cube | The Agentic FP&A Platform

The FP&A Prompt Playbook

Eight prompts for finance teams, with the controls, inputs, follow-ups, and checks that turn an AI answer into one you can use.



FOREWORD

Start with a better brief

How to use this playbook alongside the FP&A Prompt Cheat Sheet.



By Christina Ross
Founder & CEO, Cube

Most finance teams have tried AI by now. They paste a trial balance into a chat window, type a one-line request, and get back an answer that sounds confident and is hard to trust. The cheat sheet you picked up lists eight prompts that work. This playbook shows you what to put around them.

Each chapter takes one prompt from the cheat sheet and covers five things: what to give the model, how to extend the prompt, what to check before the output leaves your desk, which controls apply, and where the result fits in your close, forecast, or board cycle. The prompts are written for general-purpose AI tools, so you can try them today with whatever your company has approved.

The first chapter covers the controls that generic AI prompting leaves out: security, auditability, traceability, governance, and maintenance. They come first because they decide whether any prompt in this playbook is safe to run on real financial data. The second chapter covers the anatomy of a good finance prompt, and chapters 3 to 10 cover the eight cheat sheet prompts in depth, each with a control check. A companion book, The FP&A Prompt Library, adds more than seventy prompts and covers write-back.

“Treat a prompt like a brief for a new analyst. The more context it carries, the better the work that comes back.”

Christina

Contents

CHAPTER 1 **Prompting with controls** 4
Security, auditability, traceability, governance, and maintenance, before the first prompt

CHAPTER 2 **How to prompt for finance** 6
The six parts of a reliable prompt and three habits that improve every answer

CHAPTER 3 **Variance analysis** 8
Prompt 1. Rank what moved and separate timing from permanent change

CHAPTER 4 **Board summary** 10
Prompt 2. Turn results into five bullets a board can act on

CHAPTER 5 **Scenario forecast** 12
Prompt 3. Build three cases and find the assumptions that matter

CHAPTER 6 **Headcount planning** 14
Prompt 4. Model fully loaded cost, timing, and ramp

CHAPTER 7 **Cash and burn** 16
Prompt 5. Find the drivers of burn and flag what looks unusual

CHAPTER 8 **Anomaly check** 18
Prompt 6. Scan a GL export for duplicates, outliers, and miscoding

CHAPTER 9 **Budget kickoff** 20
Prompt 7. Draft the memo that starts the planning cycle

CHAPTER 10 **Tough questions** 22
Prompt 8. Prepare for what the board will ask about your P&L

CLOSING **About Cube** 24
The agentic FP&A platform and how finance teams use it

CHAPTER 1

Prompting with controls



Security, auditability, traceability, governance, and maintenance. Settle them before the first prompt, not after the first audit question.

A prompt that works once is a good afternoon. A prompt your finance team can run every close, defend to an auditor, and hand to a new hire is a process. The gap between the two is made of five controls that a chat window does not provide, and each one is easier to set up before the first prompt than after the first question from an auditor.

The five gaps

GAP	THE QUESTION A REVIEWER ASKS	WHAT GENERIC AI PROMPTING GIVES YOU
Security	Who can see this data, and where does it go?	A copy of your numbers in a tool outside your finance systems
Auditability	Can someone see what was asked, by whom, and what came back?	A chat history that is easy to lose and hard to produce
Traceability	Can each number be followed to the transaction?	Figures with no link to the ledger entries behind them
Governance	Who approves the prompt, the data, and the output?	Personal preferences, so two people get two answers
Maintenance	What happens when the chart of accounts or the model changes?	Prompts saved in private notes that drift out of date

What to do with generic AI tools today

- **Security.** Use an environment your company has approved, with clear terms on data retention and model training. Share the minimum the task needs. Use salary bands instead of named salaries, and masked vendor names where the analysis allows.
- **Auditability.** Keep a prompt log with the date, preparer, prompt version, source file, output, and reviewer sign-off. Save the output next to the file that produced it. Anything that reaches the board or an auditor needs this record.
- **Traceability.** Ask the model to cite the source row or account for every figure, then verify a sample. Reconcile totals to the ledger before the output goes anywhere.
- **Governance.** Name an owner for each recurring prompt and an approver for each output. Keep one definitions sheet for terms such as ARR, headcount, and adjusted margin, and paste it into every prompt that uses them.
- **Maintenance.** Version your prompts and write down what changed. When the model, the template, or the chart of accounts changes, rerun the prompt on a closed period where you know the answer. Review the library each quarter and retire what nobody uses.

