

Defending Revenue when AI is changing everything about the front office

The asymmetry rewriting the B2B
revenue engine and how to close it

June 2026



The better the question. The better the answer.
The better the world works.

Pick any benchmark. The picture is the same.

Three years of data, accelerating into 2026. The pattern is structural, not cyclical.

The cost of expanding inside an existing B2B customer rose 14% year over year through 2024 and another 18% through Q1 2026^[1]. Net revenue retention compressed from 106% in 2022 to 101% in 2025, with select cohorts already printing under 100% in early 2026^[2]. Forty percent of qualified pipeline now dies to no decision rather than to a competitor – the largest single loss category in most enterprise pipelines^[3].

Eighty-nine percent of B2B buyers report a deal stalled in the past year, up from 71% three years ago; 74% of buying groups now experience unhealthy conflict during the decision process^[4]. Quota miss has climbed from 69% of sellers to 78%; win rates have fallen by roughly ten points across the major benchmark cohorts^[5]. And the headline of the past three years: by 2026, 70% of enterprise B2B buyers report using a large language model (LLM) in the research phase of a recent purchase, compared with effectively zero three years ago^{[6][7][8]}.

The pattern names itself in pieces: An enterprise CRM firm's growth decelerated from low double digits to mid-single digits across FY24-25. A Human Capital Management (HCM) leader saw similar deceleration. Four consecutive quarters of revenue decline for a marquee creative agency. Pharma giants grappling with a multiyear normalization, with the Inflation Reduction Act (IRA) accelerating Loss of Exclusivity (LOE) cliffs on key franchise drugs. Each movement is uncomfortable individually. The pattern across them, accelerating over 36 months, is the harder thing.

A note on win rates

A falling win rate, on its own, is not a revenue story. A team with a lower win rate but bigger deals and stable pipeline coverage can still grow. The reason this constellation lands as a revenue story is that everything is moving together – across the demand surface, the deal cycle, and the postsale relationship.

Across most B2B sectors, the same patterns are showing up:



The information asymmetry that built B2B sales has reversed.

For two decades, the seller knew more than the buyer. Not anymore.

For two decades, the seller-buyer relationship in enterprise B2B rested on a stable advantage that favored the seller. The seller knew the product better. The seller understood the implementation better. The seller had the references and the comparisons that the buyer needed in order to make a defensible decision.

That information asymmetry became the core tenet of every commercial motion across most enterprise sectors and segments. The buyer paid for access to the seller's knowledge – in time spent on discovery calls, in patience with sales cycles, in willingness to surface internal politics to a vendor helping them build the business case.

That structure has been quietly inverted. Not by a competitor. Not by a single new technology. By a set of small behavioral changes on the buyer's side of the desk that, in aggregate, have made the buyer the better-informed party in the room.

The cause is everywhere now. Every individual contributor inside the buyer's organization – the analyst running the evaluation, the procurement officer redlining the contract, the champion building the internal business case, the deputy in the CFO's office stress-testing the proposal, the customer success counterpart deciding whether to renew – uses generative AI (GenAI) tools daily, in casual unstructured ways that compound.

Procurement teams run LLM-driven vendor comparisons. Champions arrive at the pre-RFP meeting with a structured three-vendor analysis prepared not by a paid analyst but via their trusted LLM of choice. Reference checks happen as model queries against case-study libraries on the public web. Pricing pushback is drafted with AI assistance.

None of this required the buyer to publish an AI strategy. It happened the way most behavioral shifts happen – through millions of small decisions made by individual contributors who found a tool useful and kept using it.

70%

of enterprise B2B buyers used an LLM during the research phase of a recent purchase by 2026. The number was effectively zero three years ago.^{[6][7][8]}

02 / the symmetry

The seller's customer-facing teams haven't had the same upgrade. Not in most companies. There are exceptions, and the exceptions are pulling away.

The asymmetry that used to favor the seller has reversed. The buyer is now the better-informed party in most enterprise sales and renewal conversations. That single change is the underlying cause of the deal-mechanic deterioration, the demand-channel collapse and the customer-relationship leakage. It is also, more importantly, the cause of the underperformance that boards will start asking chief executives about sometime in 2027.

This piece is for the chief revenue officer (CRO) and the broader front-office leadership group at multibillion-dollar B2B businesses – the chief commercial, sales, customer, marketing and revenue officers who together own the engine. The argument is that the work of the next 12 months is not adopting AI. It is closing the asymmetry.

And it is worth saying out loud, because some leaders are still framing this as a marketing problem, a sales problem or a service problem:

This is not a sales vs. service problem. It is not a business vs. technology problem. It is not a CRO vs. CMO problem. It is all of it. End-to-end revenue.

The companies that recognize this and act on it will compound an advantage that late movers will struggle to recover. The ones that do not will discover, somewhere in 2027, that the gap has widened past the point a reorganization can close.

