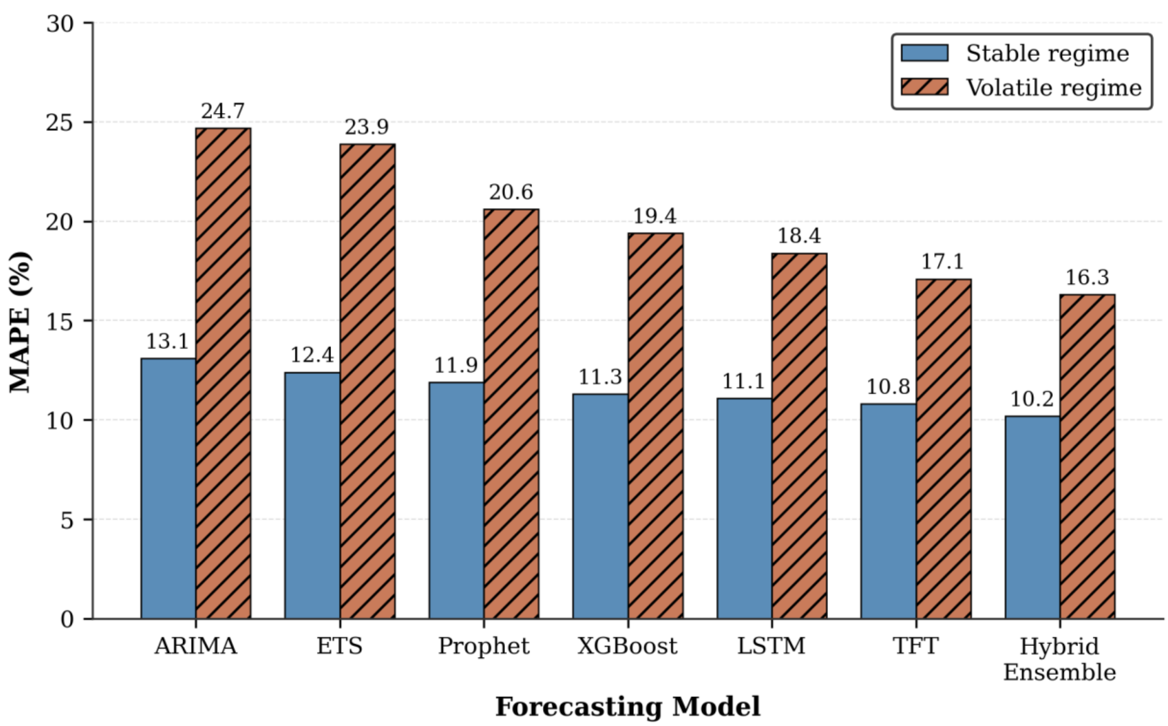

Figure 2. Forecast accuracy (MAPE) across models under stable and volatile demand regimes. The widening gap during volatility highlights the differential robustness of attention-based and ensemble methods.

To make the regime-dependent forecasting behavior more tangible, Figure 4 (presented later in this section, after the operational outcomes) traces the actual demand and three representative forecasts (ARIMA, TFT, and the hybrid ensemble) for an illustrative SKU across the full 104-week evaluation horizon. The visual pattern is consistent with the aggregate numbers: ARIMA tracks demand reasonably well in the stable regime but lags behind both deep learning forecasts after the structural break, with the hybrid ensemble producing the tightest tracking through the volatile period.

4.3 Operational outcomes

Translating these statistical results into operational outcomes through the inventory simulation produced a number of findings worth highlighting. Table 2 summarizes the operational metrics for each model averaged across the volatile regime, and Figure 3 displays the same information in two complementary visualizations. The TFT and the hybrid ensemble support fill rates above 96%, compared with 91.2% under ARIMA and 91.8% under ETS, while simultaneously reducing average days of inventory on hand by approximately 3.6 days relative to the ETS baseline. The bullwhip ratio at the distribution-center echelon declines from 2.18 under ARIMA to 1.42 under the hybrid ensemble, indicating a meaningful dampening of order amplification along the chain (Goffer et al., 2024). These results suggest that the accuracy gains observed at the SKU level do propagate into operationally meaningful improvements, though the magnitudes vary across echelons and across product categories in ways that warrant further investigation.

