

Publication Date: 30.05.2025

Mehmet A. Begen¹

1. Western University, Ivey Business School, London, Canada Correspondence: mbegen@ivey.ca | ORCID: 0000-0001-7573-0882

Inventory and Pricing with AI Forecasting: Robust vs. Adaptive Policies

Abstract



Background: AI-driven demand forecasting expands the signal space for inventory and pricing decisions, enabling faster reactions to market changes. However, forecast error, non-stationarity, and distribution shifts raise a governance question: should decisions be designed to be robust to uncertainty, or adaptive to feedback?

Methods: This review integrates inventory control, probabilistic demand forecasting, dynamic pricing, and robust optimization into a unified decision architecture. We organize prior findings around a closed-loop cycle: data ingestion, forecasting (point and distribution), policy selection (robust/adaptive/hybrid), execution, monitoring, and recalibration.

Results: Robust policies protect against tail risk by optimizing over uncertainty sets and worst-case scenarios, but may be conservative and costly in stable environments. Adaptive policies leverage frequent feedback to improve average performance, yet can become unstable under regime changes, delayed signals, or strategic customer responses. The synthesis supports a hybrid design: adaptive learning within robust guardrails (service constraints, pricing move limits, and inventory safety floors).

Conclusions: The practical frontier is not “robust versus adaptive” as a binary choice. Best practice is layered: robust feasibility and risk limits at the outer layer, with adaptive learning tuned inside auditable constraints. Future research should prioritize regime-switching demand, decision-focused learning, and explainable pricing and replenishment rules.

Keywords: AI forecasting; inventory management; dynamic pricing; robust optimization; adaptive control; demand uncertainty; distribution shift; governance

1. Introduction

Inventory and pricing decisions sit at the intersection of operational feasibility and revenue capture. AI forecasting has increased the set of usable signals, including web traffic, marketing exposures, promotions, competitor actions, and macro indicators. Yet more signals do not automatically yield better decisions. The fundamental challenge is that **forecast outputs are uncertain inputs to a control problem**: the same prediction error can have asymmetric business costs depending on where it occurs (stockouts versus excess inventory; margin loss versus demand destruction). This review focuses on a central design question: **Should decision policies be conservative (robust) to protect against forecast uncertainty and distribution shift, or adaptive to exploit new information quickly?** We argue that the most effective systems combine both logics: robustness ensures feasibility and limits tail losses, while adaptation improves performance when feedback is informative and timely.

2. Materials and Methods

This review synthesizes concepts from four domains: (i) inventory management under uncertain demand, (ii) probabilistic demand forecasting, (iii) dynamic pricing and revenue management, and (iv) robust optimization and robust control. Robust optimization formalizes uncertainty via sets (e.g., interval, polyhedral, ellipsoidal, budgeted uncertainty) and chooses decisions that remain feasible and cost-effective under adverse realizations. www2.isye.gatech.edu+1

2.1 Synthesis architecture: a closed-loop decision cycle

We organize the literature into a decision cycle that practitioners can implement:

1. **Data ingestion:** sales, prices, promotions, availability, macro signals, competitor features, and operational constraints.
2. **Forecasting:** point forecasts and predictive distributions; calibration and reliability checks.
3. **Policy selection:** robust (set-based), adaptive (feedback-driven), or hybrid (adaptive within robust guardrails).
4. **Execution:** replenishment and price deployment with operational constraints.
5. **Monitoring:** service levels, stockouts, markdowns, realized margins, and decision stability.
6. **Feedback and governance:** retraining triggers, model risk checks, price fairness rules, audit logs.

2.2 Evaluation criteria

We compare policy families using operationally meaningful criteria: service level attainment, stockout probability, inventory holding and markdown costs, margin capture, decision volatility, and robustness under regime changes.

3. Results

3.1 Robust policies: stability under severe uncertainty

Robust policies define uncertainty sets around forecasts (e.g., demand within quantile bands or within $\pm x\%$). Decisions are selected to perform acceptably across all outcomes inside the set. This is especially valuable when the cost of failure is highly asymmetric, as in service-critical items, regulated markets, or fragile supply environments. Robust optimization provides a tractable methodology for such designs when uncertainty sets are chosen carefully. Foundational work describes how robust counterparts can remain solvable for many problem classes, while giving explicit protection against parameter uncertainty. www2.isye.gatech.edu In inventory contexts, robust formulations can incorporate both demand and lead-time uncertainty, often improving worst-case service outcomes at the expense of higher buffers. Springer Nature. **Strengths:** reduces catastrophic stockouts; provides explicit risk limits; improves auditability and operational trust. **Limitations:** can be conservative; may overstock or underreact to real demand signals if uncertainty sets are too wide.

