

The definitive guide to marketing measurement

Index

Executive summary 4

1. Why measurement feels broken 6
- More information means less clarity
- The trust gap between marketing and the C-suite
- The attribution trap: asking the wrong question
- Method confusion: when every tool looks like a hammer

2. The questions that drive measurement 10
- Diagnostic questions: what happened?
- Causal questions: did it work?
- Predictive questions: what should we do next?
- Matching the right method to the right question

3. Diagnostic and descriptive analysis 13
- The foundation: funnels, KPIs and dashboards
- Cohort analysis and customer behavior
- What descriptive analysis can and can't tell you

4. Attribution models 18
- How attribution models work
- What attribution can and can't tell you
- Where attribution is headed

5. Experimentation and testing 22
- Randomized controlled trials
- A/B testing
- Geo-testing and market-level experiments
- Experimental design challenges

6. Incrementality and lift measurement 28
- What incrementality measures
- Incremental ROAS
- Conversion lift studies
- Limitations of incrementality testing

7. Marketing mix modeling (MMM) 33
- How marketing mix modeling works
- MMM versus legacy approaches
- Limitations of marketing mix modeling

8. Saturation, diminishing returns and marginal analysis 38
- Audience saturation
- Marginal ROAS, POAS and MER
- Optimal budget allocation

9. Forecasting and scenario planning 43
- Planning the future by understanding the past
- Scenario modeling with MMM
- Why simulations help with stakeholder buy-in
- Connecting measurement to planning cycles
- Limitations of forecasting

10. Financial and unit economics analysis 47
- Connecting marketing to business outcomes
- Key financial metrics
- Building the business case for marketing



11. Measuring brand effects 52
- What brand lift measures
- The long and short of brand building
- Measuring brand effects over time
- Limitations of brand measurement

12. Audience and targeting effectiveness 56
- Measuring audience targeting
- AI-driven targeting evolution
- First-party data and privacy-safe targeting

13. Uplift and persuasion modeling 60
- Defining uplift modeling
- The four customer segments
- Uplift modeling methods
- Practical applications
- Limitations

14. Qualitative and mixed-methods analysis 64
- The role of qualitative research in marketing measurement
- Mixed-method research designs
- Qualitative research methods for marketers
- Emotional testing and creative effectiveness
- Limitations

15. Measurement maturity and operating models 68
- Assessing your current state
- The maturity gap
- BCG's four accelerators
- Building the operating model

16. Advanced measurement and triangulation 73
- Why triangulation became the standard
- How modern systems reconcile competing signals
- Limitations

17. Econometrics and causal inference 78
- What econometrics is and why it matters for marketers
- The counterfactual problem
- Methods for estimating the counterfactual
- How the methods connect

18. The promise of agentic measurement 84
- The scale of the AI pivot
- Static models moving to continuous systems
- The role of human judgment
- Challenges and limitations

Measurement as a discipline 88

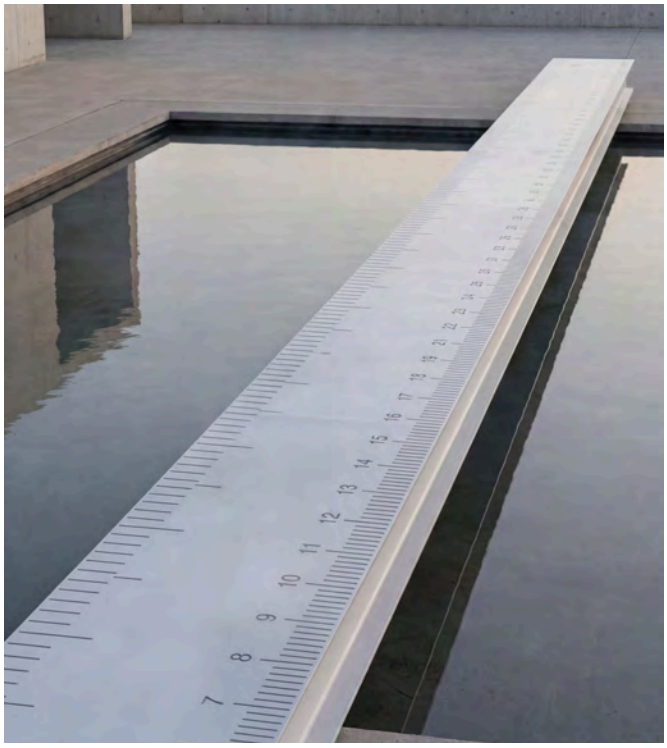
References 92

Executive summary

Marketing measurement should be the engine that drives smarter spending, stronger growth and clearer accountability. Instead, for most organizations, it's become a source of confusion and, increasingly, of conflict between marketing teams and the executives who fund them.

The reasons are structural. Marketing data has exploded in volume while clarity has not kept pace. Integration remains the single biggest martech management challenge for CMOs,[1] and more than a third admit they don't trust their own data.[2] Simultaneously, teams still rely on attribution models, particularly last-touch and platform-native attribution, to answer questions that those models were never designed to address. Attribution assigns credit among digital touchpoints but doesn't tell you if a conversion would have happened without the campaign. The distinction between credit allocation and causal measurement sits at the very heart of why marketing dashboards can look strong while growth stagnates.

If marketing can't confidently tie spend to impact, it's hard to justify more investment. As a result, marketing budgets have flatlined at 7.7% of company revenue, and 59% of CMOs report insufficient budget to execute their strategy.[3] The trust gap between marketing and finance is widening, and clarity from effective measurement has never been more important.



But that doesn't mean marketing teams should use every measurement technique all the time. That's a sure way to create more confusion and frustration. Instead, marketers should have a solid grasp on what's possible so they can execute the best types of measurement for the questions they're trying to answer. The foundational premise of this guide is that different questions require different measurement approaches, and to treat them as interchangeable is the root cause of confusion.

We therefore organize measurement questions into three categories: diagnostic (what happened?), causal (did it work?) and predictive (what should we do next?), and walk through each major methodology in depth. For every approach, we explain what it measures, where it falls short and when to use it. No single method is a silver bullet for every question a marketer needs to answer.

The most effective organizations don't choose between marketing mix modeling, attribution and incrementality testing. They build measurement systems that combine multiple methods, reconcile competing signals and connect analytical outputs to business decisions. Measurement, done well, is a discipline, not a dashboard.

Use this guide as a framework to build that discipline.



01

Why measurement feels broken





The core tension in modern business is the gap between marketing reports and the bottom line. It's a common, frustrating scenario. Marketing dashboards report high conversion rates, yet the finance team sees no corresponding growth in net revenue.

When the two perspectives clash, the immediate assumption is that someone's data is wrong. In reality, the problem is deeper: the system itself is producing conflicting versions of the truth. Marketers aren't bad at data, but their measurement challenge has outgrown the tools teams use to address it.

More information means less clarity

Marketing teams have access to more data than at any point in history. But they also have less clarity about what the data means for their business.

According to the CMO Survey and MarTech, 65.7% of CMOs cite data integration as their biggest martech management challenge.[1] Roughly 31% of senior marketers continue to face challenges connecting data from different sources.[4] And the consequences are financial: 55% of US marketers believe poorly integrated data has caused revenue loss, while 34% of CMOs say they don't trust their data at all.[2,5]

Professor Sunil Gupta of Harvard Business School observes that the starting point to know if campaigns are working is having the right metrics aligned to your objectives.[6] It sounds straightforward in theory, but in practice, marketing teams are measuring outputs (clicks, impressions, conversions) without a clear connection to business outcomes (revenue, profit, customer lifetime value). The volume of data creates an illusion of precision, while the lack of integration obscures the signal.

The trust gap between marketing and the C-suite

The data paradox, in turn, feeds a larger credibility problem. Marketing's relationship with finance is strained in most organizations.

According to Gartner, 33% of marketers admit to cherry-picking data to support preconceived choices. Another 24% rely on gut feelings, and 24% reject data-driven recommendations altogether.[7] These are self-reported numbers from the marketers themselves, so they're not entirely objective. That means the actual figures may be worse.

Meanwhile, the IPA Databank reveals that although return on investment has increased by roughly 4% since COVID, net profit is actually down 11%.[8] In other words, ROI looks better on paper while the actual financial contribution of marketing has declined. As McKinsey has found, CFOs remain especially skeptical of existing marketing ROI approaches.[9]

The budget picture reflects this skepticism. Marketing budgets are stagnant.[3] Yet 61% of CMOs say marketing is viewed as a profit center, up from 54% the prior year. The expectation is growing while the resources are not.

As Ewan McIntyre, VP Analyst at Gartner, has put it, CMOs cannot risk incremental change when the enterprise expects transformative results.[3]

The attribution trap: asking the wrong question

Beyond data complexity and trust gaps, there is a deeper structural reason measurement feels broken. One that teams don't recognize because the tool causing the problem is one they rely on every day.

Marketers still use attribution as their primary lens for performance. According to ProperExpression and OWOX BI, 41% of marketers use last-touch attribution and 44% rely on first-touch attribution.[10] Both are single-touchpoint models that ignore everything else in the customer journey. Google Analytics 4's data-driven attribution attempts to move beyond this by algorithmically weighting multiple touchpoints, but it operates as a black box with a 90-day lookback window; marketers can't see how credit is assigned or account for anything that happened outside that window.[11]

Attribution was designed to answer which touchpoints should receive credit for a conversion. So for optimizing within digital channels, comparing ad sets, evaluating keywords and testing creatives, that question is useful.

Because attribution models prioritize trackable, lower-funnel digital interactions, they often undervalue upper-funnel brand activity and broad-reach media. The narrow lens creates a blind spot for offline impact, such as TV, out-of-home and direct mail, while ignoring the compounding effects of long-term brand building. Most importantly, the models cannot account for the counterfactual: if a conversion would have occurred anyway, even without marketing exposure.

So when a CFO asks if a campaign generated revenue, the underlying question they're really asking is about effectiveness. Did this spend produce outcomes the business would not have seen otherwise? You need to run incrementality tests to find out. Attribution assigns credit to touchpoints that preceded a conversion. It doesn't distinguish between a customer who converted because of the ad and a customer who was already heading to checkout when the ad appeared.

But this is not to say attribution is wrong. It answers the question it was designed to answer. The problem is that when organizations treat attribution outputs as if they were answering the effectiveness question, they present credit allocation to finance as evidence of causal impact. What attribution measures and what stakeholders interpret it to mean is a major source of the skepticism discussed earlier. It also explains the frustrating pattern in which dashboards show strong ROAS while the business doesn't see corresponding profit growth.

As McKinsey has observed, this dynamic has historically reinforced silos. Digital marketers point to attribution models to justify search budgets while brand managers cite marketing mix models to defend TV spend, and neither approach addresses what the other measured.[12]

The confusion doesn't stem from having too much data. It stems from misapplied methods to the wrong questions.

Method confusion: when every tool looks like a hammer

The attribution trap is a symptom of a broader problem: organizations don't have a clear framework for matching measurement methods to measurement questions.

At Cannes Lions 2024, industry experts estimated that just 6% of advertising spend drives real results.[13] The challenge isn't that the other 94% is definitely wasted, it's that existing tools can't reliably distinguish between the spend that drives outcomes and the spend that captures demand that would have occurred anyway.

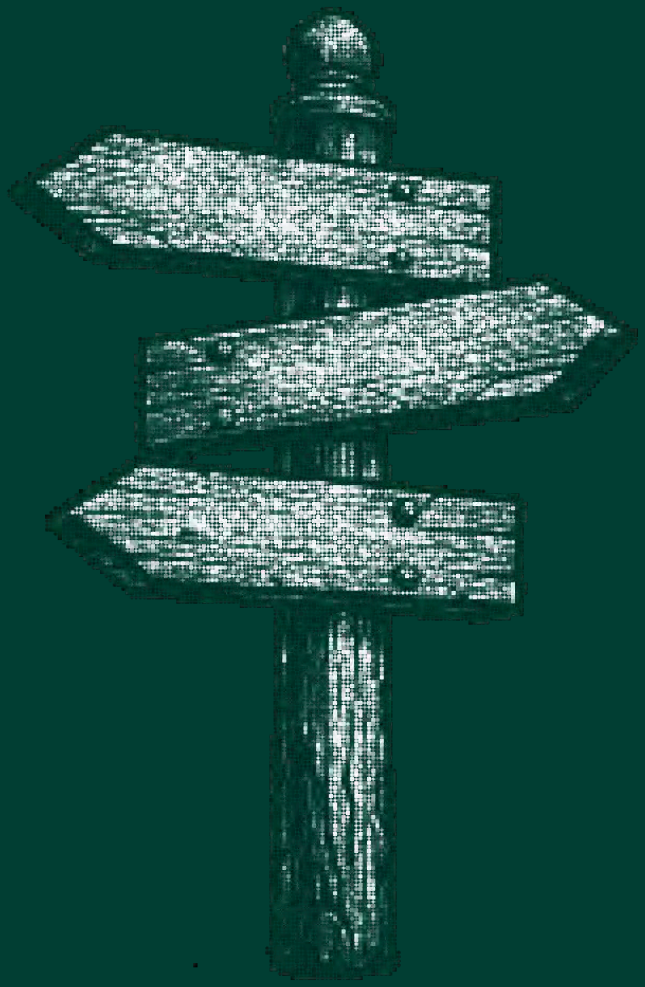
When every measurement approach is treated as interchangeable, such as when attribution reports are used to justify brand budgets, the result is confusion and eroded trust.

The rest of this guide provides a framework for solving that problem. We start by defining the three types of questions measurement needs to answer, then walk through each major methodology with a clear-eyed view of what it does, what it doesn't do and when to use it.



02

The questions that drive measurement





If the root cause of measurement confusion is misapplied methods, the fix starts with matching the right method to the right question.

Measurement questions usually fall into three categories: diagnostic, causal and predictive. Each requires different approaches, data and assumptions.

Diagnostic questions: what happened?

Diagnostic questions look backward to understand performance; they are the most common type of measurement question and the most intuitive. How many leads did we generate last month? Which channels drove the most traffic? What was our conversion rate by campaign? What was the cost per acquisition across our paid media channels?

These questions are answered through descriptive analytics; funnels, key performance indicators, dashboards and cohort analysis. And they are essential. Every marketing team needs reliable diagnostic capabilities to monitor performance and identify patterns.

But diagnostic questions have a hard ceiling; they can tell you what happened, not why it happened. Observing that paid search drove 40% of conversions doesn't tell you if those conversions were incremental or if the customers would have found you through organic search anyway. That's a causal question, and it requires a different approach.

Causal questions: did it work?

Causal questions attempt to establish if marketing activity really caused outcomes or if those outcomes would have happened anyway. Did this campaign drive incremental sales? Would these customers have converted without seeing our ads? What was the true return on our spend, not the attributed return, but the actual impact?

As such, causal questions are harder to answer than diagnostic questions because they require estimating a counterfactual: what would the world have looked like if we hadn't run the campaign? We can never observe this directly. We can only estimate it through methods designed for causal inference like randomized experiments, incrementality testing, holdout studies and quasi-experimental designs.

Predictive questions: what should we do next?

Predictive questions use historical patterns to inform future decisions. How should we allocate budget across channels next quarter? What happens if we increase TV spend by 20%? Where is the point of diminishing returns for each channel? What is the optimal media mix for next year?

The answers are found through forecasting, scenario planning, marketing mix modeling and optimization. They represent the highest-value application of measurement because they inform spending decisions directly, but they also depend on the quality of diagnostic and causal work that feeds them.

A forecast built on attributed conversions will produce different recommendations than one built on incremental impact. The method you use to answer "what happened?" and "did it work?" will give the answer to "what should we do next?"

Matching the right method to the right question

Different questions require fundamentally different measurement approaches, and treating them as interchangeable is what causes confusion. As BCG has noted, companies need to employ a suite of methodologies to get a full picture, especially if they want to predict the impact of changes in spending across individual channels.[14]

Recognizing which question you're asking and choosing the method built to answer it is the first step toward a measurement system you can trust. The rest of this guide walks through each major approach, from descriptive analytics and attribution through experimentation, incrementality, marketing mix modeling, brand measurement and beyond. For each method, we explain what it does, what it can't do and when to use it.



03

Diagnostic and descriptive analysis



Before you can measure what works, you need to know what happened. Diagnostic and descriptive analysis provides the foundation for all marketing measurement by tracking performance metrics, identifying patterns and finding anomalies that warrant deeper investigation.

Every marketing team uses some form of descriptive analytics. The challenge isn't adopting these tools. It's understanding their limits so you don't over-claim what the numbers mean.

The foundation: funnels, KPIs and dashboards

Marketing teams organize performance data around a funnel, tracking how customers move from awareness through consideration to conversion and retention. And, they measure where they drop off. The KPIs attached to each stage are what give the funnel its diagnostic power.



Leading indicators, such as website traffic, ad impressions and email open rates, provide early signals. Lagging indicators like revenue, customer acquisition cost and retention rate confirm whether those signals translated into results. A well-designed system tracks both.

The dashboard is where all of this comes together. Or falls apart. As McKinsey states, linking KPIs across channels and funnel stages to real business results allows companies to better understand the impact of their marketing and craft messages that drive stronger responses.[15] The keyword here is "linking." A dashboard that displays channel metrics in isolation, without connecting them to business outcomes, is just noise that provides no real insight.



