International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2025: 7.089

Benchmark-Guided Prescriptive Analytics for Menu and Margin Optimization in Multi-Location Restaurants

Jerry Brown[1], **Uma Ajay Kumar Reddy PS**[2]

[1]Culinary Analytics Consultation Group
Email: *jerry[at]culinaryanalytics.com*

[2]Culinary Analytics Consultation Group
Email: *umaveena.alvas[at]gmail.com*

Abstract: ***Multi-location restaurant enterprises generate large volumes of operational data from point-of-sale systems, recipes, inventory platforms, procurement records, payroll systems, delivery channels, and external benchmark sources. However, many operators still rely on retrospective dashboards or isolated forecasts that describe performance without recommending feasible action. This paper presents a benchmark-guided prescriptive analytics framework for menu and margin optimization in multi-location restaurants. The framework integrates operational and benchmark data, estimates demand, ingredient cost volatility, labor burden, and stockout risk, converts peer-relative performance gaps into robust opportunity scores, and ranks feasible interventions under pricing, service, brand, supply, and labor constraints. The restaurant decision problem is formalized as a constrained stochastic optimization problem, and the paper specifies the data architecture, feature engineering pipeline, predictive modules, benchmark-gap scoring method, recommendation objective, evaluation protocol, and operational safeguards. Because the manuscript proposes a framework rather than reporting completed field results, it defines offline backtesting, ablation analysis, and matched-store pilots as the appropriate validation path. The contribution is a practical methodology for moving restaurant analytics from descriptive reporting and standalone prediction toward explainable, benchmark-aware prescriptive decision support.***

Keywords: benchmarking, decision support systems, demand forecasting, inventory analytics, labor analytics, margin optimization, menu engineering, multi-location restaurants, prescriptive analytics, restaurant analytics

1. Introduction

Multi-location restaurant operators make repeated, interdependent decisions across menu pricing, promotional exposure, recipe execution, procurement, labor allocation, replenishment, and item availability. Each decision can affect revenue, food cost, labor cost, waste, stockouts, service speed, and guest experience. A menu item that appears profitable in one location can become fragile in another when ingredient costs rise, delivery-channel fees increase, local demand changes, or recipe execution drifts.

Most restaurant analytics programs are still organized around descriptive performance reporting. Dashboards summarize sales, food-cost percentage, labor percentage, waste, voids, discounts, and inventory variances after the fact. Some organizations add demand forecasting, but forecasts alone do not answer the managerial question that matters most: what should be changed next, in which location, by how much, and with what operational risk?

The question is especially difficult in restaurant chains because stores differ in wage structures, customer mix, daypart patterns, local competition, vendor availability, delivery-channel dependence, and execution quality. A location can grow revenue while losing margin through recipe drift or delivery-channel dilution. Another can appear labor efficient while experiencing service degradation, avoidable overtime, or hidden stockouts.

Prescriptive analytics is the natural next step after descriptive and predictive analytics because it links data and prediction to action selection under constraints. In restaurants, however, prescriptions must be operationally credible. A recommendation to increase price is incomplete unless it accounts for demand response, substitution, competitive position, channel margin, and brand rules. A recommendation to raise safety stock is incomplete unless it accounts for shelf life, storage capacity, supplier reliability, and waste risk. A recommendation to reduce labor hours is incomplete unless it accounts for service-level risk and role coverage.

This paper proposes a benchmark-guided prescriptive analytics framework for menu and margin optimization in multi-location restaurants. The central premise is that the best next action should be selected from the intersection of reconciled operational data, predictive risk estimates, peer-relative benchmark gaps, business constraints, and manager-reviewable explanations.

The contributions are fourfold:

1) A formal restaurant decision model that treats menu and margin optimization as a constrained stochastic decision problem;

2) An enterprise architecture that unifies POS, recipe, inventory, procurement, payroll, delivery-channel, and benchmark data;
3) A predictive stack for action-aware demand, ingredient cost volatility, labor burden, and stockout risk;
4) A benchmark-gap and recommendation-ranking mechanism that converts peer-relative performance gaps into explainable action portfolios.

2. Related Work

a) Restaurant Analytics, Menu Engineering, and Revenue Management

Restaurant analytics literature has examined demand forecasting, menu engineering, menu pricing, and revenue management. Recent studies show that restaurant demand forecasts improve when models use richer internal and external covariates, including product, location, time, weather, event, and disruption-related variables [5]-[7]. Menu engineering research emphasizes the joint role of item popularity, contribution margin, menu position, and substitution effects in pricing and menu placement decisions [8]-[10]. Restaurant revenue-management research extends pricing, duration, and capacity logic to restaurant contexts [11], [12].

