cube

THE AGENTIC FINANCE LAYER

How to Use AI for FP&A

The questions to ask before you put AI to work across close, forecasting, planning, and board reporting.

CATEGORY

FP&A · Agentic Finance

BY

Cube

INTRODUCTION

Finance Is the Last Function to Adopt AI

The gap is closing fast. The teams that move first will ask better questions of AI than the ones that wait.

Sales has AI writing emails. Support has AI answering tickets. Engineering has AI writing code. Finance is still closing the books in spreadsheets and writing board commentary by hand at 11pm the night before the meeting.

That gap is closing fast, and it should. Finance held out longer than other functions for good reason: the data is messy, the stakes are high, and a generic chatbot with no connection to your GL will hallucinate a plausible number with total confidence. The question is no longer whether to use AI in FP&A. It's how to use it without inheriting that failure mode. This guide is the checklist — the questions to ask before you trust AI with the numbers.

11pm	Days	1 layer
When board commentary still gets written by hand, the night before	Close-cycle time an AI-drafted first pass can give back each cycle	Where trustworthy finance AI has to start: the governed data layer

> **The purpose of a forecast is not to be right. It is to be useful. The same is true of every AI draft finance produces.**
>
> — Cube

Table of Contents

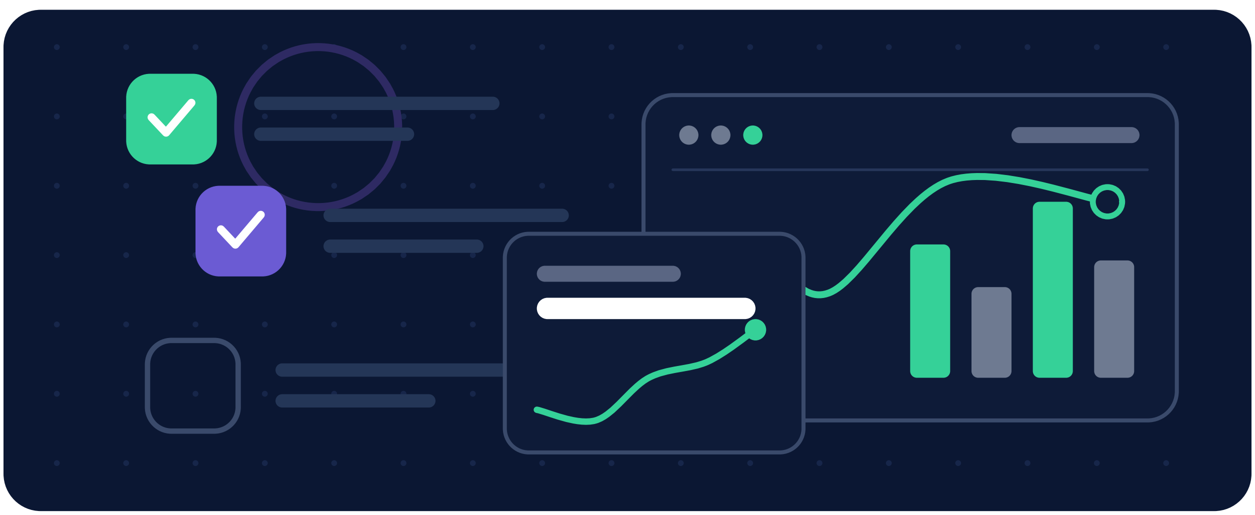

SECTION 01

Foundations and Readiness

Before any AI use case in finance works, four things have to be true. These checks fix the ground you build on.

The single biggest predictor of whether AI in finance succeeds is not which model you picked. It's whether the data underneath it is connected, governed, and traceable. Start here.

1. Do I know where AI actually earns its place in FP&A?

AI is not equally useful everywhere in finance. It earns its keep on the mechanical, repetitive, structured work that currently eats days of every cycle. It does not make the judgment calls. Knowing the difference is the first check, because it sets expectations before a single tool is bought.

HIGH LEVERAGE

Close and variance

First-draft flux commentary, anomaly flags, and reconciliation prep. Mechanical, repetitive, and currently eating days of every close cycle.

HIGH LEVERAGE

Forecasting and planning

Driver-based scenario generation, what-if modeling, and headcount and revenue build-outs that used to take a full cycle to iterate on.

