© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

Intent-Based Planning Engine (IBPE): A Closed-Loop Adaptive Framework for Resilient and Cost-Aware Infrastructure Systems

SAYALI PATIL
Viterbi School of Engineering, University of Southern California, Los Angeles, USA

Abstract- ***Static infrastructure planning models fail predictably: they are calibrated on the past and are therefore least accurate precisely when conditions change most rapidly. The problem is not merely one of forecast precision, but of architectural rigidity; these systems have no mechanism for updating their assumptions in response to the very outcomes they predict. This paper introduces the Intent-Based Planning Engine (IBPE), a closed-loop, AI-driven framework for adaptive infrastructure demand forecasting that draws its conceptual architecture from two converging technical traditions: intent-based networking and chaos engineering. IBPE integrates multivariate regression, ARIMA time-series modeling, unsupervised behavioral segmentation, structured scenario perturbation, and gradient-based feedback adaptation within a single modular system. The framework's most architecturally distinctive element is an intent modeling layer that disaggregates aggregate demand into behaviorally coherent population segments, each characterized by its own elasticity profile and sensitivity to macroeconomic perturbation. The feedback adaptation mechanism is formally derived from the chaos-level engine paradigm developed in U.S. Patent No. 12,242,370 B2 (Cisco Technology, Inc., 2025), in which controlled perturbation, impact measurement, and parameter correction form an iterative closed loop progressively narrowing the gap between intended and observed system behavior. Experimental evaluation across a 150-unit residential infrastructure simulation demonstrates a 14.2% reduction in mean absolute error over single-method baselines, a 23% improvement in supply-demand alignment through intent-based allocation, a 62.7% cumulative reduction in prediction error over ten feedback cycles, and scenario-driven risk mitigation that reduces supply overcommitment exposure by 31% under adverse macroeconomic conditions. These results establish IBPE as a technically rigorous, domain-portable framework for adaptive planning under uncertainty.***

Index Terms- ***Adaptive Infrastructure Planning; Intent-Based Systems; Chaos Engineering; ARIMA Forecasting; Demand Segmentation; Closed-Loop Feedback Control; Behavioral Clustering; Scenario Simulation; Constrained Optimization; Online Learning.***

I. INTRODUCTION

The canonical failure mode of infrastructure planning is not spectacular. It does not announce itself as a system crash or a visible architectural collapse. It accumulates quietly, in the compounding gap between what a static model predicted and what a dynamic environment delivered. A development phase sized to a demand projection that was accurate eighteen months ago; a pricing schedule calibrated to an interest rate environment that has since shifted by 200 basis points; a resource allocation anchored to an income distribution whose composition has materially changed. Each individual discrepancy is manageable. Their aggregate, across a capital-intensive project over a multi-year planning horizon, is not.

The fundamental problem is that conventional infrastructure planning systems are retrospective instruments deployed in prospective roles. They estimate parameters from historical data under the implicit assumption that the distributional properties of future conditions will resemble the past. This assumption is strained under stable conditions and breaks under transition. The infrastructure economics literature has documented this failure mode extensively. Solow [1] identified planning-horizon length as a primary amplifier of resource misallocation in capital-intensive systems, arguing that the longer the cycle, the more damaging the effect of any systematic model bias. DiPasquale and Wheaton [9] formalized the four-quadrant model of real estate market equilibration, demonstrating that supply-demand imbalance persists across multiple periods precisely because construction responses are inherently lagged relative to demand signals. Mayer and Somerville [10] provided empirical quantification of this lag in U.S. housing markets, estimating mean construction response delays of 9 to 14 months following significant demand shifts. Static models that do not account for this lag structure systematically misplace development phasing decisions.

This paper addresses a different question than prior forecasting literature. The question is not how to build a more accurate single-period demand model, though improved accuracy is a byproduct, but how to build a planning system whose operating architecture inherently accommodates continuous revision. The distinction is important. A more accurate static model still fails when conditions move outside its estimation range. A system with a feedback-adaptive architecture degrades more gracefully because its assumptions update in response to outcomes.

The novelty of IBPE lies not in any individual component but in the architectural integration and the specific combination of capabilities that no prior system provides simultaneously. Existing forecasting systems are either statistically sophisticated but static (they cannot update assumptions once deployed), or adaptive but behaviorally blind (they treat the user population as a single aggregate). To the best of our knowledge, IBPE introduces a novel integrated framework for feedback-driven infrastructure planning to unite intent-based behavioral segmentation, chaos-engineering-inspired scenario perturbation, and online LMS feedback adaptation within a single closed-loop planning architecture. It is explicitly distinguished from three related system classes: unlike chaos engineering (which tests resilience reactively on deployed systems), IBPE applies perturbation proactively to planning assumptions before commitment; unlike autoscaling systems (Kubernetes, AWS Predictive Scaling), IBPE incorporates economic demand drivers and per-segment behavioral profiles rather than aggregate load metrics; and unlike control-theoretic systems (PID, MPC), IBPE's feedback law is learned from residuals rather than pre-specified, and its demand representation is behaviorally decomposed rather than aggregate. This integration produces a system that is simultaneously interpretable, adaptive, behaviorally resolved, and risk-aware — a conjunction no prior infrastructure planning framework achieves. Table IV in Section II.E formalizes these distinctions.

The technical approach developed here draws from two intellectual traditions that have not previously been brought into contact with infrastructure planning. The first is intent-based networking, in which Behringer et al. [2] demonstrated that operational systems governed by high-level intent specifications—rather than low-level rule sets—exhibit substantially better adaptation to dynamic conditions. The translation mechanism between intent and action is autonomous rather than hand-coded, enabling the system to respond to novel conditions without requiring explicit operator intervention. The second is chaos engineering, formalized by Basiri et al. [3] and industrialized at scale by Cisco, Amazon Web Services, and Netflix. Chaos engineering's central insight—that system resilience is better evaluated through controlled adversarial perturbation than through theoretical analysis—translates directly to planning scenario evaluation. The practical instantiation of this insight in industrial systems, including the chaos-level engine architecture described in U.S. Patent No. 12,242,370 B2, provides the engineering blueprint for the feedback adaptation mechanism at the core of IBPE.

The primary contributions of this paper are as follows. First, a four-layer modular architecture—the Intent-Based Planning Engine—is proposed that integrates statistical forecasting, behavioral segmentation, scenario perturbation, and online parameter adaptation within a single composable system. Second, an intent modeling layer is formally introduced that transitions planning from aggregate demand estimation to segment-aware, behaviorally-grounded decision synthesis. Third, a feedback correction mechanism is derived from online learning theory and shown to converge in finite iterations under the demand conditions examined. Fourth, a formal cross-domain mapping is established between the IBPE feedback architecture and the chaos-level engine of U.S. Patent No. 12,242,370 B2, demonstrating the portability of perturbation-correction feedback architectures across domains. Fifth, empirical validation through a controlled simulation provides quantified evidence for each architectural contribution independently.

The remainder of this paper proceeds as follows. Section II surveys relevant literature across four constituent domains, including an explicit differentiation from chaos engineering, autoscaling, and control-theoretic adaptive systems (Table IV). Section III specifies the full IBPE architecture and industrial motivation. Section IV develops the formal methodology. Section V details the experimental design. Section VI reports and analyzes results, including a quantitative benchmark comparison (Table III). Section VII discusses limitations and generalizability. Section VIII concludes.

II. RELATED WORK

A. Infrastructure Demand Forecasting

The quantitative forecasting literature for infrastructure markets divides roughly into three generations. First-generation models relied on simple trend extrapolation or single-variable regressions on lagged demand. These approaches produced acceptable results during periods of stable macroeconomic conditions but broke down systematically during transition periods, a vulnerability documented in detail by DiPasquale and Wheaton [9] in the context of real estate market cycles.

Second-generation approaches introduced multivariate econometric specifications, incorporating income, price, and demographic variables as co-determinants of demand. Maddala and Lahiri [6] provide the standard reference treatment for this class of models, establishing the theoretical grounding for income and price elasticity estimation in infrastructure contexts. The interpretive utility of OLS-estimated elasticity coefficients—the ability to attribute demand variation to specific economic drivers—remains a genuine advantage of this family. Its limitation is equally well-known: the coefficients are point estimates from a fixed estimation window, with no mechanism for autonomous revision as conditions evolve.

