The missing variable in renewal forecasting

Why a Team of 90% Accurate Account Managers Can Still Miss the Forecast

Alex Raymond, Founder of AMplify
amplifyam.com

You sent the invite six weeks out. You built the agenda with your champion. You got your own boss on the calendar, and your boss's boss. Forty-eight hours out, the executive sponsor gets pulled into something with their CEO. Twenty-four hours out, somebody has a dentist appointment. At start time you are sitting there by yourself. One minute, nobody. Two minutes, nobody. Seven minutes in, one person joins and says, “just me today.” You run the deck anyway, because you are polite, and you thank them for their time.

Then you open the CRM, find the renewal, and type 70%.

Your QBRs are perfectly fine. Your slides are lovely. That is not the issue. You inherited this meeting from a manager who got it from their manager, and every vendor your customer works with runs the same one, so your customer shows up expecting exactly what you brought. Nobody sat down and designed this. It arrived.

What it costs is the subject of this paper.

Count something first

Before you read the rest of this, go and count something in your own book.

Take your last four business reviews. For each one, ask whether a person who controls the budget was in the meeting for the whole meeting. Not invited. Not on the forward. Not represented by somebody who said they would pass it along. In the meeting, for the duration.

Write down the number.

That count is your renewal forecast. Not the number in the CRM, the one you just wrote down. The rest of this paper explains why, and what the arithmetic does with it once ten people on the same team all write down the same low number.

What your number is for

Here is what happens to that 70% after you type it.

Your boss takes all of your seventies and adds them all up together, and then averages them with everybody else's seventies, and then they roll up to the director. The director rolls up to the VP. The VP takes the total into a meeting where senior decision makers use it to answer questions with money attached: are we heading in the right direction, how does cash come and go this quarter and this year, what can we hire against, what can we spend against. Since about 73% of your company's revenue comes from customers you already have, your number is most of the answer.

Josh Abdulla, Chief Customer Officer at Asana, holds his organization to plus or minus 3% at the start of a quarter. His version of it: “If you say $100 is going to churn at the beginning of the quarter, by the end of the quarter, you should be between $97 and $103.” Cliff Unger, who has been CRO at several growth-stage companies, found that only 20% of sales organizations forecast within 5% of their results, and he ranks the two jobs in order. Job number one is hitting the number. Job number 1A is being accurate about it.

It runs both ways, which surprises people. Beat the number by two million and the CFO smiles, then works out that if she had known three months earlier she would have spent it on headcount, pipeline, or product, and now she has lost a quarter of growth to your good news.

So the issue is not churn. Your company understands that customers churn for all sorts of reasons, many of which are totally outside your control. The issue is that you did not see it coming, and that is the credibility killer. **The enemy is surprise.**

And you get one miss. You do not get two. Miss once and you can explain it, and they will listen, because everybody in that conversation has missed a number somewhere. Your next ninety days start, and you are not allowed to miss again. Miss the second time and nobody announces anything. You find out when finance starts building its own version of your forecast.

Ten people, ninety percent

Let me ask you a question I skipped when I taught this live.

You have ten Account Managers. Every one of them is 90% accurate on their renewals. How accurate is the team's rolled-up forecast?

Most people say 90%. Some say better than 90%, because ten people's mistakes should average out. They can. But only if your ten people are making different mistakes.

Take ten Account Managers with a million dollars of renewals each, a ten million dollar book, and model a 90% accurate Account Manager as somebody whose dollar error typically runs about 10% of their book, or $100,000.

Ten Account Managers, each ~90% accurate	Portfolio variance
They are wrong about different things	3.2%
Half of what they get wrong is shared	7.4%
They are wrong about the same things	10.0%

