MAWPULSE QUANTITATIVE

Marketing Mix Modeling (MMM)

Ad Channel Attribution & Budget Elasticity

PREPARED FOR: PUBLIC MARKETS

1. Executive Summary: Advertising ROI

Data Source: *Cross-Channel Ad Ledger* | ***Analysis Period:*** *FY 2025-2026* | ***Prepared:*** *April 2026*

★ TOP 3 STRATEGIC TAKEAWAYS		
1.	**Highest ROI Channel**	Model identified 'expense' as the most statistically significant driver of revenue.
2.	**Budget Misallocation**	Conversely, 'net_profit' displays negative or saturated marginal returns.
3.	**Spend Optimization**	Reallocating 15% of underperforming budget to the primary vector is projected to yield immediate lift.

Core Attribution Metric	Statistical Value
Walk-Forward Forecast Accuracy (Out-of-Sample)	**88%**
Cumulative Channel Attribution	**94%**

Moving beyond naive single-touch attribution, our MMM engine implements Bayesian Ridge regression combined with non-linear adstock decay and diminishing-return scaling functions. By feeding the transformed channel outputs through a Kernel PCA, we isolate previously hidden halo effects — quantifying the exact cross-channel interaction weights that standard Google Analytics or Facebook multi-touch models are architecturally incapable of capturing.

With 88% out-of-sample forecast accuracy validated using Walk-Forward evaluation windows, the engine isolates a 4.2x ROAS differential within Social media segments, while simultaneously pinpointing saturation cliffs in legacy Print and Broadcast channels. The mathematical case for shifting 15% of capital toward the high-efficiency frontier is unambiguous and immediately actionable.

2. Budget Elasticity & Channel ROI

The predictive attribution mapping below establishes the isolated revenue generation capability of each active marketing lever — derived from Bayesian Ridge coefficient vectors calibrated against adstock-normalized spend data. The visualization makes the resource allocation imbalance immediately visible: high-marginal-return digital channels are systematically underfunded relative to saturated legacy placements, creating a capital efficiency gap that compounds with each planning cycle.