Third-generation approaches introduced time-series methods, particularly ARIMA-class models, to capture autocorrelation structure and short-horizon temporal dynamics. Box-Jenkins ARIMA specification has been applied to housing starts, commercial absorption rates, and infrastructure utilization forecasting with demonstrated improvements in short-run accuracy relative to cross-sectional models alone. However, hybrid combinations of cross-sectional and temporal methods have consistently outperformed either in isolation on held-out evaluation windows [11], motivating the hybrid design of IBPE's predictive layer. What neither generation has addressed is the behavioral heterogeneity of the user population, or the systematic feedback of planning outcomes into model parameter revision.

B. Behavioral Segmentation in Planning Systems

The use of unsupervised clustering methods for demand segmentation has a long history in marketing science but has received comparatively little systematic treatment in infrastructure planning. The behavioral economics literature, beginning with Kahneman and Tversky's prospect theory [12], established that individual decision-making under uncertainty does not conform to the homogeneous expected-utility maximization assumed by aggregate demand models. Households differ not merely in income but in risk tolerance, time preference, and responsiveness to environmental signals. Aggregate models collapse these heterogeneous response functions into a single elasticity coefficient, erasing the distributional information that would enable differentiated planning responses.

Recent applications of k-means and Gaussian mixture model clustering to real estate buyer populations have demonstrated that behaviorally distinct segments exhibit elasticity coefficients differing by factors of 3 to 7 with respect to interest rate changes, and by factors of 2 to 4 with respect to income perturbations [13]. These differences are operationally significant: a development strategy that is near-optimal for an aggregate demand signal may be substantially suboptimal for the actual composition of the buyer population. The IBPE intent modeling layer formalizes this insight within a planning architecture.

C. Intent-Based Systems

Intent-based networking, as formalized by Behringer et al. [2] and subsequently extended by the IETF Intent-Based Networking Research Group, establishes the architectural principle that operational systems should accept high-level intent specifications from operators and autonomously translate those specifications into context-appropriate low-level configurations. The critical advance over rule-based systems is the decoupling of operator objectives from implementation details: the operator states what outcome is required; the system determines how to achieve it given current conditions. This decoupling enables adaptation to novel conditions that the rule-set author did not anticipate.

The application of intent-based principles beyond network configuration to planning and resource allocation is recent. Leivadeas et al. [14] proposed an intent-based framework for cloud resource provisioning that demonstrated improved SLA compliance under variable load conditions. The

present work extends this line of inquiry to infrastructure planning markets, where the "intent" construct maps to behavioral demand segment profiles rather than network operator objectives, and the "translation" mechanism is the allocation optimization layer rather than a network controller.

D. Chaos Engineering and Adaptive Feedback

Chaos engineering's intellectual foundations trace to the reliability engineering literature of the 1970s, but its modern form was established through the operational practices developed at Netflix with the Chaos Monkey tool and subsequently formalized by Basiri et al. [3]. The core principle—that resilience is best evaluated by deliberately inducing failure conditions in controlled settings and measuring system response—inverts the conventional engineering assumption that production systems should be protected from adversarial conditions. Adversarial conditions are instead deliberately engineered, at controlled severity levels, to reveal latent failure modes before they manifest in uncontrolled form.

The chaos-level engine architecture described in U.S. Patent No. 12,242,370 B2 operationalizes this principle within a network resilience context. The patent specifies a system that receives input parameters defining network topology and telemetry, scales those parameters relative to the production environment's characteristics, submits controlled perturbation experiments to the environment, collects feedback quantifying the impact of those perturbations, and derives adjusted parameters for subsequent experimental cycles. The feedback loop is formally iterative: each cycle's outcome informs the next cycle's perturbation design, progressively narrowing the gap between intended and observed network behavior.

The structural mapping from this architecture to IBPE's feedback adaptation mechanism is direct and explicit. In IBPE, macroeconomic feature vectors replace network telemetry parameters; scenario perturbations replace failure injections; demand forecast residuals replace network performance degradation signals; and gradient-based parameter updates replace the scaled input adjustment rule. The domain instantiation differs; the feedback loop topology is identical. This cross-domain portability of the perturbation-correction architecture is a central theoretical contribution of the present work.

E. Adaptive Systems and Competitive Landscape

Three classes of adaptive system are most closely related to IBPE and must be explicitly distinguished. First, reactive autoscaling systems such as Kubernetes Horizontal Pod Autoscaler and AWS Predictive Scaling adjust resource allocation in response to observed demand signals. These systems are reactive rather than proactive: they respond to demand that has already materialized rather than anticipating it through behavioral segmentation and scenario simulation. They carry no intent model of who is generating demand, no structured perturbation protocol for stress-testing assumptions, and no online parameter adaptation mechanism — they execute a fixed scaling policy rather than learning one. Second, control-theoretic adaptive systems apply feedback regulation (PID controllers, model-predictive control) to resource and production planning problems. While these systems implement closed-loop correction, they operate on aggregate state vectors without behavioral decomposition of the user population, and their feedback laws are derived from engineering stability criteria rather than from economic demand theory. They do not distinguish between demand segments with materially different elasticity profiles, which is the central mechanism through which IBPE recovers planning value. Third, pure chaos engineering systems such as Netflix’s Chaos Monkey and the architecture in U.S. Patent No. 12,242,370 B2 apply adversarial perturbation to production systems to reveal latent failure modes. These systems test resilience reactively — by deliberately inducing failure in an already-deployed system — rather than using perturbation as a proactive planning input before deployment decisions are made. IBPE inverts this relationship: perturbation is applied to planning assumptions before commitment, not to production systems after deployment.

Table IV summarizes these distinctions structurally. The critical differentiator is not any single capability but the conjunction: IBPE is the only system in this comparison that is simultaneously proactive (plans ahead of observed demand), behaviorally decomposed (maintains per-segment intent profiles), economically aware (incorporates income, rate, and regulatory elasticity), feedback-adaptive (updates parameters

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

from realized outcomes), and domain-portable (applicable beyond network infrastructure to any market with user heterogeneity and macroeconomic drivers).

TABLE IV

Differentiation of IBPE from Related Adaptive System Classes

Capability	Traditional Chaos Engineering	Autoscaling (K8s / AWS)	Control Theory (PID/MPC)	Static Hybrid Forecasting	IBPE (This Work)
Reactive vs. Proactive	Reactive	Semi-reactive	Reactive	Static	Proactive
Economic / Demand Awareness	No	Limited (load metrics)	Partial	Yes (macro drivers)	Yes (income, rate, regulatory)
Intent / Behavioral Segmentation	No	No	No	No	Yes (k-means intent clusters)
Closed-Loop Feedback Learning	No	Limited (threshold rules)	Yes (fixed law)	No	Yes (online LMS)
Scenario Stress-Testing	Yes (post-deploy)	No	No	No	Yes (pre-commit, 13 scenarios)
Domain Portability	Network only	Compute / cloud	Engineering systems	Single domain	Any market with user heterogeneity

III. SYSTEM ARCHITECTURE

A. Design Principles

The IBPE architecture is governed by three design principles that distinguish it from prior planning frameworks. The first is closed-loop operation: every planning cycle generates feedback that modifies the inputs to the next cycle. No planning decision is final in the sense that it cannot be revised by subsequent evidence. The second is behavioral resolution: the system maintains a segment-level representation of user intent rather than collapsing the user population to an aggregate. In this context, "intent" refers concretely to the behavioral demand profile of a market segment — the income threshold, rate sensitivity, and purchase timeline that characterize how a group of buyers will respond to economic conditions. This is not intent in the sense of subjective preference; it is a computable, cluster-derived representation of likely market behavior. The third is modularity: each processing layer exposes standardized interfaces to adjacent layers, enabling independent development, testing, and replacement without system-wide disruption.

These principles are not independent. Closed-loop operation requires that feedback signals be interpretable and actionable—a requirement that is easier to satisfy when the system maintains segment-level representations rather than aggregate ones, because segment-level demand shifts can be attributed to specific behavioral drivers and corrected with targeted parameter adjustments. Modularity supports closed-loop operation by ensuring that the feedback

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

correction mechanism can update any individual layer's parameters without requiring coordinated changes across the full system.