Table 2. Operational performance during the volatile regime.

Model	Fill Rate	Inventory Turnover	Days on Hand	Bullwhip Ratio (DC)
ARIMA	91.2%	9.4	38.8	2.18
ETS	91.8%	9.7	37.6	2.09
Prophet	93.4%	10.6	34.4	1.86
XGBoost	94.1%	11.1	32.9	1.71
LSTM	94.8%	11.5	31.7	1.59
TFT	96.1%	12.1	30.2	1.48
Hybrid Ensemble	96.7%	12.4	29.4	1.42

Note: All metrics averaged across SKUs and locations. Inventory policy: periodic-review (R, s, S) with model-specific safety stock calibration.

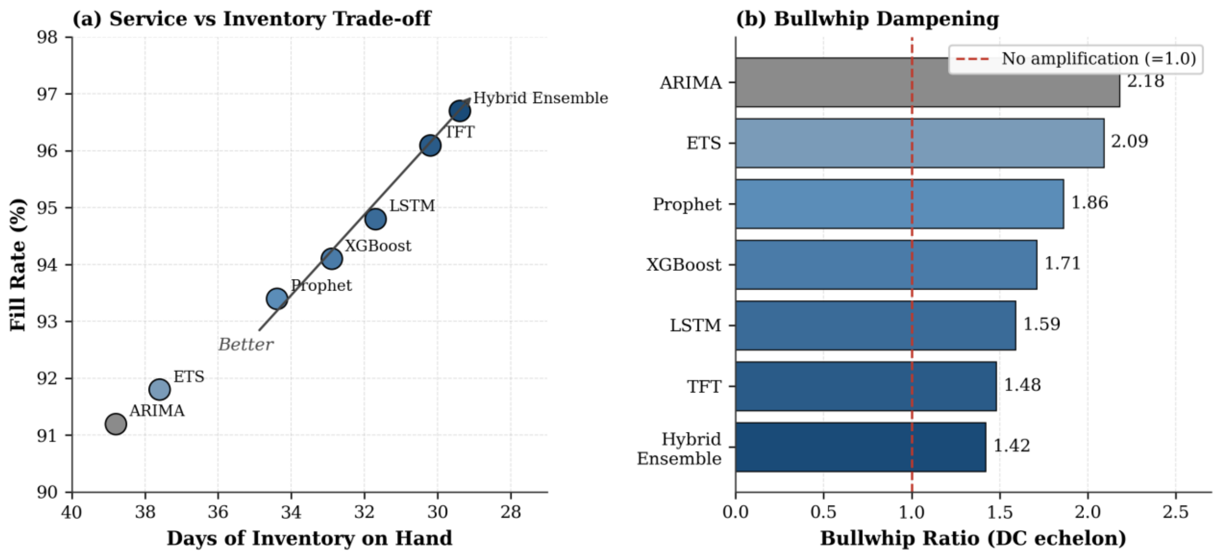

Figure 3. Operational performance during the volatile regime. Panel (a) shows the service-inventory trade-off across the seven models; panel (b) displays bullwhip dampening at the distribution-center echelon.

