SOWMYA MONI ADVISORY

ENTERPRISE GTM INSIGHTS

Causal AI and Marketing Investment Decisions

What evidence should guide your next investment?



Sowmya Moni

Enterprise GTM and Decision Intelligence Advisor

October 2026

Where should the next investment go?

A marketing budget review can explain where the money went and still leave the most important question unresolved: *which investment deserves the next commitment?*

Paid search reports conversions. An executive program appears in large opportunities. Positioning work, analyst engagement and partner enablement may influence the same accounts without receiving identifiable channel credit. The CMO must decide what to expand, what to protect and what to test.

For enterprise B2B leaders, **three distinctions matter: contribution versus causation, historical return versus the return on additional spending, and measurable channel activity versus the wider work of creating demand**. A few large deals can dominate the data without resolving those distinctions.

My view is that causal AI can improve these investment choices when the evidence supports the decision. Its role is to challenge assumptions and estimate the effects of change. Leadership must still judge commercial value, strategic relevance and the cost of learning.

7.8%	15.3%	56%
Average marketing budget as a share of company revenue	Average share of marketing budgets allocated to AI	Report insufficient budget to deliver their strategy

Gartner 2026 CMO Spend Survey [1]. 401 marketing leaders; mostly businesses above $1 billion revenue in North America, the UK and Europe. Self-reported, cross-industry results; not a benchmark for all global CMOs or B2B SaaS firms.

The leadership question is which spending change can improve a defined commercial outcome, over a realistic period, and how confidently that improvement can be estimated.

Did the marketing activity make a difference?

An account attends your event and later becomes a customer. Did the event help win the business, or was the account already likely to buy? That is the question causal analysis tries to answer.

Three approaches answer different questions:

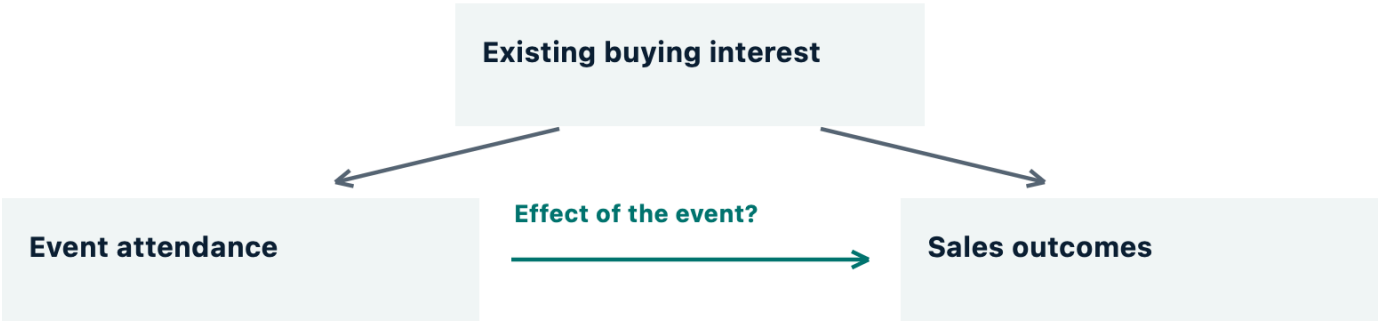
Prediction: Which accounts are likely to buy? A lead score can help prioritize them.

Attribution: Which activity gets credit for a sale? A reporting rule might credit the event because the customer attended.

Causation: Did the activity create additional business? This asks what would probably have happened without it, using a credible comparison.

Causal AI uses statistical methods and machine learning to help estimate that difference, but it still needs suitable data and a sound comparison.

FIGURE 1 WHY ATTENDANCE ALONE DOES NOT PROVE IMPACT



Existing interest can influence both attendance and sales. Compare similar account groups to estimate what the event adds. Conceptual illustration; see [5].

Some accounts may attend because they already plan to buy. Comparing attendees with nonattendees can therefore exaggerate the event's impact. A well-designed experiment offers a stronger comparison.