HIGH LEVERAGE

Reporting and narrative

Drafting board commentary, exec summaries, and ad hoc analysis answers in the time it used to take to format a slide.

> WATCH OUT
>
> What is missing from that list: judgment calls. Whether to raise the quota, how to read the board's mood, what a variance means for the business. AI drafts the first pass. You make the call. That division of labor is the throughline of everything that follows.

2. Can I trace every AI-generated number back to its source?

This is the check that separates finance-grade AI from a chatbot. If you cannot point to the GL account, the contract, or the source system behind any AI-generated figure, it isn't decision-ready. It's a guess with good formatting. The architecture that removes finance's core objection to AI is the one that reads your actual GL and shows you the transaction behind every number it produces.

TRACE TO THE SOURCE · CLAUDE + CUBE, LIVE

Why did Q2 opex come in over plan?

CLAUDE

Q2 opex closed at $6.2M against a $5.8M plan, a $400K overage. $340K of that is S and M, driven by one-time event spend in May. R and D and G and A came in within 2% of plan.

CUBE · GL 6000-6900 | Q2 ACTUALS VS PLAN | PO #4471

3. Is our data layer connected, governed, and reconciled before we build on it?

Most finance teams want to skip straight to the top tier: AI that advises on strategy. But an AI tool advising on strategy from disconnected, unreconciled data is worse than no AI at all, because it sounds confident while being wrong. A finance function has to earn its way up a hierarchy, and each layer depends on the one below it being solid.

The Strategic Finance Hierarchy of Needs

BUSINESS PARTNER
Strategic advising, scenario judgment

PLANNING
Forecasting, budgeting, what-if modeling

REPORTING
Variance analysis, board decks, analytics

DATA MANAGEMENT
GL integration, reconciliation, trace to source

BEST PRACTICE

Treat the data layer as the project. The forecasting model, the board-deck generator, the close automation are downstream features. Get the GL connection, the chart-of-accounts mapping, and the reconciliation logic right first, and every later use case gets dramatically easier.

4. Do we know who can see what, and what happens when the data is wrong?

Permissions and access rules that existed before AI still apply after it. An AI tool is not an exception to your access controls; it inherits them. And bad source data produces bad AI output faster and more confidently than it produces bad spreadsheet output, so reconciliation discipline matters more, not less. Decision-ready data has a few non-negotiable traits.

Trait	Why it matters	Common failure mode
Connected to source	AI output is only as current as its inputs	Static CSV exports that go stale within a week
Traceable to the transaction	Every number needs an audit trail back to the GL	A dashboard number nobody can explain in a board meeting
Governed access	AI sees exactly what the user is permitted to see	Broad access granted for convenience, never revisited
Consistent taxonomy	Cost centers and metric definitions must match across systems	Marketing calls it 'programs', finance calls it 'demand gen'

SECTION 02

Close and Forecasting

The close is the easiest place to prove AI out. Forecasting is where it gets the most attention and the most skepticism. Four checks for both.

The monthly close is the most repetitive, most time-pressured, least judgment-intensive part of the calendar. It's the best place to start. Forecasting is the harder, higher-stakes case. These checks keep both honest.

5. Am I using AI where the close is mechanical, not where it needs judgment?

AI reads the actual-versus-plan variance by account and drafts an explanation grounded in the underlying transactions, not a generic template. It flags a duplicate invoice or a miscoded entry before close instead of during it. It pulls supporting detail so the analyst's time goes to judgment calls instead of data assembly. None of that replaces the controller's review. It replaces the blank-page problem.

MONTH-END FLUX · OPEX BY COST CENTER ($K)

A	B	C	D	E
Cost Center	**Actual**	**Plan**	**Var**	**Commentary**
S and M	1,840	1,500	+340	One-time conference spend. PO #4471.
R and D	2,110	2,060	+50	Within plan; contractor timing into next month.
G and A	1,240	1,290	-50	Under plan on lower software renewals.
Total	5,190	4,850	+340	Overage explained by S and M event spend.

