

## USE OF AI FOR PREDICTIVE INVENTORY AND DYNAMIC PRICING IN RETAIL & E-COMMERCE

Mrs. Megha D. Bichewar

Gopikabai Sitaram Gawande Mahavidyalaya, Umarkhed, Dist- Yavatmal  
meghabichewar123@gmail.com

### Abstract

Artificial intelligence (AI) has become central to two levers of retail performance: (1) **predictive inventory**—forecasting demand to place the right products in the right locations at the right time; and (2) **dynamic pricing**—setting prices that respond to demand, supply, and competitive conditions in near-real time. This article surveys practical AI methods (from gradient boosting to deep learning and reinforcement learning), data requirements, modeling workflows, governance and MLOps, and the interplay between inventory and pricing. It proposes an implementation roadmap, key performance indicators (KPIs), and risk controls. While empirical results vary by category and channel, mature programs typically reduce stockouts and overstocks, improve gross margin return on investment (GMROI), and accelerate inventory turns, provided they combine high-quality data, interpretable models, robust experimentation, and cross-functional ownership.

### 1) Introduction

Retailers and marketplaces operate under volatile demand, long and variable lead times, and intense competition. Traditional rule-based planning (e.g., fixed safety stocks, seasonal uplift factors, periodic price changes) struggles against today's granularity—SKU  $\times$  store/FC  $\times$  day—and the speed of online price wars. AI reshapes both **how much to buy/position** and **what to charge**, by:

- Learning nonlinear demand patterns at fine granularity.
- Updating decisions as new signals arrive (clicks, weather, events, competitor prices, logistics constraints).
- Optimizing against multiple objectives (e.g., profit, sell-through, service level, inventory risk).

Two questions guide this article:

1. Which AI methods work best for predictive inventory and dynamic pricing?
2. How should retailers implement them responsibly and profitably?

### 2) Predictive Inventory: Methods & Workflow

#### 2.1 Demand Forecasting Approaches

- **Classical baselines:** ARIMA/ETS, Croston (intermittent demand), hierarchical reconciliation.
- **Machine learning:** Gradient boosting (XGBoost/LightGBM, CatBoost), random forests for tabular features.
- **Deep learning:** Sequence models (LSTM/GRU), Temporal Convolutional Networks (TCN), Transformers for long-range seasonality and promotions.
- **Probabilistic forecasting:** Quantile regression, DeepAR/DeepState, TFT—outputs full demand distributions instead of point forecasts.
- **Causal/Promo models:** Uplift/causal forests, double ML to isolate the true effect

of promotions vs. seasonality.

- **Intermittent & new items:** Hierarchical Bayesian pooling, similarity-based cold-start (content embeddings), demand classification (zero-inflated models).

#### 2.2 Features (by signal family)

- **Calendar & seasonality:** day-of-week, holidays, payday, school terms, Ramadan/Diwali/Chinese New Year.
- **Price & promo:** own price, discounts, promo type (BOGO, coupons), display, ad spend, elasticities, price ladders.
- **Product content:** attributes, images/text embeddings (category, brand, pack size).
- **Competition:** scraped competitor price/stock signals.
- **Operations & supply:** lead time, supplier fill rate, receiving/put-away capacity, DC—store transfer cost.
- **Demand shocks:** weather, major events, social trends, returns/negative reviews.

#### 2.3 From Forecasts to Decisions

1. **Safety Stock (SS)** under service level target  $\alpha$ /alpha using a probabilistic forecast:  $SS = z_{\alpha} \sigma_L$  (demand std. during lead time L)
2. **(s, S) or (R, Q) policies** tuned from forecast distributions and holding/stockout costs.
3. **Multi-echelon optimization** for DCs, stores, and online FCs: determine where to hold buffer stock.
4. **Allocation & Replenishment:** daily solver uses:

$$\min \sum (c^{\text{hold}} \cdot I_i + c^{\text{oos}} \cdot E[\text{lost sales}_i] - m_i \cdot q_i) \text{ s.t. } \\ \text{capacity, budget, vendor MOQs}$$

5. **Network constraints:** pack-size rounding, truckload, shelf space, substitution and cannibalization.

## 2.4 Evaluation

- **Forecast:** WAPE/MAPE, RMSE, P50/P90 coverage, interval calibration, bias.
- **Inventory:** Service level (% lines in-stock), fill rate, **stockouts** (units & revenue lost)
- **overstock** (weeks of supply), **turns**, GMROI.
- **Operational:** purchase order adherence, supplier variability, DC congestion.

