

Is Your AI Budget a Forecast or a Guess?

WHY THIS MATTERS

Uber ran through its 2026 AI coding budget in four months. The tool worked. Adoption success is what blew up the forecast. Most budgets miss in the same direction.

HOW TO ANSWER

From memory, in one pass. If you'd have to look it up, the honest answer is "Not sure," and that is itself the finding.

Scoring: Yes = 2 · Partly = 1 · No or Not sure = 0. Tick one box per row, total each leg out of 6, then turn the page.

LEG 1

Technology

Score, out of 6

	YES	PARTLY	NO / NOT SURE
I know what share of our AI spend is metered by usage rather than fixed per seat.	☐	☐	☐
I could tell what our highest-cost AI users spend per month, and how many users fall into that category	☐	☐	☐
Our budget includes the platform layer: orchestration, retrieval, and guardrails, not just models and infrastructure.	☐	☐	☐

LEG 2

Data

Score, out of 6

	YES	PARTLY	NO / NOT SURE
Data cleaning, integration, and governance work is funded as its own line item, not absorbed by an existing team.	☐	☐	☐
Our data readiness has been assessed against the AI use cases we approved, rather than assumed.	☐	☐	☐
We budgeted for keeping data governed on an ongoing basis, not only for the initial cleanup.	☐	☐	☐

LEG 3

People

Score, out of 6

	YES	PARTLY	NO / NOT SURE
Training and enablement appear in the AI budget, not somewhere else or nowhere at all.	☐	☐	☐
We budgeted for redesigning the workflows the tools sit inside, not just for buying the tools.	☐	☐	☐
We accounted for the temporary productivity dip while teams learn new ways of working.	☐	☐	☐

LEG 4

Process

Score, out of 6

	YES	PARTLY	NO / NOT SURE
Our budget treats AI as an ongoing run-rate cost, not a one-time project with an end date.	☐	☐	☐
We have a plan and a budget for when a model we depend on is changed or retired.	☐	☐	☐
A designated owner is accountable for AI spending forecasts and budget variances.	☐	☐	☐

Is Your AI Budget a Forecast or a Guess?

Your score

TECHNOLOGY	DATA	PEOPLE	PROCESS	TOTAL / 24

0 – 9

GUESSING

Your number is a placeholder. It will move, probably within two quarters, and probably upward. This is the most common result, and the least comfortable to discover during a budget review rather than before one.

10 – 17

TRACKING

You can see what you are spending but not what you are about to spend. Usually one or two legs carry the score while the others are assumed. Look at your lowest leg below.

18 – 24

FORECASTING

You have a defensible number. The remaining risk is concentration. A single 0 on an otherwise strong scorecard is where the surprise comes from.

Your lowest-scoring leg: This is where your next budget surprise will come from. Write it here, and name who owns it.

On “Not sure” answers. Count how many you ticked. Not knowing is not a neutral answer. It means the cost exists and nobody is watching it. Most organizations score lowest on People and Process, which also determine whether the program returns anything at all.

Four questions to take into your next budget review

1. If AI adoption doubled next quarter, can we accurately predict the impact on costs?
2. Have we fully accounted for the ongoing cost of preparing, governing, and maintaining the data that AI depends on?
3. Do we have a budget and plan for training, adoption, and workflow change management?
4. Who owns AI spending outcomes, and what metrics determine success or course correction?

One thing this self-check can't tell you. It measures whether you know your number, not whether the number is right. Getting to an actual figure means pricing all four legs against your real environment. That is a working session, not a questionnaire.

Turn the score into a number

Our AI Foundation Kickstart Workshop takes leadership teams through all four legs against your actual stack, data estate, and adoption curve. It gives your team the clarity, direction, and deliverables needed to turn AI ambitions into practical, measurable next steps.

[\[Book the AI Foundation Kickstart Workshop → \]](#)