These streams provide important building blocks, but they often remain single-domain. A forecasting model may predict item demand without recommending an inventory, labor, or pricing action. A menu-engineering model may rank items by margin and popularity without reconciling supply risk, labor burden, or benchmark gaps. A revenue-management model may optimize capacity or price without integrating recipe-level cost and inventory variance.

b) Prescriptive Analytics and Decision Systems

Prescriptive analytics extends prediction by recommending decisions under uncertainty. Reviews of the field position prescriptive analytics as a higher-maturity layer that uses optimization, simulation, and policy design to convert analytical signals into decisions [1]. Operations and decision-science research has developed data-driven prescriptions for conditional stochastic optimization, inventory control, and explainable prescriptive models [2]-[4]. These methods are directly relevant to restaurant chains because restaurants repeatedly choose prices, ordering quantities, staffing levels, and menu exposure under uncertain demand and cost conditions.

Restaurant-specific prescriptive frameworks remain limited. Existing work commonly treats prediction as the terminal product or optimizes only one domain at a time. In practice, a margin intervention can affect demand, inventory, waste, labor, channel mix, and guest perception simultaneously. This motivates an integrated framework that converts predictions and benchmarks into feasible actions.

c) Benchmarking and Efficiency Analysis

Benchmarking and efficiency analysis help managers interpret performance relative to comparable units. Data envelopment analysis and related methods have been applied to hospitality and restaurant productivity assessment across multiunit contexts [14]-[16], building on foundational efficiency measurement work [18]. Benchmarking is useful because raw ratios can be misleading without peer context: a food-cost ratio, labor ratio, or stockout rate should be interpreted relative to concept, geography, volume band, service model, wage structure, and channel mix.

The limitation is that benchmarking often stops at diagnosis. It may identify that a location trails a cohort but not whether the best next intervention is price correction, recipe audit, supplier substitution, labor reallocation, safety-stock adjustment, or menu exposure change. This paper therefore treats benchmarking as an input to prescriptive action ranking rather than as an endpoint.

3. Framework Overview

The framework is designed for enterprises with multiple restaurant locations, shared menu architecture, and access to store-level operational data. It can also be adapted to smaller groups when recipe, POS, labor, procurement, and inventory records are available. Fig. 1 shows the conceptual workflow. Source data are harmonized into a semantic operational layer, predictive modules estimate risk and response, benchmark engines define peer gaps, and prescriptive optimization generates action cards for human review and feedback.

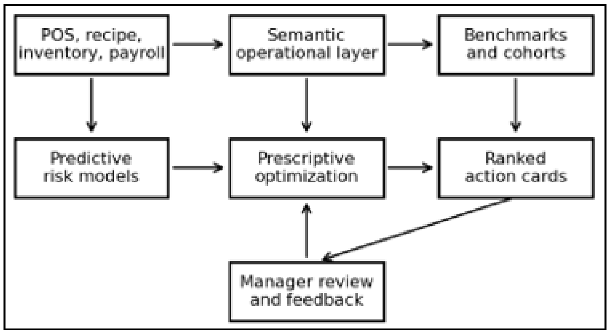

Figure 1: Benchmark-guided prescriptive analytics workflow for multi-location restaurant optimization.

The framework follows six design principles. First, modeling should be decision-first: features and predictions should support actions that managers can implement. Second, recommendation should follow reconciliation: item, ingredient, vendor, unit, modifier, channel, labor-role, and store identifiers must be standardized before optimization. Third, benchmark-awareness should guide prioritization: the system should not only ask whether a store is worse than its own history, but whether it trails a credible peer target. Fourth, uncertainty should be explicit through confidence intervals and downside-risk penalties. Fifth, prescriptions must respect operational constraints, including brand standards, supplier feasibility, service levels, labor rules, and price-change limits.

Sixth, each recommendation should be explainable as a human-reviewable action card.

