Problem Set 2 — The Private Firm CFO

The AI-Assisted Forecast (Redline Labs)

MGT 816 — Yale School of Management — Fall 2026 — Professor John M. Barrios

---

Out: Monday, September 28, 2026. **Due:** Monday, October 12, 2026, 1:00 pm (Canvas). **Points:** 40 (of the two problem sets, *both* count — 40 + 40).

Groups & format. Work in your S07 workshop group of **2–4**. Submit exactly **three files**: (1) your **Excel workbook** (the model), (2) a **one-page PDF** of your board memo (exported from `one_pager_template.pptx`), and (3) your **workbench submission sheet** (HTML) generated at the course-site PS2 workbench. **Put every group member's name on all three.** Parts A, the checksum certification, the scenario and sensitivity readouts, and the Verification Log are completed **in the workbench** (`https://barrios88.github.io/mgt816-private-firm-cfo/ps2/`); the workbench provides structure only and never checks answers.

AI policy. AI is required on this problem set. Every number in your workbook and memo must be **human-certified** — you must be able to reproduce and defend each figure without the AI. The closed-book final exam tests the same mechanics (churn, NRR, LTV, driver logic, the cash and deferred-revenue walks) **without AI**. **Any number that appears in your memo but cannot be traced to a cell in your model loses 2 points, each.**

This problem set finishes and polishes the Redline Labs FY2027 forecast you began in the Session 7 workshop. Your workshop `starter.xlsx` is the starting state; the full Model Pack (`case_facts.txt`, `drivers.csv`, `starter.xlsx`, `benchmark_sheet.pdf`, `verification_log.xlsx`, `prompt_card.pdf`, `ground_rules.pdf`, `one_pager_template.pptx`) is posted in the Canvas module *S07 Workshop Kit*, and the full Model Pack is also downloadable from the course site's Problem Sets page. Treat the workshop build as a *v0*: nothing in it is certified until you have tied each cell back to a driver and to the model's five checksum cells.

Redline Labs, in one paragraph. Redline Labs sells an AI contract-review copilot to mid-market law firms; it was founded by a former DocuSigner VP of Product. Pricing is hybrid: a $2,500/month platform fee per firm plus usage billed at $2 per document above a 1,000-document monthly allowance. Firms grow their document volume as they mature, 40% of new firms prepay a year in advance at a discount, and the cost of an AI inference falls every quarter. All the exact assumptions live in `case_facts.txt` and `drivers.csv`; use those as your source of truth.

Part A — Unit-economics warm-up (8 points)

A1 — Churn, logo retention, and NRR (5 points)

Exhibit A is Redline's **trailing FY2026 book**: the monthly recurring-revenue (MRR) movement and logo movement for the six months July–December 2026. (This is history, provided for the warm-up; it is a separate dataset from the FY2027 forecast you build in Part B.) Using **only** the figures in Exhibit A:

1. For each of the six months, compute the **monthly logo churn rate** and the **monthly logo retention rate**.
2. For each month, compute the **gross MRR churn rate** and **gross revenue retention (GRR)**.
3. For each month, compute **net revenue retention (NRR)**. Then compute the **compounded NRR for the full six-month window**. State the NRR formula you used and be explicit about what it includes and excludes.
4. In one or two sentences, explain how NRR exceeds 100% even though the book loses roughly 1.5% of its logos every month. Point to the specific line of Exhibit A that drives the result.

Exhibit A. Redline Labs — trailing FY2026 book (Jul–Dec 2026).

	MRR movement ($)					
Month	**Beginning**	**New**	**Expansion**	**Contraction**	**Churned**	**Ending**
Jul 2026	560,000	66,000	14,000	(3,500)	(8,400)	628,100
Aug 2026	628,100	71,000	15,500	(3,800)	(8,900)	701,900
Sep 2026	701,900	80,000	17,200	(4,100)	(11,300)	783,700
Oct 2026	783,700	86,000	19,000	(4,400)	(11,700)	872,600
Nov 2026	872,600	96,000	21,000	(4,700)	(14,600)	970,300
Dec 2026	970,300	106,000	23,500	(5,000)	(14,800)	1,080,000

	Logo movement (number of firms)			
Month	**Beginning**	**New**	**Churned**	**Ending**
Jul 2026	200	24	3	221
Aug 2026	221	26	3	244
Sep 2026	244	30	4	270
Oct 2026	270	32	4	298
Nov 2026	298	36	5	329
Dec 2026	329	40	5	364

A2 — Spot the AI's errors (3 points)

A group member asked an AI to summarize Exhibit A. It returned the draft below. **It contains exactly two mistakes** — one is a **formula error**, the other is a **hallucinated benchmark**. Find both.