Cohort analysis and customer behavior

Cohort analysis groups customers by shared characteristics (typically the time period when they were acquired) and tracks their behavior over time. The advantage of using cohort analysis is that it provides richer insight than aggregate metrics. For example, you can see if performance has changed for new customers versus existing ones.

Harvard Business School explains that the pitfalls of not conducting cohort analysis are significant, because aggregate customer metrics can obscure changes in retention, revenue and engagement that only become visible when you segment by when customers joined, how they were acquired or how heavily they use a product.[16]

For example, aggregate conversion rates might look stable month over month. But cohort analysis might reveal that recent cohorts convert at a lower rate and churn faster, a signal that acquisition quality is really declining even while headline numbers hold steady.

What descriptive analysis can and can't tell you

Descriptive analytics shows correlation, not causation. It describes patterns such as channel X had the highest conversion rate or cohort Y retained better than cohort Z. However, descriptive analysis can't tell you if marketing activity caused these patterns.

It's easy to draw causal conclusions from descriptive data. If paid search has the highest conversion rate, a natural reaction is to increase search spend. But if search is primarily capturing people who were already looking for your product, the high conversion rate reflects existing demand, not the impact of your marketing efforts.

Descriptive analysis is also inherently backward-looking. It tells you what happened last month. It doesn't predict what will happen next month under different conditions. And it's vulnerable to selection bias. If your reports only include audiences who have already engaged with your brand, the numbers will look better than they are because you're measuring a warm audience, not the total addressable market.

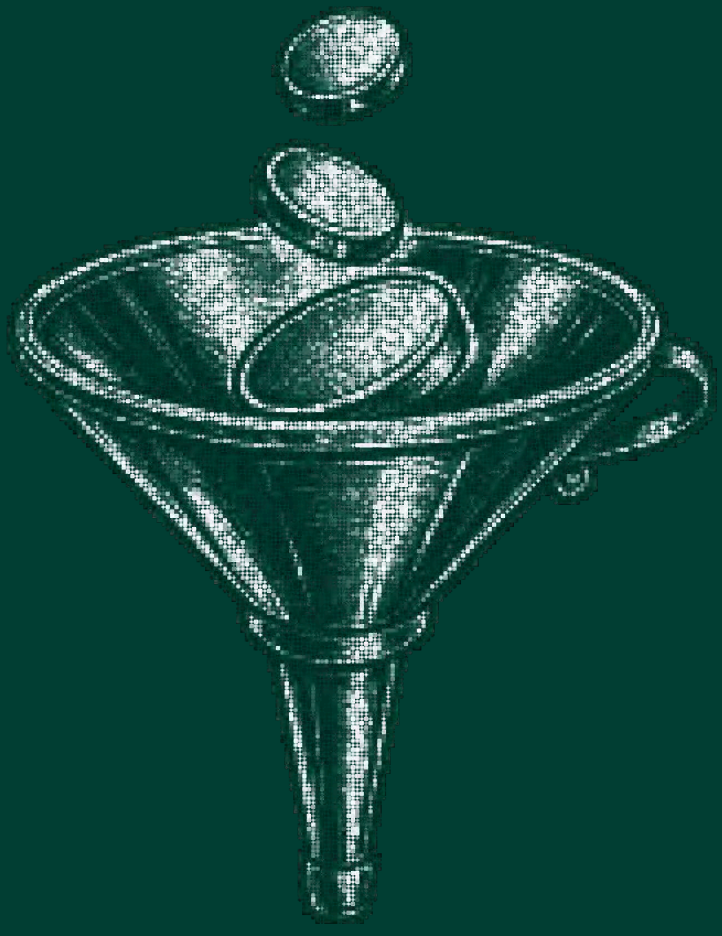
As observed by Harvard Business School, marketing success requires a balance between quantitative analysis and qualitative judgment. The metrics need to go beyond surface-level patterns to account for factors like the quality of leads and the long-run value of customers acquired through different channels.[17]

None of this means descriptive analytics is optional. It's the bedrock. Every measurement system starts here. The discipline is in recognizing where descriptive analysis ends and causal or predictive methods need to begin.

Use diagnostic and descriptive analysis when you need to monitor ongoing performance, identify patterns in customer behavior, find anomalies that need investigating or build a shared understanding of what happened before asking why.

04

Attribution models





Attribution has been the default measurement framework for digital marketing for more than a decade. It attempts to answer a specific question: when a customer converts, which marketing touchpoints along their journey should receive credit?

While the question is useful, it’s narrower than most teams realize, which is why building reporting around it is a mistake. Understanding what attribution does, what it doesn't do and where it's heading is essential for any modern measurement strategy.

How attribution models work

Rules-based attribution models assign credit according to predetermined formulas. Each model reflects a different assumption about which touchpoints matter most.

First-touch attribution gives 100% of credit to the initial interaction. It's useful for understanding which channels drive awareness and top-of-funnel discovery. The limitation is that it ignores everything that happens between awareness and conversion: the nurturing, retargeting and persuasion that can influence the final decision.

Last-touch attribution gives 100% of credit to the final interaction before conversion. It's useful for identifying which channels close deals. The limitation is that it ignores all prior influence. As ReportGarden has documented, just 4% of internet users click on advertisements, with almost no established link between click-through rates and actual purchasing patterns.[18]

Linear attribution distributes credit equally across all touchpoints. It provides a simple, balanced view but assumes every interaction is equally valuable, which is rarely true.

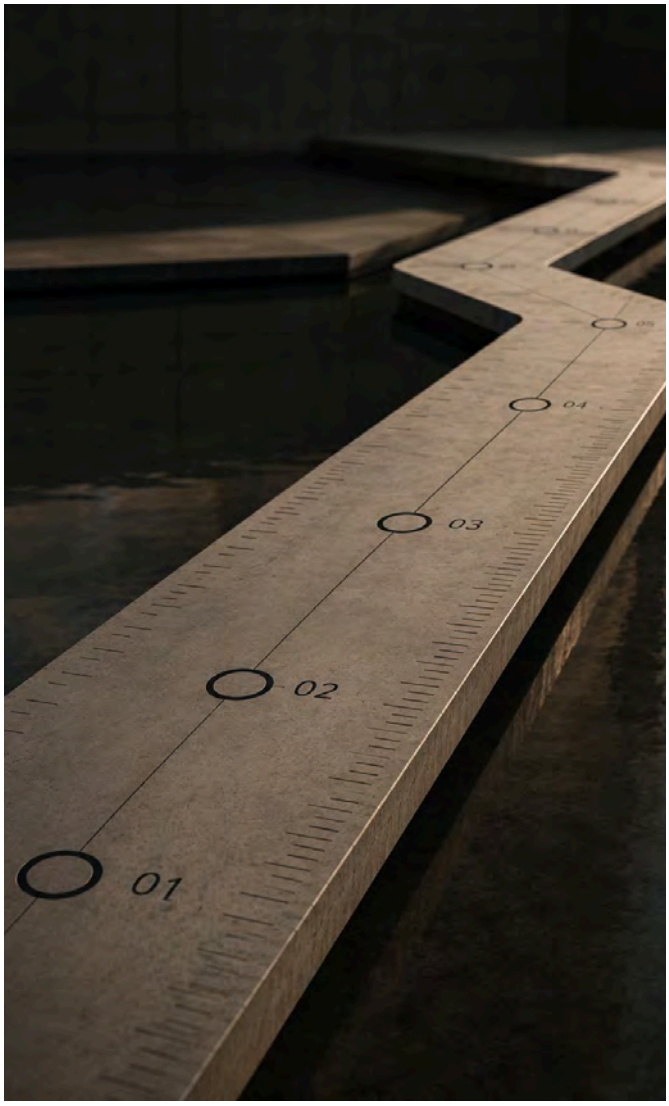
Time-decay attribution assigns more credit to interactions closer to the conversion. It works well for short sales cycles where recent touches are more influential. But it can systematically undervalue awareness activities that happened earlier in the journey.

Position-based (or U-shaped) attribution assigns 40% to the first touch, 40% to the last touch and distributes the remaining 20% across middle interactions. It attempts to balance awareness and conversion but does so through arbitrary weighting.

Each of these models represents a different hypothesis about how marketing works. None of them can be validated against ground truth because they are assigning credit, not measuring impact.

Multi-touch attribution was supposed to fix these limitations. Instead of predetermined formulas, algorithms would analyze actual conversion paths and weight touchpoints based on their observed contribution. In practice, the options have narrowed considerably. In September 2023, Google removed first-click, linear, time-decay and position-based models from both GA4 and Google Ads, citing adoption rates below 3%.[11] As Ruler Analytics and Feld-M have both argued, the deeper issue was that privacy changes and fragmented user journeys had already undermined the reliability of multi-touch attribution.[11,21]

That left marketers with two choices: last-click or Google's data-driven attribution. Google's data-driven attribution uses a counterfactual approach, comparing the probability of conversion with and without each touchpoint to assign credit.[19] In theory, this is more sophisticated than any rules-based model. In practice, it has constraints. It operates within a 90-day lookback window, which means longer customer journeys get truncated. It works as a black box; marketers can see the credit assignments but not the reasoning behind them. And it can only observe touchpoints that reach your site via the GA4 tag, so offline activity, in-platform engagement within walled gardens like Meta or TikTok and CRM pipeline interactions all fall outside its view.[11]



What attribution can and can't tell you

Think with Google has stated that relying on last-click attribution can lead to misallocated budgets and missed revenue opportunities, and that data-driven attribution provides more accurate credit assignment.[19] That's likely true for the touchpoints the model can see. But a more sophisticated algorithm distributing credit across an incomplete picture of the customer journey doesn't close the gap. It just allocates within it more precisely.

The more fundamental issue is what attribution measures even when it's working as designed. Attribution assigns credit for conversions to touchpoints. Incrementality measures whether those conversions would have happened without the marketing.[20]

And these are different questions entirely. A retargeting ad might receive last-touch credit for a sale, but if the customer was already heading to checkout, the ad intercepted the conversion rather than caused it. As McKinsey notes, attribution models depend on cookie data as an input, which limits the richness of the data set and makes it difficult to accurately attribute the importance of each touchpoint.[22] But even with perfect tracking, attribution still wouldn't answer the causal question finance teams are asking.

Where attribution is headed

The tracking infrastructure that attribution depends on is forever eroding. Less than one in five consumers always accept cookies, and future estimates suggest 90% of US browsers could become effectively cookieless, though Google has delayed full cookie deprecation in favor of user choice prompts.[23] Apple's iOS 14.5 update gave iPhone users the choice to opt out of cross-app tracking, and the majority did, disrupting the user journey data that mobile attribution relied on.[24]

The consequence is that attribution models need to see the customer journey to assign credit across it. When users decline cookies, switch between devices or browse in privacy-focused environments, parts of the journey go dark. The model doesn't know they're missing, it just assigns credit across whatever touchpoints it can still see, which skews results toward channels with better tracking (typically lower-funnel and owned digital) and away from channels that are harder to track.

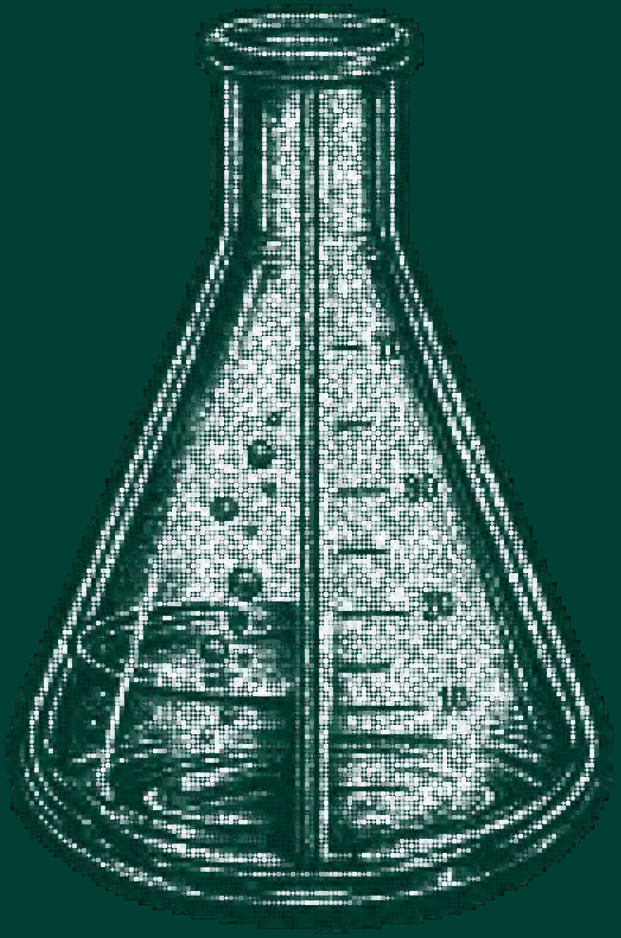
The industry has responded in two directions. Platform vendors, Google and Meta in particular, have invested in modeled attribution, using machine learning to statistically estimate the conversions and touchpoints that tracking can no longer observe directly. Yes, it fills some of the gaps, but it also means marketers are increasingly relying on the platforms' own models to evaluate the platforms' own media.

The broader marketing community is moving toward methods that sidestep user-level tracking, like incrementality testing and marketing mix modeling, to build a clearer picture. However, these approaches aren't replacements for attribution. What they do is answer different questions. But they are resilient to the privacy changes that are undermining attribution's data foundation.

Use attribution models when you need to optimize within digital channels, compare the relative performance of campaigns or touchpoints or allocate credit across a known customer journey. Do not use attribution as evidence of causal impact or as your primary metric for C-suite level reporting.

05

Experimentation and testing





Where attribution assigns credit, experimentation establishes causation. Randomized experiments are the most direct way to answer the question that attribution can't: did this marketing activity cause outcomes that would not have happened otherwise?

The logic is pretty straightforward: show ads to one group (the treatment), withhold them from another (the control) and compare the results. If the treatment group performs better by a statistically significant margin, you have evidence that the ads worked. The control group provides an estimate of the counterfactual: what would have happened without the marketing.

Randomized controlled trials

Randomized controlled trials (RCTs) are the most rigorous approach to causal inference. Hal Varian writes in his PNAS paper that randomized trials provide the gold standard for causal inference because if treatment assignment is random, the control group serves as an unbiased estimate of the counterfactual, and selection bias is eliminated by design.[25]

In marketing, RCTs take the form of holdout experiments, conversion lift studies and audience-level randomization. The underlying principle is always to randomly assign exposure and non-exposure, then measure the difference.

Yet, despite the mathematical simplicity of the method, the real-world challenge is rarely technical. It's organizational. Implementing RCTs requires a marketing team to do something that feels inherently counterintuitive: deliberately withholding marketing from a portion of their audience. It requires a commitment to a process that looks like failure before it looks like insight. As Stefan Thomke wrote in Harvard Business Review, for every online experiment that succeeds, nearly ten do not. In organizations that prioritize immediate efficiency and total predictability, these "failed" tests can look like a waste of budget.[26] In reality, however, they are the essential mechanism through which an organization learns what truly works.

This failure rate is but a reflection of the difficulty of moving the needle in a crowded marketplace. Even at the highest levels of digital maturity, the success rate for experiments remains strikingly low; only 10% to 20% of Google's and Bing's online experiments generate positive results.[27] But for these companies, this isn't a deterrent; it's a learning rate. Tech giants like Amazon, Booking.com and Microsoft annually run tens of thousands of controlled experiments to find those few winning insights. At Booking.com, this culture is so deeply embedded that 75% of its core employees participate in experimentation, possessing the autonomy to test and validate their own hypotheses independently.[28]

Ultimately, the shift toward RCTs is a shift in power dynamics. Thomke observes that nothing stalls innovation faster than a HiPPO: the highest-paid person's opinion.[29] Organizations that successfully build experimentation into their DNA create an environment where evidence consistently outweighs seniority. In these cultures, negative results are valued for the waste they prevent and the clarity they provide. The cost of testing is no longer viewed as a sunk expense, but as the necessary cost of learning.

A/B testing

A/B testing serves as the tactical workhorse of the modern marketing department. It's the most accessible and widely used form of experimentation: compare two or more versions of a digital asset by randomly assigning visitors to each and letting their behavior dictate the winner.



However, the perceived simplicity of a split test often masks the rigorous mechanics required to make it meaningful. Most marketing teams default to a 95% confidence threshold ($p = 0.05$) to determine success, but the right threshold depends entirely on a team's specific appetite for risk.[30] In the high-pressure environment of a growth team, it can be sobering to realize that "inconclusive" is the most likely outcome. Indeed, research has shown that only about 20% of experiments ever reach that 95% significance bar.[31] The reality forces teams to choose between waiting longer for a clear signal or making a directional call based on less-than-perfect data.