WATCH OUT

A model can cite a row that does not exist or attach a number to the wrong account. A citation that looks precise is still a claim. Check a sample of them against the source every time.

BEST PRACTICE

Treat each prompt you rely on like a policy document. It has an owner, an approved version, and a review date. A prompt without those three things is a personal habit.

CHAPTER 2

How to prompt for finance



The six parts of a prompt that finance teams can rely on, and the habits that separate a usable answer from a plausible one.

A finance prompt works like a brief to a new analyst. Hand over a vague request and you get a vague answer, delivered fast and with total confidence. Hand over context, data, and a format, and you get something you can edit instead of redo.

Anatomy of a finance prompt

Every prompt in this playbook has the same six parts. The cheat sheet versions are short because they assume you will fill in the rest. Here is the rest.

PART	WHAT IT DOES	EXAMPLE
Context	Tells the model who you are and what the numbers represent	200-person company. Figures in USD thousands. Fiscal year ends June.
Task	States the one thing you want done	Rank the top five variances by dollar impact.
Data	Supplies the numbers, with labels and period	Budget and actuals by department, April through June.
Format	Defines the shape of the answer	A table with columns for account, variance, and driver.
Checks	Asks the model to show its work and flag gaps	Show each calculation. Mark anything you inferred.
Controls	Limits what data you share and records what was run	Use salary bands, not names. Log the prompt version and source file.

Weak prompt, strong prompt

Here is the same request written two ways. The first produces a general essay about budgets. The second produces something an FP&A manager can review.

WEAK PROMPT

Analyze our budget.

No period, no data, no format. The model fills in the blanks with guesses.

STRONG PROMPT

Context: 200-person company, USD thousands, fiscal year ends June. Task: compare Q4 actuals to budget by department and rank the five largest variances by dollar impact. Data: the table below. Format: a table with account, variance, favorable or unfavorable, and likely driver. Checks: show each calculation and mark anything you inferred.

Every part is present, and each one narrows the range of possible answers.

Three habits that change the output

- **Ask for the calculation.** Language models can make arithmetic errors, especially across long tables. A request to show each step makes errors visible and gives you something to audit.
- **Separate fact from inference.** Tell the model to label each explanation as supported by the data or assumed. You find out fast where the data ends.
- **Iterate in the same thread.** The first answer is a draft. Follow-ups such as 'exclude one-time items' or 'rank by percentage instead' take seconds and usually improve the result more than rewriting the original prompt.

WATCH OUT

Check your company's AI policy before you paste financial data into any tool. Use an approved environment, remove personal data such as employee names and salaries where the task allows it, and never paste credentials or bank details. Chapter 1 covers the wider set of controls.

CHAPTER 3

Variance analysis



Prompt 1. Rank what moved, explain why, and separate timing from permanent change.

Variance analysis is the task most finance teams hand to AI first. It has clear inputs, a defined output, and a reviewer who knows the answer well enough to catch mistakes. That makes it a good place to learn how prompting works.

The prompt

PROMPT 1 · VARIANCE ANALYSIS

Compare actuals to budget for [period]. Rank the top 5 variances by dollar impact and write a one-line explanation for each.

From the FP&A Prompt Cheat Sheet. Replace [period] with the month, quarter, or year you are reviewing.

What to give the model

- Actuals and budget in one table, by account or department, with the period labeled in the header.
- **A materiality threshold.** Tell the model to ignore variances below a dollar amount or percentage, or the top five list fills with noise.
- **Known one-time items.** A settled legal claim or a delayed vendor invoice changes the story, so list them up front.

