The True Value of Marketing (TVM): The Blueprint for Profitable Growth

From dashboards to decisions: how to turn marketing investments into a capital allocation system.

Executive Summary

Most enterprises are not short on marketing metrics. They are short on marketing decisions.

When budgets tighten, the scrutiny lands where it always does: marketing. Not because marketing investments are uniquely wasteful but because they're uniquely hard to evaluate with a consistent, finance-grade standard. Marketing teams defend investments using proxy metrics—clicks, opens, ROAS—while finance asks for something else entirely: incrementalprofit, payback, and confidence. The result is predictable: more analysis, more debate, fewer decisive reallocations.

This paper introduces True Value of Marketing (TVM), a benchmark and blueprint for allocating marketing capital. TVM evaluates every marketing investment in a common economic unit—incremental contribution margin—and provides a repeatable decision framework for making explicit trade-offs between competing investments.

TVM evaluates every material marketing investment on three dimensions:

i. **Profitable Customer Value (PCV):** Did the investment create incremental contribution margin?

ii. **Opportunity Value (OV):** Where should incremental dollars go next, based on expected marginal return?

iii. **Complexity Cost (CC):** What value is lost to friction—tool sprawl, integration burden, latency, and rework?

The purpose is not to produce a perfect score, but rather a repeatable decision standard. A scorecard that aligns the C-suite, and improves quarter over quarter.

The meeting that explains the whole problem

The meeting is familiar.

The CMO arrives with a deck. It is thick, comprehensive, even elegant. It contains charts that show share of voice, pipeline influenced, website conversion, brand lift, ROAS, and an attribution model that appears to have an opinion on everything.

The CFO arrives with a different lens. She isn't there to debate whether marketing "worked." She is there to decide whether marketing investments deserve more capital, less capital, or different capital and where it should go.

She asks a question that lands like a gavel:

> *"If I give you the next dollar, where will it create the most value? And what should we stop funding to make room?"*

The room pauses. Not because the question is unfair, but because the marketing system wasn't built to answer it.

One shift is making this question unavoidable: AI has dramatically expanded what marketing can do. Campaigns, creative, audiences, offers, and channels can now be created and tested at a velocity that was unimaginable two years ago, and agents are beginning to execute autonomously.

Technology has a recurring habit: when it makes one resource abundant, scarcity moves somewhere else. AI is making marketing possibilities abundant. The scarce resource is becoming the ability to decide which ones deserve capital.

For decades, the constraint in marketing was production. There were never enough campaigns, audiences, or experiments to pursue every opportunity. AI reverses that. The constraint is no longer creation—it is allocation. The challenge is no longer generating possibilities. It is deciding which possibilities deserve capital. That requires a standard for distinguishing activity from value.

Two broader trends make this both necessary and possible. The mandate for profitable growth means the enterprise can no longer tolerate proxy metrics as evidence of value. And the same forces expanding the possibility set are improving the evidence available to evaluate it: richer data, better attribution, and—where feasible—controlled experimentation at scale. According to the CMO Council and Zeta Global's *CMO Intentions* study, these shifts are showing up in boardroom expectations: 54% of CMOs say senior management expects top-line revenue growth, 51% cite operational efficiency and reduced cost, and 41% cite AI-driven efficiencies.[1]

Every other use of enterprise capital—R&D, capital expenditures, acquisitions—is evaluated against the fundamentals of value creation: incremental return, profitable growth, and a long-term horizon. Marketing has been the exception, not because it is exempt from those fundamentals, but because enterprises lacked a practical way to apply them. AI makes that exception impossible to sustain. TVM provides a practical way to evaluate marketing investments against the same fundamentals of value creation that govern every other strategic use of enterprise capital. The remainder of this paper introduces that framework and shows how to build it using data most enterprises already possess.

What the research shows

Across a combination of executive surveys, industry research, and campaign-level analysis spanning telecom, insurance, retail, financial services, and SaaS, a consistent pattern emerges: marketing organizations are reporting performance but are not effective at allocating capital.

