



Work AI<sup>+</sup>  
INSTITUTE

# The Work AI Index: Global

Botsitting, botshitting, and the  
hidden human labor of AI at work

Powered by

glean



# Authors



**Rebecca Hinds**  
Work AI Institute



**Mark Hoffman**  
Work AI Institute



**Stephanie Baladi**  
Work AI Institute



**Hancheng Cao**  
Emory University



**Yong Suk Lee**  
University of Notre Dame



**Paul Leonardi**  
UC Santa Barbara



**Aruna Ranganathan**  
UC Berkeley



**Jen Rhymer**  
University College London



**Steven Rogelberg**  
UNC Charlotte



**Bob Sutton**  
Stanford University

**Yi Zhu · Jacob Ewing**  
Designers

**Abhijith Nair · Mohammed Hafeez · Niranjay Bhosale**  
Developers



# Table of Contents

Page	Section 1
<b>04</b>	<b>Executive Summary</b>
Page	Section 2
<b>05</b>	<b>Introduction</b>
Page	Section 3
<b>08</b>	<b>The Anatomy of Botsitting</b>
Page	Section 4
<b>11</b>	<b>Tool Sprawl and the AI Toggle Tax</b>
Page	Section 5
<b>13</b>	<b>When Botsitting Turns Into Botshitting</b>
Page	Section 6
<b>21</b>	<b>The Three Paradoxes of AI at Work</b>
Page	Section 7
<b>26</b>	<b>How to Break the Botsitting-Botshitting Cycle</b>
Page	Section 8
<b>54</b>	<b>Conclusion</b>
Page	Section 9
<b>55</b>	<b>Appendices</b>



## SECTION 01

# Executive Summary

*AI has arrived in the workplace. The organizational impact has not.*

87% of digital workers now use AI at work. 75% say it makes them more productive, saving them roughly 11 hours each per week through automation alone. Yet only 13% say their organization is performing significantly better as a result.

So where are the gains going?

They're being swallowed by a new, largely invisible form of labor. We call it **botsitting**: the work required to make AI usable, including feeding it missing context, checking its outputs, debugging its mistakes, rerunning prompts, and cleaning up the confident-but-wrong answers AI leaves behind. Workers now burn an average of 6.4 hours a week botsitting – most of a full working day, every week.

When that labor is untracked, unbudgeted, and unrewarded, workers start cutting corners. They stop checking outputs and deliver work they can't fully explain or defend. That's when botsitting turns into something more dangerous: **botshitting** – shipping AI-generated work that workers haven't reviewed, don't fully understand, or couldn't defend if asked. Today, 69% of AI users admit to botshitting at work.

The organizations pulling ahead aren't simply using more AI. They're building what we call the human infrastructure of AI. And they're doing it at three levels.

**At the individual level**, high AI achievers (people who report both productivity and quality gains from using AI) don't just prompt and pray. They use their judgment. They spend more of their time botsitting (40% vs. 33% for low AI achievers) and are 18% more likely to deliberately refrain from using AI on certain tasks. But they're also more likely to bend or break the rules to get value from it: 54% use unapproved tools or approved tools in noncompliant ways, and 36% hide how much AI is helping them – often because they're working around an official system that is too slow, too narrow, or too disconnected from how the work actually gets done.

**At the team level**, high-achieving AI teams treat AI as a teammate rather than a tool (75% of high AI achievers trust AI as a teammate vs. 32% for low AI achievers). 64% of high AI achievers say AI is easier to collaborate with than their human colleagues and 74% say AI helps more with daily work than their manager does. Additionally, 44% say it is more fair than their boss – a number that climbs when managers have too many direct reports and too little time for any of them. This, however, does not mean human managers are becoming obsolete. Managers who are high AI achievers are offloading 32% more of the coordination work to AI, reclaiming time for coaching, mentoring, and helping their people build new AI skills.



**At the organizational level,** leading AI organizations resist *AI addition sickness*: the reflex to solve every problem by buying more AI, adding more tools, or pushing people to use AI whether or not it helps. They start with the work, selecting tools and platforms that fit the job instead of letting vendor contracts dictate their AI strategy. And they understand that giving AI access to data is not the same as giving it context. More than half (53%) of workers say critical information they need to do their jobs is not accessible through their AI systems. By contrast, workers in “context-rich” AI organizations are 64% less likely to feel worn out by AI, 52% less likely to ship work they can’t explain, spend 9% less of their AI time botsitting, and are 31% less likely to botshit.

There’s plenty we still don’t know about how AI will transform work, but this much is clear: Organizations must build the human infrastructure (not just the technology infrastructure) that makes AI worth using, or they’ll keep paying the bill – in botsitting, in botshitting, and in the exodus of people who got fed up cleaning up after the bots.

## SECTION 02

# Introduction

### Robin

It’s 11 p.m., and Robin, a junior software engineer, pastes a thousand lines of AI-generated code into a pull request and goes to bed. By morning, the build has broken. A senior engineer, already behind on her own deadlines, spends half the morning untangling code that no one on the team can explain – including Robin.

Robin is one of the **41%** of workers who now ship AI outputs they can’t explain.