The 2019 and 2023 Marketing Science studies compared non-experimental estimates with results from Facebook advertising experiments. The methods tested often missed the effects measured in those experiments. That is why a credible comparison matters. Section 06 considers how more recent research builds on experimental evidence. These findings do not establish what works for enterprise B2B events. [2, 3]

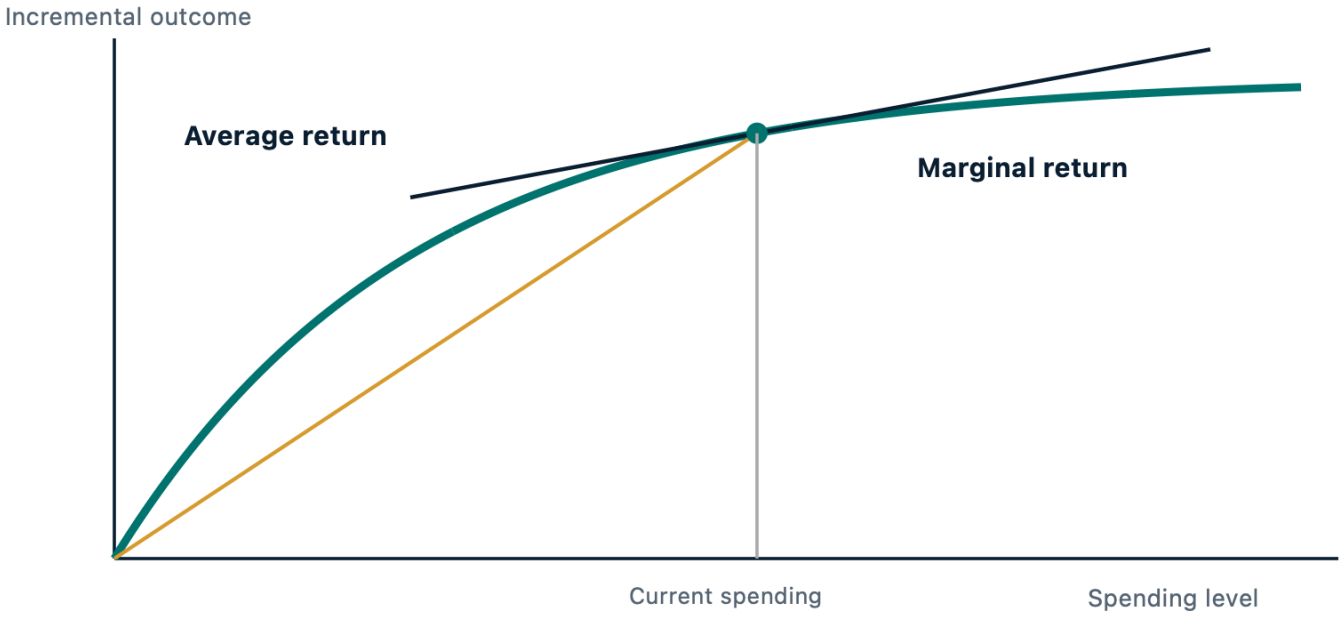
Before expanding an activity, ask how much business it adds, and not just how much business it appears alongside.

Does your best channel deserve more?

Marketing mix modeling relates aggregate outcomes to marketing and other factors over time. Google's open-source Meridian, generally available since January 2025, supports experimental calibration. Its documentation warns that predictive fit alone does not establish causal validity. [4, 5]

A channel can have a strong historical return and still be a poor destination for additional budget. Average return describes past spending; marginal return concerns the effect of a small increase. [8]

FIGURE 2 AVERAGE RETURN AND THE NEXT INCREMENT



Illustrative shape only; no company data or numerical estimates. Average return is the slope from the origin; marginal return is the tangent slope at current spending. See [8].

A material reallocation requires comparison across the proposed spending range, using comparable costs and horizons. Assess expected incremental value against what the business would forgo elsewhere. Beyond observed spending, the response curve depends on extrapolation. [8]

Invest in the next increment on its expected value, not the reputation of past returns.

Value beyond channel credit

Positioning, analyst engagement and partner enablement belong in an enterprise CMO's investment review. Their value may run through buyer understanding, confidence or access rather than a traceable channel conversion. Start by stating the mechanism you believe will change a commercial outcome.

Hypotheses for wider GTM investment

Investment	Proposed change	Evidence to examine
Positioning	Clearer buyer understanding	Comprehension tests; later commercial follow-up
Analyst engagement	More accurate market understanding	Narrative accuracy and relevant evaluations
Partner enablement	Better identification and progression of suitable accounts	Phased rollout; comparable partner and account cohorts

These are hypotheses to test, not established causal effects. A useful intermediate signal is not a substitute for incremental commercial value.

