WORKBIGLY

AI IN THE BUSINESS · A GUIDE

# What has to be in place before the AI is

Why most rollouts don't move a single number, and what the ones that do are doing differently.

WHERE THIS STARTS

# It's everywhere, and barely used

Nearly eight in ten firms already use AI, and four in five senior executives use it regularly. On average they spend 2.1 hours a week with it, up from 1.5 at the start of 2026. Adoption was fast. Usage is growing, but it's still thin.

In the same study, across the three years generative AI has been available, more than eight in ten executives saw no productivity impact at their own company, and the average gain was 0.2% a year. Three years is early for a general-purpose technology, and it's worth saying that yourself before someone in the room says it for you.

For the next three, those same firms expect 0.6% a year — triple that. It isn't blind optimism: in six months, the gain they report went from 0.1% to 0.2%. But what will separate the firms that get there from the ones that don't is the work they do in-house, not the tool.

NBER w34836, "Firm Data on AI," Revised October 2026 · survey of nearly 9,000 firms in seven countries

# The amplifier is neutral. What it amplifies isn't.

It's the principle that explains why the same tool pays off at one company and not at the one next door.

AMPLIFYING ISN'T TRANSFORMING

# AI makes the company faster at what it already does

If the process is murky, AI will make it faster and murkier. Automating a process nobody ever wrote down doesn't clarify it — it locks the ambiguity in and runs it at scale.

If your data holds three versions of the truth, you'll get all three faster and each one just as confident. And if the team doesn't know what to do with the output, you'll get output nobody uses.

Which leaves one question worth asking before any purchase: not which AI, but what does the company need to be able to do that it can't today. The model gets picked after that answer, not before.

THE BASELINE

# "No measurable impact": a verdict on the tool, or on the measurement?

Before reading it as a verdict, ask whether anyone wrote down how the work was done before. Without a starting number, any later result is an anecdote, and the argument about whether it worked becomes a contest of impressions and volume.

The baseline is the project's first deliverable, not its last. Our view: it's cheaper to take it before the rollout than to rebuild it after one nobody can defend.

You don't have to buy anything to get one. It comes out of records the company already keeps, even if they're on paper or in a spreadsheet.

WHAT TO MEASURE

# Four things you can already measure today

End-to-end time, not time on task. That's the difference between how long someone takes to draft a thing and how long the thing takes to ship.

How much work comes back, and why. Is rework counted anywhere today, and by whom?

How long something takes to get recorded after it happens. That lag decides whether a number is good for deciding or only for explaining afterward.

Where the work waits. Is it where the work is done, or somewhere else? If it's somewhere else, a tool that speeds up production changes nothing there.

THE PARADOX

# Your team feels faster, and the numbers don't show it

A study of around 750 executives documents a systematic gap between the improvement people perceive and the improvement that gets measured. Both readings can be true at once.

The perception isn't lying — it's measuring something else. What drops is the cost of starting: the blank page, the first draft, hunting for the number. That's the part that hurts most, so the relief is enormous.

But if the bottleneck sits further down the line, total time doesn't budge. Example: the draft goes from two hours to twenty minutes and the document still ships in nine days, because six of them are spent waiting on a signature. If the bottleneck is an approval, AI doesn't touch it.

NBER w34984, March 2026 · around 750 executives

# Whoever measures first has the edge over whoever buys first

And the edge compounds, because the second rollout finally has something to be compared against.

BEFORE YOU SIGN

# Three questions that keep a vendor meeting honest

Which baseline is the improvement measured against? If the answer is the vendor's estimate rather than a number from your company, there's nothing to measure yet.

What happens when it's wrong? Not if it's wrong. A vendor who answers "it doesn't get things wrong" is telling you they haven't pushed it hard enough. A useful answer describes the error, how often it happens, and who catches it.

What gets logged about how it reached its answer? If you have to explain a decision six months from now, is what the system saved enough? All three get answered with evidence from your operation, not with a demo — a demo shows you the best day it ever had.

ASSISTANT, WORKFLOW OR AGENT

# The difference is who decides the next step

An assistant answers what you ask and you pick what happens next. A workflow runs steps somebody laid out in advance. An agent decides on its own what to do next to hit a goal.

That changes where control has to sit. When an assistant gets it wrong, it hands you a bad output and somebody reads it before it's used. When an agent gets it wrong, it may already have run several steps: control moves out of the output and into the guardrails.

Neither is better. For most business processes a well-defined workflow performs better and fails predictably. Agents win where the path can't be written down in advance. Picking the most autonomous option by default is paying for risk you don't need.

THE CHECK

# Six questions before you approve the budget

Is the process the AI will automate defined, and followed the same way every day? Is the data it will learn from consistent?

Can the team read and act on the output without a specialist? Close to half of organizations haven't defined when an output needs a human to sign off on it.

Is somebody governing the decisions the system makes, and is there a record of them? Is there a path back from the output to the process and the people? Is success defined as a change in outcome rather than adoption of a tool?

A shaky answer doesn't kill the project. It kills signing before you've fixed it.

McKinsey 2025, on defined processes for human validation

## Start with one, not with five

Five never happen. One happens, and the second one takes care of itself.

**workbigly.com**