Journal of Technology Informatics and Engineering (JTIE)
Vol. 4 No. 3 December 2025
E-ISSN: 2961-9068; P-ISSN: 2961-8215, Pages 493-520
DOI: https://doi.org/10.51903/jtie.v4i3.466

A Constrained, Data-Driven Budgeting Framework Integrating Macro Demand Forecasting and Marketing Response Modeling

Yifei Lu*[1], Hailin Zhou[2], Yitian Zhang[3]
Email: yifei.lu0816@outlook.com
[1]*Computer Science, UCSD, University of California San Diego (UCSD), La Jolla, CA, USA*
[2]*Applied Analytics, Columbia University, NY, USA*
[3]*Accounting, The University of Wisconsin-Madison (UW-Madison), WI, USA*
*Corresponding Author

Abstract

Budgeting and financial planning & analysis (FP&A) increasingly require combining macroeconomic signals, channel-level marketing effectiveness, and hard accounting constraints into a single, auditable decision process. This paper proposes and empirically evaluates an end-to-end framework that (i) forecasts category-level demand from public macro data, (ii) learns diminishing-returns marketing response curves, and (iii) solves a constrained portfolio optimization problem to allocate marketing spend while satisfying SG&A and cash-flow guardrails consistent with real public-company statements. Using quarterly Personal Consumption Expenditures (PCE) components from FRED (durable goods, nondurable goods, and services) as a proxy for market demand, we compare seasonal naïve, SARIMAX, gradient boosting, and a multivariate VAR model in a rolling backtest (2018Q1-2025Q3). In parallel, we estimate marketing response from the Advertising dataset (TV, radio, and newspaper spend) via linear models, gradient boosting, and a Hill-function saturation model. We then calibrate financial constraints-gross margin, SG&A ratio, and operating cash-flow coverage-directly from Apple Inc.'s FY2025 Form 10-K filed with the SEC, and integrate all components into a Monte Carlo-evaluated budgeting optimizer. Results show that multivariate models improve total-demand accuracy (≈2.85% MAPE) and that nonlinear response curves indicate strong diminishing returns and negligible incremental value for newspaper spend. The constrained optimizer produces stable allocations that trade off expected operating profit and downside risk, and it highlights a practical insight: budgets that exactly meet a ratio-based cap under point forecasts may violate constraints under realistic demand uncertainty. The proposed workflow is fully reproducible from public data sources and provides a template for transparent, constraint-aware budgeting.

Keywords: *Budget Optimization, FP&A, Demand Forecasting, Marketing Mix.*

I. INTRODUCTION

Budgeting is one of the most repeated and most consequential management processes inside organizations. A budget is simultaneously a forecast (what is likely to happen), a commitment device (what the organization will fund), and a control system (how performance will be evaluated). In modern FP&A, the process is further complicated by rapid changes in consumer behavior, channel fragmentation, and the expectation that financial plans can be refreshed quickly as conditions change. Organizations, therefore, need tools that can be updated frequently, explain their logic clearly, and connect directly to accounting outcomes (Jamabudin Ghofur & Riyanto, 2025). While FP&A budgeting is used here as a case study, the paper's contribution is methodological: an auditable end-to-end pipeline that integrates demand forecasting, marketing

Received on August, 2025; Revised on November, 2025; Accepted on November, 2025; Published on December, 2025

response modeling, and constraint-aware optimization under uncertainty (Adowusu-Mensah et al., 2024).

A common practical tension is that budget inputs come from heterogeneous sources and at different temporal resolutions. Macroeconomic indicators and market demand signals may be published monthly or quarterly; internal sales data may be published daily; and some key financial constraints (e.g., SG&A-to-revenue targets) are reviewed on a quarterly and annual basis. A coherent pipeline must combine these signals without losing interpretability. For example, a business may decide to cut discretionary spending when demand softens. However, it must also ensure that the cut does not damage growth and still satisfy fixed commitments such as R&D and payroll.

Marketing spend is especially challenging to budget because its benefits are uncertain, nonlinear, and often delayed. Empirically, most channels exhibit diminishing marginal returns: the first dollars spent are usually the most effective, while later dollars mostly buy redundant reach. Additionally, channels can substitute for each other (e.g., radio and TV) or work synergistically. This means that linear response models, while easy to communicate, can produce unrealistic implied returns at high budgets.

Accounting constraints further constrain feasible budget actions. Even if a marketing investment has a positive expected return, it might be unacceptable if it causes SG&A ratios to exceed guidance, if it reduces operating cash flow below a threshold required for debt covenants, or if it pushes operating margin below a target. These constraints are not hypothetical: public-company filings explicitly report the components of profitability and cash flow, and internal governance often sets ratio targets that mirror those external statements (Apple Inc., 2025).

The goal of this paper is to propose a lightweight yet end-to-end framework that can be instantiated using fully public datasets and to demonstrate how forecasting, response modeling, and constraint-aware optimization can be integrated into a single budgeting workflow. Our emphasis is not on producing the best possible forecast in isolation, but rather on producing a decision recommendation that is defensible under uncertainty and aligned with accounting realities.

A useful mental model is to treat budgeting as a pipeline that turns information into commitments. The input information includes macro and market signals, internal run-rate data, strategic priorities, and constraints. The commitments are resource allocations and targets. The key challenge is that the same pipeline must be both analytically sound and organizationally legible: the plan must be explainable to executives and must be implementable by operating teams.

We use this paper's public-data instantiation as a pedagogical reference implementation. It is intentionally conservative: it uses simple but robust model classes, emphasizes backtesting and uncertainty estimation, and calibrates constraints to audited public statements. The trade-off is that some domain-specific nuance (e.g., product-level mix, marketing carryover) is abstracted away. The benefit is a clear template for linking data, models, and constraints in an auditable way.

Figure 1 summarizes the proposed end-to-end architecture. The demand module produces category-level forecasts and an empirical uncertainty distribution. The marketing module estimates a response function with diminishing returns. The constraint module converts audited financial statement ratios into feasibility constraints. Finally, the optimizer searches over channel allocations to maximize an explicit profit objective while respecting constraints and controlling risk.

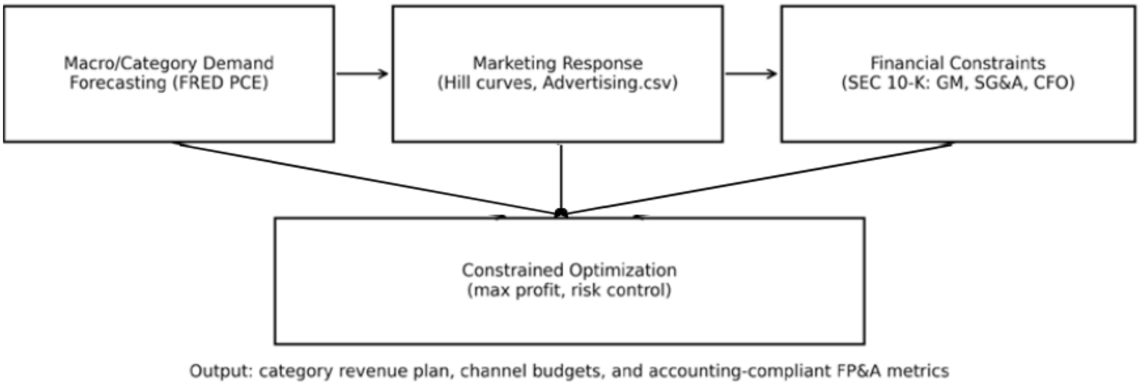

Figure 1. Integrated Forecasting-Response-Constraint Budgeting Framework

We structure our contributions in terms that map to practical FP&A deliverables:

1. Forecasting deliverable: a rolling backtest comparison that yields an explicit accuracy-complexity trade-off for macro demand proxies.
2. Marketing deliverable: marginal ROI curves by channel that identify saturation points and low-value channels.
3. Finance deliverable: a budget recommendation expressed in the same ratios used in accounting reviews (gross margin, SG&A ratio, operating cash-flow coverage).
4. Governance deliverable: Monte Carlo constraint satisfaction rates that quantify how often a plan could violate ratio caps under demand uncertainty.