But how much data is enough? In truth, there is no universal minimum sample size that guarantees a valid result. The amount of traffic required is a moving target, dictated by your baseline conversion rate and the specific lift or effect size you are attempting to detect.[32] While the industry relies on rough rules of thumb such as 5,000 visitors per variation and at least 100 conversions, these should be viewed as helpful guardrails rather than hard requirements.[33]

Despite these technical hurdles, adoption has become nearly universal. Today, 77% of organizations conduct A/B testing on their websites, 60% optimize their landing pages and a massive 93% of US companies apply these principles to their email marketing.[33] Yet, even with the right tools and traffic, timing remains the ultimate confounding factor. Effective testing requires a Goldilocks duration: usually between two and six weeks.[34] Anything shorter than two weeks risks capturing a temporary blip rather than a real trend, while stretching a test beyond six weeks invites external noise like seasonality.

There is no universal minimum sample size that guarantees a valid result.

Geo-testing and market-level experiments

Though A/B testing is perfect for the digital sandbox, it fails when marketing moves into the physical world or broad-reach media. You cannot randomly assign a passerby to see a billboard, nor can you easily split a TV audience down the middle. But by using geography as the proxy for the individual, teams can execute geo-holdout tests and purposely keep specific regions dark to serve as a baseline. At the same time, they can run campaigns in others to measure the true delta in performance.[35]

Geo-split tests are where different marketing strategies are deployed across separate regions to see which one moves the needle further. To ensure these comparisons are fair, teams rely on matched market tests, which involve carefully pairing twin markets based on shared traits and historical sales patterns.[36] The goal is to ensure that any difference in revenue is actually the result of the marketing treatment rather than just the fact that consumers in one city happen to behave differently from those in another.

Meta's GeoLift methodology, which is open source, uses synthetic control methods to construct a "synthetic" control from a weighted combination of non-treated markets. Synthetic control methods have been described as arguably the most important innovation in policy evaluation in the past 15 years.[37,38] For marketing teams, the true superpower of GeoLift is that it requires zero user-level tracking, making it an incredibly resilient strategy in an era of increasing privacy constraints.

However, it isn't a set it and forget it tool; robust power calculations are essential. Without them, there is a high probability of a false negative, where the model fails to detect a lift that actually occurred. In one case study, a leading retailer tested its catalog program with a known-audience holdout test and found the catalog delivered 14% incremental lift in revenue, far below the 40% contribution claimed by vendor attribution reporting.[39]



Experimental design challenges

Experiments represent the best way to establish causation, but they are not a frictionless endeavor. In a perfect world, every campaign would be a controlled study; in a marketing team's reality, experiments come with practical limitations that dictate how and when they can actually be deployed.

For example, required sample sizes can be large and expensive. Detecting a small improvement, say a 2% increase in conversions, requires substantial traffic volume. That means running the test longer or allocating more budget, both of which carry costs.

Then, the spillover effects between test and control groups can dilute results. In the messy ecosystem of modern digital life, a control group is rarely truly unexposed. Consumers talk to one another, share links, switch between multiple devices and encounter your brand through organic channels that an experimental design can't fully gate. Leakage from the test group into the control group inevitably dilutes the results. It also makes the measured difference appear weaker than the true effect, leading teams to potentially undervalue a winning strategy.

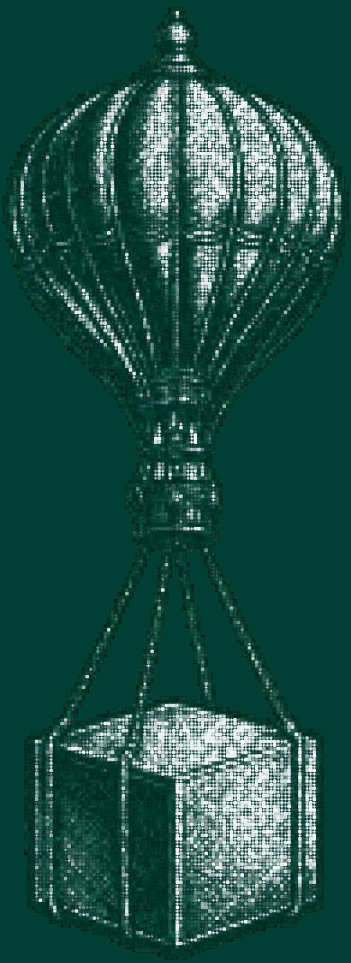
And even when a test produces a statistically significant result, the question of external validity remains. A successful test is essentially a snapshot of a specific audience reacting to specific creative during a specific week. Whether that same uplift will materialize at scale, or in a different season with different messaging, is a matter of professional judgment.

That said, the challenges do not diminish the value of experimentation; rather, they explain why it must be a component of a larger measurement system rather than a standalone solution.

Use experimentation and testing when you need to establish whether a specific marketing activity causes outcomes. Testing is particularly useful when making go/no-go decisions on campaigns, validating assumptions before scaling spend or building an evidence base that satisfies finance and executive stakeholders.

06

Incrementality and lift measurement





Where experimentation provides the framework for proving cause and effect, incrementality is the question at its center: how many of these conversions would not have happened without our marketing efforts?

As such, incrementality and lift measurement have moved from a niche analytical practice to a mainstream priority. Over half of performance marketers now use incrementality testing, and 36% plan to increase their investment in it over the coming year.[40] The reason is that incrementality provides the closest thing to a definitive answer on whether marketing is driving real, additional business value.

What incrementality measures

Incrementality identifies conversions that would not have occurred without the campaign. The methodology typically involves reserving 10% to 20% of the target audience as a holdout group that receives no ads, then comparing conversion rates between the treatment and holdout groups.[41] The difference represents the incremental effect.

Doing so is important because not all conversions are created equal. Some people who see your ad and then convert would have converted anyway through channels like organic search or word of mouth. Attribution models count these conversions as marketing-driven; incrementality testing separates the organic baseline from the true marketing effect.

According to a Google and BCG study, 80% of US senior marketing analytics professionals report that implementing insights from incremental experiments had a high impact on revenue growth.[42] The value comes from making decisions based on actual impact rather than attributed credit.

Incremental ROAS

Incremental ROAS (iROAS) applies the incrementality framework to return on ad spend. Rather than calculating ROAS from all attributed conversions, iROAS measures only the revenue generated by conversions that would not have happened without the campaign, divided by the spend.

Benchmarks for what constitutes as a strong iROAS vary significantly depending on margins, growth goals and business model. But what's important is the 1.0 threshold. An iROAS above 1.0 means the campaign is generating more incremental revenue than it costs. Below 1.0, the campaign may still be driving some additional conversions, but not enough to cover the spend—which is a signal to optimize or scale back. Recent experiments have shown iROAS ranging from 253% to 1,609% across advertisers, but benchmarks vary significantly depending on margins, growth goals and business models.[43]

Think with Google has shared an illustrative example so you can see what lift on ROAS looks like: a beauty brand ran an incrementality test for Performance Max and learned its incremental ROAS was £6. For every £1 invested, the campaign generated £6 in revenue that would not have occurred without it, a 600% return based on causal evidence rather than attributed credit.[44]

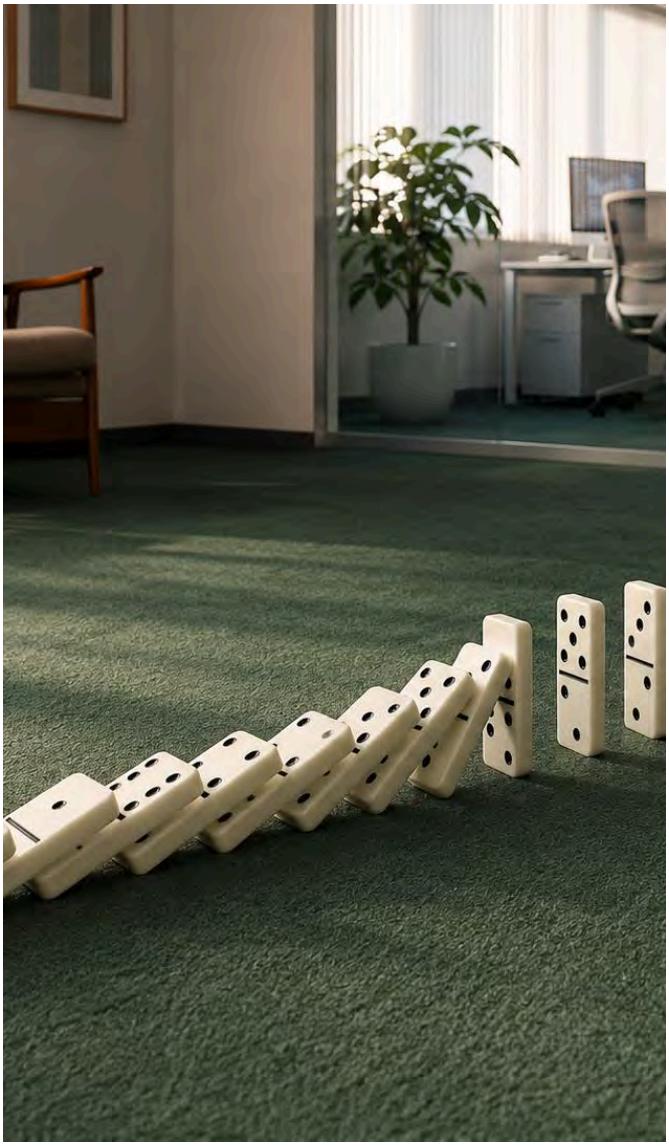
Conversion lift studies

Major platforms like Meta, Google and TikTok now offer structured Lift Studies to measure both brand lift (awareness and recall) and conversion lift (actual sales). Google, in particular, has worked to democratize these tools, lowering the minimum spend for incrementality experiments from $100,000 to just $5,000 and refining the methodology to deliver conclusive results 50% more frequently.[45]

However, these studies come with a walled garden constraint. A Google lift study measures Google; a Meta study measures Meta. Neither can fully capture the cross-platform interactions that define the modern, messy customer journey.[46]



Recent experiments have shown iROAS ranging from 253% to 1,609% across advertisers



Limitations of incrementality testing

Despite being the strongest evidence of causal impact, incrementality testing isn't without its challenges.

Withholding ads from 10% to 20% of your audience means foregoing potential revenue during the test period.[47] For large campaigns, this opportunity cost can be high, and it creates organizational resistance from teams measured on short-term volume.

Smaller advertisers may also struggle to reach statistical significance because of volume requirements.[46] If your campaign doesn’t generate enough conversions, the test may be inconclusive even if there is a real incremental effect.

Tests capture snapshots rather than trends. A single incrementality test tells you what happened during a specific window, but it doesn’t capture the compound, long-term effects of sustained marketing investment.[48]

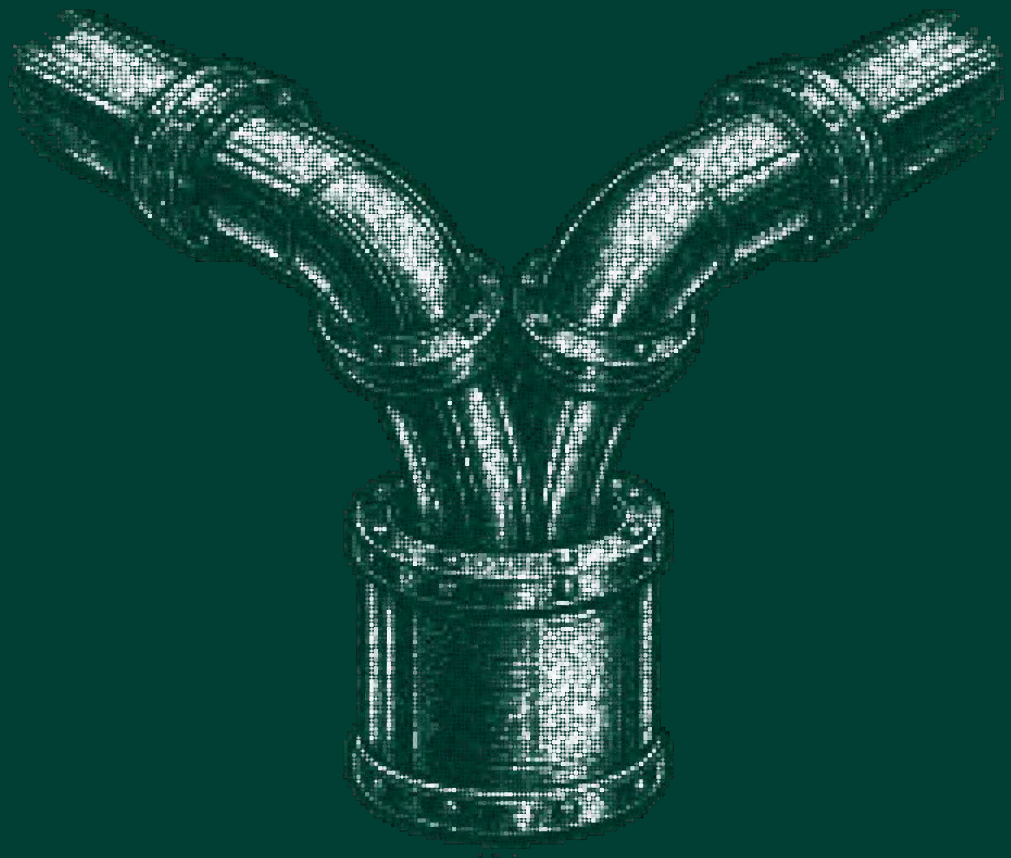
Next, platform-measured lift studies carry an inherent conflict of interest. When Google or Meta runs a lift study measuring the effectiveness of campaigns on their platform, they are both the measurement provider and the media vendor. They’re essentially grading their own homework.[46]

Finally, external factors like competitor activity, seasonality and spillover between test and control groups can skew results.[48] Incrementality testing is rigorous, but it isn’t immune to the messiness of real markets.

Use incrementality and lift measurement when you need causal evidence of marketing’s true impact. It’s useful for validating channel-level performance claims, building the business case for budget decisions and resolving discrepancies between attributed and actual results.

07

Marketing mix modeling (MMM)





Marketing mix modeling (MMM) has been around since the 1960s, but it's experiencing a resurgence right now. According to EMARKETER, 49% of marketers worldwide now use marketing mix modeling.[49] In the US, 56% of ad buyers were placing more focus on it heading into 2025.[49] The Drum reports that 76% of UK advertisers now use it.[50] And a Feld-M survey found that 61.4% of marketers want better or faster marketing mix modeling, with 30.1% believing it is the best approach for identifying business drivers.[51]

The primary reason for the revival of MMM is that privacy constraints have undermined the user-level tracking that attribution depends on. The industry needs a method that can evaluate cross-channel performance without cookies or device-level data. Marketing mix modeling fits that requirement because it works with aggregate data (total spend and total outcomes) rather than individual user journeys.

But the model's growing popularity also requires an honest assessment of what it can and can't do.

How marketing mix modeling works

At its core, MMM uses regression analysis to estimate the statistical relationship between marketing inputs (spend by channel, timing, creative variables) and business outcomes (sales, revenue, leads). The model controls for external non-marketing factors like seasonality, economic conditions, competitive activity, pricing and promotions, to isolate the estimated contribution of each marketing channel.

Unlike many modern machine learning applications, MMM does not operate on large datasets. Instead, it works with aggregated time-series data, often limited to a few hundred observations. This makes model design particularly sensitive to how many variables are included and how much independent variation exists in the data.

The output is a set of contribution estimates and ROI calculations by channel, which marketers use to evaluate past performance and inform future budget allocation.

Think with Google has reported that C-level leaders who placed high importance on marketing mix models were over two times more likely to exceed revenue goals by 10% or more.[52] The value lies in the strategic perspective; marketing mix modeling provides a cross-channel view that no single-platform tool can offer.

However, as BCG has observed, there is no one-size-fits-all measurement solution. Because marketing remains a blend of science and art, the most effective teams don't rely on a single model; instead, they use a combination of methodologies to triangulate the truth.[53] McKinsey also has noted that brand messaging remains one of the most critical determinants of ROI, often proving more important to a campaign's success than a precisely optimized marketing mix.[54] A perfect model can tell you where to spend your money, but it cannot fix a message that doesn't resonate.

Lag effects and carryover

One of the foundational concepts in marketing mix modeling, and one that distinguishes it from simpler analytical approaches, is adstock, also known as carryover or lag effects.

Advertising effects don’t stop the moment spend stops. A consumer who sees a TV ad this week might not convert until next month. A brand awareness campaign might influence purchasing decisions for months after it ends. If a model only compared this week’s spend to this week’s sales, it would systematically undervalue any activity with delayed impact, which includes most upper-funnel and brand-building work.

Marketing mix models address this by incorporating decay functions that distribute the impact of advertising over time. The model estimates how quickly the effect of spend diminishes after it occurs. Some channels have steep decay (the effect fades within days) while others have gradual decay (the effect persists for weeks or months). Brand campaigns tend to have longer carryover than performance campaigns, which is why MMM is often better suited to capturing upper-funnel effects than attribution models.

Understanding adstock doesn’t require mathematical sophistication. You only need to recognize that marketing is cumulative, that today’s impressions contribute to tomorrow’s conversions and that any model failing to account for this dynamic will produce misleading recommendations about where to invest.

MMM versus legacy approaches

To understand where the industry is heading, it's important to distinguish between the static reports of the past and the dynamic systems of today. The evolution of marketing mix modeling has changed how quickly teams can act on insights and, ultimately, how much they can trust the outputs.