COMMENTARY DRAFTED BY CLAUDE · FIGURES GOVERNED BY CUBE · EVERY LINE TRACES TO THE GL

6. Am I treating AI's commentary as a hypothesis, not a finding?

The AI draft will get the mechanics right far more often than it gets the business context right. It is a starting point you correct, not a final answer you trust blindly. The analyst's edit is where the real commentary gets written, and it's what keeps a confidently wrong sentence out of the board deck.

AI DRAFTS

The mechanics

What moved, by how much, and the transaction behind it. Variance math, anomaly flags, and a first-pass explanation grounded in the GL.

YOU OWN

The meaning

What the variance means for the business, whether it's a trend or a one-off, and what the board needs to hear about it.

> **AI INSIGHT**
> The AI draft is a hypothesis, not a finding. Don't let AI-drafted commentary go into a board deck unreviewed. The fastest way to lose the board's trust in your AI rollout is one confidently wrong sentence in a deck nobody double-checked.

7. Does every AI forecast show its drivers and a range, not a single number?

A forecast is a judgment call dressed up as a number, and handing that judgment entirely to a model is a mistake. The opportunity is real too: AI collapses the time it takes to generate and stress-test a scenario from days to minutes. The most useful AI-assisted forecasts don't output one number. They output a range, grounded in explicit assumptions you can interrogate.

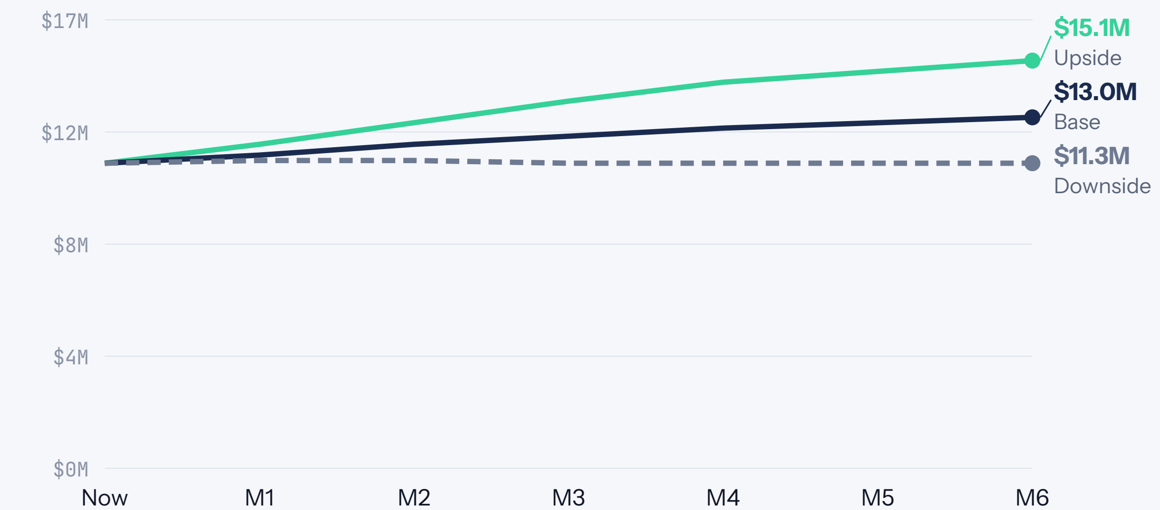

AI INSIGHT

Ask any AI-generated forecast to show its drivers before you trust it. A range with no visible assumptions is not more trustworthy than a single number with no visible assumptions. It just looks more sophisticated.

8. Have I kept the judgment calls — assumptions, macro, curation — with the human?

AI can model the consequence of a 5% higher churn rate. Deciding whether 5% is the right number to model is still your call. A model trained on your history has no view on a hiring freeze, a competitor's pricing move, or a shift in buyer behavior. And three scenarios is useful where twelve is noise. Use AI to widen the set you can afford to test, then use judgment to narrow it back down.

HUMAN

The assumptions

AI models the consequence of a driver change. Deciding which driver values are realistic is a judgment call, not a math problem.

HUMAN

The macro read

Hiring freezes, competitor moves, buyer-behavior shifts. None of it is in the historical data the model trained on.

HUMAN

The curation

Which two or three scenarios actually matter for the decision in front of the board. Curating the story is leadership, not modeling.

SECTION 03

Planning and Reporting

Annual planning is the most time-intensive ritual on the calendar. A board deck is a narrative, not a data dump. Four checks to compress one and elevate the other.