We also summarize the methodological components as:

1. Forecasting systems: rolling-origin macro demand forecasting and empirical uncertainty estimation.
2. Data-driven modeling/analytics: channel-level response modeling with diminishing returns and marginal ROI diagnostics.

3. Optimization under uncertainty: constrained budget allocation with Monte Carlo evaluation and ratio-based feasibility under stochastic revenue.
4. Decision support and auditability: a reproducible, transparent workflow linking public data, interpretable models, and accounting-calibrated constraints.

II. LITERATURE REVIEW

Forecasting for planning and control. Forecasting research emphasizes that model evaluation must reflect how forecasts are used. In planning settings, forecasts are typically produced repeatedly (e.g., monthly or quarterly), and models are re-estimated as new data arrive. Rolling-origin backtests therefore provide more realistic performance estimates than a single train/test split. Classical ARIMA/SARIMA methods remain popular because they are relatively transparent and can perform well on many economic series (Box et al., 2015). Model selection criteria such as AIC and BIC provide principled ways to trade fit and complexity (Akaike, 1974; Schwarz, 1978).

Structural breaks are especially relevant for budgeting. Events such as recessions, pandemics, and policy shocks can introduce level shifts and volatility spikes. In such environments, multivariate models may help by borrowing strength across correlated categories. VAR models are a standard tool for capturing such cross-series dynamics and are widely used in macroeconometrics (Lütkepohl, 2005). However, VAR models still assume linear dynamics and may struggle with nonlinear regime changes.

Machine learning for time series. Tree-based machine learning methods, such as gradient boosting, often perform well on tabular predictive tasks because they can capture nonlinearities and interactions. Friedman (2001) formalized gradient boosting as stagewise additive modeling. In forecasting applications, a common strategy is to convert a time series into a supervised learning problem using lagged features and rolling statistics. This approach can approximate nonlinear autoregressions and can be combined with cross-validation or walk-forward validation. It does, however, sacrifice some of the interpretability and probabilistic structure of classical time-series models.

Marketing mix and saturation modeling. Marketing response modeling has long recognized that marginal returns diminish with spend. Empirical work and applied MMM frequently use saturation functions, such as Hill curves and S-shaped logistic curves, to model concave responses and avoid unrealistic linear extrapolation (Hanssens et al., 2001). In addition, many practical MMM systems include adstock (carryover) effects, reflecting that some advertising impacts future periods. Our demonstration omits adstock because the Advertising dataset is cross-sectional, but the optimization framework readily accommodates dynamic response functions when time series data are available.

Optimization, risk, and constraints in FP&A. From a decision-analytic perspective, budget allocation resembles portfolio selection: spending in each channel produces an uncertain return with diminishing marginal gains. Mean-variance optimization (Markowitz, 1952) provides a simple template for balancing expected profit and risk. Robust optimization and stochastic programming provide stronger guarantees under uncertainty, including chance constraints that bound the probability of violating requirements (Bertsimas & Sim, 2004; Shapiro, Dentcheva, & Ruszczyński, 2014). In corporate finance, constraints are often ratio-based (e.g., SG&A/revenue), which introduces an additional challenge: the denominator is uncertain and can move against the plan.

A final theme in the literature is the gap between predictive accuracy and decision quality. A model that slightly reduces RMSE may not change the optimal decision if constraints bind or response curves saturate early. Conversely, a model that is slightly less accurate in RMSE but provides better tail risk estimates could improve constraint satisfaction. This motivates our emphasis on Monte Carlo evaluation and on reporting satisfaction rates, not only point forecasts.

III. RESEARCH METHOD

We instantiate the framework with three public data sources chosen to be easy to reproduce without proprietary inputs. The datasets were selected to cover the full pipeline: (i) a macro demand proxy for forecasting, (ii) marketing response data for response estimation, and (iii) a real, audited financial statement source for constraint calibration. Table 1 lists the data sources and variables.

Table 1. Dataset Inventory and Coverage

Dataset	Source	Variables Used	Frequency	Time Range	N
FRED PCE Components (PCDG, PCND, PCESV)	Federal Reserve Bank of St. Louis (FRED) / BEA	PCE Durable Goods, Nondurable Goods, Services (SAAR, $bn)	Quarterly	2010-01-01 to 2025-07-01	63
Advertising.csv (ISLR)	ISLR dataset (raw GitHub mirror)	TV, Radio, Newspaper spend; Sales outcome	Cross-sectional	n/a	200
Apple Inc. Form 10-K (FY2025)	SEC EDGAR filing (aapl-20250927.htm)	Net sales, cost of sales, gross margin, SG&A, R&D, operating cash flow	Annual (fiscal year)	FY ended 2025-09-27	1

A. Dataset

Macro Demand Proxy: PCE Components from FRED

The macro module uses three components of Personal Consumption Expenditures: durable goods (PCDG), nondurable goods (PCND), and services (PCESV). These series are available in the Federal Reserve Bank of St. Louis's FRED database and originate from the Bureau of Economic

Analysis' National Income and Product Accounts (Federal Reserve Bank of St. Louis, 2025a, 2025b, 2025c). Values are seasonally adjusted annual rates (SAAR) in billions of dollars.

SAAR reporting implies a simple conversion when connecting to quarterly firm revenue: a quarterly SAAR value can be divided by four to obtain the corresponding quarterly level (in the same units). We use this conversion in the calibration step that maps macro demand to firm revenue. This is an approximation; in applied settings, practitioners should ensure that the firm's revenue timing and the macro series definitions are aligned (e.g., fiscal vs calendar quarters).

We restrict to 2010Q1-2025Q3 to focus on contemporary dynamics and include a major structural break (COVID-19). The categories behave differently during 2020: durables spike (reflecting substitution toward goods), services decline sharply (reflecting mobility restrictions), and nondurables remain relatively stable. These differences are useful for testing whether multivariate forecasting adds value. Table 2 provides descriptive statistics for the demand proxy series used in forecasting.

Table 2. Summary Statistics for Quarterly PCE Component Series (SAAR, $bn)

Series	Start	End	N	Mean	Std	Min	Max
PCDG	2010-01-01	2025-07-01	63	1557.575	405.467	1021.085	2277.613
PCND	2010-01-01	2025-07-01	63	3054.254	616.332	2245.069	4269.068
PCESV	2010-01-01	2025-07-01	63	9632.057	2197.879	6835.669	14568.178

Marketing Response Data: Advertising.csv

The marketing module uses the Advertising dataset from the ISLR (James et al., 2021) package. The dataset includes 200 observations of advertising spend across three channels (TV, radio, newspaper) and their associated sales outcomes. Units are not explicitly defined in the textbook; we treat them as normalized spend and sales units. The dataset is cross-sectional rather than time-indexed, which simplifies response estimation but prevents modeling carryover (adstock) effects. Nevertheless, it remains a useful benchmark for understanding channel importance, multicollinearity, and saturation.

Because the dataset is small, we emphasize simple models with clear diagnostics. We report cross-validated accuracy and then prioritize interpretability for optimization. In real deployments, the same approach can be applied to internal channel data with time indices, allowing adstock, seasonality, and targeting variables.

Constraint Data: SEC Form 10-K

To ground constraints in audited accounting statements, we use Apple Inc.'s FY2025 Form 10-K filed with the SEC (Apple Inc., 2025). We extract the statement of operations and the statement of cash flows. The extracted values are used to compute: (i) gross margin rate, (ii) SG&A ratio, (iii) R&D ratio, and (iv) operating cash flow coverage. These ratios serve as empirical anchors for our budget constraints and for translating incremental revenue into operating profit.

Although our budgeting problem is generic and can be calibrated to any firm, using a real public filing ensures that the constraints reflect realistic magnitudes. In practice, firms would calibrate constraints to their internal statements and might also use additional constraints, such as minimum marketing presence requirements, contractual commitments, or channel capacity limits.