Megacap tech leaders are running the closest thing to symmetry on the seller's side anywhere in the 124-company sample we analyzed. Their channel partner sales teams are armed with AI-powered analysis. Customer engineering sales motions are essentially AI-native end-to-end. Many run bootcamps at scale across sales engineering. **Tech leaders integrating AI in the sales motion are also, not coincidentally, among the lowest-drift in the entire sample.**

The leading indicators have stopped being signals

Three explanations recur in our work with leadership teams.

The first is that the buyer-side change has been silent and gradual. There was no press release that procurement teams started using AI. No industry report that buying groups began aligning faster with AI assistance. No announcement that customer expectations had recalibrated to consumer-grade AI experiences. The change happened across millions of individual decisions inside thousands of customer organizations, with no single moment that demanded a corresponding response from the seller's side.



A moment from the field

One of us – Chris – recently spent time with the revenue team of a large global enterprise SaaS business. They walked us through their weekly, monthly and quarterly results dashboards. Overall, the team was bullish. Results were stable despite a rocky market. They were surviving the quarter and had visibility to the next. The room felt confident.

But everything in their forward-looking indicators suggested they were heading into significant declines they could not see in the data. The dashboards and business insights they had built were entirely instrumented for a selling motion that was quickly evaporating. In the moment, it felt like watching a real-world version of a lobster being slowly boiled – the team knew something didn't feel right, but they weren't sure what, and the dashboards on the wall were reassuring them that nothing was wrong.

The second is that the metrics most front offices monitor have stopped reporting the change. The dashboards built for the 2018 buyer are blind to the 2026 buyer. MQL volume – the count of marketing-qualified leads, traditionally a leading indicator of pipeline – is propped up by AI-paraphrased query traffic that converts at a fraction of historical rates. Last-click and multi-touch attribution cannot see the LLM citation that placed you on a short list before the buyer opened a browser. Net Revenue Retention (NRR) reported in aggregate hides cohort-level contraction that is the early signal of structural compression.

The third is the most uncomfortable. The leaders of most large front offices are the same leaders who built the playbook now under stress – the chief revenue officer, the chief marketing officer, the chief customer officer, the heads of demand generation and sales operations. Acknowledging the scale of the required adjustment is, for many of them, an implicit acknowledgement that the function they ran in 2022 is not the function the company needs in 2026. That is a difficult internal conversation to start. The companies that find a way to begin it early – frequently through a structured outside diagnostic that names the gap with data – will be in materially better shape than the ones that do not.

Where revenue actually erodes – and why it's invisible to your dashboard.

A diagnostic frame for B2B revenue exposure, scored across 124 publicly-traded businesses in 21 sectors.

The revenue drift map

Two axes. Four quadrants. Five named drifts. Read the map first – we plot the 124-company data on the next page.

The structure is a two-by-two. The horizontal axis is whether a category of revenue loss is **visible to the forecast** – whether it shows up on the chief revenue officer's Monday dashboard. The vertical axis is whether the loss has a **clear executive owner** – whether one named executive will be measured against it.

Invisible · owned Quiet erosion. A leader knows the leak is happening but the dashboard does not surface it. Cohort-level expansion problems. Quiet downsales. Deferred renewals classified as saves. Owned in name, hidden in data.	Visible · owned What every Quarterly Business Review (QBR) fights. Pipeline metrics. Win rates. Deal slippage. Visible to the forecast and owned by the chief revenue officer. The smallest of the five drifts on average – and the one with the most attention pointed at it.
Invisible · unowned Where revenue actually erodes. Demand, decision, and model drift compound here. Roughly 3× the size of the visible losses, and the leak the seller cannot see. No dashboard catches it; no executive is measured against it.	Visible · unowned Tracked. Unaccountable. The signals show up – in usage logs, in MQL volume, in customer success notes – but no single executive is measured against the integrated picture. Multiple teams see fragments. Nobody owns the whole.

You're watching the smallest leak hardest.

That is the diagnostic punchline of the revenue drift map: the trajectory from where executive attention goes runs in the opposite direction from where revenue actually leaks. Five named drifts inside the map, each a category of revenue exposure that compounds when ignored. The next pages walk through each in turn.

124 companies, analyzed.

Each dot is one publicly-traded B2B business. Position = exposure profile. Bubble size = total drift score. The diagonal from top right (low exposure) to bottom left (high exposure) is real and structural.

We scored 124 publicly-traded B2B businesses against the framework – across SEC 10-K filings, the most recent four earnings calls per company, sector analyst commentary and category-prompt panels run against four major LLM models (we may need to use these names here as its core to understanding the research sources).

The revenue drift map

All sectors. 124 publicly-traded B2B business scored against five drifts.




FIGURE 1.
AI tailwind cluster top right; highest exposure cluster bottom left.

What the drift map is telling us.

Three things stand out. The diagonal is real. The sector dispersion is wider than people expect. And a small number of companies – fewer than 5% of the sample – are visibly closing the asymmetry already.