Most of the planning cycle is coordination, not analysis. Most of a board deck is narrative, not numbers. AI is good at the coordination and the first draft. It is not good at the negotiation or the framing. Keep the line clear.

9. Do I know where the planning calendar actually bogs down?

Ask any FP&A leader where the planning cycle stalls and the answer is rarely the modeling. It's the back-and-forth: collecting department inputs, catching the cost center that submitted in the wrong format, reconciling a bottom-up headcount build against a top-down target that doesn't yet agree. AI compresses exactly that mechanical coordination — flag the checkpoints where it helps.

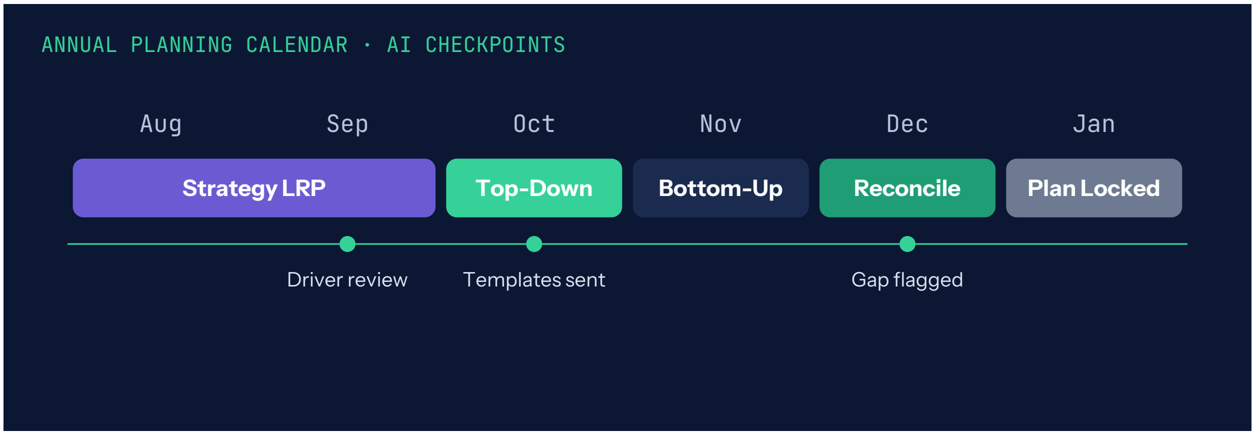

10. Is AI reconciling top-down to bottom-up, or just collecting templates?

Template collection is table stakes. The higher-value check is whether AI can surface exactly where a department's bottom-up build diverges from the top-down target, by how much, and which line items drive the gap. That's the difference between a faster inbox and a faster reconciliation.