B. Data Reliability, Revisions, and Alignment Caveats

Public macroeconomic data are often revised after initial publication. National accounts series, including PCE components, can be revised as source data are updated or methodologies change. A real-time forecasting system must therefore decide whether to use latest-vintage data (the best estimate of historical truth) or real-time vintages (what would have been known at the time). Our demonstration uses the latest vintage values available on the FRED pages we queried. This is appropriate for illustrating model comparisons, but it can overstate performance compared with a strict real-time evaluation.

Alignment across data sources also matters

The PCE series is a calendar-quarter measure, while corporate financials are reported on fiscal calendars that may not align exactly with calendar quarters. We abstract away these nuances by treating quarter boundaries as compatible and by using ratio-based mapping. In applied FP&A work, alignment should be handled explicitly-for example, by mapping macro quarters into fiscal quarters using weighted averages or by using monthly data when available.

Finally, marketing response estimation is sensitive to measurement error and omitted-variable bias. In practice, channel spend may correlate with other drivers (pricing, distribution, competitor actions). Controlled experiments, geo tests, or causal inference tools can improve identification. Our use of Advertising.csv is a stylized approximation designed to highlight the mechanics of optimization rather than to provide causal estimates.

C. Models and Evaluations

This section describes (i) forecasting models for category demand, (ii) marketing response models with diminishing returns, (iii) the calibration that links macro demand to firm revenue and links marketing response to incremental revenue, and (iv) the constrained optimization problem. We include explicit formulas and algorithmic descriptions to support auditability.

Demand Forecasting Models

We evaluate four demand forecasting approaches following Adiyatma et al. (2025) to produce one-step-ahead forecasts at a quarterly frequency. Let $y_{k,t}$denote the realized demand for the category $k$at time t. For each category, the corresponding forecast is denoted by y hat sub ented by $\hat{y}_{k,t}$. In addition to category-level forecasts, we also assess aggregate demand, defined as (1).

$$P_t = \sum_k y_{k,t} \qquad (1)$$

All forecasting models are re-estimated at each rolling forecast origin to reflect operational forecasting practice. A summary of the model classes and their corresponding specifications is provided in Table 3.

Table 3. Demand Forecasting Models and Specifications

Model	Type	Specification	Notes
Seasonal Naïve	Baseline	$\hat{y}_t = y_t - 4$ (quarterly seasonality)	No fitting; strong baseline for seasonal series
SARIMAX	Statistical	SARIMAX(1,1,1)×(0,1,1,4)	Per-series univariate; refit each step
Gradient Boosting Regressor	ML	GBR with lags 1-4 as features	n_estimators=500, learning_rate=0.05, max_depth=3, subsample=0.8
VAR (log-diff)	Multivariate	VAR on $\Delta\log(y)$ with lag selected by AIC (≤4)	Jointly forecasts 3 categories; converted back to level forecasts

Evaluation Metrics

Forecast accuracy is evaluated using the root mean squared error (RMSE) and the mean absolute percentage error (MAPE). For a test sample consisting of observations $y_t$and corresponding forecasts $\hat{y}_t$, RMSE is defined as (2).

$$RMSE = \sqrt{\frac{1}{n}\sum_t (y_t - \hat{y}_t)^2} \qquad (2)$$

MAPE is computed as (3)

$$MAPE = \sqrt{\frac{1}{n}\sum_t \left|\frac{y_t - \hat{y}_t}{y_t}\right|} \qquad (3)$$

MAPE is scale-free and directly interpretable as a percentage error. While it can place disproportionate weight on observations with small denominators, this concern is mitigated in our setting because demand levels are large and strictly positive. RMSE, by contrast, penalizes larger forecast errors more heavily, which is informative when performance during stress periods is of particular interest.

Rolling-origin evaluation

We adopt a rolling-origin design because it mirrors how forecasts are used in budgeting. In each quarter, the model is trained on data up to the prior quarter and then used to predict the next. This avoids look-ahead bias and captures how model performance evolves as more data becomes available. The design also supports uncertainty estimation by collecting an empirical distribution of forecast errors.

Algorithm 1. Rolling One-Step-Ahead Backtesting Procedure

```
Input: Time series $\{y_t\}_{t=1}^{T}$; test window $t = T_0, \dots, T$
For $t = T_0$ to $T$:
    1. Estimate model $M$ using observations $\{y_1, \dots, y_{t-1}\}$.
    2. Generate a one-step-ahead forecast $\hat{y}_t = M(1)$.
    3. Compute the forecast error $\varepsilon_t = y_t - \hat{y}_t$.
Output: Forecasts $\{\hat{y}_t\}$ and errors $\{\varepsilon_t\}$ over the test window.
```

Seasonal Naïve Baseline

The seasonal naïve benchmark is defined as (4).

$$\hat{y}_t = y_t - 4 \qquad (4)$$

which captures quarterly seasonality by projecting demand from the corresponding quarter in the previous year. Although highly parsimonious, this baseline is often competitive for seasonal economic time series. Including strong and transparent benchmarks is considered best practice in forecasting evaluation, as it provides a meaningful reference point and helps avoid overstating gains from more complex models.

SARIMAX

We estimate a SARIMAX(1,1,1)×(0,1,1,4) model. The (1,1,1) component captures short-term autoregressive and moving-average structure, while the seasonal (0,1,1,4) component captures quarterly seasonality and seasonal shocks. Differencing (d=1, D=1) helps address nonstationary level and seasonal patterns. We fix the order to keep the comparison transparent; in practice, one could select orders based on AIC/BIC or use automated procedures, but such choices can complicate auditability.

Gradient boosting regression

We construct a supervised learning dataset using lagged demand values $[y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4}]$as input features. A gradient boosting regressor is then estimated with 500 boosting iterations, a learning rate of 0.05, a maximum tree depth of 3, and a subsampling ratio of 0.8. Given the limited training sample (32 quarterly observations prior to the test period), model complexity is deliberately constrained by restricting tree depth and introducing subsampling as a regularization mechanism. This specification allows the model to capture nonlinear autoregressive patterns while maintaining robustness to noise and potential outliers.

VAR on log differences

We fit a VAR model to $\Delta\log(y)$ for the three-category vector. Log-differencing approximates growth rates and stabilizes variance. VAR then models each category's growth rate as a linear function of past growth rates of all categories. We select the lag order via AIC, with a maximum of 4 lags. This provides a balance between flexibility and overfitting in a small sample.

Forecast aggregation

To obtain forecasts of aggregate demand P_tWe compute the sum of the category-level forecasts. This aggregation-by-summation approach is simple and widely used in practice. An alternative strategy would be to model aggregate demand directly alongside category-level series and subsequently reconcile the resulting forecasts using hierarchical forecasting methods. We leave such extensions for future work.

Marketing Response Modeling

The marketing module estimates how marketing spend translates into outcomes (Santoso & Priyadi, 2024). In many organizations, the core decision question is marginal: if we move $1M from channel A to channel B, what is the expected effect on revenue or profit? Therefore, beyond predictive accuracy, we need models that provide stable marginal effects and that can be interrogated for saturation.

We fit OLS, ridge, lasso, gradient boosting, and a Hill saturation model. OLS provides a baseline and can be interpreted as a constant marginal effect. Ridge and lasso address multicollinearity by shrinking coefficients (James et al., 2021). Gradient boosting provides a flexible nonlinear benchmark (Friedman, 2001). The Hill model provides an interpretable saturation curve consistent with diminishing returns (Hanssens et al., 2001).

Cross-validation

Because the dataset is cross-sectional, we use standard K-fold cross-validation. This differs from time-series evaluation; however, it is appropriate here because there is no temporal ordering. RMSE and R^2 provide complementary views: RMSE measures absolute predictive error, while R^2 measures the proportion of variance explained.

