Research on the Application of Artificial Intelligence in Dynamic Pricing and Vacancy Early Warning for Factory Rentals in Small- and Medium-Sized Industrial Parks

Zhilin Zhang[1,a], Shuang Han[2,b]

[1]Academy of Management Sciences, Macao Polytechnic University, Macao, China

[2]Qingdao Ruidongxiang Textile Co., Ltd, Qingdao, China

[a]675072215@qq.com, [b]598686835@qq.com

Keywords: Artificial Intelligence; Dynamic Pricing; Vacancy Early Warning; Industrial Parks; SMEs; Demand Forecasting; Price Elasticity; Causal Inference; Reinforcement Learning; Real Estate Strategy

Abstract: This paper proposes a strategy-oriented AI framework for dynamic pricing and vacancy early warning in small- and medium-sized industrial parks. Centered on causal elasticity estimation, multi-source data fusion, constrained optimization, and explainable risk stratification, the framework aligns revenue maximization with occupancy stability and policy compliance. A layered architecture—data, features, forecasting, pricing optimization, and vacancy risk management—supports continuous, safe online learning and human-in-the-loop governance. A calibrated case analysis indicates that AI-driven approaches can increase revenue, reduce vacancy, and stabilize price volatility relative to rule-based baselines, offering a pragmatic pathway for operators to transition from reactive adjustments to proactive portfolio steering.

1. Introduction

1.1. Research Background and Significance

The study responds to the dual imperative of improving industrial land-use efficiency and stabilizing SME manufacturing clusters where lease churn and cash-flow sensitivity are high. Reconceptualizing rent-setting as a continuous control problem with explicit risk tolerances and resource constraints enables granular segmentation, timely interventions, and measurable trade-offs between short-term monetization and long-term occupancy health. Strategically, AI-driven pricing and vacancy warning support resilient operations during shocks, guide retrofit investments in power and shared facilities, and provide auditable evidence for governance and policy alignment.

1.2. Research Status at Home and Abroad

Internationally, dynamic pricing is mature in hospitality and airlines with robust demand forecasting and online learning, yet industrial real estate applications remain nascent, often confined to hedonic models and broker heuristics with limited causal identification. Domestically, research emphasizes policy, land supply, and agglomeration effects, while in-park pricing practices largely follow cost-plus rules and episodic reviews, lacking a closed-loop integration with vacancy early warning. Gaps persist in elasticity estimation tailored to industrial units, operational constraint modeling, systematic use of IoT utilization signals, and causal robustness under promotions and subsidies, underscoring the need for an implementable AI strategy framework.

2. Theoretical Framework and Problem Definition

2.1. Market Mechanism and Pricing Objectives

The pricing objective integrates revenue growth, occupancy stability, and compliance by defining segment-specific occupancy targets and price corridors and adjusting prices according to

  DOI: 10.25236/gemmsd.2025.074

lead velocity, conversion behavior, and in-lease health signals [1]. High-elasticity segments warrant opportunistic markups in tight markets and price-for-volume tactics in soft markets, whereas low-elasticity, load-critical bays prioritize tenure structures, service bundling, and relationship management to anchor occupancy with minimal price volatility, thereby smoothing revenue, improving tenant mix, and optimizing resource utilization.

2.2. Data-Generating Process and Identification

The identification strategy disentangles the bidirectional causality between price and demand by leveraging multi-park panels, exogenous events like electricity tariff changes and staggered policy windows, and synchronized cross-building experiments. Debiased machine learning with cross-fitting estimates heterogeneous elasticities while integrating utilization breaks, payment discipline, and sectoral indicators into vacancy risk models, closing the loop between pricing actions, operational feedback, and risk calibration so that recommendations remain interpretable and robust under regime shifts [2].

2.3. System Architecture

The system architecture is layered for operability and governance. The data layer unifies CRM and leases, billing, IoT energy and access logs, lead funnels, competitor listings, and macro-policy calendars. The feature layer standardizes unit quality indices, location frictions, utilization trajectories, market pressure metrics, and event flags. The modeling layer combines gradient boosting and multihorizon temporal models for demand, debiased and hierarchical Bayesian methods for elasticities, and survival analysis for renewal and vacancy risks. The decision layer applies model predictive control and safety-constrained contextual learning to generate executable price updates and retention interventions, augmented with explainability for managerial review and compliant deployment.