Apply a readiness test before promising causal ROI

The advertising studies cited here do not validate causal ROI for these activities. Small deal counts, overlapping programs and inconsistent CRM definitions may prevent isolation. Meridian's guidance on how much data history you need is a starting point, not a guarantee of adequacy. [6] Combine feasible tests with buyer and commercial evidence, retaining the limits with the recommendation.

Minimum viable scale. There is no universal account threshold. An analyst must assess baseline pipeline rates, the smallest worthwhile improvement, desired precision and similarity within account groups. The 40-group example in Section 07 is arithmetic, not a sample-size recommendation. Sparse deal counts may require longer or repeated comparable tests.

Absence of isolated channel credit is not evidence of no value.

05 FROM MY EXPERIENCE

Reallocating a flagship event budget

Around 2022, I was reviewing an event portfolio in North America for an organization serving clients across industries. We invested around $50,000 in a signature partner event, alongside another $60,000–$70,000 in bespoke events. The approximate total was well over $100,000 in that market.

For three consecutive years, the signature event generated conversations and visibility but little of the pipeline we expected. We began asking whether the accounts we met were progressing into commercial opportunities, rather than treating attendance as evidence of acquisition value.

In the fourth year, we stopped sponsoring the signature event and redirected investment toward smaller, focused events and sales-led roundtables. We continued to use the partner platform for visibility, while changing where we placed conversion-focused investment.

We also built a propensity-to-buy model using CRM and lead data, event conversations, sales confidence, historical opportunities, and records of event spend, leads, opportunities and closed business. This was predictive decision support; it did not estimate which conversions the event itself caused.

Following the shift, we observed better conversion, particularly through bespoke events. We did not quantify the improvement. Changes in account mix, follow-up and market conditions could also have contributed.

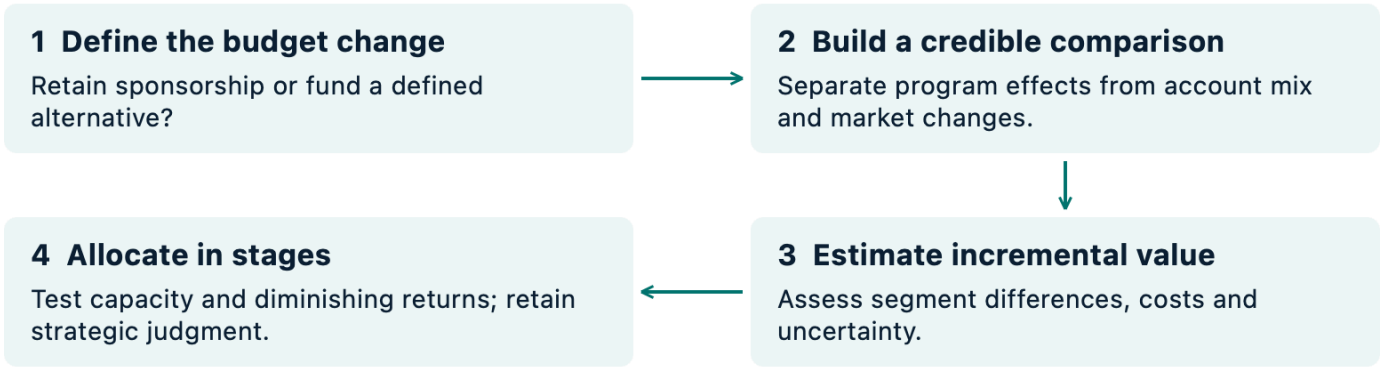
For the next allocation, I would separate the portfolio decision from the program test. Withdrawing sponsorship and offering a focused event answer different questions; one successful roundtable cannot establish the effect of the whole reallocation.

The investment decision changed when we separated visibility from the acquisition outcome the budget was expected to support.

What I would test next time

I would use the four-step workflow below to test the next spending change.

PROPOSED DECISION WORKFLOW



Arrows show the decision sequence. Reassess as new evidence accumulates.

1. **Define the budget change.** Specify the alternative mix, eligible accounts, full costs and commercial outcome. Record the decision, comparison, observation window and expansion rule before launch.

2. **Build a credible comparison.** Randomly assign eligible account groups where feasible, or assess a phased rollout across comparable markets. Record baseline account readiness, demand and sales coverage.

3. **Estimate incremental value.** Analyze everyone assigned, including nonattendees, and track consistently defined opportunities and later wins. Allow for group-level uncertainty and exposure across groups. See the worked example in Section 07.