Hill saturation model

We parameterize the marketing response function as (5)

$$Sales(s) = \beta_0 + \sum_i \beta_i h(s_i; \alpha_i, \gamma_i) \quad (5)$$

where $h(\cdot)$denotes the Hill function. Model parameters are estimated using bounded nonlinear least squares. Parameter bounds are imposed to ensure nonnegative coefficients $\beta_i$and economically plausible values for the saturation parameters $\alpha_i$and γ_i. While additional components such as adstock effects or interaction terms can be incorporated, the baseline saturation specification is often sufficient to identify regions of diminishing returns. A comparative evaluation of alternative marketing response specifications is reported in Table 4, based on five-fold cross-validation using the Advertising dataset.

Table 4. Marketing Response Model Comparison (5-Fold CV)

Model	CV_RMSE	CV_R2
GBR	0.661	0.984
Lasso	1.716	0.891
OLS	1.722	0.891
Ridge	1.722	0.891
Hill	1.449	0.922

Table 5 reports Hill-model parameter estimates fit to the full Advertising dataset.

Table 5. Estimated Hill Saturation Model Parameters

Parameter	Estimate
b0	0.0
b_TV	37.95
b_Radio	31.7322
b_Newspaper	0.0
alpha_TV	0.7283
gamma_TV	499.9991
alpha_Radio	1.4732
gamma_Radio	91.9223
alpha_Newspaper	3.0
gamma_Newspaper	300.0

Figure 4 shows the estimated response curves (varying one channel at a time, with the others held at the mean).

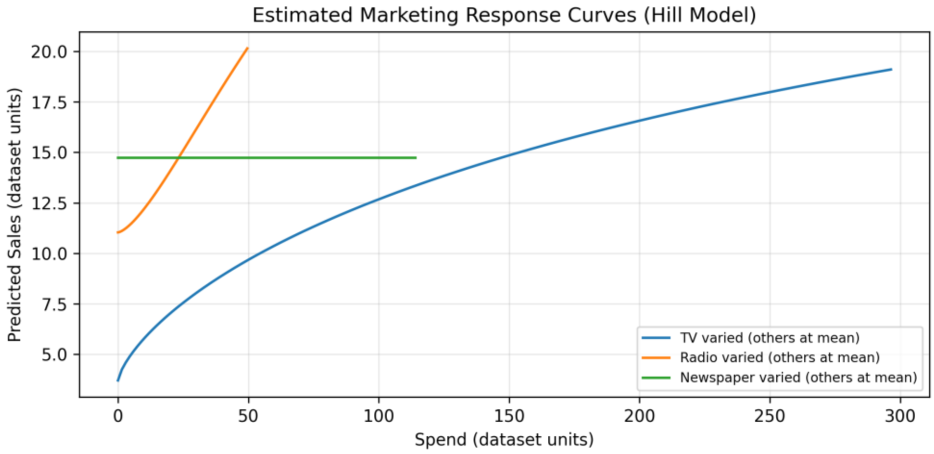

Figure 4. Hill-Model Response Curves (Others Held at Mean)

Figure 5 shows marginal response, emphasizing diminishing returns and weak newspaper contribution.

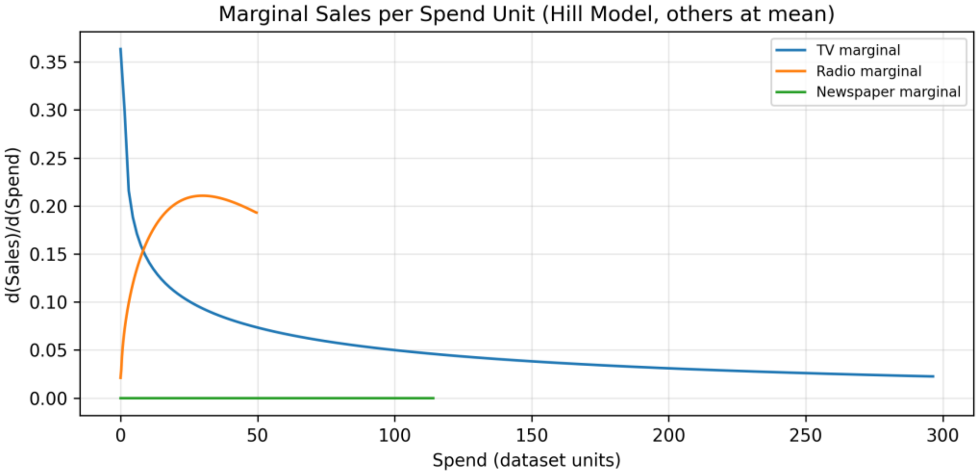

Figure 5. Hill-model Marginal Response (Others Held at Mean)

Calibrating Profitability and Liquidity Constraints

We calibrate profitability and liquidity parameters using data from Apple Inc.'s FY2025 Form 10-K. The statement of operations provides information on net sales, cost of sales, and operating expenses, while the statement of cash flows reports net cash from operating activities. Based on these disclosures, we compute the following ratios: the gross margin rate $g = \text{Gross Margin}/\text{Revenue}$; the selling, general, and administrative (SG&A) ratio $s = \text{SG\&A}/\text{Revenue}$; the research and development ratio $r = \text{R\&D}/\text{Revenue}$; and the operating cash flow ratio $c_{\text{FO}} = \text{CFO}/\text{Revenue}$.

The resulting estimates are $g \approx 46.9\%$, $s \approx 6.63\%$, $r \approx 8.30\%$, and $c_{\text{FO}} \approx 26.8\%$(Apple Inc., 2025). These calibrated values are summarized in Table 6. In practice, financial constraints would be derived from a firm's own accounting data; Apple is used here as a publicly available and auditable reference case.

Table 6. FY2025 Accounting Metrics Used for Constraint Calibration (Apple 10-K)

Metric	FY2025 (USD Millions)	Ratio to Revenue
Net sales (Revenue)	416161.0	100.0
Cost of sales	220960.0	53.09
Gross margin (Gross profit)	195201.0	46.91
Selling, general & administrative (SG&A)	27601.0	6.63
Research & development (R&D)	34550.0	8.3
Net cash provided by operating activities (CFO)	111482.0	26.79

Marketing as SG&A. Because SG&A aggregates multiple costs, we allocate a portion of SG&A headroom to marketing as a modeling choice. We set the marketing budget cap to 1.5% of revenue. This implies that the remaining SG&A (non-marketing) is 5.13% of revenue, matching the 6.63% total SG&A ratio observed in the filing. This choice is conservative and creates a feasibility region consistent with an audited cost structure.

In applied settings, firms might choose a different marketing cap based on strategic priorities or on peer benchmarks. Importantly, the framework can treat the cap as a scenario parameter and can evaluate how optimal budgets and satisfaction rates change as the cap changes. We present such sensitivity analyses in the Results and Discussion section.

Linking Modules and Solving the Budget Allocation Problem

Link 1: Macro demand to firm revenue

We link the macro demand proxy $P_t$to firm-level revenue $R_t$through a constant market share parameter ρ. The parameter $\rho$is calibrated using the most recently observed macro demand and the firm's reported annual revenue, defined as (6).

$$\rho = \frac{Revenue_{annual}}{P_{last}} \quad (6)$$

Quarterly firm revenue is then approximated as (7).

$$\hat{R}_{t+1} = \rho \frac{\hat{P}_{t+1}}{4} \quad (7)$$

This mapping assumes a stable revenue share and a proportional relationship between macro demand and firm sales. In applied settings, ρcould alternatively be estimated by regressing firm revenue on macro demand indicators or by incorporating industry-level market share information.

Link 2: Marketing Response to Incremental Revenue

The marketing response model based on the Advertising dataset produces sales outputs in abstract units. To express these outputs in monetary terms, we calibrate a scaling factor such that a baseline marketing expenditure—set at 1% of quarterly revenue—corresponds to a baseline revenue uplift of 3% of quarterly revenue. This normalization ensures numerical coherence in the optimization while maintaining a conservative calibration. The resulting scaling factor can be interpreted as encoding an assumed average return on investment at the baseline spending level.

