


Getting Team Buy-In

How to bring your people with you, not leave them behind



THE RESISTANCE MYTH

Only 4% of non-technical workers actively distrust and avoid AI. Nearly seven in ten are enthusiastic or open to exploring it.

Deloitte AI Institute, 2026 | State of AI in the Enterprise

Jacobs & Xu

Applied AI consultancy

The Licences Were Never the Problem

The subscriptions are bought. The announcement went out. Everyone nodded in the meeting. Six months later the usage figures are dismal and nobody can quite explain why. If that's familiar, there's some comfort in knowing it's the most common AI story in business, and that it has almost nothing to do with the software.

It also has less to do with resistance than people assume. The instinct is to conclude that the workforce is sceptical or set in its ways. The evidence points somewhere much more useful: people are broadly willing, and nobody ever showed them where the tool fits into an ordinary Tuesday.

This guide covers what the research says about why good tools sit unused, who's actually in the room when you announce something new, why trust breaks the way it does and the specific design decisions that turn availability into daily use.


50% growth in worker access to approved AI tools in a year, with the share of them using it daily barely changed (Deloitte, 2026)	30% of staff use AI on average at UK organisations that have adopted it (DSIT, 2026)	4% of non-technical workers actively distrust and avoid AI, against 13% actively seeking it out (Deloitte, 2026)
---	--	--

What's Inside

- 01 The Gap Between Access and Use
- 02 Who's Actually in the Room
- 03 Why Trust Breaks, and How to Protect It
- 04 Designing for Use, Not Awareness
- 05 The Question Everyone Is Thinking
- 06 Measuring Adoption Properly

The Gap Between Access and Use

Start with the single most useful statistic in this whole field, because it reframes the problem entirely.



Access Surged and Take-Up Didn't Follow

Worker access to approved AI tools grew by 50% in a single year. Among the people who have that access, fewer than 60% use it in their daily work, a figure essentially unchanged from the year before (Deloitte, 2026).

That flat line is the important part. Far more people can reach these tools than a year ago, and the share of them who actually use it has barely shifted, which tells you availability alone was never going to be enough. More than four in ten people who have been given access still don't use it in their daily work, and a year of effort hasn't changed that.

The Same Pattern Inside UK Organisations

It shows up domestically too. Among UK organisations that have adopted AI, an average of 30% of staff are using it, and 77% report that fewer than half their people do (DSIT, 2026).

Read that alongside the fact that 53% of those organisations describe their own use as constant. Adoption isn't half-hearted everywhere at once. It's concentrated: intense among a few people, absent for most, in the same organisation, with the same tools available to everybody.

Nobody Wants a Tool

People don't arrive at work wanting to use software. They want the report finished, the query answered, the backlog cleared, the meeting written up. Until a tool is the obvious route to one of

those things, at a specific moment, in a specific task, it stays an icon somebody means to explore when things quieten down.

Things never quieten down. That's the whole mechanism, and it explains why awareness sessions and launch emails move the numbers so little.

This Is Where Value Leaks Away

More than 80% of organisations report no tangible effect on enterprise-level earnings from generative AI, even though 78% were using AI somewhere in the business when McKinsey surveyed them (McKinsey, 2025). Scattered private use by willing individuals doesn't aggregate into anything an organisation can see.

Adoption isn't a soft concern that follows the real work of choosing technology. It's the mechanism by which the technology becomes worth having at all.

Diagnostic Questions

If a rollout has stalled, these usually locate the problem within an hour.

- Can you name the specific task this tool is meant to improve, in one sentence?
- Does anyone's normal working process now route through it, or does it sit beside their process?
- Who was the first person to use it successfully, and did anyone else hear about it?
- What did people have to stop doing when this started? If the answer is nothing, that's the problem
- Are you counting licences issued, or people using it this week?


THE HABITS ALMOST NOBODY BUILDS

Fewer than a third of organisations follow most of the twelve adoption practices tested. Only a quarter of large organisations have a defined roadmap for driving use, and half that among smaller ones.