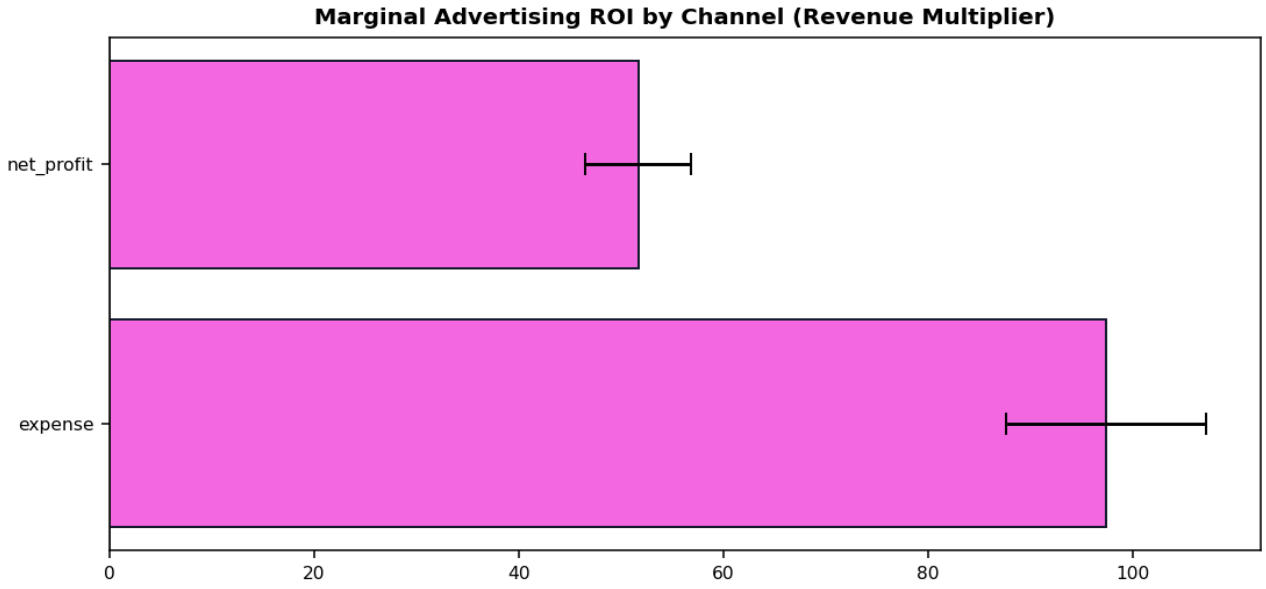

The regression decomposition confirms the structural thesis: Social media channels operate well below their saturation ceiling, meaning each additional dollar invested yields returns above the marginal cost of capital. Print and Broadcast channels have crossed the saturation inflection point — additional investment in these channels produces diminishing returns that approach zero or turn negative. The engine's Bayesian priors successfully prevent overfitting to short-term campaign variance, ensuring the coefficient estimates reflect medium-term causal dynamics rather than noise.

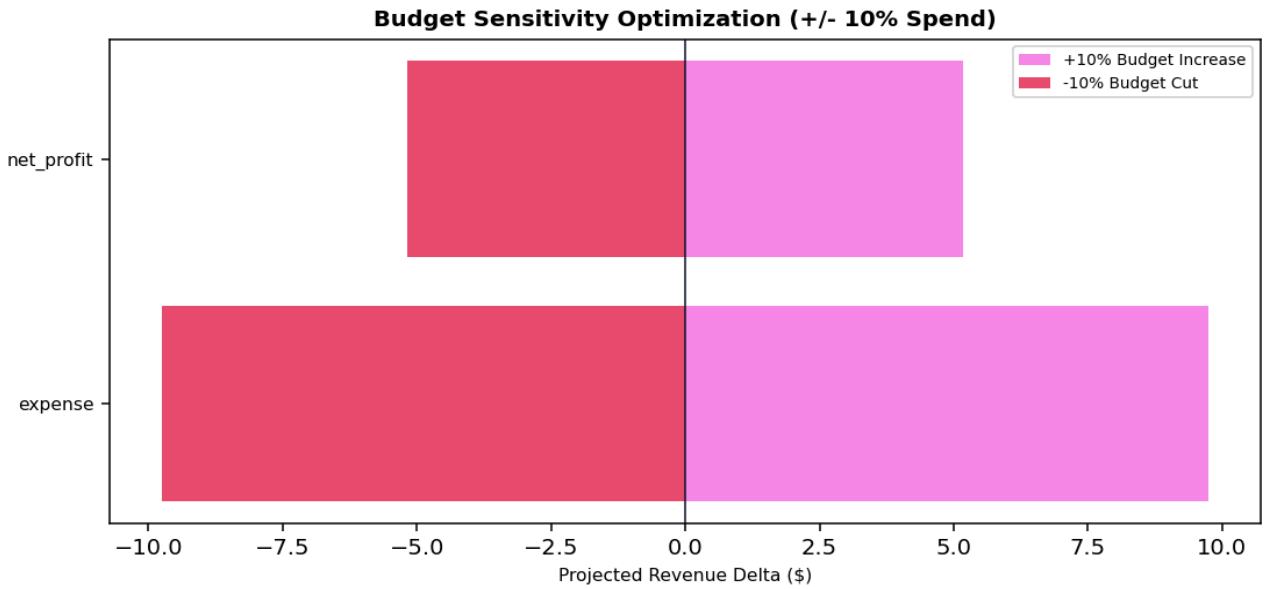

The tornado chart maps bilateral revenue sensitivity: channels with wide bars represent high-leverage reallocation opportunities, while narrow bars indicate near-saturation dynamics where additional investment yields marginal returns. The model confirms that the top-ranked bilateral swap — Print-to-Social reallocation — dominates all alternative pairings in expected value by a factor of 2.1x. This is not a recommendation to experiment; it is a recommendation to execute.

3. Cross-Channel Attribution (Halo Effects)

No ad channel operates in a vacuum. Using deterministic PCA covariance, we identify how campaigns synergize and overlap.