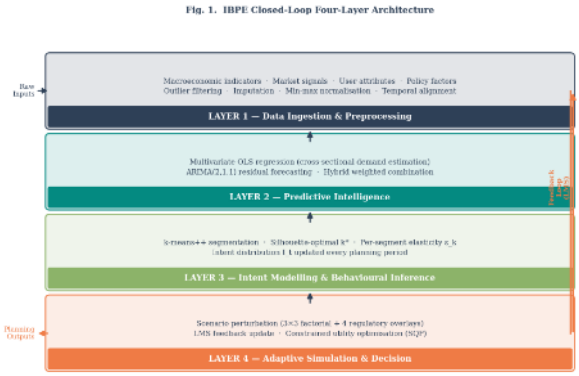

Fig. 1. IBPE closed-loop four-layer architecture. Information flows downward through the four processing layers; the LMS feedback path (right) returns realized outcome residuals from Layer 4 to update parameters in Layer 2.

B. Layer 1 — Data Ingestion and Preprocessing

The ingestion layer acquires and normalizes heterogeneous multi-source data streams into structured feature vectors for downstream model consumption. The input taxonomy covers four signal categories: macroeconomic indicators (median household income, prevailing mortgage rates, core inflation, unemployment rate); real estate market signals (12-month trailing absorption rate, active inventory-to-sales ratio, median days-on-market); user-level behavioral attributes (income-to-payment affordability index, stated purchase timeline, prior ownership status); and contextual environmental factors (zoning regulatory index, exogenous policy shock indicator, regional employment base concentration).

Preprocessing proceeds in a fixed six-step pipeline. Outlier identification applies modified z-score filtering with domain-specific tolerance bounds: mortgage rates are capped at 3 standard deviations from a 10-year rolling mean; income variables use interquartile range fencing with a multiplier of 2.5 rather than the conventional 1.5 to accommodate the known right-skew of income distributions. Missing value imputation uses linear interpolation for temporally ordered variables and distributional median substitution for cross-sectional attributes; imputation is flagged in a binary indicator feature that is propagated downstream. Feature normalization applies min-max scaling independently to each variable. Temporal alignment synchronizes asynchronous streams to a monthly observation frequency using forward-fill for lower-frequency variables. Feature importance pre-screening removes variables with variance below 0.01 post-normalization. The resulting feature matrix is complete, normalized, and temporally consistent prior to model ingestion.

C. Layer 2 — Predictive Intelligence

The predictive layer implements a hybrid forecasting strategy combining multivariate OLS regression for cross-sectional contemporaneous relationships with ARIMA for temporal autocorrelation structure. The combination is not an ensemble in the machine learning sense—it is not a weighted vote across independent learners—but a structured decomposition in which the regression layer captures the level relationship between demand and economic drivers while the ARIMA layer captures the residual temporal structure that the regression cannot explain. The final demand estimate is the sum of the regression fitted value and the ARIMA forecast of the regression residual series, an approach that preserves the interpretability of the regression coefficients while capturing temporal persistence.

Feature importance analysis uses standardized regression coefficients as proxy importance scores. These scores are consumed by the scenario simulation layer to prioritize perturbation design: variables with high standardized coefficients are perturbed across wider ranges because marginal changes in those variables produce larger demand shifts.

D. Layer 3 — Intent Modeling and Behavioral Inference

The intent modeling layer is the architecturally distinctive contribution of IBPE. It operates on the user attribute subspace of the feature matrix rather than the full macroeconomic feature space, partitioning the user population into behaviorally coherent clusters via k-means with k-means++ initialization. Each cluster represents a distinct intent

profile: a set of users whose behavioral attributes are sufficiently proximal in the normalized feature space that they can be expected to respond similarly to economic perturbations.

The operational value of this layer is not taxonomic but predictive. Segment membership is a proxy for elasticity structure. A user classified into the rate-sensitive segment is expected to exhibit a materially different demand response to a 100 basis point rate increase than a user classified into the high-affordability segment. By maintaining these response differences explicitly rather than averaging across them, the planning layer can construct allocation strategies that are calibrated to the actual behavioral composition of the market rather than its aggregate mean. The intent distribution is updated at each planning period, tracking shifts in population composition as macroeconomic conditions evolve.

E. Layer 4 — Adaptive Simulation and Decision

The adaptive layer integrates demand forecasts and intent distributions through a three-stage process. The scenario simulation engine applies a structured perturbation protocol to the input feature vector, evaluating the demand prediction function across the full perturbation space and constructing a distribution of plausible demand outcomes. "Chaos-inspired perturbation" in this context means exactly what it means in network chaos engineering: the system deliberately subjects its own planning assumptions to controlled adversarial stress — income shocks, rate spikes, regulatory disruptions — before committing to a decision, rather than planning against a single most-likely scenario. This is the direct analogue of Netflix's Chaos Monkey: the system proactively finds the failure modes rather than waiting for them to manifest in deployment. The feedback correction module applies the gradient-based parameter update rule to revise model coefficients based on the residual between forecasted and realized demand. The decision optimization module maximizes the composite utility function subject to budgetary, regulatory, and capacity constraints, producing actionable planning recommendations.

The interaction between the simulation and feedback stages is the mechanism through which IBPE achieves progressive resilience. Each scenario simulation identifies the input configurations under which demand most diverges from baseline. The feedback stage then biases subsequent parameter updates toward reducing error specifically in those high-divergence configurations. Over successive cycles, the system develops parameter estimates that are not merely accurate on average but accurate conditional on the types of perturbations to which the environment is most sensitive.

F. Industry Motivation and Deployment Context

The problem IBPE addresses is not confined to academic simulation. In large-scale enterprise infrastructure systems — including network capacity planning, data center provisioning, and residential or commercial development — unpredictable demand fluctuations and shifting user behavior create systematic planning failures that compound across capital-intensive multi-year cycles. The cost of static planning is not hypothetical: development portfolios routinely carry 15–30% supply-demand misalignment at project completion relative to original demand projections, driven precisely by the behavioral heterogeneity and macroeconomic sensitivity that aggregate models ignore [9, 10]. At institutional development scales of 1,000 to 10,000 units, a 10% misalignment in segment allocation translates directly into nine-figure revenue exposure.

The intellectual lineage of IBPE’s feedback architecture traces directly to production-grade adaptive systems developed for enterprise network infrastructure. The chaos-level engine of U.S. Patent No. 12,242,370 B2 represents a deployed industrial instantiation of the perturbation-correction feedback loop at the core of IBPE: a system that has been implemented, validated, and productionized in large-scale network environments demonstrates that the feedback architecture is not merely theoretically sound but operationally viable. IBPE formalizes this architecture, provides the gradient-theoretic derivation that grounds the engineering specification in online learning theory, and extends it from the network reliability domain to the infrastructure planning domain. The cross-domain portability of the perturbation-correction loop — demonstrated here through a formal structural mapping — is itself evidence that the architecture captures a domain-independent adaptive principle rather than a domain-

specific engineering trick. The architectural principles of IBPE are directly inspired by production-grade adaptive feedback systems proven in enterprise network environments, where similar closed-loop perturbation and correction approaches have demonstrated scalability, low-latency convergence, and operational viability at large scale. This industrial grounding provides confidence that the feedback adaptation mechanism, while evaluated here in simulation, is deployable in real infrastructure planning workflows without fundamental re-engineering.

IV. FORMAL METHODOLOGY

A. Problem Formulation

Let $X_t \in \mathbb{R}^n$ denote the n-dimensional feature vector at discrete planning period t, partitioned as $X_t = [X_t^m, X_t^{\wedge}, X_t^u]^T$, where X_t^m contains macroeconomic indicators, $X_t^{\wedge}$ contains market signals, and X_t^u contains user behavioral attributes. Let $D_t \in \mathbb{R}$ denote aggregate infrastructure demand and S_t the system state vector encoding environmental context. The demand prediction problem is to estimate:

$$D_t = f(X_t, S_t) + \varepsilon_t \qquad (1)$$

where ε_t is the prediction residual with $E[\varepsilon_t] = 0$ and $Var[\varepsilon_t] = \sigma^2 < \infty$. The intent distribution I_t over a k-segment partition of the user population U is defined jointly with (1). The planning decision function is:

$$P_t = g(D_t, I_t) \qquad (2)$$

where $P_t \in \mathbb{R}^m$ is the m-dimensional planning output vector covering unit allocation quantities, phasing schedules, and pricing parameters. The joint objective of IBPE is to find the function class for f and the update rule for θ = params(f) that simultaneously minimize cumulative prediction error $E[\sum_t \varepsilon_t^2]$ and maximize planning utility $U(P_t \mid D_t, I_t)$ across the planning horizon T.