The first signal is the diagonal. Low-exposure companies cluster top right; high-exposure companies cluster bottom left. The diagonal holds across 21 sectors and 124 companies – suggesting that exposure is driven by structural factors (pricing model, product modality, customer accountability) rather than by management quality. Two businesses with the same revenue and similar growth can sit on opposite ends of this map. The implication for a board: you cannot out-execute a structurally exposed revenue model with better sales leadership. The shape of the underlying business has to change first.

The second signal is sector dispersion. Two of our sectors have wider internal variance than the gap between sectors – Enterprise Software and Semiconductors are bimodal, with both AI-tailwind winners and structurally pressured incumbents inside the same category. Sector membership predicts exposure in some cases (Pharma, Telco, Media) and is almost meaningless in others (Tech). The widest spread inside any sector in our sample is in Enterprise Software, where the gap between the lowest-exposure and highest-exposure company is nearly 12 points on a 25-point scale – a 3× exposure gap inside a single category. Don't tell a client “your sector is fine.” Score them.

The third signal is concentration at the extremes. The middle of the chart is sparse. Companies are either visibly responding (the bottom-left tailwind cluster), visibly compressing (the top-right pocket), or in transitional territory that will resolve into one of the two within 12 months. The most useful predictive question for a board is not “where are we today” – it is “which way are we trending.”

Most exposed: where the structural pressure compounds.

Media/Advertising/MarTech 1

avg 18.8 / 25

Hours-billed creative work is the most direct casualty of generative AI. Advertising and MarTech all sit in the high-exposure pocket. The unit being sold – billable creative output – is being substituted by a combination of AI powered SaaS provides and LLM's.

Telecommunications 2

avg 18.2 / 25

Wireline decline, satellite direct-to-cell pressure on consumer wireless and AI-mediated B2B account targeting.

Legacy telecoms all clustered. Model drift averages 4.8 – the highest single drift score in any sector.

Pharma & Life Sciences 3

avg 17.0 / 25

IRA price negotiation, LOE cliffs, and Pharmacy Benefit Manager (PBM) consolidation drive the contraction. Decision Drift averages 5.0 across the sector – buying-group complexity is now the largest single leak in pharma commercial.

Where the structural protection holds – and who is closing the gap.

The bottom of the exposure distribution is anchored by sectors with structural insulation, and by a small number of companies whose commercial motion is already AI-symmetric. The pattern across them is unambiguous.

Least exposed: where structure protects the revenue model.

Energy/Utilities

1

avg 7.0 / 25

Regulated, mandatory spend and AI tailwind from data center build-out. Power generation utilities are the lowest-exposure sector in the entire sample. Demand cannot drift away from a regulated utility – and AI infrastructure demand is the single biggest commercial tailwind in the dataset.

Aerospace & Defense

2

avg 12.5 / 25

Government-direct, AI-resistant systems and multiyear contract structures. Defense contractors and Aerospace all appear in the low-drift cluster. The decision unit is structurally insulated – DoD procurement does not run vendor evaluations through LLM's.

Logistics & Transportation

3

avg 12.6 / 25

Rail and freight benefit from AI-driven demand for hardware and infrastructure shipment. Logistics and transportation leaders cluster low. The product is physical movement, not advisory or per-seat license – model drift cannot easily substitute it.



The common thread across the lowest-exposure sectors is a structurally-insulated commercial motion: contracts are long, decision units are small or regulated, and the AI substitution risk to the unit being sold is low or negative. Sector wins the comparison here – not management.

Responding well today: the four companies anchoring the AI tailwind cluster.

Software and SaaS leader

Embedded across the entire commercial motion – every seller has an AI assistant. AI seller productivity is now a measured Key Performance Indicators (KPIs) inside the Go-To-Market (GTM) org, not a project.

AI-first defense contractor

Bootcamps run as sales engineering at scale. Customers are trained on the deployment by engineers, not by sales reps. Decision drift cannot compound when the seller is inside the implementation.

AI semiconductor leader

Customer engineering motion is AI-native end-to-end. Direct relationships with the buyer's technical decision-maker; no traditional pipeline drift because the seller is codesigning the deployment.

Utility and energy tech leader

Regulated utility riding the AI infrastructure tailwind. Not a commercial-motion case – a structural one. The lesson: pricing model and customer accountability often beat seller skill.

The pattern across these four is consistent: Aggressive seller-side AI investment, consumption- or contract-anchored pricing, and direct customer engagement that closes the asymmetry. The companies in the high-exposure pocket score low on all three of those dimensions.

How to use this page in your own diagnostic. Three questions for the executive reader. Which of the three structural protections do we already have, and which can we acquire? If our commercial motion looks more like the high-exposure pocket, what is the 12 month plan to migrate? And which of our peers is moving fastest toward the bottom-left cluster – because that is the gap we are most likely to lose to.

The visible deal losses dominating the QBR.