This aligns with broader industry findings. McKinsey notes that while organizations have significantly increased investment in data and marketing technology, many still struggle to translate that investment into measurable financial outcomes.[2]

Three structural gaps are frequently mentioned:

A measurement gap. Most organizations cannot evaluate marketing investments across channels using a common economic unit, such as incremental contribution margin. Instead, they rely on channel-specific metrics that are not comparable. That measurement problem is visible in operations: 40% of respondents say improving ROI and demonstrating attribution is one of their top priorities for improving marketing operations.[3]

A trade-off paralysis. Marketing reviews rarely result in meaningful budget reallocation. Spend tends to remain "sticky," driven more by historical allocation than marginal return. Indeed, a recent McKinsey report highlights this disconnect between marketing activity and financial decision-making, emphasizing the need for stronger CMO–CFO alignment.[4]

A complexity tax. Fragmented marketing technology stacks and identity systems slow execution and reduce realized performance. That friction is not theoretical: 31% cite proving the ROI of existing marketing technology investments as a challenge, 29% cite usability and complexity, 27% cite demonstrating value over total cost of ownership, and 26% cite legacy integration challenges.[5]

The three components of TVM discussed next are aimed at closing these three gaps by putting every investment into a standard unit, supporting explicit trade-offs, and evaluating the friction that quietly destroys value.

The TVM Scorecard: a framework that clarifies value creation

The TVM Scorecard is the one-page standard that evaluates every major marketing investment on the profitable customer value created (PCV), the trade-off among opportunity values (OV), and the complexity cost (CC)—to support a Scale/Tighten/Stop decision. You do not need to include every marketing investment: apply the Pareto rule and include the 20% of investments that account for 80% of the value.

Each row (marketing investment opportunity) is examined along three dimensions: the incremental contribution margin the investment actually created (PCV), the expected return on the next incremental dollars (OV, expressed here as expected marginal return), and the friction cost the opportunity would impose on the system (CC).

Given that the measurement of PCV and expected marginal return is often difficult, 'Confidence' reflects the strength of the evidence behind these measures: High indicates testing or strong quasi-experimental evidence, Moderate indicates modeled estimates with controls, and Directional indicates sparse data or expert judgment.

Expected marginal return estimates the incremental contribution margin created by the next $1 at today's spend level, reflecting the reality that returns typically diminish as spend scales.

This scorecard can become the standing agenda for the monthly CMO–CFO review, providing a common language between marketing and finance.

The TVM Scorecard (illustrative)

Marketing investment	Current spend	PCV	Expected marginal return	Confidence	Action
Retention – churn prevention	$0.8M	$2.1M	$1.85 per $1	High	**Scale**
Acquisition – paid search	$1.5M	$3.2M	$1.40 per $1	High	**Tighten**
Growth – lifecycle cross-sell	$0.4M	$0.68M	$1.20 per $1	Moderate	**Scale**
Acquisition – paid social	$2.0M	$2.8M	$0.65 per $1	High	**Tighten**
Acquisition – display prospecting	$1.2M	$0.96M	$0.30 per $1	Moderate	**Stop**
Brand – sponsorship	$0.5M	$1.1M (est.)	$0.90 (est.)	Directional	**Test**

What TVM is—and what it is not

Marketing exists to acquire, retain, and grow profitable customers. That has always been true. What has not been true, until recently, is that an enterprise could evaluate the profitability marketing investments generate in a consistent, comparable, finance-grade way.

TVM can offer such a measurement approach and a path to act on it. It is a benchmark, not an audit. It does not require atomic precision to be useful. It requires a common unit—incremental contribution margin—and a consistent discipline for applying it across investments.

The first time an organization runs TVM, the estimates will be rough. But rough estimates can be very useful for decision making (Frank, Magnone, and Netzer 2022). Even rough estimates that are on the same scale allow organizations to compare investments, make explicit trade-offs, and track improvement quarter over quarter. Each TVM evaluation cycle is likely to tighten inputs. The benchmark becomes a blueprint.

For directional context, TVM can be expressed in dollars:

$$\text{TVM\$} = \text{PCV\$} + \text{OV\$} - \text{CC\$}$$

For tracking improvement over time, TVM can also be expressed as an index, with 100 representing the baseline period. An organization at 110 has improved the economic value of its marketing portfolio by 10% against its own baseline—a number a board can follow quarter over quarter without relitigating the underlying methodology.

But the power of TVM is not the math. The power is the discipline: evaluating every material investment on the same three dimensions toward decisions.

