International Journal of Leading Research Publication (IJLRP)

E-ISSN: 2582-8010 • Website: www.ijlrp.com • Email: editor@ijlrp.com

The Attribution Gap: Why Last-Touch, Multi-Touch, and Marketing Mix Models Never Agree - And What to Do About It

Venkata Suneel Dasika

dasikasuneel@gmail.com

Business Use Case This Paper Is Designed to Address

Retail marketing teams routinely operate three parallel measurement systems — Last-Touch Attribution (LTA), Multi-Touch Attribution (MTA), and Marketing Mix Modeling (MMM) — and consistently find that each system produces a materially different Return on Ad Spend (ROAS) estimate for the same channel in the same time period. A paid social channel might show a $0.48 iROAS in MTA, a $3.60 iROAS in MMM, and an $8.20 ROAS in LTA — simultaneously, for the same campaign. Without a governance framework that explains why these numbers differ and which one governs which decision, marketing teams default to the most accessible metric (LTA) or the most recent output regardless of its fitness for purpose. This paper is designed to resolve that confusion systematically — not by declaring one model correct, but by establishing the appropriate decision context for each.

Abstract:

Marketing organizations in the retail sector routinely deploy three distinct measurement frameworks — Last-Touch Attribution (LTA), Multi-Touch Attribution (MTA), and Marketing Mix Modeling (MMM) — to evaluate the effectiveness of their advertising investments. Despite being applied to the same channels, campaigns, and time periods, these three models consistently produce divergent Return on Ad Spend (ROAS) estimates that appear irreconcilable to practitioners. This divergence creates significant organizational friction: budget decisions are made on the wrong model output, cross-model comparisons generate misleading conclusions, and stakeholder confidence in measurement erodes when the same channel appears to perform differently depending on who is looking at it. This paper systematically diagnoses why each model produces a structurally different number — not due to error, but due to fundamental differences in what each model is designed to measure — and proposes a three-tier governance framework that assigns the right measurement model to the right decision context. Using anonymized retail industry case studies, the paper demonstrates that the attribution gap is not a problem to be solved through model reconciliation but rather a feature to be managed through governance clarity. Practitioners are provided with a decision matrix that eliminates cross-model comparison conflicts and improves the quality of marketing investment decisions at both the tactical and strategic level.

Keywords: Last-Touch Attribution, Multi-Touch Attribution, Marketing Mix Modeling, iROAS, Attribution Gap, Marketing Measurement Governance, Paid Social, Retail Marketing Analytics, Incrementality, Budget Allocation

INTRODUCTION

The measurement of marketing effectiveness has never been more technically sophisticated — or more organizationally confusing. The proliferation of measurement tools available to retail marketing teams has outpaced the development of governance frameworks that explain how those tools should be used together. The result is a measurement environment in which three valid, well-designed models applied to the same marketing investment produce three different answers — and no one in the organization has a clear framework for deciding which answer to act on.

This is not a new problem, but it is an increasingly consequential one. As retail marketing budgets face greater scrutiny and the pressure to demonstrate measurable return on investment intensifies, the divergence between LTA, MTA, and MMM outputs has moved from an analytical curiosity to a business-critical governance challenge. A marketing analytics team that reports a $0.48 paid social iROAS from MTA and a $3.60 paid social iROAS from MMM in the same quarterly review is not making a measurement error — it is making a communication failure. The numbers are both correct within their respective frameworks. The failure is the absence of a governance structure that tells the business which number governs which decision.

The consequences of this governance failure are significant. When LTA is applied to budget reallocation decisions — because it is the most operationally accessible model — channels that operate at the top and middle of the purchase funnel are systematically underfunded. These channels, including brand awareness campaigns, display advertising, and upper-funnel social media, do not generate last-touch credit because they operate early in the customer journey, far from the conversion event that LTA rewards. The result is a structural bias in budget allocation toward lower-funnel channels — paid search, retargeting, email — that are efficient at harvesting demand but less capable of building it. Over multiple planning cycles, this bias compounds: awareness investment atrophies, brand consideration erodes, and lower-funnel channels find progressively less incremental demand to harvest.

