From Specification to Regularization: A Bayesian Seasonal Modulator for Marketing Mix Models

Emanuele Fedrigolli, Nicola Spiller, Marco Mandolfo, Giuliano Noci
Politecnico di Milano

Abstract: Allocating advertising budgets across channels and over the year requires reliable estimates of channel effectiveness. Marketing mix models provide such estimates from aggregate data, but two established empirical regularities complicate the task: the relationship between spending and outcomes is nonlinear, with diminishing returns, and advertising effectiveness itself varies over time. Existing frameworks address these phenomena separately, and recent work shows that the two are often conflated (a manifestation of the broader endogeneity concern in observational marketing data), producing indistinguishable model fit yet divergent budget recommendations. We address this through a theory-driven specification. Drawing on the Elaboration Likelihood Model, we argue that seasonal demand peaks shift the consumer pool toward higher involvement, increasing average advertising effectiveness in a predictable, periodic pattern. This motivates a Bayesian seasonal modulator that scales channel contributions multiplicatively, producing time-varying effectiveness estimates constrained to seasonal structure. Simulations show over 50% reduction in channel attribution error, and three firm-level applications demonstrate that the modulator self-regulates, strong where seasonality is high and converging toward unity where it is low.

Keywords: *Marketing Mix Modeling, Advertising Effectiveness, Bayesian Regularization*

Track: *Track 5 — Marketing Analytics & Research Methods*

Paper presented at the EMAC Fall Conference 2026, University of Bremen, 16-18 September 2026.
Conference version.

1. Introduction

Marketing mix models (MMMs) estimate the effectiveness of marketing channels from aggregate time-series data on spending and business outcomes. Two empirical regularities are established in the academic literature: the relationship between spending and outcomes is nonlinear, with increasing expenditure yielding diminishing marginal returns (Little, 1970; Jin et al., 2017), and advertising effectiveness is not stable over time (Winer, 1979; Bruce, 2008; Guhl et al., 2018). Both regularities coexist in practice, yet existing MMM frameworks address them separately: nonlinear models treat coefficients as static, while time-varying models embed response nonlinearity in a fixed functional form (Hanssens et al., 2001; Ng et al., 2022). Neither approach explicitly separates the seasonal modulation of advertising channel effectiveness from nonlinear saturation within a single identified structure.

Marketing strategy research has increasingly recognized that observational data on firm decisions present an endogeneity concern: firms do not allocate resources at random but respond to demand conditions, generating omitted-variable bias when the same conditions also drive outcomes (Rutz & Watson, 2019). The empirical magnitude of this concern is substantial, with observational advertising estimates documented to diverge from experimental benchmarks by up to a factor of three (Gordon et al., 2019). Within MMMs specifically, this concern manifests as the conflation of nonlinear and time-varying effects identified by Dew et al. (2024): when seasonal demand drives both spending and effectiveness, a flexible time-varying coefficient can absorb the channel × season interaction that a static specification leaves unmodeled, producing two models that fit equally well yet imply divergent budget recommendations.

We argue that the Elaboration Likelihood Model (Petty et al., 1983) provides a theoretical basis for constraining this separation. Seasonal demand peaks increase the share of consumers with an active category need, raising their felt involvement with category-related information (Celsi and Olson, 1988) and thereby strengthening advertising persuasion. At the aggregate level, this compositional shift implies that advertising effectiveness varies periodically with demand conditions, a prediction supported empirically by Gijsenberg (2017). This periodic structure motivates a regularization approach: rather than choosing between a nonlinear or a time-varying specification, we constrain the time-varying component to seasonal patterns, so that both effects coexist within an identified structure. We operationalize this through a Bayesian seasonal modulator that scales advertising channel contributions via a Fourier basis, producing effectively time-varying coefficients whose variation is restricted to periodic patterns. Because the saturation function depends on spending level alone and the modulator on calendar time alone, the two components operate on structurally distinct dimensions, limiting their ability to absorb each other's signal.

This paper makes three contributions. First, we bridge individual-level involvement theory and aggregate marketing mix modeling: drawing on the Elaboration Likelihood Model, we derive the

prediction that seasonal demand conditions modulate advertising effectiveness through a compositional shift in consumer involvement, and we operationalize this prediction within a Bayesian modeling framework. Second, we propose the seasonal modulator, a parsimonious mechanism requiring only a small number of additional Fourier parameters that produces time-varying ROAS estimates for seasonal budget allocation. Third, we provide evidence through simulations with known ground truth, where the model reduces channel contribution error by more than 50%, and three empirical applications spanning low, medium, and high seasonality, where the modulator self-regulates: strong where seasonality is high, converging toward unity where it is low.