For decades, legacy MMM was defined by its black box nature and glacial pace. These were typically massive, annual projects: expensive engagements managed by external agencies or consulting firms that often took months to complete. By the time a single set of results finally reached decision-makers, the data was outdated. Because the methodologies were proprietary and opaque, in-house teams had little visibility into the how behind the numbers, making the cadence far too slow to keep up with the volatile nature of modern digital auctions.

Modern MMM operates on an entirely different philosophy. Rather than a periodic post-mortem, it's treated as a living system. In December 2025, the IAB published a set of best practices for modernizing the discipline. They recommend that teams move toward weekly data refreshes and monthly model retrains to maintain accuracy.[55] The shift has been accelerated by the democratization of high-level tools; open-source libraries like Meta's Robyn and Google's Meridian have moved the power back in-house because they allow internal data teams to build, audit and maintain their own models.

And transition is critical because marketing conditions simply do not sit still. A model trained on last year's data is effectively a map of a different world; it cannot account for evolving audience behaviors or the rapid-fire changes in platform algorithms. By treating MMM as a continuous feedback loop rather than a static report, modern organizations ensure that their strategic big picture is always in sync with their tactical reality.

Limitations of marketing mix modeling

While MMM provides the macro view essential for strategic planning, it's not a crystal ball that can show you what will definitely happen. Like any statistical tool, it has structural limitations that teams must respect to avoid drawing the wrong conclusions.

Marketing mix modeling is inherently correlation-based not definitively causal. Because the model relies on historical data where spend was not randomly assigned, it can sometimes confuse coincidence with cause.[56] If you always increase search spend during your busiest sales month, the model might over-attribute success to that spend rather than the natural demand. This is compounded by collinearity. In a typical marketing plan, channels will rarely move in isolation. Teams scale up search, social and programmatic simultaneously during a product launch, which makes it difficult for the model to tease apart each channels individual impact, potentially muddying the ROI of one with the success of another.[56]

The model also needs sufficient variation in spend across each time period to produce stable estimates. With only a few hundred weekly observations and potentially dozens of variables, there is a tradeoff between complexity and reliability. Brands with limited spend variation or shorter histories will struggle to separate each channels contribution. For smaller organizations, the cost of traditional modeling has historically been a barrier, though open source tools and lighter-weight approaches have made the method increasingly accessible.

Even for larger brands, the model lacks the micro granularity required for daily or campaign-level optimization. Marketing mix modeling is an excellent tool for deciding how to balance your budget between YouTube and Meta next quarter, but it's virtually useless for deciding which specific ad set to pause this afternoon.

Finally, there is the gap between hindsight and foresight. As BCG has found, most traditional measurement focuses heavily on calculating the ROI of past spending, yet what marketing analysts truly crave are forward-looking simulations. In fact, 80% of marketers surveyed stated that robust scenario-planning capabilities, the ability to ask "What happens if I move 10% of my budget tomorrow?" would be the primary driver for wider MMM adoption.[57] Without that predictive element, a model is just a rearview mirror in a fast-moving car.



Use marketing mix modeling when you need a cross-channel view of marketing performance, strategic guidance on budget allocation across channels or a privacy-resilient alternative to user-level tracking. Pair it with incrementality testing to validate causal claims and with scenario planning to make the outputs forward-looking.

08

Saturation, diminishing returns and marginal analysis





Every marketing channel has a breaking point, a moment where the next dollar spent generates less value than the last one. Identifying exactly where that point lies is the analytical foundation of a sophisticated budget. While many teams focus on total volume, the most effective marketers focus on the edges of their spend.

We'll explore the three interconnected concepts that define these edges: **saturation** (the profitability limit), **diminishing returns** (the degradation of performance) and **marginal analysis** (the metric that reveals what your next dollar is actually doing).

Audience saturation

Saturation is the point where the next increment of spend no longer generates at least a dollar of return. Where marginal return falls below 1.0 and additional investment becomes unprofitable.



As spend increases on any given channel, performance typically follows a curve rather than a straight line. The first thousand dollars might generate 100 conversions, but the next thousand might generate only 80, and the thousand after that might generate 50.[58] Cost per acquisition rises as spend increases because you start running out of new people to reach, you show ads to the same people with increasing frequency and the most responsive audiences convert first.

And that pattern applies across virtually all marketing channels. It's a function of finite audiences and declining marginal attention, not a flaw in the campaign or the platform.

So the real question to ask isn't if saturation exists, but where it begins. Channels with larger addressable audiences saturate more slowly. Channels targeting niche segments saturate faster. And saturation points move over time as audiences evolve, competition changes and creative fatigue sets in.

Marginal ROAS, POAS and MER



Understanding marginal returns is essential for budget decisions, and it's a distinction that teams can overlook.

Return on ad spend (ROAS) is revenue divided by advertising spend. It's a commonly used metric for determining which campaigns drive the most revenue. Profit on ad spend (POAS) looks at actual profit rather than revenue, and includes factors such as product margin and variable costs in the calculation. It's a helpful metric for measuring channel efficiency when you have mixed-margin products. The marketing efficiency ratio (MER) is total revenue divided by total ad spend. It tells you how the entire budget performed in aggregate and is useful when making scaling decisions.

Marginal ROAS is the extra revenue generated by the next increment of spend divided by that incremental spend. It tells you what the next dollar is actually doing.

The problem is that ROAS can look healthy even when marginal dollars are losing money. Consider a channel where you spend $1,000 and generate $3,000 in revenue; the ROAS is 3.0, or 300%, meaning your campaign generates $3 in revenue for every $1 spent. So, you increase spend by $500, expecting to get the same 300% return. But revenue only increases by $200, not the $1500 you anticipated. Marginal ROAS on that incremental $500 is 0.4, or $0.40 per dollar. The extra spend is unprofitable even though the overall return still looks strong.

The same principle applies to POAS and MER. Marginal POAS shows you the marginal increase in profit from incremental spend. Marginal MER looks at when the next ad dollar spent will stop making you more money. When marginal MER starts trending downward, investing more marketing dollars will result in a decreasing return.

Teams that optimize on average ROAS, POAS or MER often overspend on saturated channels without realizing it. Without looking at the marginal versions of these metrics, spend keeps scaling while each additional dollar drives less and less revenue. If you're not measuring marginal differences, whether that's marginal ROAS, marginal POAS or marginal MER, the fact that the most recent dollars are dragging down performance gets buried in a healthy-looking average.

Optimal budget allocation

The ultimate goal of marginal analysis is to equalize marginal returns across your entire portfolio. If one channel is sitting at a marginal ROAS of 0.5 (saturated) and another is at 2.5 (undersaturated), you are leaving money on the table. The solution is to reallocate funds from the saturated channel, where the next dollar costs more than it returns, to the undersaturated one, where there is still headroom to perform.

In practice, this means using response curves from marketing mix modeling to identify channels approaching saturation and channels with remaining headroom. The reallocation requires shifting spend toward where the marginal return is highest.

As British supermarket train Waitrose described their measurement journey in The Drum, they learned that there isn't a single source of truth, but a suite of truths. The value comes from triangulating different sources to land on a number they could act on. They stopped focusing on digital attribution models and worked toward framing measurement in the context of business challenges.[59]

That framing applies directly to saturation analysis. The question isn't "what's our ROAS?" It's "where is the next dollar most productive, and where have we already passed the point of diminishing returns?"

Use saturation and marginal analysis when you need to optimize how much to spend on each channel, identify where additional budget will be wasted and make the case for reallocating spend from saturated to undersaturated opportunities.

09

Forecasting and scenario planning





Dashboards tell you what happened last month. Attribution models assign credit for conversions that have already occurred. The more valuable questions are forward-looking: what happens if we increase brand spend by 20%? Where should we invest the next thirty thousand dollars? What's the likely range of outcomes if we shift budget from paid search to connected TV?

Forecasting and scenario planning bridge the gap between "what happened" and "what should we do next." They transform measurement outputs from retrospective reports into tools for decision-making.

Planning the future by understanding the past

Once you understand how different channels contribute to outcomes through marketing mix modeling, you can simulate what would happen under different conditions.Calibrating the models with incrementality testing strengthens the results by grounding them in causal evidence.

And this is where measurement starts earning its keep with finance teams. According to BCG, 80% of survey respondents said that forward-looking simulations would widen their use of marketing mix models.[57] There is an appetite for models that help answer "what if?" questions before money is committed.

Yet only a third of respondents currently track the long-term impact of investments on brand equity.[60] This gap means organizations are making forward-looking budget decisions without accounting for the compounding effects of brand investment, which can take months to materialize.



Scenario modeling with MMM

Marketing mix modeling estimates response curves for each channel. Those curves show diminishing returns at different spend levels, which means you can simulate outcomes before committing budget.

The capabilities include "what if" analysis that tests the projected impact of budget increases or cuts before they happen, channel reallocation scenarios that identify where to shift spend for maximum impact and saturation forecasting that predicts when additional spend on a given channel will stop generating meaningful returns.

The goal isn't to predict the future with certainty. No model does that. The goal is to give decision-makers a range of likely outcomes so they can make informed trade-offs. A forecast that says "brand TV spend of $2M will generate between $5M and $7M in incremental revenue over 12 months, based on these assumptions and subject to these constraints" is far more useful than a single-point ROAS estimate with no range and no caveats.

As such, best-in-class simulations use a variety of effectiveness metrics and can isolate the impacts of investments with both short-term and long-term goals.[60]

Why simulations help with stakeholder buy-in

One of the biggest challenges in marketing measurement is organizational. Finance teams and executives distrust marketing's numbers because they've seen too many dashboards that don't connect to business outcomes.

Scenario planning changes the conversation. Instead of presenting a single ROI number and hoping for buy-in, marketers can show a range of outcomes under different assumptions. This makes the uncertainty explicit and gives stakeholders something they can engage with.

Only 11% of advertisers use shared KPIs between marketing and finance.[61] And CFOs remain especially skeptical of existing marketing ROI approaches because those approaches don't connect to financial planning cycles.[9]

When marketing can show finance a simulation like "if we invest an additional $500K in brand campaigns, we expect incremental revenue of $1.2M to $1.8M over 12 months, based on these assumptions," the conversation goes from debating if last quarter's ROAS was accurate to evaluating whether the projected return justifies the investment. A far more productive dialogue.

Connecting measurement to planning cycles

For these models to be effective, they must be embedded into the company's operational heartbeat. Sophisticated organizations refresh their models weekly or even daily so the outputs can feed directly into budget discussions. As MarTech suggests, if a model doesn't influence a decision, it's simply an expensive system providing no strategic direction.[62]

So, this also requires a cultural change. The analytics team must deliver data on a cadence that matches financial timelines, translating technical tables into the language of trade-offs and assumptions that executive stakeholders can actually use.

Limitations of forecasting

It is important to remember that forecasts are informed estimates, rather than ironclad commitments. They are based on historical patterns that can be invalidated by external factors.

Therefore, the best teams treat forecasts as living documents; they present them with clear confidence intervals and immediately update them when market conditions change. Using forecasting and scenario planning allows you to build a bulletproof business case, ensuring that your budget decisions are rooted in the best available evidence of what will drive growth tomorrow.

Use forecasting and scenario planning when you need to translate historical measurement outputs into forward-looking budget decisions, build the business case for marketing investment with finance teams or evaluate trade-offs between different allocation strategies before committing spend.

10

Financial and unit economics analysis





Marketing measurement and financial measurement have historically operated in parallel using different metrics, different timelines and different assumptions. This disconnect is a primary driver of the trust gap between marketing and the C-suite.

Bridging it requires marketing teams to adopt financial language and financial teams to understand the nature of marketing investments.

Connecting marketing to business outcomes

The IPA Bellwether report found that only 53% of marketers formally track financial return.[63] The other 47% are measuring activity, engagement or attributed conversions without translating those metrics into financial outcomes.

The IPA has also issued a cautionary note about the metric most teams default to: ROI is a potentially dangerous way to measure effectiveness in the marketing context because one easy way to boost short-term ROI is to cut the marketing budget.[64] When you reduce spend, attributed ROI improves in the short term because the remaining budget is concentrated on the most efficient (usually lower-funnel) activities. But you're depleting brand equity and future demand in the process. It becomes a trade that looks good on this quarter's dashboard and terrible on next year's P&L.

Key financial metrics

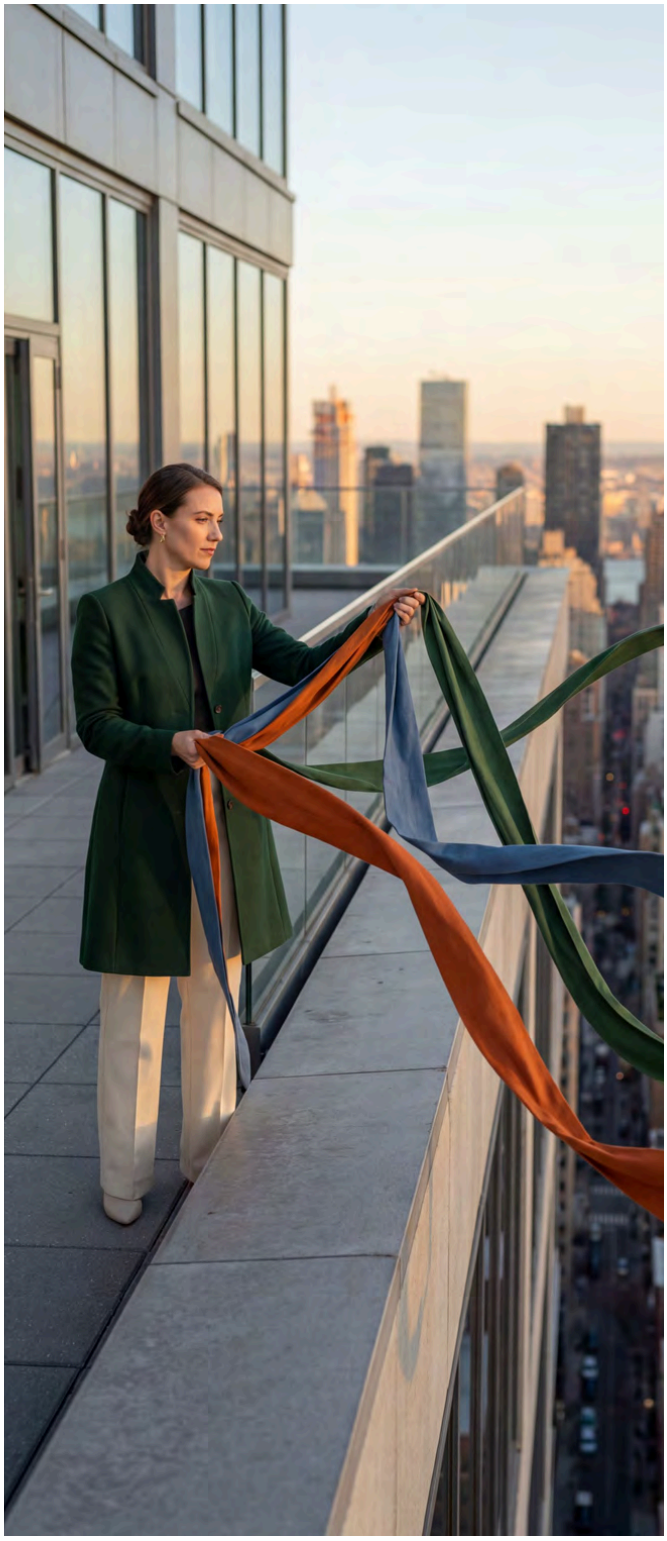
Several financial metrics help connect marketing activity to business outcomes.

Customer acquisition cost (CAC) measures the total cost of acquiring a new customer: marketing spend divided by the number of new customers acquired. It provides a unit economics view that's more tangible to finance teams than channel-level ROAS.

Customer lifetime value (CLV) estimates the total revenue or profit a customer generates over their relationship with the brand. The CAC-to-CLV ratio is one of the most important metrics for evaluating whether acquisition spend is sustainable. A ratio below 1:3 often signals that you're paying too much to acquire customers relative to what they're worth.

The payback period measures how long it takes for a customer's revenue to cover their acquisition cost. Finance teams care deeply about payback because it affects cash flow forecasting and capital allocation.

Contribution margin tells you what percentage of revenue from marketing-acquired customers remains after variable costs. A high ROAS on a low-margin product may contribute less to the business than a moderate ROAS on a high-margin product.



Building the business case for marketing

According to BCG/Google research, as reported in The Drum, only 11% of advertisers use shared KPIs between marketing and finance, and only 18% actively bring effectiveness insights into cross-team decision-making.[61]

When marketing and finance operate with different metrics, every budget conversation becomes a translation exercise. Marketing talks about ROAS, impressions and attributed conversions. Finance talks about contribution margin, payback period and cash-flow impact. The numbers don't translate, and trust erodes.

The solution is not for marketing to abandon its metrics, but to build a financial translation layer. When a CMO can present performance through the lens of unit economics, connecting channel-level data to customer acquisition cost (CAC), customer lifetime value (CLV), payback periods and incremental contribution margin, the dynamic of the conversation changes. You are no longer justifying spend; you are presenting a case for optimizing an investment.