Extend it

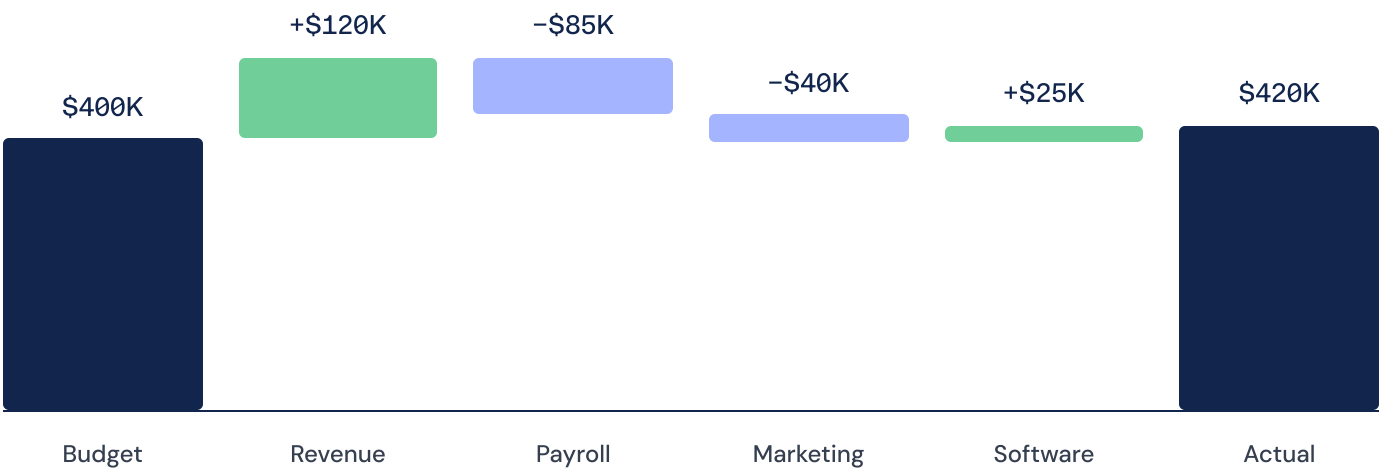
EXTENDED PROMPT

Compare actuals to budget for **[period]** by department. Rank the top 5 variances by dollar impact. Label each favorable or unfavorable and mark whether it looks like timing or a permanent change. Write a one-line explanation for each, then list one question to ask the budget owner.

The timing label and the owner question turn a summary into an agenda for the variance review.

Read the output the way you would read an analyst's first draft. The ranking is easy to verify. The explanations deserve more scrutiny, because the model can only reason from the numbers you gave it. The bridge below shows the walk from budget to actual that a good response reconstructs.

ILLUSTRATIVE EXAMPLE · BUDGET TO ACTUAL OPERATING INCOME ($K)



WATCH OUT

The model does not know that a contract slipped to next quarter or that a campaign launched early. Explanations that sound specific may be inferred. Ask for the evidence behind each one, and confirm the top two with the budget owner.

Before you use it

- Recalculate the largest variance by hand and confirm the sign convention. Models often flip favorable and unfavorable on expense lines.
- Check that every account in your table appears in the analysis.

Control check. Account and department totals usually carry no personal data, which makes this a low-risk first prompt. Log the prompt version and the budget file you used, and keep the output with them.

CHAPTER 4

Board summary



Prompt 2. Turn results into five bullets a board member can act on.

Board members read the first screen of your update and decide how closely to read the rest. A five-bullet summary has to carry the quarter on its own.

The prompt

PROMPT 2 · BOARD SUMMARY

Turn these results into a 5-bullet CFO summary: what happened, why, and what we are doing about it.

Paste your variance analysis, KPI table, or results slide below the prompt.

Give each bullet a job

Five bullets with no assigned purpose drift into five versions of the same sentence. Assign each one a role in the prompt.

BULLET	JOB	STARTS WITH
1	State the headline result against plan	The number and the comparison
2	Name the biggest driver	The driver and its dollar effect
3	Name the second driver or a risk	The item the board will ask about
4	Describe the action underway	The owner and the date
5	State the ask or the outlook	What you need, or what changes next quarter

Extend it

EXTENDED PROMPT

Turn these results into a 5-bullet CFO summary for the board. Bullet 1 is the headline result against plan. Bullets 2 and 3 are the two biggest drivers with dollar impact. Bullet 4 is the action underway, with an owner. Bullet 5 is the outlook or the ask. Put the number first in each bullet and keep each under 30 words. Mark any claim the data does not support.

Specific roles and a word limit produce text you can paste into a deck.

BEST PRACTICE

Write your own five bullets for one quarter before you run the prompt, then compare the two versions. The gaps show you which context the model lacks, and that context belongs in the prompt next time.

AI INSIGHT

A model writes in the same confident tone whether the evidence is strong or thin. The instruction to mark unsupported claims gives you a way to find the weak bullets before the board does.