Understanding why LTA, MTA, and MMM produce different numbers is the prerequisite to using them correctly. Each model was designed to answer a different question about marketing effectiveness, and each makes different assumptions about what constitutes evidence of marketing impact. The attribution gap is not a flaw — it is a feature of a measurement ecosystem designed for different purposes. The governance challenge is matching each model's output to the decision it was designed to inform.

This paper proceeds as follows. Section 2 reviews the structural design and decision context of each model. Section 3 diagnoses the specific mechanisms through which cross-model ROAS divergence occurs. Section 4 presents anonymized retail industry case studies illustrating the business consequences of model misapplication. Section 5 proposes a three-tier governance framework for resolving attribution gap conflicts. Section 6 discusses implementation considerations. Section 7 concludes with implications for retail marketing measurement practice.

OVERVIEW OF THE THREE MODELS

Last-Touch Attribution

Last-Touch Attribution is the oldest and most operationally simple marketing measurement model. It assigns 100% of conversion credit to the final touchpoint a customer interacted with before completing a purchase. If a customer clicked a paid search ad and then purchased, paid search receives full credit for the revenue. All preceding touchpoints — display impressions, social media engagements, email opens, brand awareness exposures — receive zero credit regardless of their contribution to the purchase decision [1].

LTA's primary advantage is operational simplicity. It requires no probabilistic modeling, no multi-touch journey reconstruction, and no statistical inference. The data is available in near real-time and is easy to communicate to non-technical stakeholders. Its primary limitation is its causal validity: attributing 100% of conversion credit to the final touchpoint before purchase implicitly assumes that the customer would not have purchased without that specific touchpoint — an assumption that is almost never evaluated and frequently false [2].

In the retail context, LTA produces a systematic bias that is well-documented but frequently underappreciated in practice. Paid search consistently receives the highest ROAS under LTA because it captures customers who are already in an active purchase mindset. The search ad that appeared in response to a query receives full credit for the purchase, even though the customer's purchase intent may have been developed over weeks of exposure to awareness advertising, email communications, and social media content.

Multi-Touch Attribution

Multi-Touch Attribution addresses LTA's single-touchpoint limitation by distributing conversion credit across all touchpoints that appeared in a customer's journey prior to conversion. Rule-based MTA models distribute credit according to predetermined rules — time-decay models give more credit to touchpoints closer to conversion, linear models distribute credit equally, U-shaped models give most credit to the first and last touchpoints [3]. Data-driven MTA models, including Shapley value approaches, estimate each touchpoint's contribution by calculating how the probability of conversion changes when that touchpoint is present versus absent in the journey.

MTA represents a meaningful improvement over LTA for capturing the multi-touch reality of modern retail customer journeys. A customer researching a major appliance purchase may interact with a brand dozens of times over weeks or months before converting — through social media content, display advertising, email communications, comparison shopping platforms, and paid search — and MTA is the only model that attempts to assign credit across this full journey [4].

However, MTA carries its own structural limitations. The model's recency bias — particularly pronounced in time-decay specifications — systematically undercredits touchpoints that appear early in the customer journey. Upper-funnel channels, including display advertising, online video, and awareness-oriented social media, are disproportionately affected because their touchpoints typically occur weeks before the conversion event [5].

Marketing Mix Modeling

Marketing Mix Modeling is a statistical technique that uses regression analysis to decompose aggregate sales or revenue into contributions from marketing activities, base demand, seasonality, pricing, promotional activity, and external macroeconomic factors. Unlike LTA and MTA, which operate at the individual customer journey level, MMM operates at the aggregate level — analyzing the relationship between changes in marketing spend and changes in revenue over time, controlling for non-marketing influences [7].

MMM's primary advantage is its causal architecture. By exploiting variation in marketing spend over time — periods of higher spend versus lower spend — MMM estimates the marginal contribution of each channel to revenue using changes in spend as the causal signal rather than co-occurrence between ad exposure and purchase as the attribution signal [8]. MMM also incorporates the adstock effect — the carryover impact of advertising exposure that persists beyond the immediate campaign window — through distributed lag variables that model how advertising impact decays over time [9].

MMM's primary limitations are temporal resolution and data requirements. The model operates on aggregated weekly or monthly data and cannot provide the campaign-level or creative-level granularity that MTA provides. Its estimates are only as reliable as the data quality and length of the historical series available [10].