3. Data and Feature Engineering

3.1. Data Sources and Integration

Internal sources include rent and discount histories, lead and deal logs, unit physical and power-capacity parameters, contract terms, receivables and collections, high-frequency energy and access readings, and service and maintenance records [3]. To strengthen observability, CRM notes and call-center transcripts can be distilled into structured tags (e.g., move-in urgency, subsidy requests), while ticketing systems provide timestamps for service-level adherence. External sources span nearby park listings and transactions, purchasing manager indices, freight rates, electricity tariffs, industry policies, and environmental curtailments, complemented by web-scraped competitor availability, logistics congestion indicators, and local permit data. Integration relies on master data governance with canonical IDs for units, buildings, and tenants, weakly supervised entity matching for fuzzy names and phone changes, and time harmonization to a weekly cadence aligned to decision cycles. Data quality routines include schema validation, referential integrity checks, and latency SLAs, with Kalman or EM imputation for missing values, variance-stabilizing transforms for noisy IoT streams, and anomaly detection to filter structural outliers from batch relocations or meter resets. A reproducible data lineage and versioning system (e.g., data lake with medallion layers) ensures analytic consistency for downstream pricing and risk decisions and supports auditability and rollback during governance reviews.

3.2. Feature Construction

Feature construction emphasizes comparability and dynamism by creating cross-building relative quality scores that blend floor loading, ceiling height, power redundancy, truck access, and proximity to gates, and time-varying supply–demand intensity indices based on active listings, inquiry velocity, and conversion rates. Utilization signals prioritize slopes and volatility over levels to capture early movements in production activity, using rolling gradients, weekend/overnight ratios, and change-point flags from energy and access data. Renewal and vacancy risk features focus on remaining term, payment discipline trajectories, service experience proxies such as ticket backlog and resolution times, and exposure to sectoral stress via external indices. Market pressure is represented by competitive density, price undercutting frequency, and responsiveness, with lagged policy shocks and decay functions to model fading effects of incentives. Interaction terms—such as relative price×lead velocity or utilization slope×time-to-expiry—encode managerial levers, while monotonicity and leakage checks enforce business plausibility [4]. All features are standardized per segment to support stable model training and fair cross-park comparisons.

3.3. Labeling and Validation Protocol

Labeling defines weekly leads and signed leases for demand and churn within horizon h and time-to-vacancy for risk, with renewal outcomes tied to contract maturity windows and grace periods to avoid mislabeling early exits or negotiated extensions. To reduce hindsight bias, labels are frozen with event-time stamps and feature lags, and exclusion windows handle data gaps around meter swaps or major refurbishments. Evaluation uses rolling-origin tests with park-level blocking and non-overlapping folds across macro regimes to assess temporal and spatial generalization. Validation encompasses strategy metrics such as compliance rates of price recommendations, occupancy volatility, shifts in the revenue–occupancy Pareto frontier, and calibration and stability of risk tiers, alongside traditional predictive metrics (WAPE/RMSE for demand, AUC/PR-AUC and Brier score for risk) [5]. Uncertainty diagnostics (prediction intervals coverage, calibration curves) and decision-focused tests (uplift in expected revenue under policy simulation, regret analysis under counterfactual pricing) ensure that modeling gains translate into measurable operational outcomes and governance-ready evidence, with error budgets and escalation thresholds guiding safe deployment.