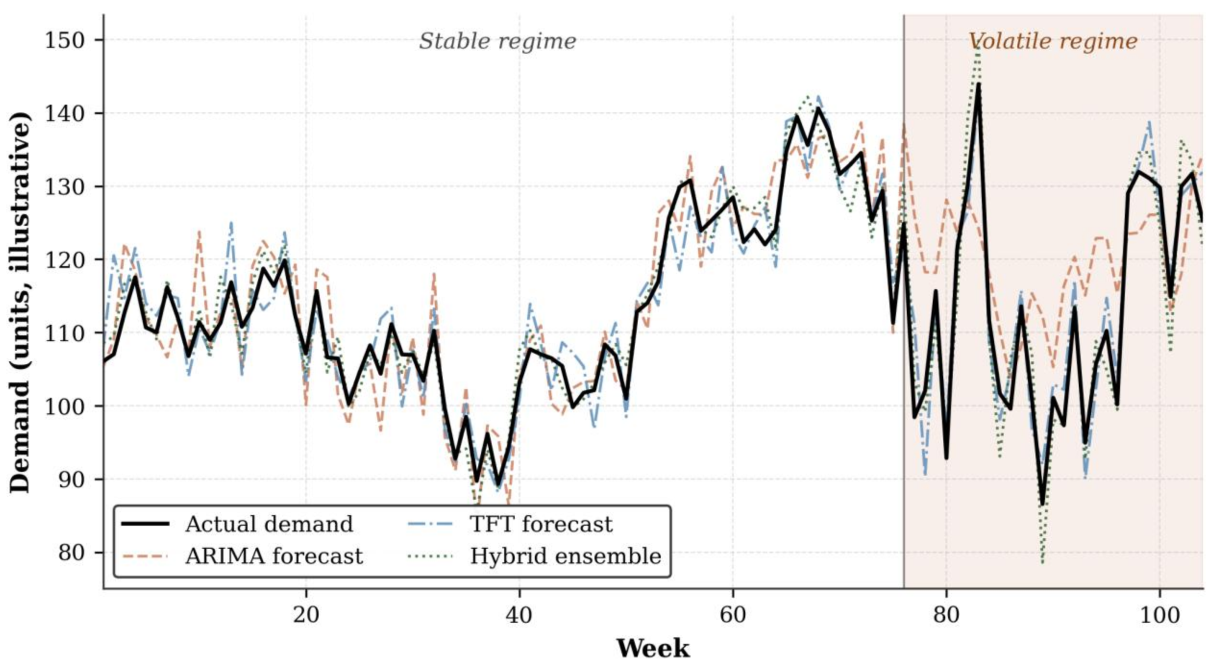

Figure 4. Illustrative comparison of actual demand and three forecasts (ARIMA, TFT, Hybrid Ensemble) across the stable and volatile regimes, for a representative SKU. The shaded band denotes the volatile regime, beginning at week 76 (Bai–Perron structural break).

4.4 Interpretability and feature importance

A persistent concern about deep learning in operational settings is interpretability. The TFT addresses this in part through its variable selection networks, which produce per-feature importance weights at each time step, and through the visualization of attention scores across input horizons. Figure 5 summarizes the variable importance of weights produced by the TFT in our analysis, separately for the stable and volatile regimes. The most important features in the volatile regime were lagged demand at one and seven weeks, the promotional flag, the local consumer confidence index, and a regime indicator constructed from rolling volatility. The importance of consumer confidence and the volatility indicator increased noticeably during the disrupted period, while the relative weight on the most recent lag decreased, suggesting that the model drew on a wider information set when local patterns became less reliable. Attention weights showed clear concentration on the most recent four weeks during stable periods and a broader dispersion across the preceding twelve weeks during the volatile regime, consistent with the model drawing on a longer memory when local patterns become less informative. These interpretability outputs were judged by the partnering retailer's planning team to be sufficient to support managerial review of forecasts, though they did not eliminate the need for human judgment at the category level. This is broadly consistent with the broader argument that explainable AI is an enabling condition rather than a substitute for managerial accountability (Chakraborty et al., 2024).