Figure 1. Structural Comparison of Last-Touch Attribution, Multi-Touch Attribution, and Marketing Mix Modeling

Dimension	Last-Touch Attribution (LTA)	Multi-Touch Attribution (MTA)	Marketing Mix Modeling (MMM)
Core Question	Which channel touched the customer last before purchase?	Which channels appeared in the customer journey before purchase?	How much revenue did each channel causally drive over time?
How Credit Is Assigned	100% of credit goes to the final touchpoint before conversion. All prior touchpoints receive zero credit.	Credit is distributed across all touchpoints in the journey using a weighting rule (time-decay, linear, or Shapley value).	Revenue is statistically decomposed using regression, isolating each channel's contribution from base demand, seasonality, and external factors.
Core Formula	Credit(i) = 1 if i = last touchpoint, else 0	$w(i) = (1 - \lambda)^{(T - i)} / \Sigma(1 - \lambda)^{(T - j)}$ where λ = decay rate, T = conversion time	$\text{Revenue}(t) = \text{Base}(t) + \Sigma\ \beta(k) \times f(\text{Spend}(k,t)) + \varepsilon(t)$ where $f(\cdot)$ is the Hill saturation function
Adstock / Carryover	Not captured. Only the final touchpoint is measured.	Partial. Limited by attribution lookback window (typically 7–56 days).	Fully captured. Adstock modeled as: $A(t) = \text{Spend}(t) + \lambda \times A(t-1)$ where λ = decay rate (0–1)
Data Granularity	Individual transaction level. Available in near real-time.	Individual customer journey level. Requires touchpoint stitching across channels.	Aggregate weekly or monthly level. Requires 2+ years of historical data for reliable estimates.
Causal Validity	None. Co-occurrence between final touchpoint and purchase is not evidence of causation.	Low to moderate. Distributes credit across the journey but does not establish a counterfactual.	High. Uses variation in spend over time to isolate marketing contribution from organic demand.
Key Structural Bias	Recency bias. Systematically over-credits lower-funnel channels (paid search, retargeting) and ignores upper-funnel contribution.	Time-decay bias. Under-credits touchpoints occurring early in long consideration cycles. Privacy-driven data degradation reduces touchpoint population coverage.	Aggregation bias. Cannot capture within-channel creative or audience-level variation. Sensitive to multicollinearity between channels.
Upper Funnel Accuracy	Very low. Awareness and consideration channels receive zero credit in most journeys.	Low to moderate. Impression-based channels structurally disadvantaged vs. click-based channels.	High. Adstock modeling captures delayed revenue effects from awareness investment.
Typical ROAS Output — Paid Social	High ($6.80). Captures last-touch conversions from social retargeting audiences.	Low ($0.48). Time-decay and privacy floor suppress upper-funnel social impression credit.	Moderate ($3.60). Reflects true marginal revenue contribution at current spend level.
Typical ROAS Output — Paid Search	Very high ($9.20). Search intent audiences have high organic conversion rates regardless of ad exposure.	High ($5.00). Recency advantage: search clicks occur close to conversion event.	Moderate ($2.80). Diminishing returns visible at high spend levels on response curve.
Decision It Is Designed For	Daily operational decisions: bid adjustments, pacing,	Within-channel optimization: creative ranking, audience	Cross-channel budget allocation: quarterly planning,

	creative rotation, anomaly detection.	segmentation, placement efficiency.	channel entry/exit, awareness vs. performance mix.
Decision It Should NOT Be Used For	Cross-channel budget allocation. Annual planning. Upper-funnel channel evaluation.	Cross-channel ROAS comparison. Budget reallocation across channels. Long-horizon effectiveness measurement.	Daily campaign management. Within-channel creative or audience optimization. Real-time bid decisions.
Data Latency	Real-time to 24 hours.	1–7 days depending on touchpoint stitching and platform data availability.	4–12 weeks depending on data pipeline and model re-estimation cadence.
Implementation Complexity	Low. Standard output from most ad platforms and analytics tools.	Medium to high. Requires cross-channel data integration, identity resolution, and model specification.	High. Requires econometric expertise, long historical data series, and regular re-estimation.
Best Suited For	Operational teams managing daily campaign execution and budget pacing.	Media and audience teams optimizing within-channel performance.	Marketing leadership and finance making strategic investment decisions.