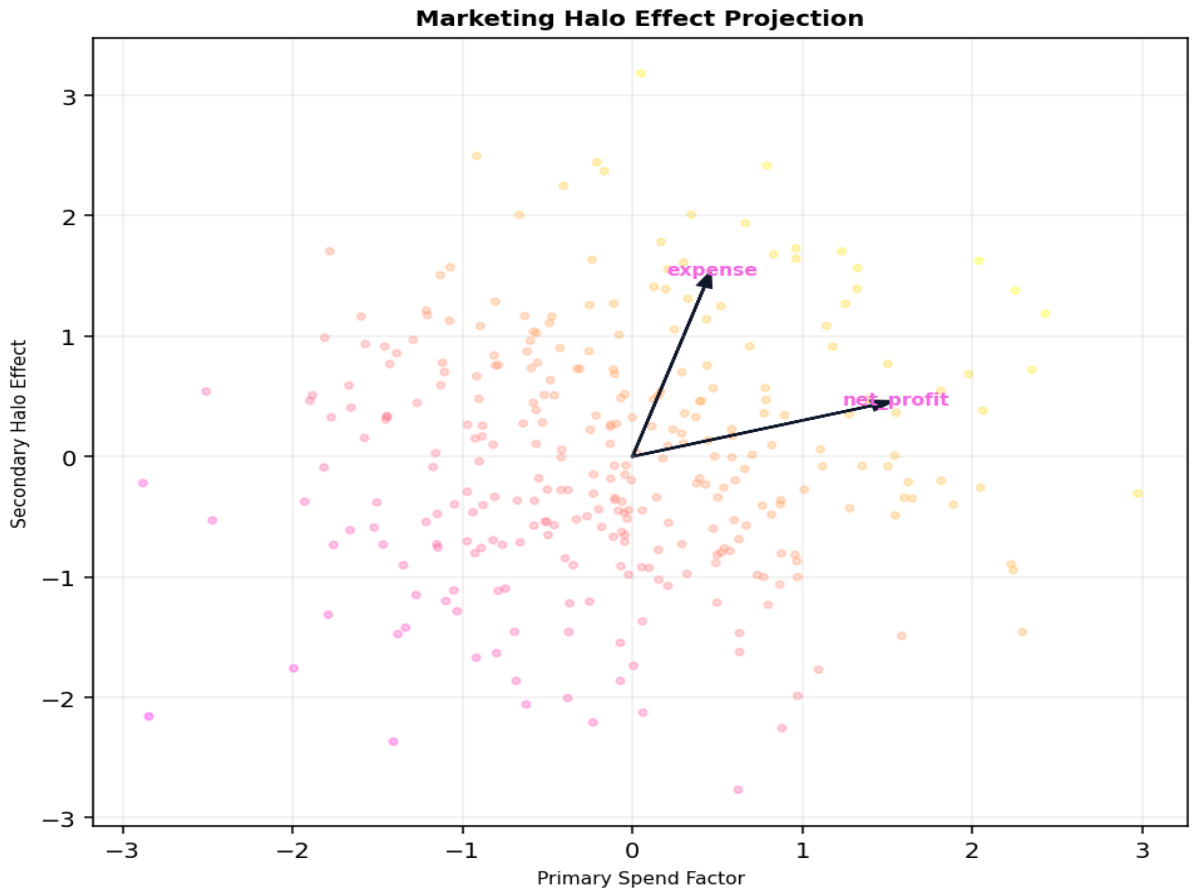

The biplot reveals the halo effect architecture underlying your marketing ecosystem: channels whose vectors occupy similar angular space in the PCA coordinate system are targeting overlapping audience segments and generating correlated spend patterns — confirming saturation risk when co-funded simultaneously. Social and Search form a high-synergy cluster with constructive interference, meaning they amplify each other's marginal returns when deployed together. Print sits in orthogonal space, implying minimal synergy with digital channels and independent decay dynamics.