2. Theoretical Background

2.1 The nonlinearity-time-variance trade-off

With aggregate time-series data of modest length, typically 100 to 200 weekly observations, simultaneously estimating a flexible nonlinear response function and freely evolving time-varying coefficients is intractable without structural restrictions (Hanssens et al., 2001), yet both shape advertising budget decisions.

Beyond degrees of freedom, the trade-off has a causal substrate. Firms typically concentrate advertising spending during peak-demand periods, generating co-movement between spending and seasonal demand (Rutz & Watson, 2019). When advertising effectiveness itself varies with the same seasonal conditions for theoretical reasons we develop in Section 2.2, a static-coefficient specification cannot disentangle the channel × season interaction from average channel effectiveness, biasing β upward for channels whose timing is most concentrated in peaks.

Dew et al. (2024) demonstrate that this trade-off produces a deeper problem. They identify seasonality as a paradigmatic case: "if seasonality determines marketing effectiveness, and spending on marketing is set following such seasonality, then seasonality is a parent process" (p. 13) that confounds the two specifications. A nonlinear model prescribes spending that balances marginal returns along a stable curve; a time-varying model prescribes spending that adapts to current elasticity. These recommendations diverge under intervention, even when both models fit observed data equally well. Dew et al. propose guardrails including multi-market data and targeted experiments. Our approach is complementary: rather than enriching the data, we constrain the model structure so that the time-varying component cannot approximate the nonlinear one (Section 2.3).

2.2 Why effectiveness varies seasonally

Among the sources of time-varying effectiveness, seasonal demand conditions are particularly relevant for aggregate MMMs. The Elaboration Likelihood Model (ELM; Petty et al., 1983) predicts that consumers process persuasive messages more thoroughly when their personal involvement is high, following a central route that produces stronger and more durable attitude change. Involvement, in

turn, depends on the situational relevance of the message: when consumers carry an active category need, this need acts as a situational source of personal relevance that raises their felt involvement with category-related information (Celsi and Olson, 1988). During seasonal demand peaks, a larger share of consumers in the market carries such an active need. This compositional shift toward higher average involvement implies, via the ELM, that advertising messages are elaborated more deeply and produce stronger persuasion effects at the aggregate level. Gijsenberg (2017) provides direct empirical support: across 252 brands in 61 consumer packaged goods categories, advertising elasticities are significantly stronger during intrayear demand peaks.

Because the ELM-based mechanism operates through seasonal demand cycles, it predicts that the time-varying component of advertising effectiveness follows a periodic pattern. In practice, firms also concentrate advertising budgets during peak-demand periods, creating the co-movement between spending and effectiveness that enables conflation (Section 2.1). This motivates constraining the time-varying component to periodic structure derived from theory rather than leaving it unconstrained.

2.3 How periodic constraints reduce conflation

For advertising budget allocation, it matters whether higher effectiveness during peak seasons reflects genuine seasonal modulation or is an artifact of diminishing returns at higher spending levels. The saturation function Sat(x) maps spending to a concave, monotonically increasing output with no dependence on calendar time. The seasonal modulator M(t), constructed from a low-order Fourier basis, produces smooth periodic patterns with no dependence on spending level. A periodic function of time is unlikely to replicate the monotone concave-in-x shape of a saturation function, which depends on how much is spent, not when. Conversely, a time-invariant saturation function cannot capture periodic fluctuations. When observed spending is not itself purely periodic, spending level and calendar time are not perfectly collinear, and the two components remain distinguishable. This is not a claim of perfect identification: when spending is highly seasonal, residual ambiguity may persist. However, the structural constraints substantially reduce the space of observationally equivalent models relative to unconstrained specifications. We verify this computationally in Section 4.

3. Modeling Framework

3.1 Model structure

The observed outcome y(t) is decomposed as:

$$y(t) = Base(t) + \Sigma_c \left[\beta_c \cdot M(t) \cdot Sat_c(Adstock_c(x_c(t))) \right] + \varepsilon(t) \qquad (1)$$

where epsilon(t) ~ Normal(0, sigma). Base(t) captures non-marketing drivers through a smooth trend and additive Fourier seasonality, which models periodic variation in the base demand level. Carryover effects are captured via geometric adstock with retention rate alpha in (0,1). The saturation function

Sat(x) = log(kx) captures nonlinear diminishing returns (e.g., doubling TV spend does not double sales) as a concave, monotonically increasing function of spending alone.

