Journal of Business Administration Online
ARKANSAS TECH UNIVERSITY College of Business HTTPS://JBAO.ATU.EDU

Spring 2026, Vol. 20 No. 1

The Adaptive Edge: Navigating Shifting Intermittent Demand with Structural Shifts

Matthew Lindsey, Stephen F. Austin State University

Robert Pavur, The University of North Texas

ABSTRACT

This article addresses the critical "Forecaster's Dilemma" within supply chain management: the strategic conflict between stability and responsiveness. Specifically, it explores how traditional, static forecasting models which are often embedded as defaults in Enterprise Resource Planning systems might fail to account for shifts in products characterized by intermittent demand, leading to significant financial losses through either stockouts or inventory obsolescence. The research proposes and evaluates a move toward adaptive forecasting, specifically utilizing the Adaptive Croston Procedure. This methodology functions as a mathematical logic gate that monitors Interval Deviation which is the variance between expected and actual arrival times of orders. Unlike traditional methods that utilize a fixed smoothing constant, the adaptive system dynamically adjusts its learning speed when it detects a structural shift in the Time Between Demands. The performance of this model was validated through 500 replications of a 900-period simulation, comparing Root Mean Square Error across Single Exponential Smoothing, standard Croston, and Adaptive Croston models under varying demand probabilities. The simulation data reveals that the Adaptive Croston method provides its most significant utility during severe demand transitions. Specifically, the model achieved up to a 28% reduction in forecasting error during the transition to a slow-demand state. The findings demonstrate that while low-alpha systems provide necessary stability during steady-state growth, they remain blind to the onset of obsolescence. Conversely, the adaptive model successfully identifies widening gaps in demand cycles, triggering a smoothing override that aligns inventory levels with the new market reality faster than manual human intervention or static arithmetic means. The article concludes that distributional awareness is a prerequisite for modern inventory resilience. In intermittent markets, where data is sparse, and zeros dominate the dataset, managers must shift from a reactive posture to a proactive, structurally aware strategy. The research culminates in the Five Pillars of Adaptive Inventory, a practitioner framework designed to bridge the gap between high-level liquidity goals and daily warehouse operations. By automating the detection of state changes, firms can significantly mitigate the risk of dead capital and ghost demand, transforming sparse data into a sustainable competitive advantage in increasingly volatile sectors such as consumer electronics, automotive legacy parts, and medical technology.

Spring 2026, Vol. 20 No. 1

INTRODUCTION

In the landscape of modern commerce, Big Data is often touted as the ultimate solution for supply chain efficiency. However, for many managers, the real challenge isn't an abundance of data but its scarcity. This is the realm of intermittent demand, where zeros are the most common data point and a single sale can represent weeks or months of revenue.

Items characterized by intermittent demand, such as high-value aircraft components, specialized medical hardware, or aging consumer electronics, carry a disproportionate amount of risk. As highlighted by Mishahi, Falatouri, Nasserj, Brandtner & Oplatkova (2026), modern supply chains are transitioning toward a Two-Step Framework for managing this uncertainty. This cutting-edge research suggests that practitioners must first identify the nature of the underlying demand state before selecting a forecasting tool. If a manager over forecasts because they failed to recognize a state shift, the company is left with dead capital and the looming threat of obsolescence. This article proposes a move toward Adaptive Forecasting, a hybrid-ready method that allows the math to evolve as quickly as the market does, fulfilling the 2026 mandate for more structurally aware inventory systems.

The Forecaster's Dilemma: Stability vs. Responsiveness

At the heart of every supply chain failure lies a fundamental conflict: the Forecaster's Dilemma. This is not merely a mathematical problem; it is a strategic and psychological battle between two competing goals: Stability and Responsiveness. In the context of intermittent demand, where data points are few and far between, this dilemma is magnified to a degree that can paralyze decision-making. Syntetos, Boylan and Croston (2005) used demand classification to help identify the optimal forecasting methodology.

The dilemma begins with the smoothing constant (Alpha, or a). This parameter dictates the learning speed of the forecast. A low Alpha values history; a high Alpha values the present.

The Case for Stability (The Anchor): Managers naturally crave stability. A stable forecast allows for predictable logistics, consistent labor scheduling, and lean warehouse operations. By setting a low Alpha (e.g., 0.05), the system assumes that any sudden spike in demand is likely a "random walk" that is a one-time fluke that should not disrupt the long-term plan. The danger here is Inertia. When the market truly shifts, a stable system is an anchor that prevents the ship from turning.

The Case for Responsiveness (The Pivot): In fast-moving sectors like tech or automotive, being late is often more expensive than being wrong. Managers in these fields lean toward responsiveness (High Alpha). They want to catch the wave of a new trend or the drop of a dying product. The danger here is nervousness. An over-responsive system hunts for trends in the dark, treating every random sale as a structural revolution.

The Psychological Trap of Stability

Stability is the anchor of supply chain management. When a product has been slow-moving for months, a manager's natural instinct is to trust the long-term history over a sudden, singular data point. They set their system to a low learning speed (Alpha, usually 0.05). This provides stability; the system will not jump or trigger a massive purchase order just because one random customer placed a small order.