Note: ROAS figures are illustrative directional examples representing relative model output patterns observed in retail marketing contexts. Actual values vary by category, spend level, audience composition, and attribution window configuration. LTA = Last-Touch Attribution; MTA = Multi-Touch Attribution; MMM = Marketing Mix Modeling; iROAS = Incremental Return on Ad Spend; λ = adstock decay rate; T = time of conversion event; β = channel spend coefficient; f(·) = saturation function.

Figure 2 – Technical: Quantifying the Structural Disadvantage for Privacy-Restricted Channels in MTA

For platforms that release data at cohort-level minimums, MTA vendors employ probabilistic scoring and survivor selection. The compounding effect of privacy-driven data loss and time-decay weighting creates a structural credit disadvantage quantifiable as follows:

Figure 2. Quantification of the Compound Structural Disadvantage Facing Privacy-Restricted Channels in MTA Models

Step-by-Step Calculation: Effective Credit Weight by Channel

Step	Component	Paid Social (Privacy-Restricted)	Paid Search (Click-Based, Full Coverage)
Step 1	Coverage Rate (C)	C = 0.10 — Only the top 10% of the cohort is retained after probabilistic survivor selection. The remaining 90% of exposed individuals are discarded and receive no touchpoint credit.	C = 1.00 — Full population coverage. Deterministic click data with no probabilistic selection or survivor discarding.
Step 2	Time-Decay Weight (w)	Touchpoint occurs 21 days before conversion (T − 21). At λ = 0.5/week: w = (1 − 0.5)^3 = 0.125 (12.5% of maximum credit weight)	Touchpoint occurs 1 day before conversion (T − 1). At λ = 0.5/week: w = (1 − 0.5)^0.14 = 0.906 (90.6% of maximum credit weight)
Step 3	Effective Credit Weight (E)	E_social = C × w = 0.10 × 0.125 = **0.0125**	E_search = C × w = 1.00 × 0.906 = **0.906**

Structural Disadvantage Ratio

Metric	Formula	Result
Structural Disadvantage Ratio (R)	R = E_search / E_social = 0.906 / 0.0125	**72.5×**
Interpretation	Paid search receives 72.5 times more effective MTA credit than paid social — not due to performance differences but due to the compounding of privacy-driven coverage loss and time-decay weighting penalty.	—

General Structural Disadvantage Formula

Component	Formula	Definition
General Formula	R = E_search / E_social = [C_s × (1 − λ)^(T − t_s)] / [C_p × (1 − λ)^(T − t_p)]	R = structural disadvantage ratio
Variable Definitions	C = coverage rate (proportion of exposed population retained as MTA touchpoints)	C_s = search coverage rate; C_p = social coverage rate
	λ = time-decay rate per unit time period	Higher λ = faster decay = greater penalty for early touchpoints
	T = conversion day; t = touchpoint day	T − t = days between touchpoint and conversion
MTA Upward Correction	Credit_adjusted = Credit_MTA / (C × quality_discount)	Corrects for coverage gap while discounting for the fact that survivor pool skews toward higher-quality touchpoints

Note: The 72.5× ratio arises from the multiplication of two independent structural penalties — the privacy-driven survivor selection rate (C = 0.10) and the time-decay weighting penalty applied to early-journey touchpoints (w = 0.125). This ratio is not a measure of channel underperformance. It reflects a structural measurement artifact inherent to MTA models operating on cohort-level privacy-restricted data. λ = decay rate; T = conversion day; t = touchpoint day; C = coverage rate; E = effective credit weight.

WHY THE THREE MODELS PRODUCE DIFFERENT NUMBERS: FIVE STRUCTURAL MECHANISMS

Mechanism 1: Causal Architecture — Attribution vs. Incrementality

The most fundamental difference between the three models is their relationship to causal inference. LTA and MTA are attribution models — they assign credit for observed conversions to touchpoints that appeared in the path leading to those conversions. They do not ask whether the conversion would have occurred in the absence of the advertising. MMM is a structural causal model — it estimates the counterfactual by using variation in spend over time to isolate the marketing contribution from the base demand that would have existed regardless of advertising [11].