Operating profit objective

We define operating profit as a function of the marketing budget vector b as (8)

$$\Pi(b) = (g - o)\big(R + \Delta R(b)\big) - \sum_{o} b_i \qquad (8)$$

where $g$denotes the gross margin rate and $o = r + s_{\text{other}}$represents the non-marketing operating expense rate. Based on the calibration described above, $g - o \approx 0.335$. Under this formulation, incremental revenue generated by marketing activity contributes directly to operating profit.

Constraints

Two constraints are imposed. First, an SG&A budget constraint limits total marketing expenditure to a fixed share of revenue use formula (9),

$$\sum_{i} b_i \leq 0.015\, R \qquad (9)$$

Second, a liquidity constraint restricts marketing spend to the portion of revenue supported by operating cash flows (10),

$$\sum_{i} b_i \leq (c_{FO} - c_{min})R, \qquad (10)$$

where $c_{\min} = 0.25$denotes the minimum cash flow ratio required for operations. Given the calibrated values, $c_{\text{FO}} - c_{\min} \approx 0.0179$, implying that the SG&A constraint is binding in the present setting.

Chance Constraints and Robustness Interpretation

Ratio-based budget constraints are inherently stochastic when revenue is uncertain. Consider a planner who selects a marketing budget Bbased on forecast revenue $\hat{R}$. The SG&A constraint requires as (11)

$$B \leq \kappa R \qquad (11)$$

where κdenotes the allowable spending ratio (e.g., 1.5%) and $R$is realized revenue. When realized revenue falls below its forecast ($R < \hat{R}$), this constraint may be violated even if the planned budget satisfies $B \leq \kappa\hat{R}$.

A natural way to formalize the required safety buffer is through a chance constraint define as (12),

$$\mathbb{P}(B \leq \kappa R) \geq 1 - \delta \qquad (12)$$

where δ represents an acceptable probability of constraint violation (e.g., 5%).

Assuming a multiplicative revenue error model $R = e\,\hat{R}$ with $e > 0$, the chance constraint can be rewritten as (13).

$$B \leq \kappa\,\hat{R}\,q_\delta \qquad (13)$$

where q_δ denotes the δ-quantile of the revenue error factor e. This formulation highlights that setting the budget at the deterministic cap $B = \kappa\hat{R}$ implicitly corresponds to a violation probability of approximately $\delta \approx \mathbb{P}(e < 1)$, which can be substantial when forecast errors are roughly symmetric.

In our empirical revenue error distribution, $\mathbb{P}(e < 1)$ is close to 40%, indicating that spending at the deterministic ratio cap would breach the constraint in roughly 40% of realizations. Consequently, the optimizer selects a budget strictly below $\kappa\hat{R}$, implicitly enforcing a more conservative chance constraint.

We implement robustness through Monte Carlo evaluation rather than an explicit quantile constraint. This approach is transparent: decision-makers can see how many simulated scenarios violate constraints and set λ or explicit satisfaction thresholds accordingly. In more sophisticated implementations, one could embed chance constraints or robust uncertainty sets directly into the optimization (Bertsimas & Sim, 2004; Shapiro et al., 2014).

Search procedure

Because the budget decision has three channels and a simplex constraint, a grid search is practical and transparent. We discretize allocation shares and budget utilization in 5% increments from 20% to 100% of the cap. This yields 3,927 candidate portfolios per scenario. Grid search ensures global coverage and avoids local-optimization failures on nonlinear response surfaces. The computational cost is modest: each portfolio is evaluated with 500 Monte Carlo scenarios using vectorized computation.

Experimental Design

Demand Forecast Backtest

We perform a rolling one-step-ahead backtest from 2018Q1 to 2025Q3. At each origin, models are fit on an expanding window starting at 2010Q1. This yields 31 forecast evaluations per series. We report RMSE and MAPE and also evaluate the aggregate total PCE proxy. The expanding window is appropriate because macro series often benefit from longer histories, but the rolling design still captures structural changes.

Hyperparameter choices

For SARIMAX, we fix the order for transparency. For VAR, we select the lag order via AIC, with a maximum of 4 lags to avoid overfitting. For gradient boosting, we use a moderate number of estimators and shallow trees, reflecting the small sample size. These choices emphasize robustness over maximal fit.

Marketing Model Evaluation

Marketing response models are evaluated with 5-fold cross-validation on the Advertising dataset. We use CV_{RMSE} and CV_{R^2} to compare predictive accuracy. Because our downstream optimizer requires smooth marginal returns, we fit the Hill model with bounded parameters and use it for budget allocation even when more flexible models exist.

Monte Carlo and Scenario Design

We create three macro scenarios by scaling the baseline revenue forecast by -5%, 0%, and +5%, representing recession, baseline, and boom conditions. Within each scenario, we model forecast uncertainty by bootstrapping multiplicative error ratios from the VAR aggregate-demand backtest. We draw $N = 500$ scenarios for Monte Carlo evaluation. For each candidate budget, we compute the mean and standard deviation of operating profit and the SG&A constraint satisfaction rate.

Risk-aversion parameter

We evaluate λ in $\{0, 0.25, 0.5, 1.0, 2.0\}$. While λ has no universal meaning, it can be interpreted as the number of dollars of expected profit the planner is willing to sacrifice to reduce profit volatility by one dollar. In an organizational context, λ can be tuned to align with risk appetite or to meet governance constraints.

IV. RESULT

Table 7 reports category-level forecasting accuracy, and Table 8 reports accuracy for the aggregate demand proxy. The results highlight three patterns. First, the seasonal naïve baseline is difficult to beat for the smoother nondurable series, suggesting that complex models may provide limited incremental value in stable categories. Second, durables and services exhibit larger errors, reflecting greater volatility and structural change. Third, multivariate VAR improves aggregate accuracy, consistent with cross-category substitution during shocks.

Table 7. Demand Forecasting Accuracy by Category (2018Q1-2025Q3). RMSE in \$bn; MAPE in %

Series	RMSE GBR	RMSE SARIMAX	RMSE SeasonalNaive	RMSE VAR	MAPE GBR	MAPE SARIMAX	MAPE SeasonalNaive	MAPE VAR
PCDG	113.25	104.47	168.46	97.24	3.53	3.19	5.87	3.08
PCESV	461.74	966.45	866.79	731.35	3.01	3.77	6.74	3.76
PCND	126.59	83.27	229.49	120.47	2.29	1.65	5.22	2.11

Table 8. Total Demand PROXY accuracy (Sum of Three PCE Components). RMSE in $bn; MAPE in %

RMSE	MAPE
1191.33	6.06
1090.42	2.92
631.3	2.86
860.06	2.85

Figure 2 plots total demand forecasts versus actuals. Errors widen around 2020, as expected, but the forecasts recover as the system adapts to post-shock dynamics. Figure 3 shows the durable series, which features the strongest pandemic-era discontinuity. In this series, multivariate and nonlinear models offer an advantage because they can learn the new regime more quickly than a purely seasonal baseline.

