SIGNAL TO STRATEGY · WEEK 2 OF 4

Stress-test the bet

Signal is not a decision. Today you build the layer that turns truth into a call you can defend.

Heather Gawel + Babajide Okusanya · Wednesday September 16, 2026 · 12:00–1:30 ET

Signal is not a decision.

You can be completely right about what customers said and still make the wrong call.

What you have now

A finding. Something true about your customers that you can evidence, cross-referenced across sources, with numbers attached.

What a decision needs

That finding tested against what you already committed to — your strategy, your architecture, your quarter — and against the people who will push back.

What's missing

The layer in between. Most teams don't have one, so the loudest finding becomes the roadmap and nobody notices until the room goes quiet.

A finding is an observation. A bet is an observation plus everything you chose not to do instead.

THE ARC

Signal → Decision → Strategy

Week two is the one that changes how you argue, not just how you analyse.

1

Automate the signal

Done. Your intake runs on your own inputs.

2

Stress-test the bet

Today. Test a real decision against your strategy and the people who push back.

3

Defend the call

ROI measured the way your org already defines success.

4

Close the loop

In-market monitoring on the tools you already run.

Same rule as last week. Your product, your stack, your bet. Or follow along the fictitious LaunchPad example

TODAY

Ninety minutes, one real bet

You leave with your decision system, run end to end on something you're actually carrying.

Sort the signal

Confirms, contradicts, flatters, silent. 14 min.

Watch it get built

The stress-test layer on a live bet, from scratch. 22 min.

Run your own

Your bet, your strategy doc, our help. 24 min.

Kill it cheaply

Draft the test that could prove you wrong. 12 min.

Bring one you're not sure about. The bets you're certain on make for a boring ninety minutes.

PART ONE

The bet that agrees with you

Your own framing is the most sycophantic thing in the room — and it was there before the AI was.

THE TRAP

You framed the bet, then went looking for support

And you found it, because you were looking in the place you already work.

The model agrees

Ask "is this a good idea?" and it says yes. We covered this in the Push Back session — sycophancy is a writing style, not judgment.

Your evidence agrees

You collected it after you had the idea. Signal gathered to support a bet is not the same as signal that produced one.

Your discipline agrees

You called it an editing problem because you were looking at the editor. Every craft turns the problem into the shape of its own tools.

Confident is not correct. And a finding that confirms everything you already believed is the weakest thing in your deck.

THE INSTRUMENT

Four verdicts. Every signal gets exactly one.

Not a score. A verdict — against the strategy you already committed to.

Confirms

The signal and the strategy point the same way. Useful, and the least informative of the four.

Contradicts

They point in opposite directions — or the bet quietly breaks a commitment you already made. This is the valuable one.

Flatters

It looks like support. It's only support because you framed the bet to match something you already believed.

Silent

Nothing says anything, and you've been filling the gap with assumption. Most silences are a phone call, not a quarter.

If everything confirms, that's not validation. That's a bet you framed to win. Go find what contradicts it.

THE ONE TO LOOK AT

Flatters is where the money is

It doesn't fail loudly. It fails as agreement.

What the data said

Flows in churned accounts sit untouched for 86 days. In healthy accounts, 18. Every account idle past 60 days is churned or at risk — eight out of eight.

What we read into it

"Flows are hard to edit." Nothing in the data says that. The data says flows are not being edited. The bridge between those two sentences is a belief.

What it cost

One reading costs nine engineering weeks. The other costs an email. Nobody asked which was true before writing the memo.

Agreement is not evidence. When your design peer nods, ask whether they're agreeing with the principle or the diagnosis.

THE LOOP

Five moves, one bet, one chat

Each one is a move, not a magic string. You're copying the reasoning, not the words.

1

Restate

Make it prove it understood you — and name the assumptions you never stated.

2

Critique

Make it find the holes a skeptical reviewer would. Highest-leverage move.

3

Constrain

Tie it to real money, time and headcount. Make it say what has to give.

4

Compare

Top three paths side by side. No straw men. The fact that would flip the pick.

5

Challenge

Argue the strongest case against your own number one — and for doing nothing.

Almost everyone stops at Critique. Challenge is where the surprises live. Don't let it soften the case against itself.

THE LENSES

Don't invent a stakeholder. Feed it the real one.

"Pretend to be a skeptical CFO" gets you a stock CFO. Your CFO is more specific and much harder.

CEO

Reads risk, and when they'll know it worked. "If the answer is renewal rates in twelve months, that's not a Q4 project — that's a hope."

CFO

Reads payback, and what it displaces. Gives you the test he'll apply, out loud, before you've asked.

Eng lead

Reads scope, and what gets built twice. Usually knows about a commitment your memo forgot.

Your own peer

Reads whether you reframed the problem into your own surface. The most uncomfortable and most useful lens.

Use transcripts, not adjectives. You already have the material — last quarter's review, the Slack thread, your notes.

PART TWO

Walking a bet through the system

A real strategy document, a real memo, and a decision that looks completely reasonable until you test it.

THE SCENARIO

Launchpad turned last week's finding into a bet

Same company. Same data. The team wrote up what they want to build in Q4.

The finding

When the admin who built the onboarding flows leaves, nobody else will touch them. Eight of eleven churned or at-risk accounts had an admin change. $112,900 of ARR behind it.

The bet

Build a guided flow-editing mode — a safe, role-scoped way for someone who didn't build a flow to change it — and ship it in Q4.

The ask

All nine usable engineering weeks and all three weeks of design. The analytics work slips to Q1. It has already slipped from Q2 and Q3.

Read it and it sounds right. It targets the biggest number in the data and it maps to a principle the design team already holds.

