Plan reference: `docs/strategy/EXECUTION_AND_MARKETING_PLAN_2026_08_17.md` §3 M-8.
Status: Source-of-truth Markdown. The rendered PDF is at `docs/sales/FinAdvantage_Outcome_Pricing_Q3_2026_Trial.pdf` (also mirrored to `frontend/public/docs/sales/FinAdvantage_Outcome_Pricing_Q3_2026_Trial.pdf`). CEO + finance to sign off the rendered PDF before sending to a customer.
Why this matters (plan §8.3 sharpening): Of 56 companies in the Aug 18 research file, only **2 publicly ship outcome-based pricing** — HighRadius (*"Mutually Agreed Success Criteria"*) and FloQast (*"No Per-User Fees. Pricing Built Around Value, Not Headcount."*). With this trial, we'd be the **3rd in-market on outcome-based pricing**, anchored on a published QA benchmark (the ±0.5% reproduction rule) — HighRadius and FloQast don't contest that wedge because they don't sell on QA on a fixed narrow deliverable.

The proposal (one page)

For selected customers, on the Quality-of-Earnings Adjusted-EBITDA Bridge deliverable:

- **Implementation fee:** $0.
- **Per-deliverable fee:** $X (one figure, fixed, in proposal).
- **Acceptance criterion (both must hold for the deliverable to count as billable):**
    The Adjusted-EBITDA bridge reproduces within ±0.5% of our published golden-output benchmark for the run (golden = `4ee56aa0`; values 1500 / 1075 / 1389 / 1536 across the four bridge cells), **AND**
    The cell-level provenance review opens cleanly in Excel for every source-trail page — every cell in the deliverable shows the source document, source column, literal SQL chain, and the source-text fragment that produced the cell's value.
- **If the deliverable does not meet both criteria on first delivery, no charge for that run.**
- **If binding fails (deliverable did not arrive), no charge for that run.** *(Aug 19 update from PR `ef9b4ff1`: empty runs surface as `NEEDS_REVIEW`, not `COMPLETED`. An empty result is materially different from a wrong result — both are unbillable.)*
- **Pilot:** 3 customers, 1 quarter each, with a published before/after metric at the close of Q3 2026.
- **Risk to FinAdvantage:** We ship the deliverable without charging if our output is wrong. Our internal defect rate on this deliverable is the basis for the bet — see `docs/reliability/DELIVERABLE_QUALITY_2026_Q3.md`.
- **Risk to customer:** Zero, beyond giving us read-only access to their workbook for one quarter. Standard NDA template.

What this signals to the market

We price on a **measurable promise nobody else in the wedge makes** — *the bridge will reproduce within a published tolerance on our own benchmark*.

Peer	What they publish	Why that's not what we publish
Finsider, Termina	Speed	Not accuracy
Pennylane, Ramp	Workflow breadth	Not QA

Peer	What they publish	Why that's not what we publish
HighRadius	Outcome-based on working capital	Not a judgment deliverable
FloQast	Value-based on broad product	Not a published tolerance
KPMG, Big-4	Enterprise engagement fee	Not a published tolerance on QA

The closest published comparisons to ours are FloQast's value-based pricing (which works on a much broader product) and KPMG's enterprise engagement. Neither is a published-tolerance-on-QA-quoted deliverable. **The covenant we offer is genuinely different.**

What "binding failure" means in this contract (Aug 19 addition)

A binding failure is when the deliverable did not arrive at all — i.e., one or more of the input data files failed to bind to the pipeline's required shape (column names, types, or row structure). After Aug 19 PR `ef9b4ff1`, the executor surfaces this as `NEEDS_REVIEW` (not `COMPLETED`), with a clear payload listing which inputs failed.

For the outcome-pricing trial, this is *not* a defect rate event. A binding failure is materially different from a wrong number — we shouldn't charge for it, but we also shouldn't count it against our defect rate, because the deliverable didn't arrive to be wrong.

If binding fails for > 30% of runs in a customer's quarter, the trial pauses and the customer gets a credit on the next quarter. We don't promise binding works on messy data; we promise when binding works, the bridge is right.

Constraints (per plan §3 M-8)

- **The pricing pilot must not undercut existing agreements.**
- **One** new customer, **one** existing customer who we expand with, **one** customer who walked away on price and comes back.
- No public list of trial customers. The published before/after metric at Q3 close uses aggregate numbers.

Sign-off

- [ ] CEO sign-off
- [ ] Finance sign-off
- [ ] Legal review of NDA template
- [ ] One customer identified in the pipeline to send to
- [ ] Render to PDF and host at `docs/sales/FinAdvantage_Outcome_Pricing_Q3_2026_Trial.pdf`

Provenance

- **Where the ±0.5% reproduction rule came from:** E-2 PR-1 (per-cell provenance sidecar), Appendix D.1 of the plan. The rule itself is a public commitment: any deliverable must reproduce within ±0.5% of the published golden values across re-runs with the same input.

- **Where the binding-failure gate came from:** Aug 19 PR `ef9b4ff1` (empty runs surface as `NEEDS_REVIEW`).
- **Where the "3rd in-market" framing came from:** Plan §8.3 — only 2 of 56 companies publicly ship outcome-based pricing.
- **Where the defect-rate baseline came from:** `docs/reliability/DELIVERABLE_QUALITY_2026_Q3.md` — 0.0% on 31 re-runs of the QoE Adjusted-EBITDA bridge in July 2026.