Pipeline drift sits in the top-right quadrant of the drift map – visible to the forecast, owned by the CRO. It is the leak the Monday revenue meeting is fighting hardest. It is also, in our 124-company analysis, the smallest of the five drifts on average. The dashboards see it; the playbook is built around it; the front office knows what to do.

Take for example a leader in SaaS. Revenue growth decelerated from low double digits to mid-single digits across FY24-25, with the company's own commentary citing elongating evaluation cycles in enterprise CRM (Customer Relationship Management). A major semiconductor leader sits here too – data center share losses to competitors and custom hyperscaler silicon are well tracked and showing up in the financials in real time. The HCM segment core is showing the same deceleration pattern. Growth has fallen to single digits as customers consolidate to alternatives.

None of these movements are explained by deterioration in seller competence. They are explained by a buyer-side change the sales motion has not adapted to. Pipeline drift is what shows up first in the data because it is the most visible – but it is also the symptom, not the cause.

What to look for in your own business: Deceleration of growth vs. plan over four to six quarters; Sales and marketing (S&M) as a percentage of revenue rising at flat or falling growth; sales cycle length lengthening in earnings commentary; deal slippage as a share of forecasted close.

By the numbers

78%

Win rate, down from 29% in 2022.^[5]

78%

Sellers missing quota, up from 69%.^[5]

36%

Deals slipping past close, up from 28%.^[5]

11×


Top-rep velocity vs. bench, up from 9%.^[5]

Sector spotlight: enterprise SaaS

Of the 21 enterprise software companies in our sample, seven scored 4 or higher on pipeline drift – concentrated in per-seat-licensed categories where AI agent compression is most visible. The pricing model alone explains most of the variance.

Sector spotlight: healthcare services

Healthcare services has an average expansion Drift score of 3.2 across the sample – the highest of any sector. Member churn under Medicare Advantage (MA) reimbursement pressure retail pharmacy traffic erosion, and PBM regulatory uncertainty all flow through to expansion-revenue compression. Chief customer officer (CCO) is structurally a revenue role here, and the org chart has not caught up.



Customers shrinking before they churn.

Expansion drift sits on the boundary between visible and invisible – semi-tracked, fragmented in ownership. It is the leak that nobody owns end to end. The CCO sees usage trends. The CRO sees contracts. The CFO sees bookings. None of them sees the integrated picture in time to act on it.

Most contraction shows up as shrinkage long before it shows up as churn. Customers drop seats. Downgrade modules. Defer renewals. Negotiate "right-sizing" agreements the renewal owner books as a save while the underlying revenue picture deteriorates. By the time the customer formally churns – the moment the dashboard catches it – the contraction has been visible in usage data for two quarters. Three different teams saw three different signals. None of them owned the integrated view.

One company in our sample illustrates the pattern clearly: pressure from member churn combined with declining quality ratings is beginning to impact reimbursement, with signals visible in isolation but no single owner of the integrated picture across the full franchise. A similar dynamic is playing out elsewhere. In enterprise software, net revenue retention is under pressure as collaboration and marketing-focused cohorts contract. In identity management, consolidation is occurring as large enterprise customers standardize on a single platform.

What to look for: NRR by cohort (not aggregate); seat count trends in the two quarters before each renewal cycle; "right-sizing" language in renewal outcome data; the integration of usage, commercial intent and renewal posture in a single owned view.

By the numbers

106% → 101%

NRR compression from 2022 to 2025; under 100% for select cohorts in early 2026.^{[2] [20]}

+14%

Expansion Customer Acquisition Cost (CAC) year over year through 2024.^[1]

5x-7x

Cost to acquire a new customer vs. retain and expand an existing one.^{[14] [15]}

The dark funnel widening as buyers move into the model.

Demand drift sits in the lower-right quadrant – invisible to traditional attribution, contested between marketing and product. It is the leak getting worse fastest.

The dark funnel is widening. AI-mediated search is reducing organic referral traffic. High-intent buyers are bypassing forms, researching inside large language models and arriving later in the journey, if they arrive at all. For most attribution systems, that activity reads as zero.

The pressure is clearest in categories where discovery, comparison and recommendation used to flow through search, marketplaces or agency-led demand channels. In search advertising, AI-generated answers are beginning to compress click-through behavior in commercially relevant categories. The same platforms that built the search economy are now re-mediating it through AI-generated summaries, recommendations and answer engines. The downstream agency and media-services layer faces a parallel challenge: buyers can now use AI tools to generate creative concepts, compare vendors, plan campaigns and bring work in-house that historically required external partners. The result is not a single disruption event. It is a chain reaction across the demand ecosystem.

Travel offers another version of the same pattern. AI assistants can now draft itineraries, compare options, recommend providers and increasingly connect to booking workflows through partner integrations. Discovery that once moved through online travel platforms is shifting into conversational interfaces. Incumbents can respond by building more integrated customer journeys, but the strategic question remains: how much demand will they still originate, and how much will be mediated before the customer ever reaches their environment?