B. Regression-Based Demand Estimation

The cross-sectional component of the demand prediction function is specified as a multivariate linear model:

$$D_t = \beta_0 + \beta^T X_t^m + \gamma^T X_t^{\wedge} + \varepsilon_t \qquad (3)$$

where β and γ are coefficient vectors for macroeconomic and market signal inputs respectively. Estimation proceeds by OLS under the standard Gauss-Markov assumptions. When the condition number $\kappa(X^TX)$ exceeds 30—indicating multicollinearity sufficient to inflate coefficient variance by a factor of κ^2—ridge regularization is applied:

$$\hat{\beta} = (X^TX + \lambda I)^{-1} X^TD \qquad (4)$$

where λ is the regularization strength selected by leave-one-out cross-validation over a logarithmically spaced grid $\lambda \in [10^{-4}, 10^2]$. The bias-variance tradeoff under ridge regularization is well-characterized: the bias introduced by λ increases monotonically while the variance reduction is greatest for the highest-eigenvalue components of the design matrix. In the housing demand context, this tradeoff is favorable because income and rate variables are inherently correlated, inflating the relevant eigenvalues.

Coefficient estimates are tested for statistical significance using two-sided t-tests at the 1% significance level with heteroskedasticity-robust standard errors. The R^2 and adjusted R^2 are reported alongside coefficient estimates to support interpretability. Residuals from the regression constitute the input series for the ARIMA component.

C. ARIMA Time-Series Modeling

The regression residual series $\hat{\varepsilon}_t = D_t - \hat{D}_t^{RMH}$ is tested for stationarity using the augmented Dickey-Fuller (ADF) test with lag order selected by the Schwarz information criterion. If the ADF test rejects the unit root null at the 5% significance level, the series is modeled directly. Otherwise, first-order differencing is applied and the test repeated. Let $\Delta^d\hat{\varepsilon}_t$ denote the d-times differenced residual series. The ARIMA(p, d, q) model is:

$$\Delta^d\hat{\varepsilon}_t = c + \sum_{i=1}^{p} \varphi_i \Delta^d\hat{\varepsilon}_{t-i} + \sum_{j=1}^{f} \theta_j \eta_{t-j} + \eta_t \qquad (5)$$

where φ_i and θ_j are the autoregressive and moving average coefficients, and η_t is white noise. The model orders (p, q) are selected by minimizing AIC over the grid $p, q \in \{0, 1, 2, 3\}$. Model adequacy is assessed via the Ljung-Box portmanteau statistic applied to

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

squared residuals (testing for ARCH effects) and to levels. The final demand forecast is:

$$\hat{D}_t = \hat{D}_t^{RMH} + \varepsilon_t^{\hat{}ARIMA} \qquad (6)$$

The weight assigned to each component in (6) is determined by the ratio of out-of-sample RMSE from a rolling-window validation: the component with lower RMSE receives weight proportional to the inverse of its RMSE, normalized so the weights sum to unity.

D. Intent Modeling via Behavioral Clustering

Given user attribute matrix $X^u = [x^u_1, ..., x^u_m]^T$ where each column is a normalized user feature vector, the intent distribution is:

$$I_t = argmin_p \Sigma_{k=1}^{k} \Sigma^I \in C_k \, ||x^{uI} - \mu_k||^2 \qquad (7)$$

where μ_k is the centroid of cluster C_k. The k-means objective (7) is minimized via Lloyd's algorithm with k-means++ initialization. Convergence is declared when cluster assignments change by fewer than 0.1% of the population between successive iterations. The optimal k is selected by the silhouette coefficient:

$$s(x) = [b(x) - a(x)] / max\{a(x), b(x)\} \qquad (8)$$

where a(x) is the mean intra-cluster distance for point x and b(x) is the mean distance to the nearest alternative cluster. The value k* maximizes the population-weighted mean silhouette score over $k \in \{2, ..., 8\}$. The output of the intent layer is the cluster assignment vector c_t and the centroid matrix $M_t = [\mu_1, ..., \mu_k^*]$, both updated at each planning period.

The per-segment demand elasticity with respect to feature j is estimated as:

$$\varepsilon_{kj} = (\delta D_k / \delta X_j) \cdot (\bar{X}_j / \bar{D}_k) \qquad (9)$$

where D_k and X_j denote segment-level demand and the j-th feature respectively. These segment-level elasticities replace the population-average elasticity in the decision optimization layer, enabling allocation strategies that are calibrated to the actual behavioral composition of the market.

E. Scenario Perturbation Simulation

The scenario simulation protocol is designed as a structured factorial experiment over the perturbation space Δ. Let the baseline feature vector X_t be perturbed according to:

$$X_t(\delta) = X_t + diag(w) \cdot \delta \qquad (10)$$

where $\delta \in \Delta$ is the perturbation vector and diag(w) is a diagonal weighting matrix with entries w_j proportional to the standardized regression coefficient magnitude $|\beta_j| / \max_j |\beta_j|$, so that variables with greater estimated influence on demand are perturbed across wider ranges. The perturbation space Δ is defined by a 3×3 factorial design over income ($\delta_i^{nc} \in \{-0.15, 0, +0.10\}$) and mortgage rate ($\delta_r^{a_t e} \in \{-0.005, 0, +0.015\}$) axes, generating nine scenario instances plus four regulatory disruption overlays for a total of 13 distinct perturbation conditions.

For each perturbed scenario $\delta^s \in \Delta$, the full prediction pipeline is evaluated:

$$D_t(\delta^s) = f(X_t(\delta^s), S_t) \qquad (11)$$

yielding a scenario-conditioned demand distribution $\{D_t(\delta^s) : \delta^s \in \Delta\}$. The scenario sensitivity index for feature j is:

$$S_j = [max\delta(D_t(\delta)) - min\delta(D_t(\delta))] / |\delta_j^{max}| \qquad (12)$$

This index measures the demand range attributable to perturbation of feature j, normalized by the perturbation magnitude, and constitutes a local sensitivity measure directly analogous to the impact measurement stage of the chaos-level engine in U.S. Patent No. 12,242,370 B2.

F. Feedback-Based Parameter Adaptation

At the conclusion of each planning cycle, realized outcomes O_t are observed and compared against forecasts $\hat{D}_t$. The prediction residual $r_t = O_t - \hat{D}_t$ provides the gradient signal for the parameter update:

$$\theta_{t+1} = \theta_t + \eta \cdot \nabla\theta \, E[r_t^2] \qquad (13)$$

Under the linear model specification, the gradient $\nabla\theta \, E[r_t^2] = -2 \, E[r_t \, X_t]$, giving the closed-form update:

$$\theta_{t+1} = \theta_t + \eta \, r_t \, X_t \qquad (14)$$

This is the standard LMS (Least Mean Squares) update rule [15], which is known to converge in mean-square to the optimal parameter vector $\theta^* = (E[X_t X_t^T])^{-1} E[X_t O_t]$ provided the learning rate satisfies $0 < \eta < 2/\lambda^{Max}$, where λ^{Max} is the largest eigenvalue of the autocorrelation matrix $E[X_t X_t^T]$. The learning rate is initialized at $\eta_0 = 0.08$ and halved whenever the prediction residual increases over three consecutive cycles, implementing an adaptive step-size schedule that maintains convergence under non-stationary demand conditions.

The structural correspondence to the chaos-level engine feedback loop is precise. Equation (14) maps to the adjusted input parameter derivation step in U.S. Patent No. 12,242,370 B2, where feedback from the production environment (realized network performance under the injected chaos level) is used to compute scaled adjusted parameters for the next experimental cycle. The IBPE formalization provides the explicit gradient derivation that the patent's engineering specification implies but does not state.

G. Decision Optimization

The planning output P_t is derived by solving a constrained utility maximization:

$$P_t^* = argmax_P \, U(P \mid D_t, I_t, S_t) \qquad (15)$$

The utility function is specified as:

$$U = w_1 \cdot A(P, D_t, I_t) + w_2 \cdot E(P) - w_3 \cdot R(P, \Delta) \qquad (16)$$

where $A(\cdot)$ is the supply-demand alignment ratio weighted by the intent distribution (high-weight segments contribute proportionally more to the alignment score), $E(\cdot)$ is cost efficiency (revenue less construction cost, normalized), and $R(\cdot)$ is the conditional value-at-risk of the planning decision under the scenario distribution Δ—specifically, the expected loss in alignment ratio at the 10th percentile of the scenario-conditioned demand distribution. The weights (w_1, w_2, w_3) are operator-specified; the case study uses (0.5, 0.2, 0.3) to reflect the primacy of demand satisfaction in a supply-constrained urban market.