The expectation for this mindset is already reflected in the C-suite. Gartner's 2025 data shows that 61% of CMOs now see their departments viewed as profit centers, up from 54% just a year prior.[3] Marketing is no longer a cost to be managed, but a primary engine for revenue. Meeting this expectation requires measurement that speaks the language of the business.

Ultimately, grounding your strategy in financial and unit economics analysis allows you to evaluate if your customer acquisition is truly sustainable. By building a budget case based on payback periods and incremental contribution rather than just attributed conversions, you provide the finance team with the evidence they need to treat marketing as a predictable, scalable asset.



Use financial and unit economics analysis when you need to connect marketing performance to business outcomes in terms that finance teams understand, evaluate whether customer acquisition is sustainable or build a budget case grounded in contribution and payback rather than attributed conversions.

11

Measuring brand effects





The greatest challenge in marketing measurement is that the most powerful effects are usually the ones that refuse to show up in this week's dashboard. Brand equity, built through awareness, consideration and preference, is a slow-burning asset. It accumulates over months and years, influencing purchase decisions long after the original ad exposure has faded from memory. And it creates a measurement trap; the very activities that drive the highest long-term value look the most inefficient when viewed through the narrow lens of short-term metrics.

Measuring brand effects is one of the hardest unsolved problems in the industry. The methods exist in theory, but applying them reliably remains a genuine challenge for most organizations. Understanding what's available and what's realistic is important for any team trying to make the case for sustained brand investment.

What brand lift measures

Brand lift studies measure changes in consumer perceptions (awareness, ad recall, message association, favorability, consideration and purchase intent) by comparing an exposed group to a control group.

Brand awareness itself operates at several levels. Unaided awareness measures whether consumers can name your brand unprompted when asked about a category. Aided awareness measures whether they recognize you from a list. Top-of-mind awareness, the first brand a consumer thinks of in a category, is the most competitively valuable form.

Platform lift studies from Google and Meta provide a structured way to measure these changes within their ecosystems. The methodology compares survey responses from users who were exposed to the campaign against a control group.

The limitation of brand lift studies is that they measure stated intent rather than actual behavior. A consumer who says they are "more likely to consider" your brand after seeing an ad may or may not change their purchasing behavior. So be aware that lift studies provide a directional signal, not a guarantee of commercial impact.

The long and short of brand building

The most influential body of research on brand measurement comes from Les Binet and Peter Field's analysis of the IPA Effectiveness Databank. Their central finding is that the optimal balance between brand building and sales activation is approximately 60:40.[65]

The reasoning is grounded in how each type of activity works. Activation (search, retargeting, promotions and direct response) captures existing demand. Its effects are immediate but decay quickly. Brand (TV, broad social, sponsorship and content) builds mental availability. Its effects compound over time but are harder to measure in the short term.

The IPA data support several related findings. Long-term brand investment delivers approximately twice the profit of short-term activation alone. Also worth noting: emotional campaigns are almost twice as likely to drive top-box profit growth compared to purely rational messaging.[65]

The implication for measurement is clear even if acting on it is difficult: if you only measure short-term ROAS, you will undervalue brand activity and overinvest in activation. Over time, this depletes the brand equity that makes activation effective, and you end up falling into a cycle that produces short-term metrics that are improving and declining long-term results.



Measuring brand effects over time

Because brand effects unfold over months and years, measuring them requires approaches designed for longer time horizons. None of these are easy, and most organizations are still working to implement them well.

Tracking studies, which survey consumers monthly or quarterly to monitor changes in awareness, consideration and preference, provide a longitudinal view that point-in-time lift studies miss. But they require consistent investment and disciplined methodology over extended periods.

Marketing mix modeling can incorporate brand metrics as input variables to estimate how brand health contributes to sales. But this requires consistent measurement of brand metrics alongside business outcomes.

Econometric lag analysis can estimate the time delay between brand investment and sales impact. This lag can be six months or longer, which creates an accountability challenge in organizations with quarterly reporting cycles.

Limitations of brand measurement

Brand measurement relies on surveys and modeling rather than conversion tracking. The data is softer, the confidence intervals are wider and the results take longer to materialize.

Lift studies measure what people say, not what they do. The correlation between stated purchase intent and actual purchasing behavior is positive but imperfect.

Brand effects can take six to twelve months to materialize, creating accountability challenges with short-term reporting cycles. A brand campaign that will drive significant commercial impact over the next year may look like a poor investment on this quarter's dashboard.

And the 60:40 rule, while supported by extensive data, is a guideline rather than a prescription. The right ratio varies by category, brand maturity, growth stage and competitive dynamics. Emerging brands in competitive categories may need to weigh brand investment even higher. Mature brands with strong awareness may benefit from a different balance.

Use brand measurement when you need to capture long-term effects that short-term metrics miss, evaluate the balance between brand building and sales activation or build the case for sustained investment in brand equity. Recognizing that measuring these effects remains one of the industry's hardest challenges, and that imperfect measurement is better than no measurement at all, is itself a sign of measurement maturity.

12

Audience and targeting effectiveness



Measurement doesn't just tell you whether your marketing worked. It should also tell you if you reached the right people. Audience targeting effectiveness is the bridge between media strategy and measurement outcomes. If you're reaching the wrong audience, no amount of creative genius or analytical wizardry can save your ROI.

Measuring audience targeting

The data on global targeting accuracy is, frankly, sobering. According to Nielsen Digital Ad Ratings, the average "on-target" rate for digital ads in North America sits at just 65%.[66] That means that over one-third of digital advertising budgets are reaching people outside the intended audience. Nielsen's research also indicates that unique reach is a strong indicator of ROI performance. Wasted impressions on the wrong audience don't just fail to convert; they actively dilute campaign efficiency.

The Nielsen One Ads Meta Analysis also reveals what happens when marketers try to plug these leaks.[67] Between 2022 and 2024, the share of campaigns using closed (highly specific) targeting doubled from 15% to 30%. However, reaching the right person is only half the battle; you also have to reach them at the right cadence.

The analysis found a massive spread in campaign frequency, ranging from a negligible 1.01 to an exhausting 48.06. This spread shows a gap in frequency management; many brands are either whispering into the wind or shouting until their audience hits the mute button. Furthermore, only 18% of digital campaigns utilize multiple platforms or publishers, suggesting that the vast majority of digital marketing still happens within the safety (and limitations) of single-platform silos.

When a team manages this manual complexity perfectly, the results are transformative. A Samsonite India case study showed that by maintaining an 85% on-target rate and optimizing frequency to a precise 3.38, the brand achieved a 3x cost optimization.[67] But as the digital ecosystem grows more complex, achieving this level of manual control is becoming increasingly difficult for the average advertiser.



AI-driven targeting evolution

The mechanics of ad delivery have been rebuilt around automated signal processing. Platform algorithms are increasingly determining who sees ads, based on conversion data and creative signals rather than audience segments defined by the advertiser.

Meta's Advantage+ campaigns are now the primary growth engine for 65% of US advertisers.[68] Meta advertisers who consolidate campaigns into fewer campaign structures are seeing a 32% drop in cost per acquisition.[69] Also, broad targeting with strong creative and conversion optimization shows 113% higher ROAS compared to lookalike audiences.[69]

Google AI Max, launched in May 2025, extends this pattern. Advertisers upload creative assets, conversion data and landing pages while Google's AI handles targeting through predictive modeling that matches ads to audience intent signals. Early adopters reported up to 27% higher conversion volume with no increase in spend.[70]

When the platform's algorithm decides who to target, defining audiences and measuring on-target delivery becomes less relevant. The measurement question you need to ask goes from "did we reach the right people?" to "did the algorithm find more valuable people than we would have targeted manually?"

McKinsey reports that 71% of consumers now expect personalized interactions.[71] Algorithmic targeting can help here, but it's important to track the impact rather than assuming the AI is doing its job. For the modern analyst, the task is to evaluate whether this AI-driven personalization is driving true incremental value, or if it's simply creating a smoother user experience that doesn't actually change the final commercial outcome.

First-party data and privacy-safe targeting

As third-party tracking declines, first-party data becomes the foundation for targeting. Data clean rooms, privacy-safe environments where partners can match audiences and measure outcomes without exposing individual-level data, are gaining adoption, led by the IAB Tech Lab.

Google's Privacy Sandbox APIs represent the platform-level response. The Topics API supports interest-based advertising without sharing the specific sites a person has visited. The Protected Audience API supports remarketing and custom audiences via on-device auctions designed to limit cross-site tracking.

These approaches preserve some targeting capability in a privacy-constrained world, but they introduce new measurement challenges. Aggregated targeting data is harder to evaluate than user-level data, and the effectiveness of privacy-safe alternatives relative to legacy approaches is still being established.

Use audience and targeting effectiveness analysis when you need to evaluate if your campaigns are reaching the right people, assess the impact of AI-driven targeting changes, measure frequency and reach optimization or quantify the waste from off-target delivery.

13

Uplift and persuasion modeling





Most marketing measurement treats the audience as a single group and asks, “Did the campaign work?” Uplift modeling asks a more specific question: for which individuals did the campaign actually change behavior?

This is an advanced technique, but the concept is intuitive. Not everyone responds to marketing the same way. Some people will buy regardless. Some will never buy regardless. The people who actually changed their behavior because of the marketing, the persuadables, are the ones who matter for incremental impact.

Defining uplift modeling

Uplift modeling, also called incremental modeling or true-lift modeling, estimates the causal impact of a marketing intervention on individual outcomes. The goal is to identify customers who convert only because they were marketed to, as opposed to those who would have converted anyway.

At the heart of uplift modeling is a statistical paradox first articulated by Donald Rubin in 1974. Rubin noted that it's impossible to observe the causal effect of a treatment on a single unit.[72] You can't simultaneously observe what a customer does when they see your campaign and what the same customer does when they don't. This is the fundamental problem of causal inference, and uplift modeling attempts to estimate what can't be directly observed.

The four customer segments

Uplift modeling classifies customers into four conceptual segments based on how they respond to marketing.

Sure Things will convert regardless of marketing exposure. They have high baseline conversion rates and low or no uplift from campaigns. Targeting them with ads is wasteful because you're paying to reach people who would have converted for free.

Lost Causes won't convert regardless of marketing. They have low baseline conversion rates and no response to campaigns. Spending on them generates no return.

Persuadables convert only when marketed to. They have low baseline conversion rates but respond to campaigns. These are the customers who generate incremental value, and they are the primary target of uplift modeling.

Sleeping Dogs are actively turned off by marketing. Their conversion probability actually decreases when they see a campaign, giving them a negative uplift. Marketing to them is worse than wasteful; it's counterproductive.

If you can identify the Persuadables and avoid spending on Sure Things, Lost Causes and Sleeping Dogs, you allocate budget far more efficiently than blanket targeting allows.

Uplift modeling methods

Several modeling approaches have been developed to estimate uplift at the individual level.

The **Two-Model** approach trains separate prediction models on people who saw the campaign (treatment) and people who didn't (control). For each person, you subtract the two predictions to estimate the campaign's extra impact.[73]

Causal Forest uses a random-forest-style method that groups people into segments where the campaign effect differs. It averages many causal trees to produce more stable uplift estimates.[74]

Meta-Learners offer several variants.[75] The S-learner uses a single model that includes campaign exposure as an input. The T-learner trains separate models for treated and control groups. The X-learner first estimates the missing counterfactual outcomes and then learns uplift, which is helpful when treatment and control group sizes are uneven. The R-learner removes baseline patterns before estimating uplift.

Each approach makes different trade-offs between simplicity, accuracy and data requirements. The right choice depends on data volume, the degree of treatment effect heterogeneity and the computational resources available.

Practical applications

The most direct application is targeting. If you can score individuals by predicted uplift, you can prioritize spend on Persuadables and suppress ads for Sure Things and Sleeping Dogs.

Personalization extends this logic to selecting the best message, offer or channel for each individual when you have multiple treatments to choose from.

Cost optimization combines uplift estimates with cost data to choose actions based on incremental value minus cost. As an example, you might decide whether a discount, a paid media impression or an email is the most efficient way to move a specific customer.

Limitations

The fundamental validation problem is the same one that defines uplift modeling's purpose: because you can never observe the true uplift for any individual, you can't validate the model the way you would a standard prediction task.[76] You're estimating a quantity that is inherently unobservable at the individual level.

Evaluation metrics like Qini curves, uplift curves and the Area Under the Uplift Curve provide relative comparisons between models, but they don't confirm that estimated uplift values match reality at the individual level.

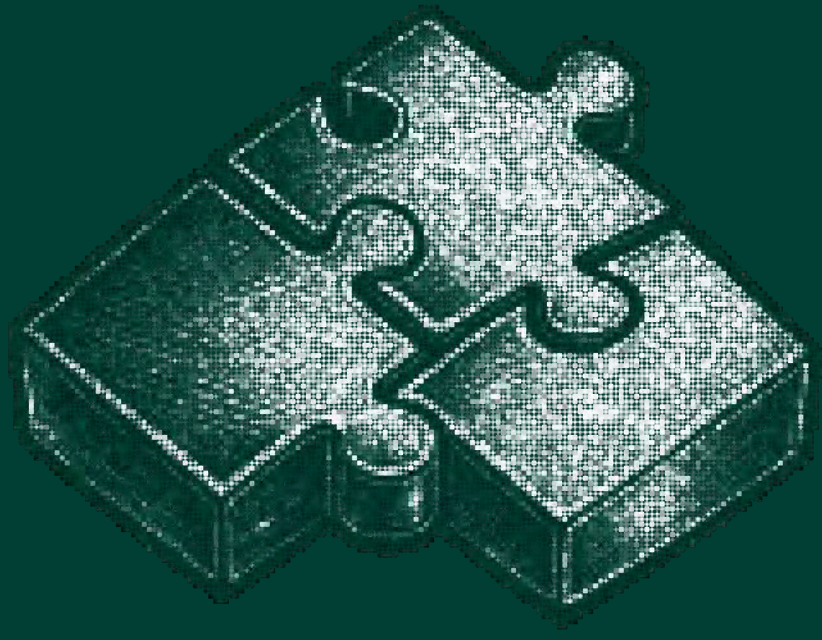
Uplift modeling also requires experimental data; you need a randomized treatment and control group to train the models. And the models are sensitive to data quality, treatment assignment integrity and the assumption that the experimental conditions reflect real-world deployment.



Use uplift and persuasion modeling when you need to identify which customers are truly persuadable by your marketing, optimize targeting to avoid wasting spend on people who would convert anyway or personalize campaigns based on individual-level causal response estimates.

14

Qualitative and mixed-methods analysis





Quantitative measurement tells you what happened and, with the right methods, whether your marketing caused it. What it can't easily tell you is why.

Why did that creative resonate? Why did customers in one segment respond differently than another? Why did the brand campaign generate awareness but not consideration? Answering "why" questions requires qualitative research, and integrating qualitative insights with quantitative measurement produces a more complete picture than either approach alone.

The role of qualitative research in marketing measurement

Qualitative research answers "why" and "how" questions rather than "what," "where" and "when." It uncovers motivations, emotions and experiences that quantitative data can't reveal. And it provides context to numbers and helps marketers understand why a metric moved rather than just that it moved.

According to Harvard Business School, marketing success requires a balance between quantitative analysis and qualitative judgment. Essential qualitative factors include the volume of a channel and the quality of leads, including the long-run value of customers acquired through different channels.[17]

Mixed-method research designs

To move beyond gut feeling, sophisticated teams use mixed-method research (MMR) to formalize how they blend the two worlds. Greenbook identifies three primary designs.[77]

Explanatory Sequential Design starts with quantitative data to identify trends and patterns, then follows with qualitative research to explore those trends in depth. This is useful for understanding why a survey found a particular result.

Exploratory Sequential Design begins with qualitative research to uncover themes and hypotheses, then quantifies those themes through surveys or behavioral data. This is useful for new product development or entering unfamiliar markets.

Convergent Parallel Design conducts both types of research simultaneously and integrates results during the analysis phase. This is useful when you need comprehensive findings on a tight timeline.

Qualitative research methods for marketers

Focus groups bring a small group of four to twelve participants together with a moderator to discuss a topic. They are useful for observing how participants interact with ideas and with each other, and they can inform the design of larger quantitative studies.[78]

In-depth interviews (IDIs) provide a one-on-one setting for detailed personal insights. They capture individual perspectives and experiences that might not surface in a group dynamic. The typical duration is 30 to 90 minutes.[79]

Ethnographic research involves observing how consumers interact with products in real environments. Watching from a distance can reveal usage patterns that aren't visible in conversations because consumers often use products in ways the brand hadn't imagined.[80]



Emotional testing and creative effectiveness

Some creative pre-testing methods measure emotional response to give an early signal of whether an ad is likely to drive short-term response and long-term brand building.

System1's methodology, for example, produces a Spike Rating (intended to indicate short-term sales potential based on the intensity of emotion and strength of branding) and a Star Rating (intended to indicate long-term brand-building potential based on measured emotional response).[81]

Emotional signals can provide directional guidance on creative effectiveness before a campaign launches. They should be triangulated with other measurement methods, including lift studies and business outcomes, to assess if the emotional response translates into commercial impact.