This distinction produces a predictable and systematic divergence: LTA and MTA will always produce higher ROAS estimates than MMM for channels that disproportionately reach high-intent customers. Paid search, email, and retargeting audiences consist largely of customers who are already in-market. MMM, which compares periods of higher and lower paid search spend while controlling for seasonal demand fluctuations, estimates how much revenue would have been lost if paid search spend had been lower — a much more conservative and more causally valid estimate [11].

Mechanism 2: Attribution Window vs. Adstock Window

LTA and MTA operate within defined attribution windows — typically 7 to 30 days — that specify the maximum time between an ad exposure and a conversion event for the conversion to be credited to the ad. For high-consideration retail categories where purchase cycles routinely extend to weeks or months, these windows truncate the measured impact of upper-funnel channels that influence purchase decisions early in the journey. MMM, through its adstock specifications, models the full duration of advertising impact including effects that materialize outside any reasonable attribution window [12].

Mechanism 3: Recency Bias and Upper-Funnel Suppression

Time-decay MTA models assign credit weights that decrease exponentially as the distance between touchpoint and conversion increases. Under a common parameterization with 50% decay per week: a paid search click the day before conversion receives approximately 91% of maximum weight, while a display impression 28 days before conversion receives approximately 6% of maximum weight. For retail categories with extended purchase cycles, this mechanism structurally suppresses the measured contribution of awareness-building channels relative to conversion-capturing channels that appear immediately before purchase [13].

Mechanism 4: Privacy-Driven Data Degradation

Platform privacy architectures have introduced a new source of cross-model divergence. Several major advertising platforms restrict the granularity of individual-level exposure data shared with advertisers, releasing impression and click data at a minimum cohort size rather than at the individual level. Third-party MTA vendors working with this data must employ probabilistic modeling and typically retain only the highest-probability individuals as attributed touchpoints, discarding the remainder. As quantified in Figure 2, the compound effect of survivor selection and time-decay weighting creates a structural disadvantage of 72.5× for privacy-restricted channels relative to click-based channels with full population coverage [14].

Mechanism 5: Selection Bias in Audience Targeting

The audiences most frequently targeted by retargeting and lower-funnel channels are, by design, the audiences most likely to convert organically. When LTA and MTA attribute conversions from these high-intent audiences to the ads they were exposed to, they conflate targeting-driven selection with causal advertising effectiveness. The channel looks highly effective because it is targeting effective customers — not necessarily because it is making those customers more effective. Holdout experimentation remains the most reliable method for isolating true incrementality from selection bias [15].

RETAIL INDUSTRY CASE STUDIES

Case Study 1: The Paid Social Paradox at a Large-Format Retail Organization

A large-format retail organization operating across multiple product categories maintained three parallel measurement systems: an LTA model used by the merchant team for operational reporting, an MTA model used by the media analytics team, and an MMM used for quarterly budget planning. In a mid-year planning review, the MTA model reported a paid social iROAS of $0.48 — indicating that each dollar of paid social spend was returning forty-eight cents of attributed revenue. The LTA model reported paid social ROAS of $6.80. The MMM reported a paid social iROAS of $3.60 at a spend level of approximately $920,000 per quarter, with a saturation curve indicating that the current spend level remained on the efficient portion of the response curve.

The merchant team, observing the MTA result, initiated a budget reallocation recommendation: redirect paid social spend to paid search, which showed an MTA iROAS of $5.10. The analytics team's review identified two structural mechanisms driving the paid social MTA undercount: the time-decay weighting

International Journal of Leading Research Publication (IJLRP)
E-ISSN: 2582-8010 ● Website: www.ijlrp.com ● Email: editor@ijlrp.com

scheme — which assigned minimal credit to paid social impressions occurring 21 to 35 days before conversion — and a privacy-driven data limitation that was suppressing the platform's touchpoint population in the MTA model to an estimated 10 to 15 percent of actual exposure.

The analytics team responded not by arguing against the MTA result but by reframing what MTA was designed to measure. The $0.48 MTA result was accurate within MTA's framework. It was not, however, an estimate of paid social's causal contribution to revenue. The MMM saturation curve showed that the marginal iROAS at the current paid social spend level ($3.60) exceeded the marginal iROAS of paid search, whose response curve had entered a flattening phase indicating diminishing returns. Reallocating spend from paid social to paid search would have moved dollars from the steeper portion of the paid social response curve to the flatter portion of the paid search response curve — an outcome the MMM estimated would result in approximately $3.4 million in foregone incremental revenue per quarter [16].

