Closing Foundry.

Forecast Accuracy: Why B2B Forecasts Miss

Written by
Douglas Mancini
Senior Operating Partner

Reviewed by
Laurie Mascott
Senior Operating Partner

Published 23 Mar 2026 Read time 8 min read

Key Points

- Most B2B forecasts reflect seller confidence rather than buyer evidence, which is why they miss.
- The root causes are upstream: poor qualification, stage gates that don't enforce buyer actions, and CRM data that records activity rather than buyer commitment.
- Forecast accuracy is a proxy for the health of the whole sales system as well as a number to report to the board.
- Good forecast accuracy requires stage exits tied to buyer actions and a weekly review that inspects evidence rather than asks for rep updates.
- The fix is sequential: qualification standard first, then stage discipline, then CRM evidence quality, then how the week is run.
- Teams that forecast well don't use better prediction tools. They use better information, built from consistent qualification and inspection.

What is forecast accuracy?

Forecast accuracy is the measure of how closely a sales team's predicted revenue for a given period matches the revenue actually closed. A team with high forecast accuracy can tell the board at the start of a quarter what will close, and be right within a tight margin. A team with low forecast accuracy can't.

In most B2B sales organisations, the forecast isn't a prediction. It's a summary of seller confidence. The sales leader asks the team what will close. Each rep gives a number based on how they feel about their deals. Those numbers are rolled up, adjusted by the manager's instinct, and presented as

a forecast. This isn't forecasting. It's organised optimism.

Forecast accuracy is a proxy for how well the whole sales system is running. A team that forecasts reliably has stage discipline, qualification standards, buyer evidence in the CRM, and a weekly review rhythm that catches problems before they become misses. A team that can't forecast reliably usually has none of those things. Hope isn't a forecast.

Why B2B forecasts miss

The most common explanation for a missed forecast is "deals slipped." That isn't a cause. It's a description of what happened. The actual causes sit further upstream.

The first is qualification failure. Deals that should never have entered the forecast are included because no one asked the right questions early enough. The buying group wasn't mapped. The economic buyer never engaged. The timeline was assumed, not confirmed. These deals were never real, but they looked real on the forecast.

The second is stage discipline that exists on paper. Most CRM setups have stages with entry criteria. Very few teams actually enforce them. A deal moves to Proposal because a proposal was sent, not because the buyer confirmed a decision timeline, named their stakeholders, and described what success looks like. Stage movement reflects seller activity, not buyer commitment.

The third is a CRM that records effort rather than evidence. Notes say "good call, positive" rather than "buyer confirmed budget is held by the CFO, decision needed by end of Q2, three other stakeholders identified." One of those entries is useful for forecasting. The other isn't.

The fourth is how the week is run. If the weekly pipeline review is a number update rather than an evidence review, the forecast never improves. The manager asks where are we on this deal. The rep says still progressing, should close this month. No one asks: what did the buyer do last week? What is the next committed step? Who else is involved in the decision? Without those questions, problems stay hidden until the last week of the quarter.

What it actually costs

The direct cost of a missed forecast is obvious: missed targets, disappointed boards, and the pressure that follows. The less visible costs compound over time.

Missed forecasts destroy planning accuracy across the business. If sales can't tell finance what will close, finance can't plan headcount, spend, or runway. The entire operating plan becomes a guess layered on top of another guess.

Persistent forecast inaccuracy signals something specific to the board: that the sales motion isn't under control. It's one of the clearest indicators that the team is selling on hope rather than process. Boards use forecast accuracy as a health signal. Revenue leaders who can't explain why their forecast was wrong lose credibility faster than those who miss the number but can explain the mechanism.

At the deal level, the cost shows up in wasted sales effort. Reps and managers spend time on deals that will never close. Late-stage work happens on opportunities that were never properly qualified. That time comes at the expense of deals that could have closed if given the attention.

What good looks like

A team with strong forecast accuracy operates within a plus or minus 10 per cent margin against the called number, consistently. Not every quarter. Consistently, over a run of quarters.

What makes that possible isn't a better forecasting tool. It's the quality of the data going into the forecast. Deals in the forecast have confirmed buyer evidence at each stage. The economic buyer is known and has engaged. The decision timeline is confirmed, not assumed. The business case has been presented and tested. There is a mutual action plan with agreed next steps, and the buyer has completed at least one of them.

Stage gates mean something. When a deal is in Commit, specific things have happened on the buyer's side as well as the seller's. If those things haven't happened, the deal isn't in Commit, regardless of when the rep thinks it will close.

The weekly review inspects evidence, not confidence. The manager reviews what the buyer did last week, not what the rep plans to do next week. That shift separates the teams that forecast well from the teams that guess well.

Forecast calls are graded. Each deal is categorised: Commit (high confidence, evidence-backed), Upside (possible but not certain), and Pipeline (working deals not expected to close this period). The rolled-up number is built from those categories, not from rep optimism.

