



# BRIDGING THE GAP

*Why marketing reports never reach the metric the deal team paid for.*

---

Claymore Partners — May 2026

**Claymore**<sup>®</sup>  
PARTNERS

Photo: Mitchell Luo / Unsplash

## SECTION 01

# The QBR

---

Quarterly business review. Conference room. Bad coffee.

The agency goes first. They have been working on the deck for two weeks. There are seventy-three slides. The first sixty-one are about how the campaign performed. Impressions are up. Click-through rate is up. Engagement is up. There is a slide showing a ten-point lift in unaided brand recall in the test markets. There is a slide showing the reach of an influencer activation that the CFO has never heard of and is mentally cataloguing as the kind of thing she will ask the CEO about later.

The CMO speaks for the next twenty minutes. She is translating. Marketing-qualified leads are up 14 percent. Sales-qualified leads are up 9 percent. Customer acquisition cost is flat. Pipeline coverage is at 3.4x against a target of 3.0. She is competent. She knows what the agency report is missing and she is filling it in. None of what she is saying maps cleanly onto anything in the deal model, but she has been in enough of these meetings to know that the room needs the translation done in real time or it will not get done at all.

The operating partner speaks last. He has three questions. The first is how much of the 14 percent MQL lift converted to revenue. The CMO says the conversion data lags by 60 to 90 days and they will have it for the next QBR. The second is whether the cost per closed-won customer has moved. She says the data is in two different systems and the integration project is in the IT roadmap for Q3. The third is whether any of this is showing up in EBITDA.

There is a pause in the room that everyone present will, in their own way, remember.

The CMO says the marketing investment is a leading indicator. The operating partner says the EBITDA number is not. The CFO says nothing because she has heard this exchange in some form at every QBR for nine months and she has stopped finding it interesting.

The meeting ends. The agency invoices for the quarter. The retainer is paid on the fifteenth. Everyone agrees that the next QBR should focus on attribution, which is the same thing everyone agreed at the last QBR.

This is the gap.

*"Three audiences in the same room, three different definitions of what the meeting was about, and no shared metric to settle which one was right."*

The portfolio company writes a cheque every month for marketing services. The agency reports against the metrics that justify the cheque. The CMO reports against the metrics that justify her continued employment. The sponsor reports against the metrics that justify the deal. Three sets of books. Three sets of incentives. None of them reconcile.

## SECTION 02

# Three audiences, three languages

---

Almost no overlap between what is being reported and what is being asked for. The translation has to happen somewhere, and most of the time it does not happen at all.

Be precise about what each side is actually saying when they say marketing is working.

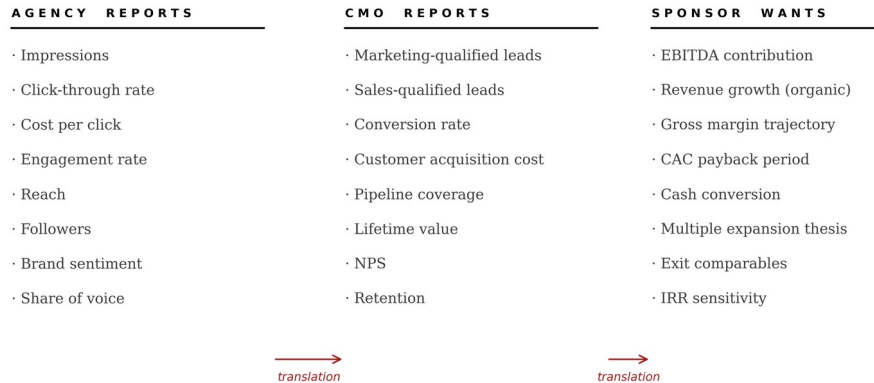
The agency is reporting on activity. Impressions, clicks, engagement, reach, share of voice, sentiment. These are the metrics agencies have always reported on because they are the metrics agencies can demonstrate without making claims they cannot defend. An agency can prove it ran the campaign. It cannot prove the campaign worked, in the sense that anyone outside the agency uses the word work.