The opposite pattern is visible in categories where AI models consistently recognize and recommend a small set of dominant providers. Run a category prompt across several leading models – for example, “best AI infrastructure,” “best enterprise data platform” or “best cloud migration partner” – and some brands appear repeatedly near the top of the answer set. Others do not appear at all. In an LLM-mediated journey, that distinction is increasingly binary: either the model knows your name in the category context, or it does not. There is very little middle ground.

What to look for: LLM citation rate against your priority category prompts, run quarterly across a fixed model panel; branded search volume versus category search volume; organic traffic trajectory on your domain; and the share of pipeline your attribution dashboard cannot explain.

By the numbers

70%

NRR compression from 2022 to 2025; under 100% for select cohorts in early 2026.^{[6] [7] [8]}

35%

Average organic referral decline in B2B AI Overview categories, Q1 2026.^[9]

1.6:1

Branded search vs. category search ratio in tracked enterprise software vs. 1:1.5 three years ago.^{[10][11]}

Sector spotlight: Media, Advertising, MarTech

The MarTech sector has the highest average demand drift in the sample – 4.2 across WPP, Omnicom and the agency holding companies. The structural pressure is the same dynamic that compresses the client’s search business: when buyers research inside an LLM, the channels that fed agencies for 30 years stop delivering demand. Telco and Wealth Management cluster nearby for adjacent reasons.



No-decision losses and buying-group conflict.

Decision drift sits in the lower-right quadrant alongside demand – invisible to the CRM, unowned and in our 124-company analysis, the largest single revenue exposure on the map. The committee structure that was supposed to bring rigor to enterprise procurement has, in many companies, become the leading reason revenue does not happen.

The buying group has gotten bigger. A decade ago a typical enterprise B2B purchase involved three or four people in a single function. Today it involves 5 to 16 people across as many as four functions. The expansion is sharpest in regulated industries – large healthcare and life sciences purchases now routinely involve clinical, compliance, IT security, finance and procurement; large financial services purchases bring in risk, compliance, technology, business unit and procurement leaders.

A global ERP leader is a textbook case. Major migration procurement cycles take 12 to 24 months at most large enterprises and involve the CFO, the COO, the chief information officer, finance leadership, procurement leadership, regional general managers and consulting partners. The same dynamic exists in the HCM segment. These are the median enterprise software purchases – not edge cases.

The buyer-side AI augmentation makes the buying group both more aligned (every IC arrives with a structured analysis ready) and more conflicted (each function brings its own AI-supported analysis to the table, expanding the surface area for disagreement). The result is the rise of the do-nothing default – the deals that never close-lost because they never closed at all.

What to look for: Percentage of close-lost coded as “no decision”; sales cycle length trajectory; the integrated record of every stakeholder touched in the deal (CRM contact data alone is not enough); the buying-group orchestration motion from stage two onward.

By the numbers

40%

Of qualified pipeline now lost to “no decision” – not to a competitor.^[3]

89%

Of buyers report a deal stalled in the past year, up from 71% three years ago.^[4]

5-16

People in a typical enterprise B2B buying group, across up to four functions.^[7]

Sector spotlight: Pharma & Life Sciences

Pharma is the cleanest archetype for decision drift. Every pharma company in our sample scored 5 (the maximum) on decision drift, reflecting the formulary-plus-payer-plus-provider buying-group complexity that makes oncology and specialty-drug decisions structural multi-stakeholder negotiations. Aerospace and Defense and ERP show the same dynamic.

No-decision losses and buying-group conflict.

Model drift sits in the lower-left quadrant – most invisible, most unowned, most existential. It is the structural risk that the product or revenue model itself becomes obsolete. Different mechanisms; same pattern. The category, the pricing model or both are being rewritten faster than the business can transition.

In the Advertising and MarTech sector agency hours-billing under direct AI substitution is the cleanest model drift case anywhere in the sample, and these companies score 5 (the maximum) on this dimension. One leaders' revenue declined four consecutive quarters through 2024 and into 2025. The agency hours-billing model is being directly substituted by LLMs that handle copywriting, image generation, media planning and increasingly campaign execution.

The single sharpest finding in our sample: **the pricing model alone is a near-deterministic predictor of model drift exposure.** Per-seat licensing in any category exposed to AI agent substitution – sales, service, marketing, HR – scores 4. Consumption-priced businesses where every AI workload deployed by the customer drives net-new vendor revenue score 1-2. There is no overlap.

What to look for: Percentage of revenue priced on inputs an AI-augmented customer can compress (seats, hours, transactions); category-level analyst forecasts of your underlying market; channel disintermediation risk; revenue mix evolution toward AI-resistant or AI-vulnerable segments.

By the numbers

4 of 4

Per-seat enterprise SaaS companies scored 4 on model drift uniformly.