4. **Allocate in stages.** Compare additional contributions with additional costs, capacity and diminishing returns. Preserve a separate objective for partner visibility. DoWhy can test robustness, but cannot prove observational assumptions. [10]

A study published in Management Science in October 2026 uses 2,226 Meta ad experiments to learn patterns in causal effects. Its Predicted Incrementality by Experimentation (PIE) approach learns from repeated experiments; it does not validate the event portfolio in Section 05 or establish transfer to enterprise B2B. [7]

Disclosure: Moakler works in Ads Research at Meta. Gordon and Zettelmeyer were formerly part-time Facebook employees and now hold concurrent Northwestern and Amazon appointments. The paper describes work performed at Northwestern, not associated with Amazon. [7]

Test the offer you can isolate; do not claim it proves the value of every change in the portfolio.

07 A WORKED INVESTMENT DECISION

Did the extra business justify the cost?

Illustrative example only. These numbers are invented; they are not results from my experience.

Start with a fair test. Randomly assign 40 account groups (25 accounts each): 20 groups receive the focused-event offer and 20 continue with usual activity. Each side contains 500 accounts. Count accounts with at least one new qualified sales opportunity within 180 days, counting each account once. The offer costs $30,000 more than usual activity.

HYPOTHETICAL EXAMPLE · SAME-SIZED GROUPS

FOCUSED EVENT OFFER

80 accounts with a new opportunity

out of 500 accounts (16%)

USUAL ACTIVITY

50 accounts with a new opportunity

out of 500 accounts (10%)

SUBTRACT THE BUSINESS THAT MAY HAVE HAPPENED ANYWAY

80 – 50 = 30

estimated extra accounts with an opportunity · 6 percentage points

Extra program cost ÷ estimated extra accounts with an opportunity

$30,000 ÷ 30 = $1,000 per extra account

Why subtract 50? The comparison group helps estimate how many accounts might have progressed anyway. Crediting all 80 accounts' progression to the event would overstate its value.

Does it pay? A new opportunity is a potential deal, not a sale. Business won from the estimated 30 additional accounts must contribute more than $30,000 after customer delivery costs to cover the extra event spend.

My decision. Before spending more, I would ask whether the gap is reliable or could reflect chance. The marketing analyst must assess uncertainty using the 40 randomized groups. If the result is too uncertain, I would extend or refine the test.

Who benefits most? Only then would I examine which customer groups respond better. The Lemmens review and EconML methods support this question; they do not validate the event portfolio in Section 05. [9, 11]

Spend more when the extra business is credible and its expected value covers the extra cost.

Build the evidence before optimizing

I would begin with the systems already in use: connect CRM, event, marketing and finance records into one account history.

Data element	What to record	Why
Account identity	Account, parent, contact and opportunity IDs	Connect activity to outcomes; avoid double counting.
Program assignment	Eligibility, invitations, attendance, comparison group, date and format	Retain nonattendees; analyze assignment rather than attendance.
Commercial history	Pre-activity status; dated opportunity stages, wins and margin	Establish the baseline and follow commercial value.
Sales and context	Coverage, follow-up, other campaigns and material market changes	Examine alternative explanations and delivery differences.
Full costs	Sponsorship, delivery, travel, agency and staff effort	Compare economics and the activity forgone elsewhere.
Experiment records	Hypothesis, assignment, outcome window, sample size and analysis plan	Make the test interpretable; check power and cross-group exposure.

A practical stack. Use an existing warehouse or managed database, SQL and statistical analysis. Add causal tools where justified; show incremental outcomes, costs and uncertainty in a decision dashboard.

Test feasibility. Use the readiness checks in Section 04; the 40-group illustration in Section 07 is not a minimum sample size.

Effort and cost. My illustrative clean-data allowance is 5–10 analyst days plus 3–5 operations days, valued at fully loaded internal rates. This excludes event delivery, infrastructure and sales follow-up. It is a scoping assumption, not a benchmark; unresolved joins can materially increase effort.

Ownership. Marketing operations owns data quality; sales operations account history; finance the economics; an analyst with experimental-design skills the evaluation. Name the budget decision-maker.

If the test cannot distinguish a commercially useful improvement from noise, fund better evidence before a larger commitment.

Lift must earn its place in the budget

A positive effect is only one part of an investment decision. Compare incremental value with fully loaded program and measurement costs, implementation time and the business forgone elsewhere. For recurring revenue, state retention, margin and cash-timing assumptions rather than substituting headline contract value for profit.