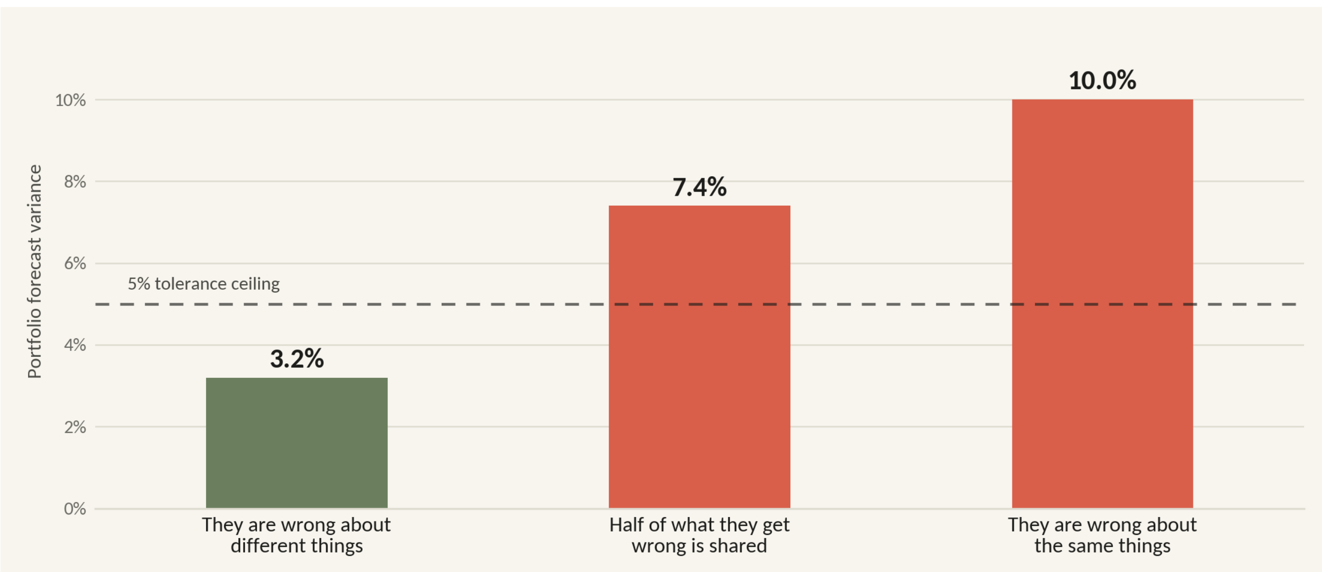

Same ten people, same individual accuracy. Only the team wrong about different things clears the 5% ceiling.

Same ten people. Same individual accuracy. One number is the gold standard and one number gets you the conversation with finance.

When one of your people forecasts too high and another forecasts too low, the two of them cancel each other out, and enough of that across a portfolio makes the total surprisingly accurate even though no individual is. When all ten are wrong about the same thing in the same direction, nothing cancels. Nobody got worse at their job, and your variance more than tripled.

> That is the variable almost nobody in Post-Sale manages. Are your people wrong separately, or are they wrong together?

Why ten people miss the same thing

Look at what your Account Managers read to build a forecast. Product usage, which tells them what somebody clicked. Health scores, built out of that same usage. Satisfaction scores, which tell them how somebody felt three weeks ago. CRM notes, which tell them what they already knew.

All four come from inside your own company. The customer never has to say a word.

And none of them answers the only question the forecast needs answered, which is whether the person controlling the customer's budget intends to keep paying you. Your customer can be in the product every day while procurement builds a consolidation case. Your day-to-day contact can love you while their executive sponsor has stopped defending the line item. The health score can be green a month after finance decided the budget does not survive next year.

So your Account Manager guesses. And because they like their customers and believe in the work they have done, they guess high. They type 70%, the universal yeah, kind of, maybe, I think so, I hope so rating, against a renewal that is going to be a yes or a no.

One person doing that is uncertainty. Ten people doing it for the same reason is a forecast that is voted on rather than measured, and it goes to the board formatted to look like analysis.

Two things that will not fix it

Headcount will not fix it. If the errors were independent, a bigger team would average them away beautifully. Shared error does not average away. At 90% individual accuracy with half the error shared, ten Account Managers put you at 7.4%, forty put you at 7.2%, and four hundred put you at 7.1%. Hire 390 people and buy three tenths of a percentage point.

Standardization is the more dangerous one, because it feels responsible. After a bad quarter you roll out the same template, the same scoring model, the same questions, the same definitions, and you do get consistency. But if you train ten people to talk to the same day-to-day contact, read high usage as a strong signal, and treat attendance as commitment, you have not removed the blind spot. You have written it into the operating model, and now your whole team misses at the same time.

> Standardize the evidence. Never standardize the conclusion.

A standard tells your people what they have to come back with. Does the problem we solve still matter to this customer next year, in a funded way. Are we visibly the answer to it. Does somebody who controls the budget believe it is worth the money. Is anything structural in the way, like a reorg, an acquisition, a consolidation, a competitive evaluation. How does this renewal get decided, by whom, and when.

A formula tells them what it means. Executive attended, usage is high, call it 90%. The standard sends ten people to ten customers and brings back ten different answers. The formula hands ten people the same answer before anybody has spoken to a customer.

Teach your team what a non-answer sounds like, too, because most Account Managers accept one and write it down as evidence.