3.2 Seasonal modulator

The seasonal modulator is the key extension of this modeling framework. Unlike additive seasonality, which shifts the base demand level, the modulator operates multiplicatively on advertising channel contributions, capturing the involvement-based prediction that seasonal demand conditions, by shifting the composition of consumers toward those with active category need, amplify or attenuate channel effectiveness (Section 2.2). It is constructed from a Fourier basis passed through a positive-valued transformation g:

$$M(t) = g(\Sigma_{n=1}^{N} [a_n \cdot sin(2\pi nt / P) + b_n \cdot cos(2\pi nt / P)]) \quad (2)$$

The effective coefficient at time t is:

$$\beta^{eff}_c(t) = \beta_c \cdot M(t) \quad (3)$$

The coefficients are effectively time-varying, but the variation is constrained to periodic patterns through the Fourier basis and to positive values through g. The static beta_c captures average channel effectiveness; M(t) captures how seasonal conditions modulate it, yielding time-varying ROAS estimates by channel. Because M(t) depends only on calendar time and Sat depends only on spending level, the two components remain structurally distinguishable (Section 2.3). M(t) is applied after saturation, preserving this separation. In our implementation, M(t) is shared across channels for parsimony; channel-specific extensions are discussed in Section 6.

3.3 Estimation

All parameters are estimated jointly via MCMC (NUTS, 4 chains, 2,000 draws). Channel coefficients receive positive priors calibrated on advertising elasticities (Korkames et al., 2025).

4. Simulation Study

We evaluate the model through a simulation with known ground truth, allowing direct assessment of advertising channel-level decomposition accuracy.

4.1 Design

We generate five years of weekly data (T = 260: 208 training, 52 test) from the model in Section 3, with six advertising channels representative of a typical media mix (TV, Display, Search, Social, Radio, and Out-of-Home), each with distinct spending patterns reflecting common media planning practices such as seasonal flighting and budget concentration. The data-generating process uses logarithmic saturation, geometric adstock, a constant baseline, additive Fourier seasonality, and a shared modulator with two harmonics. One channel (OOH) has a true effectiveness coefficient of zero.

We compare two specifications: the *Modulated* model (Section 3, modulator active) and a *Standard* ablation that removes the modulator entirely, retaining only additive seasonality. All other components are identical.

4.2 Results

Aggregate fit. Both models achieve comparable prediction accuracy (MAPE 1.97% Modulated vs. 2.03% Standard in-sample; 2.08% vs. 1.98% out-of-sample), confirming that prediction accuracy alone cannot distinguish which model produces more accurate advertising budget recommendations, consistent with the conflation premise (Dew et al., 2024). The estimated modulator correlates 0.894 with the true modulator, both in-sample and out-of-sample, indicating that the model recovers the seasonal effectiveness pattern and generalizes without overfitting.

Contribution accuracy. Although both models fit the aggregate data equally well, they produce substantially different decompositions. Table 1 reports the MAPE of estimated advertising channel contributions relative to ground truth. The Modulated model recovers total contributions with MAPE 6.9% vs. 14.2% for Standard, a reduction of more than 50%. The improvement is largest for TV (from 15.3 to 7.6%), the highest-contributing channel. Radio is the only exception: its spending pattern closely mirrors the modulator's seasonal shape, creating local confounding that channel-specific modulators would address. Both models correctly assign negligible contributions to OOH (true beta = 0), confirming that the modulator does not spuriously activate inactive channels. These contribution errors translate directly into biased ROAS estimates and suboptimal advertising budget allocation.

Table 1. Channel contribution MAPE (%), training set. Lower is better.

Channel	Modulated	Standard
TV	**7.6**	15.3
Display	**12.1**	17.4
Search	17.2	16.5
Social	**27.4**	39.0
Radio	83.1	**73.2**
Total	**6.9**	**14.2**

5. Empirical Application

5.1 Data and setup

We estimate both models on weekly marketing mix data from three Italian companies: (i) a non-profit charity (outcome: donations, 5 paid advertising channels including TV, social, and search), characterized by strong seasonal peaks around the holiday period; (ii) a telecommunications company (outcome: fiber connections, 4 paid advertising channels), with moderate seasonality; and (iii) an

insurance company (outcome: auto insurance, 8 paid advertising channels spanning TV, video, display, social, and search), with weak seasonality. Training periods range from 117 to 132 weeks, with 26 weeks held out for testing. All specifications follow Section 3. All models converge without divergences (R-hat = 1.0, ESS > 2,500).