However, stability is a double-edged sword. If that singular order is actually the first signal of a new trend such as a competitor's product failing or a sudden market resurgence that is a stable system will be blind to it for months. By the time the low-alpha forecast finally reflects the new reality, the company has already suffered through months of backorders and lost market share.

The Chaos of Over-Responsiveness

Conversely, the nervous forecaster sets the system to be highly responsive (High Alpha). This manager wants to catch every trend. But in intermittent environments, this leads to the Bullwhip Effect (Lee, Padmanabhan & Whang, 1997) Because intermittent demand is naturally lumpy, a high-alpha system mistakes random noise for a structural shift. It over-orders after a minor spike, only to be left with excess inventory when the demand inevitably returns to zero the following week. This creates a cycle of constant warehouse churn and inefficient logistics costs.

The Lifecycle Conflict

The dilemma is fundamentally tied to where a product sits in its journey. A fixed system is like a thermostat that only works at 70°F. It is perfect when the weather is mild, but is a disaster during a heatwave or a blizzard. During the Growth and Maturity phases, demand is relatively predictable, and stability is the goal. However, once a product enters the Decline phase, the intervals between sales widen unpredictably. (See Figure A) A static system, that is one that uses the same Alpha for the entire journey is structurally guaranteed to fail at the turn of the curve. It will either over-order at the peak of Maturity or hold too much stock during the slide into Obsolescence.

Visualizing Intermittent Volatility

To understand why standard math fails, a manager must visualize the difference between Smooth demand and Intermittent demand. Traditional models are often the default in large Enterprise Resource Planning (ERP) systems which assume a continuous Smooth Demand flow (as shown in Panel A below). This environment is predictable: every period has a sale, and updates happen consistently. Single exponential smoothing (SES) is a dependable technique to forecast when stable demand rates are present. Croston's (1972) variation of SES for intermittent demand cases remains the principal procedures for forecasting slow-moving items.

Spring 2026, Vol. 20 No. 1

Figure A

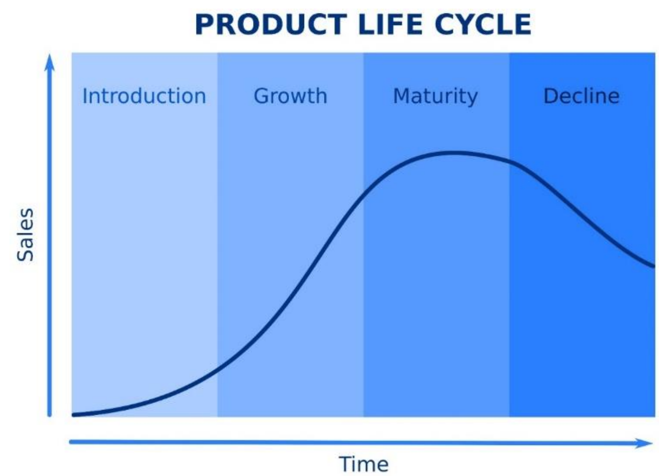

Advantages of Croston's Method are identified by Willemain, Smart, Shockor, and DeSautels (1994). Johnston and Boylan (1996), Syntetos and Boylan (2005) demonstrate the dominance of forecasting intermittent demand with Croston's method. Investigations of the impact of seasonality, promotions, irregular events, trends and correlated demands on the accuracy of forecasts have been conducted (Altay, Litteral and Rudisill, 2012; Lindsey and Pavur, 2008).

***Figure B Panel A: Smooth Demand (ERP Default*)**: Constant, stable, active demand in every time period (Periods 1 through 50).

Figure B Panel B: Intermittent/Lumpy Demand (Obsolescence Risk): Sparse, tall, volatile 'lumpy' demand events (Red Bars at Periods 4, 12, 25, 26, 48). The vast majority of the 50 periods show gray lines at exactly zero.

