CareerSwami · Business Analytics & Intelligence

FREE TWO-WEEK INTERNSHIP

Project 1 — Six months of growth, and the month the business stops working

Revenue quadrupled in six months. So did the cost. In month six the business earns eighty thousand and spends seventy-nine. Every chart points up and the gap has nearly closed. Your job is to find out which part of this is real.

2 weeks	1 project	6–8 hrs
START TO SUBMIT	ONE SUBMISSION	PER WEEK

WHAT IS IN THIS BRIEF

1. The situation — the problem you are being handed
2. The four traps in this dataset — the cases your work must survive
3. What you are given — your data and starting files
4. What finished looks like — the two things you hand in
5. How the two weeks run — week by week
6. Before you start — rules, tools and your starter kit
7. How the work is judged — the marking scheme
8. How to submit — one link, pasted into your dashboard

Read 1 to 4 before you do anything. They take five minutes and they are what the marking scheme is built on. Sections 5 to 8 are the how — come back to them once you know what you are solving.

1 The situation

Treat this as a real client brief. The subscription club is fictional, but the dataset is built around the commonest misreading in business analytics — growth that is not profitable and cannot be told apart from growth that is.

A subscription club has six months of data: new users, active users, paying users, revenue, cost and support tickets. New users have gone from two hundred a month to four hundred and fifty. Active users from a hundred and sixty to eight hundred. Revenue from twenty thousand to eighty thousand.

Cost has gone from twenty-six thousand to seventy-nine thousand. The business lost six thousand in month one and makes a thousand in month six, which looks like a turning point. But look at paying users: a hundred and forty in month four, a hundred and fifty in month five, a hundred and sixty in month six. The paying base has nearly stopped growing while active users added two hundred and forty.

Support tickets are the other signal: eighteen in month one, a hundred and thirty in month six. Per active user that is roughly flat; per paying user it has more than doubled. Which of those two is the right

denominator is a judgement, and your KPI tree has to make it.

> WHAT YOU ARE BEING ASKED TO DO
>
> **Build the KPI tree and the definitions that make this data readable, then say which of the growth signals are real and which are artefacts of a denominator nobody chose carefully.**

2 The four traps in this dataset

Each of these produces a confident wrong answer. Your metric definitions exist to make the right reading possible, and your memo should name every one you avoided.

Conversion to paying is collapsing

Paying 40 → 75 → 105 → 140 → 150 → 160 against active 160 → 800.

Why the obvious reading is wrong — Paying users are growing, so the chart looks fine. As a share of active users it has gone from 25% to 20%, and the last two months added ten paying users each against hundreds of actives. The product is acquiring people who do not pay.

Cost is total, not acquisition spend

The file gives one cost number per month.

Why the obvious reading is wrong — Dividing total cost by new users gives a figure that looks like customer acquisition cost and is not — it includes serving everyone else. Call it cost per new user if you like, but never call it CAC, and say why in the dictionary.

No cohorts, so no retention

Active users is a monthly snapshot, not a tracked group.

Why the obvious reading is wrong — You cannot tell whether month six's eight hundred actives include month one's hundred and sixty or have entirely replaced them. Retention and lifetime value are therefore not computable from this data. Any figure you produce for them would be invented.

Month six looks like break-even

Revenue 80,000, cost 79,000.

Why the obvious reading is wrong — A thousand rupees on eighty thousand of revenue is noise, not a turning point. And the trend that produced it — costs rising with actives while revenue rises only with payers — points the other way next month. Read the trajectory, not the single month.

The honest headline is that the business is buying users who do not convert, and is one month from the lines crossing the wrong way. Everything in the dashboard says growth; the KPI tree is what lets you say otherwise.

3 What you are given

Six rows of data and no cohorts. The definitions, the denominators and the honesty about what is missing are the entire assignment.

The columns you are given

`month` `new_users` `active_users` `paying_users` `revenue` `cost` `tickets`

Two metrics, worked, so the denominator problem is concrete

`Month 1`	160 active · 40 paying · 18 tickets
`Month 6`	800 active · 160 paying · 130 tickets
`Paying ÷ active`	25% → 20% — conversion is falling
`Tickets ÷ active`	0.11 → 0.16 — modest rise
`Tickets ÷ paying`	0.45 → 0.81 — nearly doubled
`Which is right?`	Your judgement — but you must state it and justify it
`Cannot be computed`	Retention, LTV — no cohort data exists

Four things the KPI tree must connect

`Acquire` `Activate` `Convert` `Serve`

All figures are fictional. Cost is total cost, not acquisition spend, and must never be presented as customer acquisition cost. No cohorts are supplied, so true retention and lifetime value cannot be inferred and must not be estimated. Where you choose a denominator, state the choice and the reason.

4 What finished looks like

At the end of two weeks someone should be able to read your definitions and compute every number you quote, without asking you what you meant.

- **A brief, a KPI tree and a metric dictionary** — the decision this analysis supports and who makes it; a tree showing how the top-level outcome breaks into the metrics that drive it, so you can see which branch is failing; and a dictionary defining every metric by its exact numerator, denominator and period.
- **A workbook with the metrics calculated and their definitions attached** — every metric computed from the raw data by formula, with the definition visible next to it rather than buried in a separate document.