TABLE I. CORE NOTATION

Symbol	Meaning
L, I, K, R, T	Sets of locations, menu items, ingredients, labor roles, and time periods.
$s_{l,t}$	Operational state for location l at time t.
$a_{l,t}$	Decision/action vector for location l at time t.
$\Pi_{l,t}$	Operating profit after revenue, variable cost, waste, stockout, channel, discount, and implementation costs.
$G_{l,t}$	Weighted benchmark-gap score; larger positive values indicate peer-relative underperformance.
$D_{i,l,t}$	Random demand for item i at location l and time t.
$\pi^{so}_{k,l,t}$	Ingredient stockout probability.
$\rho_i, \tau, \bar{\pi}$	Price-change bound, service-level requirement, and stockout-risk threshold.

4. Formal Decision Model

Let l in L denote restaurant locations, i in I menu items, k in K ingredients, r in R labor roles, and t in T decision periods. The operational state for a location combines transaction, recipe, inventory, labor, procurement, channel, and benchmark information:

$$s_{l,t} = \{x^{pos}_{l,t}, x^{rec}_{l,t}, x^{inv}_{l,t}, x^{lab}_{l,t}, x^{chan}_{l,t}, x^{bench}_{l,t}\} \quad (1)$$

The decision vector contains operational levers available to management:

$$a_{l,t} = \{\Delta p_{i,l,t}, \Delta y_{i,l,t}, q_{k,l,t}, h_{r,l,t}, m_{i,l,t}, u_{i,l,t}\} \quad (2)$$

where Delta p denotes price adjustments, Delta y recipe yield or portion corrections, q procurement quantities, h labor-hour allocations, m menu exposure or promotion decisions, and u item activation or substitution decisions. Operating profit is represented as

$$\Pi_{l,t} = R_{l,t} - C^{food}_{l,t} - C^{labor}_{l,t} - C^{waste}_{l,t} - C^{stockout}_{l,t} - C^{channel}_{l,t} - C^{disc}_{l,t} \quad (3)$$

The enterprise objective is to maximize expected profit while closing benchmark-relative gaps over a finite horizon:

$$\max_{\mu} E\left[\sum_{h=0}^{H-1} \beta^h (\Pi_{l,t+h} + \lambda \Delta G_{l,t+h})\right] \quad (4)$$

subject to the stochastic state transition

$$s_{l,t+1} = T(s_{l,t}, a_{l,t}, \xi_{l,t}) \quad (5)$$

where mu is the decision policy, beta is the discount factor, Delta G is expected benchmark-gap closure, and xi captures demand, supplier, labor, and execution uncertainty. Operational constraints include price-change, recipe, stockout-risk, service-level, and brand guardrails:

$$|\Delta p_{i,l,t}| \le \rho_i, \quad y_i^{min} \le y_{i,l,t} \le y_i^{max} \quad (6)$$

$$\pi^{so}_{k,l,t}(a) \le \bar{\pi}, \quad svc_{l,t}(a) \ge \tau, \quad brand(a)=1 \quad (7)$$

This formulation intentionally separates prediction from prescription. Predictive models estimate uncertain outcomes; the prescription layer selects feasible actions that optimize expected value while controlling downside risk and implementation cost.

5. Data Architecture and Feature Engineering

1) *Semantic Operational Layer*

Prescriptive logic should operate on reconciled operational entities rather than raw source-system fields. A delivery-platform menu item may need to be mapped to the internal recipe, adjusted for modifiers, associated with channel fees, and converted into usable ingredient quantities before true contribution margin is visible. Payroll records should be aligned to store, date, shift, daypart, and labor role. Procurement records should normalize vendor items, pack sizes, units of measure, usable yield, taxes, rebates, and freight when available.

TABLE II. PROPOSED SYSTEM ARCHITECTURE

Layer	Purpose
Ingestion	Collect POS, recipe, inventory, purchasing, vendor, payroll, delivery-channel, and benchmark feeds.
Harmonization	Standardize identifiers, modifiers, units, pack sizes, vendors, roles, stores, and channels.
Semantic layer	Build canonical item-location-daypart, ingredient-location-day, labor-shift, and channel entities.
Feature store	Persist time-aligned features, targets, action histories, and prediction outputs.
Predictive services	Estimate demand, cost volatility, labor burden, stockout risk, and uncertainty bands.
Benchmark engine	Assign peer cohorts, compute robust targets, and score performance gaps.
Prescriptive engine	Generate actions, simulate outcomes, enforce constraints, and rank recommendations.
Workflow layer	Present action cards, capture approvals and overrides, and log realized outcomes.