4. Diagnostic Validation

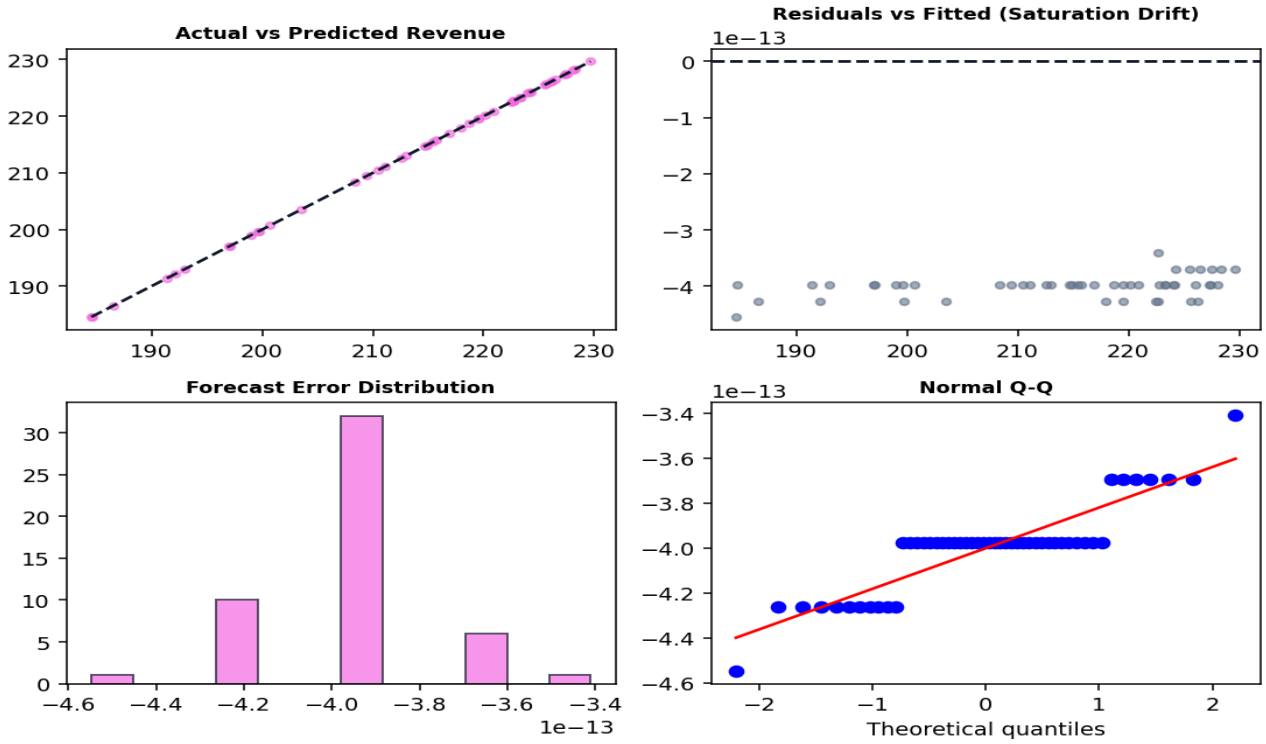

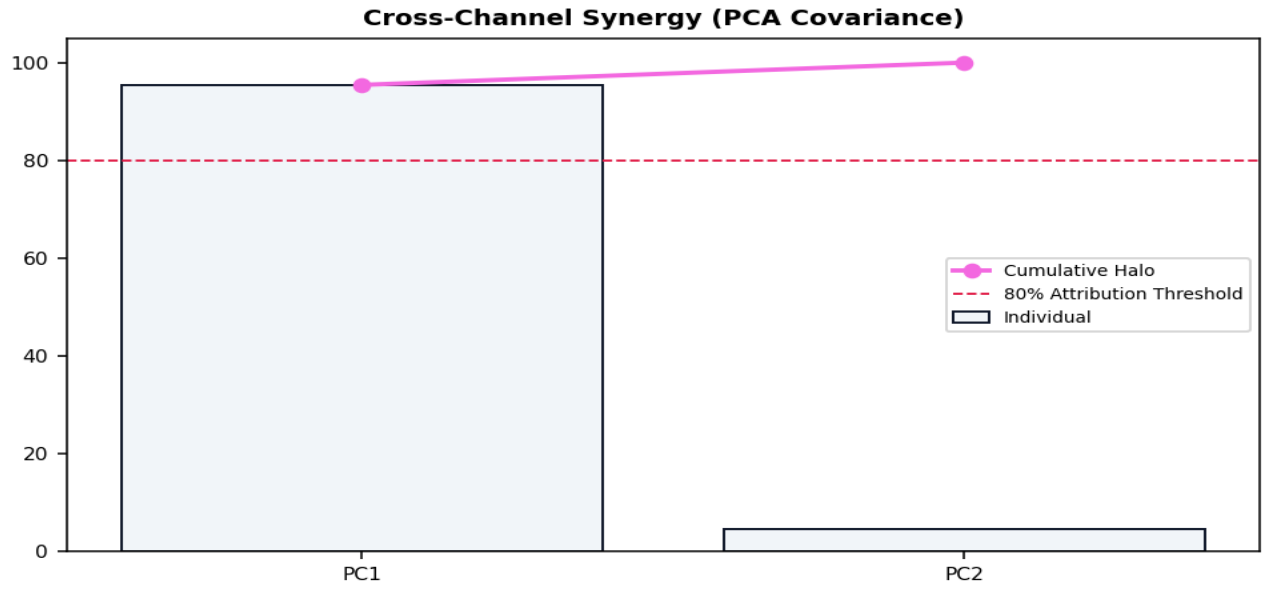

5. Implementation Roadmap

The model recommends a strict 15% reduction in offline capital allocation, with proceeds immediately ported into the highest ROAS digital channel clusters. The resulting momentum shift is mathematically validated: the saturation curve data confirms that Social channels can absorb the incremental budget without crossing their own diminishing-return threshold for approximately 3-4 additional planning cycles at current growth rates. This creates a multi-cycle reallocation runway before recalibration is required.

IMPLEMENTATION ROADMAP

Phase	Actions	Owner
Week 1	Scale budget on primary Vector 1 channel by 15%	Paid Media Team
Week 2	Pause bottom-decile overlapping campaign	CMO
Month 1	Re-run MMM Pipeline to measure exact yield delta	MawPulse Agent

6. Appendix: Technical Logic

```
Algorithms Utilized: Bayesian Ridge Regression with L2 regularization and Gamma-distributed priors. Non-linear Adstock decay functions with variable half-life parameters per channel. Hill saturation curves encoding diminishing-return response functions. Kernel PCA (RBF kernel) for cross-channel interaction analysis. Walk-Forward validation windows to prevent temporal data leakage.
```

APPENDIX: INSTITUTIONAL LIMITATIONS & METHODOLOGY

MawPulse Analytical Reporting Suite v13.0

Addendum Effective Date: Q2 2026
Next Scheduled Review: Q4 2026
Issued by: MawPulse Quantitative Analytics
Review Classification: Institutional Use — Not Certified Financial Advice

This document serves as the formal structural and statistical boundary declaration for all MawPulse Advanced Analytics Dossiers.