The CMO is reporting on funnel mechanics. MQLs, SQLs, conversion rates, CAC, pipeline coverage, lifetime value, retention. These metrics are closer to the business than the agency's, but they sit one or two layers above the financial outcomes the deal team underwrote. A 14 percent lift in MQLs translates into revenue, sometimes. The translation is mediated by sales conversion, deal size, sales cycle length, and a half-dozen other variables that move independently of marketing.

The sponsor is reporting on financial outcomes. EBITDA, organic revenue growth, gross margin trajectory, CAC payback period, cash conversion, multiple expansion. These are the numbers in the model. They are the numbers the LPs see. They are the numbers the deal team underwrote. They are the numbers that determine whether the deal is good or bad.

### Three audiences, three sets of metrics

*Almost no overlap. None of these languages translate cleanly into the next.*



*Both translations happen inside the portfolio company. Neither is anyone's job.*

**FIGURE 1.** *Three audiences, three sets of metrics. The translation between any two of them is mediated by assumptions that nobody owns and nobody verifies.*

The trick is that the agency's metrics need two translations to reach the sponsor's metrics. Agency metrics translate into CMO metrics, which translate into sponsor metrics. Each translation introduces uncertainty. Each uncertainty compounds. By the time you have moved from impressions to EBITDA contribution you have stacked four or five layers of assumption, and the error bars on the final number are wide enough to drive a deal team through.

Nobody in the chain is responsible for the translation.

The agency is not paid to translate its own metrics into the sponsor's. The CMO is paid to translate, but only as a secondary task layered onto her primary job, which is running the marketing function. The sponsor expects the translation to arrive done, and reads the result through the lens of an underwriting model that was built before the marketing data existed.

So the translation gets done badly, or partially, or not at all. Most often it gets done by the CMO in the QBR room while the operating partner is half-listening because he is checking his email under the table.

SECTION 03

## The agency problem

---

Agencies do not report on revenue because revenue is a metric they can fail at.

Impressions cannot be failed at. The campaign ran. The slot was filled. The audience was reached. There is no version of the report where the impression number is missed and the agency has to explain why. The metric is unfalsifiable by design, and agencies report on it because it cannot be used against them.

Revenue is different. Revenue can be measured against a target, and the target can be missed, and the miss can be traced to the work that was supposed to produce the lift. Agencies do not want this. They want the version of the QBR where the work is reported in metrics that nobody can use to fire them.

This is the centre of the agency problem. It is not a measurement gap. It is a deliberate choice not to be measured against the thing that matters.

Every other party in the QBR has skin in the financial outcome. The CEO's bonus is tied to EBITDA. The CMO's tenure depends on the CEO. The operating partner's carry depends on the deal hitting the model. The sponsor's next fundraise depends on a track record. The agency is paid a retainer that renews on the strength of its own reports.

The agency is the only party in the room whose revenue does not depend on the portfolio company growing. It is also the party choosing the metrics it gets graded on.

**Agency revenue and portfolio revenue, three years post-close**

*The agency is doing fine. The business it is paid to grow has not grown.*



Source: Claymore Partners composite, drawn from agency contracts and revenue ledgers across nine portfolio companies.

**FIGURE 2.** Agency monthly fees and portfolio organic revenue, indexed to close. The agency grew. The portfolio did not.

Look at what happens to the agency over the three years. Revenue from the account compounds. The first year it is the original retainer. The second year it is the original retainer plus a paid social pilot, plus a content programme, plus a website rebuild. The third year it is all of that plus a CRM implementation that the agency has subcontracted to a partner firm and is marking up. By month 36 the agency is billing 30 percent more per month than it was at close.

Look at what happens to the portfolio over the same three years. Organic revenue is up 2 percent. EBITDA is flat. The deal model called for both to be considerably higher.

The agency is on a growth trajectory. The portfolio is not. The agency grew because it sold its services well, not because the services worked.