Figure B

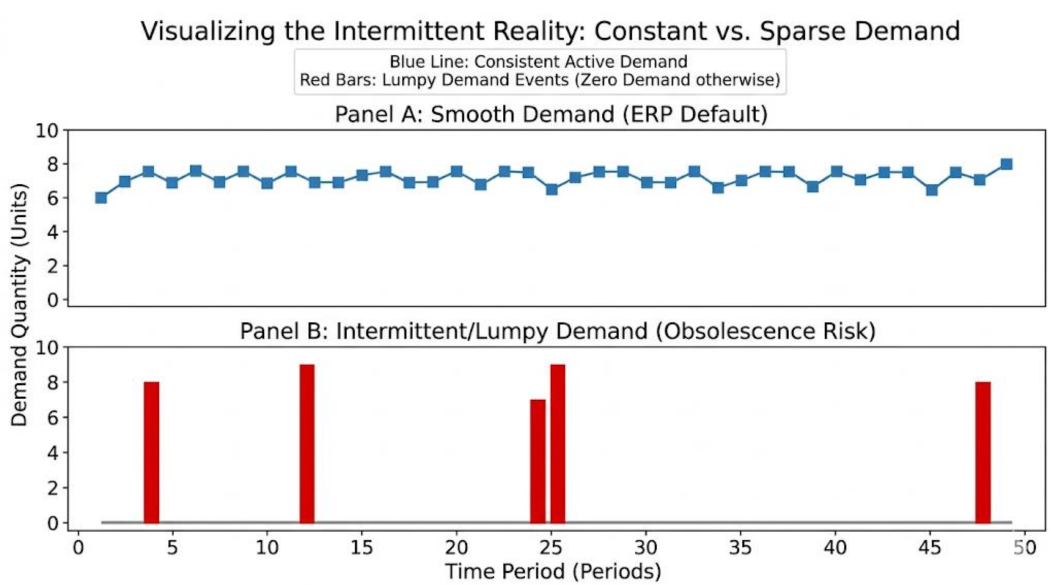

Standard forecasting looks at Panel B and lowers the mathematical "average" slightly for every empty period. An Adaptive System looks at the pattern of the gaps (the Time Between Demands, or TBD) themselves. It asks: "Is the rhythm of the market fundamentally changing?" If the rhythm of the empty spaces (the gaps between red bars) has structurally shifted, the system knows it must pivot its learning sensitivity immediately, rather than waiting for the arithmetic mean to slowly catch up.

Selecting the appropriate methodology requires that the practitioner select the correct demand pattern (Syntetos, Boylan, and Croston, 2005; Boylan, Syntetos and Karakostas, 2006). Only by appropriately categorizing demand can the best available forecasting method be selected (Syntetos, et al 2005, Rožanec, Fortuna, and Mladenić, 2022). Theoretically, an optimal technique exists for specific conditions, but the literature does not indicate what the measure should be to make the decision. Syntetos, Boylan and Croston (2005) segmented groups based on the time between demands and the coefficient of variation of the demand amount squared, but no consensus exists that it is the best methodology for classifying demand.

Assumed Distributions

Literature traditionally assumes a geometric distribution for the interval between positive demands (Croston, 1972; Segerstedt, 1994; Willemain et al., 1994; Bagchi et al., 1983; Levén and Segerstedt, 2004). Regarding demand size, the normal distribution is most common (Croston, 1972; Segerstedt, 1994; Willemain et al., 1994; Bagchi et al., 1983), though lognormal and Erlang distributions are also utilized (Willemain et al., 1994; Levén and Segerstedt, 2004). However, the geometric assumption falters during transitions between intermittent and regular demand, and its efficacy is unclear when demand rates decrease toward obsolescence.

Obsolete Inventory

Research into obsolescence within intermittent demand remains limited. Teunter, Syntetos, and Babai (2011) introduced specific inventory techniques for this scenario, while Babai, Syntetos, and Teunter (2014) analyzed the impact of smoothing constants. Further methods were presented by Babai et al. (2019), who noted that results were often conflicting and failed to account for intermittency, and non-stationarity. Pince and Dekker (2011) examined continuous review systems, while Van Jaarsveld and Dekker (2011) focused on risk determination. Sanguri, Patra, and Punia (2023) utilized structural connections across aggregation levels. Their findings indicate that while temporal aggregation generally improves intermittent forecasts, traditional Croston techniques underperform when seasonality or trends are present as in obsolescence.

Intermittent Demand Forecasting Method

The seminal work for forecasting intermittent demand is Croston (1972) who created two time series, one for demand and one for the time between demand. Traditionally, SES, a weighted moving-average technique that uses exponentially decreasing weights over time to forecast demand, is recommended for stationary time series. Willemain, et al (1994) should be consulted for a full explanation of Croston's method.

Proposed Croston Adaptive Method

This method integrates adaptive exponential smoothing into Croston's procedure to better handle non-stationary demand. Unlike standard Croston's, which uses a fixed smoothing constant, this approach adjusts the smoothing constant based on the variability of the time between demands. A larger alpha allows the model to react quickly to structural shifts in demand patterns, while a smaller alpha maintains stability during sporadic fluctuations.

The heuristic adjustment occurs only when a demand occurs:

- **Compute Interval Deviation (Diff_Q)**
  Calculate the interval deviation, denoted as Diff_Q, as the absolute difference between the most recent inter-demand interval and the preceding inter-demand interval. Formally,

  Diff_Q=$|Q_t - Q_{t-1}|$, where Qt represents the current time between demands and Q_{t-1} represents the previous interval.

- **Smooth the Interval Deviation (SM_Diff_Q)**
  Apply single exponential smoothing to Diff_Q using the same base smoothing constant employed in the Croston procedure. Denote the resulting smoothed value as SM_Diff_Q. Let Prev_SM_Diff_Q represent the most recent prior smoothed value available at the time of a demand.

- **Adjust the Smoothing Constant**
  Compute the absolute deviation between the current interval deviation and its smoothed historical value, namely, |Diff_Q−Prev_SM_Diff_Q|.