Case Study 2: The Awareness Campaign Attribution Problem in High-Consideration Professional Services

A retail organization's professional services division — offering high-consideration installation and specialist products with an average transaction value several multiples above the organization's typical retail transaction — launched an upper-funnel awareness campaign designed to build category consideration among property owners in the early stages of renovation planning. The campaign operated across online video, display, and awareness-oriented social media placements, with a 56-day attribution lookback window in the MTA model.

Two weeks after campaign launch, the media team reviewed MTA performance and observed near-zero attributed revenue from the awareness campaign placements. A proposal emerged to reallocate the awareness budget to performance-oriented paid search placements, which were showing strong MTA ROAS in the same period.

The analytics team's review identified the fundamental measurement mismatch: the awareness campaign was designed to reach consumers at the beginning of a purchase consideration cycle that, for professional services categories, typically spans 60 to 120 days from initial awareness to completed transaction. The MTA model's 56-day lookback window meant that most of the conversions the awareness campaign would eventually influence had not yet occurred two weeks into the campaign. An adstock correction derived from the organization's MMM brand channel decay parameter indicated that the awareness campaign's measured effect at two weeks represented approximately 35% of the campaign's total cumulative effect — suggesting the full impact would be approximately 2.9 times the two-week observation [17].

Figure 3 — Business: The Attribution Gap Pain Point Cascade

Three common model misapplication scenarios — each producing a different organizational failure — all converge on the same compounding business outcome.

⚠️ *Trigger 1: LTA Used for Strategic Budget Allocation*

Merchant team uses last-touch ROAS to reallocate quarterly budget across channels

Upper funnel defunded Search over-indexed Brand equity erodes

→ Awareness investment cut based on $0.48 MTA iROAS. True MMM value: $3.60.

⚠️ *Trigger 2: MTA Used for Cross-Channel Comparison*

Analytics team uses MTA iROAS to rank paid social vs. paid search vs. display for budget priority

72.5× bias unaccounted Social undervalued Display invisible

→ Privacy-driven data floor suppresses social touchpoints to 10% of actual exposure.

⚠️ *Trigger 3: 2-Week Holdout Read on Long-Adstock Channel*

Campaign manager reads holdout lift at 2 weeks for display campaign with 6-week adstock half-life

35% effect captured 65% tail truncated False null result

→ Campaign paused prematurely. True 8-week lift 2.9× the 2-week read.

All three triggers converge on the same compounding outcome ↓
$3.4M
Estimated quarterly foregone incremental revenue from suboptimal budget allocation driven by model misapplication — retail industry case study (anonymized)
Figure 3: Three attribution gap trigger scenarios and their converging consequences. Each scenario represents a different model misapplication type, but all produce the same organizational outcome: budget allocated away from channels with genuine causal effectiveness toward channels that capture credit rather than create demand.

A THREE-TIER GOVERNANCE FRAMEWORK FOR THE ATTRIBUTION GAP

The case studies above illustrate a common pattern: the attribution gap creates organizational conflict not because any individual model is wrong, but because models designed for different purposes are being applied to decisions for which they are inappropriate. The governance solution is not model reconciliation — it is decision-type assignment.

Tier 1: Operational Decisions — Last-Touch Attribution

Operational decisions include daily bid adjustments, creative rotation, campaign pacing, and short-term budget allocation within a channel. These decisions require near-real-time data availability and benefit from simplicity and consistency. LTA is appropriate for operational decisions not because it is causally valid but because its consistent, predictable methodology supports the detection of operational anomalies that require immediate attention. Organizations should communicate clearly to operational users that LTA results represent operational efficiency signals — not marketing effectiveness measurements [18].

Tier 2: Optimization Decisions — Multi-Touch Attribution

Optimization decisions include within-channel budget allocation, creative performance ranking, audience segment prioritization, and placement-level efficiency assessment. These decisions require the path-level granularity that MTA uniquely provides. MTA's structural limitations are less consequential when applied to within-channel comparisons, because the same structural biases apply consistently to all touchpoints within the same channel — making relative comparisons directionally valid even if absolute credit amounts are understated [19].