This is the conversation agencies will not have. Ask any agency to put a material portion of its fees at risk against revenue and the answers fall into a small set of patterns. Revenue is too lagged. Revenue is influenced by too many other variables. Revenue cannot be cleanly attributed. Revenue is the responsibility of the sales team, not marketing. Revenue is not the right metric for early-stage funnel work. Each of these is partially true and entirely self-serving. They are also the same arguments the agency uses to avoid every other form of accountability.

Agencies that genuinely believe in their own work would accept some version of revenue accountability. The ones that do not, do not.

The buyer wants performance-based pricing. The agency is paid to argue against it during the negotiation. The agency wins this argument in nine out of ten engagements I have observed. The version that actually works, where the agency is at material financial risk if the portfolio does not hit a financial outcome, is rare enough that I have seen it three times and only one of those survived a renewal.

If you are the deal partner reading this, the line to put in the 100-day plan is the one about the agency contract. Inherited agency relationships become sunk costs the moment the operating partner is in seat. The window to restructure them closes faster than most people realise.

The harder truth is that the agency model is structurally broken for PE-backed companies, and the long-term answer is to bring the function in-house. Performance contracts are a patch. The fix is a marketing team that owns its own data, runs its own measurement, and does not need an agency to translate impressions into revenue. That team is the subject of the next section.

#### SECTION 04

## The CMO trap

---

The CMO seat is winnable. Most CMOs do not win it because they are recruited for the wrong job, briefed against the wrong metrics, and equipped with the wrong infrastructure.

The right CMO at a PE-backed company is not running marketing. She is running the translation between marketing and finance. She builds the dashboards, negotiates the agency contracts, owns the attribution methodology, walks the operating partner through what each line item is producing. She treats the translation as the deliverable and the marketing function as the input.

That CMO can hold the agency accountable, can defend the marketing line in IC, and can produce the financial outcomes the deal team underwrote. The seat is not impossible. It is just rarely set up to succeed.

Three things go wrong. The first is the recruiting profile. The PE-backed CMO seat gets filled by candidates pulled from public-company CMO functions, where the role is materially different. They have spent their careers reporting against brand metrics, with an FP&A team handling the financial translation. They arrive at a smaller function with worse reporting infrastructure and shorter patience for ramp time, and they default to the metrics they know.

The second is the brief. Most CMOs are hired with a marketing brief: build the function, run the campaigns, hit the lead targets. The brief that matches the seat is a translation brief: every dollar of marketing tied to a financial outcome the deal team will recognise, every QBR opening with the EBITDA contribution, every agency contract structured against revenue. The CMO who writes that brief into her own job description in week one tends to survive. The CMO who waits to be told tends not to.<sup>1</sup>

The third is the infrastructure. The architecture decision in Move 4 is what makes the CMO's job tractable. Without it, even the right CMO with the right brief is doing manual translation work in a spreadsheet at midnight before every QBR. With it, the data already speaks the deal team's language and the CMO is interpreting rather than fabricating.

The CMOs who succeed in PE have done this before. They know the QBR rhythm, they have built attribution dashboards from scratch, they have walked away from an agency retainer over revenue accountability. They are rarer than the industry needs and the recruiting market does not select for them. Hiring one is the closest thing to a single decision that closes the gap inside a portfolio company.

---

<sup>1</sup>Average tenure data for portfolio company CMOs is patchy. Spencer Stuart's 2024 CMO Tenure Study puts the median across all Fortune 500 companies at 4.2 years. The PE-backed sample I draw from is closer to 20 months. The discrepancy is partly real (PE tenure is genuinely shorter) and partly a sampling artefact (I only see CMOs who were in seat at entry or were hired post-close, not the ones who left before I arrived).

## Halberd Services

---

Halberd's marketing function is doing more work than it has ever done before, and the business is in worse shape than it was at close.