3.2 Adaptive policies: responsiveness through feedback

Adaptive policies update forecasts and actions frequently, using forecast errors and new signals for recalibration. In stable regimes with clean feedback, this can reduce excess inventory and improve margin capture. Adaptive pricing and inventory control also align with online learning approaches where policies evolve as data arrives.

However, adaptivity can fail under:

- **Regime shifts** (structural breaks, viral demand, sudden supply shocks),
- **Delayed or censored feedback** (lost sales during stockouts),
- **Strategic behavior** (customers responding to algorithmic pricing),
- **Model drift** (forecasting model degradation over time).

3.3 Hybrid decision architecture: adaptive learning inside robust guardrails

The empirical and conceptual synthesis points to a hybrid: **robust guardrails** define feasibility and risk ceilings, while **adaptive tuning** operates within these limits. This creates “bounded learning”: the system adapts, but cannot violate service constraints, inventory safety floors, or maximum price change rules.

Practically, guardrails can include:

- Minimum service constraints and safety stock floors,
- Price move caps (absolute and relative),
- Fairness and transparency checks for algorithmic pricing,
- Escalation rules when drift or instability is detected,
- “Kill-switch” governance and human override.

This approach often dominates because it blends resilience (robustness) with efficiency (adaptation).

4. Discussion

4.1 Why forecasting accuracy is not enough

Forecast metrics such as MAPE and RMSE can be misleading for decision-making because they treat errors symmetrically and ignore asymmetric costs. Decision quality depends on the **cost of error**, not just error magnitude. For example, under-forecasting for high-margin, high-service items can be far more costly than over-forecasting for slow movers. Therefore, evaluation should be decision-centered: expected cost, stockout penalties, markdown losses, and stability over time. Decision-focused learning (where models are trained to improve downstream decisions, not only forecast accuracy) is increasingly relevant in AI-enabled operations.

4.2 Governance and trust: pricing as a sensitive operational policy

Dynamic pricing can create reputational risk if perceived as unfair or opaque. For frequent algorithmic updates, governance must include transparency principles, customer communication standards, and audit logs that explain why a price changed and whether it violated internal rules.

4.3 Research agenda


Future research should focus on:

1. Regime-switching demand and distribution shift detection,
2. Robust–adaptive co-design (guardrail selection + learning dynamics),
3. Decision-focused ML for inventory and pricing (predict-and-optimize vs end-to-end),
4. Explainability and auditability of learning-based pricing.

5. Conclusions

Robust and adaptive policies are complements, not substitutes. Robustness provides reliability and explicit protection against forecast failure and tail risk. Adaptation provides responsiveness when feedback is informative and regimes are learnable. The strongest operational designs combine both: **adaptive learning within robust guardrails**, supported by governance that ensures feasibility, stability, transparency, and auditability.

Figure 1



Decision cycle for AI forecasting with robust guardrails and adaptive tuning.
 Data → Forecast (point + distribution) → Guardrails (robust constraints) → Adaptive policy tuning → Execution (inventory + pricing) → Monitoring (service, cost, stability) → Feedback and retraining triggers.

Table 1. Robust versus adaptive inventory and pricing policies: practical comparison

Dimension	Robust policy	Adaptive policy	Operational implication
Uncertainty handling	worst-case / set-based	error-driven updating	robust lowers tail risk; adaptive improves average performance
Decision stability	high	medium to low	adaptivity can create volatility without guardrails
Data dependency	lower (works with bounds)	higher (needs clean feedback)	noisy or delayed feedback harms adaptive control
Stockouts	lower probability	depends on learning speed	robust typically protects service
Costs	higher buffers possible	lower buffers possible	adaptive reduces holding but can amplify spike losses
Governance	easier to audit	harder to audit	hybrid improves auditability and performance

Note: Hybrid designs often dominate: adaptive learning within robust service and pricing guardrails.