  Update the smoothing constant (α\alphaα) according to the following thresholds:

  - If |Diff_Q−Prev_SM_Diff_Q|≤1.5, retain the current smoothing constant.
  - If |Diff_Q−Prev_SM_Diff_Q|>1.5, set α=0.20.
  - If |Diff_Q−Prev_SM_Diff_Q|>2.0, set α=0.30.
  - If |Diff_Q−Prev_SM_Diff_Q|>3.0, set α=0.40.
  - If |Diff_Q−Prev_SM_Diff_Q|>4.0, set α=0.50.

These thresholds are evaluated sequentially, with larger deviations triggering higher smoothing constants to increase model responsiveness. The initialization of the system within conservative bounds (a = 0.08) is directly grounded in the empirical findings of the intermittent demand literature (e.g., Syntetos & Boylan 2005), which demonstrate that smoothing parameters are optimized on the lower end—typically between 0.05 and 0.20 to prevent excessive forecast variance in sparse data environments. By honoring this baseline constraint during steady states, the proposed logic only bypasses these literature-validated low constants when severe structural shifts emerge, scaling up to more responsive alpha levels (0.20 to 0.50) exclusively to form a rapid correction before dropping back to baseline levels.

While adaptive smoothing (Ekern, 1981) is conceptually intuitive, it involves a trade-off. It significantly enhances performance if the underlying demand distribution is truly shifting. But it may decrease performance if changes in demand intervals are merely noise rather than a fundamental shift in the distribution.

Detailed Mechanics: How the Adaptive System "Thinks"

The Adaptive Croston Procedure functions like a digital logic gate. It constantly monitors the "Interval Deviation" which is the difference between when an order should have arrived and when it actually arrived.

- **The Logic Gate: Interval Deviation**
  The system calculates the Interval error in real-time. For example:
  - **The Expected State:** Based on history, an order is expected every 2 weeks.
  - **The Observation:** The current order arrives after 10 weeks.
  - **The Decision:** The system calculates a deviation of 8 weeks.

Instead of merely averaging this 10-week gap into the old 2-week history, the adaptive system triggers a smoothing override. It recognizes that a deviation of this magnitude (5x the average) is unlikely to be a random fluke but it is likely a State Change. It temporarily raises the learning speed (Alpha) from 0.05 to 0.40. This allows the system to forget the old, high-volume data and align itself with the new, slower market reality. Once the intervals become consistent again at the new pace, the gate closes, and the system returns to its stable, low-alpha state.

Real-World Applications

To understand the power of an adaptive logic gate, it is helpful to categorize products by the specific type of change in state they undergo.

1. Consumer Electronics: Planned Obsolescence

Consider replacement screens for an older smartphone. As newer models launch, repair centers shift from weekly bulk orders to sporadic, individual requests. A standard system remains stuck in the weekly mindset for months, over-stocking parts that will eventually expire in a warehouse. The adaptive system identifies the widening gap in repair center orders within a few cycles and slashes the forecast, protecting the firm’s inventory investment.

2. Automotive and Industrial: Transition to Legacy Parts Obsolescence

Alternators for a decade-old truck model move from periodic dealership stock-up orders to sporadic independent mechanic shop purchases. The interval deviation trigger recognizes when these bulk orders stop. By pivoting to a high-alpha state, the system realizes the product has entered its long tail phase, allowing managers to move stock to lower cost storage facilities sooner.

3. Medical and Technology Supplies: Technology Upgrade Obsolescence

When a hospital upgrades to a newer robotic surgery system, demand for Version 1.0 sterile robotic surgery kits drop. Because these items have strict expiration dates, over-forecasting equals 100% loss of all products that were on hand when the upgrade occurred. The adaptive Croston method acts as a safety valve, preventing the automatic replenishment cycle that often plagues hospital and other supply chains during technology transitions.

Simulation Methodology

To evaluate the efficacy of the proposed adaptive method, a simulation study was designed to model environments where the distribution of the time between demands undergoes structural shifts. To facilitate a clear analysis, the experiment utilizes two distinct demand states: a fast demand state and a slow demand state. While the standard assumptions of the Croston method are maintained, this study introduces the variable of state transitions from fast to slow. The interval between demands is modeled using a geometric distribution with a fixed probability, while demand magnitude remains independent of timing and is modeled via a normal distribution.

The simulation experiment defines the fast demand state by a relatively high probability of occurrence, tested at levels of 0.90, 0.75, and 0.65. Conversely, the slow demand state is defined by a low probability of occurrence, fixed at 0.25. In the fast state, demand follows a normal distribution with a mean of 500 and a standard deviation of 100. The slow state reflects a diminished market presence, with a mean of 200 and a standard deviation of 20. Each cycle consists of 50 periods of fast demand followed by 100 periods of slow demand. To reflect real-world conditions where the timing of state changes is unknown to the practitioner, the forecasting methods were evaluated over 900 time periods, representing six full cycles. Each data set was replicated 1000 times to ensure statistical robustness.

Smoothing constants for the procedures were set at 0.08, 0.05, and 0.02. These conservative values align with the literature's recommendations for stationary data and provide a baseline for the system's stable state. Furthermore, using these low base values allows the adaptive logic to utilize significantly higher smoothing constants, such as 0.2 through 0.5, when the system detects changing conditions.