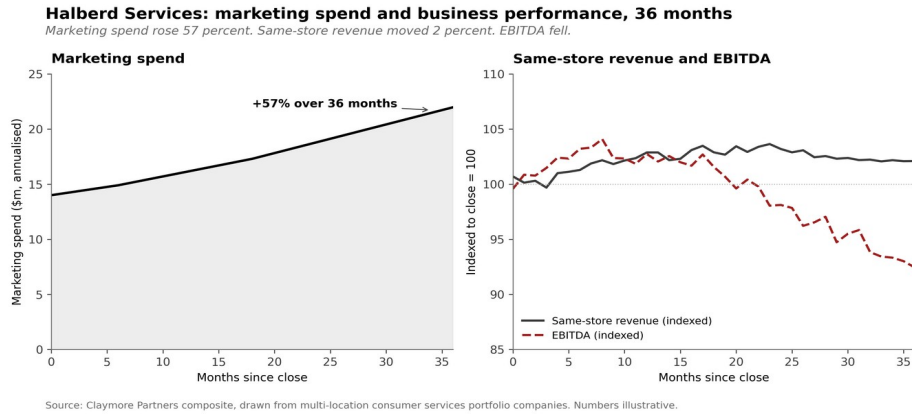
Halberd Services is a composite. The numbers are illustrative, drawn from patterns observed across more than a dozen multi-location consumer services portfolio companies. None are named.

Halberd is a multi-location consumer services business at \$320m of revenue at close, operating in 47 markets. EBITDA is \$52m. The deal closes at an 11x multiple. The marketing function at close is mature for the category: a VP of Marketing, three regional managers, a \$14m annual budget split between paid digital, brand, and field marketing, and a roster of three agencies handling paid media, creative, and analytics. The deal thesis includes a digital lead-generation transformation supposed to add 4 to 6 points of organic revenue growth annually, on top of the new-unit pipeline.

In the first 100 days the new operating team commissions a marketing audit. The audit recommends consolidating the agency roster, hiring a CMO with digital expertise, building out an attribution infrastructure that ties spend to bookings at the location level, and increasing paid digital spend. All of this is reasonable. None of it is wrong on its own terms.

The CMO is hired in month four. She comes from a Fortune 500 retailer where she ran a marketing function five times the size of Halberd's, with a centralised data team that does not exist at Halberd. The agency roster is consolidated to a single holding-company shop in month six. The new agency is more expensive than the prior three combined and runs a more sophisticated programme. The attribution platform goes live in month nine, four months behind schedule, and the integration with the booking system is incomplete until month fourteen.

By month twelve, marketing spend is \$17m annualised. By month twenty-four it is \$20m. By month thirty-six it is \$22m, up 57 percent from close.



**FIGURE 3.** Halberd Services. Marketing spend rises 57 percent. Organic revenue is flat. EBITDA falls.

Look at what the function is producing. The agency's QBR decks are excellent. Cost per lead in paid digital is down 22 percent. Booking-page conversion is up. Branded search volume in the target markets is climbing. The CMO presents a plausible story at every quarterly review.

Look at what the business is doing. Same-store revenue growth in year one is 1.7 percent against a deal model of 5 percent. Year two is 0.4 percent. Year three is negative 0.6 percent. EBITDA is flat in year one, down 4 percent in year two, down 8 percent in year three. The 4 to 6 points of incremental organic growth that were in the deal thesis have not appeared. Customer acquisition cost at the location level has risen 31 percent over the three years.

The marketing function is doing more work than it has ever done before, and the business is in worse shape than it was at close.

When the deal team does its post-mortem at exit, the marketing function is one of several places they look. They find that the spend trajectory was managed properly, the agency relationship was properly contracted, the CMO was experienced, the attribution infrastructure was built. They find no failure of execution at the function level.

They also find that none of the work the function produced ever got translated into a financial outcome that was visible in the deal model. The cost-per-lead number never connected to the revenue number. The revenue number never moved. The EBITDA number went the other way.

The audit team writes a memo concluding that marketing was, in their phrase, hard to attribute.

The phrase is correct. It is also the heart of the problem. The audit team is treating attribution as a property of the marketing programme, as if the marketing programme failed to be attributable. Attribution is a property of the measurement infrastructure that surrounds the programme. The programme was not attributable because the firm had not built the

infrastructure to attribute it. This was true at close. It was true at exit. It was true at every QBR in between.