Before you use it

- Match every number in the summary to the source table.
- Remove adjectives the data has not earned, such as 'strong' or 'healthy'.
- Read it out loud. If you cannot say it to the board without caveats, rewrite it.

Control check. The board relies on this text. Record who reviewed each number, and keep the source table with the final summary so the figures can be reproduced.

CHAPTER 5

Scenario forecast

Prompt 3. Build base, upside, and downside cases and find the assumptions that matter most.



A single-point forecast invites an argument about the number. A set of scenarios invites a conversation about the assumptions behind it, which is where the useful work happens.

The prompt

PROMPT 3 · SCENARIO FORECAST

Build base, upside, and downside forecasts for [metric] using [assumptions]. Flag which assumptions move the number most.

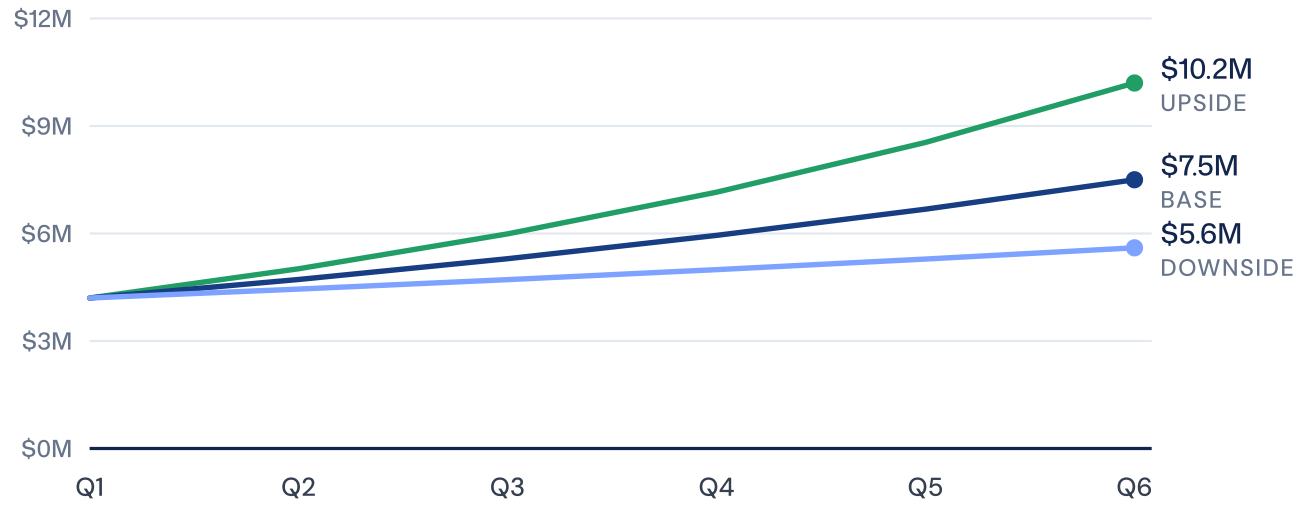
Replace [metric] with revenue, cash, or headcount cost. State each assumption as a plain sentence with a value.

Write the assumptions down

The model can only flex what you name. Give each driver a base value and a value for the other two cases. A small table works well. These inputs are illustrative.

DRIVER	BASE	UPSIDE	DOWNSIDE
New customer wins per quarter	40	50	28
Average contract value ($K)	30	33	27
Annual churn	8%	6%	12%
Sales hiring start	Q1	Q1	Q3

ILLUSTRATIVE EXAMPLE · REVENUE SCENARIOS ($M)



Find the sensitivity

The second sentence of the prompt asks the model to rank your assumptions by their effect on the result. That is a sensitivity analysis. In a model like the one above, churn and win rate often outrank contract value because they compound across quarters. Ask the model to change one assumption at a time and report the effect on the year-end number, so the ranking comes from calculation.

cube

Scenarios are only as current as the actuals under them. Cube's bi-directional connectivity keeps actuals, drivers, and scenarios in one governed place, so the base case starts from the current ledger and not from last week's export.

WATCH OUT

Models tend to produce scenarios that look reasonable and sit close together. Ask for a downside that reflects your worst plausible case, and give it specific inputs.

Before you use it

- Confirm all three cases start from the same actuals and differ from base by the amounts you specified.
- Ask the model to restate its assumptions in the answer, then compare them to your list.