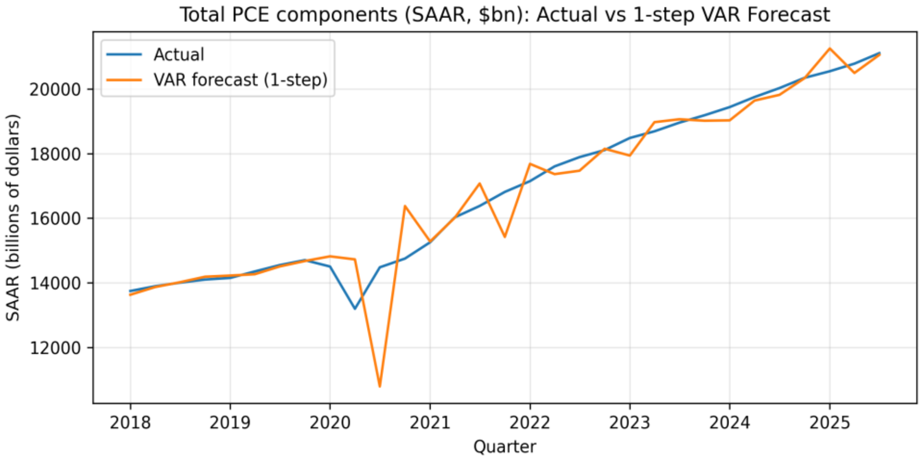

Figure 2. Total Demand Proxy: Actual vs VAR 1-Step Forecasts (Test Window)

Demand Forecasting Results

Category interpretation. Nondurable goods show the lowest MAPE, reflecting smoother consumption patterns. Services show higher errors around the pandemic due to the abrupt collapse and recovery. Durables show both a pandemic shock and post-shock volatility. From a planning perspective, this suggests that budgeting buffers should be larger in categories whose demand is more discretionary or more sensitive to shocks.

Table 4 compares marketing model accuracy. Gradient boosting provides the best predictive fit, but the Hill model provides a structural decomposition into saturating channel contributions. In the fitted Hill parameters (Table 5), the estimated coefficient for newspaper is essentially zero, and marginal response plots confirm that newspaper provides negligible incremental sales over the observed range. This aligns with the common interpretation of the Advertising dataset in ISLR (James et al., 2021).

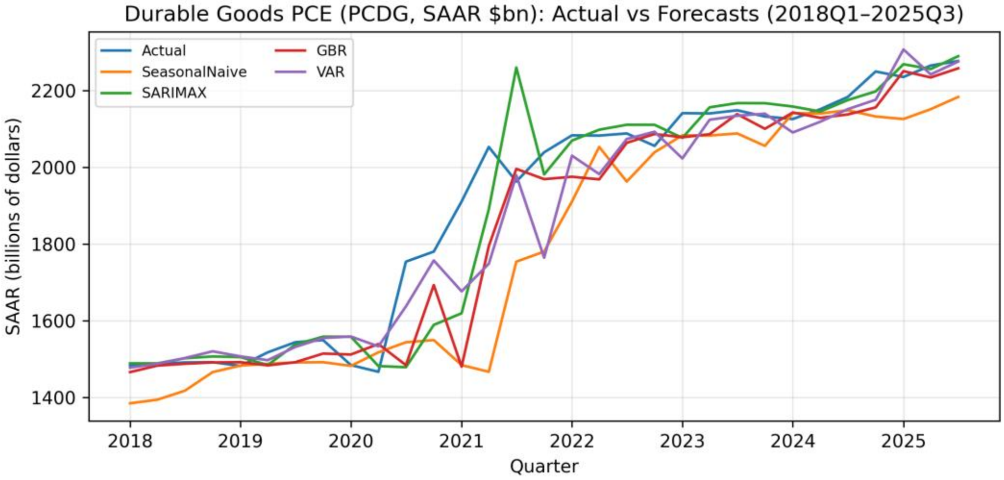

Figure 3. Durable Goods PCE (PCDG): Actual vs Forecasts by Model (Test Window)

Marketing Response Results

From a budgeting perspective, the main use of the response model is not to predict sales perfectly but to estimate marginal ROI curves. Figure 5 shows that radio has the highest marginal returns at low spend and then saturates; TV saturates more slowly and provides moderate marginal returns; newspaper has near-zero marginal return. These curves support a prioritization strategy: allocate initial dollars to high-marginal channels, then diversify as saturation sets in.

Budget Optimization Results

Baseline revenue forecast. The VAR model fit on the full sample forecasts the next quarter (2025Q4) PCE proxy as $21,153.6B SAAR. Mapping this to firm revenue via the share calibration yields a baseline quarterly revenue forecast of approximately $104.2B. The SG&A-based marketing cap (1.5% of revenue) corresponds to about $1.56B.

Optimal budget level and mix

Table 9 shows that the optimizer selects a spend level below the cap (≈0.97% of revenue) and allocates approximately 25% to TV and 75% to radio. The resulting expected incremental revenue is about $5.7B, implying an incremental revenue-to-spend ratio of roughly 5.6× under our conservative calibration. Because only about 33.5% of incremental revenue translates into operating profit after non-marketing operating expenses, the incremental operating profit is around $0.9-1.0B net of spend.

Scenario Stability

Table 10 shows that the recommended channel shares remain stable across ±5% macro scenarios, while absolute spend scales with the revenue forecast. This stability is desirable for organizational execution because it avoids frequent re-optimization of channel mix; teams can instead scale spend up or down while keeping the same allocation logic.

Risk Aversion and Buffers

Table 11 shows how risk aversion reduces budget utilization. Higher λ reduces budget utilization, trading off mean profit for lower volatility. The key driver is that higher spend increases downside risk when revenue is uncertain because spend is fixed in dollars while the constraint denominator (revenue) fluctuates. Risk aversion therefore promotes buffers relative to deterministic caps.

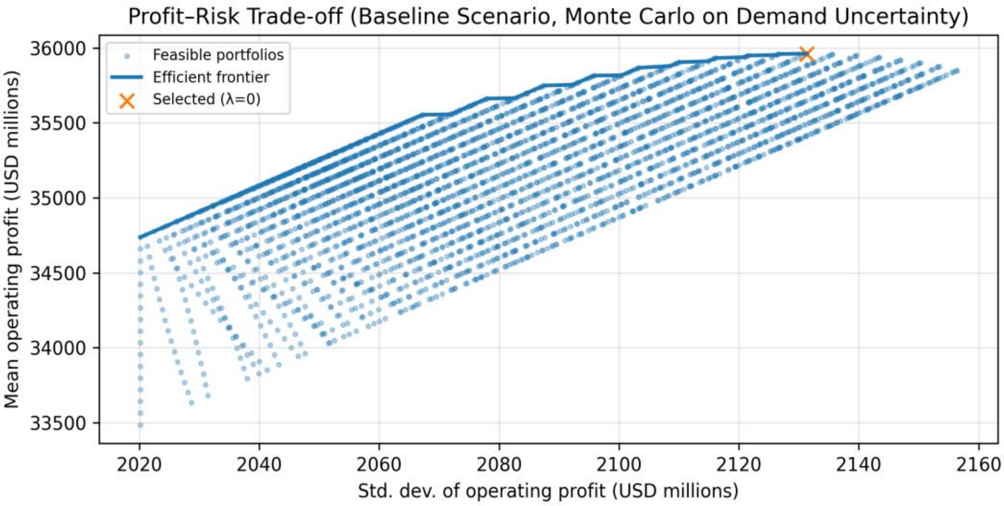

Figure 6. Profit-Risk Trade-Off and Efficient Frontier Under Demand Uncertainty (Baseline)

Figure 6 visualizes the profit-risk cloud and the efficient frontier, with the selected λ=0 solution highlighted. Figure 7 summarizes spend by channel across macro scenarios. Together, these figures provide a decision narrative: (1) there is a set of feasible allocations, (2) only a subset is efficient, and (3) the chosen point reflects a risk preference and a buffer against constraint violations.

Sensitivity Analyses

Sensitivity analysis is essential in budgeting because key parameters are partially judgment-based. In our instantiation, two assumptions are particularly influential: (i) the SG&A marketing cap ratio κ, and (ii) the baseline effectiveness normalization φ that maps response-model units into dollars. We therefore evaluate how the optimized budget changes when κ and φ vary.