The dictionary is the deliverable people underestimate. "Active user" means three different things in most companies, and the arguments that follow cost more time than the analysis did.

5 How the two weeks run

Work through it in two parts — but you submit **once**, at the end, not once per week.

PART 1 · WEEK 1

Decide what you are measuring and why

Start from the decision, not the data. Build the KPI tree downward from the outcome that matters, and define every metric precisely enough that two people would compute it identically.

Produce

- A brief: the decision, the decision-maker, what would change their mind
- A KPI tree from the top-level outcome down to the inputs that move it
- A metric dictionary: numerator, denominator, period, exclusions

> Every metric needs an explicit denominator. "Tickets per user" is not a definition until you say which users. Your dictionary must state that cost is total, not acquisition spend.

PART 2 · WEEK 2

Compute it, and read what it says

Build the workbook with every metric as a formula from the raw data, then look at the trends properly — especially the ones that move in the opposite direction to the totals.

Produce

- A workbook with every KPI calculated by formula from raw data
- Metric definitions attached next to the calculations
- A note on which metrics cannot be computed from this data, and why

> Your note must state that retention and lifetime value are not computable without cohort data. Any figure presented for either, however hedged, is a fabrication.

6 Before you start

- Define every metric with an explicit numerator, denominator and period before you compute anything.
- Do not call total cost divided by new users customer acquisition cost. Name it for what it is.
- Retention and lifetime value are not available here. Saying so is a finding; estimating them is a fabrication.

- Read trajectories, not single months. Month six's thousand-rupee profit is noise on eighty thousand of revenue.
- Where two denominators give different stories, present both and say which you would use for this decision and why.

Your tools

Google Sheets or Excel. Python with pandas is fine if you prefer, but the workbook has to be readable by someone who does not code. No paid tool is required, and no ID document is ever required.

CS_business-analytics-intelligence_Starter_Kit.zip — subscription_club.csv, a KPI tree template, a metric dictionary template, a calculation workbook, an evidence log, and a file index. https://careerswami.com/wp-content/uploads/careerswami-projects/CS_business-analytics-intelligence_Starter_Kit.zip

7 How the work is judged

WHAT IS ASSESSED	WEIGHT
Problem and requirements — did you understand what you were solving	**20%**
The core artefact — the thing you actually made	**40%**
Validation and evidence — how you know it works	**25%**
Clarity and reproducibility — can someone else open and follow it	**15%**

Strong	Complete, correct, evidence-backed and easy to inspect.
Developing	A useful partial result, with the gaps explained.
Needs revision	Core deliverable missing, claims unverified, or the work cannot be opened.

Credit reflects demonstrated work, not visual polish. A small, honest result beats a large, unsupported one.

What counts as evidence

expected result → actual result → the workbook cell and the metric definition → conclusion. Then say what the evidence does not establish.

Why this one is worth doing. Analytics interviews hand you a chart that goes up and ask what you think. The weak answer describes the chart. The strong answer asks what the denominator is, notices that the paying base stopped growing three months ago, and says which number is not computable from what you were given. That third part — refusing to produce a figure you cannot support — is the one that gets remembered.

8 How to submit

One link, pasted into your dashboard

1. Put your work somewhere with a link — Google Drive, Figma, GitHub, Notion, Loom. Whatever fits the work.
2. **Set it so anyone with the link can view.** A locked link is the single most common reason a submission comes back.
3. Include a short README or cover page saying what to open first, and where each deliverable lives.
4. Open your dashboard, paste the link into the submit box, and add a line or two on anything you want the reviewer to know.

Name files `CS_business-analytics-intelligence_YourName_v1`. Keep your own copy. A screenshot on its own is not enough for anything interactive. A video is optional — annotated screenshots or narrated slides are fine.

Submitted and heard nothing, or the link came back? Email **support@careerswami.com** with the email address you enrolled with.

Work honestly. Use fictional data and your own work. Credit anything you reused, and disclose AI assistance — including what you checked yourself. Never invent interviews, users, tests or results. Nothing here requires you to spend money, contact real organisations, or deploy anything publicly. Never upload identity documents, real customer records or credentials.

Stuck, confused, or something is broken? Write to us.

`support@careerswami.com`

Or WhatsApp **+91 7383138418**. We answer on working days. Use it for anything at all — a part of this brief that does not make sense, a starter file that will not open, a dashboard link that has stopped working, a deadline you need moved, or a submission you are not sure about.

To get a useful answer fast, spend twenty minutes narrowing the problem down first, then tell us four things: what you expected, what actually happened, what you already tried, and your one specific question. Include your track name and the email you enrolled with.

Asking is not penalised. Submitting partial work with a clearly stated blocker always scores better than a finished answer you copied, and better than a silent no-show.

Optional stretch

Optional, and only once the main work is done — cap it at an hour. Model what month seven looks like if the current trends hold, or specify exactly what cohort data you would need to compute retention. Label it clearly as a stretch.