Beyond qualitative pre-testing, quantitative creative evaluation methods, such as attention metrics, recall scores and engagement benchmarks, can provide scalable signals about creative performance. And as McKinsey has noted, brand messaging is one of the most critical determinants of ROI, often more important than a precisely optimized media mix.[54] The creative itself matters as much as where you place it.

Limitations

Qualitative research involves small sample sizes that may not represent the broader population. Interpretation varies by researcher, and participants may not always report their motivations accurately, either because they don't understand their own behavior or because social dynamics influence their responses.

Furthermore, rich qualitative insights don't translate directly into media spend decisions. Knowing why a creative resonated doesn't automatically tell you how much to invest.

And good qualitative research also takes time. Focus groups and ethnographic studies operate on longer timelines than digital analytics, which can create friction in organizations that demand rapid insights.

Lastly, there's the challenge of integration; quantitative tells you what is happening at scale while qualitative fills the gap by revealing why people act the way they do. Combined, they create a richer narrative with enhanced credibility.[77] The challenge is to design research that brings both together rather than treating them as separate workstreams.

Use qualitative and mixed-methods analysis when you need to understand why metrics moved, not just that they moved. It's useful for creative development, customer experience improvement, brand perception research and explaining quantitative findings to stakeholders.

15

Measurement maturity and operating models



Organizations know they need better measurement. However, few know where to start or how to assess whether they're making progress. Measurement maturity frameworks provide a map for navigating from basic reporting to sophisticated, multi-method systems.

Assessing your current state

Step one to measurement maturity: honestly assess your current state. LiveRamp's four-stage maturity model provides a technical and organizational roadmap for this journey.[82]



At the Nascent stage, teams rely on siloed reporting from individual platforms, with limited integration and no cross-channel view. Measurement is reactive and focused on individual campaign reports rather than strategic insights.

At the Building stage, organizations begin integrating data sources, implementing consistent KPIs across channels and running initial tests. There's a recognition that platform-reported metrics don't tell the full story.

At the Developing stage, teams deploy multiple measurement methods like marketing mix modeling, incrementality testing and attribution. They begin connecting them to business outcomes. Data infrastructure supports regular refreshes and scenario planning.

At the Leading stage, measurement operates as a continuous system. Methods are triangulated, models are refreshed frequently, insights feed directly into planning cycles and there's organizational alignment between marketing, finance and executive leadership.

The maturity gap

While the technical roadmap is straightforward, the reality for most organizations is a disconnect between what they know and what they do. Companies are scaling their technical capabilities at a rate their organizational culture simply cannot handle. According to BCG, 76% of advertisers across Europe, the Middle East and Africa remain stuck in the early "Nascent" or "Building" stages, and of the organizations that invested in experiments, 68% failed to use the results to change a single decision.[83]

The stats suggest the primary barrier is due to a lack of governance. Organizations are buying the engine but forgetting to build the steering wheel.



BCG's four accelerators

To bridge this gap, BCG identifies four accelerators that separate leaders from laggards. These are divided into technical muscle and organizational willpower: [84]

- **The first-party data virtuous cycle** involves collecting, organizing and activating owned data to improve targeting and measurement over time. Each cycle of data collection improves the next iteration's accuracy.
- **End-to-end measurement with predictive models** means deploying marketing mix modeling, incrementality testing and scenario planning as an integrated system rather than isolated projects.
- **Agile performance loops** create feedback mechanisms that connect measurement insights to execution decisions on a rapid cadence: weekly or monthly rather than annually.
- **New skills and resources** refer to building the analytical talent and organizational capabilities needed to design, maintain and interpret measurement systems.

An interesting finding from BCG's research is that technical accelerators have approximately twice the impact of organizational accelerators on business outcomes. Investing in data infrastructure and modeling capabilities generates more measurable value than organizational restructuring alone.

But this doesn't mean organizational factors are completely irrelevant. As Joanna Stringer of BCG told The Drum, if your finance team doesn't believe, your insights might as well be dust.[85]

Building the operating model

Gartner has observed that marketing is entering a new era of productivity where CMOs are turning to data and AI to unlock new efficiencies. The top productivity driver Gartner identifies is leveraging data, analytics and measurement to optimize performance.[3]

Building this into an operating model requires several elements.

Clear ownership: someone is accountable for the measurement system, not just individual analyses.

Defined cadence: models are refreshed on a schedule that aligns with planning cycles.

Integration with decision-making: outputs are structured for the audiences who use them, whether that's a media team optimizing spend or a CFO evaluating investment options.

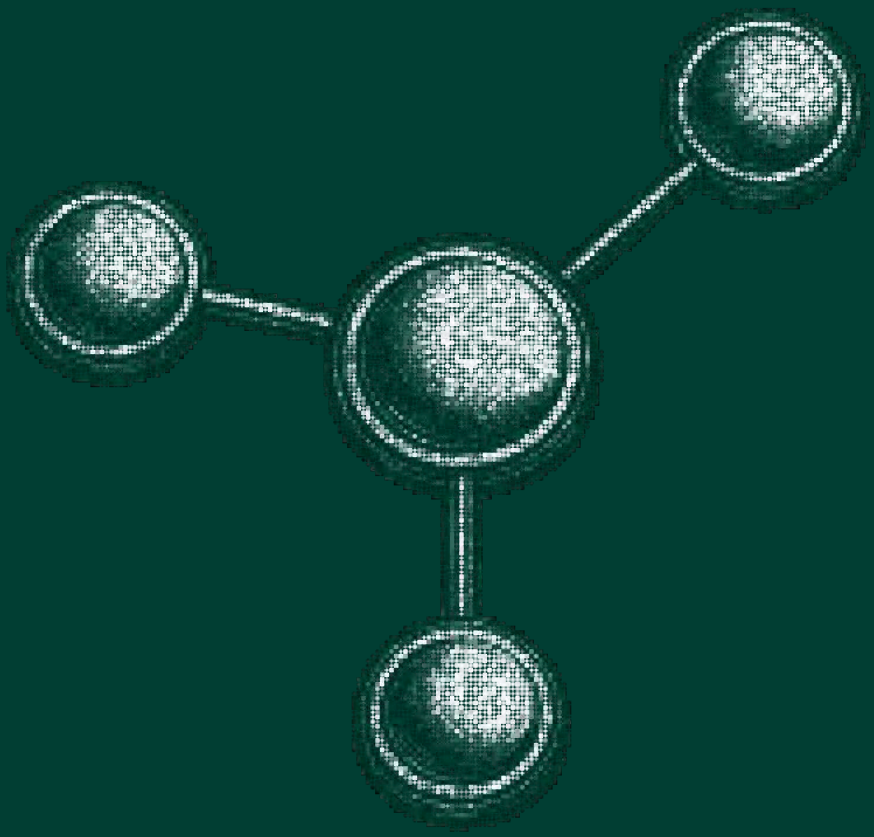
Continuous learning: the system improves as more data flows through it and as experiments validate or challenge model assumptions.

When these four pillars are in place, you have the ultimate productivity hack for the modern CMO: a measurement system that doesn't just count the money, but actively helps you make more of it.

Use measurement maturity assessment when you need to evaluate where your organization stands, identify the highest-impact investments in capabilities or build a roadmap for advancing from reactive reporting to proactive, multi-method measurement systems.

16

Advanced measurement and triangulation





If there is one idea that has defined the evolution of marketing measurement over the past decade, it's triangulation: the practice of combining multiple measurement methods to validate findings, fill gaps and produce more reliable insights than any single approach can provide.

No single measurement method answers all questions. Attribution assigns credit but doesn't prove causation. Incrementality testing establishes causation but captures only a snapshot. Marketing mix modeling provides strategic allocation guidance but depends on observational data. Each method has blind spots, and the blind spots of one method are often the strengths of another.

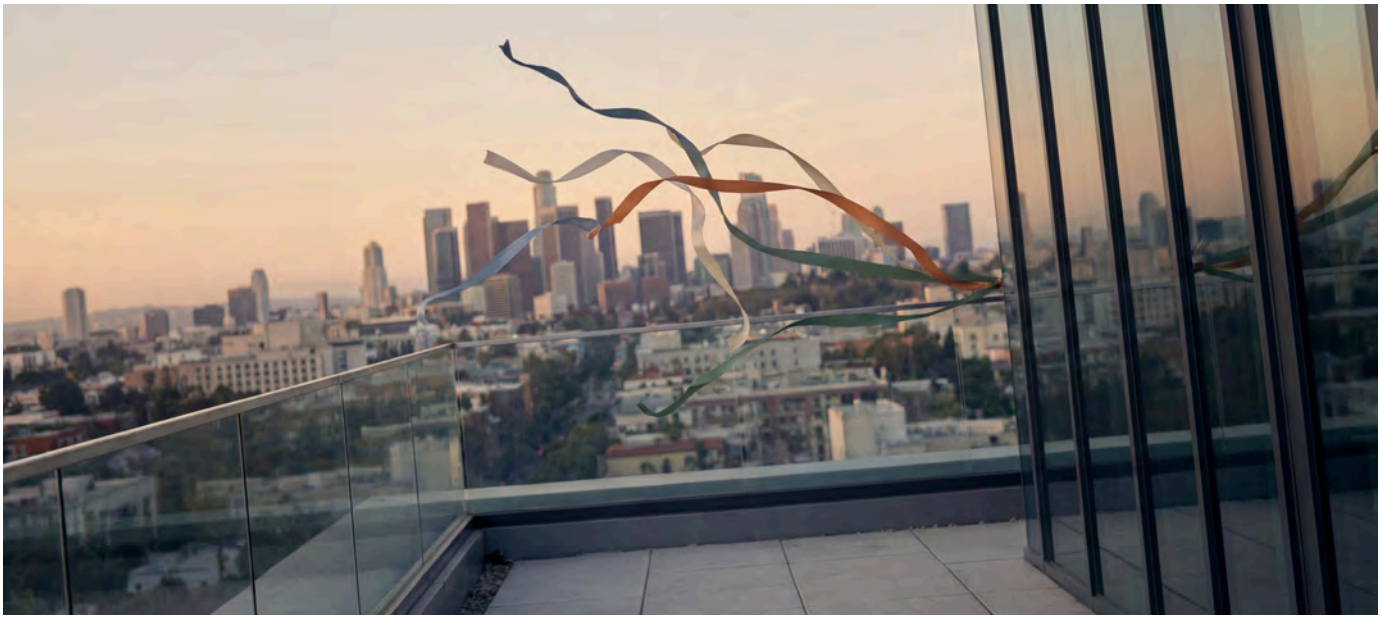
Why triangulation became the standard

Google's Modern Measurement framework articulates triangulation through a three-part structure. Attribution provides micro-level credit assignment across touchpoints. Incrementality provides meso-level causal impact measurement. Marketing mix modeling provides macro-level strategic allocation guidance.

Think with Google has stated that truly understanding media effectiveness requires a multi-step measurement setup that includes marketing mix modeling, incrementality testing and attribution, and that this combination unlocks the highest ROI for marketing dollars.[86] But they also acknowledge that implementing this layered approach is difficult, particularly since there is no one-size-fits-all measurement solution.

BCG extends this into a four-legged stool framework: marketing mix modeling for strategic allocation, incrementality experiments for causal validation, customer insights for behavioral understanding and execution metrics for operational monitoring. BCG describes this framework as providing a powerful and comprehensive method to triangulate and validate marketing ROI.[87]

The conceptual logic is sound. But mature organizations have learned that triangulation isn't simply running multiple models and hoping the answers converge. In practice, models frequently produce conflicting signals, and reconciling those signals requires structured processes.



How modern systems reconcile competing signals

When marketing mix modeling suggests one channel is underperforming while an incrementality test shows it drives significant lift, you have a conflict. Simply averaging the two answers is intellectually lazy. Discarding one is wasteful. The discipline lies in understanding why the models disagree and using structured reconciliation to arrive at a defensible answer.

Mature measurement systems address this challenge through several mechanisms.

These mechanisms transform triangulation into an operational discipline. Measurement is a system that requires governance, not a dashboard or a single model that runs itself.



1. Informed starting assumptions

Informed starting assumptions, sometimes called Bayesian priors in technical contexts, allow models to incorporate directional insights from other methods as starting points that are then updated by historical data. For example, a marketing mix model might use incrementality test results to set initial expectations for a channel's contribution, rather than starting from zero. This creates coherence across methods without forcing agreement.

2. Multi-criteria model evaluation

Multi-criteria model evaluation moves beyond selecting the model with the best statistical fit in isolation. Instead, mature systems evaluate models against multiple criteria simultaneously: statistical accuracy, business consistency, alignment with experimental results and reasonableness of implied channel contributions. A model that fits the data perfectly but suggests an implausible allocation (such as assigning 80% of contribution to a channel that receives 5% of spend) should raise flags, even if its R-squared (the predictive accuracy of the statistical model) is high.

3. Calibration and governance

Calibration and governance represent the organizational side of reconciliation. Strong measurement systems include structured processes for comparing marketing mix modeling outputs with experimental findings and granular attribution signals. When discrepancies arise, there is a defined process for investigating why and for deciding which signal to prioritize for a given decision.

Limitations

Triangulation requires more budget, more expertise and more organizational coordination than any single method. Running multiple methods means maintaining multiple data pipelines, analytical capabilities and governance processes.

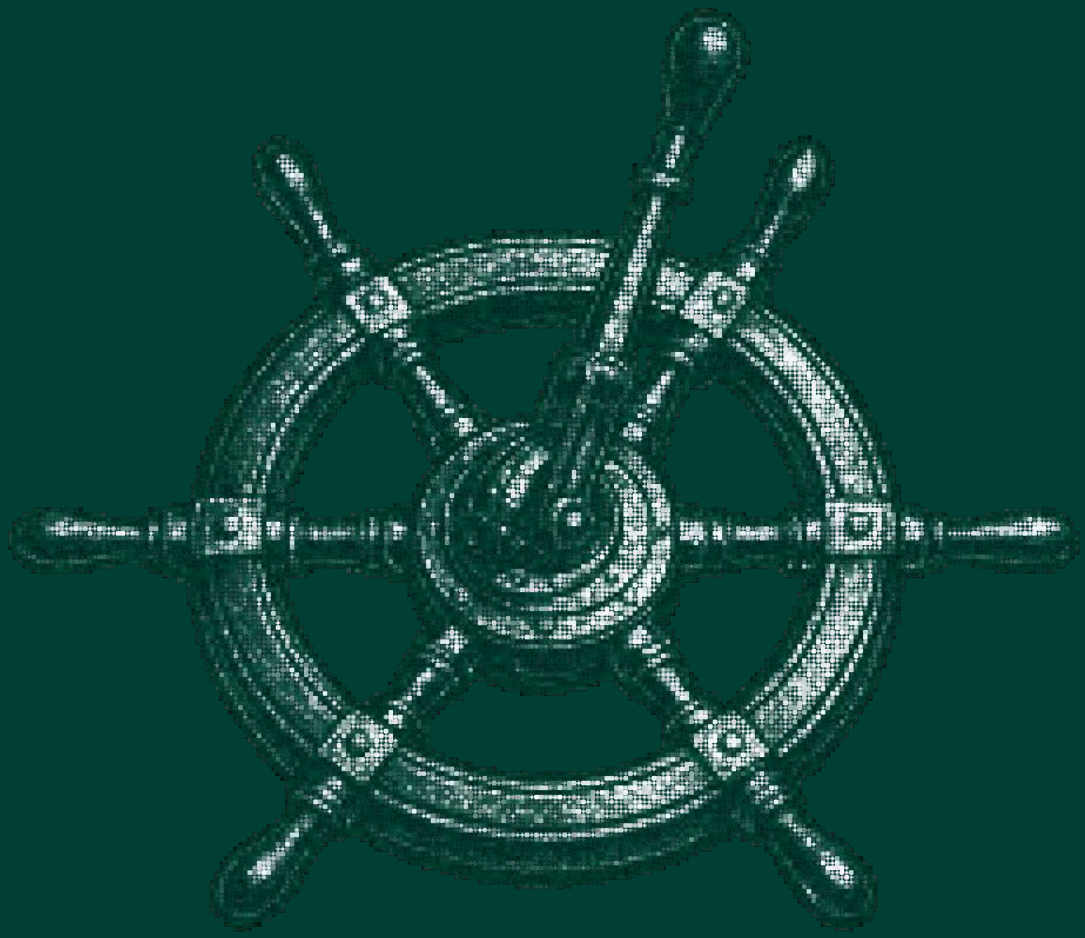
Conflicting signals still require human judgment to interpret. Triangulation doesn't eliminate ambiguity, but it does provide a better-informed basis for navigating it.

And some stakeholders find the absence of a single source of truth unsatisfying. The instinct to ask "which model is right?" persists even in organizations that intellectually accept that no single model is complete. Managing this expectation is part of the leadership challenge.

Use advanced measurement and triangulation for strategic budget decisions that require causal validation, cross-channel allocation, stakeholder alignment and governance of competing analytical signals.

17

Econometrics and causal inference





If marketing measurement is a vehicle, econometrics is the physics of the engine. Many of the sophisticated tools marketers use, such as MMM, geo-experiments and saturation analysis, are simply different ways of applying econometric techniques.