2) *Transaction-to-Ingredient Explosion*

POS line items are normalized by location, daypart, channel, and modifier state, then exploded through the recipe graph to estimate theoretical ingredient usage:

$$u^{th}_{k,l,t} = \sum_i r_{i,k} q_{i,l,t} \quad (8)$$

where $r_{i,k}$ is the recipe quantity of ingredient k in menu item i and $q_{i,l,t}$ is observed sales volume.

3) *Effective Food Cost and Margin*

Normalized item food cost should account for usable yield and vendor conversion:

$$c_{i,l,t} = \sum_{k \in K_i} (r_{i,k} / y_{i,k}) \tilde{c}_{k,l,t} \quad (9)$$

where $y_{i,k}$ is usable yield and $\tilde{c}_{k,l,t}$ is the effective unit cost after conversion, vendor normalization, and procurement adjustments. Contribution margin should then be computed by channel:

$$CM_{i,l,t,c} = P_{i,l,t,c} - c_{i,l,t} - fee_{i,l,t,c} - pack_{i,l,t,c} \quad (10)$$

where c indexes dine-in, pickup, direct delivery, or marketplace delivery. Channel-level treatment prevents high-revenue, fee-heavy transactions from appearing more profitable than they are.

4) *Inventory Reconciliation*

The residual between actual and theoretical inventory movement is

$$\delta^{rec}_{k,l,t} = (OH_{k,l,t-1} + R_{k,l,t} - OH_{k,l,t}) - u^{th}_{k,l,t} \quad (11)$$

where OH denotes on-hand inventory and R denotes receipts. Persistent residuals may indicate waste, over-portioning, spoilage, shrinkage, recipe drift, transfer error, or inaccurate counts. The residual should therefore affect recommendation confidence and may trigger audit actions before commercial optimization.

5) Labor and Benchmark Features

Payroll and time-clock records are aligned to store, date, shift, and daypart. Burdened labor cost should include base wages, overtime premiums, payroll taxes, benefits, and agency or contractor costs when available. Representative features include revenue per labor hour, contribution margin per labor hour, items per labor hour, overtime ratio, role-mix entropy, service-time proxies, and labor density by daypart.

Benchmark cohorts should be defined by concept, service model, geography, volume band, average check range, wage environment, and channel mix. Benchmark features can include food-cost ratio, burdened labor ratio, gross margin, inventory variance, stockout rate, menu concentration, average ticket, waste rate, and contribution margin per labor hour.

6. Predictive Modules

1) Action-Aware Demand Forecasting

Demand should be modeled at the item-location-daypart-channel level. Because recommendations may change prices, availability, or menu exposure, the demand model must be action-aware:

dhat_i,l,t = f_d(x_i,l,t) exp(epsilon_i Delta p_i,l,t + sum_{j in S_i} gamma_i,j Delta p_j,l,t + phi_i m_i,l,t) (12)

Here f_d(.) is the baseline demand model, epsilon_i is own-price elasticity, gamma_i,j captures substitution or complementarity, and phi_i captures menu visibility or promotional effects. This structure permits counterfactual simulation of candidate actions before ranking them.

2) Ingredient Cost Volatility

Current food cost alone is insufficient when ingredient prices are unstable. A menu item can be high margin today but fragile if exposed to volatile ingredients. Let cost volatility exposure for ingredient k be approximated by a conditional quantile spread:

v_k,t = Q_0.9(c_k,t+h | x) - Q_0.5(c_k,t+h | x) (13)

Menu-level volatility exposure is

V_i,l,t = sum_{k in K_i} (r_i,k / y_i,k) v_k,t (14)

This permits the engine to distinguish items that are profitable but fragile from items that are profitable and cost-stable.

3) Labor Burden and Service Risk

Expected labor burden can be modeled as

b_l,t = (W^base_l,t + W^OT_l,t + W^tax_l,t + W^benefit_l,t) / (S_l,t + eta) (15)

where S_l,t denotes sales and eta avoids instability at low volume. The target can be extended to estimate service degradation under alternative staffing plans. This is important because reducing labor cost may harm throughput, wait time, review scores, or repeat demand.

4) Stockout Risk

Ingredient stockout probability can be estimated from on-hand inventory, replenishment, forecast demand, lead-time uncertainty, and safety stock:

pi^so_k,l,t = P(OH_k,l,t + Q_k,l,t - sum_i r_i,k dhat_i,l,t < SS_k,l,t) (16)

where Q is incoming replenishment and SS is safety stock. The probability may be estimated using simulation or a direct classification model. High stockout risk on high-margin items should suppress promotions and trigger replenishment or approved-substitution actions.