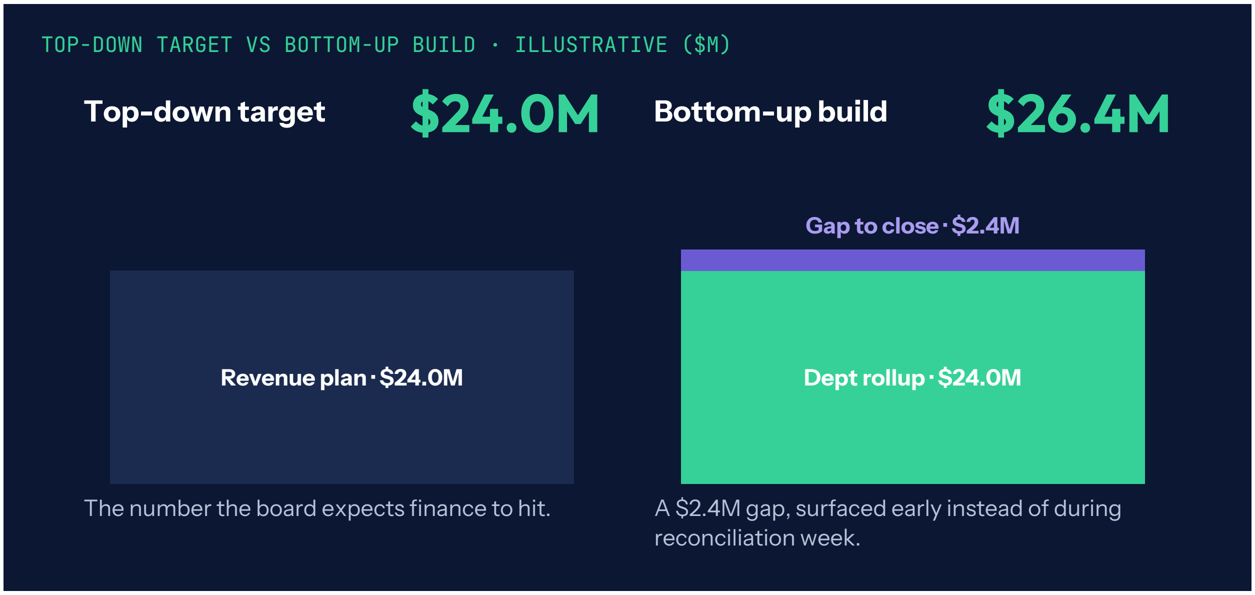

> **CUBE**
> Cube's Planner agent drafts a first-pass departmental budget from prior-year actuals and current-year drivers, then reconciles it against the top-down target automatically, surfacing the gap instead of leaving it for a manual line-by-line comparison.

11. Does the board narrative still carry the CFO's voice?

A board deck is a narrative: what happened, why, and what it means for the decisions in front of the board. The mechanical work — pulling actuals, formatting variance tables, building trend charts — is exactly what AI grounded in your GL handles well. The narrative layer still needs a human voice. Start from a sourced first draft, not a blank slide, then edit.

BOARD-DECK FIRST DRAFT · CLAUDE + CUBE, LIVE

Draft the revenue section of the Q2 board update.

> **CLAUDE**
> Q2 revenue closed at $14.2M, 4% ahead of plan. Net new ARR was $1.8M, driven by expansion in the enterprise segment. Logo churn ticked up slightly in SMB, worth flagging for the retention discussion.

CUBE · GL REVENUE ACCOUNTS | Q2 ACTUALS VS PLAN | READY FOR CFO EDIT

12. Does every variance explanation carry a forward-looking implication?

A variance explanation without a forward-looking implication is a history lesson, not finance leadership. AI can draft the history. The implication is yours to write. Three things stay in the CFO's voice every time, because they depend on board dynamics an AI tool has no visibility into.

CFO VOICE

The ask

What the CFO needs from the board, more runway, a hire, sign-off on pricing, is a strategic decision, not a drafting task.

CFO VOICE

The framing

Whether to lead with the win or the risk depends on board dynamics AI has no visibility into.

CFO VOICE

The tone

Confidence, caution, urgency. The register of the message is a leadership call every time.

SECTION 04

Operating Model and Governance

Adopting AI in FP&A isn't a tool purchase. It's a redesign of how the function spends its time. Three checks, plus a maturity self-assessment.

The function that wins the next decade isn't the one with the most AI tools. It's the one that used AI to clear the mechanical work off its plate fast enough to spend the time it freed up on judgment. These last checks are about the redesign.

13. Have I redesigned where the team's time goes?

Most FP&A teams today spend the majority of their time on data assembly and formatting, and a minority on analysis and business partnering. An AI-native function inverts that ratio. The mechanical work compresses. The strategic work expands to fill the time it frees up. If the ratio hasn't moved, the tools haven't changed anything that matters.