As marketing becomes increasingly agentic, measurement alone is not enough. An intelligent system also needs a definition of a good decision. TVM provides that value standard: PCV defines the value created, OV identifies the highest-value use of the next dollar, and CC accounts for the friction required to capture it. The intelligence can improve; TVM tells it what to improve toward.

This is what makes TVM urgent now rather than someday. AI changes more than the volume of marketing activity. It changes the quality of evidence available to marketers. Controlled experiments improve the precision of PCV by isolating the incremental contribution of individual marketing investments. Over time, experiments across multiple investments improve estimates of Opportunity Value by revealing where the next dollars generate the greatest marginal return. Experimentation is therefore not an end in itself—it is the mechanism through which TVM becomes progressively more accurate. AI expands the possibility set. Experimentation converts possibilities into evidence. TVM converts evidence into allocation. And allocation directs the next round of experiments—which is how the system compounds.

Profitable Customer Value (PCV)

> *The question: Did this investment create incremental profit or did it just get credit for what would have happened anyway?*

Many CMOs still struggle to quantify marketing's impact in financial terms. PCV is how TVM closes that gap—not with a perfect number, but with a defensible one expressed in the unit finance already uses: incremental contribution margin.

PCV focuses marketing measurement on how much incremental contribution margin this specific marketing investment creates, across acquisition, retention, and growth.

The first decision in measuring PCV is deciding on the time horizon. Retail often uses 12 months. Telecom typically needs 24. Insurance often needs 36 to capture persistence. Whatever horizon you pick, apply it consistently—because comparability is crucial.

Once a time horizon has been decided, PCV should break down profitability into the three customer lifetime value (CLV) components (Gupta and Lehmann 2005):

1. **Acquire:** incremental customers gained × contribution margin over the horizon
2. **Retain:** incremental churn avoided × remaining contribution margin preserved
3. **Grow:** incremental expansion margin from cross-sell and upsell, net of what would have occurred anyway

You don't need perfect measurement on day one. You need to be honest with yourself and the rest of the leadership team and provide a meaningful confidence label. After all, when the finance team creates the projection for next year's profitability, it makes many assumptions, like the conversion rate of the dollar to the euro a year from now.

The strongest evidence for incrementality comes from controlled experiments, and this is where AI has changed the game: holdout tests that were once slow and expensive to run can now be designed, deployed, and read at a fraction of the historical cost—particularly in digital channels. Experiments are the preferred evidence where feasible. Where controlled experiments are impractical—sponsorships, broad media, distribution—TVM relies on the strongest available evidence: matched markets, marketing mix models, and structured estimates with explicit confidence levels. That is precisely what the confidence label is for: an honest account of where each number sits on the evidence ladder, and a plan to climb it.

Operationally, PCV starts with the top 3–5 marketing investments—the largest spend categories—using the best data you have and an agreed margin definition with finance. Then you expand.

What data do you actually need?

Three things you already have, just not joined:

- **Exposure** (who saw what)
- **Behavior** (what happened)

- **Economics** (what it was worth)

The barrier is integration, not data collection. The data problem is not solely owned by marketing. In our study, 59% of business leaders report working with the CIO on customer data and 58% with the IT organization, reinforcing that PCV begins with a shared data layer.[6]

Finally, PCV becomes the fuel for OV: once you know what investments produced incremental margin at a given spend level, you can estimate where the next dollars are likely to earn the highest marginal return across investment opportunities.

In an AI-accelerated marketing organization, accurate measurement of PCV becomes more feasible: controlled experiments can isolate the incremental contribution of individual investments with a precision that was previously impractical at scale. And as agents take on more execution, PCV becomes the measure the system is held against; without it, agents reward whatever proxy they were given and compound the error at machine speed.

Opportunity Value (OV)

> *The question: Where should the next dollars go?*

If PCV is the rear-view mirror, OV is the windshield.

It tells you where the next dollars should go, and why. In a survey we conducted with XXX CFOs, fewer than one in four CFOs report that marketing reviews produce material budget shifts inside the quarter. This finding is consistent with Corstjens & Umblijs (2012) who found across industries that the largest marketing vehicles tend to keep the largest budget shares—not because marginal returns justify it, but because they always have.