Use consistent horizons. A short experiment can miss delayed benefits or costs, while a long forecast may depend heavily on assumptions. An intermediate outcome such as qualified opportunity creation is an early signal; follow the cohort through to revenue and economics before calling it a realized return.

My preference is to agree to the next spending limit with finance before seeing the result. A model recommendation becomes useful when the company knows what evidence would release, defer or stop that commitment.

Measurement has an opportunity cost too. Estimate whether the decision is large and repeatable enough to justify the analysis. A simpler test or a staged commitment may be more useful than a costly model that cannot change the action.

A practical leadership agenda

Period	Leadership action
First 30 days	Agree the decision, outcome and horizon with finance and sales. Check definitions, baseline information and test feasibility.
Days 31–60	Choose a feasible comparison. Assess power and measurement cost. Record what would support expansion, continuation or stopping.
Days 61–90	Review implementation and early evidence. Make a bounded decision where justified; schedule later commercial follow-up.

Treat days 61–90 as an early read. The commercial outcome in Section 07 needs up to 180 days. Assign an owner to the decision, and record the rationale when judgment overrides a model recommendation.

Make the recommendation reviewable

Five questions can make the discussion more useful:

1. What specific investment change are we considering?
2. What would happen under the alternative, and what makes that comparison credible?
3. Are we measuring incremental commercial value, attributed activity or an intermediate signal?
4. Which assumptions could reverse the recommendation?
5. What is the smallest practical next step that would improve the decision?

Causal AI can help estimate what an intervention changes. Leadership still needs to decide whether that change is commercially valuable, affordable and consistent with the company's growth strategy.

Three options belong in the discussion

Scale	Learn	Hold
Evidence and economics support a defined expansion.	A bounded follow-up can resolve a material uncertainty.	Preserve the current allocation when change is not justified.

My role is to help leadership frame the GTM decision, identify which evidence is useful and translate the findings into commercial choices. A measurement specialist may be needed for experimental design or model implementation. The advisory value is a decision the business can understand, challenge and act on.

Request a 30-minute marketing investment decision review. Bring one allocation choice and a summary of your CRM, event and cost data. We will identify the evidence gap and a practical next step: sowmyamoni.com/contact

RESEARCH NOTES

Sources and evidence boundaries

1 Gartner CMO Spend Survey

11 May 2026. Public summary of a cross-industry survey of 401 marketing leaders. Self-reported; not a causal study.

2 A Comparison of Approaches to Advertising Measurement

Gordon et al. Marketing Science 38(2), 193–225, 2019. Peer-reviewed; 15 US Facebook ad experiments.

3 Close Enough?

Gordon, Moakler and Zettelmeyer. Marketing Science 42(4), 768–793, 2023 issue; online 7 November 2022. Peer-reviewed; 663 Facebook experiments.

4 Meridian is now available to everyone

Google, 29 January 2025. Product announcement used for availability and capabilities, not independent effectiveness.

5 MMM as a causal inference methodology

Google for Developers, updated 1 September 2026. Technical documentation on assumptions and causal interpretation.

6 Data collection and adequacy

Google for Developers. Collect and organize your data; see also linked Amount of data needed guidance. Checked 25 October 2026.

7 Predicted Incrementality by Experimentation for Ad Measurement

Gordon, Moakler and Zettelmeyer. Management Science, published online 8 October 2026; accepted 8 April 2026. Peer-reviewed; 2,226 Meta ad experiments; not B2B validation.

8 ROI, marginal ROI and response curves

Google for Developers. Technical documentation; definitions, diminishing returns and limits of extrapolation. Checked 25 October 2026.

9 EconML: causal effect estimation

Microsoft Research. Software overview on estimating treatment effects and differences between groups; not evidence that this event reallocation worked. Checked 25 October 2026.

10 DoWhy: refuting causal estimates

PyWhy documentation. Robustness and sensitivity checks for causal estimates; passing checks does not establish that all assumptions hold. Checked 25 October 2026.

11 Personalization and targeting: how to experiment, learn & optimize

Lemmens et al. International Journal of Research in Marketing 43(2), 537–559, June 2026 issue. Peer-reviewed review of causal test-and-learn methods; no new field experiment. DOI: 10.1016/j.ijresmar.2025.07.004.

Source titles are clickable. Case figures are approximate and the conversion improvement is qualitative. No fitted causal model or controlled causal estimate is presented. Proposed experiments and the Section 07 example are author illustrations, not results.