The buyer of Halberd at exit will inherit the same problem. Marketing spend that produces excellent activity reports and ambiguous financial outcomes. They will run a marketing audit in their own first 100 days. The audit will recommend an attribution platform.

| *"The cycle does not close. It just moves to the next fund."*

## SECTION 06

# The measurement vacuum

---

Marketing measures most where it matters least, and least where it matters most.

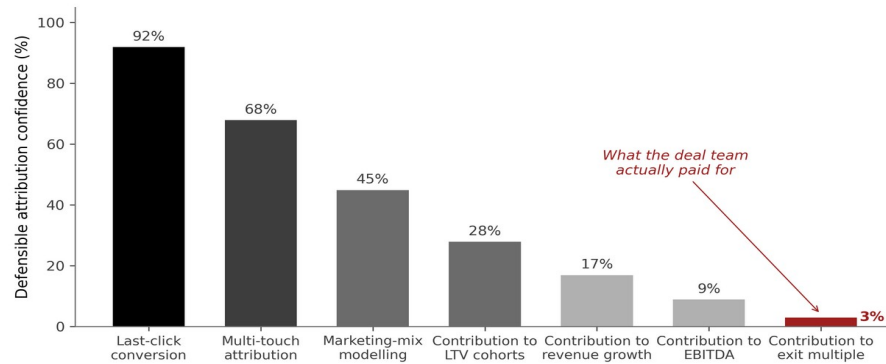
Most of what marketing produces is measured at the wrong end of the funnel. Last-click conversion data is the most reliable thing the function produces and the least valuable, because by the time the customer has clicked the last click they were going to convert at a materially higher rate anyway. Multi-touch attribution improves on this but introduces methodological assumptions the deal team is not equipped to evaluate. Marketing-mix modelling improves further but requires data sets most portfolio companies do not have. Each step toward more useful measurement reduces confidence.<sup>2</sup>

By the time you reach the question the deal team actually wants to answer, which is what fraction of the multiple expansion is attributable to marketing, the answer cannot be defended at any level of confidence an underwriter would accept.

---

<sup>2</sup>The honest version is that some portion of the spend is working and some is not, and we cannot tell which. Marketing-mix modelling can produce a directionally useful answer over a long enough horizon. Three years is not long enough for most of the businesses we see, and by the time the model has enough data to run, the deal is already at exit.

**Where attribution confidence collapses**  
*Marketing measures most where it matters least and least where it matters most.*



Source: Claymore Partners assessment, based on standard methodology robustness across attribution layers.

**FIGURE 4.** *The defensibility of marketing attribution at each layer of the analytical stack. Confidence collapses as the question moves closer to the financial outcome that matters.*

This is not a technical problem that will be solved by a better dashboard. It is a structural feature of the analytical chain. The agency measures impressions because impressions can be measured. The CMO measures funnel mechanics because the platforms produce funnel data. The sponsor wants enterprise value contribution because that is what the deal is for. The measurement that is happening sits at the wrong end of the chain from the question that needs answering.<sup>3</sup>

This is the measurement vacuum. Marketing is in the same position public pensions were in before modern underfunding accounting forced the gap into the open. The infrastructure that would close it does not yet exist.

SECTION 07

## What to actually do

There is no clean answer to this. There are four moves that change the shape of the problem.

The diagnosis in the rest of this paper does not need a sequel. The QBR is the way it is because the contracts, the recruiting, the underwriting, and the reporting cadence all hold it in place.

<sup>3</sup>There is a counter-argument that the deal team is paying for option value, the possibility that marketing produces a result, not a guarantee that it will. I accept this in principle. In practice, deal teams underwrite point estimates, not options, and the marketing line is held to the same standard as everything else.

Change those four things and the QBR changes. Leave them and the QBR runs forever, on its own momentum, in a hundred different conference rooms with different bad coffee.

ONE. Separate the marketing thesis from the rest of the EBITDA bridge, so a marketing miss does not unwind a deal that worked for other reasons.