The study compares four forecasting methodologies. Single exponential smoothing and the standard Croston method serve as the primary benchmarks. The proposed adaptive approach is evaluated as Croston Adaptive, and its performance with the Syntetos Boylan Approximation bias correction is labeled as Croston Adaptive Bias_C. Table 1 and 2 summarizes the root mean square error for each method, measuring the accuracy in estimating the true mean demand across both fast and slow demand periods.

An acknowledged limitation of this simulation design is its exclusive focus on comparing the Adaptive Croston procedure against standard static benchmarks (SES and classical Croston). Prominent alternative risk-mitigation frameworks, most notably the TSB (Teunter, Syntetos, & Babai, 2011) method, offer robust protection against obsolescence by updating demand probabilities in every period, including those with zero demand. Because our study aimed specifically to isolate the behavior of an interval-driven mathematical logic gate, a direct empirical comparison with TSB was omitted from this computational run. This remains a critical limitation; future extensions of this research should deploy both the Adaptive Croston and TSB procedures within the same non-stationary simulation matrix to evaluate their relative trade-offs across various lifecycle decay trajectories.

Performance Analysis: When Does Adaptivity Win?

Through 1000 replications of a 900-period simulation, we measured the Root Mean Square Error (RMSE). The lower the RMSE, the closer the forecast is to reality, as shown in Table 1. Table 1 reports results for environments in which demand transitions from a fast state to a slow state. Across all parameter settings, several consistent patterns emerge.

First, the standard Croston's method exhibits the highest RMSE values in nearly all scenarios, indicating relatively poor performance when structural shifts are present. This is expected, as the method relies on a fixed smoothing parameter and cannot rapidly adjust to changing demand intervals.

Second, Single Exponential Smoothing (SES) performs better than Croston's method in most cases, particularly when the smoothing constant is relatively high (e.g., $\alpha = 0.08$). However, its performance deteriorates as the smoothing constant decreases, reflecting its inability to adapt efficiently to sudden changes in demand patterns.

Third, the proposed Croston Adaptive method consistently improves forecasting accuracy relative to standard Croston's method and, in many cases, also outperforms SES. The magnitude of improvement is most pronounced under more severe demand transitions (i.e., when the difference between fast and slow demand states is greatest). For example:

- At $\alpha = 0.05$ and fast demand probability = 0.90, RMSE decreases from 174.26 (Croston) to 122.58 (Adaptive), representing a substantial reduction in error.
- At $\alpha = 0.02$ and fast demand probability = 0.90, RMSE improves from 185.97 (Croston) to 147.48 (Adaptive).

These results indicate that adaptive adjustment of the smoothing constant allows the model to respond more effectively to widening intervals between demand occurrences.

Fourth, the bias-corrected version (Croston Adaptive Bias_C) performs nearly identically to the uncorrected adaptive method, with only marginal improvements in RMSE. This suggests that while bias correction stabilizes forecasts, the primary performance gains stem from the adaptive adjustment mechanism itself.

Overall, the results demonstrate that adaptive forecasting provides the greatest benefit in environments characterized by non-stationarity and structural shifts in demand. Table 2 presents results for scenarios without a transition to a slow demand state (i.e., no structural shift). In contrast to Table 1, the findings here show a different pattern.

In stable environments:

- Standard Croston's method consistently outperforms both SES and the adaptive variants across all parameter settings.
- The Croston Adaptive method exhibits significantly higher RMSE values, particularly at lower smoothing constants.

For example:

- At α = 0.02 and fast demand probability = 0.75, RMSE increases from 19.78 (Croston) to 86.98 (Adaptive).
- At α = 0.05 and fast demand probability = 0.90, RMSE increases from 26.33 (Croston) to 42.11 (Adaptive).

These results indicate that the adaptive mechanism introduces unnecessary variability when the underlying demand process is stable. In such cases, the adaptive model's responsiveness leads to over-adjustment, reducing forecast accuracy.

Table 1. Four forecasting methods' performance using root mean squared error over shifting demand

Smoothing constant	Slow Demand Prob	Fast Demand Prob	Single Exponential Smoothing	Croston	Croston Adaptive	Croston Adaptive Bias_C
0.08	0.25	0.90	**115.38**	159.18	116.77	116.77
0.08	0.25	0.75	**96.67**	133.71	101.63	101.42
0.08	0.25	0.65	**84.20**	115.67	91.85	91.45
0.05	0.25	0.90	138.37	174.26	**122.58**	122.65
0.05	0.25	0.75	113.97	144.08	107.92	**107.82**
0.05	0.25	0.65	**97.70**	123.33	98.75	98.50
0.02	0.25	0.90	175.19	185.97	**147.48**	147.66
0.02	0.25	0.75	142.82	151.80	**124.84**	124.90
0.02	0.25	0.65	121.26	128.95	113.18	**113.17**

Table 2. Four forecasting methods' performance using root mean squared error consistent with 'fast' intermittent demand