1-2

Consumption-priced peers scored 1-2 uniformly. No overlap

Sector spotlight: telecommunications

Telco's model drift score averaged 4.2 across our five-carrier sample. Wireline is dying structurally. Cable TV bundling is in cord-cutting freefall. Satellite direct-to-cell pilots threaten the wireless core in years 2027+.

Six findings from scoring 124 large B2B businesses.

The patterns across the sample are not abstract. They are the empirical foundation of the playbook below.

01

Decision drift is universally elevated.

Average 3.7 across 124 companies. Higher than demand drift (2.9). Buying-group complexity is the most uniformly exposed drift in B2B, more so than AI search. The story most leaders are telling – AI changed how buyers find you – is true but it is not the highest exposure leak. The bigger leak is AI changed how buying groups decide.

02

Per-seat = high model drift; consumption = low.

Per-seat enterprise SaaS firms score 4 uniformly. Consumption-based cloud and hyperscalers score 1-2 uniformly. There is no overlap. The pricing model alone is a near-deterministic predictor.

03

Tech is polarized – winners and losers diverge.

TMT holds both the lowest and the highest a mix of SaaS and tech leaders. Spread is wider than any other sector. Who you are inside tech matters more than that you are in tech.

04

Hours-billing is the most consistent model drift signal.

Every services-heavy business – agencies, IB advisory, oilfield services, consulting-adjacent IT services – scored 4-5 on model drift. The hours-billing pattern matters more than the industry. The way you charge the customer is dying across multiple categories simultaneously.

05

Government-direct anchors the AI tailwind cluster.

Defense contractors, transportation and utilities all in the lowest-drift cluster. The pattern: mandatory spend + direct relationships + AI-resistant systems = lowest exposure. Defense and utilities feel old-economy until you put them on the framework.

06

The danger and tailwind lines.

The high-exposure cluster is led by hourly billing centric firms, and a cluster of SaaS, Pharma, and Healthcare one step below.

Decision drift is the universal exposure.

Read the heatmap top to bottom: AI tailwind sectors anchor the top, agency and pure-services sectors anchor the bottom.

	Pipeline	Expansion	Demand	Decision	Model	Total
Media/Advertising/MarTech	3.5	3.0	4.2	3.5	4.5	18.8
Telecommunications	3.2	3.2	3.4	3.6	4.8	18.2
Healthcare Services	3.2	3.2	3.2	4.5	3.8	17.8
Pharma & Life Sciences	2.9	2.8	3.0	5.0	3.3	17.0
TMT - Hardware/Infrastructure	3.6	3.0	2.8	3.6	3.8	16.8
Wealth & Asset Management	3.0	2.8	2.5	3.5	4.2	16.0
Energy Services	2.8	3.0	2.0	4.0	4.0	16.0
Travel & Hospitality	2.8	2.8	3.2	3.8	3.2	15.8
Industrial OEM	3.0	2.6	2.9	3.8	3.2	15.4
Banking & Capital Markets	3.0	2.6	2.7	3.7	3.1	15.1
Healthcare Distribution	3.0	2.0	2.0	4.0	4.0	15.0
Enterprise Software/Cloud	2.8	2.5	2.9	3.6	3.0	14.7
Energy	3.0	3.0	2.0	3.0	3.0	14.0
Insurance	2.3	2.4	2.6	3.6	3.0	13.9
Payments & Networks	2.3	2.0	2.3	3.7	3.0	13.3
Semiconductors	2.7	2.2	2.2	3.6	2.3	13.0
Information Services	2.0	2.0	2.0	4.0	3.0	13.0
Logistics & Transportation	2.6	2.4	1.8	3.4	2.4	12.6
Aerospace & Defense	2.3	2.2	1.7	4.2	2.2	12.5
Industrial OEM/Tech Conglomerate	2.0	2.0	2.0	3.0	3.0	12.0
Energy/Utilities	1.0	1.0	1.0	3.0	1.0	7.0

FIGURE 2.


Sector-by-drift heatmap. Average drift score by sector across the 124-company sample. Decision drift glows red across nearly every sector; the consumption-vs. per-seat split shows up as bimodal model drift; government-direct sectors anchor the top of the chart.

The picture you can read in three seconds:

Most front offices are exposed on most drifts, and the column that lights up most consistently is decision drift. The widest variance – the bimodal model drift column – is where pricing-model architecture becomes destiny.

Cyclical vs. structural – the same data, a different question.

Pipeline drift (cyclical pressure) on the X-axis. Model drift (structural pressure) on the Y-axis. Same 124 companies – different strategic frame. The quadrant a business sits in determines whether the right response is execution lift or transformation plan.



The diagnostic value is in the differentiation. Cyclical pressure with a sound model needs an execution lift. A structural model issue with a stable current pipeline needs a transformation plan. The two should not be confused – and a single bubble chart, with all 124 companies clustering at integer coordinates, hides the distinction. Naming the territory – and naming who sits in it – is the diagnostic.

Pipeline is what you watch. Model is what kills you.