4. Modeling Methodology

4.1. Demand Forecasting

The demand forecasting module is designed as a two-stage pipeline that balances accuracy with operational stability. In stage one, gradient boosting models (e.g., LightGBM/XGBoost) learn nonlinear interactions among static unit attributes, market pressure indicators, and event features including holidays, policy windows, and maintenance or supply disruptions [6]. Feature sets explicitly capture lead velocity, view-to-inquiry conversion, days on market, competitor price indices, and rolling utilization trends (kWh per square meter, access counts), while target leakage is mitigated through strict time-based splits and lagged covariates [7]. In stage two, multihorizon temporal models—such as Temporal Fusion Transformers or N-BEATS—generate coherent forecasts for 1–12 week horizons with exogenous regressors, capturing seasonality and cross-segment co-movements [8]. Hierarchical reconciliation (e.g., MinT or bottom-up with shrinkage) ensures consistency across unit–building–park aggregates, preventing inventory-level divergence and enabling portfolio-level S&OP alignment. Post-forecast calibration employs isotonic regression or Platt scaling to correct probability-like outputs (e.g., weekly lease win rates) and apply dampening factors so that small signal fluctuations do not trigger outsized price changes. Explicit event modeling operationalizes holiday effects, temporary policy incentives, and supply shocks through regime flags and decay kernels, while scenario overlays quantify forecast envelopes under optimistic, base, and conservative assumptions. The result is a forecast stack that is auditable, stress-testable, and directly consumable by pricing, with guardrails such as uncertainty-aware price adjustments, minimum signal thresholds for action, and exception queues when forecast dispersion exceeds predefined limits.

4.2. Price Elasticity and Causal Effects

Elasticity estimation framed as a causally robust, heterogeneity-aware workflow that converts observational data into actionable local demand curves. First, specify a structural equation in which realized demand depends on price, promotions, unit and tenant characteristics, macro-competition conditions, and unobservables, with endogeneity arising from managerial responses and tenant selection. To address endogeneity, apply double/debiased machine learning with cross-fitting to learn nuisance functions (propensity for price/promotion assignments and outcome regressions), and pair with instrumental variables such as exogenous electricity tariff changes, staggered municipal subsidies, or cross-building price shocks induced by unrelated maintenance. Next, estimate conditional average treatment effects and semi-parametric elasticities using causal forests or Bayesian additive regression trees, stratified by area bands, power capacity tiers, and quality grades. Resulting models deliver interpretable local response curves and elasticities with uncertainty intervals, enabling policy simulations such as "what revenue–occupancy shift results from a 5% price cut in 200–500 m², 100 kW bays?" Then, conduct diagnostics—overlap checks, sensitivity to instrument strength, placebo tests on pre-periods, and event-study plots—to verify identification stability. Finally, operationalize outputs through segment playbooks: high-elasticity, high-substitutability units receive wider exploration ranges and dynamic discount ladders; medium-elasticity segments adopt measured micro-adjustments tied to lead velocity thresholds; low-elasticity, mission-critical bays prioritize non-price levers (term length, step rents, service SLAs) to safeguard occupancy. Governance rules cap week-over-week impacts from elasticity re-estimation on prices, trigger human review for large implied elasticities, and archive counterfactual analyses to ensure compliance and auditability.

4.3. Dynamic Pricing Optimization

The pricing optimizer solves a constrained, rolling-horizon problem that trades off near-term revenue with occupancy stability under enforceable guardrails. The objective maximizes expected rent and fee income net of incentives across H weeks, penalizing vacancy risk and excessive volatility. Constraints include price floors/ceilings, week-over-week change limits, peer comparability within like-for-like clusters, fairness and anti-discrimination policies, and capacity constraints such as power quotas or logistics dock limits. The execution architecture has three layers. Baseline pricing sets anchor rents per segment using long-run demand, cost benchmarks, and compliance requirements. Windowed micro-adjustments apply small, rule-governed changes (e.g., ±1–3% weekly) driven by forecast deltas, lead funnel health, and competitor moves, with uncertainty-weighted step sizes that shrink under high variance. Event-driven shifts activate during predefined windows—policy incentives, major tenant move-ins/outs, infrastructure downtime—using temporary deviation corridors that automatically revert according to decay schedules. Optimization is computed via model predictive control with mixed-integer components for discrete incentives (free months, deposit terms), while uncertainty is incorporated through chance constraints or robust optimization to avoid overreacting to noisy signals. An online learning layer—typically safety-constrained contextual Thompson sampling over a discretized price grid—facilitates exploration in data-sparse segments and accelerates adaptation under regime changes, bounded by hard compliance limits and monotonicity rules to prevent whipsawing. Every recommendation ships with an explanation card (key drivers, constraint bindings, expected impact, confidence band) and an override workflow that logs human decisions for downstream policy refinement.