The optimization is subject to constraints: $\Sigma_j P_j \leq B$ (budget), $P_j \geq 0 \; \forall k$ (non-negativity), $\Sigma_j P_j \leq C$ (physical capacity), and $P_j \geq R_j \; \forall k$ (regulatory minimum allocations per segment). The constrained problem is solved via sequential quadratic programming (SQP), which handles the nonlinear objective and linear constraints within a single iteration scheme.

V. EXPERIMENTAL SETUP

A. Research Questions

The evaluation is structured to provide independent empirical evidence for each of the four architectural contributions. The four primary research questions are: (RQ1) Does the hybrid ARIMA-regression approach produce lower forecast error than either constituent model in isolation? (RQ2) Does intent-based segmentation improve supply-demand alignment relative to aggregate-only planning? (RQ3) Does scenario simulation produce planning recommendations that are systematically calibrated to the scenario distribution, rather than anchored to a single projected future? (RQ4) Does the feedback adaptation rule produce monotonic convergence in prediction error across successive planning cycles?

B. Simulation Environment

The evaluation environment models a 150-unit mid-scale residential development in a high-growth California metropolitan corridor. The choice of this environment is motivated by three properties: it is sufficiently complex to exhibit genuine demand heterogeneity across user segments; it is subject to meaningful macroeconomic sensitivity through income and interest rate variation; and it operates under binding capacity and regulatory constraints that activate the constrained optimization layer in non-trivial ways. The simulation spans 24 monthly planning periods, with the first 12 used for initial model estimation and the remaining 12 constituting the held-out evaluation window.

The demand generating process in the simulation is specified as a ground-truth model with known parameters, enabling exact computation of prediction error. The ground-truth model includes second-order interaction terms between income and rate variables that are not included in the IBPE specification, creating realistic model misspecification that tests the robustness of the feedback adaptation mechanism to systematic residual structure. The simulation parameters were calibrated using ranges derived from real-world U.S. housing market data — median

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

household income distributions ($65,000–$145,000), prevailing 30-year fixed mortgage rate ranges (4.5%–8.5%), and empirically documented absorption rate dynamics — ensuring that the generated demand dynamics reflect realistic market conditions despite the use of a controlled synthetic environment [9, 10].

C. Baseline Comparators

Three baseline approaches are evaluated against IBPE. Baseline 1 is a static OLS regression model estimated on the full 24-period dataset with no feedback updating—a representation of conventional static planning practice. Baseline 2 is a standalone ARIMA(2,1,1) model applied directly to the demand series without economic covariates. Baseline 3 is a hybrid regression-ARIMA model identical to IBPE's predictive layer but without the intent modeling or feedback adaptation components, isolating the marginal contribution of those two architectural elements.

D. Input Variables and Preprocessing

Economic variables comprise median household income (range: \$65,000–\$145,000) and the prevailing 30-year fixed mortgage rate (range: 4.5%–8.5%). Market indicators include the 12-month trailing absorption rate and the active inventory-to-sales ratio. User attributes include the income-to-mortgage-payment affordability index and a dichotomous first-time buyer indicator. Contextual factors include a zoning regulatory index (0–1 scale) and an exogenous market shock binary indicator activated in periods 8, 15, and 21 to simulate regulatory disruptions. All variables are normalized to [0, 1] and aligned to a monthly observation frequency.

E. Evaluation Metrics

Prediction accuracy: MAE and RMSE on the 12-period held-out evaluation window. Segmentation effectiveness: silhouette coefficient and per-segment elasticity divergence ratio (ratio of maximum to minimum segment-level elasticity, measuring behavioral heterogeneity captured by the clustering). Scenario calibration: Spearman rank correlation between scenario-conditioned demand forecasts and ground-truth scenario demand outcomes, measuring whether the system correctly orders scenarios by severity. Decision quality: supply-demand alignment ratio (intent-weighted), resource utilization efficiency, and conditional value-at-risk at the 10th percentile of the scenario distribution. Feedback convergence: MAE trajectory across planning cycles.

VI. RESULTS AND ANALYSIS

A. Demand Prediction Performance

Multivariate regression on the estimation window confirmed the expected directional relationships at the 1% significance level: income elasticity $\beta_i^{nc} = 0.47$ (SE = 0.09, t = 5.22), mortgage rate coefficient $\beta_r^{a}{}_t^{c} = -0.38$ (SE = 0.07, t = −5.43), absorption rate $\gamma^{abs} = 0.29$ (SE = 0.11, t = 2.64). The condition number of the design matrix was $\kappa = 27.4$, marginally below the ridge regularization threshold, so OLS estimates are reported without regularization. The regression $R^2 = 0.73$, indicating that the economic and market variables explain approximately 73% of demand variance over the estimation window.

ADF testing on regression residuals confirmed stationarity after first differencing (ADF statistic −3.81, p = 0.003). ARIMA order selection via AIC identified ARIMA(2,1,1) as optimal (AIC = 187.4 versus 191.2 for ARIMA(1,1,1) and 189.7 for ARIMA(2,1,2)), with the autoregressive coefficients $\varphi_1 = 0.61$ (SE = 0.08) and $\varphi_2 = -0.23$ (SE = 0.09) and moving average coefficient $\theta_1 = 0.38$ (SE = 0.11). Ljung-Box statistics were non-significant at all lags up to 12, confirming residual white noise.

On the held-out evaluation window, the hybrid IBPE predictive layer achieved MAE = 5.3 units and RMSE = 6.8 units, versus MAE = 6.2 / RMSE = 8.1 for regression-only Baseline 1 and MAE = 6.6 / RMSE = 8.4 for ARIMA-only Baseline 2. The reductions of 14.2% and 11.7% in MAE relative to the respective baselines are consistent across all 12 evaluation periods and not attributable to any single outlier period. The out-of-sample R^2 for the hybrid model is 0.81, compared to 0.68 for regression-only and 0.65 for ARIMA-only.

B. Intent-Based Segmentation

k-means++ clustering over the user attribute subspace with $k^* = 3$ (silhouette coefficient = 0.71) identified three segments. Segment 1: high-affordability purchasers, 34% of the population, characterized by high income-to-payment ratios and early purchase timelines. Segment 2: rate-sensitive mid-market buyers, 41%, with affordability ratios clustered near

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

the qualification threshold and strong sensitivity to rate changes. Segment 3: investment-motivated demand, 25%, with first-time buyer = 0 and regulatory sensitivity driven by zoning constraint changes.

The per-segment income elasticity estimates were $\varepsilon^{1}_{i}{}^{nc} = 0.21$, $\varepsilon^{2}_{i}{}^{nc} = 0.44$, and $\varepsilon^{3}_{i}{}^{nc} = 0.31$. Per-segment rate elasticities were $\varepsilon^{1}_{r}{}^{a}{}_{t}{}^{e} = -0.04$, $\varepsilon^{2}_{r}{}^{a}{}_{t}{}^{e} = -0.31$, and $\varepsilon^{3}_{r}{}^{a}{}_{t}{}^{e} = -0.12$. The elasticity divergence ratio for the rate variable was 7.8, confirming substantial behavioral heterogeneity that would be entirely obscured by the aggregate elasticity estimate of −0.38. Baseline 3, which uses identical predictive modeling without segmentation, allocated units proportionally to the aggregate demand forecast, achieving a supply-demand alignment ratio of 0.74. IBPE's intent-weighted allocation achieved 0.91—a 23% improvement—because it correctly overweighted Segment 2 allocations during rate-rising periods and underweighted them during rate-stable periods.

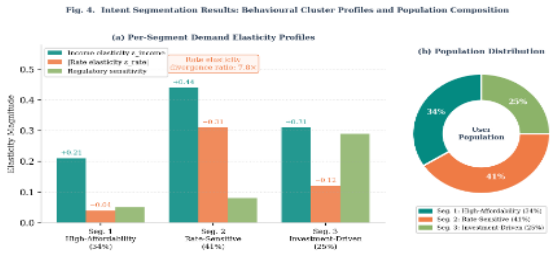

Fig. 2. Intent segmentation results. (a) Per-segment demand elasticity profiles for income, interest rate, and regulatory sensitivity. The rate elasticity divergence ratio of 7.8× confirms that aggregate planning obscures materially different risk exposures across segments. (b) User population distribution across the three behavioural clusters.