The cleanest version of this move is to refuse to underwrite a marketing thesis you cannot defend, and price the deal lower. This is the right answer in theory and a losing answer in a competitive auction. The deal partner who prices honestly loses to the deal partner who does not, and the LPs do not reward honesty over winning. Anyone who tells you otherwise has not run a process recently.

The practical version is to ring-fence the marketing line in the deal model. Build the EBITDA bridge so that 0 percent attainment of the marketing thesis still produces an acceptable IRR. Treat the marketing contribution as upside rather than base case. The model still includes it, the IC still sees it, the management team is still given a target, but the deal does not depend on it.

This is uncomfortable to do during a process because the marketing thesis is often the thing that gets you to the bid price. Separating it means accepting that the bid price was high. The compensation is that you stop carrying a structural underwriting error into the hold period.

The version that does not work, which is the version most firms run, is to underwrite the marketing thesis as a point estimate inside the EBITDA bridge using methodology nobody can defend, then hold the operating team accountable for it three years later when the methodology was never going to support the number.

TWO. Build the marketing team in-house, with the metrics and the data infrastructure owned by the company, not the agency.

The agency model is the wrong shape for a PE-backed business and the long-term answer is to stop using it. The case is straightforward. An in-house team operates against the company's revenue numbers, not against an agency report. It owns the data infrastructure from Move 4 directly, rather than asking an agency to feed back into it. It costs less per dollar of output once it is at scale. It cannot leave at the end of the contract.

The version that works is a small senior team. A head of marketing operations who owns the data plumbing. A demand-generation lead who owns paid acquisition. A creative lead, often part-time or contract, who delivers against a defined brief rather than running open-ended creative. The team is metrics-driven by design. Every dollar it spends is tagged at source.

Every campaign it runs reports against revenue, not impressions. The team can be hired in 90 to 120 days for less than the cost of the agency it replaces.

Most buyers do not do this because the agency relationship is inherited at close, the operating partner is not authorized to spend on hiring before the new CMO arrives, and by the time the CMO is in seat the agency is doing the work nobody else can do. The window closes.

The transitional patch, while the in-house team is being built, is to restructure the agency contract so material agency revenue is at risk against a financial outcome the portfolio company cares about. A base retainer that covers operating cost, plus a performance kicker that pays out only against revenue growth in a defined cohort, CAC payback in a target segment, or closed-won pipeline by quarter. The metric is chosen jointly, but the agency does not write it.

The patch is a 12-to-18 month bridge, not a destination. The agency relationship sunsets when the in-house team is operational. Most engagements never plan for the sunset, which is why they never reach it.

THREE. Recruit the CMO for the seat that exists, not the seat the recruiter knows how to fill.

The diagnosis is in the previous section. The fix is the recruiting brief. Filter for prior PE-backed CMO experience, not for the size of the public-company logo. Source from operating roles inside other PE portfolios, from CFO-adjacent commercial roles, occasionally from deal teams. Pay above the public-company benchmark, because the candidate pool is smaller and the seat is harder.

The CMO from Move 3 is the person who builds the in-house team in Move 2 and operates the data architecture in Move 4. Hiring her first sequences the rest of the structure correctly. Hire her second and the architecture decisions will be made by people who do not know what to ask for.

FOUR. Build the portfolio company on a measured-data foundation. Every marketing dollar tagged to a revenue cohort, at source, not retrofitted later.

This is the move that does the most work, and the one that most firms get wrong because they treat it as a reporting problem when it is an architecture problem. The reporting layer is the easy part. The hard part is the plumbing underneath, where every marketing transaction, every campaign, every spend item, every contact in the database is tagged at the moment it enters the system with the cohort it belongs to and the revenue it produced.

This is doable. It is not even technically hard. The reason it does not happen is that nobody owns the decision to insist on it before the spend starts. The CMO inherits whatever instrumentation is in place. The agency runs against whatever the CMO inherited. The deal team accepts the data they are given because the alternative is rebuilding the system, which costs money the operating team has not been authorised to spend.

The 100-day plan has a moment, before the agency is briefed and before the first campaign runs, where the architecture decision can be made. That moment passes in week three of most deals.