> "Yeah, this is really important to us" is present tense and says nothing about next year.
>
> "The team loves working with you" is likeability.
>
> "We have no plans to change anything" is inertia.
>
> "Not that I'm aware of," from somebody who would not be aware.

One rule that looks like a formula is worth keeping, and it runs in the opposite direction: if no senior decision maker attends, the account goes red. If the senior person leaves ten minutes in, the account goes red. That rule turns a missing conversation into a recorded unknown instead of a hopeful guess.

What the meeting has to produce

Now go back to the count you wrote down at the start.

The old-school QBR cannot produce any of the evidence above, and it is not close. You present a status deck to somebody who does not control the budget, everyone nods, and you leave with an attendance record and a nice feeling. So the 70% sits there for another quarter, on ten accounts at once, on ten Account Managers' books at once.

The Growth Review is the meeting where the customer's own senior decision makers show up and stay, and where you hear the answers from the person who controls the money. What matters to them next year. Whether we are producing something worth continuing to fund. Whether the budget survives their planning cycle. Who decides. What they have not told you they are worried about.

Each of those answers belongs to a single account, which is the one property your internal signals do not have. That is why the Growth Review matters to the forecast: it is where the shared blind spot gets broken, one account at a time.

Guy Rubin, founder of Ebsta and now at Fullcast, analyzed thousands of deals and billions of dollars of renewals, and he shared what came out of it on my podcast. When the last two business reviews before a renewal were with C-level executives, companies were seven times more likely to open a cross-sell or upsell opportunity. When those same two meetings were with somebody more junior, they were four times more likely to churn the customer. That is an association rather than a proven cause, and it is measured on outcomes rather than on forecast quality. It still tells you which of those two Account Managers knows where they stand.

Anthony DeShazor rebuilt the executive review process at a donation platform where any customer could leave on any day, with no contract and no notice period. Over eight quarters it was basically binary. His words:

> "We didn't lose a single customer who had an EBR... If they had an EBR, we know we were not going to lose them. It was just that clear."
>
> Anthony DeShazor, Account Management Secrets podcast

An Account Manager who gets that conversation stops typing 70%. They type 90%, or they type 10%. Some of your accounts are leaving, and that is not a failure of forecasting. It is the thing you need to know first. A 10% typed in June is worth more to your company than a 70% typed in November, because you can still work a 10%. You cannot work a surprise.

Find out which problem you have

Two tests. Neither one takes long.

Pull the last four quarters of forecast against actuals and look at the direction of every miss, Account Manager by Account Manager. If the misses scatter, some high and some low, your people are mostly wrong separately, individual accuracy is your constraint, and coaching works. If the misses cluster on the same side quarter after quarter, your people are missing the same thing at the same time, and coaching them one at a time changes nothing while they are all reading the same four signals.

Then open the current forecast, take the ten largest renewal probabilities in it, and ask two questions about each one. Has anybody on this team spoken with the person who controls the budget? What did that person say, in their own words, that supports the number sitting in our CRM?

Where the answer is no, and where nobody can produce the words, you are not looking at a forecast. You are looking at a guess with a decimal point on it.

> Go get the read while you can still do something with it.

For readers who want the math

This is a modeled argument rather than a study of before-and-after review performance. For the example above, 90% accuracy is modeled as individual dollar forecast error with a standard deviation equal to 10% of the Account Manager's book.

For ten equal books, aggregate forecast uncertainty is:

$$\sigma \times \sqrt{[n + \rho n(n - 1)]}$$

where σ is individual forecast error, *n* is the number of Account Managers, and ρ represents how correlated their errors are. At zero correlation, the ten-person team lands at approximately 3.2% portfolio variance. At $\rho = 0.5$ it lands at approximately 7.4%. At full correlation it lands at 10%, which is exactly as wrong as any one person, with more zeroes on it.

As *n* grows with ρ held fixed, portfolio variance approaches $\sigma\sqrt{\rho}$ as a percentage of an individual book. At $\rho = 0.5$ and $\sigma = 10\%$, that floor is 7.07%, which is why headcount cannot get you under it. A ten-person team at 90% accuracy reaches 5% once ρ drops below about 0.17.

The model assumes errors are unbiased around the truth. Consistent optimism produces something worse, because the whole forecast would also be off in the same direction on top of the correlation, so treat these figures as the conservative version.

AMplify works with Account Management and Customer Success teams on the three jobs of a Post-Sale organization: Keep Customers, Grow Accounts, and Deliver No Surprises. This paper is about the third one.

Alex Raymond is the founder of AMplify and the author of *The Growth Department*. He hosts the Account Management Secrets podcast, where the research cited in this paper first appeared.