Smoothing constant	Slow Demand Prob	Fast Demand Prob	Single Exponential Smoothing	Croston	Croston Adaptive	Croston Adaptive Bias_C
0.08	No Slow Demand	0.90	36.10	**33.34**	41.60	42.74
0.08	No Slow Demand	0.75	47.43	**40.37**	57.41	58.20
0.08	No Slow Demand	0.65	51.25	**41.01**	60.21	62.31
0.05	No Slow Demand	0.90	28.23	**26.33**	42.11	43.02
0.05	No Slow Demand	0.75	37.14	**31.79**	62.56	64.38
0.05	No Slow Demand	0.65	40.09	**32.10**	66.00	67.14
0.02	No Slow Demand	0.90	17.49	**16.46**	55.86	56.46
0.02	No Slow Demand	0.75	23.06	**19.78**	86.98	88.00
0.02	No Slow Demand	0.65	24.83	**19.78**	86.21	87.20

The Strategic Takeaway: The Adaptive Croston method provided its greatest benefit that is a 28% reduction in error when the shift in demand was most severe. In business terms, this means the adaptive system is your best defense during the most dangerous parts of a product's lifecycle, specifically during the transition into obsolescence.

Intermittent demand is asymmetric. In a normal bell curve world, high and low errors cancel each other out. In the intermittent world, the zeros create a massive skew as shown in Figure C. Managers must move toward Distributional Awareness. This means recognizing that your risk is lopsided. An adaptive system helps re-center your distribution faster than a human planner can manually adjust ERP settings.

Figure C

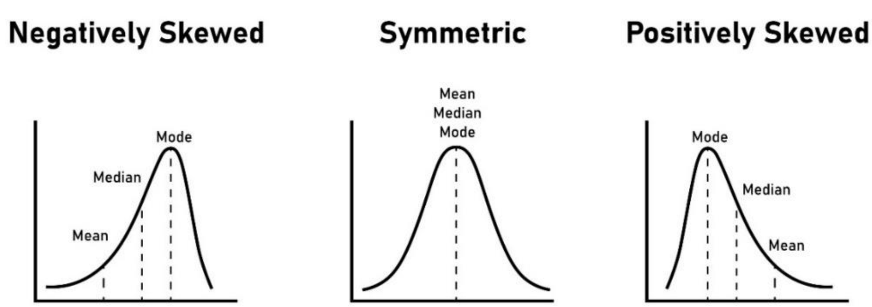

Practitioner Framework: The 5 Pillars of Adaptive Inventory

Implementing an adaptive forecasting system requires more than just a software update; it requires a shift in organizational strategy. For most firms, the primary barrier to efficient inventory management is not the math itself, but the rigidity of the process. In a static environment, inventory settings are often set and forgotten during the product launch, leaving the supply chain vulnerable as the product matures and eventually fades.

The following framework shown as the Five Pillars of Adaptive Inventory, provides a roadmap for managers to transition from a reactive emergency-order posture to a proactive, distributionally aware strategy. This framework is designed to bridge the gap between high-level executive goals like liquidity and service levels and daily warehouse operations. By using the Adaptive Croston logic as a change of state detector, managers can automate the most difficult part of the forecaster's job: knowing exactly when to pivot.

1. Category Audit

Segment your catalog. Stable items (like basic hardware) should stay on standard SBA. Volatile items (like tech components or seasonal parts) should move to adaptive logic.

2. The "0.05" Rule

Always start with a low base smoothing constant. This prevents the system from hunting for trends in noise. Adaptivity should be a reserve power that only kicks in during structural shifts.

3. Monitor "Gap Volatility"

Instruct your analysts to look at the Time Between Demands (TBD). If the standard deviation of your TBD is rising, your market state is shifting. This is the metric that precedes a financial loss.

4. Pair with Bias Correction

Always pair adaptive methods with the SBA (Syntetos-Boylan Approximation) correction to prevent the forecast from drifting too high during low-demand periods if shifting is unlikely.

5. Strategic De-stocking

Use the adaptive forecast as an early warning system. Highly adaptive Alphas, combined with widening gaps, are a signal to begin aggressive markdowns before the product hits the obsolescence cliff. While the Adaptive Croston method handles interval volatility, the future of the framework involves pairing these mathematical pivots with external signals, like Macro-Indexes, to catch obsolescence even faster (Singh, Sima & Agarwal, 2024).

Synthesizing the Framework: A Strategy for Resilience

The Framework for Adaptive Inventory is designed to provide a structured, evidence-based approach to a problem that many managers currently address with a "gut feeling." By combining categorization (Pillar 1) with specific mathematical triggers (Pillars 2-4), the framework removes the subjectivity from inventory planning.

For example, a manager at a medical device company might apply this framework to "Version A" replacement sensors as "Version B" launches. Instead of waiting for the warehouse to fill with unmoving Version A parts, the Adaptive System (Pillar 5) would detect the widening gap between orders within the first few months of the new launch, automatically lowering the forecast and triggering a reduction in safety stock. This transition ensures that the distributor remains liquid and responsive, rather than burdened by the "ghost demand" of a previous market state. In general, this framework moves the organization from a reactive posture to a proactive one, using math to bridge the gap between human intuition and warehouse reality.