If the moment is taken, the rest of the structure changes. Every dollar that flows out is tagged. Every lead that flows in carries the tag. Every closed-won deal carries the original marketing source, the campaign, the cohort, and the cost. The CMO does not have to translate, because the data already speaks the deal team's language. The agency cannot hide behind unfalsifiable metrics, because the system is producing falsifiable ones automatically. The deal team can hold the marketing line to the same standard as every other line in the model, because the standard now exists.<sup>4</sup>

This is not a firm-level capability that needs to be built once and reused. It is a portfolio-company-level decision that needs to be made within the first 100 days, and then enforced. The firms that do this get a portfolio company at exit where the marketing line is defensible. The firms that do not, get the QBR I described at the start of this paper.

This is the move I have most often seen work. It is also the one most often refused on the grounds that it is too expensive at entry. The expense is small. The cost of not doing it is the entire deal thesis.

These four moves do not solve the problem. They reshape it. The QBR still happens. The agency still presents the deck. The CMO still translates. The operating partner still has three questions at the end. The difference is that the answers, when they arrive, mean something.

## CONCLUSION

# The phrase that survives nothing

---

---

<sup>4</sup>I am being slightly utopian about the data architecture here. The reality is messier. Channels overlap. Attribution is probabilistic, not deterministic. Some customer journeys span months and touch six channels. Even with clean tagging, there are judgement calls about which campaign gets credit. The point is not that the data becomes perfect. The point is that the data becomes interrogable. The deal team can challenge the methodology, the CMO can defend it, the agency can be held to it. None of that is possible when the data does not exist at all, which is the starting point in most portfolio companies.

Every PE firm has a value creation deck that says marketing must be aligned with financial outcomes. Every PE firm has portfolio companies where it is not.

The phrase is correct. It is also the start of the problem. Saying that marketing should be aligned with financial outcomes is the first move in a chain of questions the firm has not built the infrastructure to answer. Aligned how. Measured by whom. Reported at what frequency. With what attribution methodology. Defended against what counterfactual. Recontracted on what trigger.

These are the questions that decide whether the alignment exists. The phrase covers all of them and decides none of them.

The four moves above are the closest thing to a fix that I have seen work. None of them is easy. All of them require somebody in the chain to absorb a near-term cost in exchange for a longer-term outcome that may not materialise within their tenure. The asymmetry is what keeps the gap open.

The firms that close the gap will be the firms that build the portfolio company on a measured-data foundation from the first 100 days, recruit the CMO for the seat that exists, build the marketing team in-house with the data infrastructure the company owns, and ring-fence the marketing line in the deal model so it does not unwind a deal that worked for other reasons. Each of these is a decision somebody has to make and force through, before the spending starts.

In most firms, that decision has not been made. The QBR runs on its own momentum, every quarter, in a different conference room, with the same three audiences talking past each other.

*"The four moves are not difficult. They are sequenced wrong, made by the wrong person, or skipped because somebody assumed the agency or the CMO would handle it. None of those is structural. They are all decisions."*



#### AUTHOR

**Lee McCabe**  
Founder, Claymore Partners

Lee McCabe is the founder of Claymore Partners, a value-creation consultancy for PE-backed portfolio companies. He has 25 years of operating experience across big tech, marketplaces, and PE-backed portfolios.

---

#### ABOUT CLAYMORE PARTNERS

Claymore Partners is a growth execution firm for private equity sponsors and their portfolio companies. We install revenue clarity, operating cadence, and execution ownership inside businesses where the gap between strategy and outcome is costing multiple points of EBITDA. The views expressed represent Claymore Partners' analysis as of May 2026.

[claymorepartners.com](http://claymorepartners.com) | [lmccabe@claymorepartners.com](mailto:lmccabe@claymorepartners.com)

This document is published by Claymore Partners for informational and thought-leadership purposes only. It does not constitute investment advice, financial advice, or a recommendation to enter into any transaction. Past performance is not indicative of future results. This document may not be reproduced or distributed without prior written consent of Claymore Partners.