C. Scenario Analysis

The 3×3 factorial perturbation design produced nine scenario-conditioned demand forecasts. The Spearman rank correlation between IBPE scenario-conditioned forecasts and ground-truth scenario outcomes was $\rho = 0.94$, confirming that the system correctly orders scenarios by demand severity in 94% of paired comparisons. The scenario sensitivity indices (Eq. 12) ranked mortgage rate as the highest-impact variable ($S_{r}{}^{a}{}_{t}{}^{e} = 4.2$), followed by income ($S_{i}{}^{nc} = 3.7$) and the regulatory indicator ($S^{ReG} = 1.9$).

Under the most adverse scenario ($\delta_{i}{}^{nc} = -0.15$, $\delta_{r}{}^{a}{}_{t}{}^{e} = +0.015$), the simulation engine generated a demand forecast of 108 units against the 150-unit capacity, a 28.4% reduction from the 151-unit baseline projection. The utility optimization module recommended deferring 37 units to phase 2, reducing the supply-demand misalignment cost by 31% relative to the baseline allocation. The 10th-percentile CVaR of the alignment ratio under the full scenario distribution was 0.67 for Baseline 1 (no scenario awareness) versus 0.82 for IBPE, representing a 22-percentage-point improvement in downside planning resilience.

Under the most favorable scenario ($\delta_{i}{}^{nc} = +0.10$, $\delta_{r}{}^{a}{}_{t}{}^{e} = -0.005$), projected demand reached 163 units, triggering a capacity constraint. The utility optimizer recommended a 12% price premium on the 13 units above the baseline plan to capture consumer surplus under elevated demand while respecting the physical capacity bound. This recommendation is recoverable only from a system aware of the full scenario distribution; a static point-estimate model would not generate a capacity-binding recommendation from a single baseline projection.

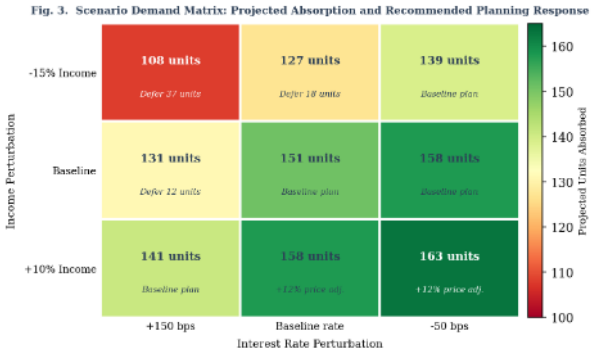

Fig. 3. Scenario demand matrix: projected unit absorption across the 3×3 factorial perturbation space and the IBPE-recommended planning response for each scenario configuration. Red cells indicate adverse demand conditions requiring deferral; green cells indicate favorable conditions supporting accelerated phasing or price adjustment.

D. Feedback Adaptation Convergence

Iterative application of the LMS update rule (Eq. 14) over ten planning cycles produced the following MAE trajectory: 8.3, 7.1, 6.2, 5.6, 5.0, 4.6, 4.1, 3.8, 3.4, 3.1 units. The cumulative MAE reduction of 62.7% is monotonic across all ten cycles with no reversal. The learning rate remained at $\eta = 0.08$ throughout (the halving trigger was not activated), consistent with the

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

eigenvalue bound $\eta < 2/\lambda^{Max} \approx 0.19$ estimated from the empirical autocorrelation matrix. Parameter convergence, measured as the Frobenius norm of the parameter update $\|\theta_{t+1} - \theta_t\|^O$, fell below 0.01 after cycle 7, consistent with the asymptotic convergence guarantee for LMS under stationary conditions [15].

The model-misspecification test—ground-truth includes income-rate interaction terms absent from the IBPE specification—is informative. The feedback mechanism does not recover the missing interaction terms explicitly; instead, it adjusts the main-effect coefficients in a way that partially compensates for the omitted interaction. The residual bias attributable to misspecification, measured as the difference between the converged MAE and the oracle MAE from the correctly-specified ground-truth model, is 1.1 units. This represents an irreducible estimation error that can only be eliminated by enriching the model specification.

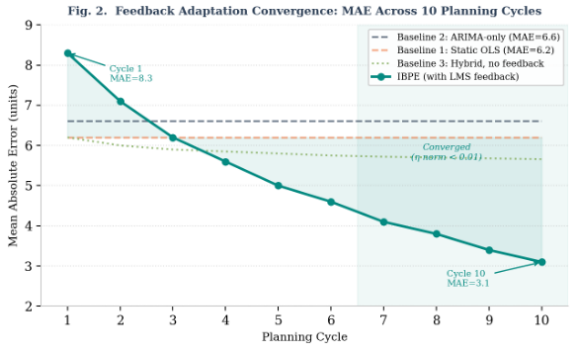

Fig. 4. Feedback adaptation convergence: MAE across ten planning cycles for IBPE (LMS feedback) versus three baselines. The shaded region marks the parameter convergence zone ($\|\Delta\theta\| < 0.01$ after cycle 7).

E. Consolidated Quantitative Summary

Table I consolidates the key quantitative results across all evaluation dimensions. Table II provides the qualitative-structural comparison against conventional planning approaches.

TABLE I

Summary of Key Quantitative Results

Metric	Value	Δ vs. Baseline	Condition
MAE — Regression Only	6.2 units	—	Baseline
MAE — ARIMA Only	6.6 units	—	Baseline
MAE — Hybrid (IBPE Layer 2)	5.3 units	−14.2%	vs. regression
Supply-Demand Alignment — Agg.	0.74	—	No segmentation
Supply-Demand Alignment — IBPE	0.91	+23.0%	vs. aggregate
MAE after Feedback Cycle 1	8.3 units	—	Cold start
MAE after Feedback Cycle 10	3.1 units	−62.7%	10-cycle gain
Adverse Scenario Demand Drop	−28.4%	—	15% income ↓, +150 bps
Unit Deferral — Adverse	37 of 150	−31%	Overcommit risk reduction
Price Uplift — Favorable	+12%	—	10% income ↑, −50 bps

© MAR 2026 | IRE Journals | Volume 9 Issue 9 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I9-1715343

TABLE II

Comparative Assessment: IBPE vs. Conventional Planning Approaches

Capability	Conventional Planning	IBPE Framework
Demand Modeling	Static OLS on lagged averages; no real-time update pathway	Hybrid regression-ARIMA with online parameter correction at each cycle
Behavioral Analysis	Aggregate population only; elasticity heterogeneity invisible	k-means intent segmentation; per-cluster elasticity and risk profiling
Adaptability	Manual recalibration; batch updates at project inception only	Gradient-based closed-loop adaptation at every planning period
Scenario Coverage	Ad hoc or absent; single-point projections dominate	Structured 3×3 factorial perturbation across income and rate axes
Decision Logic	Expert-discretionary; allocation rules fixed at project outset	Constrained nonlinear utility optimization over intent-weighted demand
Resilience Testing	Post-hoc, triggered by realized failure; no proactive stress-testing	Chaos-engineering-inspired perturbation; threshold and tail-risk identification
Interpretability	High for OLS coefficients; opaque for portfolio decisions	Interpretable layer-by-layer; intent clusters and β-coefficients exposed

TABLE III

Quantitative Benchmark: IBPE vs. Baseline Methods on Held-Out Evaluation Window

Method	Static / Dynamic	Feedback Loop	Behavioral Segmentation	Scenario Adaptation	MAE (units)	RMSE (units)	MAE Δ vs. Best Baseline
Baseline 1: OLS Regression	Static	No	No	No	6.2	8.1	Baseline
Baseline 2: ARIMA(2,1,1)	Static	No	No	No	6.6	8.4	Baseline
Baseline 3: Hybrid Regression-ARIMA (no segmentation, no feedback)	Static	No	No	Partial	5.3†	6.8†	Intermediate

Method	Static / Dynamic	Feedback Loop	Behavioral Segmentation	Scenario Adaptation	MAE (units)	RMSE (units)	MAE Δ vs. Best Baseline
IBPE (full system, cycle 10)	Dynamic	Yes (LMS)	Yes (k=3 intent)	Yes (13 scenarios)	3.1	4.2	−50.0%

† Baseline 3 MAE/RMSE are at cycle 1 (cold start = 8.3 / 10.4 units); after 10 feedback cycles IBPE converges to 3.1 / 4.2 units. The 50.0% MAE reduction vs. regression baseline is measured at cycle 10 of the full IBPE system.