References

1. Daci, E., & Rexhepi, B. R. (2024). The role of management in microfinance institutions in Kosovo: Case study Dukagjini region. *Quality – Access to Success*, 25(202). <https://doi.org/10.47750/QAS/25.202.22> ORCID
2. Murtezaj, I. M., Rexhepi, B. R., Dauti, B., & Xhafa, H. (2024). Mitigating economic losses and prospects for the development of the energy sector in the Republic of Kosovo. *Economics of Development*, 23(3). <https://doi.org/10.57111/econ/3.2024.82> ORCID
3. Murtezaj, I. M., Rexhepi, B. R., Xhaferi, B. S., Xhafa, H., & Xhaferi, S. (2024). The study and application of moral principles and values in the fields of accounting and auditing. *Pakistan Journal of Life and Social Sciences*, 22(2). <https://doi.org/10.57239/PJLSS-2024-22.2.00286> ORCID
4. Rexhepi, B. R., Mustafa, L., Sadiku, M. K., Berisha, B. I., Ahmeti, S. U., & Rexhepi, O. R. (2024). The impact of the COVID-19 pandemic on the dynamics of development of construction companies and the primary housing market: Assessment of the damage caused, current state, forecasts. *Architecture Image Studies*, 5(2). <https://doi.org/10.48619/ais.v5i2.988> ORCID
5. Rexhepi, B. R., Rexhepii, F. G., Xhaferi, B., Xhaferi, S., & Berisha, B. I. (2024). Financial accounting management: A case of Ege Furniture in Kosovo. *Quality – Access to Success*, 25(200). <https://doi.org/10.47750/QAS/25.200.09> ORCID
6. Ben-Tal, A., & Nemirovski, A. (2002). Robust optimization—Methodology and applications. (Survey). www2.isye.gatech.edu
7. Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research*, 52(1), 35–53.

8. Talluri, K. T., & Van Ryzin, G. J. (2004). *The Theory and Practice of Revenue Management*. Springer.
9. Gallego, G., & Van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8), 999–1020.
10. Arrow, K. J., Harris, T., & Marschak, J. (1951). Optimal inventory policy. *Econometrica*, 19(3), 250–272.
11. Scarf, H. (1958). A min-max solution of an inventory problem. In *Studies in the Mathematical Theory of Inventory and Production* (pp. 201–209). Stanford University Press.
12. Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and Production Management in Supply Chains* (4th ed.). CRC Press.
13. Zipkin, P. H. (2000). *Foundations of Inventory Management*. McGraw-Hill.
14. Fisher, M., & Raman, A. (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research*, 44(1), 87–99.
15. Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts.
16. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 36(1), 54–74.
17. den Boer, A. V. (2015). Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1), 1–18.
18. Lattimore, T., & Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.
19. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
20. Zhang, R., & Luo, X. (2024). Deep reinforcement learning-based dynamic pricing (application study). *ScienceDirect* article. ScienceDirect
21. Yu, Y., Wang, X., Zhong, R. Y., & Huang, G. Q. (2016). E-commerce logistics in supply chain management: Practice perspective. *International Journal of Production Economics*, 179, 179–197.
22. Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66(3), 1025–1044.
23. Elmachtoub, A. N., & Grigas, P. (2022). Smart “predict, then optimize.” *Management Science*, 68(1), 9–26.
24. Bertsimas, D., & Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54(1), 150–168.
25. Bienstock, D., & Özbay, N. (2015). Robust inventory control under demand and lead time uncertainty. *Annals of Operations Research*. Springer Nature

26. Qin, Y., Zhang, R., & Zhou, X. (2023). An end-to-end deep learning model for the data-driven newsvendor problem. *European Journal of Operational Research*. ScienceDirect
27. Oroojlooyjadid, A., & Snyder, L. V. (2017). Applying deep learning to the newsvendor problem. (Technical report). Rossin College of Engineering
28. Chan, H. K., et al. (2024). Machine learning and deep learning models for demand forecasting in supply chain management: A review. *MDPI* review article. MDPI
29. Demand forecasting. (2025). In *Wikipedia*. Retrieved December 21, 2025.
30. Oracle. (2025). *AI in demand forecasting: Overview, use cases, & benefits*. Retrieved December 21, 2025.