5.2 Results

Fit and modulator. Table 2 summarizes the results. In-sample, both models achieve comparable fit in all three cases. Out-of-sample, the Modulated model improves prediction where seasonality is meaningful (charity: -1.7pp WMAPE; telecom: -1.0pp) and performs comparably where it is low (insurance: +0.3pp). The estimated modulator exhibits a gradient across the three cases: strong in the charity case (range 0.31-2.63, ratio 8.4x, peak in December), moderate in telecom (0.65-1.25, 1.9x, peak in September), and nearly flat in insurance (0.92-1.09, 1.2x). This gradient is consistent with the ELM-based prediction developed in Section 2.2: where seasonal demand variation is strongest, the compositional shift toward higher-involvement consumers is largest, and the modulation of advertising elasticities is most pronounced. When seasonal modulation is absent, M(t) converges toward unity and the model reduces to the Standard specification.

Table 2. Empirical results across three cases.

	Charity		**Telecom**		**Insurance**	
	Mod	Std	Mod	Std	Mod	Std
WMAPE train (%)	13.3	15.3	5.5	5.6	8.4	8.5
WMAPE test (%)	**18.1**	19.7	**5.1**	6.1	9.3	**9.0**
Modulator range	0.31-2.63	—	0.65-1.25	—	0.92-1.09	—
Modulator ratio	8.4x	—	1.9x	—	1.2x	—
Paid media share (%)	58.6	74.0	41.2	40.6	24.3	24.8

Attribution divergence. The most consequential differences are in advertising channel attributions. In the charity case, the Standard model attributes 74.0% of the outcome to paid media vs. 58.6% for the Modulated model (delta 15.4pp). Without a multiplicative modulator, the Standard model absorbs seasonal amplification into channel coefficients, inflating paid media contribution and underestimating CPA (overstating ROAS) for upper-funnel channels (TV, Social) by 23 to 37%. Such distortion can lead advertisers to over-invest in channels whose effectiveness is inflated by unmodeled seasonality. In the telecom case, CPA estimates diverge by up to 11% despite comparable aggregate fit. In insurance, differences remain below 7%, consistent with the nearly flat modulator.

Self-regulation. The gradient across the three cases confirms that the modeling framework adds complexity only where the data support it. This is independently corroborated by the Standard model's own additive seasonality, which exhibits a parallel gradient (range spanning 106% of mean outcome in charity, 28% in telecom, 3% in insurance), confirming that the three cases represent distinct levels of seasonal intensity.

Read alongside the diagnostic developed in Sections 1 and 2.1, this gradient is itself informative: the magnitude of the Modulated–Standard divergence tracks the magnitude of seasonal modulation in the data, as expected if the two specifications differ chiefly in how they accommodate the channel × season component that drives the omitted-variable concern. Where seasonality is substantial (charity), the Standard model absorbs it into static coefficients and overstates paid-media contribution; where seasonality is weak (insurance), there is little to absorb, and the two specifications converge.

6. Discussion and Conclusion

This paper proposed a Bayesian marketing mix model in which nonlinear saturation and time-varying advertising effectiveness coexist within an identified structure, grounded in the Elaboration Likelihood Model (Petty et al., 1983).

The paper offers three contributions. First, a theoretical contribution: we bridge the ELM, an individual-level theory of persuasion, with aggregate marketing mix modeling. The ELM predicts that higher involvement leads to deeper message processing and stronger persuasion. We derive an aggregate-level implication: seasonal demand peaks shift the composition of the consumer pool toward those with active category need and higher felt involvement (Celsi and Olson, 1988), increasing average advertising effectiveness. This micro-to-macro bridge translates a well-established consumer behavior mechanism into a testable, operationalizable prediction for aggregate models. The empirical gradient across our three cases, where the modulator ranges from 8.4x in a high-seasonality setting to near unity where seasonality is absent, is consistent with this prediction. Second, a methodological contribution: the seasonal modulator operationalizes the ELM-derived prediction through a parsimonious Fourier-based mechanism that produces time-varying ROAS estimates without requiring freely evolving coefficients. Because saturation and modulation operate on structurally distinct dimensions, they remain distinguishable without forcing the analyst to choose between them. The simulation confirmed that this separation reduces advertising channel contribution error by more than 50%. Third, a practical contribution for marketing analysts and media planners: the modulator indicates when advertising is more or less effective, enabling seasonal budget allocation based on time-varying rather than annual ROAS. Without such adjustment, standard models absorb seasonal amplification into static coefficients, overstating effectiveness at peaks and misallocating budgets. The modulator range serves as a diagnostic: when close to unity, seasonal modulation is not material, and the standard specification is sufficient.