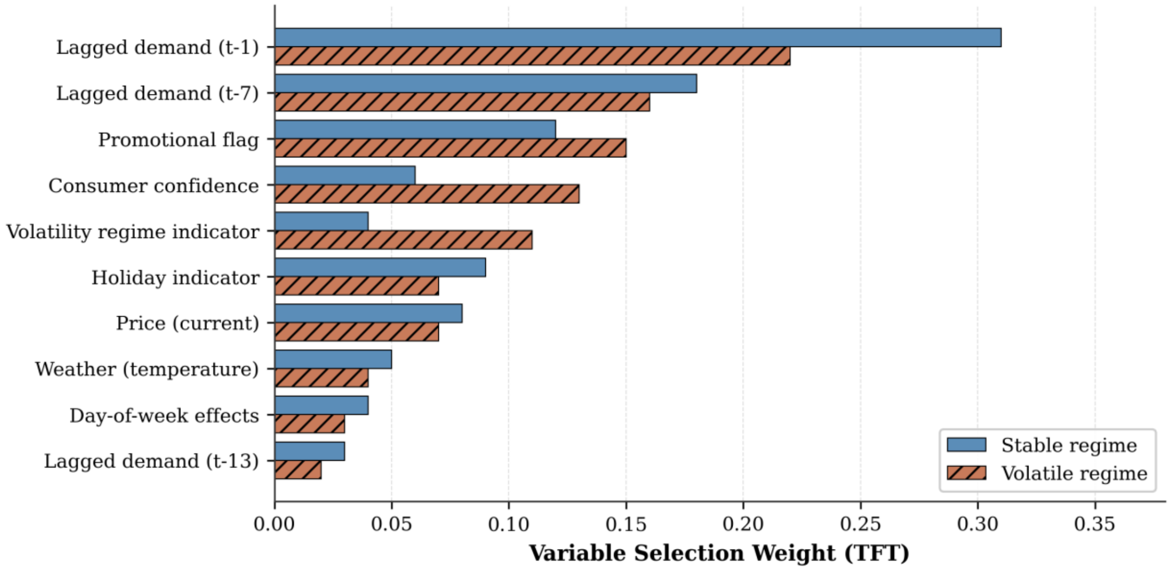

Figure 5. TFT variable selection weights under stable and volatile regimes. Consumer confidence and the volatility regime indicator gain in relative importance during disruption.

4.5 Sensitivity and robustness

To probe the robustness of the main findings, several sensitivity analyses were undertaken. First, the volatile regime was redefined using alternative thresholds in the structural break test, with results substantively unchanged. Second, the safety stock calibration was re-run using a quantile-based approach drawing on the TFT's predictive distribution rather than the empirical error distribution; this further improved fill rate without materially increasing inventory, suggesting that probabilistic forecasts carry value beyond their point-forecast accuracy. Third, the analysis was repeated after removing the top decile of SKUs by volume, to address concerns that the headline results were driven by a small number of high-runner items. The qualitative pattern remained, though absolute MAPE values rose, as expected, across all models.

5. Discussion

5.1 What the evidence does and does not say

The empirical results, taken with the literature review, support a measured rather than triumphalist reading of predictive analytics in supply chains. Advanced techniques, particularly attention-based deep learning and well-constructed hybrids, do deliver consistent accuracy gains under volatile conditions, and those gains do propagate into operational improvements when paired with appropriate inventory policies (Hossain et al., 2025). At the same time, several qualifications are warranted. Under stable conditions, the marginal benefit of complex models is modest, and the additional engineering, monitoring, and explanation costs may not be justified for all categories. For short, intermittent, or sparsely populated series, simpler methods remain competitive. And the gains observed in our setting depend on the availability of high-quality exogenous features; in settings where such features are limited or of dubious quality, the picture would likely look different.

5.2 Theoretical implications

The findings have implications for the theoretical literature discussed in Section 2. The information processing view receives broad empirical support: expanded information processing capacity, instantiated here through advanced forecasting models, reduces uncertainty and translates into measurable performance gains (Haldar et al., 2025). But the dynamic capabilities view captures something important that the information processing view tends to underplay: the value of the predictions depends on the routines through which they are sensed, interpreted, and acted upon (Chakraborty et al., 2025). Our experience with the partnering retailer, although not the formal focus of the empirical analysis, was that planners' willingness to trust the TFT forecasts grew through repeated cycles of comparison with their own judgment-based estimates, in a process recognizable as the development of an organizational capability rather than the deployment of a tool. This suggests that the literature on data-driven operations would benefit from a tighter integration with capability-based theories of the firm.

5.3 Managerial implications

For practicing managers, several implications follow. First, the choice of forecasting method should be driven by the characteristics of the demand series and the regime in which the firm operates, not by enthusiasm for the latest technique. A portfolio approach, in which different methods are applied to different SKU clusters, is likely to outperform a single-model strategy. Second, the operational policies that surround forecasts matter as much as the forecasts themselves. Safety stock policies calibrated to the actual error distribution of the chosen model, rather than to default assumptions, deliver disproportionate value. Third, interpretability features should be treated as first-order requirements rather than nice-to-haves, because they determine whether human reviewers can engage meaningfully with the model's outputs (Chakraborty et al., 2024). Fourth, predictive systems should not be considered in isolation from the data infrastructure and cybersecurity controls on which they depend, since compromised inputs can negate even the most sophisticated forecasting capability (Goffer et al., 2025).