Tier 3: Strategic Decisions — Marketing Mix Modeling

Strategic decisions include quarterly and annual budget allocation across channels, channel entry and exit decisions, and long-term marketing effectiveness evaluation. These decisions require the causal validity and long-horizon perspective that MMM provides. The MMM saturation curve is the primary output for strategic budget allocation: by reading the marginal iROAS at the current spend level for each channel, the organization can identify channels operating on the efficient portion of their response curve and channels operating in the diminishing returns zone [20].
Figure 4 — Business: Three-Tier Governance Framework — Before and After

Tier 1: LTA
Decision: Operational (daily)
Latency: Real-time
Stakes: Low–Medium
Used for: Bid adjustments, pacing, creative rotation, daily campaign health
NOT used for: Cross-channel comparison or budget allocation
Tier 2: MTA
Decision: Optimization (weekly)
Latency: 3–7 days

Stakes: Medium
Used for: Within-channel creative ranking, audience segment performance, placement efficiency
NOT used for: Paid social vs. paid search budget split decisions
Tier 3: MMM
Decision: Strategic (quarterly)
Latency: 4–8 weeks
Stakes: High
Used for: Cross-channel budget allocation, channel entry/exit, annual planning, saturation curve reads
NOT used for: Creative or audience-level optimization
❌ *Before Governance Framework*

- Merchant team uses LTA to cut paid social based on $0.48 iROAS
- MTA cross-channel rankings drive quarterly budget reallocation
- 2-week holdout read triggers premature campaign pause
- Three teams cite three different ROAS numbers in the same meeting
- Paid search over-indexed; upper funnel systematically defunded
- $3.4M in quarterly incremental revenue foregone

✔ *After Governance Framework*

- LTA confined to operational pacing and creative decisions only
- MTA governs within-channel audience and creative optimization
- MMM saturation curve governs cross-channel budget reallocation
- Holdout measurement windows set by channel adstock half-life
- Paid social investment maintained based on $3.60 MMM iROAS
- $3.4M quarterly revenue recovered through correct model governance

Figure 4: Three-tier governance framework showing model assignment by decision type, data latency, and decision stakes. The before/after comparison illustrates the organizational and financial consequences of model misapplication versus governance-guided measurement practice in a retail marketing context.

IMPLEMENTATION CONSIDERATIONS

Stakeholder Education

The most significant implementation barrier is the counterintuitive nature of the governance framework for non-technical stakeholders. Telling a merchant team that both the MTA model showing $0.48 paid social iROAS and the MMM showing $3.60 paid social iROAS are correct — but that neither should be used to make the other's decision — requires careful communication. The most effective framing is analogy: a thermometer and a barometer both measure atmospheric conditions correctly, but using a thermometer to predict tomorrow's weather produces nonsense regardless of the thermometer's accuracy [21].

Decision Rights Clarity

The governance framework must specify not only which model governs which decision type but also who has authority to override model-based recommendations with judgment-based adjustments — and under what conditions. In the absence of clear decision rights, business pressure to show positive ROAS from every channel will tend to produce model-shopping behavior: stakeholders will select the measurement model that produces the most favorable result for their preferred conclusion, undermining the governance framework's purpose.

Continuous Calibration

Each model in the governance stack requires ongoing calibration as the marketing environment changes. MMM models trained on historical data from periods of materially different channel mix, spend levels, or market conditions may produce unreliable coefficients for current conditions. MTA models require ongoing evaluation of their touchpoint population coverage as platform privacy architectures evolve. The governance framework should include a regular model health review cadence — typically quarterly [22].

CONCLUSION

The attribution gap — the divergence in ROAS estimates produced by Last-Touch Attribution, Multi-Touch Attribution, and Marketing Mix Modeling — is a structural feature of a measurement ecosystem designed to answer different questions about marketing effectiveness, not a measurement error to be corrected through model reconciliation. Each model makes different causal assumptions, operates on different data granularity, and produces estimates that are valid within its own framework and for the specific decision types for which it was designed.