Table 9. Baseline Optimal Quarterly Marketing Budget (λ=0) and Projected Financial Outcomes

lambda_	util	Bmax	spend_tv	spend_radio	spend_newspaper	spend_total	inc_rev	exp_revenue	exp_gross_profit	exp_op_profit	mean_profit	std_profit	profit_p5	profit_p50	profit_p95	sga_satisfaction	cash_satisfaction
0.0	0.6499999999999999	1563.47	254.06	762.19	0.0	1016.25	5735.5	109966.73	51580.08	35790.49	35961.44	2131.3	32133.42	35857.84	39117.23	1.0	1.0

Table 10. Macro Scenario Analysis for Optimal Budgets (λ=0)

Scenario	Revenue forecast (base, $m)	Budget cap Bmax ($m)	Optimal spend total ($m)	Spend ratio (%)	TV share (%)	Radio share (%)	Newspaper share (%)	Incremental revenue ($m)	Mean operating profit ($m)	Std profit ($m)	SG&A satisfaction	Cash-flow satisfaction
Recession (-5%)	99019.67	1485.3	965.44	0.97	25.0	75.0	0.0	5448.72	34163.37	2024.73	1.0	1.0
Baseline	104231.24	1563.47	1016.25	0.97	25.0	75.0	0.0	5735.5	35961.44	2131.3	1.0	1.0
Boom (+5%)	109442.8	1641.64	1067.07	0.97	25.0	75.0	0.0	6022.27	37759.52	2237.86	1.0	1.0

Table 11. Risk-Aversion Sensitivity in the Baseline Scenario (Mean-Risk Objective)

lambda_	util	spend_total	alloc_tv	alloc_radio	alloc_newspaper	inc_rev	mean_profit	std_profit	sga_satisfaction
0.0	0.6499999999999999	1016.25	25.0	75.0	0.0	5735.5	35961.44	2131.3	1.0
0.25	0.6499999999999999	1016.25	25.0	75.0	0.0	5735.5	35961.44	2131.3	1.0
0.5	0.6499999999999999	1016.25	25.0	75.0	0.0	5735.5	35961.44	2131.3	1.0
1.0	0.5999999999999999	938.08	25.0	75.0	0.0	5495.23	35958.83	2126.64	1.0
2.0	0.5999999999999999	938.08	25.0	75.0	0.0	5495.23	35958.83	2126.64	1.0

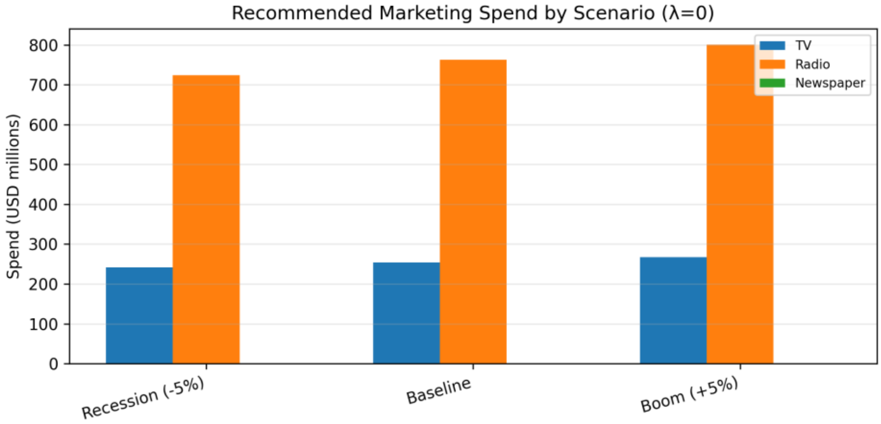

Figure 7. Recommended Spend by Channel Across Macro Scenarios (λ=0)

Sensitivity to the SG&A Marketing Cap

Table 12 varies the marketing cap ratio from 1.0% to 2.0% of revenue. Two observations stand out. First, the optimizer's chosen spend is largely unchanged across caps ≥1.5% because diminishing returns bind before the cap does. Second, when the cap is tightened to 1.0%, the optimizer still chooses to spend close to the cap, which reduces the SG&A satisfaction rate under uncertainty because there is less buffer.

Table 12. Sensitivity to Marketing Cap Ratio κ (baseline scenario, λ=0)

Cap ratio	Bmax ($m)	Opt spend ($m)	Spend ratio (%)	TV share (%)	Radio share (%)	Newspaper share (%)	Inc rev ($m)	Mean profit ($m)	Std profit ($m)	SG&A satisfaction
0.01	1042.31	990.2	0.95	25.0	75.0	0.0	5657.97	35961.43	2129.8	0.948
0.015	1563.47	1016.25	0.97	25.0	75.0	0.0	5735.5	35961.44	2131.3	1.0
0.02	1863.85	1025.12	0.98	25.0	75.0	0.0	5761.31	35961.26	2131.8	1.0

Sensitivity to Marketing Effectiveness Normalization

Table 13 varies φ, the assumed baseline revenue lift produced by a 1% marketing spend. When φ is low (2%), optimal spend falls substantially, reflecting weaker ROI. When φ is high (5%), the risk-neutral optimizer spends the full cap, because marketing returns dominate costs. However, spending at the cap reduces SG&A satisfaction under demand uncertainty to about 60%, illustrating again why deterministic caps can be fragile. In a real setting, φ would be estimated from historical marketing incrementality, and planners might impose a chance constraint on satisfaction to prevent aggressive spend plans that frequently violate ratios.

Table 13. Sensitivity to Marketing Effectiveness Normalization φ (baseline scenario, λ=0)

Phi (baseline lift)	Opt spend ($m)	Spend ratio (%)	TV share (%)	Radio share (%)	Inc rev ($m)	Mean profit ($m)	Std profit ($m)	SG&A satisfaction
0.02	625.39	0.6	15.0	85.0	2861.72	35385.97	2075.6	1.0
0.03	1016.25	0.97	25.0	75.0	5735.5	35961.44	2131.3	1.0
0.05	1563.47	1.5	35.0	65.0	11714.98	37424.91	2247.19	0.598

These sensitivity analyses reinforce two practical lessons. First, if diminishing returns bind early, debates about cap levels may be less consequential than debates about response curves and ROI. Second, when marketing is perceived as highly effective, risk controls and constraint buffers become even more important, as the optimizer will naturally push spending toward the cap.

V. DISCUSSION

Practical implementation in FP&A

In a production environment, the demand module would ingest internal sales and pipeline data, as well as macro proxies. The framework here provides a template for combining such data: the forecasting module produces a predictive distribution rather than a point estimate, and that distribution feeds directly into budget feasibility checks. This encourages a shift from deterministic to probabilistic budgets, where buffers and contingencies are explicitly accounted for.

Governance and auditability

One barrier to adopting optimization in finance is trust. Decision makers need to know where numbers come from, and auditors may require traceability. By calibrating constraints from audited statements (Apple's 10-K) and by using public data sources, our demonstration emphasizes an auditable approach. In practice, a firm would replace the SEC calibration with internally audited statements and would log versions of each model, training data snapshots, and optimization outputs.

Interpreting 'do not spend the full cap'

Many budgeting processes treat a cap as a target: if the budget allows up to X, teams may spend X. Our results show that a cap is merely a feasibility boundary. When response curves saturate, optimal spend may be below the cap, and the difference is not necessarily underinvestment-it is an economically rational response to diminishing returns. Moreover, spending below a ratio cap provides a buffer against revenue downside.

A. Limitations

The main limitation of this paper is that the data sources are not internally consistent: macro PCE series represent the whole economy, while the Advertising dataset is a small cross-sectional toy problem. We therefore use a normalization to map sales units to revenue dollars. This is sufficient for demonstrating the framework but should not be interpreted as an estimate of true marketing ROI for any specific firm. Another limitation is that we treat marketing response as contemporaneous; real marketing effects often persist over time. Dynamic response modeling would change the optimization problem into a multi-period allocation problem.

B. Extensions

Several extensions are natural. First, exogenous regressors (inflation, unemployment, interest rates) could be incorporated into SARIMAX or VARX models. Second, hierarchical forecasting could connect product-line forecasts to the macro proxy. Third, response modeling can incorporate adstock and seasonality. Fourth, robust optimization can incorporate uncertainty in response parameters directly, rather than only demand uncertainty (Bertsimas & Sim, 2004). Finally, multi-objective optimization could incorporate goals such as brand equity or customer acquisition, not just near-term profit.