7. Benchmark-Gap Scoring

Benchmarking is useful only when cohort definitions are credible and scoring is robust to outliers. Let m in M denote performance metrics, and let delta_m equal +1 when higher values are better and -1 when lower values are better. For location l belonging to peer cohort g(l), define a robust benchmark target as

B_m,g,t = Q_0.75(z_m,.,t | g) if delta_m=+1; Q_0.25(z_m,.,t | g) if delta_m=-1 (17)

where z_m,l,t is the observed metric. The standardized benchmark-gap score is

g_m,l,t = w_m clip(delta_m(B_m,g,t - z_m,l,t)/(MAD_m,g,t + eps), -3, 3) (18)

where w_m is a metric weight and MAD is median absolute deviation. The overall score is

G_l,t = sum_m g_m,l,t = G^menu_l,t + G^food_l,t + G^labor_l,t + G^inventory_l,t (19)

A positive value indicates peer-relative underperformance. Decomposition is operationally important because two locations can have the same total gap for different reasons. One may need recipe calibration while another needs daypart staffing adjustment or stockout prevention. When peer coverage is thin or cohort match quality is low, benchmark weights should be reduced and enterprise historical baselines should receive greater weight.

8. Prescriptive Recommendation Engine

The recommendation engine converts predictive states and benchmark gaps into ranked action portfolios. Candidate actions include selective price changes, recipe yield corrections, safety-stock updates, supplier substitution, prepbatch changes, labor reallocation across dayparts, menu exposure changes, and temporary deactivation of fragile low-margin items.

For feasible actions a in A(s_t), the engine solves

max_a E[Delta Pi(a)] + lambda_1 Delta G(a) - lambda_2 CVaR_alpha(-Delta Pi(a)) - lambda_3 C_impl(a) - lambda_4 H(a) (20)

where Delta Pi(a) is expected profit uplift, Delta G(a) is expected benchmark closure, CVaR penalizes downside risk, C_impl is implementation cost, and H penalizes action instability or oscillation. Expected benchmark closure is

Delta G(a) = G_l,t - G_prime_l,t(a) (21)

where G_prime is the simulated gap score after applying action a. The output should be an explainable action card rather than an opaque score. A complete card should include action description, affected store and menu items, expected profit uplift, expected benchmark closure, confidence band, primary risk, constraints checked, operational rationale, and manager feedback fields.

TABLE III. ILLUSTRATIVE CANDIDATE ACTIONS

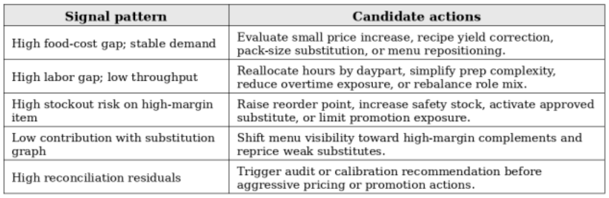

Signal pattern	Candidate actions
High food-cost gap; stable demand	Evaluate small price increase, recipe yield correction, pack-size substitution, or menu repositioning.
High labor gap; low throughput	Reallocate hours by daypart, simplify prep complexity, reduce overtime exposure, or rebalance role mix.
High stockout risk on high-margin item	Raise reorder point, increase safety stock, activate approved substitute, or limit promotion exposure.
Low contribution with substitution graph	Shift menu visibility toward high-margin complements and reprice weak substitutes.
High reconciliation residuals	Trigger audit or calibration recommendation before aggressive pricing or promotion actions.

9. Evaluation Protocol

Because the present work proposes a framework, evaluation should be treated as a protocol rather than as evidence of realized field uplift. Evaluation should measure both predictive performance and prescription quality.

a) *Predictive Metrics*

Demand forecasts can be evaluated using weighted absolute percentage error:

WAPE = sum |y - yhat| / sum y (22)

Cost-volatility models can be assessed with pinball loss and interval coverage. Labor-burden forecasts can use mean absolute error or WAPE. Stockout models should be evaluated using AUROC, Brier score, calibration error, and recall at operationally relevant thresholds.

b) *Prescriptive Metrics*

Profit uplift and benchmark closure can be defined as

Uplift = (Pi_test - Pi_baseline) / Pi_baseline (23)

BCR = (G_pre - G_post) / (G_pre + eps) (24)