OV replaces precedent or even the individual assessment of PCV with marginal economics. PCV tells you what value an investment creates. OV asks what matters next:

> *Is the next $100K in this investment better than the next $100K in another?*

That's the role of the expected marginal return column in the TVM scorecard on top of PCV. Marketing behaves like a dimmer, not a switch: the first dollar often performs well, then returns fade as spend scales. OV makes diminishing returns visible—and makes reallocation defensible.

The challenge: how do you estimate expected marginal return without pretending you have perfect curves? Start with what you can observe:

- **Variation over time:** the same investment run at different spending levels across quarters.
- **Variation across markets:** the same investment deployed with different intensity across regions.
- **Controlled variation:** matched-market tests when the stakes are high, and—increasingly—coordinated experiments across multiple investments, which reveal not just whether each investment works but where the next dollar earns the most across them.
- **Portfolio models that capture curvature:** marketing mix models deserve particular mention here. Modern data and AI have made MMMs substantially faster to build and refresh than their earlier generations, and they are often the strongest available method for estimating marginal returns across investments—particularly those that cannot be experimented on individually.
- **Exogenous shocks:** leverage exogenous shocks (e.g., technical glitches or organizational restructures) that created a reduction or increase in investment.
- **Structured expert judgment** when data is thin—labeled directional and treated as test candidates, not scale candidates.

OV doesn't require academic precision to be useful. It requires consistent discipline: every major investment gets an expected marginal return estimate and a confidence label, so leadership can make explicit trade-offs.

Two final realities matter in practice:

i. **Amplifiers:** some investments increase the yield of others. Awareness can lift conversion. Lifecycle messaging can raise paid efficiency. These should be flagged and decided jointly.

ii. **Redundancies:** overlapping investments aimed at the same audience can cannibalize incrementality. These should be consolidated or reduced.

Expanded experimentation is OV's accelerant. Experiments across multiple investments reveal where the next dollar earns its highest return—which means the organization that experiments most learns the shape of its marginal return curves fastest. Where experiments aren't feasible, modern mix models fill the gap. Agents given OV reallocate intelligently; without it, they saturate channels and miss the opportunities the evidence was gathered to find.

Complexity Cost (CC)

> *The question: What value is the system destroying—even when the strategy is right?*

Many marketing organizations can run strong investments and still underperform because the machine is slow, fragmented, and redundant. CC exists to price that friction, because complexity doesn't change until it has a number.

Here's why CC matters: it's the hidden tax that turns good programs into mediocre outcomes. When an investment launches late, misses the churn window, or can't reach the right customer because identity is broken, the cost shows up as lost contribution margin—not as a line item in a budget spreadsheet.

There's also a broader structural reality: Gartner has reported that enterprises often use only a fraction of their marketing technology stack's capabilities, even as a majority of CMOs report budget pressure.[7] Complexity is consuming resources and slowing learning at the same time.

The complexity tax shows up in multiple ways: usability and complexity, legacy integration challenges, and demonstrating value over total cost of ownership all rank among the most common barriers to marketing technology adoption.

CC is priced in four places most organizations can measure quickly:

- **Latency:** time from insight to launch.
- **Rework:** manual hours spent stitching data, rebuilding audiences, redoing QA, and re-running approvals.
- **Redundancy:** overlapping tools that duplicate capabilities and create training, maintenance, and context-switching costs.
- **Identity fragmentation:** reach and performance loss when the same customer can't be recognized consistently across channels.

The important point is not methodological purity. It's operational clarity. CC gives leadership a way to see where simplification will unlock returns—and why a tech cleanup is actually a margin initiative. At Amazon, for example, for every investment opportunity, leaders ask what are the 'dogs not barking' or 'paper cuts' that you may not be thinking about.[8]

A practical starting point: in the first TVM cycle, you don't need perfect CC. You need to lay out the biggest sources of friction and obtain directional truth:

- Price latency for the top five investments
- Run a one-week time audit for rework
- Perform a stack inventory for redundancy
- Estimate identity fragmentation using match rates and reach gaps

Once CC is visible in dollars and days, it becomes actionable—and combined with PCV and OV it can serve as the basis for decisions.

Complexity throttles experiment velocity. Every redundant tool, broken integration, and identity gap slows the loop from hypothesis to evidence—and in an environment where agents operate at machine speed, that friction compounds. CC is what makes the throttle visible and priceable.