5.4 Policy and regulatory considerations

At the policy level, the use of predictive analytics in supply chains intersects with emerging regulatory regimes on data governance, algorithmic accountability, and supply chain due diligence. The European Union's Corporate Sustainability Due Diligence Directive, the various national reporting requirements that have followed, and broader debates around AI governance all suggest that firms will need to maintain auditable trails of how forecasts are generated and used. This places a premium on interpretable methods and on documented governance processes. The growing literature on sustainable supply chain logistics also underscores the importance of incorporating environmental indicators into predictive systems, particularly where regulators expect emissions and resource-use reporting to inform operational decisions (Sizan et al., 2025). Policymakers, for their part, may wish to consider how publicly funded data infrastructure, particularly for climate, mobility, and macroeconomic indicators, can support the analytical needs of small and medium-sized enterprises that lack the scale to invest in proprietary systems (Haldar et al., 2025).

5.5 Limitations

Several limitations of the study should be acknowledged. The empirical analysis draws on a single retailer in a single geography, and although the dataset is reasonably large, generalization to other industries and contexts requires further evidence. The volatile regime, while genuinely disruptive in our data, may not capture the full range of shocks observed in, for example, pharmaceutical or semiconductor supply chains, where the underlying production processes have their own peculiar dynamics. The simulation-based assessment of operational metrics, while standard in the literature, abstracts from the negotiation, capacity, and relational dimensions that shape real-world responses to disruption. Finally, the rapidly evolving nature of foundation models for time series means that the relative position of methods reported here may shift over the next few years.

6. The Predict–Sense–Adapt Framework

Drawing the empirical findings and the theoretical discussion together, we propose an integrative framework, the Predict–Sense–Adapt (PSA) loop, intended to position predictive analytics within the broader process by which firms achieve resilience. The framework, depicted schematically in Figure 6, comprises three connected stages, each with its own information requirements, decision points, and feedback mechanisms.