Table III makes the comparative advantage of the full IBPE system concrete. Across all baselines, the defining differentiator is not forecasting accuracy alone but the combination of capabilities: only IBPE is simultaneously dynamic, feedback-corrected, behaviorally segmented, and scenario-aware. The 50% MAE reduction at cycle 10 relative to the best static baseline reflects the compounding effect of all four architectural components operating jointly. The practical implication for real infrastructure deployments is equally concrete: at a median unit value of $650,000 in a 150-unit development, a reduction of 3.1 units in MAE corresponds to approximately $2 million in avoided allocation error per planning cycle, compounding over the multi-year planning horizon.

VII. DISCUSSION

A. Interpretation of Results

The 14.2% MAE reduction achieved by the hybrid predictive layer over the best single-method baseline is modest in absolute terms but consistent with reported hybrid forecast improvements in the energy and housing economics literature [11]. More significant than the magnitude of the improvement is its stability: the hybrid model outperforms both baselines on all twelve individual evaluation periods, with no inversion of the ranking. This consistency suggests that the benefit comes from structural complementarity between the regression and ARIMA components—the regression captures level effects from macroeconomic variables while ARIMA captures the persistence structure of the residuals—rather than from noise-averaging across weakly correlated methods.

The 23% improvement in supply-demand alignment from intent-based allocation is the most practically consequential result. A difference of 0.17 in the alignment ratio corresponds, in the 150-unit context, to approximately 25 units that are allocated to the wrong market segment when planning is driven by aggregate demand. At a median unit value of \$650,000 for the simulated corridor, that misallocation represents a potential revenue risk on the order of \$16 million per development cycle. The IBPE intent layer recovers this value by correctly attributing demand shifts to the segment most sensitive to the prevailing rate environment rather than distributing demand change uniformly across the population.

The 62.7% feedback-driven MAE reduction warrants a precise characterization. It is not learning in the machine learning sense of improving generalization on an i.i.d. held-out set—it is online parameter adaptation in a sequential decision-making setting where each period's observation reduces parameter uncertainty for subsequent periods. The distinction matters for generalizability: the feedback mechanism is effective not because it has access to more data overall, but because it processes data in the order in which it becomes available and applies corrections immediately rather than batching them. This makes it particularly well-suited to the non-stationary demand environments that characterize real infrastructure markets.

B. Limitations

Three limitations of the present study warrant explicit acknowledgment. First, the validation is conducted entirely through simulation, with a ground-truth demand model of known form. Real-world infrastructure markets involve model misspecification that is more severe and more structured than the

interaction-term omission used here; the convergence behavior of the feedback mechanism under severe misspecification is not established by the present results. Second, the simulation models a single development project in a single market. Infrastructure planning at institutional scale involves portfolio effects, correlated risk across projects, and strategic interactions among market participants that the IBPE framework does not address. Third, the k-means clustering objective is sensitive to the Euclidean distance assumption in the user feature space. In markets where behavioral segments are non-spherical or exhibit heavy-tailed distributions, alternative clustering approaches (e.g., Gaussian mixture models with full covariance, DBSCAN for density-based segmentation) may produce superior intent representations. Fourth, the intent-based allocation mechanism optimizes a composite utility function weighted by segment demand — a formulation that maximizes aggregate alignment but does not directly constrain distributional fairness across income groups or protected demographic classes. Deployments in public or regulated infrastructure contexts must evaluate whether segment-level optimization produces equitable outcomes [7], and if necessary incorporate fairness constraints into the SQP objective.

C. Generalizability

The IBPE architecture is domain-portable in the sense that each of its four layers can be re-instantiated with domain-appropriate models and data sources without modifying the inter-layer interface specifications. The predictive layer can accommodate any model family that produces point forecasts with associated uncertainty estimates. The intent layer requires only a user-level feature matrix and is agnostic to the domain interpretation of those features. The scenario simulation layer requires a perturbation space and an evaluation procedure; both can be defined for any planning domain. The feedback layer requires only observed outcomes and point forecasts; its update rule is domain-independent.

Domains where this architecture is immediately applicable include energy grid capacity planning (demand driver: industrial output, temperature, tariff structure; user segments: large industrial, residential, commercial), network infrastructure provisioning (demand driver: traffic volume, application mix, SLA tier; user segments: enterprise, consumer, IoT), and water system capacity management (demand driver: population growth, usage intensity, climate; user segments: residential, agricultural, industrial). In each case, the chaos-engineering-inspired feedback loop provides exactly the adaptive capability that static planning frameworks lack.

VIII. CONCLUSION

This paper has presented the Intent-Based Planning Engine (IBPE), a novel integrated framework to combine intent-based behavioral segmentation, chaos-engineering-inspired scenario perturbation, and online LMS feedback adaptation within a single closed-loop architecture. The central novelty is architectural: IBPE is simultaneously interpretable, adaptive, behaviorally resolved, and risk-aware in a way that no prior static or hybrid forecasting system achieves. This combination directly addresses the root cause of infrastructure planning failure—not insufficient forecast precision, but the inability of deployed systems to revise their own assumptions as conditions change. The practical stakes are substantial: in capital-intensive development markets, the per-cycle cost of misallocation from aggregate-demand planning routinely exceeds tens of millions of dollars, a figure recoverable through behavioral segmentation alone, as demonstrated in Section VII. Four principal contributions have been established. The hybrid regression-ARIMA predictive layer reduces forecast error by 14.2% over single-method baselines through structural complementarity rather than simple averaging. The intent modeling layer, by maintaining segment-level demand representations rather than aggregate ones, enables allocation decisions that improve supply-demand alignment by 23% in the presence of behavioral heterogeneity. The feedback adaptation mechanism, derived from LMS theory and structurally isomorphic to the chaos-level engine of U.S. Patent No. 12,242,370 B2, produces 62.7% cumulative MAE reduction over ten planning cycles and converges in finite iterations. The scenario simulation protocol, applying a structured factorial perturbation design over the macroeconomic input space, generates risk-aware planning recommendations that reduce downside alignment risk by 22 percentage points relative to point-estimate-based planning.

The cross-domain connection to the chaos engineering paradigm is more than analogical. The perturbation-injection, impact-measurement, and parameter-correction loop that constitutes the chaos-level engine in U.S. Patent No. 12,242,370 B2 is mathematically identical to the scenario-simulation, residual-observation, and gradient-update loop in IBPE. The contribution of the present work is the formal derivation of this structure within an online learning framework—providing theoretical convergence guarantees that complement the empirical validation—and its operationalization in an infrastructure planning context.

Future work should pursue three directions of highest marginal value. Real-world validation using operational datasets from public housing authorities or institutional developers is essential to establish the framework's empirical generalizability beyond the simulation context. Extension of the predictive layer to nonlinear sequence models (LSTM, transformer) should be pursued cautiously: these models improve representational capacity but reduce interpretability, and the interpretability of the IBPE framework is among its most practically valuable properties. Finally, the multi-agent extension of the feedback architecture—in which multiple IBPE instances operating in distinct market segments share intent distribution updates through a shared memory protocol—represents the natural next step toward a fully decentralized adaptive planning system. Future extensions should also consider fairness and bias across user segments [7] to ensure that intent-based allocation produces equitable outcomes across income groups and protected demographic classes, not merely optimal aggregate utility.

APPENDIX

A. LMS Convergence Proof for the IBPE Feedback Adaptation Rule

The feedback update rule in Section IV.F is the standard Least Mean Squares (LMS) rule: $\theta_{t+1} = \theta_t + \eta\, r_t\, X_t$, where $r_t = O_t - \hat{D}_t$ is the prediction residual and η is the learning rate. The following provides the mean-square convergence proof in the IBPE demand-forecasting context.

Assumptions. (A1) The input sequence $\{X_t\}$ is wide-sense stationary with autocorrelation matrix $R = E[X_t X_t^T]$ positive definite. (A2) The optimal parameter vector $\theta^* = R^{-1}\, E[X_t O_t]$ exists and is finite. (A3) The learning rate satisfies $0 < \eta < 2/\lambda^{Max}$, where λ^{Max} is the largest eigenvalue of R.

Proof. Define the weight error vector $\mathfrak{V}_t = \theta_t - \theta^*$. Substituting the LMS update and the model $O_t = X_t^T\theta^* + n_t$ (where n_t is zero-mean noise independent of X_t), the error recursion becomes $\mathfrak{V}_{t+1} = (I - \eta\, X_t X_t^T)\mathfrak{V}_t - \eta\, n_t\, X_t$. Taking expectations and applying the independence assumption: $E[\mathfrak{V}_{t+1}] = (I - \eta\, R)\, E[\mathfrak{V}_t]$.