Additional metrics include food-cost reduction, labor-burden reduction, waste reduction, stockout-rate reduction, recommendation acceptance rate, precision at top-k, implementation latency, manager override rate, customer complaint movement, and action stability over time.

c) *Backtesting, Pilots, and Ablations*

Offline backtesting should compare the framework against descriptive threshold dashboards, predictive-only forecasts with manual rules, traditional menu engineering, and single-domain heuristics such as food-cost or labor-percentage triggers. Online evaluation should use matched-store pilots or cluster-randomized tests with difference-in-differences analysis to isolate realized impact. Ablation studies should remove benchmark guidance, stockout-risk modeling, labor burden, substitution-aware demand, and action-stability penalties to determine which components drive performance.

10. Operational Safeguards

A production system must handle operational edge cases explicitly. Cold-start stores and new items should use hierarchical shrinkage from category, concept, region, and cohort priors. Modifier-heavy menus and combo items should be decomposed to recipe and margin level. When inventory reconciliation residuals remain elevated, the system should reduce recommendation confidence and prioritize audit actions. When vendor fill rates or lead times deteriorate, resilience actions should be prioritized over aggressive promotion.

The framework should also model dine-in, pickup, direct delivery, and marketplace delivery separately because margins, fees, elasticity, and guest expectations differ by channel. Commissary and shared-prep constraints should propagate upstream capacity and ingredient dependencies so local recommendations do not create network-level infeasibility. Guest and brand guardrails should cap price changes, monitor complaints and ratings, and require human review for actions that may harm perceived value. Labor optimization should include service quality, compliance, overtime burden, role coverage, and schedule stability rather than treating labor only as a cost ratio.

Governance is also required. The system should minimize personally identifiable information, restrict access to sensitive labor and sales data, log decisions, record manager overrides, monitor drift, and maintain audit trails. Benchmark data should be anonymized or aggregated when external peer data are used.

11. Limitations and Threats to Validity

The framework has several limitations. First, recommendations derived from historical enterprise data remain observational unless paired with experimentation, causal effect estimation, or carefully designed quasi-experimental evaluation. Second, benchmark quality depends on cohort construction, data comparability, update frequency, and peer coverage. Third, price elasticity, substitution behavior, supplier reliability, and stockout response can drift during economic shocks, promotions, competitor actions, local events, or menu changes. Fourth, important outcomes such as guest satisfaction, brand equity, perceived fairness, sustainability, and employee well-being may be only partially captured in structured operational data. Fifth, system performance depends on recipe fidelity, inventory-count discipline, modifier mapping, vendor-item mapping, and payroll granularity. Finally, local recommendations can interact through shared commissaries, regional promotions, vendor allocations, and cross-location guest behavior.

12. Conclusion and Future Work

This paper presented a benchmark-guided prescriptive analytics framework for menu and margin optimization in multi-location restaurants. The central argument is that restaurant analytics should not end with dashboards or forecasts. Unified POS, recipe, inventory, procurement, payroll, channel, and benchmark data should be translated into action-aware predictions, peer-relative performance gaps, and constrained recommendations that specify what to change, where, why, and with what expected risk.

The framework contributes a formal decision model, an enterprise architecture, feature-engineering logic, predictive risk layers, robust benchmark-gap scoring, and a recommendation engine that ranks feasible actions by profit uplift, benchmark closure, downside risk, implementation cost, and stability. Its novelty lies in connecting peer benchmarking to practical recommendation ranking across menu, food cost, labor, inventory, and guest-risk domains.

Future work should extend the framework with causal uplift modeling, randomized pilots, receding-horizon control, reinforcement learning under safety constraints, privacy-preserving benchmark computation, and broader objective functions that include guest satisfaction, sustainability, workforce fairness, and brand equity alongside financial margin optimization.

References

[1] K. Lepenioti, A. Bousdekis, D. Apostolou, and G. Mentzas, "Prescriptive analytics: Literature review and research challenges," International Journal of Information Management, vol. 50, pp. 57-70, 2020, doi: 10.1016/j.ijinfomgt.2019.04.003.

[2] D. Bertsimas and N. Kallus, "From predictive to prescriptive analytics," Management Science, vol. 66, no. 3, pp. 1025-1044, 2020, doi: 10.1287/mnsc.2018.3253.

[3] S. Punia, S. P. Singh, and J. K. Madaan, "From predictive to prescriptive analytics: A data-driven multi-item newsvendor model," Decision Support Systems, vol. 136, Art. no. 113340, 2020, doi: 10.1016/j.dss.2020.113340.