## 3) Dynamic Pricing: Methods & Workflow

### 3.1 Price Elasticity & Demand Models

- **Own/competitive elasticity** via log-log or semi-log demand models.
- **Hierarchical Bayes** to pool elasticities across SKUs/categories.
- **ML demand models:** GBMs/NNs with price as a feature, incorporating promos, seasonality, competition, availability.
- **Assortment interactions:** cross-elasticities for substitutes/complements (graph-based or multivariate demand).

### 3.2 Policies

- **Rule-based:** price ladders, MAP compliance, promo cadence.
- **Optimization:** constrained profit maximization with guardrails (floor/ceiling, price endings, competitor response).

- **Reinforcement learning (RL):** contextual bandits for A/B-like exploration; full RL for multi-period inventory+price coupling (learns state transitions with stock levels and lead times).
- **Markdown optimization:** determine price path to deadline (season end) maximizing sell-through & margin with remaining weeks and demand decay.

### 3.3 E-Commerce Specifics

- **High-frequency repricing:** respond to competitor changes, search ranking, and ad auctions.
- **Personalization** (non-discriminatory): tailor discounts or free-shipping thresholds by session intent while complying with fairness and legal constraints.
- **Marketplace dynamics:** Buy Box optimization (latency, stock, shipping speed, seller ratings).

### 3.4 Evaluation

- **Experimentation:** randomized price tests or geo-split to estimate causal lift; CUPED/variance reduction.
- **KPIs:** gross margin, profit per visitor (PPV), conversion rate, average selling price, win rate vs competitor, price perception (NPS/"value for money").

4) The Interplay: Joint Inventory–Price Optimization  
Inventory affects price (scarcity ↑) and price affects demand (hence inventory). Joint policies consider state

$s_t = \{I_t, P_{t-1}, \text{signals}\}$  and choose price  $P_t$  and order  $q_t$  to maximize discounted profit:

$$\max \mathbb{E} \left[ \sum_t \gamma^t (P_t \cdot Q_t(P_t, \mathbf{x}_t) - c \cdot q_t - h \cdot I_t - \text{stockout cost}) \right]$$

Solutions:

- **Stochastic dynamic programming** for small state spaces.
- **Approximate dynamic programming / RL** for realistic settings.
- **Two-stage heuristics:** inventory from P50/P90 forecasts + pricing from short-horizon elasticities; iterate daily.

## 5) Data, Architecture & MLOps

### 5.1 Data Foundation

- **Granularity:** SKU × location × day (or hour online).
- **Tables:** sales, inventory snapshots, purchase orders, transfers, promos, price

history, competitor prices, product catalog, web/app events, supply lead times, returns, reviews.

- **Quality:** correct backfills for stockouts (censoring), returns, and phantom inventory; de-dup campaigns; promo causal flags.
- **Feature store** with point-in-time joins to avoid leakage.

### 5.2 Modeling System

- **Pipelines:** Auto-regress & aggregate features, holiday calendars, weather APIs, embeddings for text/images.
- **Training & validation:** rolling origin splits; per-cluster models for long-tail SKUs.

- **Serving:** low-latency APIs for pricing; batch jobs for nightly forecasts; streaming updates on stock or competitor events.
- **Orchestration & monitoring:** data drift, forecast bias, price constraint violations, guardrails, rollback.
- **Explainability:** SHAP/feature attributions; elasticity dashboards.

### 5.3 Optimization Layer

- **Solvers:** MILP/LP for allocation; non-linear solvers for price; combinatorial search for markdown ladders; RL policies.
- **Constraints:** MAP/MSRP, legal limits, price endpoints (.99), step sizes, fairness rules, and vendor deal terms.

#### 6) Governance, Ethics & Compliance

- **Fair pricing:** avoid discriminatory outcomes by protected attributes; use cohort-level policies rather than per-person unless legally vetted.
- **Transparency:** clear promo rules, return policies.
- **Consumer protection:** price surge caps; disaster/emergency gouging restrictions; unit pricing disclosure.
- **Data privacy:** consent management, differential privacy or aggregation for sensitive signals.
- **Operational guardrails:** caps on daily price movement; human-in-the-loop approvals for sensitive categories.