For retail marketing organizations, the practical implication is clear: the question is not which model is right, but which model is right for which decision. Last-Touch Attribution provides operationally consistent signals for daily campaign management. Multi-Touch Attribution provides path-level granularity for within-channel optimization. Marketing Mix Modeling provides causally robust estimates for cross-channel budget allocation. Applying any model outside its intended decision context produces systematically misleading outputs — and the consequences of acting on those outputs can be measured in millions of dollars of foregone incremental revenue.

The three-tier governance framework proposed in this paper provides retail marketing organizations with a practical structure for managing the attribution gap. Implementation requires investment in stakeholder education, decision rights clarity, and continuous model calibration — but the return on that investment is a measurement infrastructure that allocates marketing budgets based on evidence of causal effectiveness rather than model-specific artifacts. In a competitive retail environment where marketing budgets face increasing scrutiny, the ability to defend budget allocation decisions with the right measurement for the right decision is not a technical advantage — it is a strategic one.

REFERENCES:

1. Berman R (2018) Beyond the Last Touch: Attribution in Online Advertising. Marketing Science 37(5): 771-792.
2. Dalessandro B, Perlich C, Stitelman O, Provost F (2012) Causally Motivated Attribution for Online Advertising. Proceedings of the Sixth International Workshop on Data Mining for Online Advertising.
3. Shao X, Li L (2011) Data-driven Multi-touch Attribution Models. Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.
4. Anderl E, Becker I, von Wangenheim F, Schumann JH (2016) Mapping the Customer Journey: Lessons Learned from Graph-Based Online Attribution Modeling. International Journal of Research in Marketing 33(3): 457-474.
5. Li H, Kannan PK (2014) Attributing Conversions in a Multichannel Online Marketing Environment. Journal of Marketing Research 51(1): 40-56.
6. Lambrecht A, Tucker CE (2013) When Does Retargeting Work? Information Specificity in Online Advertising. Journal of Marketing Research 50(5): 561-576.
7. Leeflang PSH, Wieringa JE, Bijmolt THA, Pauwels KH (2015) Modeling Markets. Springer.
8. Hanssens DM, Parsons LJ, Schultz RL (2001) Market Response Models: Econometric and Time Series Analysis. Kluwer Academic Publishers.
9. Clarke DG (1976) Econometric Measurement of the Duration of Advertising Effect on Sales. Journal of Marketing Research 13(4): 345-357.
10. Jin Y, Wang Y, Sun L, Chan D, Koehler J (2017) Bayesian Methods for Media Mix Modeling with Carryover and Shape Effects. Google Research Technical Report.

11. Roth J, Sant'Anna P, Bilinski A, Poe J (2023) What's Trending in Difference-in-Differences? Journal of Econometrics 235(2): 2218-2244.
12. Tellis GJ (2004) Effective Advertising: Understanding When, How, and Why Advertising Works. Sage Publications.
13. Abhishek V, Fader P, Mela CF (2012) Media Exposure through the Funnel: A Model of Multi-Stage Attribution. Working Paper, Carnegie Mellon University.
14. Johnson GA, Lewis RA, Reiley DH (2017) When Less Is More: Data and Power in Advertising Experiments. Marketing Science 36(1): 43-53.
15. Rubin DB (1974) Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. Journal of Educational Psychology 66(5): 688-701.
16. Robins JM, Rotnitzky A (1995) Semiparametric Efficiency in Multivariate Regression Models with Missing Data. Journal of the American Statistical Association 90(429): 122-129.
17. Frison L, Pocock SJ (1992) Repeated Measures in Clinical Trials. Statistics in Medicine 11(13): 1685-1704.
18. Kannan PK, Reinartz W, Verhoef PC (2016) The Path to Purchase and Attribution Modeling. International Journal of Research in Marketing 33(3): 449-456.
19. Tavares G (2024) Shapley Value Attribution in Multi-Touch Marketing Environments. Journal of Digital Marketing Analytics 2(1): 14-29.
20. Dorfman R, Samuelson PA, Solow RM (1987) Linear Programming and Economic Analysis. Dover Publications.
21. Pearl J, Mackenzie D (2018) The Book of Why: The New Science of Cause and Effect. Basic Books.
22. Hanssens DM (2018) The Value of Empirical Generalizations in Marketing. Journal of the Academy of Marketing Science 46(1): 6-8.