[4] P. M. Notz and R. Pibernik, "Explainable subgradient tree boosting for prescriptive analytics in operations management," European Journal of Operational Research, vol. 312, no. 3, pp. 1119-1133, 2024, doi: 10.1016/j.ejor.2023.08.037.

[5] B. K. Chae, C. Sheu, and E. O. Park, "The value of data, machine learning, and deep learning in restaurant demand forecasting: Insights and lessons learned from a large restaurant chain," Decision Support Systems, vol. 184, Art. no. 114291, 2024, doi: 10.1016/j.dss.2024.114291.

[6] T. Tanizaki, T. Hoshino, T. Shimmura, and T. Takenaka, "Demand forecasting in restaurants using machine learning and statistical analysis," Procedia CIRP, vol. 79, pp. 679-683, 2019, doi: 10.1016/j.procir.2019.02.042.

[7] T. Tanizaki, T. Hoshino, T. Shimmura, and T. Takenaka, "Restaurants store management based on demand forecasting," Procedia CIRP, vol. 88, pp. 580-583, 2020, doi: 10.1016/j.procir.2020.05.101.

[8] B. M. Noone and G. Cachia, "Menu engineering re-engineered: Accounting for menu item substitutes in pricing and menu placement decisions," International Journal of Hospitality Management, vol. 87, Art. no. 102504, 2020, doi: 10.1016/j.ijhm.2020.102504.

[9] B. Ozdemir and O. Caliskan, "A review of literature on restaurant menus: Specifying the managerial issues," International Journal of Gastronomy and Food Science, vol. 2, no. 1, pp. 3-13, 2014, doi: 10.1016/j.ijgfs.2013.12.001.

[10] H. B. J. Lai and S. Karim, "Do-it-yourself menu management and pricing," Journal of Revenue and Pricing Management, vol. 22, no. 6, pp. 431-445, 2023, doi: 10.1057/s41272-023-00419-9.

[11] S. E. Kimes, R. B. Chase, S. Choi, P. Y. Lee, and E. N. Ngonzi, "Restaurant revenue management: Applying yield management to the restaurant industry," Cornell Hotel and Restaurant Administration Quarterly, vol. 39, no. 3, pp. 32-39, 1998, doi: 10.1177/001088049803900308.

[12] L. Bujalance-Lopez, L. Gonzalez-Serrano, M. P. Lechuga Sancho, and P. Talon-Ballestero, "Restaurant revenue management: A systematic literature review and future challenges," British Food Journal, vol. 127, no. 6, pp. 2169-2196, 2025, doi: 10.1108/BFJ-08-2024-0816.

[13] G. M. Thompson, "Labor scheduling, Part 1: Forecasting demand," Cornell Hotel and Restaurant Administration Quarterly, vol. 39, no. 5, pp. 22-31, 1998, doi: 10.1177/001088049803900507.

[14] D. Reynolds and G. M. Thompson, "Multiunit restaurant productivity assessment using three-phase data envelopment analysis," International Journal of Hospitality Management, vol. 26, no. 1, pp. 20-32, 2007, doi: 10.1016/j.ijhm.2005.08.004.

[15] D. Reynolds, "Hospitality-productivity assessment using data-envelopment analysis," Cornell Hotel and Restaurant Administration Quarterly, vol. 44, no. 2, pp. 130-137, 2003, doi: 10.1016/S0010-8804(03)90026-7.

[16] P. Alberca and L. Parte, "Operational efficiency evaluation of restaurant firms," International Journal of Contemporary Hospitality Management, vol. 30, no. 3, pp. 1959-1977, 2018, doi: 10.1108/IJCHM-09-2016-0547.

[17] Y. Liu, D. Kalaitzi, M. Wang, and C. Papanagnou, "A machine learning approach to inventory stockout

prediction," Journal of Digital Economy, vol. 4, pp. 144-155, 2025, doi: 10.1016/j.jdec.2025.06.002.

[18] A. Charnes, W. W. Cooper, and E. Rhodes, "Measuring the efficiency of decision- making units," European Journal of Operational Research, vol. 2, no. 6, pp. 429-444, 1978, doi: 10.1016/0377-2217(78)90138-8.

[19] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," Journal of Risk, vol. 2, no. 3, pp. 21-41, 2000, doi: 10.21314/JOR.2000.038.