#### 7) Implementation Roadmap (12–20 Weeks to First Value)

##### Phase 0: Alignment (1–2 wks)

- Use cases: replenishment for top categories; markdown for seasonal; competitive repricing online.
- Define KPIs: service level, OOS %, weeks of supply, gross margin, sell-through, PPV.
- Governance: pricing guardrails, data privacy, experimentation policy.

##### Phase 1: Data Readiness (2–4 wks)

- Build SKU–location–day panel; fix stockout censoring and returns; backfill promos & price.
- Create feature store with point-in-time joins; set quality alerts.

##### Phase 2: Modeling (4–6 wks)

- Baseline forecasts (boosting + quantiles); elasticity estimation per cluster.
- Pilot pricing optimizer with constraints; pilot replenishment with P90 service target.

##### Phase 3: Online Tests (4–6 wks)

- Geo-split or store-split for replenishment; A/B for price.

- Monitor bias/variance, safety stock impacts, cannibalization.

#### Phase 4: Scale & Integrate (ongoing)

- Expand SKUs/regions; add RL for coupled price–inventory decisions; integrate supplier collaboration (VMI, OTIF incentives).

8) Results: What Organizations Typically Observe (*Illustrative expectations; actual impact depends on category, seasonality, lead times, and competitive intensity.*)

- **Inventory health:** 10–30% reduction in overstocks, 20–40% reduction in stockouts for piloted categories; 0.5–2.0 turn improvement.
- **Financial:** 1–5% gross margin uplift on impacted SKUs from optimized price/promos; higher GMROI through better mix and markdown timing.
- **Operational:** smoother DC congestion, fewer emergency transfers; improved forecast bias and calibration.
- **Digital:** improved Buy Box win rate and PPV via responsive pricing and availability.

*How to validate:* require pre-registered test plans, holdout controls, and post-mortems; report effect sizes with confidence intervals, not just point estimates.

#### 9) Practical Patterns & Pitfalls

##### Patterns That Work

- **Cluster then specialize:** cluster SKUs by velocity/elasticity; long-tail uses pooled models; top SKUs get bespoke models.
- **Probabilistic**
  - **Guardrails before RL:** codify floors/ceilings, daily change caps, and fairness rules so learning can't breach policy.
  - **Joint dashboards:** price moves, forecast deltas, stock risk, and competitor signals in one pane.

##### Pitfalls to Avoid

- Treating **sales = demand** (censoring during stockouts).
- Ignoring **returns** and negative reviews (future demand decay).
- Over-fitting to short promo histories; failing to **separate causal promo lift** from seasonality.
- **Static lead times**—supplier variability must be modeled.
- Launching personalized pricing without legal review or customer trust measures.

#### 10) Summary & Conclusion

AI enables retailers to **predict demand more**

accurately and set prices more intelligently, especially when demand is volatile and competition is fast-moving. Predictive inventory uses probabilistic forecasting with features spanning calendar, price/promo, competition, and supply signals, then converts forecasts into replenishment and allocation decisions under multi-echelon constraints. Dynamic pricing estimates elasticities, optimizes within guardrails, and—where appropriate—employs bandits/RL for continuous learning. The **joint optimization** of price and inventory captures the feedback loop between availability and demand.

Sustainable success depends on a strong data foundation, interpretable models, disciplined experimentation, and governance that protects consumers and brand trust. With these capabilities, retailers typically see fewer stockouts and overstocks, faster turns, and healthier margins. The highest returns come from a **portfolio approach**: start with replenishment and markdown optimization, add constrained price optimization, then advance toward integrated RL that treats inventory and price as two sides of the same decision.

Appendix: Minimal KPI Scorecard (for pilots)

- **Forecast:** WAPE  $\leq 25\%$  (category-dependent), P90 coverage within  $\pm 5$  pp of target.
- **Inventory:** In-stock  $\geq 95\%$  on A-SKUs; weeks of supply within  $[\text{target} \pm 1]$ .
- **Pricing:** Profit per visitor  $\uparrow$ ; price change incidence within daily caps; MAP violations = 0.
- **Customer:** NPS/value-for-money steady or  $\uparrow$ ; complaint rate steady or  $\downarrow$ .
- **Ops:** OTIF from suppliers  $\uparrow$ ; emergency transfers  $\downarrow$ .

Appendix: Team & Roles

- **Product** (pricing & replenishment), **Data Engineering** (feature store), **Applied Science** (forecasting, elasticity, RL), **Operations** (planning, buying), **Legal/Compliance**, **Experimentation/Analytics**, **SRE/MLOps**.

