

How Revenue Plans Break.

A field guide to the 23 failure modes.

What each failure mode looks like in the wild, what it costs, and the question to ask your own plan before the year answers it for you.

403

B2B revenue plans
analyzed, 2022 to 2025

43 pts

attainment gap between
the strongest and weakest
plans

\$329,840

lost attainment per quota-
carrying rep, per year

THE PREMISE

Most plans miss before the year starts

When a revenue team misses its number, the autopsy usually lands on execution: coaching, pipeline discipline, market headwinds. Our analysis of 403 B2B revenue plans points somewhere earlier. The strongest plans in the dataset hit 91% of quota and the weakest hit 49%, with the 51 plans weak on all four drivers falling further still, to 42%. That gap was visible on the day the plans were deployed, before a single rep had worked a single deal.

Four structural drivers, each measurable at deploy, predicted the gap. Ranked by predictive strength:

DRIVER	STRONGEST THIRD	WEAKEST THIRD	GAP
ICP Coherence	86% attainment	49%	37 pts
Territory Equity	84%	55%	29 pts
Coverage Adequacy	87%	51%	36 pts
Capacity Realism	80%	56%	24 pts

The drivers also compound. Each additional weak driver drags attainment further down, and the fall is steep:

WEAK DRIVERS AT DEPLOY	PLANS	AVERAGE ATTAINMENT
None	78 (19%)	91%
One	91	72%
Two	96	59%
Three	87	49%
All four	51	42%

60% of the plans in the dataset deployed with two or more drivers already in the weakest third. On a 30-rep organization, the difference between a strong plan and a weak one is roughly \$9.9M a year.

Dataset: 403 B2B SaaS annual revenue plans deployed 2022 to 2025, organizations of 30 to 300 quota-carrying reps, average AE quota \$760,000. Full methodology and findings: orbytal.ai/research/annual-plan-reality-check

Plans break twice

The four drivers above describe how plans break at deploy, when the structure itself carries a defect the year will eventually expose. But plans also break a second way: gradually, while the year runs. Pricing erodes deal by deal, retention leaks cohort by cohort, target lists drift while reps keep working them. We call this plan decay, and it explains why a plan that scored well in January can still miss in October.

Mapping both kinds of failure across the full revenue engine produces six pillars and 23 named failure modes. Some are structural, meaning they can be confirmed from the plan and current state on day one. Others are behavioral, meaning they emerge in the data as the year runs. Every one of them costs money that a board deck will eventually have to explain.

PILLAR	THE QUESTION IT ANSWERS	MODES
Capacity Realism	Can the team we actually have carry the quota we actually assigned?	5
Coverage Adequacy	Is there enough qualified pipeline, well run, to make the number?	6
ICP Coherence	Are we selling to the right accounts, and does our model still know who those are?	4
Territory Equity	Are books balanced, and is the opportunity where the quota is?	4
Pricing	Are we holding the price we planned to?	1
Retention & Expansion	Are we defending and growing the base we already won?	3

The next four pages walk through all 23. Each entry follows the same pattern: what the failure looks like from the inside, why it costs money, and the question to ask your own plan. We lead with Capacity Realism, which ranks fourth in predictive strength but first in false confidence, because it is the failure most teams are certain they do not have.

Can the team we actually have carry the quota we actually assigned?

Phantom capacity rising

The plan counts reps and quota that do not exist yet: unfilled reqs carrying full targets, mid-ramp hires counted at full productivity. The number was approved against capacity that was never real, and the gap surfaces two quarters later looking like an execution problem.

ASK: How much of our assigned quota is carried by people who are not hired or not ramped?

Quota over-assignment cliff

Total assigned quota sits well above the operating plan to create cushion, and the cushion hides an expected shortfall from day one while quietly teaching reps that targets are fiction.

ASK: What is our assigned-quota-to-plan ratio, and can we defend it from attainment history?

Ramp time exceeding plan

New hires reach productivity slower than the ramp curve assumes, so planned capacity arrives light and late. The research benchmark for full AE ramp is 5.3 months; plans routinely assume less.

ASK: What was our last cohort's actual time to productivity, and does the plan use that number?

Attrition without backfill

A departure with no backfill req behind it leaves quota structurally uncovered while the forecast keeps counting it, and every week of delay pushes the recovery further into next year.

ASK: For every departure this year, is there an open req with a dated start behind it?

Hiring-manager bandwidth collapse

Open reqs sit unworked because the managers who own them have no interviewing capacity, which means the hiring plan will slip regardless of what the recruiting dashboard says.

ASK: Are our open reqs actually moving through interviews, or just open?

Is there enough qualified pipeline, well run, to make the number?

Coverage erosion from a capacity gap

Coverage thins because the capacity that was supposed to generate it is short, yet the response is usually a pipeline push aimed at a symptom while the root cause sits in the hiring plan.

ASK: When coverage dips, do we check the hiring plan before we check marketing?

Single-threaded deal majority

Most open deals rest on a single contact, and single-threaded deals win at materially worse rates. The pipeline number looks fine while its composition quietly does not.

ASK: What share of our open pipeline has exactly one name on it?

Late-stage stall

Late-stage deals age past the point where history says they slip or die, and the forecast keeps counting them at full weight until the quarter ends and they roll or vanish.

ASK: How many committed deals are older in stage than our closed-won deals ever were?