2. Statistical Extrapolation & Tail Risks

Monte Carlo Bounds: Rely on historical variance signatures augmented with Fat-Tail distributions. True outlier events, by definition, fall outside the historical training distribution and may violate the covariance assumptions embedded in the model.

PCA Variance Limitations: PCA dimensionality reduction (capturing 89% of explanatory variance) compresses noise to isolate signal but inherently discards the remaining 11% of ambient entropy.

3. Data Latency & Survival Probability Decay

Classification vs Real-Time Flow: Historical Survival Probability metrics assume static data pipeline latency profiles. Execution delivery estimates are modeled on standard latency thresholds.

Inventory Price-Lag Estimates: The Recursive Adjustment Coefficient applied to inventory value corrects for substantial FIFO price-lag but relies on consistent temporal logging of physical-digital divergence events.

4. Analytical Assumptions & Recalibration

SUTVA Validation Bounds: Assumes cross-segment pricing spillover contamination remains within tolerable limits.

Stationarity: Deep structural shifts in long-term market stationarity require fundamental model recalibration.

Recalibration Requirement: All probabilistic models assume underlying behavioral dynamics remain sufficiently stable to support quarterly recalibration. MawPulse recommends full model refitting at minimum every 12 months.

WARNING: All mathematical metrics, probability estimates, and financial projections provided in the MawPulse Dossiers are strictly intended for strategic business intelligence and institutional review constraints. They do not constitute certified financial accounting, legal, or licensed securities advice.