For a senior leader, the goal is to understand the machine's logic. The following section will give you the vocabulary to ask better questions of your analytics teams without needing a degree in statistics.

What econometrics is and why it matters for marketers

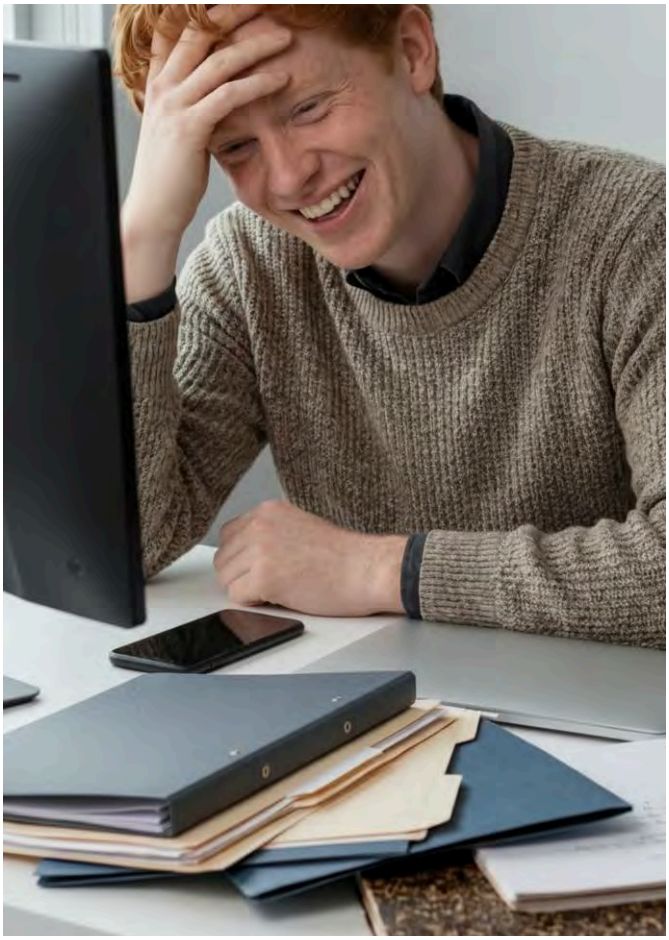
In plain English, econometrics is the application of statistical methods to economic data to test a hypothesis. In marketing, it is the filter that strips away the noise of the outside world, like weather, the economy or competitor moves, to isolate the signal of your marketing's actual impact.

The core challenge in marketing measurement is answering a question that can never be observed directly: what would have happened if we hadn't run that campaign? This is the counterfactual, and as Varian noted in his influential PNAS paper, we can never see what sales would have been in a world where we didn't advertise, so we have to estimate it.[25]

Every measurement method is, at its foundation, an attempt to estimate the counterfactual using different data and different assumptions. Understanding this common foundation helps explain why different methods can produce different answers, and why that isn't necessarily a problem.

The counterfactual problem

The problem with simple analysis is that if you plot sales against ad spend, you will find a relationship. But that relationship doesn't prove that ads caused the sales.



Varian uses an instructive example: cities with a high interest in hockey tend to have both high advertising spend on hockey equipment and high hockey equipment sales. A simple regression would show a positive relationship between ad spend and sales, but the underlying cause is hockey interest, not advertising. The regression overestimates the effect of ads because it's capturing underlying demand rather than the causal impact of the marketing.[25]

So this is the endogeneity problem. Advertising spend isn't randomly assigned; companies spend more where they expect demand to be higher, in seasons when sales naturally increase and on channels where they've seen past success. Without accounting for these factors, statistical relationships between spend and outcomes will be biased.

Companies spend more where they expect demand to be higher — so of course ad spend and sales correlate.

Methods for estimating the counterfactual

Different situations call for different methods. Angrist and Pischke's framework from Mostly Harmless Econometrics outlines the primary approaches.[88]

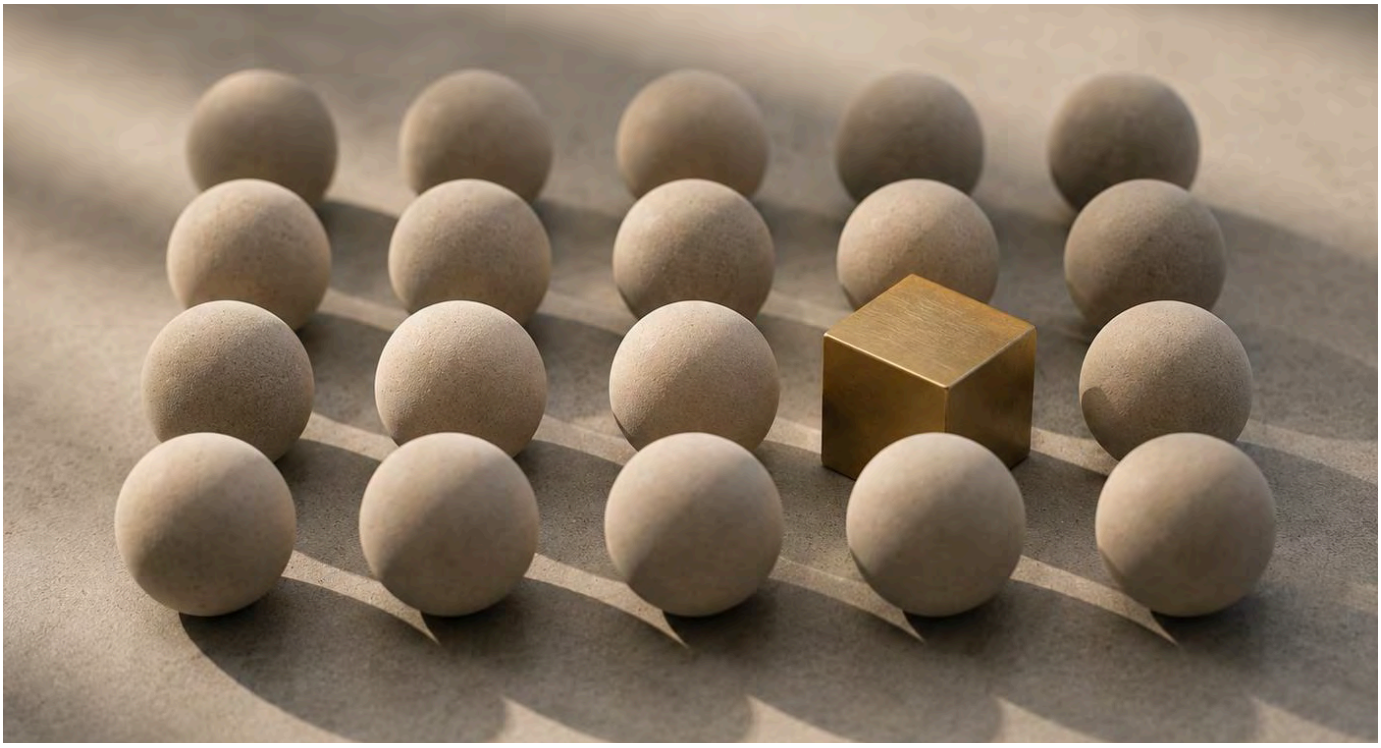
Randomized controlled trials are the most direct approach, where treatment and control groups are assigned at random to compare outcomes. In marketing, this logic powers A/B tests, conversion lift studies and holdout experiments. The idea is to use randomization to eliminate selection bias and ensure the control group serves as a valid proxy for the counterfactual.

When full randomization isn't possible, marketers turn to difference-in-differences, which compares the change in outcomes between a treated group and a control group over time. By focusing on the delta before and after a campaign, this method removes the influence of permanent, underlying differences between the two groups.

Similarly, regression discontinuity exploits specific thresholds, such as a $100 spending limit to trigger a discount, to compare customers just above and just below the cutoff. As such, it allows researchers to isolate the causal effect of the incentive at the point of decision.

Instrumental variables find a variable that affects treatment but doesn't directly affect outcomes, and use it to isolate the causal effect. In marketing, weather might serve as an instrument for store traffic. It affects whether people visit the store, but doesn't directly affect their purchasing intent.

Lastly, synthetic controls construct a synthetic version of the treated unit by weighting untreated units to match its pre-treatment characteristics. This method is the foundation of geolift and geo-experiments.



How the methods connect

Every method covered in this guide, marketing mix modeling, incrementality testing, geo-experiments and saturation analysis, is a different application of the econometric techniques described above. Marketing mix modeling uses regression. Incrementality testing uses RCTs or synthetic controls. Geo-experiments use difference-in-differences. The methods differ in their assumptions and trade-offs, but the underlying problem is always the same: estimating the counterfactual.

Remember that these methods aren't competing with each other. Each makes different trade-offs between feasibility, precision and generalizability. The most credible measurement programs use them together, calibrating marketing mix models with experimental results from geo-tests or lift studies to improve the causal credibility of both.[89]

Use econometrics and causal inference as the analytical foundation for evaluating any model's causal claims. Understanding counterfactual reasoning helps you assess when a model is measuring correlation rather than causation, and when to demand experimental validation before acting on the results.

18

The promise of agentic measurement





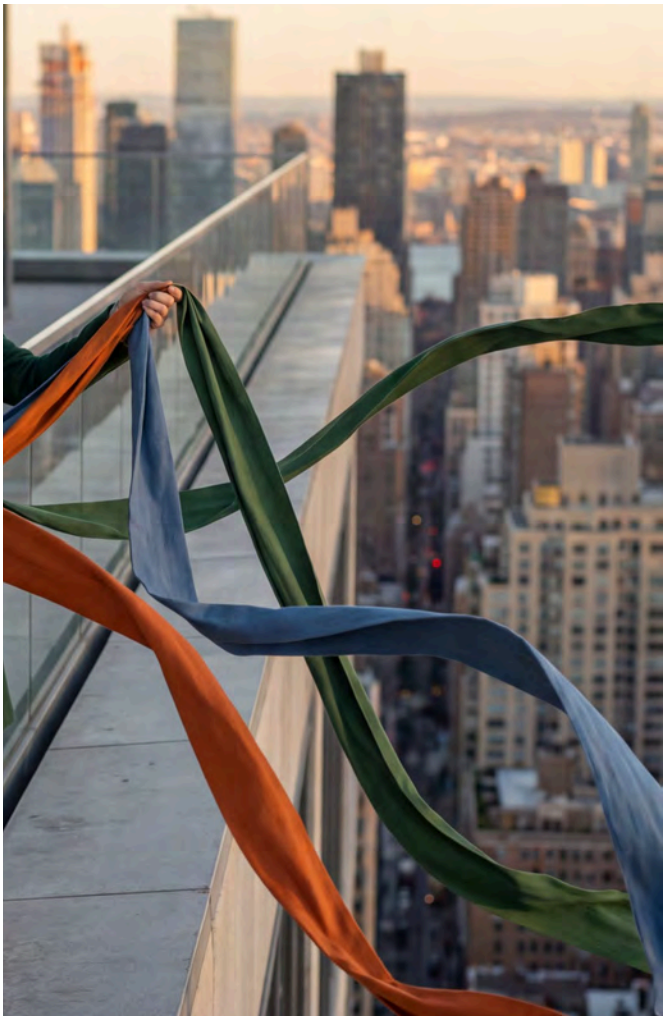
The measurement methods covered in this guide have traditionally operated as periodic projects. Between refreshes, market conditions change and spend patterns and model assumptions can drift away from reality.

Artificial intelligence, and specifically agentic AI, has the potential to change how measurement systems are built, maintained and governed. The shift can transform measurement from a periodic analysis into a continuous, adaptive system.

The scale of the AI pivot

The global agentic AI market is projected to grow from $7.55 billion in 2025 to $199.05 billion by 2034. That's a compound annual growth rate of 43.84%.[90] In marketing specifically, 95% of CMOs in 2024 reported that generative AI investments are a priority,91 and 65% believe advances in AI will dramatically transform their role within the next two years.[92]

BCG reports that early adopters are freeing 15% to 20% in costs and accelerating analytical cycles.[93] However, there is a catch: only 5% of marketing leaders who use generative AI solely as a tool report significant gains on business outcomes.[91] The implication is that AI delivers value when it's embedded in systems and workflows but not when it's bolted on as an isolated capability.



Static models moving to continuous systems

The most meaningful application of AI is the transformation of static measurement models to continuously evaluated systems.

Traditional marketing mix models are static by nature. A team builds a model, runs it, delivers results and waits for the next cycle. Between refreshes, the assumptions baked into the model drift from reality. By the time the next refresh arrives, the recommendations may already be outdated.

Modern AI-enabled systems can evaluate many possible model configurations continuously rather than settling on a single fixed model. They assess which combination of variables, lag structures and channel interactions best explains the data, and they keep reassessing as new data arrives. Models are scored against multiple criteria simultaneously, not just statistical accuracy but also business consistency. A model that fits the historical data well but implies that a channel receiving 5% of spend is responsible for 80% of revenue gets flagged, even if its statistical fit looks strong.

When market conditions change enough that a model's accuracy degrades, what analysts call model drift, the system detects it and triggers recalibration. In the past, retraining the model would have required a consulting engagement on a monthly or even weekly basis.

So we have a genuine step forward. But it introduces a question that organizations need to answer before they adopt it: where does automation end and human judgment begin?

The role of human judgment

Agentic systems can automate model search, scoring and recalibration. The systems can flag anomalies, suggest optimizations and run scenario analyses faster than any human team. What they can't do, and shouldn't be expected to do, is define business objectives, set acceptable risk thresholds, interpret trade-offs between competing priorities or make judgment calls when the data is ambiguous.

The key here is not to "agentify everything." There's a rich opportunity in the middle, where AI makes decisions within structured workflows, guardrails and human oversight.[94]

The middle ground is the sweet spot for measurement. A system that automatically recalibrates marketing mix models based on incoming data is valuable. A system that automatically reallocates millions in spend based on those recalibrations, without human review of the assumptions, is risky.

The organizations that will benefit most from agentic measurement are those that build governance into the system from the start. There should be clear rules about what AI can decide autonomously, what requires human approval and how to escalate when the system encounters situations it wasn't designed for.

Challenges and limitations

Agentic AI is moving beyond conceptual discussion into production infrastructure. PubMatic launched AgenticOS in December 2025, a system-level layer that enables agents to plan, transact and optimize programmatic advertising.[95] In January 2026, the IAB Tech Lab published an agentic roadmap extending existing protocols (OpenRTB, AdCOM and VAST) for agent-based systems.[96]

But the challenges are proportional to the ambition. AI-driven measurement systems inherit the data quality problems of the data they consume. As Stanford's AI research has noted, AI is only as good as the data fed into it.[97] If your data infrastructure is fragmented, inconsistent or incomplete, automating the analysis doesn't fix the problem; it accelerates it.

The black box problem, already a concern with platform attribution models, gets worse as measurement systems become more autonomous. When humans can't explain why a model produced a specific recommendation, trust erodes rather than builds.[98] As a result, accountability gaps emerge alongside. When an agent reallocates budget based on flawed assumptions, who is responsible? The data team that built the system, the marketing leader who approved its use or the vendor who sold it?

Furthermore, vendor lock-in presents a practical concern, too. Organizations that build their measurement systems on a single vendor's agentic platform may find it difficult and expensive to switch if the technology underperforms or the vendor's incentives diverge from the client's.

And it bears repeating: AI is still early-stage technology. The marketing industry has a history of over-indexing on new capabilities before the evidence base catches up. Agentic AI has genuine potential to improve measurement, but it doesn't replace the discipline. It operationalizes it. The objectives, governance and skepticism still come from humans. The speed and scale come from the system.

Use agentic measurement capabilities when you need to move from periodic model refreshes to continuous evaluation, monitor model drift in real time or scale the volume of scenario analyses beyond what manual processes can support. But maintain human oversight for strategic decisions, assumption validation and governance.

Measurement as a discipline

Marketing measurement is not a problem to be solved once. It's a discipline to be built and maintained.

Different methods answer different questions, carry different assumptions and have different limitations. Attribution assigns credit. Incrementality establishes causation. Marketing mix modeling enables strategic planning. Experimentation proves what works. Brand measurement captures what short-term metrics miss. No single method covers the full picture, and expecting one to provide that level of insight is what causes confusion in the marketing department and beyond.

The organizations that get measurement right have learned to work with this complexity rather than fight it.

Matching the method to the question is the starting point. A CFO asking if a campaign generated incremental revenue needs causal evidence from an incrementality test rather than an attribution report showing which touchpoints received credit. A marketing VP planning next year's media mix needs scenario modeling from a marketing mix model, not a dashboard of last month's conversion rates. When the question is clear, the right method usually follows.

Embracing uncertainty is just as important. Disagreement between models is normal and often productive. When MMM and an incrementality test produce different estimates of a channel's contribution, the gap contains information. Understanding why the models disagree is more valuable than picking a single number and moving on.

Building systems rather than silos is what separates mature measurement organizations from everyone else. Triangulation, integrating multiple methods, reconciling competing signals and maintaining structured governance, is the closest thing to a gold standard in marketing measurement right now. Informed starting assumptions, multi-criteria model evaluation and structured calibration processes turn triangulation from a concept into an operational discipline.