Pulling it together: minimum viable TVM (the first two cycles)

The first two TVM cycles are not a research project. They are a focused effort to establish a baseline using data that largely already exists. The first cycle should follow an initial meeting with the CFO team to introduce the approach and the second to refine following input from the CFO team. While the specific timing will depend on the enterprise's business cycle, many organizations should be able to complete both cycles in roughly 90 days.

Weeks 1–2: Select the portfolio. Apply the Pareto rule: pick the 3–5 major marketing investments that account for the largest share of spend—the 20% of investments driving 80% of the value—and pull 12–24 months of program metadata into a consistent record.

Weeks 3–4: Align economics with finance. Agree on contribution margin definition and horizon. This is the credibility hinge—skip it and nothing downstream holds.

Weeks 5–6: Estimate PCV. Measure incremental contribution margin for each investment using the best available evidence; assign confidence.

Weeks 7–8: Populate expected marginal return. Estimate "next $100K" marginal returns using variation, experiments where possible, and structured directional judgment where necessary.

Weeks 9–10: Price CC. Run a time audit, stack audit, and latency reconstruction to produce a first CC estimate in dollars and days.

Weeks 11–12: Run the first TVM review. Bring the TVM Scorecard to the CMO–CFO meeting. Make at least one meaningful reallocation decision and one simplification commitment.

Iterate.

The TVM flywheel

The first TVM Scorecard is a rough benchmark. The fourth is a precision instrument. That improvement is not incidental—TVM is built as a flywheel:

Experimentation generates evidence. Every test—and every mix-model refresh where tests aren't feasible—adds a data point to the campaign library. Over time, the library grows richer and more reliable.

Evidence sharpens PCV. Better incrementality measurement means higher-confidence estimates of the value each investment actually created.

PCV tightens OV. More reliable value data means tighter estimates of diminishing returns, better identification of amplifiers and redundancies, and higher-confidence marginal return estimates.

OV drives allocation. Reallocating dollars toward higher-yield investments and away from diminishing-returns investments improves next-period PCV.

Allocation funds the next turn. Reallocated dollars and reduced complexity buy more experiments, faster cycles, and better evidence—and the flywheel accelerates.

CC reduction lubricates every stage. Lower latency means campaigns execute on time. Simpler systems mean less rework and faster test cycles. Lower complexity frees budget for the highest-yield experiments.

Agents run this flywheel at machine speed—which is why the standard matters more, not less, as execution automates.

We see TVM as a benchmark and blueprint rather than a report. A report tells you what happened. A benchmark tells you where you stand. A flywheel compounds.

The economic payoff is more than better measurement. It is learning leverage. Every allocation decision produces an outcome; every outcome adds evidence; and every new piece of evidence can improve the next decision. The organization is not simply accumulating campaign history. It is accumulating knowledge about what creates value, at what level of investment, for which customers, and under what conditions. The thousandth allocation decision should be better than the first.

The governance layer

Goodhart's Law suggests that one risk with a single metric is that it will be gamed. Because TVM evaluates investments on three dimensions simultaneously—value created, marginal return available, and friction cost—using different data sources for each, it should be less susceptible to such gaming. Optimizing PCV by inflating attribution is checked by the confidence band. Inflating marginal returns is checked by next-period PCV actuals. Ignoring CC is checked by the existence of frictions post-adoption of the investment.

Low-confidence results earn test budget—not scale budget.

The governance council—CMO, CFO/FP&A, analytics leader, data platform leader—owns the definitions: contribution margin spec, horizon, evidence standards, confidence methodology, and the method for pricing CC. Their job is not to adjudicate every campaign. It is to ensure that the standard is trustworthy even when the answers are uncomfortable.

Monthly CMO–CFO TVM Review

- **Inputs:** TVM Scorecard; PCV waterfall; expected marginal return estimates; CC dashboard
- **Decisions:** every major investment receives a verdict: Scale / Tighten / Stop
- **Outputs:** explicit reallocations, a one-page variance note, and top CC reduction actions

Case studies

TVM is easiest to understand in motion—when it changes the decisions leaders make. We have implemented TVM within our companies and with several of our clients. Each case includes a brief inset and a fuller narrative.