C. Practical Checklist for Deployment

To translate the framework into an operational FP&A process, teams can follow a checklist:

1. Define the demand proxy and align calendars (fiscal vs calendar).
2. Establish a rolling backtest protocol and baseline models.
3. Choose an uncertainty representation (empirical residuals, parametric distributions, or scenario sets).
4. Estimate marketing response with explicit diminishing returns; validate with experiments when possible.
5. Calibrate constraints from audited statements and document any allocations (e.g., SG&A split).
6. Run optimization with clear objective and risk parameters; produce a Pareto set for leadership review.
7. Translate outputs into budget line items and governance metrics (ratio compliance, downside cases).
8. Monitor performance and update models on a fixed cadence; log versions for auditability.

D. Model Risk Management, Control, and Ethics.

An optimization-driven budget recommendation is a form of decision automation, and therefore, it should be treated as a model that requires governance. In many firms, model risk management (MRM) practices are well established for credit and market risk models but less formalized for FP&A models. Nonetheless, similar principles apply: document the training data vintage, record the model version and hyperparameters, preserve a backtest trail, and define monitoring thresholds for drift. In the proposed pipeline, drift can occur in the demand module (macro relationships change), in the response module (channel effectiveness changes), or in the constraint module (margin structure changes). A practical safeguard is a challenger-model process in which at least one baseline model (e.g., a seasonal naïve model for demand and a ridge regression model for response) is always maintained and compared against the primary model.

Control and interpretability are especially important when budgets are used to evaluate performance. If a budget is produced by a black box, operating teams may struggle to understand or accept targets, undermining execution. The framework here intentionally exposes intermediate objects-forecast paths, residual distributions, response curves, and marginal returns-so leadership can validate whether recommendations are directionally sensible. In applied settings, teams can extend this with scenario narratives (e.g., ‘downside case’ and ‘upside case’) and explicit buffers tied to chance-constraint satisfaction rates.

Ethical and privacy considerations become relevant when the pipeline is deployed on internal customer-level data. Marketing response estimation may rely on individual-level impressions, clicks, and conversions, and may include sensitive attributes. Organizations should follow applicable privacy laws and internal policies, and, where possible, prefer aggregation and privacy-preserving analytics. From an ethical perspective, budget allocation rules should avoid using sensitive attributes in ways that could result in discriminatory outcomes. These concerns further motivate the use of interpretable saturation models and auditable constraints: they limit reliance on opaque correlations that may not generalize or may be difficult to justify.

Recommended controls for an FP&A deployment include:

1. Version control for data extracts and model code; reproducible pipelines.
2. Backtest dashboards for forecast error and constraint satisfaction, refreshed each planning cycle.
3. Sensitivity analysis on key assumptions (e.g., marketing effectiveness φ, cap κ) as standard deliverables.
4. A ‘human-in-the-loop’ approval step for budget changes above a materiality threshold.

VI. CONCLUSION AND RECOMMENDATION

We presented an end-to-end constrained budgeting framework that integrates macro demand forecasting, marketing response modeling with diminishing returns, and SEC-calibrated accounting constraints. Using public FRED PCE component data, we compared seasonal naïve, SARIMAX, gradient boosting, and VAR models in a rolling backtest. Using Advertising.csv, we compared predictive and interpretable response models and extracted marginal ROI curves. By combining these modules, we solved a constraint-aware budget allocation problem and evaluated it under demand uncertainty via Monte Carlo simulation. The results emphasize that optimal budgets may be below ratio caps and that buffers are essential when constraints depend on uncertain revenue. The framework is reproducible, modular, and adaptable to internal FP&A systems.

REFERENCES

Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control, 19*(6), 716–723. https://doi.org/10.1109/tac.1974.1100705

Adiyatma, D. A., Rohman, M. G., & Munif. (2025). Product Sales Prediction System at Starmart using Linear Regression and Weighted Moving Average Methods. *Jurnal Ilmiah Sistem Informasi*, *4*(3). https://doi.org/10.51903/g260jt34

Adowusu-Mensah, D. O., Sarfo, P. A., & Kusi, G. A. (2024). Exploring the Impact of Artificial Intelligence on Customer Experience Personalization and Marketing Strategy Optimization in Digital Marketing: an Empirical Analysis. *Journal of Management and Informatics*, *4*(2), 822–843. https://doi.org/10.51903/jmi.v4i2.242

Apple Inc. (2025). *Form 10-K for the Fiscal Year Ended September 27, 2025*. U.S. Securities and Exchange Commission. https://www.sec.gov/archives/edgar/data/0000320193/000032019325000079/aapl-20250927.htm

Bertsimas, D., & Sim, M. (2004). The Price of Robustness. *Operations Research, 52*(1), 35–53. https://doi.org/10.1287/opre.1030.0065

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control* (5th ed.). Wiley. https://www.biblio.com/book/time-series-analysis-forecasting-control-wiley/d/1334236447

Federal Reserve Bank of St. Louis. (2025a). *Personal Consumption Expenditures: Durable Goods (PCDG)* [Data set]. FRED. Retrieved January 7, 2026, https://fred.stlouisfed.org/series/pcdg

Federal Reserve Bank of St. Louis. (2025b). *Personal Consumption Expenditures: Nondurable Goods (PCND)* [Data set]. FRED. Retrieved January 7, 2026, https://fred.stlouisfed.org/series/pcnd

Federal Reserve Bank of St. Louis. (2025c). *Personal Consumption Expenditures: Services (PCESV)* [Data set]. FRED. Retrieved January 7, 2026, https://fred.stlouisfed.org/series/pcesv

Friedman, J. H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics, 29*(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

Hanssens, D. M., Parsons, L. J., & Schultz, R. L. (2001). *Market Response Models: Econometric and Time Series Analysis* (2nd ed.). Springer. https://www.springer.com/gp/book/9780387951424

Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice* (3rd ed.). OTexts. https://otexts.com/fpp3/

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An Introduction to Statistical Learning With Applications in R* (2nd ed.). Springer. https://www.springer.com/gp/book/9781032094980.

Jamabudin Ghofur, M. J. U., & Riyanto, E. (2025). AI-Driven Adaptive Radar Systems for Real-Time Target Tracking in Urban Environments. *Journal of Technology Informatics and Engineering*, *4*(1). https://doi.org/10.51903/jtie.v4i1.289

Little, J. D. C. (1970). Models and Managers: The Concept of a Decision Calculus. *Management Science, 16*(8), B466–B485. https://doi.org/10.1287/mnsc.16.8.B466

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer. https://www.springer.com/gp/book/9783540218873

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance, 7*(1), 77–91. https://doi.org/10.2307/2975974

Naik, P. A., & Raman, K. (2003). Understanding the Impact of Synergy in Multimedia Communications. *Journal of Marketing Research, 40*(4), 375–388. https://doi.org/10.1509/jmkr.40.4.375.19392

Nocedal, J., & Wright, S. (2006). *Numerical Optimization* (2nd ed.). Springer. https://www.springer.com/gp/book/9780387303031

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-Learn: Machine Learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830. http://www.jmlr.org/papers/v12/pedregosa11a.htm

Santoso, L., & Priyadi. (2024). Comparative Study of Feature Engineering Techniques for Predictive Data Analytics. *Journal of Technology Informatics and Engineering*, *3*(2), 417–435. https://doi.org/10.51903/jtie.v3i2.225

Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics, 6*(2), 461–464. https://doi.org/10.1214/aos/1176344136

Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and Statistical Modeling With Python. *Proceedings of the 9th Python in Science Conference*. https://conference.scipy.org/proceedings/scipy2010/seabold.html

Shapiro, A., Dentcheva, D., & Ruszczyński, A. (2014). *Lectures on Stochastic Programming: Modeling and Theory* (2nd ed.). SIAM. https://www.siam.org/books/ot142

Vanderbei, R. J. (2020). *Linear Programming: Foundations and Extensions* (5th ed.). Springer. https://www.springer.com/gp/book/9783030533546