Investing in capabilities that compound is what makes the discipline sustainable. Measurement maturity requires data infrastructure, analytical talent, organizational alignment and leadership commitment. Technical accelerators have twice the impact of organizational ones. But organizational alignment determines whether insights translate into action. If your finance team doesn't believe, your insights might as well be dust.

And preparing for continuous evolution keeps the discipline current. AI is transforming measurement from periodic analysis to continuously evaluated systems. Agentic capabilities can automate model search, scoring and recalibration. But human judgment remains essential for setting objectives, defining constraints and interpreting the trade-offs that no algorithm can resolve on its own.

Once you transform your mindset from seeking a universal truth to identifying which method is best equipped to solve your immediate problem, your path forward becomes remarkably clear.

Measurement is a discipline, not a dashboard. Build it accordingly.

Marketing measurement should be the engine that drives smarter spending



References

1 **CMO Survey, MarTech. CMOs, CEOs and marketers are all struggling with martech data issues. 2024.** *https://martech.org/cmos-ceos-and-marketers-are-all-struggling-with-martech-data-issues/*

2 **Funnel. Marketing measurement survey. 2024.** *https://funnel.io/blog/marketing-measurement*

3 **Gartner. 2025 CMO Spend Survey reveals marketing budgets have flatlined. 2025.** *https://www.gartner.com/en/newsroom/press-releases/2025-05-12-gartner-2025-cmo-spend-survey-reveals-marketing-budgets-have-flatlined-at-seven-percent-of-overall-company-revenue*

4 **NIQ. 2025 CMO Outlook report. 2024.** *https://nielseniq.com/global/en/news-center/2024/niq-releases-2025-cmo-outlook-report/*

5 **Funnel. Marketing measurement: data integration and revenue loss. 2024.** *https://funnel.io/blog/marketing-measurement*

6 **Gupta S. How to measure marketing effectiveness. Harvard Business School Online** *https://online.hbs.edu/blog/post/how-to-measure-marketing-effectiveness*

7 **Gartner. Marketing analytics and decision-making practices. 2024.** *https://www.gartner.com/en/newsroom/press-releases/2022-09-15-gartner-survey-reveals-marketing-analytics-are-only-influencing-53-percent-of-decisions*

8 **IPA. Go big or go home: IPA Databank analysis.** *https://ipa.co.uk/news/go-big-or-go-home*

9 **McKinsey & Company. Performance branding and how it is reinventing marketing ROI.** *https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/performance-branding-and-how-it-is-reinventing-marketing-roi*

10 **ProperExpression/OWOX BI. Unified marketing measurement guide. As cited in Lifesight.** *https://lifesight.io/blog/unified-marketing-measurement-guide/*

11 **Ruler Analytics. Last-touch vs. multi-touch attribution. 2024.** *https://www.ruleranalytics.com/blog/click-attribution/last-touch-vs-multi-touch-attribution/*

12 **McKinsey & Company. Performance branding and the silo mentality.** *https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/performance-branding-and-how-it-is-reinventing-marketing-roi*

13 **Funnel. Marketing mix modeling. 2024.** *https://funnel.io/blog/marketing-mix-modeling*

14 **Boston consultancy group: Marketing Measurement Done Right, 2019** https://www.bcg.com/publications/2019/marketing-measurement-done-right

15 **McKinsey & Company. Why every business needs a full-funnel marketing strategy.** https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/why-every-business-needs-a-full-funnel-marketing-strategy

16 **Ofek E. Customer management dynamics and cohort analysis. Harvard Business School Technical Note 520-122. 2020.** https://store.hbr.org/product/customer-management-dynamics-and-cohort-analysis/520122

17 **Harvard Business School Online. How to measure marketing effectiveness.** https://online.hbs.edu/blog/post/how-to-measure-marketing-effectiveness

18 **ReportGarden. Marketing attribution models. 2024**. https://reportgarden.com/post/marketing-attribution-models

19 **Think with Google. ROI and AI-powered measurement. 2024.** https://business.google.com/us/think/measurement/roi-and-ai-powered-measurement/

20 **Dataslayer. Incrementality becomes the primary KPI for retail media advertisers. 2024.** https://www.dataslayer.ai/blog/incrementality-becomes-the-primary-kpi-for-retail-media-advertisers

21 **Feld-M. Google Meridian, triangulation and the future of marketing measurement. 2024.** https://www.feld-m.de/en/blog/google-meridian-triangulation-future-marketing-measurement/

22 **McKinsey & Company. Using marketing analytics to drive superior growth.** https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/using-marketing-analytics-to-drive-superior-growth

23 **Forrester. Google Finally Scraps Its Cookie Deprecation Plans 2024.** https://www.forrester.com/blogs/google-finally-scraps-its-cookie-deprecation-plans/

24 **Lifesight. Trends that are shaping the future of marketing measurement. 2024.** https://lifesight.io/blog/trends-that-are-shaping-the-future-of-marketing-measurement/

25 **Varian HR. Causal inference in economics and marketing. Proc Natl Acad Sci. 2016;113(27):7310-7315.** https://www.pnas.org/doi/10.1073/pnas.1510479113

26 **Thomke SH. Building a culture of experimentation. Harvard Business School.** https://www.hbs.edu/faculty/Pages/item.aspx?num=57805

27 **Thomke SH. Interview: the surprising power of business experiments. InnovationManagement. 2020.** https://innovationmanagement.se/2020/11/12/interview-with-stefan-h-thomke-author-of-experimentation-works-the-surprising-power-of-business-experiments/

28 **Harvard Business School Working Knowledge. Creating the experimentation organization.** https://hbswk.hbs.edu/item/creating-the-experimentation-organization

29 **Thomke SH. Stefan Thomke on experimentation. The CEO Magazine.** https://www.theceomagazine.com/business/innovation-technology/stefan-thomke-experimentation/

30 **Nielsen Norman Group. A/B testing. 2024.** https://www.nngroup.com/articles/ab-testing/

31 **MetricsWatch. What is statistical significance in A/B testing. 2024.** https://metricswatch.com/insights/what-is-statistical-significance-in-ab-testing

32 **Optimizely. Sample size calculator. 2024.** https://www.optimizely.com/sample-size-calculator/

33 **TrueList. A/B testing statistics. 2024.** https://truelist.co/blog/ab-testing-statistics/

34 **GuessTheTest. Calculating sample size in A/B testing: everything you need to know. 2024.** https://guessthetest.com/calculating-sample-size-in-a-b-testing-everything-you-need-to-know/

35 **SH. Interview** https://lifesight.io/blog/geo-experiments-guide/

36 **SH. Interview** https://www.rightsideup.com/blog/guide-to-marketing-incrementality-testing

37 **Facebook. GeoLift methodology documentation. 2024.** https://facebookincubator.github.io/GeoLift/docs/Methodology/

38 **Athey S, Imbens GW. The state of applied econometrics: causality and policy evaluation. 2017.** https://www.aeaweb.org/articles?id=10.1257/jep.31.2.3

39 **Measured Inc. Holdout test case study. 2024.** https://www.measured.com/faq/holdout-test/

40 **INCRMNTAL. 2025 predictions retrospective.** https://www.incrmntal.com/resources/2025-predictions-retrospective

41 **Gartner. Marketing analytics: holdout testing methodology. 2024.** https://www.gartner.com/en/marketing/topics/marketing-analytics

42 **Google/BCG. Incrementality experiments and revenue growth. 2024.** https://support.google.com/google-ads/answer/16719772

43 **Moloco. Proving incremental ROAS in retail media advertising. 2024.** https://www.moloco.com/blog/proving-incremental-roas-retail-media-advertising

44 **Think with Google. Incrementality testing. 2024.** https://business.google.com/us/think/measurement/incrementality-testing/

45 **Google. Incrementality experiments: reduced minimum spend. 2024.** https://support.google.com/google-ads/answer/16719772

46 **Improvado. Incrementality guide. 2024.** https://improvado.io/blog/incrementality-guide

47 **Sellforte. What is incrementality testing. 2024.** https://sellforte.com/blog/what-is-incrementality-testing

48 **Prescient AI. Incrementality testing guide. 2024.** https://prescientai.com/blog/incrementality-testing-guide

49 **EMARKETER. GenAI, programmatic advertising and marketing mix modeling adoption. 2024.** https://www.emarketer.com/content/genai-will-take-over-programmatic-advertising-2026-agentic-ai-isn-t-far-behind

50 **The Drum. Measure what matters: where maths meets magic in modern marketing. 2025.** https://www.thedrum.com/news/2025/05/02/measure-what-matters-where-maths-meets-magic-modern-marketing

51 **Feld-M. Google Meridian, triangulation and the future of marketing measurement. 2024.** https://www.feld-m.de/en/blog/google-meridian-triangulation-future-marketing-measurement/

52 **Think with Google. ROI and AI-powered measurement: MMM and C-level revenue goals. 2024.** https://business.google.com/us/think/measurement/roi-and-ai-powered-measurement/

53 **BCG. A four-legged approach to understanding marketing ROI.** https://www.bcg.com/x/the-multiplier/four-legged-approach-to-understanding-marketing-roi

54 **McKinsey & Company. Marketing return on investment.** https://www.mckinsey.com/capabilities/growth-marketing-and-sales/how-we-help-clients/marketing-return-on-investment

55 **IAB. Modernizing MMM: best practices for marketers. December 2025.** https://www.iab.com/guidelines/modernizing-mmm-best-practices-for-marketers/

56 **Measured. What is marketing/media mix modeling (MMM)? 2024.** https://www.measured.com/faq/what-is-marketing-media-mix-modeling-mmm/

57 **BCG. CMOs need new measurement strategies to cut through the noise. 2024.** https://www.bcg.com/publications/2024/cmos-need-new-measurement-strategies-to-cut-through-the-noise

58 **Taboola. Diminishing returns in advertising. 2024.** https://www.taboola.com/blog/diminishing-returns-advertising

59 **The Drum. How to turn marketing measurement into a growth lever: lessons from Waitrose. 2024.** https://www.thedrum.com/news/how-turn-marketing-measurement-growth-lever-lessons-waitrose

60 **BCG. Brand equity tracking and best-in-class simulations. 2024.** https://www.bcg.com/publications/2024/cmos-need-new-measurement-strategies-to-cut-through-the-noise

61 **BCG/Google, as reported in The Drum. Shared KPIs between marketing and finance. 2025.** https://www.thedrum.com/news/2025/05/02/measure-what-matters-where-maths-meets-magic-modern-marketing

62 **MarTech. Marketing mix modeling has a usage problem, not a tech problem. 2024.** https://martech.org/marketing-mix-modeling-has-a-usage-problem-not-a-tech-problem/

63 **IPA/Marketing Society. The pursuit of effectiveness.** https://www.marketingsociety.com/the-library/pursuit-effectiveness

64 **IPA/Marketing Society. ROI as a measurement metric: cautionary note.** https://www.marketingsociety.com/the-library/pursuit-effectiveness

65 **Binet L, Field P. The long and the short of it. IPA Effectiveness Databank analysis.** https://ipa.co.uk/knowledge/effectiveness-research-analysis/les-binet-peter-field

66 **Nielsen. Digital Ad Ratings: on-target rate benchmarks. 2024.** https://global.nielsen.com/solutions/audience-measurement/digital-ad-ratings/

67 **Nielsen/MMA Global. 2024 Nielsen One Ads Meta Analysis: key trends in digital advertising.** https://mmaglobal.com/articles/2024-nielsen-one-ads-meta-analysis-key-trends-digital-advertising

68 **Insider Intelligence. Advantage+ campaigns as primary growth engine. 2024.** https://www.insiderintelligence.com/

69 **Madgicx. Meta ads trend analysis. 2024.** https://madgicx.com/blog/meta-ads-trend-analysis

70 **Go Fish Digital. Digital marketing prediction: audiences are everything. May 2025.** https://gofishdigital.com/blog/digital-marketing-prediction-audiences-are-everything/

71 **McKinsey & Company. The value of getting personalization right — or wrong — is multiplying. 2025.** https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/the-value-of-getting-personalization-right-or-wrong-is-multiplying

72 **Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. J Educ Psychol. 1974;66(5):688-701.** https://www.jstor.org/stable/2529684

73 **Gutierrez P, Gérardy J-Y. Causal inference and uplift modeling: a review of the literature. Proceedings of The 3rd International Conference on Predictive Applications and APIs. 2017.** https://proceedings.mlr.press/v67/gutierrez17a/gutierrez17a.pdf

74 **Athey S, Imbens GW. Estimating treatment effects with causal forests: an application. Stanford Graduate School of Business.** https://www.gsb.stanford.edu/faculty-research/publications/estimating-treatment-effects-causal-forests-application

75 **Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. 2017.** https://arxiv.org/abs/1706.03461

76 **Towards Data Science. Evaluating uplift models. 2024.** https://towardsdatascience.com/evaluating-uplift-models-8a078996a113/

77 **Greenbook. Mixed-method marketing research: a complete guide. 2024.** https://www.greenbook.org/insights/research-methodologies/mixed-method-marketing-research-a-complete-guide

78 **Drive Research. Difference between qualitative and quantitative market research. 2024.** https://www.driveresearch.com/market-research-company-blog/difference-between-qualitative-and-quantitative-market-research/

79 **Sago. Multi-method research for consumer insights. 2024.** https://sago.com/en/resources/blog/multi-method-research-consumer-insights/

80 **Attest. Qualitative market research. 2024.** https://www.askattest.com/blog/articles/qualitative-market-research

81 **System1. Methodology: Spike Rating and Star Rating. 2024.** https://system1group.com/methodology

82 **LiveRamp. Measurement maturity model. 2024.**

83 **BCG. CMOs need new measurement strategies to cut through the noise. 2024.** https://www.bcg.com/publications/2024/cmos-need-new-measurement-strategies-to-cut-through-the-noise

84 **BCG. CMOs need new measurement strategies: four accelerators. 2024.** https://www.bcg.com/publications/2024/cmos-need-new-measurement-strategies-to-cut-through-the-noise

85 **Stringer J (BCG), as quoted in The Drum. Measure what matters. 2025.** https://www.thedrum.com/news/2025/05/02/measure-what-matters-where-maths-meets-magic-modern-marketing

86 **Think with Google. ROI and AI-powered measurement: modern measurement framework. 2024.** https://business.google.com/us/think/measurement/roi-and-ai-powered-measurement/

87 **BCG. A four-legged approach to understanding marketing ROI.** https://www.bcg.com/x/the-multiplier/four-legged-approach-to-understanding-marketing-roi

88 **Angrist JD, Pischke J-S. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press; 2009.** https://www.mostlyharmlesseconometrics.com/

89 **Number Analytics. The key methods to causal effects in econometrics.** https://www.numberanalytics.com/blog/the-key-methods-to-causal-effects-in-econometrics

90 **Precedence Research. Agentic AI market forecast. 2025.** https://www.precedenceresearch.com/agentic-ai-market

91 **Gartner. Survey finds 65% of CMOs say advances in AI will dramatically change their role. 2024.** https://www.gartner.com/en/newsroom/press-releases/2024-11-17-gartner-survey-finds-65-percent-of-cmos-say-advances-in-ai-will-dramatically-change-their-role-in-the-next-two-years

92 **Gartner. CMOs on AI transformation of their role. 2024.** https://www.gartner.com/en/newsroom/press-releases/2024-11-17-gartner-survey-finds-65-percent-of-cmos-say-advances-in-ai-will-dramatically-change-their-role-in-the-next-two-years

93 **BCG. CMOs who move first in agentic marketing will win. 2025.** https://www.bcg.com/publications/2025/cmos-who-move-first-in-agentic-marketing-will-win

94 **Brinker S. How agentic AI is changing the future of marketing. MarTech. 2025.** https://martech.org/how-agentic-ai-is-changing-the-future-of-marketing/

95 **WebProNews. Agentic dawn: how AI agents will automate programmatic ads in 2026.** https://www.webpronews.com/agentic-dawn-how-ai-agents-will-automate-programmatic-ads-in-2026/

96 **WebProNews. IAB Tech Lab agentic roadmap. January 2026.** https://www.webpronews.com/agentic-dawn-how-ai-agents-will-automate-programmatic-ads-in-2026/

97 **Stanford University HAI. Data-centric AI: AI models are only as good as their data pipeline.** https://hai.stanford.edu/news/data-centric-ai-ai-models-are-only-good-their-data-pipeline

98 **EMARKETER. GenAI and agentic AI: trust and transparency in decision-making. 2024.** https://www.emarketer.com/content/genai-will-take-over-programmatic-advertising-2026-agentic-ai-isn-t-far-behind

99 **Gartner. Gartner identifies top three priorities for CMOs to deliver marketing excellence in 2025. 2024.** https://www.gartner.com/en/newsroom/press-releases/2024-12-03-gartner-identifies-top-three-priorities-for-cmos-to-deliver-marketing-excellence-in-2025

100 **Gartner. 2025 Marketing Symposium Xpo: day 1 highlights.** https://www.gartner.com/en/newsroom/press-releases/2025-05-12-gartner-marketing-symposium-xpo-day-1-highlights

