

AI for FP&A — Ground Rules

Redline Labs forecast build (PS2) • MGT 816 • The AI is your analyst; you are the model’s owner and its only certifier.

The four roles — and each one’s blind spot

1. Analyst

driver extraction

Blind spot. Invents assumptions the case never states, quietly filling gaps instead of flagging them.

Your check. Trace every extracted assumption back to the case text. Anything not in the text belongs under OPEN QUESTIONS, not in the model.

2. Skeptic

assumption benchmarking

Blind spot. Cites external benchmarks it cannot source — a hallucinated “median churn” becomes a driver and compounds for 12 months.

Your check. Every external number is traced to the benchmark sheet or a checkable source, or it is stamped **UNVERIFIED** in the log.

3. Builder

schedule logic drafting

Blind spot. Ramp, lag, and hire-early timing slips — an off-by-one month that every roll-forward still balances around.

Your check. Hand-recompute one month, then confirm beginning + adds – churn = ending across all 12 months.

4. Board Red-Team

forecast critique

Blind spot. Softens, or misses internal inconsistencies unless pushed — revenue vs. sales headcount, CAC vs. spend, support ratio vs. customer count.

Your check. Test the “most sensitive driver” claim with a one-variable data table, and recompute December top to bottom.

The protocol (memorize it)

context → role → constraint → output format → verification demand. The prompts change with the tools; this protocol does not.

The Verification Log (graded)

Every AI claim you rely on gets a row. Columns:

prompt used	AI claim	check performed	result (confirmed / corrected / rejected)	change made to model
-------------	----------	-----------------	---	----------------------

Two rules that are not optional

- **Provenance.** Every external benchmark the AI cites is traced to the professor’s benchmark sheet or a checkable source, or it is stamped **UNVERIFIED**. An unsourced statistic never becomes a driver.
- **The AI never touches the .xlsx.** Paste facts in, paste logic out. The AI hands you schedule logic in words — never cell formulas. You build every cell; you own the spreadsheet.

The integrity frame. AI is **required** on PS2. **Every number is human-certified** — you sign off on each figure, not the AI. The final exam (200 of 400 points, **closed book**) tests this same material **without AI**. That is why the skill you are building is verification, not prompting.

PS2 out Mon Sep 28, 2026 • due Mon Oct 12, 2026, 1:00 pm • groups of 2–4 • submit the Excel workbook + the memo one-pager PDF + the workbench sheet on Canvas.