Control check. Save the assumptions table with the output and note who approved each assumption. A scenario nobody owns is hard to defend later.

CHAPTER 6

Headcount planning



Prompt 4. Model the fully loaded cost of new hires, including benefits, timing, and ramp.

Salary is the smallest-looking part of a hire and the only part most first-pass models include. Headcount planning needs the full cost and the timing.

The prompt

PROMPT 4 · HEADCOUNT PLANNING

Model the fully loaded cost of [X] hires in [department] starting [month], including benefits and ramp time.

Replace [X], [department], and [month]. The start month matters more than most people expect.

What fully loaded means

COMPONENT	INCLUDE	COMMON MISS
Base pay	Salary or hourly rate by role and level	Using one midpoint for every hire
Taxes and benefits	Your actual burden rate	Accepting a generic percentage from the model
Equipment and software	Laptop, licenses, and seat costs	Leaving out one-time setup costs
Recruiting	Agency fees or internal recruiter cost	Treating referrals as free
Ramp time	Months to full productivity by role	Counting a new hire at full output on day one

BEST PRACTICE

Timing drives the result. A class of 20 engineers starting in January carries about four times the in-year expense of the same class starting in October.

Extend it

EXTENDED PROMPT

Model the fully loaded cost of **[X]** hires in **[department]** starting **[month]**. Use a **[percentage]** burden rate for taxes and benefits and add **[amount]** per hire for equipment and software. Assume **[number]** months to full productivity. Show monthly cost for the next 12 months in a table, then show the total in-year cost.

Supplying the rates yourself keeps the model from inventing them.

WATCH OUT

Models invent benefit rates when you leave them out. Provide your burden rate and your salary bands, or the output will look precise and be wrong.

Before you use it

- Spot-check one month by hand.
- Confirm each hire starts on the date you gave and not on the first of the month.
- Ramp time reduces expected output. Check that the model applied it to productivity and left cost unchanged.

Control check. Headcount data carries pay. Use salary bands and role counts instead of named employees, and limit who can run this prompt.

CHAPTER 7

Cash and burn



Prompt 5. Summarize what drove cash burn and flag anything that looks unusual.

Burn is a summary number, and a summary number hides its causes. The useful question is which movements drove it and which of them will repeat.

The prompt

PROMPT 5 · CASH AND BURN

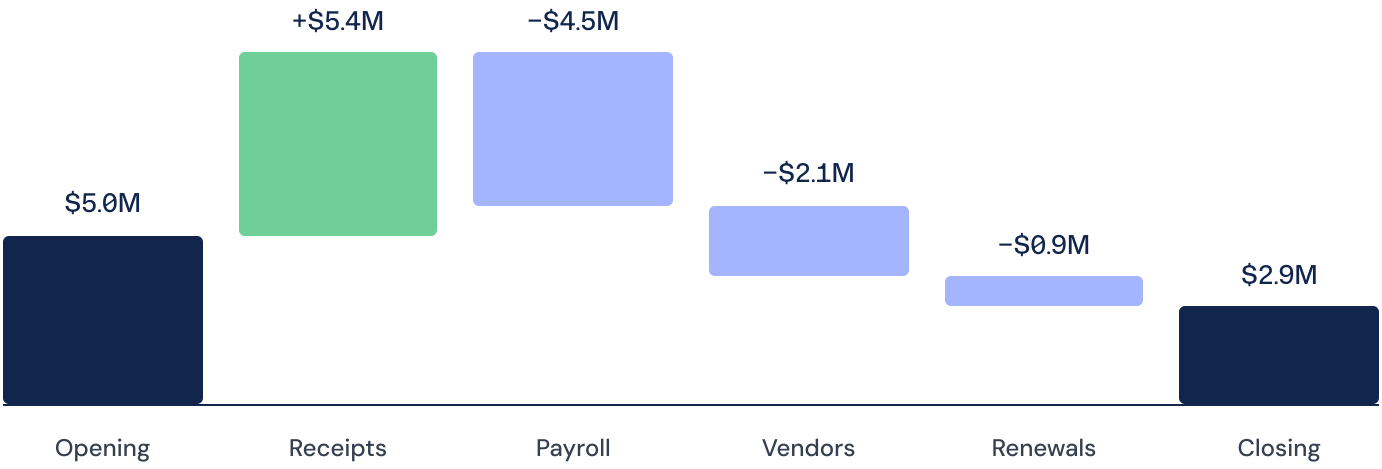
Summarize the key drivers of cash burn over the last 3 months and flag anything unusual.