How to build it

Start with the qualification standard. If your team doesn't have a consistent set of questions that every deal must answer before entering the forecast, that is the first thing to fix. What is the business impact the buyer is trying to address? Who owns the budget? What is the decision timeline? Who else is involved in the decision? These aren't optional questions. They're the foundation of a reliable forecast.

Next, tighten stage exits. Go through your current CRM stages. For each one, write down what a buyer must have done, not what the seller must have done, for the deal to advance. If a stage can be passed by sending an email, it doesn't give you forecasting signal. Stage exits need buyer actions: attending a stakeholder call, sharing a decision timeline, confirming budget, completing a step on a mutual action plan.

Then address what goes into CRM. A note that says "good call" isn't forecast data. Build a habit, or a CRM field, that captures the specific evidence behind each stage. What did the buyer confirm? What did they agree to do next? Who was on the call? What objections came up and how were they addressed? This takes discipline to install. It pays back in forecast accuracy.

Finally, change how the week is run. The weekly pipeline review should be evidence-first. The manager's questions should be: what did the buyer do last week? What is the next committed action on the buyer's side and when is it happening? If the rep can't answer those questions, the deal shouldn't be in the forecast.

This doesn't need to be done all at once. A phased approach works: qualification standard first, then stage discipline, then CRM evidence quality, then the inspection rhythm. Each one builds on the last.

Common mistakes

Forecasting from the pipeline number rather than the evidence. A large pipeline feels safe. It isn't. Pipeline coverage is useful as a planning metric but it tells you nothing about deal quality. A pipeline full of poorly qualified deals with no buyer engagement will miss the forecast regardless of size. Busy pipeline isn't progress.

Confusing forecast calls with deal reviews. A forecast call should take the committed number and check whether the evidence behind each deal has changed. It isn't a deal review. If the team spends forecast calls going deal by deal in detail, the meeting is doing the wrong job. Deal reviews happen separately, focused on the evidence quality of specific opportunities.

Using a single forecast number. A single number obscures risk. A forecast with Commit, Upside, and Pipeline categories gives the business a realistic view of the range. The board can plan against Commit and understand that Upside and Pipeline represent probability-weighted additions.

Fixing the symptom rather than the cause. Many teams respond to forecast misses by asking reps to update their close dates more frequently. That doesn't improve accuracy. It improves the frequency of inaccuracy. The fix is upstream: qualification, stage discipline, and evidence standards. Updating a poorly maintained CRM record more often isn't forecasting rigour.

How to tell if it's working

The clearest signal is variance. Track the called number at the start of the month against the actual close. Over three to four months, the variance should be narrowing. If it isn't, the inputs aren't improving.

Deals that slip from Commit should be unusual. If more than one or two Commit deals slip in a quarter, the Commit standard isn't tight enough. Stage velocity should be consistent. If deals are sitting in the same stage for a long time, either the stage isn't meaningful or the buyer isn't progressing. Both are useful signals.

Reps should be able to articulate buyer evidence, not rep activity, when asked about a deal. If the answer to "where is this deal?" is "I sent the proposal last week," the evidence standard isn't embedded yet.

The weekly review should produce fewer surprises over time. Early-stage problems should be visible earlier. Late-stage risks should surface before the last week of the quarter. If surprises are still appearing in the final days, the inspection rhythm isn't doing its job.

Further reading

From Gut Feel to Ground Truth: Operationalising Sales Forecast Accuracy A practical guide to the operational changes that move a forecast from opinion to evidence.

Forecast Accuracy: The Founder's Discipline and How to Reach Plus or Minus 10 Per Cent How founder-led teams can build the discipline behind a reliable number.

From Revenue Surprises to Early Intervention How to spot at-risk deals earlier and act before they become misses.

Win Rate Is the Cleanest Signal of Sales Health for a Start-Up Why win rate and forecast accuracy are connected, and how improving one reinforces the other.

How to Create a Repeatable Sales Process The process foundation that makes forecast accuracy possible at scale.

Related terms

- Mutual Action Plan: A shared timeline agreed with the buyer that maps steps to a decision date.
- Pipeline Coverage: The ratio of pipeline value to quota, used to assess target attainability.
- Forecast Accuracy: The degree to which predicted revenue matches actual closed revenue.
- Stage Discipline: The practice of only advancing deals when objective criteria are genuinely met.
- Economic Buyer: The person with budget authority and final sign-off power in a deal.
- Buyer Evidence: Proof points and documentation that validate a buyer's decision to purchase.
- Forecast Call: A regular meeting where sales leaders review deal status and revenue projections.
- Business Case: The documented justification for a purchase, including ROI and strategic rationale.

Related service: MEDDPICC training
Source: https://www.closingfoundry.com/insights/forecast-accuracy