McKinsey & Company, 2025 | The State of AI

Who's Actually in the Room

Rollout plans are almost always designed for the wrong group. Knowing the actual distribution changes what you do.



The Four Groups

Around 13% of non-technical workers are actively enthusiastic and seeking AI out. 55% are at least open to exploring it. 21% would rather not but will if required. Just 4% actively distrust and avoid it (Deloitte, 2026).

Read that properly and the picture is encouraging. Roughly two-thirds are willing. The group actively avoiding it today is one person in twenty-five. Whatever's blocking adoption in your organisation, it almost certainly isn't hostility.

Why Plans Built for Enthusiasts Fail

The enthusiastic 13% are the people who volunteer, attend the optional demonstration, ask good questions and give encouraging feedback. It's entirely natural to build around them, because they're visible and they make the project feel successful.

But they were always going to adopt it. Designing for them tells you nothing and moves nobody. Meanwhile the willing majority, people who'd use it happily if someone sat with them for twenty minutes and showed them one thing relevant to their actual job, quietly carry on exactly as before.

Design for the Willing Majority

This group needs something specific rather than something general. Not a demonstration of what the tool can do, but a demonstration of what it does for the thing they're doing on Thursday. Not a training session, but twenty minutes alongside someone in their own work.

The reluctant 21% mostly follow the majority once use is normal. They rarely need convincing separately, and pressuring them early tends to harden a position that would otherwise have softened on its own.

What Each Group Needs

Group	Roughly	What they need	What to avoid
Enthusiastic	1 in 8	Somewhere to share what they've found, and a role helping others	Mistaking their energy for organisational progress
Open but waiting	Over half	One relevant example in their own work, and twenty minutes alongside someone	General awareness sessions with no connection to their week
Reluctant but willing	1 in 5	Time, visible normality and no pressure to be early	Singling them out, or making early adoption a performance measure
Actively resistant	1 in 25	An honest conversation about what's actually worrying them	Building the whole programme around winning them over

SECTION 03

Why Trust Breaks, and How to Protect It

People judge AI by rules quite different from the ones they apply to colleagues. A decade of experimental research points to three mechanisms behind that, and each one implies something specific about how you introduce a tool.

HOW TRUST BREAKS AND RETURNS

People are more sensitive to an algorithm's mistake than a person's. One failure can be generalised to every AI system, but giving people control to modify the output restores confidence.

Yalcin Williams and Lim, 2024 | Review of a decade of experimental research

People Assume It Can't See the Particulars

The first mechanism is a belief that an algorithm can't account for qualitative or unique information, which makes its conclusions feel less fair and reduces willingness to act on them (Yalcin Williams and Lim, 2024). The same research finds people struggle to internalise even a favourable decision when a system made it.

The practical implication is about framing. Present AI as something that handles the general case and hands the particulars to a person, because that's both true and precisely the reassurance people are looking for.

One Mistake Can End It

The second mechanism is error sensitivity. People are far more sensitive to mistakes made by a system than by a person, driven by two beliefs: that a colleague will learn from a mistake while a system won't, and that one tool's failure predicts how all such tools behave (Yalcin Williams and Lim, 2024).

The consequence is stark. A colleague who gets something wrong is forgiven within the hour. A tool that gets something wrong once is quietly written off, along with the whole category, and the person rarely tells you they've stopped using it. So set expectations low enough to survive contact. Say the word draft constantly, say plainly that it will sometimes be confidently wrong, and make catching it part of the job rather than a sign the thing has failed.

Give People Control and Show the Human

The third mechanism is opacity. People believe they understand human decisions better than algorithmic ones, largely because they can project their own reasoning onto another person and can't do the same with a system (Yalcin Williams and Lim, 2024). Two things measurably help.

People are more willing to use a system they can modify, and making the human contribution visible in AI-assisted output improves how readily it's accepted. If a person shaped the instructions, reviewed the result or set the boundaries, say so. Concealing the human involvement makes the output feel less trustworthy, not more impressive.