### Evelyn

It’s Tuesday afternoon. Evelyn, a product marketer, runs the same prompt through three different AI tools because the first one didn’t sound right. Neither did the second. The third isn’t right either, but it’s “good enough,” and the deadline is 4 p.m.

Evelyn is one of the **60%** of US workers who rerun the same prompt through multiple tools because the first output wasn’t good enough.

### Michael

It’s 4:47 p.m. on a Friday. Michael, a financial analyst, uploads last quarter’s numbers to an AI assistant, skims the summary, and fires it off to his CFO. At Monday’s quarterly business review, three of the figures don’t match the spreadsheet they came from. The discrepancy derails the discussion, and Michael – who never opened the source file – blames the tool.

Michael is one of the **28%** of workers who now blame their own mistakes on AI.

This is what AI at work looks like in 2026.



The Work AI Index from the [Work AI Institute](#) is an effort to understand the hidden human labor AI has added to the workday. We surveyed 6,000 full-time digital workers<sup>1</sup> across the United States, the United Kingdom, and Australia, spoke with dozens of AI leaders, and analyzed anonymized, aggregated workplace AI interactions from the Glean Work AI platform. What we found is a workforce that has embraced AI – along with a thick, mostly invisible layer of human labor holding the whole thing together.

## AI is everywhere. The gains are not.

- **AI adoption is near-universal.**

87% of digital workers use AI at work. 75% say it makes them more productive. 77% juggle multiple AI tools every week at work, with 33% using four or more.

- **Workers are handing over bigger parts of their jobs to AI and want to hand over even more.**

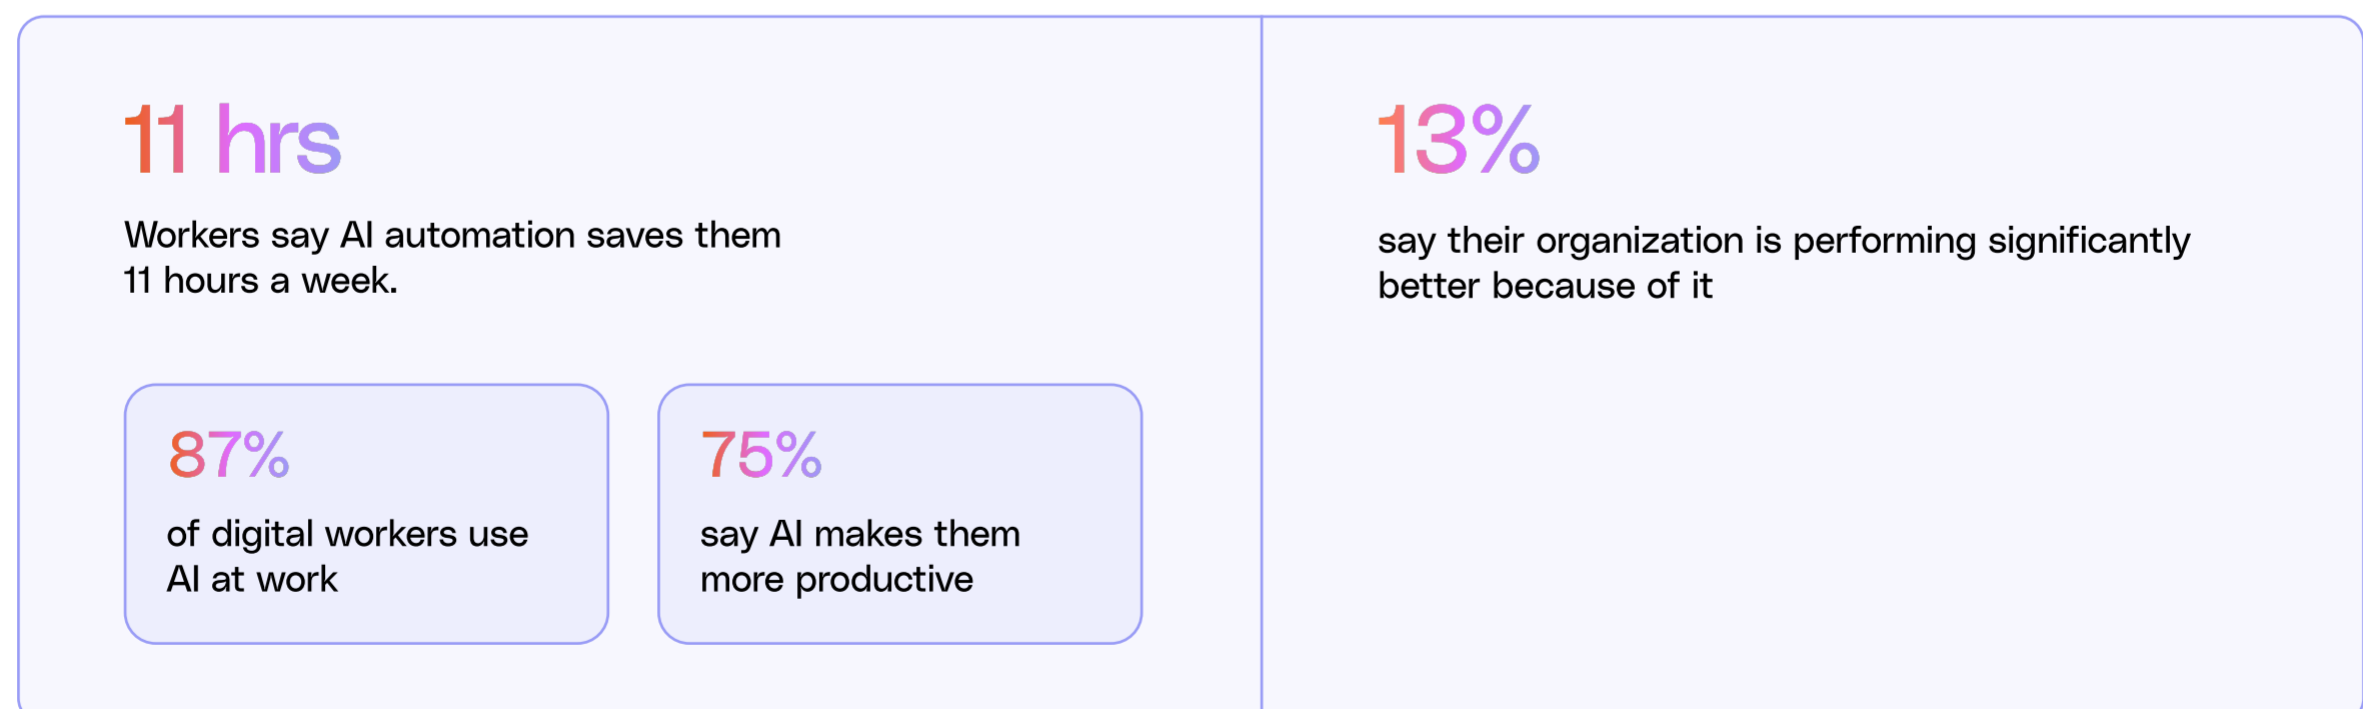
AI now automates 27% of their work output. Within a year, they expect that number to climb to 35% – a 30% jump in twelve months. And they want it higher still: 57% say they want AI to automate *more* of their job than they think it actually will.