While the Five Pillars of Adaptive Inventory framework provides a systematic roadmap for unpredictable components, its mathematical foundations limit its universal applicability. Practitioners must recognize two critical limitations where the framework, if not adjusted, may underperform:

1. **The Retail 'Phantom Stockout' Trap:** In consumer-facing retail environments, an extended string of zero-demand periods may not represent a structural market downshift or product decline. Instead, it may indicate a failure to restock or an unrecorded stockout. If the Adaptive Croston model is applied to such items, the widening interval gap will trigger a smoothing override, dropping the forecast to zero and permanently limiting replenishment orders. This creates an undesired situation where the retailer inadvertently surrenders market share to competitors. Therefore, the framework is explicitly restricted to environments where zero-demand is structural and demand is pull-driven (e.g., aerospace, legacy automotive parts, and repair parts).

2. **Interaction with Seasonality:** The current iteration of the Adaptive Croston procedure assumes non-seasonal, non-stationary demand states. In markets governed by predictable seasonal variations (like agricultural components), a widening or narrowing of the Time Between Demands may just reflect standard seasonality rather than obsolescence. If implemented without modifications, the system would introduce forecast shakiness during low-season dips. To tie this framework into seasonal operations, analysts must first deseasonalize the demand intervals prior to calculating the Interval Deviation or pair it with a seasonally adjusted aggregation layer as suggested by Sanguri et al. (2023).

The visualization below (Image D) maps the relationship between the actual demand patterns (from Fast to Slow) and the specific time points where each of the 5 Pillars must be applied for a resilient inventory strategy.

Image D

Mapping Lifecycle Volatility and Practitioner Interventions

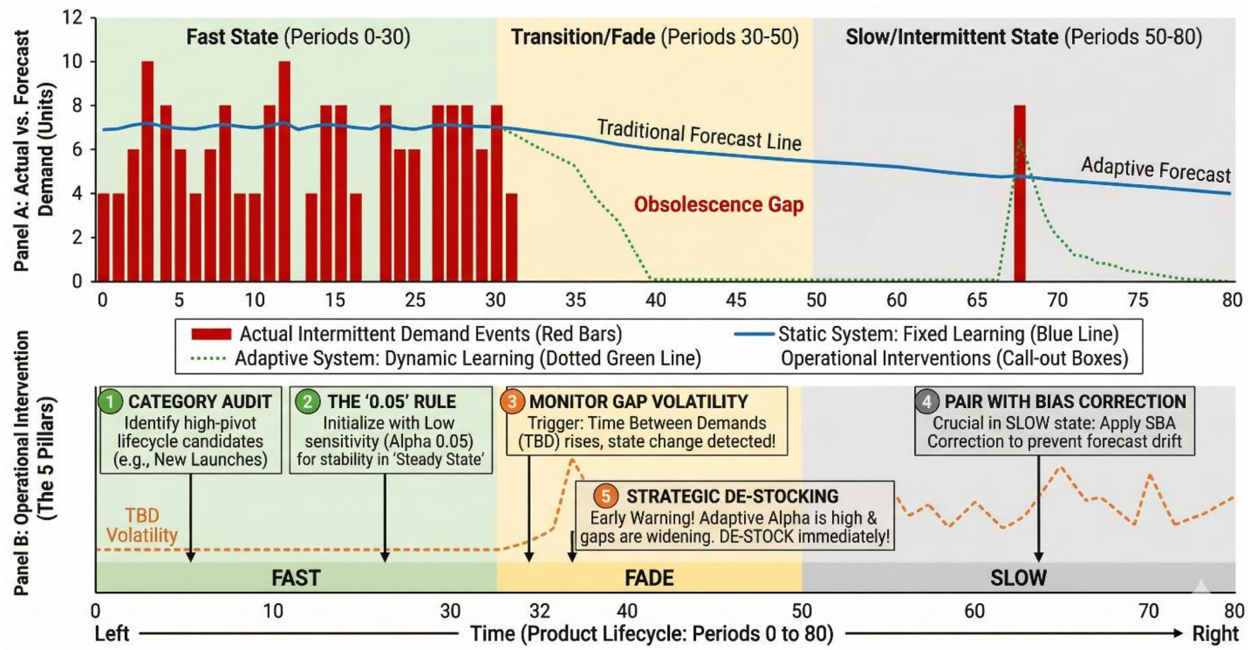

Top Panel (A: Fast to Slow Lifecycle Transitions): This panel demonstrates how a static forecasting system (the traditional fixed 0.05 learning speed) is fundamentally blind to state changes. The demand events (Red Bars) are frequent in the "Fast State" (e.g., Periods 0-30), creating a massive 'Obsolescence Gap' labeled in red once the "Transition/Fade" (Periods 30-50) occurs. The Adaptive forecast, in contrast, recognizes the growing "empty spaces" and drops to zero, engineering a soft landing. In the "Slow/Intermittent State" (Periods 50-80), demand is sparse, but handled responsively by the adaptive model.

Bottom Panel (B: When to Apply the 5 Pillars of the Practitioner Framework): This panel maps each specific intervention point:

1. **CATEGORY AUDIT (Periods 0-10):** Apply this during product launch/growth.