Five drifts, decoded. One publicly-traded exemplar each – every score independently verifiable from FY25 disclosures.

Visible - Owned →		← Invisible - Unowned	
01 Pipeline drift	The visible deal losses dominating the QBR. Win rates collapsing, deals slipping, quota miss rising. The drift the dashboard catches – and the smallest of the five on average.	Exemplar Enterprise Software/Cloud Leader PIPELINE DRIFT: high exposure Growth decelerated from low-double digits to mid-single digits across FY24-25. Company commentary cited elongating evaluation cycles and AI-agent substitution in core CRM.	
02 Expansion drift	Customers shrinking before they churn. NRR compression, seat downgrades, deferred renewals. The drift nobody fully owns – usage data, contracts and bookings live in three different teams.	Exemplar Enterprise Software/Cloud Leader EXPANSION DRIFT: elevated exposure HCM core decelerating; large customers consolidating onto AI-native HR platforms. Per-employee pricing under direct AI-agent compression.	
03 Demand drift	AI mediating how buyers research. Buyers research with LLMs before opening a browser. Last-click attribution can't see the citation that placed you on a short list three weeks earlier.	Exemplar Enterprise Software/Cloud Leader DEMAND DRIFT: elevated exposure Both author and victim of the disruption. AI Overviews cannibalize the high-margin search-ad business. Cloud and AI products partially offset, but structural pressure is real.	
04 Decision drift	Buying-group complexity that nobody sees. Average B2B buying group is 11+ people, up from 6.8 in 2017. AI-assisted analysis lets every stakeholder build a counter-thesis before the seller meets them.	EXEMPLAR Pharma & Life Sciences Leader DECISION DRIFT: highest exposure Multiyear normalization. IRA accelerating LOE franchise medications are the headline; underneath, hospital and PBM buying-group complexity has grown beyond what historical commercial playbooks address.	
05 Model drift	When AI substitutes the unit you sell. Per-seat SaaS, hours-billed services and human labor are all under direct substitution from agents. The company's revenue model itself is the leak.	Exemplar Media/Advertising/Media Technology Leader MODEL DRIFT: high exposure in sample Highest-exposure company in the 124-firm sample. Hours-billed creative work is the most direct casualty of generative AI.	

The five drifts compound left to right: ignore pipeline at the top of the funnel and you accelerate expansion, demand, decision, and model.

Four moves to make immediately.

None require a reorganization, a rebrand or a new technology stack. All require acknowledging the previous playbook has stopped working.

Four moves to make immediately.

Each one is high-leverage on its own. Sequenced together, they close most of the asymmetry inside four quarters.

01 Run a revenue exposure diagnostic.

Score your business against the five drifts. Identify where revenue is actually exposed, which engines are under-resourced and which operating-model decisions are still calibrated to the 2018 buyer. Output: a sized capital reallocation plan with named owners and a 90-day cadence. Most front offices respond to the AI shift by putting AI on the perimeter. A real diagnostic rewires the underlying motion.

02 Audit your share of model – and feed the model what it will cite.

Run a fixed prompt set of category-relevant buyer queries across a mix of leading LLM's once a quarter. Track citation rate and short list appearance. Companies tracking share of model are **3-5× more likely to appear on buyer short lists**^[10]. Then ship proprietary research and named-author POVs at **4-7× the rate of generic content**^{[10][9]}. Retire the listicle program. **The model is the new reader.**

03 Move the sales motion upstream.

Signal-based account intelligence. Executive sponsor activation. Multithreading and orchestration of the buying group. Customer success co-engagement on six-figure deals from stage two. Sales teams that have made the change close **23%-37% more pipeline at the same headcount**^{[12][13]}. **The single most important asymmetry-closing move.**

04 Convert existing customers into your demand engine.

Stand up a structured customer advocacy and reference program designed to feed next year's pipeline. Case studies the models will cite advocates. Advocates posting in dark social channels. Reference contacts invocable in late-stage deals. Companies running formal customer advocacy programs see **25%-40% more late-stage pipeline conversion** influenced by peer references^{[21][22]}.


The cost of doing nothing.

A \$10b B2B business, exposed to the four invisible-and-unowned drifts at the average rate observed in our sample, is leaking 9% to 13.5% of theoretical revenue per year – **\$900m to \$1.35b** annually. The visible piece (pipeline drift) accounts for \$150m-\$250m. Decision drift, the largest single leak, runs at **\$300m-\$500m**. **The piece you are fighting hardest is the smallest, by a factor of three.**

EY Studio+ owns the customer agenda – end to end.

Strategy, design, technology and outcomes. The whole front office, run as one team. AI-native by design.

EY Studio+ is the premier front-office consultancy inside the EY organization – built to sit at the table with chief revenue, chief customer and chief marketing officers before the brief is written. Seven integrated capabilities, applied to the sectors where revenue motion is being rewritten fastest. **FOSS – Front Office Strategy & Solutioning** – is one of those seven, and the most natural starting point for the conversations this paper sparks.