Assisting a Manager Is Not the Same as Being One

One finding deserves particular care. Where an algorithm acts in a managerial capacity, people perceive themselves as having lower status and react negatively to the job itself (Yalcin Williams and Lim, 2024).

There's a real and defensible line between AI helping a manager prepare, and AI appearing to allocate, evaluate or judge. Cross it and you'll damage something considerably more valuable than a rollout.

How to Introduce It

A framing that works. Adapt the wording, keep the five moves.

- Name the specific task it's for, not the technology. "This drafts the first version of the case notes"
- Say what it does badly, before anyone finds out. "It will occasionally state something confidently that's wrong"
- Define the human role as essential rather than remedial. "Your job is the judgement. It handles the typing"
- Say plainly what this means for jobs. Whatever the honest answer is, say it before people ask
- Ask for their help improving it, and mean it. "Tell me what it gets wrong and we'll fix the setup"

SECTION 04

Designing for Use, Not Awareness

Three habits separate organisations where AI gets used from organisations where it gets announced. All are within your control and none are expensive.

THE SKILLS EQUALISER

A short overview of how to prompt lifted task completion to 93%, against 91% for the tool alone and 82% without AI. The consultants who gained most were those scoring lowest before it started, narrowing the gap between skill levels rather than widening it.

Dell'Acqua et al., 2026 | Controlled experiment with 758 consultants

Skills Are the Barrier, and Almost Nobody Is Fixing Them

Insufficient worker skills is named as the biggest barrier to integrating AI into existing workflows, yet fewer than half of organisations are making significant adjustments to talent strategy, and only 20% consider themselves highly prepared on talent, the lowest preparedness score in the study (Deloitte, 2026).

UK organisations say the same thing about themselves: 54% of those already using AI report that limited skills and expertise have held back wider adoption (DSIT, 2026). This is the constraint. It's also the cheapest one on the list to do something about.

Protect the Time to Learn

Organisations achieving real value expect to upskill more than half their people, against roughly a fifth at those getting nothing back, and they're four times more likely to carve out protected time for structured learning. The result is 50% more staff using AI in their daily work (BCG, 2025).

Protected means protected. Learning time that competes with a full workload loses every time, and the loss isn't a scheduling problem. It's a signal about what actually matters here.

Make It Role-Based, Not General

A session showing what AI can do in general produces interest and no behaviour change. A session showing a case worker how to draft their specific report, using a real example from last week, produces a user.

This is also where a small amount of instruction goes a long way. In the controlled experiment, consultants given a short overview of how to prompt effectively produced better results than those given the tool alone, completing 93% of tasks against 91% and scoring higher on quality (Dell'Acqua et al., 2026). A short piece of relevant technique measurably changed output.

Involve People in Reshaping the Work

Organisations achieving value involve their workforce in redesigning workflows around twice as often as everyone else (BCG, 2025). This isn't consultation for appearances. The person doing the task daily knows which exception breaks everything.

The adoption benefit is the larger one though. People who helped design something use it. People handed a finished system work around it, politely and indefinitely.

Watch Who Adopts Early

One caution from the research: early voluntary adopters skew male and more technologically confident (Dell'Acqua et al., 2026). Left to run itself, adoption can widen the gaps already present in an organisation rather than closing them.

That's easily addressed once you're watching for it. Look at who's using it after three months. If the pattern is uneven, go directly to the people who haven't started rather than waiting for enthusiasm to spread on its own.

Thirty, Sixty, Ninety

Period	What to do	What good looks like
First 30 days	Name the task. Set expectations honestly. Answer the jobs question. Run one role-based session with real examples. Choose a small first group who want to try	A handful of people using it on real work, and everyone else knowing what it's for
Days 31 to 60	Protect learning time in the diary. Gather what's awkward and fix the top two. Let the first users show colleagues. Involve people in changing the surrounding steps	Use spreading by demonstration rather than instruction, and the workflow starting to change shape
Days 61 to 90	Check who hasn't started and go to them directly. Remove the steps the tool made unnecessary. Measure weekly use rather than licences	It's becoming the normal way the task is done, rather than an initiative

SECTION 05

The Question Everyone Is Thinking

People will be wondering about their jobs from the moment you mention AI. Whether you address it or not, it will shape everything that follows.

