Stress-Test Prompt Pack

Signal to Strategy — Session 2 · Momentum Product Co.

Eight prompts, in order, in one chat. Two versions of each: the general one with brackets to swap, and the Launchpad one you can paste as-is.

Before you start: upload your strategy doc, your bet, your signal and your stakeholder notes into one chat. If you're using the Launchpad pack, that's all four files — `00_launchpad_2026_strategy.md`, `01_the_bet.md`, `02_signal_summary.md`, `03_stakeholder_source.md`.

The four verdicts, which prompt 3 uses:

- **Confirms** — signal and strategy point the same way
- **Contradicts** — they point opposite ways, or the bet breaks a commitment you already made
- **Flatters** — it looks like support, but only because you framed the bet to match something you already believed
- **Silent** — nothing says anything, and you've been filling the gap with assumption

---

The prompts

1 · Load it

> I'm about to make a bet and I want it stress-tested, not improved. This workspace has my strategy doc, the write-up of the bet, the customer signal behind it, and transcripts showing how my stakeholders actually talk.
>
> Read all of it before you say anything. Then tell me what decision you think I'm making and what I seem to be optimizing for. Don't recommend anything yet, and don't build a framework.
>
> One rule for this whole conversation: ground every claim in a line you can point to. Never invent a quote or a number. If you need something that isn't here, say it's missing.

2 · Restate

> Restate this bet in your own words. What's actually being optimized, what constraints did the write-up give you, and what assumptions am I making that I did NOT state? List the assumptions separately and make them specific — "we want to reduce churn" is not an assumption, it's filler.

Then pick one assumption, tell it that assumption is wrong, and re-run. That's where the answer moves.

3 · Sort it against the strategy

Go through my strategy document line by line against this bet. Every piece of signal and every claim the bet makes gets exactly one verdict: CONFIRMS, CONTRADICTS, FLATTERS or SILENT.

CONFIRMS — the signal and the strategy point the same way
CONTRADICTS — they point opposite ways, or the bet breaks a commitment I already made
FLATTERS — it looks like support, but only because the bet was framed to match something I already believed
SILENT — the strategy and the data say nothing, and I've been filling the gap with assumption

Quote the specific line or number you're relying on. Put the bet's central claim in CONTRADICTS or FLATTERS if that's where it belongs — don't protect it.

4 · Name the trade-off

In one paragraph: what's the trade-off I've been avoiding naming? Not a generic trade-off — the specific thing this bet gives up that the write-up doesn't admit.

5 · Do the arithmetic

Run the numbers my CFO would run, using the figures in these files. Show the working. Then give me the strongest argument FOR this bet that isn't in the write-up.

6 · Constrain it

Re-run this against my real limits: [time], [budget], [headcount], [technical or compliance constraints], [what's already committed]. If something has to give, say explicitly what you'd trade and why that trade rather than another. Then name one thing I should NOT trade even under pressure.

7 · Stakeholder lenses

Using only how each person actually talks in the transcripts, give me the three hardest questions [the CEO], [the CFO], [the engineering lead] and [a peer in my own discipline] will each ask about this bet — in their own register, with their own numbers. Quote them where the transcript gives you the hook, and never invent a quote. For each question, tell me honestly whether the bet as written can answer it.

8 · Challenge it, then kill it

Argue against this bet. Build the strongest case for doing the opposite, and for doing nothing and waiting two quarters. What would have to be true for this to be the wrong call, what single fact would change the answer, and what option have you been underweighting?

Then draft the cheapest tests that could kill it. Order by risk retired per dollar. For each: the assumption, the test, what result would KILL the bet, cost and elapsed time, and a decide-by date working back from [the date the decision locks]. At least two must run in under a day. At least one must use data we already have rather than talking to customers.

The Launchpad versions

Paste these as-is. Same eight prompts, no brackets to fill.

1 · Load it — same as above, no changes needed.
2 · Restate — same as above.
3 · Sort it — same as above.
4 · Trade-off — same as above.