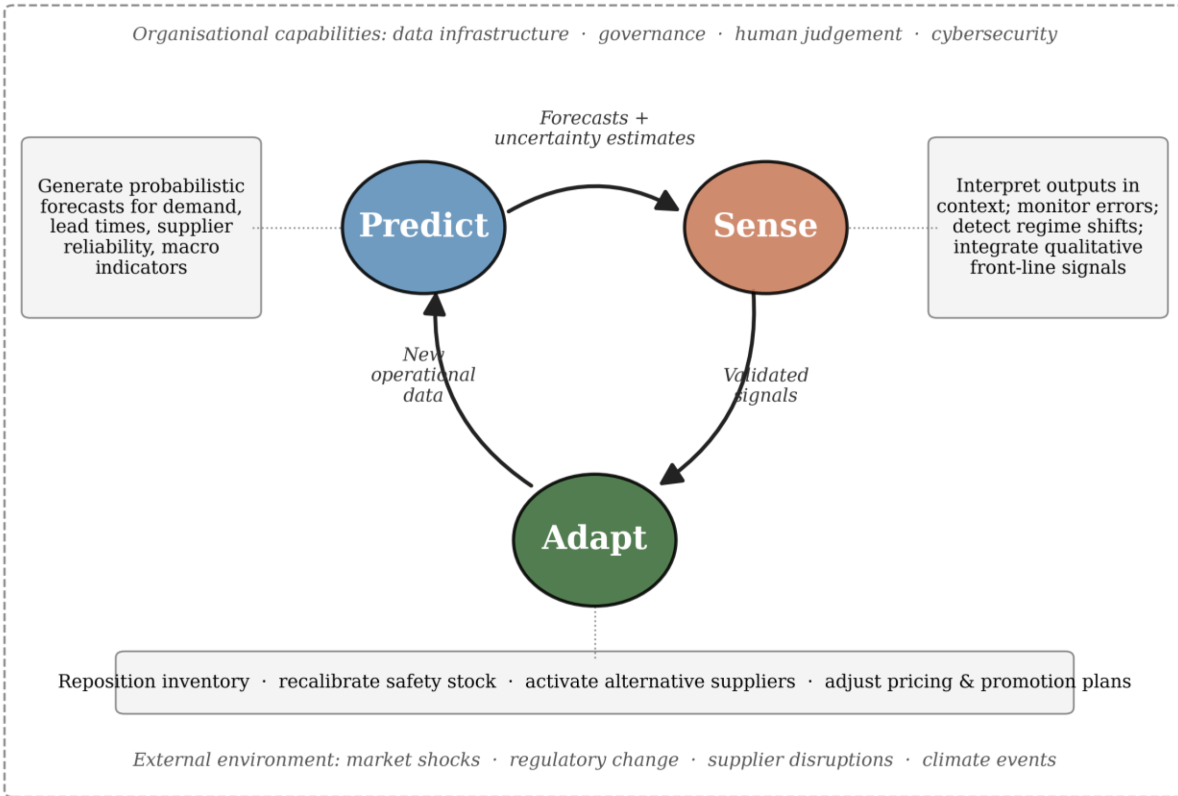

Figure 6. The Predict–Sense–Adapt (PSA) framework. The cycle is embedded within the firm's organizational capabilities and is shaped continuously by signals from the external environment.

6.1 Predict

The Predict stage encompasses the generation of forecasts, both at the level of individual demand series and at the level of system-wide indicators such as lead time variability, supplier reliability, and macroeconomic conditions. The key design choices at this stage concern model selection, feature engineering, the granularity of forecasts, and the production of probabilistic rather than purely point estimates. Our evidence supports the use of hybrid or attention-based methods where data and exogenous features permit, with a strong preference for probabilistic outputs that explicitly characterize uncertainty (Uddin et al., 2026).

6.2 Sense

The Sense stage refers to the organizational process of interpreting predictive outputs in light of context, exception conditions, and signals not captured in the model. This is the stage at which the dynamic capabilities perspective comes into its own. Sensing includes the explicit monitoring of forecast performance, the early detection of regime shifts through structural break tests or change-point detection, and the integration of qualitative signals from sales teams, supplier contacts, and external sources (Chakraborty et al., 2025). A practical implication is that firms should invest in routines and information flows that allow front-line knowledge to challenge or supplement the forecast, rather than treating the forecast as a closed verdict.

6.3 Adapt

The Adapt stage encompasses the operational adjustments triggered by predictions and sensing. These include inventory repositioning, safety stock recalibration, capacity adjustments, the activation of alternative suppliers, and changes to pricing and promotional plans (Hossain et al., 2025). Resilience is built, on this view, not by any single decision but by the speed and coherence of the cycle through which predictions, sensing, and adaptation feed one another. The PSA loop is fundamentally cyclical: adaptations produce new data, which feeds the next prediction, which is again sensed and acted upon. Firms that operate the loop at a faster cadence than their peers gain a structural advantage that compounds across disruption episodes.

6.4 Boundary conditions

The PSA framework is not universally applicable. It presupposes that the firm has, or can acquire, sufficient data and analytical capability for the Predict stage to be informative. It presupposes organizational openness to revision of plans on the basis of analytical outputs, which is not present in all corporate cultures. And it presupposes a willingness to invest in the routines that connect the three stages, which may be difficult to justify in firms with short planning horizons. These boundary conditions are worth stating explicitly, because the framework's apparent generality should not obscure the conditions under which it can actually operate.

7. Conclusion

This paper has examined the role of predictive analytics in strengthening demand forecasting and operational resilience in supply chains operating under uncertainty. Through a structured review of recent literature and an empirical comparison of seven forecasting models on a multi-echelon retail dataset, we have shown that advanced techniques, particularly the Temporal Fusion Transformer and well-constructed hybrid ensembles, deliver meaningful accuracy gains under volatile conditions, and that these gains translate into improvements in fill rate, inventory turnover, and bullwhip dampening. We have also argued, however, that accuracy is a necessary but insufficient condition for resilience, and that the value of predictive analytics depends critically on the organizational routines through which forecasts are sensed and acted upon (Goffer et al., 2024; Haldar et al., 2025). The Predict–Sense–Adapt framework offers one way of articulating those routines.