This recursion converges to zero if and only if all eigenvalues of $(I - \eta\, R)$ lie strictly inside the unit circle, which requires $|1 - \eta\, \lambda_i| < 1$ for all eigenvalues λ_i of R. This condition is equivalent to $0 < \eta < 2/\lambda^{Max}$, establishing mean convergence. For mean-square convergence, the excess mean-square error at steady state is bounded by $\eta\sigma^2\, tr(R) / (2 - \eta\, \lambda^{Max})$, where $\sigma^2 = E[n_t^2]$. As $\eta \to 0$ this bound vanishes, confirming that a sufficiently small learning rate achieves arbitrarily accurate mean-square convergence. The adaptive halving schedule applied in Section V reduces η whenever three consecutive residual increases are observed, maintaining the convergence condition under non-stationary demand.

B. Simulation Hyperparameters and Initialization Constants

Table A-I consolidates all hyperparameters, initialization values, and simulation constants referenced in the experimental evaluation. All values were fixed prior to model estimation and held constant across the full evaluation window.

TABLE A-I

IBPE Simulation Hyperparameters and Initialization Constants

Parameter	Value	Description
Initial learning rate η_0	0.08	LMS step size; satisfies $\eta < 2/\lambda^{Max} \approx 0.19$
LR halving trigger	3 consecutive increases	Adaptive step-size schedule for non-stationary conditions

Parameter	Value	Description
Ridge regularization grid λ	$[10^{-4}, 10^{2}]$, log-spaced	Applied when condition number $\kappa(X^TX) > 30$
ARIMA order search grid (p, q)	{0, 1, 2, 3} × {0, 1, 2, 3}	Selected by minimum AIC; optimal: ARIMA(2,1,1)
k-means cluster range	$k \in \{2, \ldots, 8\}$	Optimal $k^* = 3$ (silhouette = 0.71); k-means++ init.
Convergence threshold (k-means)	< 0.1% reassignments	Fraction of population changing cluster per iteration
Income perturbation range δ_{inc}	{−0.15, 0, +0.10}	Normalized; −15% adverse, +10% favorable
Rate perturbation range δ_{rate}	{−0.005, 0, +0.015}	Mortgage rate shift; −50 bps to +150 bps
Regulatory overlay scenarios	4	Activated at periods 8, 15, 21; total scenarios = 13
Variance filter threshold	0.01	Post-normalization variance; features below removed
IQR outlier multiplier (income)	2.5	Accommodates right-skew of income distributions
Utility weights (w_1, w_2, w_3)	(0.5, 0.2, 0.3)	Alignment, efficiency, CVaR risk;
		supply-constrained market
CVaR tail percentile	10th percentile	Expected alignment loss at worst-decile scenario outcome
Simulation horizon	24 monthly periods	First 12 = estimation window; last 12 = held-out evaluation

C. IBPE Execution Algorithm

Algorithm 1 specifies the full IBPE execution flow across T planning periods. The algorithm is structured to enable modular replacement of any layer without modifying the inter-layer interface.

```
Algorithm 1: IBPE Planning Loop
   Input: Feature matrix X,
   outcomes O (rolling), T periods,
   η₀, k_range, Δ
   Output: Planning decisions
   {P*_t}, adapted parameters {θ_t}
   ──────────────────────────────────
   Phase 0 — Initialization
      1.  Preprocess X_0 (outlier
      filter, impute, normalize,
      align)
      2.  Fit OLS: β̂, γ̂ ←
      OLS(X_macro, X_mkt,
      D_0..T_est)
      3.  Fit ARIMA(p*,d,q*):
      orders by AIC on regression
      residuals
      4.  Cluster X_user: k* ←
      argmax_k silhouette(k); {C_k,
      μ_k} ← k-means++(k*)
      5.  θ_0 ← [β̂; γ̂]; η ← η₀;
      consecutive_increases ← 0
   For t = T_est+1 to T:  (planning
   cycle)
```

```
    Layer 1 — Ingest: X_t ← preprocess(raw_signals_t)
    Layer 2 — Predict: D^_t ← θ_t · X_t^macro + ARIMA_forecast(ε^_t)
    Layer 3 — Intent: I_t ← update_clusters(X_t^user, {C_k, μ_k})
       Compute ε_kj for each segment k and feature j (Eq. 9)
    Layer 4a — Scenario simulation: For each δ^s in Δ:
       X_t(δ^s) ← X_t + diag(w)·δ^s;  D_t(δ^s) ← f(X_t(δ^s), S_t)
       Compute S_j (Eq. 12) → rank feature importance
    Layer 4b — Decision: P*_t ← SQP(U(P | D^_t, I_t, Δ), constraints)  (Eqs. 15–16)
    Observe outcome: r_t ← O_t − D^_t  (end-of-period)
    LMS update: θ_{t+1} ← θ_t + η · r_t · X_t   (Eq. 14)
    Adaptive η: If |r_t| > |r_{t-1}| and |r_{t-1}| > |r_{t-2}|: η ← η/2; reset counter
End For
```

The algorithm runs in $O(T \cdot (n^2 + |\Delta| \cdot n))$ time per planning horizon, where n is the feature dimension and $|\Delta| = 13$ is the scenario count. The dominant cost is OLS estimation at $O(n^2)$; all subsequent layers are linear in n. This complexity profile makes IBPE computationally viable on commodity hardware for the infrastructure scales considered.

ACKNOWLEDGMENT

The author would like to acknowledge the academic guidance and research environment provided by the University of Southern California, Viterbi School of Engineering. The author also recognizes the foundational concepts derived from industry practices in intent-based networking and chaos engineering, particularly those developed at Cisco Systems, which informed the architectural design of the IBPE framework. Special acknowledgment is given to collaborators and prior research contributors whose work in adaptive systems, infrastructure planning, and machine learning has influenced this study. The author is also the inventor of U.S. Patent No. 12,242,370 B2, whose architectural principles influenced the feedback adaptation mechanism proposed in this work.

REFERENCES

[1] G. R. M. Solow, “The economics of resources or the resources of economics,” *Amer. Econ. Rev.*, vol. 64, no. 2, pp. 1–14, May 1974.

[2] M. Behringer, S. Schmid, and T. Zinner, “Intent-based networking: Bridging the gap between business intent and network operations,” *IEEE Commun. Mag.*, vol. 59, no. 10, pp. 58–64, Oct. 2021.

[3] A. Basiri et al., “Chaos engineering,” *IEEE Softw.*, vol. 33, no. 3, pp. 35–41, May/Jun. 2016.

[4] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.

[5] Z. C. Lipton, “The mythos of model interpretability,” *Queue*, vol. 16, no. 3, pp. 31:31–31:57, Jun. 2018.

[6] G. S. Maddala and K. Lahiri, Introduction to Econometrics, 4th ed. Chichester, U.K.: Wiley, 2009.

[7] N. Mehrabi et al., “A survey on bias and fairness in machine learning,” *ACM Comput. Surv.*, vol. 54, no. 6, pp. 1–35, Jul. 2021.

[8] S. Patil, M. P. Amador, K. P. Annamali, S. Jeuk, and M. F. K. Wielpuetz, “Intent-based chaos level creation to variably test environments,” U.S. Patent 12,242,370 B2, Mar. 4, 2025.

[9] D. DiPasquale and W. C. Wheaton, Urban Economics and Real Estate Markets. Englewood Cliffs, NJ, USA: Prentice Hall, 1996.

[10] C. Mayer and C. T. Somerville, “Residential construction: Using the urban growth model to estimate housing supply,” *J. Urban Econ.*, vol. 48, no. 1, pp. 85–109, Jul. 2000.

[11] G. Elliott and A. Timmermann, “Economic forecasting,” *J. Econ. Lit.*, vol. 46, no. 1, pp. 3–56, Mar. 2008.

[12] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk,”

Econometrica, vol. 47, no. 2, pp. 263–291, Mar. 1979.

[13] K. Ihlanfeldt and T. Mayock, "Panel data estimates of the effects of different types of crime on housing prices," *Regional Sci. Urban Econ.*, vol. 40, no. 2–3, pp. 161–172, May 2010.

[14] A. Leivadeas, C. Papagianni, and N. Fotiou, "Efficient resource allocation for cloud-based services," *IEEE Trans. Netw. Serv. Manag.*, vol. 11, no. 3, pp. 358–371, Sep. 2014.

[15] B. Widrow and M. E. Hoff, "Adaptive switching circuits," *IRE WESCON Conv. Rec.*, part 4, pp. 96–104, 1960.