5 · Do the arithmetic

> Marcus stated the test he'd apply in the Q3 business review. Run it, show the working, and tell me whether this bet clears it. Then give me the strongest argument FOR the bet that isn't in the write-up.

6 · Constrain it

> Re-run this against the real limits: nine usable Q4 engineering weeks after the December freeze, three weeks of design capacity from a designer who is 0.6 FTE and shared with the Scale squad, no new headcount, no research budget, and the February architectural commitments. Say explicitly what has to be traded and why that trade rather than another. Then name one thing that should NOT be traded even under pressure.

7 · Stakeholder lenses

> Using only how each person actually talks in the transcripts, give me the three hardest questions Dana, Marcus, Priya and Tomas will each ask about this bet — in their own register, with their own numbers. Quote them where the transcript gives you the hook, and never invent a quote. For each question, tell me honestly whether the bet as written can answer it.

8 · Challenge and kill — same as above, with "working back from October 1, when Q4 planning locks."

What the Launchpad run found

Not an answer key to read first. Check it after, to see whether your run landed. Your counts won't match ours — two runs of the same pack sorted 25 items and 45. What matters is where the bet's central claim lands, and whether the points below turn up somewhere.

The flatterer, and the point of the whole exercise. Flows sit idle 86 days in churned accounts versus 18 in healthy ones. The data says flows aren't being edited. The bet says flows are hard to edit. Nothing joins those two sentences — the bridge is a belief. One reading costs nine engineering weeks. The other costs an email. Your run may put the 86-day number itself in Silent or even Contradicts, since it's the strongest predictor in the folder and the bet doesn't use it — and put the bet's central claim in Flatters. That's the same finding. The number is fine. The reading is what flatters.

The two other flatterers. "It maps to Experience Principle 5, which the design team already believes" — agreement cited as evidence. And "safe to try," where the principle's actual mechanism is undo; the bet swapped in simulation and kept the principle's name as cover for an unscoped environment.

The contradictions. The new "Flow Editor" role breaks a February commitment signed by Engineering, Security and Product — no new roles or permission scopes before Q2 2027. A new mode inside the editor breaks a second one, so it gets built twice. And the ask puts all of Q4 into priority 3 while priority 2 slips for the third quarter running and priority 1 gates every Scale deal.

The "why now" that might come back Silent. "Q4 is the last quarter that can move the number" often lands in Silent, because the pack has no renewal dates. That's fair. Push on it anyway: the nine weeks are what's left after the December freeze, so a late-Q4 ship is unlikely to change a December 31 number.

The arithmetic. 30% of $112,900 is $33,870 a year, against $190,000 of loaded cost — about a 5.6-year payback. At 100% recovery it's 1.7 years, so it only looks tolerable at the recovery rate Marcus rejects. Marcus never names a payback period he'd accept, so that part is missing, not failed.

The steelman the memo left out. The board tracks logo churn, not revenue. Fourteen per cent to under ten, across ~340 customers, is about fourteen logos, and these eleven accounts are most of that gap — if they're still savable. On the metric the board actually watches, the bet looks far better. That's the bet's best argument and it isn't in the bet.

The number that's true twice. $112,900 is 9 Growth accounts at $36,900 and 2 Scale accounts at $76,000. So the pattern is Growth by headcount and Scale by money — both statements in the memo are true, and they point at different accounts.

The silences, all cheap to fill. Why a non-admin won't touch a flow — nobody was asked. Whether the idle accounts still have an active admin — one query against data already in hand. Whether the 11 admin-change accounts and the 8 idle accounts are the same accounts — one join. When the at-risk accounts renew — a CRM pull. What was different about the four healthy accounts that survived an admin change — four phone calls.

The cheapest killer. Under an hour with Engineering, then confirmed in writing with Security: can this ship without a new role or a new editor surface? A written no ends two of the bet's three pieces this week.