The revenue drift map, the diagnostic and the moves in this paper are FOSS work – a way to start the conversation. The execution that follows pulls in the capabilities above as a single team. That is the unlock: most front-office problems are cross-cutting, and most consulting firms are not.

Start with FOSS

The revenue exposure diagnostic.

Six weeks. Five drifts scored against your data, your dashboards and your buyer behavior. Sector benchmark from the 124-company analysis. Top-three prioritized moves. A costed roadmap. EY Studio+ executes what the diagnostic prescribes.

Start at FOSS@studio.ey.com.



The companies that begin this work this year will compound an advantage their competitors will spend the back half of the decade trying to close.

The ones that don't will discover, somewhere in 2027, that the disruption they were watching at a distance was already inside their P&L.

The team behind this analysis.

Four partners inside EY Studio+ working with chief revenue, customer, marketing, and commercial officers on the structural rebuilds the next twelve months will require.



Chris Day

Principal, EY Studio+ Front Office Strategy & Solutioning Leader

Ernst & Young LLP

chris.day@studio.ey.com

Chris has led GTM strategy and front-office transformation at multiple global enterprises with a relentless focus on value.





Salima Eboo

Principal, EY Studio+ Marketing and Adobe Alliance Leader

Ernst & Young LLP

salima.eboo@studio.ey.com

Salima has led MarTech transformation and experience strategy for global brands with a deep focus on personalization and growth.





Brian Smith

Principal, Americas EY Studio+ Leader

Ernst & Young LLP

brian.smith@studio.ey.com

Brian leads EY Studio+ in the Americas – EY enterprise design and strategy studio for technology-led transformation.





Kevin Salsberry

Principal, EY Studio+ Commercial Leader

Ernst & Young LLP

kevin.salsberry@studio.ey.com

Kevin leads Commercial GTM at EY Americas – go-to-market design, sales transformation and revenue motion redesign.



References.

1. "2025 SaaS Performance Metrics Benchmark Report," *Benchmarkit*, 2025.
2. "Q1 2026 SaaS Metrics Update," quarterly cohort tracking, *Benchmarkit*, 2026.
3. Dunford, A., *Sales Pitch: How to Craft a Story to Stand Out and Win* (Ambient Press, 2024).
4. "B2B Buying Survey 2025: Stalled Deals, Buying Group Conflict, and the Rise of No Decision," *Gartner*, 2025.
5. "2025 GTM Benchmarks Report," *Ebsta and Pavilion*, 2025.
6. "2026 B2B Buying Behavior Survey," *Gartner*, 2026.
7. "2026 B2B Buyer Journey Research," *Forrester*, 2026.
8. "Buyer Experience Report 2026," *6sense*, 2026.
9. "AI Catalyst Q1 2026 Field Study: Organic Referral Decline in B2B Categories," *BrightEdge*, 2026.
10. "Share of Model: Enterprise Software Citation Tracking 2026," *Profound*, 2026.
11. "Branded Search Trends and ZeroClick Search Behavior 2024-2025," *SparkToro*, 2025.
12. Account Engagement Platform Customer Outcomes 2025," *6sense*, 2025.
13. "Account-Based Experience Benchmark Study," *Demandbase*, 2025.
14. "Customer Experience Index and Retention Economics Benchmarks 2025," *Forrester*, 2025.
15. "Customer Success Benchmarks 2025," *Gartner*, 2025.
16. "Topline Growth and Operating Efficiency Benchmark Report," *ICONIQ*, 2025.
17. Public software-as-a-service comparison analysis 2025-26, multiple sources, including ChartMogul NRR-to-multiple regressions.
18. "Operating Model Benchmarks 2026: Revenue per FTE by Scorecard Discipline," *Forrester*, 2026.
19. "Customer Success Benchmarks 2025," *Gainsight*, 2025.
20. "SaaS Retention Report 2025: Net Revenue Retention by Cohort," *ChartMogul*, 2025.
21. "Buyer Behavior Report 2026," *G2*, 2026.
22. "Customer Reference Research 2025," *Gartner*, 2025.

In addition to the published research above, this paper draws directly on the FOSS analytical workbook scoring 124 publicly-traded B2B businesses against the Revenue Drift Map. SEC 10-K filings and the most recent four earnings call transcripts for each company in the sample are the primary public-data inputs; the analytical methodology is documented in a separate appendix available on request. The figures cited are triangulated from the most defensible benchmark in each domain. Where studies conflict, the figure used is the one most consistent with adjacent data. Most figures are within a defensible range rather than point estimates.

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is a client-serving member firm of Ernst & Young Global Limited operating in the US.

© 2026 Ernst & Young LLP.
All Rights Reserved.

2605-10529-CS
US SCORE no. 31550-261US

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, legal or other professional advice. Please refer to your advisors for specific advice.

ey.com