Paste monthly cash flow, a bank activity summary, or GL detail for your cash accounts.

Separate operating burn from one-time items

A cash bridge makes the answer visible. It starts from opening cash, adds receipts, subtracts each major outflow, and lands on closing cash. One-time items such as annual software renewals or tax payments appear as their own bars.

ILLUSTRATIVE EXAMPLE · THREE-MONTH CASH BRIDGE ($M)



Define unusual

The model has no baseline for unusual unless you give it one. Say what you mean in the prompt.

- **Against history.** Flag any category more than a set percentage above its trailing three-month average.
- **Against plan.** Flag any outflow above its budgeted amount.
- **By timing.** Flag payments made earlier or later than the usual cycle.

Extend it

EXTENDED PROMPT

Summarize the key drivers of cash burn over the last 3 months. Separate operating burn from one-time items. Flag any category more than **[percentage]** above its trailing average. For each flag, say whether it is likely to repeat. End with months of runway at the current burn rate.

The repeat question is the one that changes how you plan the next quarter.

DATA PRINCIPLE

Reconcile the cash balance in your export to the bank statement before you ask the model to explain it. An explanation of the wrong balance is worse than no explanation.

WATCH OUT

Runway depends on which burn rate you use. Ask the model to show runway at the three-month average and at the most recent month, and to say which one it used.

Control check. Cash data is among the most sensitive you hold. Share summaries instead of account numbers, and keep the bank reconciliation on file with the analysis.

CHAPTER 8

Anomaly check



Prompt 6. Scan a GL export for duplicates, unusual amounts, and miscoded accounts.

Month-end close already has reviewers, but they read thousands of lines under deadline. A first-pass scan by a model can surface the lines worth a person's attention.

The prompt

PROMPT 6 · ANOMALY CHECK

Scan this GL export for duplicates, unusual amounts, and miscoded accounts. List each with the reason it was flagged.

Use a transaction-level export with date, account, vendor, amount, and description.

What the model can flag

ANOMALY	WHAT TO ASK FOR	WHY IT MATTERS
Duplicates	Same vendor, amount, and date or invoice number	Double payments and double-posted entries
Unusual amounts	Amounts far outside the account's typical range	Typos, misplaced decimals, one-off charges
Miscoded accounts	Descriptions that do not match the account	Misstated categories and department results
Round numbers	Even-dollar entries in accounts that rarely have them	Estimates or manual entries worth confirming
Timing	Postings on weekends or after period close	Control gaps and late adjustments

Extend it

EXTENDED PROMPT

Scan this GL export for duplicates, unusual amounts, miscoded accounts, round-dollar entries, and postings after period close. For each item, give the row, the amount, the reason it was flagged, and a confidence level. Tell me how many rows you analyzed.

The row count tells you whether the model read the whole file.

WATCH OUT

Large exports can exceed what a model reads in one pass, and it may not say so. Compare the row count in the response to the row count in your file. If they differ, split the file and run it in parts.

cube

A flagged line needs a path back to the entry that created it. Cube lets you trace to the transaction from any number, so confirming an anomaly takes one step instead of a search through exports.

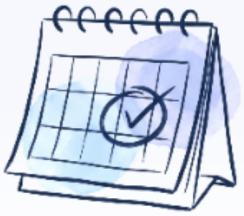
BEST PRACTICE

Treat the output as a review list. Each item needs a person to confirm it, correct it, or clear it, and the outcome should be logged.

Control check. A GL export holds vendor names and amounts. Mask what the task allows, and log each flagged item with who cleared it and why.

CHAPTER 9

Budget kickoff



Prompt 7. Draft the memo that starts the planning cycle.

The budget kickoff memo sets the tone for the next three months. Department heads who receive a clear timeline, a template, and guardrails submit better budgets and need fewer revision rounds.

The prompt

PROMPT 7 · BUDGET KICKOFF

Draft a budget kickoff memo for department heads with a timeline, templates, and guardrails.

Add your planning calendar and company assumptions below the prompt.

What the memo needs

Timeline

Dates for kickoff, first drafts, finance review, executive review, and lock, with the reviewer at each step.