WHERE FP AND A TIME GOES · TODAY VS AI-NATIVE

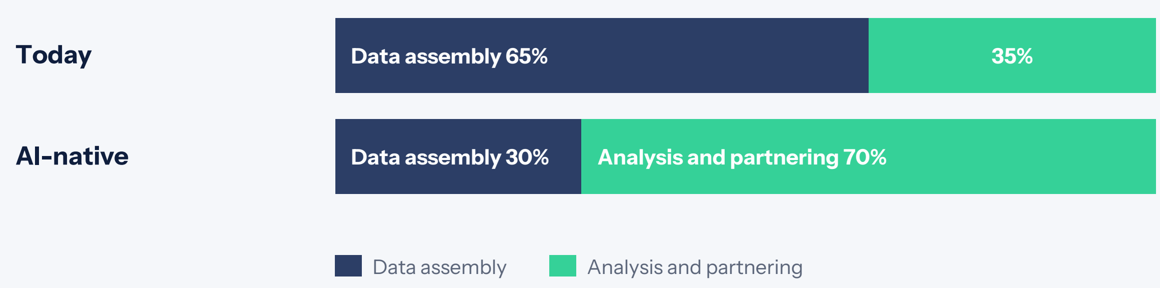

14. Is our AI work organized around a clear model on one governed layer?

Scattered AI point tools recreate the fragmentation finance is trying to escape. The alternative is a small set of agents, each mapped to a layer of the work finance actually does, all built on a single governed data layer so every draft traces back to the transaction behind it. Cube organizes this around four FP&Agents.

The Data Manager
DATA LAYER · GOVERNANCE

Connects to the GL, HRIS, and billing systems. Maintains the chart-of-accounts mapping and governs who can see what.

The Analyst
AD HOC ANALYSIS · ANOMALY

Answers direct questions against live actuals, flags anomalies before they reach the close, and traces every figure to its source.

The Planner
PLANNING · FORECASTING

Drafts first-pass forecasts and budgets, reconciles top-down targets against bottom-up builds, regenerates scenarios on demand.

The Business Partner
BOARD DECKS · NARRATIVES

Drafts board commentary and exec summaries grounded in the GL, ready for a CFO's edit, not a CFO's verification.

15. Are we starting with one workflow, with a human on every external-facing draft?

The teams that stall try to boil the ocean. The teams that win pick the highest-friction workflow, prove it out, and expand from there. Three governance principles keep the rollout safe while it scales.

PRINCIPLE

Start with one

Close commentary or planning reconciliation. Pick the highest-friction workflow and prove it out before expanding to five.

PRINCIPLE

Inherit access

The AI layer respects the same permissions your finance systems already enforce. Access is inherited, not rebuilt.

PRINCIPLE

Human on exit

Board commentary, investor updates, anything leaving the function gets a human edit before it leaves.

Maturity Self-Assessment

Where does your finance function sit today? Be honest — most teams start in Tier 1, and the goal is deliberate movement toward Tier 3, one workflow at a time.

TIER 1

Cautious

AI is banned or used off to the side in a generic chatbot with no GL connection. Every number is still assembled and checked by hand.

TIER 2

Assisted

AI drafts flux commentary and first-pass scenarios on a connected data layer. Outputs are trusted only after a manual trace.

TIER 3

AI-native

The function is redesigned around agents on one governed layer. Every draft traces to the transaction, and time has shifted from assembly to judgment.

CUBE

Finance teams with clean, connected, consistent data can move from Tier 1 to Tier 3 in quarters, not years, because the hard part is the data layer, not the model. That is what Cube, the Agentic Finance Layer, is built to deliver: decision-ready data and AI-native workflows for close, planning, and reporting, all traceable back to the source transaction.

About Cube

The Agentic Finance Layer — built by a three-time CFO, for finance teams.

Cube is the Agentic Finance Layer for the office of the CFO. It connects directly to your GL, HRIS, and billing systems, and gives finance teams decision-ready data and AI-native workflows for close, planning, and reporting, all traceable back to the source transaction.

Cube's four FP&Agents, the Data Manager, the Analyst, the Planner, and the Business Partner, work on a single governed data layer so every draft, forecast, and board narrative they help produce can be traced to the transaction behind it. No black-box numbers. No re-keying the same data into five tools.

Finance-grade AI

Built for the accuracy and auditability finance requires, not adapted from a general-purpose chatbot.

Bi-directional connectivity

Reads from and writes back to the GL, HRIS, and billing systems you already run on.

Trace to the source

Every insight maps back to the GL. No number without a transaction behind it.

The FP&Agent Suite	What It Does
The Data Manager	GL integration, reconciliation, data-quality governance, and trace to the transaction
The Analyst	Ad hoc analysis, trend detection, anomaly flagging, and instant report generation

The FP&Agent Suite	What It Does
The Planner	Planning cycles, rolling forecasts, scenario modeling, and variance tracking
The Business Partner	Board decks, investor narratives, and cross-functional financial communication

Decision-ready data, everywhere you work.

Cube is the Agentic Finance Layer built for finance teams who are serious about AI.

cube

THE AGENTIC FINANCE LAYER