Case 1: Top 5 Telecom — When retention beats acquisition at the margin

At a glance

- *What TVM revealed:* acquisition investments were less incremental than assumed; expected marginal returns were stronger in retention; latency was destroying value in the churn window.
- *Decision:* cap marginal acquisition expansion, reallocation to retention, simplified workflows to reduce time-to-launch.
- *Result:* margin efficiency improved 12–18% within two quarters; retention investments delivered 2–4× the marginal yield of the lowest-performing acquisition expansions; cycle time improved 25–35%.

The case

The telecom CMO faced an uncomfortable truth: acquisition campaigns looked strong on ROAS, yet net subscriber growth had slowed and profitability was deteriorating. Promotions were getting richer, handset subsidies rising, and churn remained stubborn. The marketing organization could demonstrate activity. It could not demonstrate a consistent economic rationale for where to invest next.

The CFO's question was blunt: "If we add $25 million to marketing, where does it generate incremental contribution margin? And what do we cut to fund it?"

The company adopted TVM with a 24-month horizon, reflecting the payback dynamics of customer lifetime value and the cost-to-serve realities of the category.

PCV surfaced a surprise. Incrementality testing across acquisition and retention flows showed that a meaningful share of acquisition spend was reaching customers who would have converted anyway—while retention investments created higher incremental margin per customer in high-value segments.

OV made the trade-offs explicit. The TVM Scorecard made expected marginal returns visible across investments. Two patterns emerged: acquisition in certain channels had hit diminishing returns, while proactive churn interventions showed higher expected marginal return with tighter confidence.

CC explained the execution gap. Campaign-to-launch timelines were long due to fragmented identity resolution and manual segmentation. When CC was priced in dollars and delay, latency stopped being an inconvenience and started reading like what it was: lost incremental margin because churn interventions arrived after the churn window had already closed.

The decisions followed. In the first two TVM cycles, the CMO and CFO jointly reallocated budget from marginal acquisition expansion to retention and lifecycle investments, with a defined cap on acquisition channels past diminishing returns. In parallel, they funded simplification—reducing audience-building steps and consolidating identity workflows—to cut time-to-launch.

The lesson was not a single campaign win. It was a new operating dynamic: the CFO stopped asking for receipts and started asking for yield. Marketing became a portfolio.

Case 2: Top 5 Insurer — When brand investment becomes the rational choice

At a glance

- *What TVM revealed:* low-cost leads were often low-value cohorts; cross-sell delivered higher marginal returns; brand became comparable and rational when measured on the correct horizon; underwriting friction inflated cost-to-serve.
- *Decision:* reduce low-quality lead-gen scale, reallocate to persistency and cross-sell, run a disciplined brand investment with an evidence plan, simplify marketing-to-underwriting handoffs.
- *Result:* contribution margin per acquired customer improved 10–15%; retention improved 2–3 points in targeted segments; cycle time improved 20–30%.

The case

Insurance marketing lives inside a classic tension: brand builds trust and consideration, but performance dominates because it appears measurable. This insurer had built a formidable performance machine—lead gen, quote flows, retargeting, and affiliate partnerships. Proxy metrics looked strong.

The business story was harder. Persistency was uneven. Customer quality varied across cohorts. Cross-sell rates stagnated in key segments. The CFO's complaint wasn't that marketing wasn't working; it was that marketing's wins didn't reliably show where the board cared: contribution margin over time.

The CEO's question cut through the noise: "Are we optimizing for the right customers—or for the easiest conversions?"

The company implemented TVM with a 36-month horizon aligned to renewal economics and lifetime profitability. That governance decision made performance and brand investments comparable over a period that captured value.

PCV redefined "a good lead." Some high-volume lead sources delivered customers who looked great on Day 1 but churned early, carried lower margins, and generated higher servicing costs. Once measured in incremental contribution margin over the horizon, these sources were revealed as high volume, low value.

OV reframed the budget. The TVM Scorecard made expected marginal returns visible across investments. Retention and renewal save journeys, cross-sell programs, and selective brand investments rose to the top when measured over the correct horizon and compared in the same unit.

CC exposed the hidden tax between marketing and underwriting. Complexity wasn't just tool sprawl, it was operational friction: fragmented customer views, manual handoffs, delayed decisions, and repeated requests. That tax showed up in lower conversion and higher cost-to-serve.