THE INPUTS

Four files, and the one nobody loads

The customer data is not the interesting input today. The commitments are.

The strategy

Outcomes, priority order, architectural commitments, experience principles. The thing you agreed to in February and haven't reread since.

The bet

The memo as the team wrote it. Optimistic, well-argued, and quietly missing three things.

The signal

Last week's findings, carried forward with their numbers intact.

The transcripts

How four stakeholders actually talk. Business review, architecture review, design review, a Slack message from support.

Everyone loads the customer data. Almost nobody loads what they already committed to. That's where the contradictions are.

WHAT THE SORT FINDS

All three things it wants to build are blocked

And the memo asks for the weeks. It doesn't ask for the two exemptions it actually needs.

2

signed architectural commitments the bet breaks without mentioning either one

3

pieces of signal that flatter the bet rather than confirm it

4

questions nobody asked, all of them a phone call or a query away

0

in-quarter metrics — the first thing the CEO said she would ask for

The memo asks for nine weeks. It doesn't ask for the two exemptions it needs to spend them. That's what a stress test is for.

THE BEAT

“We decided this was an editing problem because we’re looking at the editor.”

A principal designer said this in the review. Nobody followed up. The support lead had already called those accounts — the people weren’t blocked. Two of them didn’t know the flows could be changed at all.

THE SUBTRACTION

Nobody does the arithmetic in the room

The CFO told you the test he'd apply. Run it before he does.

$190k

loaded cost of nine engineering weeks, his number, said out loud in the review

30%

the recovery rate he said he'd assume, “because that's what actually happens”

$33.9k

what that recovers against $112,900 of at-risk ARR

5×

underwater — and it doesn't clear at 100% recovery either

Then give it the steelman. The board tracks logos, not ARR — and on that metric the bet looks far better. It's not in the memo.

THE OUTPUT

Name the trade-off you've been avoiding

Not a generic one. The specific thing this bet gives up that the write-up doesn't admit.

We'd spend the whole quarter that gates the upmarket move on the lowest-value tier, pay for it with a fourth permission system we already committed to deleting, and not know whether it worked until after the deadline that was the only reason we called it urgent.

Every leg of that is in the files. None of it is in the memo.

If you can't finish that sentence, you haven't stress-tested the bet. You've described it.

PART THREE

Kill it before you build it

The cheapest version of being wrong is finding out in a day instead of a quarter.

THE RULE

If no result could kill it, it isn't a test

It's a demo with a feedback form attached.

Write the kill first

Before you design the test, write the result that would make you abandon the bet. If you can't write one, you're not testing — you're building support.

Cheapest risk first

Order by risk retired per dollar, not by what's fun to build. The test that ends the bet on a commitment you already signed costs ninety minutes.

Date every one

A test with no decide-by date is a research project. Put the date planning locks at the top of the canvas and work backwards.

Most validation isn't validation. It's a prototype shown to people who were never going to say no.

Five columns, and the third one does the work

Assumption · The test · What kills it · Cost · Decide by.

Ninety minutes

Ask Eng and Security, in writing, whether it can ship without a new permission scope. A written no ends the bet on day three.

Under a day

Query whether the frozen accounts even have an active admin. If there's nobody there, the whole design has no user.

Half a day, no calls

Re-read the tickets and NPS you already have, classified for the mechanism: blocked, afraid, didn't know, not my job.

Two design days

Prototype it and watch five people try to publish a flow they didn't build — but only run this if the cheap tests point at fear.

Under 5% of the build cost. And every kill decision lands before planning locks. That's the case you take to your exec.

YOUR TURN

Your bet, your strategy, twenty-four minutes

You will not finish. One contradiction you didn't have this morning is the goal.

Lane A

You have a strategy doc and a live bet. Load both. Run the sort, then moves 1, 2 and 5. Get to the trade-off.

Lane B

You have a bet but no written strategy. Use what exists — a vision statement, experience principles, last quarter's goals. That is a strategy doc.

Lane C

You can't get to either right now. Use the Launchpad pack and change the stakeholder — run it as the agency pitching the client, not the team building it.

No customer PII, and no perfect inputs. A vision statement and three principles is enough to find a contradiction.

PRESSURE-TEST IN PAIRS

Say the trade-off out loud to another person

Not “how did it go.” One sentence, the thing you've been avoiding.

1

State the bet

One sentence. What you'd build and what you believe it changes.

2

Name the trade-off

The specific thing it gives up that your write-up doesn't admit.

3

Take the hardest question

Your partner asks the one your CFO or your eng lead would. You answer or you concede.

If you can't name one, your bet hasn't met a contradiction yet. Go back to the sort and look harder at Flatters.

THIS WEEK

Stress-test, ship one test, post the trade-off

The project is the same shape as the session, run properly instead of in twenty minutes.

1

Run the full loop

Your real bet, all five moves, your own stakeholder transcripts loaded as lenses.

2

Ship one validation test

The cheapest one on your canvas. Actually run it this week — not design it, run it.

3

Post the trade-off

One sentence in the cohort channel. The thing you'd been avoiding naming.

Everything from today goes up tonight. Slides, the working guide, the canvas and the Launchpad bet pack. Free, same as always.

WEEK 3

Defend the call

You'll have a bet that survived a contradiction. Next week you price it in your org's own language.

How your org defines success

Not ROI in the abstract. The metric your leadership already uses, and where your bet lands in it.

What it displaces

Every bet cancels something. Next week you make the displaced thing visible instead of a footnote.

The in-quarter read

What you'll see in thirty days that tells you it's working — before the number anyone's tracking moves.

Bring the trade-off you named today. That sentence is the input to week three.

NEXT STEP

The signal tells you what's true.

It doesn’t tell you what you’re willing to give up. That part is still yours.