4.4. Vacancy Early Warning and Intervention

The vacancy risk engine forecasts renewal probability and time-to-vacancy by fusing contractual, behavioral, and operational signals into an interpretable risk stratification. Core models combine survival analysis (Cox, Weibull AFT) for time-to-event with gradient boosting or recurrent models for classification, incorporating features such as tenure and time-to-expiry, payment timeliness and arrears trajectory, service ticket frequency and severity, utilization slopes and volatility from energy and access data, and sectoral stress indicators [9]. Early-warning signals detect pre-churn patterns like sustained utilization decline, rising overnight inactivity, or sudden reductions in access diversity, triggering alerts before formal termination notices. Risks are tiered into green–yellow–red bands with calibrated thresholds that balance precision and recall across intervention capacities. Each tier is bound to SLA-based playbooks: green receives lightweight check-ins and loyalty nudges; yellow prompts account reviews, tailored term restructuring, and power quota optimization; red escalates to executive outreach, temporary fee buffers, or sublease enablement with rapid decision cycles. Interventions are prioritized by expected vacancy-days averted per unit cost, and their outcomes feed back into uplift models to refine treatment effect estimates [10]. A closed-loop governance process audits alert accuracy, intervention adherence, and realized impact, updates thresholds via Bayesian calibration, and enforces ethical constraints such as non-retaliatory pricing and consistent treatment across similar tenants. Operationally, dashboards surface tenant-level drivers, counterfactual risk trajectories with and without action, and capacity-aware scheduling to prevent over-commitment of retention resources while maximizing overall occupancy stability.

5. Empirical Evaluation and Case Study

5.1. Experimental Setup

The evaluation adopts rolling-origin tests with park-level blocking to prevent temporal leakage and cross-site contamination, using 8–12 week forecast/pricing horizons and weekly decision cadences that mirror operational practice. We benchmark against three baselines—static price tables, quarterly review adjustments, and sales-driven heuristic discounting—under identical inventory and policy constraints to ensure fairness. Treatment arms compare model predictive control and safety-constrained contextual learning, both fed by the same demand forecasts, elasticity estimates, and risk signals. Data are split into pre-tuning, validation, and locked test windows across distinct macro regimes, with hyperparameters selected via nested cross-validation. Scenario stress tests inject exogenous shocks including temporary policy incentives, electricity tariff step-ups, competitor capacity expansions, and ±20% demand swings to probe robustness and governance boundaries. All runs produce immutable audit trails capturing inputs, constraints active at solve time, driver attributions, uncertainty bands, and human overrides, enabling post-hoc explainability and evaluation of override effectiveness and drift.

5.2. Results

AI strategies improve revenue, reduce vacancy, and temper price volatility relative to baselines, delivering consistent gains across parks with heterogenous mixes of unit sizes and power capacities. Performance uplift concentrates in high-substitutability, mid-quality segments where micro-adjustments capture transient demand, while load-critical, low-elasticity bays exhibit steadier occupancy through term and service levers rather than discounts. Price change smoothness remains within governance-defined limits, with week-over-week variance notably below heuristic methods. Compliance adherence rises due to embedded guardrails, and proxy tenant satisfaction indicators—complaint rates, negotiation cycle length, and renewal intent—improve alongside financial metrics. Under stress tests, both MPC and safety-constrained bandits sustain advantage, with MPC excelling in stable regimes and bandits adapting faster during abrupt shocks, indicating complementary operational roles.

5.3. Ablation and Sensitivity

Removing IoT utilization features weakens early-warning sensitivity and erodes pricing foresight, increasing regret and prompting reactive discounting during latent downturns; adding back even coarse energy trends recovers a substantial portion of performance. Omitting causal debiasing biases elasticities toward over-discounting in soft markets and under-pricing scarcity in tight markets, compressing long-run value and elevating vacancy volatility. Tightening fairness and smoothing constraints modestly reduces revenue uplift but significantly lowers perceived volatility and complaint risk, highlighting governance as a strategic asset; conversely, loosening constraints raises short-term gains at the expense of override frequency and reputational exposure. Sensitivity to forecast dispersion shows that uncertainty-aware dampening mitigates whipsawing without materially sacrificing opportunity capture.