Templates

Which file to use, where to submit it, and how to fill it in. Link one completed example.

Guardrails

Assumptions the company has already set, such as growth target, headcount limits, and merit increase.

ILLUSTRATIVE EXAMPLE · BUDGET CALENDAR

SEP	OCT	NOV	DEC	JAN
Kickoff	**Dept drafts**	**Review**	**Exec & board**	**Plan locked**
Memo sent	Drafts due		Board	

Extend it

EXTENDED PROMPT

Draft a budget kickoff memo for department heads. Company context: **[growth target]**, **[headcount limit]**, **[merit increase]**. Include a week-by-week timeline starting **[date]** with a final lock date of **[date]**, a list of templates with submission instructions, and five guardrails. Keep it under 400 words and use a direct tone.

Your dates and assumptions replace the placeholders. The model supplies the structure and the wording.

BEST PRACTICE

Send the draft to one department head before you send it to all of them. Their first question is usually one the memo should have answered.

WATCH OUT

The model invents dates and guardrails when you leave them out. Provide your calendar and your assumptions, and check every date before the memo goes out.

Control check. The memo sets company assumptions, so have the CFO approve the guardrails before it goes out. Save the version that was sent.

CHAPTER 10

Tough questions



Prompt 8. Prepare for the questions a board will ask about your P&L.

The questions that cause trouble in a board meeting are rarely the ones you prepared for. A model can produce the list you would write on your best preparation day, plus a few you would miss.

The prompt

PROMPT 8 · TOUGH QUESTIONS

List the 10 hardest questions our board could ask about this P&L, with suggested answers.

Paste the P&L, the prior period, and your plan.

Use perspectives

A board is a group of people with different concerns. Run the prompt once per perspective and the lists will differ.

PERSPECTIVE	TENDS TO ASK ABOUT
Lead investor	Growth rate, efficiency, runway, and the path to the next round
Audit chair	Revenue recognition, reserves, controls, and one-time items
Operator director	Unit economics, hiring pace, and spend that is not producing results
Lender	Covenants, liquidity, and downside cases

Extend it

EXTENDED PROMPT

Act as a [role] on our board. List the 10 hardest questions you would ask about this P&L, ranked by how uncomfortable they would be to answer. For each, give a suggested answer that cites the specific line items, and mark any answer where the data is incomplete.

Ranking by discomfort puts the questions you most need to rehearse at the top.

cube

Board answers hold up when every number traces to the source. In Cube, any figure on a board page traces to the transaction behind it, so you can answer 'where does that come from?' in the room.

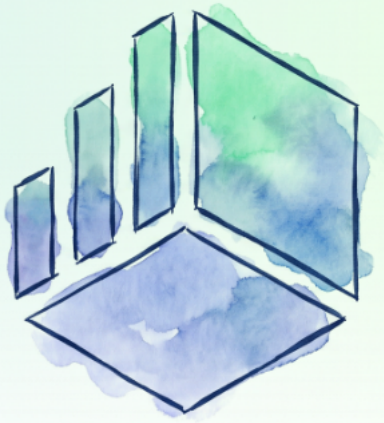
BEST PRACTICE

Rehearse the three most uncomfortable answers out loud with your CEO before the meeting.

Control check. A P&L with board context is confidential. Run it in an approved environment, and keep suggested answers marked as drafts until you have verified each figure.

Ready to put this guide into practice?

Learn more about Cube



The agentic FP&A platform that helps you tell the story behind the numbers. Founded by three-time CFO Christina Ross.

Start your free trial →

Or get a demo

We believe the future of business is finance-driven. Cube connects your business data into one governed model and brings it everywhere decisions happen.

How finance teams use Cube

Close and consolidation

Multi-entity rollups and a faster month-end.

Budgeting and planning

Budgets, headcount, and revenue plans.

Forecasting and scenarios

Rolling forecasts and what-ifs in minutes.

Reporting and board decks

Packs in Excel, Sheets, and slides.

Variance analysis

Drivers explained and traced to the GL.

KPIs and performance

One set of metrics for the business.

AI Agents: Your Always-On Digital Finance Team

Charlie, the superagent, routes work to four FP&Agent teams. People stay in control.

The Data Managers
Integrate and reconcile

The Analysts
Analyze and report

The Planners
Forecast and model

The Business Partners
Shape decisions