The decisions were counterintuitive and therefore decisive. The insurer reduced funding to high-volume lead sources that looked efficient on proxies but were inferior on PCV. It reallocated dollars toward persistency and cross-sell investments with higher expected marginal return. It reinstated disciplined brand investment with an explicit horizon and evidence plan. In parallel, it simplified marketing-to-underwriting workflows to reduce handoff friction.

The lesson was the strategic shift. Marketing stopped being about leads and started being about profitable customer value. Brand became rational when measured with the right horizon and evaluated in the same unit.

Case 3: Leading Specialty Retailer — When ROAS lies and returns tell the truth

At a glance

- *What TVM revealed:* return-adjusted contribution economics turned several high-ROAS investments into value losers; lifecycle and loyalty had higher marginal returns; tool overlap and manual QA slowed learning velocity.
- *Decision:* cap saturated acquisition spend, reallocate to lifecycle and loyalty, simplify tooling and QA, incorporate returns propensity into targeting.
- *Result:* incremental margin per marketing dollar improved 15–25% within two quarters; return-adjusted profitability improved materially in targeted categories; insight-to-launch improved 30–40%.

The case

The retailer was proud of its performance marketing sophistication. Paid social, search, lifecycle email, and onsite personalization were humming. ROAS looked healthy. Profit did not. Returns were high, discounts were intense, and customer support costs were climbing.

The CFO summarized the situation in one sentence: “We’re buying revenue. We’re not buying profitable growth.” TVM was deployed with a 12-month horizon aligned to purchase cycles and repeat behavior.

PCV forced an uncomfortable reframe: returns are not a rounding error. Once incremental contribution margin was measured by segment and category and adjusted for returns and cost-to-serve, several high-ROAS acquisition investments were revealed as value-eroding. They were driving low-margin baskets with high return propensity.

OV made the reallocation case obvious. The TVM Scorecard made expected marginal returns visible across the portfolio. Certain acquisition expansions were past diminishing returns and producing marginal customers with weak return-adjusted economics. Lifecycle investments, category cross-sell, and loyalty programs delivered higher expected marginal return over the horizon.

CC explained why improvements weren’t compounding. Tool overlap and fragmented workflows created friction: multiple personalization layers, split systems for onsite and lifecycle, and manual QA steps that delayed iteration. Slow iteration meant slow learning, which meant more spend allocated based on outdated assumptions.

The decisions followed the economics. The retailer capped spending where expected marginal return had collapsed and reallocated budget toward lifecycle, loyalty, and cross-sell investments with higher expected yield. It simplified the machine—rationalized overlapping tools, streamlined QA, and tightened the path from insight to launch. It incorporated returns propensity and cost-to-serve into targeting so the business stopped buying unprofitable customers by design.

The lesson: ROAS is not a business model. Profit is. Once investments were evaluated in the unit that matters and marginal return guided incremental allocation, the organization stopped optimizing for a proxy and started optimizing for value.

Closing

Marketing is becoming the epicenter of enterprise experimentation—the first function transformed by AI, the first where tests run at machine velocity, and the first where agents are beginning to execute autonomously. But experimentation is not the destination; allocation is.

The question is no longer whether marketing can experiment. It is whether the enterprise has a standard for converting all that experimentation into capital allocation. TVM is that standard. And as each decision produces new evidence, the standard becomes the basis for something more powerful: a marketing system that learns.

With TVM, a company can baseline the value marketing delivers, improve its decision-making regarding marketing investment, and manage it as the vital growth driver of the business that marketing has the promise to be.

The failure of marketing as a practice to align marketing performance with business outcomes isn't a dashboard problem. It's a capital allocation problem.

For thirty years, marketing measured what it did. TVM measures what it's worth.

References

1. CMO Intentions 2024: Fueling Martech Innovation Through AI
2. Rewiring martech into a growth engine | McKinsey
3. CMO Intentions 2024: Fueling Martech Innovation Through AI
4. The CMO's Comeback: Aligning the C-Suite
5. CMO Intentions 2024: Fueling Martech Innovation Through AI
6. CMO Intentions 2024: Fueling Martech Innovation Through AI
7. Gartner Marketing Technology Research
8. https\://devbizops.medium.com/two-years-at-amazon-3e8328a2343b