2. **THE '0.05' RULE (Periods 10-30):** The system must initialize and run at low sensitivity in the steady state.

3. **MONITOR GAP VOLATILITY (Period 32):** This is the trigger. The visualization shows a dashed orange line for TBD Volatility (Time Between Demands) spiking dramatically as the Fade state begins. This is when the state change is mathematically detected.

4. **STRATEGIC DE-STOCKING (Period 35):** This is the urgent intervention. Early warning of rising volatility requires immediate, aggressive markdowns.

5. **PAIR WITH BIAS CORRECTION (Periods 50-80):** In the sparse "Slow" state, bias correction prevents long-term drift in the forecast.

This lifecycle visualization provides a non-technical manager with the visual proof that precision in the average is useless without accuracy in the state. The adaptive approach bridges the gap between the predictable and the sporadic.

Conclusion

Traditional forecasting is like looking in the rearview mirror to steer a car. It works fine on a straight road, but when the market "turns," that is when a product moves toward obsolescence or a competitor launches a superior version, which is when you need a system that can turn the wheel.

As established by Mirshahi et al. (2026), the future of inventory management lies in hybrid analytics capable of distinguishing between temporary lulls and structural shifts. Adaptive forecasting provides this capability by aligning the model's learning speed with underlying market volatility. This alignment enables firms to reduce the risk of stockouts during growth phases while minimizing costly write-offs during periods of decline.

In dynamic environments, success is determined not by the volume of available data but by the ability to respond effectively to changing demand conditions. By adopting a distributionally aware approach, organizations can transform sparse, intermittent demand data into a more resilient and sustainable competitive advantage. Future research should seek to formalize the optimization of the alpha thresholds and evaluate the adaptive logic gate's performance alongside probability updating variations like the TSB model across diverse retail and industrial supply networks.

REFERENCES

Altay, N., Litteral, L. A., & Rudisill, F. (2012). Effects of correlation on intermittent demand forecasting and stock control. *International Journal of Production Economics*, 135(1), 275-283.

Babai, M. Z., Dallery, Y., Boubaker, S., & Kalai, R. (2019). A new method to forecast intermittent demand in the presence of inventory obsolescence. *International Journal of Production Economics*, 209, 30-41.

Babai, M. Z., Syntetos, A., & Teunter, R. (2014). Intermittent demand forecasting: An empirical study on accuracy and the risk of obsolescence. *International Journal of Production Economics*, 157, 212-219.

Bagchi, U., Hayya, J. C., & Ord, J. K. (1983). The Hermite distribution as a model of demand during lead time for slow-moving items. *Decision Sciences*, 14(4), 447-466.

Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2006). Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, 59, 473-481.

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

Ekern, S. (1981). Adaptive Exponential Smoothing Revisited. *Journal of the Operational Research Society*, 32, 775-782.

Johnston, F. R., & Boylan, J. E. (1996). Forecasting for items with intermittent demand. *Journal of the Operational Research Society*, 47, 113-121.

Lee, H. L., Padmanabhan, V., & Whang, S. (1997). Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science*, 43(4), 546-558.

Levén, E., & Segerstedt, A. (2004). Inventory control with a modified Croston procedure and Erlang distribution. *International Journal of Production Economics*, 90(3), 361-367.

Lindsey, M. D., & Pavur, R. (2008). A Comparison of Methods for Forecasting Intermittent Demand with Increasing or Decreasing Probability of Demand Occurrences. *Advances in Business and Management Forecasting*, 5, 115-132.

Mirshahi, S., Falatouri, T., Nasseri, M., Brandtner, P., & Oplatkova, Z. K. (2026). Two-Step Framework for Intermittent Demand Forecasting. *Procedia Computer Science*, *278*, 648-654.

Pinçe Ç., & Dekker, R. (2011). An inventory model for slow moving items subject to obsolescence. *European Journal of Operational Research*, 213(1), 83-95.

Rožanec, J. M., Fortuna, B., & Mladenić, D. (2022). Reframing demand forecasting: a two-fold approach for lumpy and intermittent demand. *Sustainability*, 14(15), 9295.

Sanguri, K., Patra, S., & Punia, S. (2023). Forecast reconciliation in the temporal hierarchy: Special case of intermittent demand with obsolescence. *Expert Systems with Applications*, *218,* 119566.

Segerstedt, A. (1994). Inventory control with variation in lead times, especially when demand is intermittent. *International Journal of Production Economics*, 35(1/3), 365-372.

Singh, A. K., Simha, J. B., & Agarwal, R. (2024). Prediction of Intermittent Demand Occurrence using Machine Learning. *EAI Endorsed Transactions on Internet of Things*, *10*.

Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, 21(2), 303-314.

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.

Van Jaarsveld, W., & Dekker, R. (2011). Estimating obsolescence risk from demand data to enhance inventory control—A case study. *International Journal of Production Economics*, 133(1), 423-431.

Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: A comparative evaluation of Croston's method. *International Journal of Forecasting*, 10(4), 529-538.