THE FORCE MULTIPLIER

One executive described giving agents to existing workers rather than automating their jobs, calling it a force multiplier. In the early phase, overseeing the agents was more work, not less.

Deloitte AI Institute, 2026 | State of AI in the Enterprise

Silence Is Read as Bad News

Avoiding the subject doesn't keep it out of the room. It concentrates it in corridors, where the least reassuring version wins by default. Whatever your honest answer is, saying it early costs you far less than letting people construct their own.

It's worth knowing the broader picture is more mixed than the headlines suggest. Time saved is most often spent on entirely new activities, though the same research finds headcount reduction

is one of the attributes most strongly associated with value realised, and expectations of reduction concentrate in service operations and supply chain (McKinsey, 2025).

Be Honest Where the Answer Is Uncomfortable

36% of leaders expect at least 10% of roles to be fully automated within a year and 82% expect it within three, while 84% of organisations have not redesigned jobs around AI at all (Deloitte, 2026).

That gap is where anxiety lives. If roles in your organisation will change, say which and by when as soon as you honestly can. People handle change considerably better than uncertainty, and they can tell the difference between a leader who doesn't know yet and one who isn't saying.

More People Using It, Not Fewer

Among UK organisations already using AI, 48% expect more of their employees to be using it within one to two years, 50% expect the same number and just 2% expect fewer (DSIT, 2026).

That's a useful number to have in your pocket, as long as you're clear about what it measures. It counts the people expected to be using AI, not the people expected to be employed. It tells you the tool is spreading, not that nothing about roles will change.

Mind the Route Into the Profession

The tasks automated first, data entry, reconciliation and first-line queries, are often where careers begin (Deloitte, 2026). Remove them without thinking and you've closed an entry route without intending to.

The answer isn't to protect the work. It's to be deliberate about what a first year looks like instead, moving junior people into reviewing and judgement sooner than you otherwise would. Done well that's a better job than the one it replaced, and it's worth saying so.

Answers Worth Preparing

Five questions you will be asked. Decide the answers before you need them.

- "Is my job at risk?" Answer with what you actually know, including that you don't know everything yet
- "Am I being measured on using it?" Say plainly whether use is expected, encouraged or optional. Ambiguity here breeds performative adoption
- "What if I'm slower with it at first?" Confirm that's expected and that the time is protected
- "Who sees what I put into it?" Have a clear answer about data before anyone asks, not after
- "What happens to the work I'm known for?" The most common unspoken worry. Name what their role becomes, not just what it stops being

Measuring Adoption Properly

Most organisations measure the wrong thing and reach comfortable conclusions. A better measure is easy to gather and tells you what's actually happening.

Count Use, Not Access

Licences issued is a purchasing figure. Weekly active users on real work is an adoption figure. The gap between them is the entire subject of this guide, and the reason a 50% expansion in access left the share of people using it daily exactly where it was (Deloitte, 2026).

Count people who used it on genuine work this week. Break it down by team and by role, because an average conceals exactly the unevenness you need to see.

Know What Good Looks Like

It helps to have a benchmark. Among UK organisations that use AI, 53% describe their use as constant and 80% use it at least weekly (DSIT, 2026). That's the shape of a tool that has landed inside the work.

Anything monthly or occasional is a tool sitting beside the work, being reached for when someone remembers. That's not a discipline problem in your people. It's a sign the workflow hasn't changed to include it.

Ask What's Still Awkward

The most useful adoption data isn't a number. Every few weeks, ask users what's still clumsy and what they've stopped using it for. The second question matters most, because abandonment is silent and people rarely report it unprompted.