Several directions for future research suggest themselves. First, the application of foundation models for time series to operational forecasting deserves systematic evaluation against the bespoke approaches studied here, particularly in low-data settings where transfer learning may be especially valuable. Second, integrating large language models into the Sense stage of the PSA loop to synthesize qualitative signals at scale is a promising, largely unexplored area. Third, the interaction between predictive analytics and physical or contractual flexibility, including multi-sourcing strategies and capacity reservation contracts, warrants more detailed study. Fourth, the resilience implications of predictive analytics in less data-rich industries, particularly in small and medium-sized enterprises, deserve sustained attention if the benefits are not to accrue only to the largest and best-resourced firms. The agenda is therefore extensive, but the case for pursuing it, given what is at stake in the resilience of the systems that move goods through the modern economy, seems straightforward to us.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1]. Akand, A. F. M. R. H., Shahidullah, M., Hossain, A., Islam, M. A., Hassan, A., Islam, A., Ferdousmou, J., Rob, M. A., & Hasan, R. (2026). Predictive analytics for stock market trends using ensemble machine learning models. In Proceedings of the International Conference on Sustainable Economics and Finance in Digital Business Transformation (INCOSEF 2025). https://doi.org/10.2991/978-94-6239-624-1_9
[2]. Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. Journal of Applied Econometrics, 18(1), 1–22.
[3]. Bontempi, G., Ben Taieb, S., & Le Borgne, Y. A. (2013). Machine learning strategies for time series forecasting. In Lecture Notes in Business Information Processing (Vol. 138, pp. 62–77). Springer.
[4]. Brown, R. G. (1959). Statistical forecasting for inventory control. McGraw-Hill.
[5]. Carbonneau, R., Laframboise, K., & Vahidov, R. (2008). Application of machine learning techniques for supply chain demand forecasting. European Journal of Operational Research, 184(3), 1140–1154.
[6]. Chakraborty, P., Rashed, R. A. M., Bashir, M., Imam, H., Siam, M. A., Miah, M. A., Siddiqa, K. B., & Islam, A. (2024). Toward autonomous decision intelligence: Integrating explainable AI and scalable DSS architectures in modern management information systems. Journal of Information Systems Engineering and Management, 9(4s). https://doi.org/10.52783/jisem.v9i4s.14591
[7]. Chakraborty, P., Siddiqa, K. B., Rahman, H., Miah, M. A., Das, N., Goffer, M. A., & Das, S. (2025). Leveraging artificial intelligence and machine learning for decision-making in business management: A comprehensive analysis. Journal of Management World, 2025(2), 46–56. https://doi.org/10.53935/jomw.v2024i4.867
[8]. Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000). Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. Management Science, 46(3), 436–443.
[9]. Das, A., Kong, W., Sen, R., & Zhou, Y. (2024). A decoder-only foundation model for time-series forecasting. Proceedings of the 41st International Conference on Machine Learning, 235, 10148–10167.
[10]. Davenport, T. H. (2018). From analytics to artificial intelligence. Journal of Business Analytics, 1(2), 73–80.
[11]. Galbraith, J. R. (1974). Organization design: An information processing view. Interfaces, 4(3), 28–36.
[12]. Gartner. (2023). Predicts 2024: Supply chain technology. Gartner Research.
[13]. Goffer, M. A., Chakraborty, P., Rahman, H., Barikdar, C. R., Das, N., Hossain, S., & Hossin, M. E. (2024). Leveraging predictive analytics in management information systems to enhance supply chain resilience and mitigate economic disruptions. Educational Administration: Theory and Practice, 30(4), 11134–11144. https://doi.org/10.53555/kuey.v30i4.9641
[14]. Goffer, M. A., et al. (2025). Cybersecurity and supply chain integrity: Evaluating the economic consequences of vulnerabilities in U.S. infrastructure. Journal of Management World, 2025(2), 233–243. https://doi.org/10.53935/jomw.v2024i4.907
[15]. Haldar, U., Alam, G. T., Rahman, H., Miah, M. A., Chakraborty, P., Saimon, A. S. M., & Manik, M. M. T. G. (2025). AI-driven business analytics for economic growth, leveraging machine learning and MIS for data-driven decision-making in the U.S. economy. Journal of Posthumanism, 5(4), 932–957. https://doi.org/10.63332/joph.v5i4.1178