5.4. Managerial Insights

Operators should institutionalize pricing and retention within a recurring S&OP cadence, supported by tri-level dashboards for price, risk, and compliance, and clearly defined human override thresholds with reason codes to reinforce accountability. Investment priority should favor data quality, lead funnel integrity, and utilization signals over expansive but noisy external scraping, yielding faster payback and more stable policies [11]. Strategically, apply flexible pricing to mid-quality, substitutable units where elasticity is high, while locking value in scarce, high-power bays through term structures, step rents, and service bundles to protect mission-critical tenants. Align price windows with marketing campaigns and policy programs to amplify lift, and run disciplined post-action reviews that compare realized outcomes to counterfactuals, feeding continuous refinement of models, constraints, and playbooks.

6. Conclusion

This paper presents an AI strategy framework that, through causally robust elasticity estimation, constrained optimization, and layered risk governance, coordinates pricing, occupancy, and compliance for factory rentals in small- and medium-sized industrial parks. The approach delivers revenue growth, vacancy control, and explainable, auditable operations suitable for real-world deployment. Future work will pursue cross-park transfer and meta-learning to address cold starts, richer exogenous instruments and quasi-experiments for stronger identification, multimodal integration of text and imagery from inquiries and inspections, embedding carbon and energy pricing into rent optimization to support low-carbon transitions, and randomized field trials to validate external generalizability and organizational adoption.

References

[1] Tsai W H , Hsu J L , Chen C H .Integrating activity-based costing and revenue management approaches to analyse the remanufacturing outsourcing decision with qualitative factors [J]. International Journal of Revenue Management, 2010, 1(4):367-387.DOI:10.1504/IJRM.2007.015539.

[2] Chernozhukov V , Newey W K , Singh R .Automatic Debiased Machine Learning of Causal and Structural Effects[J]. Econometrica, 2022, 90.DOI:10.48550/arXiv.1809.05224.

[3] Lee K I , Choi D G .SYSTEM AND METHOD FOR BUILDING INTEGRATION KNOWLEDGE BASE BASED A PLURALITY OF DATA SOURCES:KR20170181783 [P]. KR20190079805A[2025-08-28].

[4] Buckley J .Grain market report - Looser supply trims wheat price again [J]. Milling and Grain, 2023.

[5] Gray A R , Macdonell S G .A comparison of techniques for developing predictive models of software metrics[J].Information and Software Technology, 1997, 39(6):425-437.DOI:10.1016/S0950-5849(96)00006-7.

[6] Atkinson E J , Therneau T M , Iii L J M ,et al. Assessing fracture risk using gradient boosting machine (GBM) models[J].Journal of Bone & Mineral Research, 2012, 27(6):1397-1404. DOI:10.1002/jbmr.1577.

[7] Nguyen,N,T,et al.Trends in utilization and outcomes of laparoscopic versus open appendectomy [J].American Journal of Surgery, 2004.

[8] Ji Z, Xun J, Yin C, Guo Z, Tian G. Interpretable forecasting of traction energy consumption based on Nbeats and Temporal Fusion Transformers[C]//2024 IEEE 13th Data Driven Control and Learning Systems Conference (DDCLS). Kaifeng, China: IEEE, 2024: 1828-1835. DOI: 10.1109/DDCLS61622.2024.10606795.

[9] Yan J .Survival Analysis[J].Publications of the American Statistical Association, 2004, 99(467):2. DOI:10.1198/jasa.2004.s348.

[10] Rogal S S , Neely B , Merante M ,et al.Getting to Implementation for HIV Pre-Exposure Prophylaxis (GTI-PrEP): A data-driven approach to PrEP prescribing[J].Implementation Science Communications, 2025, 6(1).DOI:10.1186/s43058-025-00749-2.

[11] Villacampa Portuguez P L .Savings, Investment, Terms of Trade, and Sustainability of the Internal Approach: 1993 - 2019[J].Economía (02544415), 2024, 47(93).DOI:10.18800/economia.202401.004.