Top-of-funnel inflow declining

New qualified pipeline falls below its seasonal baseline now, which projects a coverage shortfall roughly two quarters out, long after the cause has been forgotten.

ASK: Is this quarter's pipeline creation funding the quarter after next?

Outbound effectiveness erosion

Reply rates, meeting rates, and outbound conversion decay gradually, eroding the one pipeline lever the team controls most directly while sequences keep running on autopilot.

ASK: Are outbound conversion rates trending against their own baseline, or just being reported?

Competitive displacement

Losses to a named competitor rise at a rate that deserves its own fix, and lumping them in with ordinary slipped deals hides a pattern that battlecards and enablement could address.

ASK: Do we know our win rate against each named competitor, and which direction it is moving?

ICP Coherence

4 FAILURE MODES • STRONGEST PREDICTOR IN THE RESEARCH

Are we selling to the right accounts, and does our model still know who those are?

Named-account list drift

Accounts change while the target list stands still, so reps spend the year working names that no longer fit the profile the list was built on.

ASK: When was the named list last re-scored against current account data?

ICP definition drift

The fit model diverges from what actually wins, and every downstream decision that reads it, from routing to territory design, degrades with it.

ASK: Does our win rate still climb with fit score, or has the relationship gone flat?

Marketing-routing fit collapse

Routing sends leads below the real fit bar into rep queues, burning seller time on accounts the model already said no to. This is a configuration fact, checkable in an afternoon.

ASK: Does the routing threshold match the ICP definition we actually sell against?

Account engagement misalignment

Rep effort concentrates on low-fit accounts while high-fit accounts sit untouched, a silent waste that never shows up in activity totals because the totals look healthy.

ASK: If we split activity by fit band, where is the effort actually going?

Territory Equity

4 FAILURE MODES

Are books balanced, and is the opportunity where the quota is?

Top-rep book concentration

Too much opportunity sits in too few books, a fragility that costs nothing today and a great deal the day a top rep resigns.

ASK: What share of total opportunity do our top three books carry?

Account-tier crossing

Accounts that grew or shrank past their tier stay assigned as if nothing changed, over-served on one end and under-served on the other.

ASK: How many accounts crossed a tier threshold since assignments were last set?

Under-built territories

A book without enough addressable opportunity to cover its quota cannot hit the number regardless of rep effort, and it will still be the rep who gets managed out.

ASK: Does every territory hold enough opportunity to plausibly cover its quota?

Are we holding the price we planned to?

Discount drift

Average discount slides past the guardrail one deal at a time, segment by segment, and net realization erodes quarter after quarter without any single deal looking like the problem. One failure mode, but it compounds against everything the other pillars recover.

ASK: What is our average discount by segment against the guardrail we set in planning, and which direction has it moved since?

Retention & Expansion

Are we defending and growing the base we already won?

Renewal pipeline not built

Renewals approach their date with no opportunity and no motion behind them, revenue the plan counts as safe sitting completely exposed.

ASK: For renewals inside the next 90 days, does each one have an owner and an open opportunity?

Expansion motion missing

The installed base holds whitespace the plan is often already counting, but no expansion pipeline exists against it, which turns a growth assumption into a hope.

ASK: How much of our net revenue retention target has actual pipeline behind it?

Cohort retention degrading

Each customer cohort renews slightly worse than the one before it, a slow compounding leak that annual averages smooth over until the trend is expensive to reverse.

ASK: Is each cohort retaining better or worse than the cohort before it?

A note on what this guide leaves out. Each of these failure modes has detection logic behind it: baselines, data-quality gates, and backtests that separate a real pattern from noise. That instrumentation is Orbytal's product, and it is deliberately absent here. What no instrument can replace is the willingness to ask these questions of your own plan.

Score your plan before the year does

Five questions, one afternoon, using the same thresholds the research measured. Answer them honestly and you will know which of the four deploy-day drivers you are weak on, and the compounding table on page 2 tells you what that weakness historically costs.

1 Coverage

What share of your segments enter the year at 3x pipeline coverage? In the research, plans meeting that threshold across segments sat in the strongest third of the coverage driver.

2 Territory

How evenly is opportunity spread across books? If your top book holds several multiples of your median book, you are carrying concentration risk the plan has not priced, and it comes due the day that rep resigns.

3 ICP

What share of your named accounts score above your own stated ICP definition? This was the single strongest predictor of attainment in the dataset.

4 Capacity

Add your phantom capacity (quota on unfilled seats) to your unramped capacity (quota on hires inside the 5.3-month ramp benchmark). What share of total assigned quota is that?

5 Compounding

Count how many of the four came back weak. The dataset average falls from 91% attainment at zero weak drivers to 42% at four, and most plans deploy with at least two.

The 78 plans that deployed with zero weak drivers shared one habit more than any other trait: they treated the plan as a living instrument to be checked against reality, rather than a document to be defended.

The audit above is a snapshot. Plans decay continuously.

Orbytal watches all 23 failure modes against your live data, names each risk in dollars, and gives you a governed way to fix it, starting the first day you connect.

[Book a demo at orbytal.ai](#) · [Score your plan in 10 minutes](#)

Passed this along by a colleague? Get your own copy at [orbytal.ai/how-revenue-plans-break](#) · © 2026 Orbytal. Share freely with attribution.