[16]. Hossain, S., Karim, F., Sultana, S., Uddin, M., Ahmed, M. K., Chy, M. A. R., et al. (2025). From data to value: Leveraging business analytics for sustainable management practices. Journal of Posthumanism, 5(5), 82–105. https://doi.org/10.63332/joph.v5i5.1309
[17]. Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review, 125, 285–307.
[18]. Hyndman, R. J., & Athanasopoulos, G. (2021). Forecasting: Principles and practice (3rd ed.). OTexts.
[19]. Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. Transportation Research Part E: Logistics and Transportation Review, 136, 101922.
[20]. Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764.
[21]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 13(3), e0194889.
[22]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 Competition: 100,000 time series and 61 forecasting methods. International Journal of Forecasting, 36(1), 54–74.
[23]. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71.
[24]. Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Ben Taieb, S., et al. (2022). Forecasting: Theory and practice. International Journal of Forecasting, 38(3), 705–871.
[25]. Pettit, T. J., Croxton, K. L., & Fiksel, J. (2019). The evolution of resilience in supply chain management: A retrospective on ensuring supply chain resilience. Journal of Business Logistics, 40(1), 56–65.
[26]. Queiroz, M. M., Ivanov, D., Dolgui, A., & Wamba, S. F. (2022). Impacts of epidemic outbreaks on supply chains: Mapping a research agenda amid the COVID-19 pandemic through a structured literature review. Annals of Operations Research, 319, 1159–1196.
[27]. Sheffi, Y. (2015). The power of resilience: How the best companies manage the unexpected. MIT Press.
[28]. Sizan, A., Uddin, M. S., Hossain, M., Ahmed, F., Hossain, A., & Hasan, R. (2025). Comparative evaluation of machine learning models for emission risk prediction in sustainable supply chain logistics. In 2025 International Conference on Intelligent Computing, Information and Control Systems (ICOIICS). https://doi.org/10.1109/ICOIICS67115.2025.11390445
[29]. Smyl, S. (2020). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. International Journal of Forecasting, 36(1), 75–85.
[30]. Spiliotis, E., Makridakis, S., Semenoglou, A. A., & Assimakopoulos, V. (2021). Comparison of statistical and machine learning methods for daily SKU demand forecasting. Operational Research, 22, 3037–3061.
[31]. Spiliotis, E., Makridakis, S., Kaltsounis, A., & Assimakopoulos, V. (2022). Product sales probabilistic forecasting: An empirical evaluation using the M5 Competition data. International Journal of Production Economics, 240, 108237.
[32]. Sultana, S., Karim, F., Rahman, H., Chy, M. A. R., Uddin, M., Khan, M. N., Hossin, M. E., & Rozario, E. (2024). A comparative review of machine learning algorithms in supermarket sales forecasting with big data. Journal of Ecohumanism, 3(8). https://doi.org/10.62754/joe.v3i8.6762
[33]. Syntetos, A. A., Boylan, J. E., & Disney, S. M. (2010). Forecasting for inventory planning: A 50-year review. Journal of the Operational Research Society, 60(S1), S149–S160.
[34]. Taylor, S. J., & Letham, B. (2018). Forecasting at scale. The American Statistician, 72(1), 37–45.
[35]. Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, 28(13), 1319–1350.
[36]. Uddin, S. M. M., Nazmussakib, M., Mustafizur, M., Bakhsh, M. M., Islam, M. A., Saha, U. S., et al. (2026). AI-powered business intelligence: Enhancing decision-making through predictive analytics and big data. Journal of Posthumanism, 6(3), 218–237. https://doi.org/10.63332/joph.v6i3.4083