## References

1. **Apte, M., Kale, K., Datar, P., & Deshmukh, P. (2024).** “Dynamic Retail Pricing via Q-Learning — A Reinforcement Learning Framework for Enhanced Revenue Management.” *arXiv, November 2024*. Shows simulated RL-based pricing yields superior revenue vs. traditional methods.
2. **Boston Consulting Group (2024).** “Overcoming Retail Complexity with AI-Powered Pricing.” Finds AI-driven pricing increases gross profit by 5–10%, improves revenue while maintaining customer value.
3. **Business Insider (2025).** Reports on Target, Walmart, Home Depot using AI for inventory forecasting—improving coverage, stock tracking, and regionally tailored stocking.
4. **Data-driven Dynamic Pricing & Inventory Management of an Omnichannel Retailer in an Uncertain Demand Environment.** *Computers & Industrial Engineering, 2023*. Surveys joint pricing and inventory decisions, topics like multi-period optimization, RL, substitution and perishable goods.
5. **Gallego, G. & Özer, Ö. (2001).** “Integrating Replenishment Decisions with Advance Demand Information.” *Management Science*. Examines integration of inventory replenishment with demand forecasting.
6. **Gallego, G., & Van Ryzin, G. (1994).** “Optimal Dynamic Pricing of Inventories with Stochastic Demand Over Finite Horizons.” *Management Science*. Classic work on inventory-sensitive dynamic pricing models.
7. **Grid Dynamics / Forrester (Date unspecified).** Survey: pricing and inventory optimization are top priorities; outlines challenges, needs, and readiness levels for AI adoption.
8. **Kopalle, P. K., Pauwels, K., Akella, L. Y., & Gangwar, M. (2023).** “Dynamic Pricing: Definition, Implications for Managers, and Future Research Directions.” *Journal of Retailing*. A comprehensive literature review on dynamic pricing, managerial implications, and future research paths.
9. **Le Monde (October 2024).** Discusses widespread adoption of dynamic pricing across industries (e-commerce, retail, energy), consumer backlash, and emerging regulatory scrutiny.
10. **Liu, J., Zhang, Y., Wang, X., Deng, Y., & Wu, X. (2019).** “Dynamic Pricing on E-commerce Platform with Deep Reinforcement Learning: A Field Experiment.” *arXiv, December 2019*. Field experiments on Alibaba’s Tmall show DRL methods outperform manual expert pricing.
11. **Selukar & Jain; Giannoccaro & Pontrandolfo; Dogan et al.; Rana & Oliveira; Kara & Dogan; De Moor. (Various years).** Research on perishable goods inventory



- i. and pricing using RL (Q-learning, Sarsa), plus CNN-LSTM hybrid forecasting with optimization via Grey Wolf Optimization (GWO), especially for fresh vegetables. *Symmetry (MDPI)*, 2024.
12. **StartUs Insights (2025)**. “AI in Retail: A Strategic Guide [2025–2030].” Highlights AI-enabled demand sensing (e.g., IKEA), dynamic pricing adoption rates, benefits of AI across forecasting and pricing.
13. **The Australian (2025 rollout)**. Coles implementing AI-based forecasting and automated ordering for liquor based on events, weather, promotions.
14. **The Sun (2025)**. Fashion retailers Boohoo and PrettyLittleThing implementing AI dynamic pricing—raises surge pricing concerns during peak demand events. *Foundational Theories & Models*
15. **Wang, Y., Chen, X., Chang, X., & Ge, D. (2020)**. “Uncertainty Quantification for Demand Prediction in Contextual Dynamic Pricing.” *arXiv*, March 2020.
  - i. Addresses constructing statistically valid confidence intervals for demand in dynamic pricing contexts.
16. **Yang, C., Feng, Y., & Whinston, A. (2021)**. “Dynamic Pricing and Information Disclosure for Fresh Produce: An Artificial Intelligence Approach.” *Production and Operations Management*. Explores AI-mediated dynamic pricing and supply of fresh produce.
17. **Zheng, Y., Li, Z., Jiang, P., & Peng, Y. (2024)**. “Dual-Agent Deep Reinforcement Learning for Dynamic Pricing and Replenishment.” *arXiv*, October 2024.
  - i. Proposes dual-agent DRL handling pricing and inventory across different decision time scales.