> ***AI-generated draft — Redline retention, FY2026 book (for review).***
>
> *Redline's book expanded steadily through the back half of 2026. Logo churn held near 1.5% per month across the window, so logo retention averaged roughly 98.5%. Gross revenue retention stayed just under 98%, because a handful of accounts contracted and a few churned, but usage expansion from maturing firms more than offset both. December net revenue retention was **111.3%**, calculated as ending recurring revenue ($1,080,000) divided by beginning recurring revenue ($970,300) — a very strong result driven by the tenure-based usage curve. For context, according to the OpenView 2026 SaaS Benchmarks Report (Table 9, p. 47), the median net revenue retention for Series A legal-technology companies was precisely **118.6%** last year, so Redline is tracking modestly below the peer median but well above the 100% line that separates expanding from shrinking books.*

1. Identify the two errors. Name which is the formula error and which is the hallucinated benchmark.
2. Correct each one. For the formula error, give the corrected figure and the correct formula. For the hallucinated benchmark, state precisely why it cannot stand and what you would do with it (cite the provenance rule from `ground_rules.pdf`).

Part B — The build (18 points)

Complete the FY2027 monthly (Jan–Dec 2027) driver-based model you started in the workshop, in your Excel workbook. Build directly on the `starter.xlsx` skeleton (tabs `Assumptions`, `Customer & Revenue`, `Headcount`, `Income Statement`, `Cash Walk`, `Deferred Revenue`, `Benchmarks`, `Checks`). Every driver comes from `case_facts.txt` / `drivers.csv`; honor every rounding convention stated there so your build ties out exactly.

B1. Customer and headcount schedules (6 pts). Build the monthly customer schedule (beginning firms + new logos − churned = ending firms), splitting new firms into monthly-billing and annual-prepay

cohorts, and tracking each cohort up the three-step tenure usage curve. Build the monthly headcount schedule — AEs (with the one-month ramp), customer-success managers (at the stated ratio, hired one month early), engineers (with the 60/40 support-vs-core split), and execs. Your `Checks` tab has five named checksum cells (`CHK1`–`CHK5`); `CHK1` (December ending active firms) and `CHK3` (December total headcount) must tie out exactly. *The workshop starter is unverified — audit every row against its driver and against the stated ratios before you trust it.*

B2. Income statement and cash walk (6 pts). Build the monthly income statement (platform revenue, usage revenue, COGS — inference, hosting, and CSM payroll — gross profit, S&M, R&D, G&A, operating income) and the monthly cash walk from the $6.6M opening balance (collections in, cash costs out, including commissions paid quarterly). `CHK2` (FY2027 total revenue) and `CHK5` (December ending cash) must tie out.

B3. Deferred-revenue roll-forward (3 pts). Build the monthly deferred-revenue schedule for the annual-prepay cohort (beginning balance + new prepay billings − revenue recognized = ending balance). `CHK4` (December deferred-revenue balance) must tie out. In one sentence, explain the cash-vs-revenue wedge this account creates.

B4. Model hygiene (3 pts). All assumptions live on the `Assumptions` tab and are referenced by cell — **no hardcoded numbers buried inside formulas**; every rate and quantity carries a labeled unit; schedules read left-to-right by month with clean beginning/ending links. A grader should be able to change one assumption cell and watch the model re-flow.

Part C — Scenarios, sensitivity, and the board memo (10 points)

C1. Three scenarios, through the drivers (4 pts). Build **base**, **upside**, and **downside** scenarios by changing **only** these four drivers on your `Assumptions` tab: (i) new-logo quota per AE, (ii) logo churn rate, (iii) documents/firm/month, and (iv) inference cost/document. Scenarios must flow **through the drivers** so the model re-emits every output; do **not** plug scenario numbers into the output rows. For each scenario, report end-of-FY2027 cash and months of runway, and give a one-sentence, case-grounded rationale for each driver value (no new facts).

C2. Sensitivity (2 pts). Produce a one-variable sensitivity table (or tornado) showing how December ending cash responds to each of the four drivers, holding the others at base. **Rank the four drivers** by their effect on December cash, and confirm the ranking with your own Excel data table — not with the AI's assertion.

C3. One-page board memo (4 pts). Using `one_pager_template.pptx`, write a **one-page** board memo. State the **headline**, the **base case**, the **single biggest risk**, and a **pre-agreed re-forecast trigger** (the driver and threshold that would make you re-plan). Attach the one-pager exhibit: the active-firm/ARR trajectory chart, the base/upside/downside end-of-year cash table, the top-five assumptions with sensitivity ranks, and your six-KPI dashboard row. **Every number in the memo must exist in a model cell (the −2 rule applies).**

Verification Appendix (4 points)

Complete the log in the workbench (the `verification_log.xlsx` template is the offline fallback and shows the same columns). The log documents, row by row: the prompt used → the AI's claim → the check you performed → the result (confirmed / corrected / rejected) → the change you made to the model.

1. **Completeness (2 pts).** The log covers the Five Manual Checks: extraction fidelity, roll-forward integrity, benchmark provenance, sensitivity replication, and the December column recompute.
2. **Documented catches (2 pts).** At least **two** entries show a real AI error you caught and corrected (a hallucinated benchmark, an invented assumption, a ramp/lag slip, or an internal inconsistency), *or*

affirmative evidence of a check run clean if the AI happened not to err on that line.

Reminder. Three files per group — your Excel workbook, a one-page board-memo PDF, and your workbench submission sheet — each carrying every member's name, uploaded to Canvas by Monday, October 12, 2026, 1:00 pm. AI is required; every number is yours to certify.