The proposed framework contributes to the broader effort of addressing endogeneity in observational marketing data, which Rutz and Watson (2019) characterize as a concern for which "no 'perfect' solution exists", with progress coming through best practices adapted to the structure of each problem. The seasonal modulator is one such best practice for MMMs with seasonal demand: a parametric structural restriction, motivated by consumer-behavior theory, that mitigates the seasonal component of the omitted-variable bias most pervasive in this setting.

The framework has limitations consistent with this positioning. The modulator addresses only the periodic component of the omitted-variable concern; non-periodic confounders such as competitor actions, pricing changes, or macroeconomic shocks remain unmodeled, and reverse causality within a season is not separately identified. The framework is complementary to, not a substitute for, experimental measurement: where field experiments are feasible, they remain the gold standard for causal identification (Gordon et al., 2019). The shared modulator does not capture channel-specific seasonal patterns, an extension natural to pursue via hierarchical priors. Our evidence is based on three Italian datasets; broader generalization requires further investigation across markets and product categories.

References

Bruce, N. I. (2008). Pooling and dynamic forgetting effects in multitheme advertising: Tracking the advertising-sales relationship with particle filters. *Marketing Science*, *27*(4), 659–673. https://doi.org/10.1287/mksc.1070.0317

Celsi, R. L., & Olson, J. C. (1988). The role of involvement in attention and comprehension processes. *Journal of Consumer Research*, *15*(2), 210–224.

Dew, R., Padilla, N., & Shchetkina, A. (2024). Your MMM is broken: Identification of nonlinear and time-varying effects in marketing mix models. *SSRN*. https://doi.org/10.2139/ssrn.4928268

Gijsenberg, M. J. (2017). Riding the waves: Revealing the impact of intrayear category demand cycles on advertising and pricing effectiveness. *Journal of Marketing Research*, *54*(2), 171–186.

Gordon, B. R., Zettelmeyer, F., Bhargava, N., & Chapsky, D. (2019). A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook. *Marketing Science*, *38*(2), 193–225. https://doi.org/10.1287/mksc.2018.1135

Guhl, D., Baumgartner, B., Kneib, T., & Steiner, W. J. (2018). Estimating time-varying parameters in brand choice models: A semiparametric approach. *International Journal of Research in Marketing*, *35*(3), 394–414.

Hanssens, D. M., Parsons, L. J., & Schultz, R. L. (2001). *Market response models: Econometric and time series analysis* (2nd ed.). Kluwer Academic Publishers.

Jin, Y., Wang, Y., Sun, Y., Chan, D., & Koehler, J. (2017). *Bayesian methods for media mix modeling with carryover and shape effects* (Technical Report). Google Inc.

Korkames, J., Stanley, T. D., & Stremersch, S. (2025). Meta-analysis of advertising effectiveness: New insights from improved bias corrections. *International Journal of Research in Marketing*. Advance online publication.

Little, J. D. C. (1970). Models and managers: The concept of a decision calculus. *Management Science*, *16*(8), B466–B485.

Ng, E., Wang, Z., & Dai, A. (2022). Bayesian time varying coefficient model with applications to marketing mix modeling. In *Proceedings of the KDD ADKDD Workshop*.

Petty, R. E., Cacioppo, J. T., & Schumann, D. (1983). Central and peripheral routes to advertising effectiveness: The moderating role of involvement. *Journal of Consumer Research*, *10*(2), 135–146.

Rutz, O. J., & Watson, G. F., IV. (2019). Endogeneity and marketing strategy research: An overview. *Journal of the Academy of Marketing Science*, *47*(3), 479–498. https://doi.org/10.1007/s11747-019-00630-4

Winer, R. S. (1979). An analysis of the time-varying effects of advertising: The case of Lydia Pinkham. *The Journal of Business*, *52*(4), 563–576.