Fix the two most common answers each time. Visible responsiveness does more for adoption than any communications plan, because it demonstrates that using the thing and complaining about it is welcome rather than disloyal.

Watch for Overconfidence Too

Adoption isn't uniformly good. AI-assisted work scores higher on persuasiveness and coherence whether or not the answer underneath is correct (Dell'Acqua et al., 2026), which means enthusiastic use can quietly increase risk if the review habit has slipped.

So track a quality measure alongside the usage one. Sample outputs occasionally and review the reasoning rather than the writing. Rising use with falling scrutiny is the one pattern worth catching early.

Adoption Measures Worth Tracking

Measure	How often	What it tells you
Weekly active users on real work	Weekly	Whether the tool is inside the work or beside it
Spread across teams and roles	Monthly	Whether adoption is widening or concentrating among the confident
Tasks it's being used for	Monthly	Whether use is drifting towards things it isn't good at
What people have stopped using it for	Every six weeks, by asking	Silent abandonment, which no dashboard will show you
Sampled output quality	Monthly	Whether review habits are holding as confidence grows

The Day Nobody Mentions It

Adoption was never a personality trait of a workforce. Nearly everyone is willing, most are simply waiting to be shown one thing that matters to them, and the small number actively avoiding it are far fewer than the anxiety around this subject suggests.

What's needed isn't persuasion. It's design: an honest introduction, one relevant example, protected time, a real say in how the work changes and a leader who answered the jobs question before being asked.

Get that right and the tool stops being the new thing. It just becomes how good work gets done here, and the day nobody mentions it any more is the day it worked.

Rolling something out and want it to land properly? That's a conversation we love.

Jacobs & Xu

Applied AI. Results You Can Measure.

contact@jacobsandxu.com | +44 (0)203 105 7544

Sources

References are given in the text by author and year. Full details are below, in Harvard format.

Boston Consulting Group (BCG) (2025) *The Widening AI Value Gap: Build for the Future 2025*. Boston: Boston Consulting Group. Survey of 1,250 senior executives and AI decision makers across 68 countries and 25 or more sectors. Published September 2025.

Deloitte AI Institute (2026) *State of AI in the Enterprise: The Untapped Edge*. Deloitte Insights. Survey of 3,235 board, C-suite and director-level respondents across 24 countries, fielded August to September 2025. Published January 2026.

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F. and Lakhani, K. (2026) 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality', *Organization Science*, 37(2), pp. 403 to 423. Preregistered randomised controlled experiment with 758 consultants.

Department for Science, Innovation and Technology (DSIT) (2026) *AI Adoption Research*. London: DSIT. Research conducted by IFF Research and Technopolis Group. Telephone survey of 3,500 UK businesses with five or more employees, fielded 12 February to 2 May 2025, with 100 follow-up interviews. Published January 2026, updated February 2026. Available at: <https://www.gov.uk/government/publications/ai-adoption-research> (Accessed: 12 August 2026).

McKinsey & Company (2025) *The State of AI: How Organizations Are Rewiring to Capture Value*. Survey of 1,491 participants across 101 nations, fielded July 2024. Published March 2025.

Yalcin Williams, G. and Lim, S. (2024) 'Psychology of AI: How AI Impacts the Way People Feel, Think, and Behave', *Current Opinion in Psychology*, 58, article 101835. A review of experimental research published between 2014 and 2024.

Figures drawn from executive surveys are self-reported by the organisations taking part. Where this guide describes what high-performing organisations do differently, it describes a pattern found among them rather than a guaranteed result. Four points of context are worth carrying with the numbers. Every Deloitte respondent already runs at least one AI implementation, so its access and usage figures are an upper bound rather than an average. BCG's sample is weighted towards large enterprises, so its findings read better as principles than as benchmarks. McKinsey's fieldwork was carried out in July 2024 and adoption has risen since. DSIT figures cover UK businesses with five or more employees and measure deliberate deployment of a recognised AI technology, which is a narrower definition than most commercial surveys use.