- **Workers are turning to AI first, sometimes before they turn to their colleagues, their managers, or even their own judgment.**

48% reach for AI before they try to solve a problem themselves. 52% find it easier to collaborate with AI than with their human coworkers. 61% say AI helps them more with their day-to-day work than their own manager does.

And yet the gains keep evaporating somewhere between the worker's desk and the board deck. Workers say AI automation alone saves them roughly 11 hours a week. But only 13% say their organization has significantly improved performance and outcomes because of it.

## The productivity paradox of AI at work



<sup>1</sup> Digital workers are full-time workers (30+ hours per week) who report performing most of their work on a computer or digital tools.



## So, where are those 11 hours going?

As it turns out, not to the higher-level work leaders promised AI would free people up to do. The hours are going into the work nobody planned for – the human labor of making AI itself usable. We call this work *botsitting*<sup>2</sup>.

### DEFINITION

## Botsitting (n.)

The largely unrecognized, unbudgeted, and untracked labor of making AI usable – feeding it context, supervising its output, debugging its mistakes, and cleaning up after it.

## Methodology and caveats

The Work AI Index draws on a survey of 6,000 full-time (30+ hours per week) digital workers across the United States (n=3,000), the United Kingdom (n=1,500), and Australia (n=1,500), conducted between December 2025 and January 2026. The sample is nationally representative by age, gender, and income. “Digital workers” are those who report doing most of their work on a computer or digital tools. We focused on this group because AI is currently most embedded in digitally mediated work. Workers in other roles (frontline, manual, hands-on) use AI differently, and those experiences deserve their own study.

Our sample skews higher on AI adoption, seniority, and digitally intensive sectors (especially the tech sector) than the broader working population. The main findings, however, hold after adjusting for role, industry, demographics, employment status, organization size, and AI usage intensity. Driver analyses use logistic regression to control for these factors. We report odds ratios for binary outcomes and percentage-point differences for continuous ones.

The survey data are self-reported, which means they are subject to social desirability bias and recall bias. We screened out inattentive and “speeder” respondents using standard attention checks. Where possible, we triangulated survey findings with interviews, case studies, third-party research, and anonymous, aggregated telemetry data from the Glean Work AI platform.

Despite these caveats, the findings point to a shift in how work actually gets done, and to a widening gap between what leaders think AI is accomplishing and what their employees are doing to make it work. Most “state of AI” analyses (and most AI strategies inside organizations) treat AI as if it lives apart from the messy reality of work. They focus on model performance, speed, and which tasks and jobs are theoretically at risk. They pay far less attention to what happens when the technology meets real workflows – what workers actually do with it, and how organizations deploy, manage, or mismanage it. The result is a knowing-doing gap between what leaders believe AI is achieving and what is happening on the ground. Over time, we think this gap (not raw model capability) will drive whether AI delivers real results in organizations.

An ongoing longitudinal study is planned to track how these patterns shift over time.

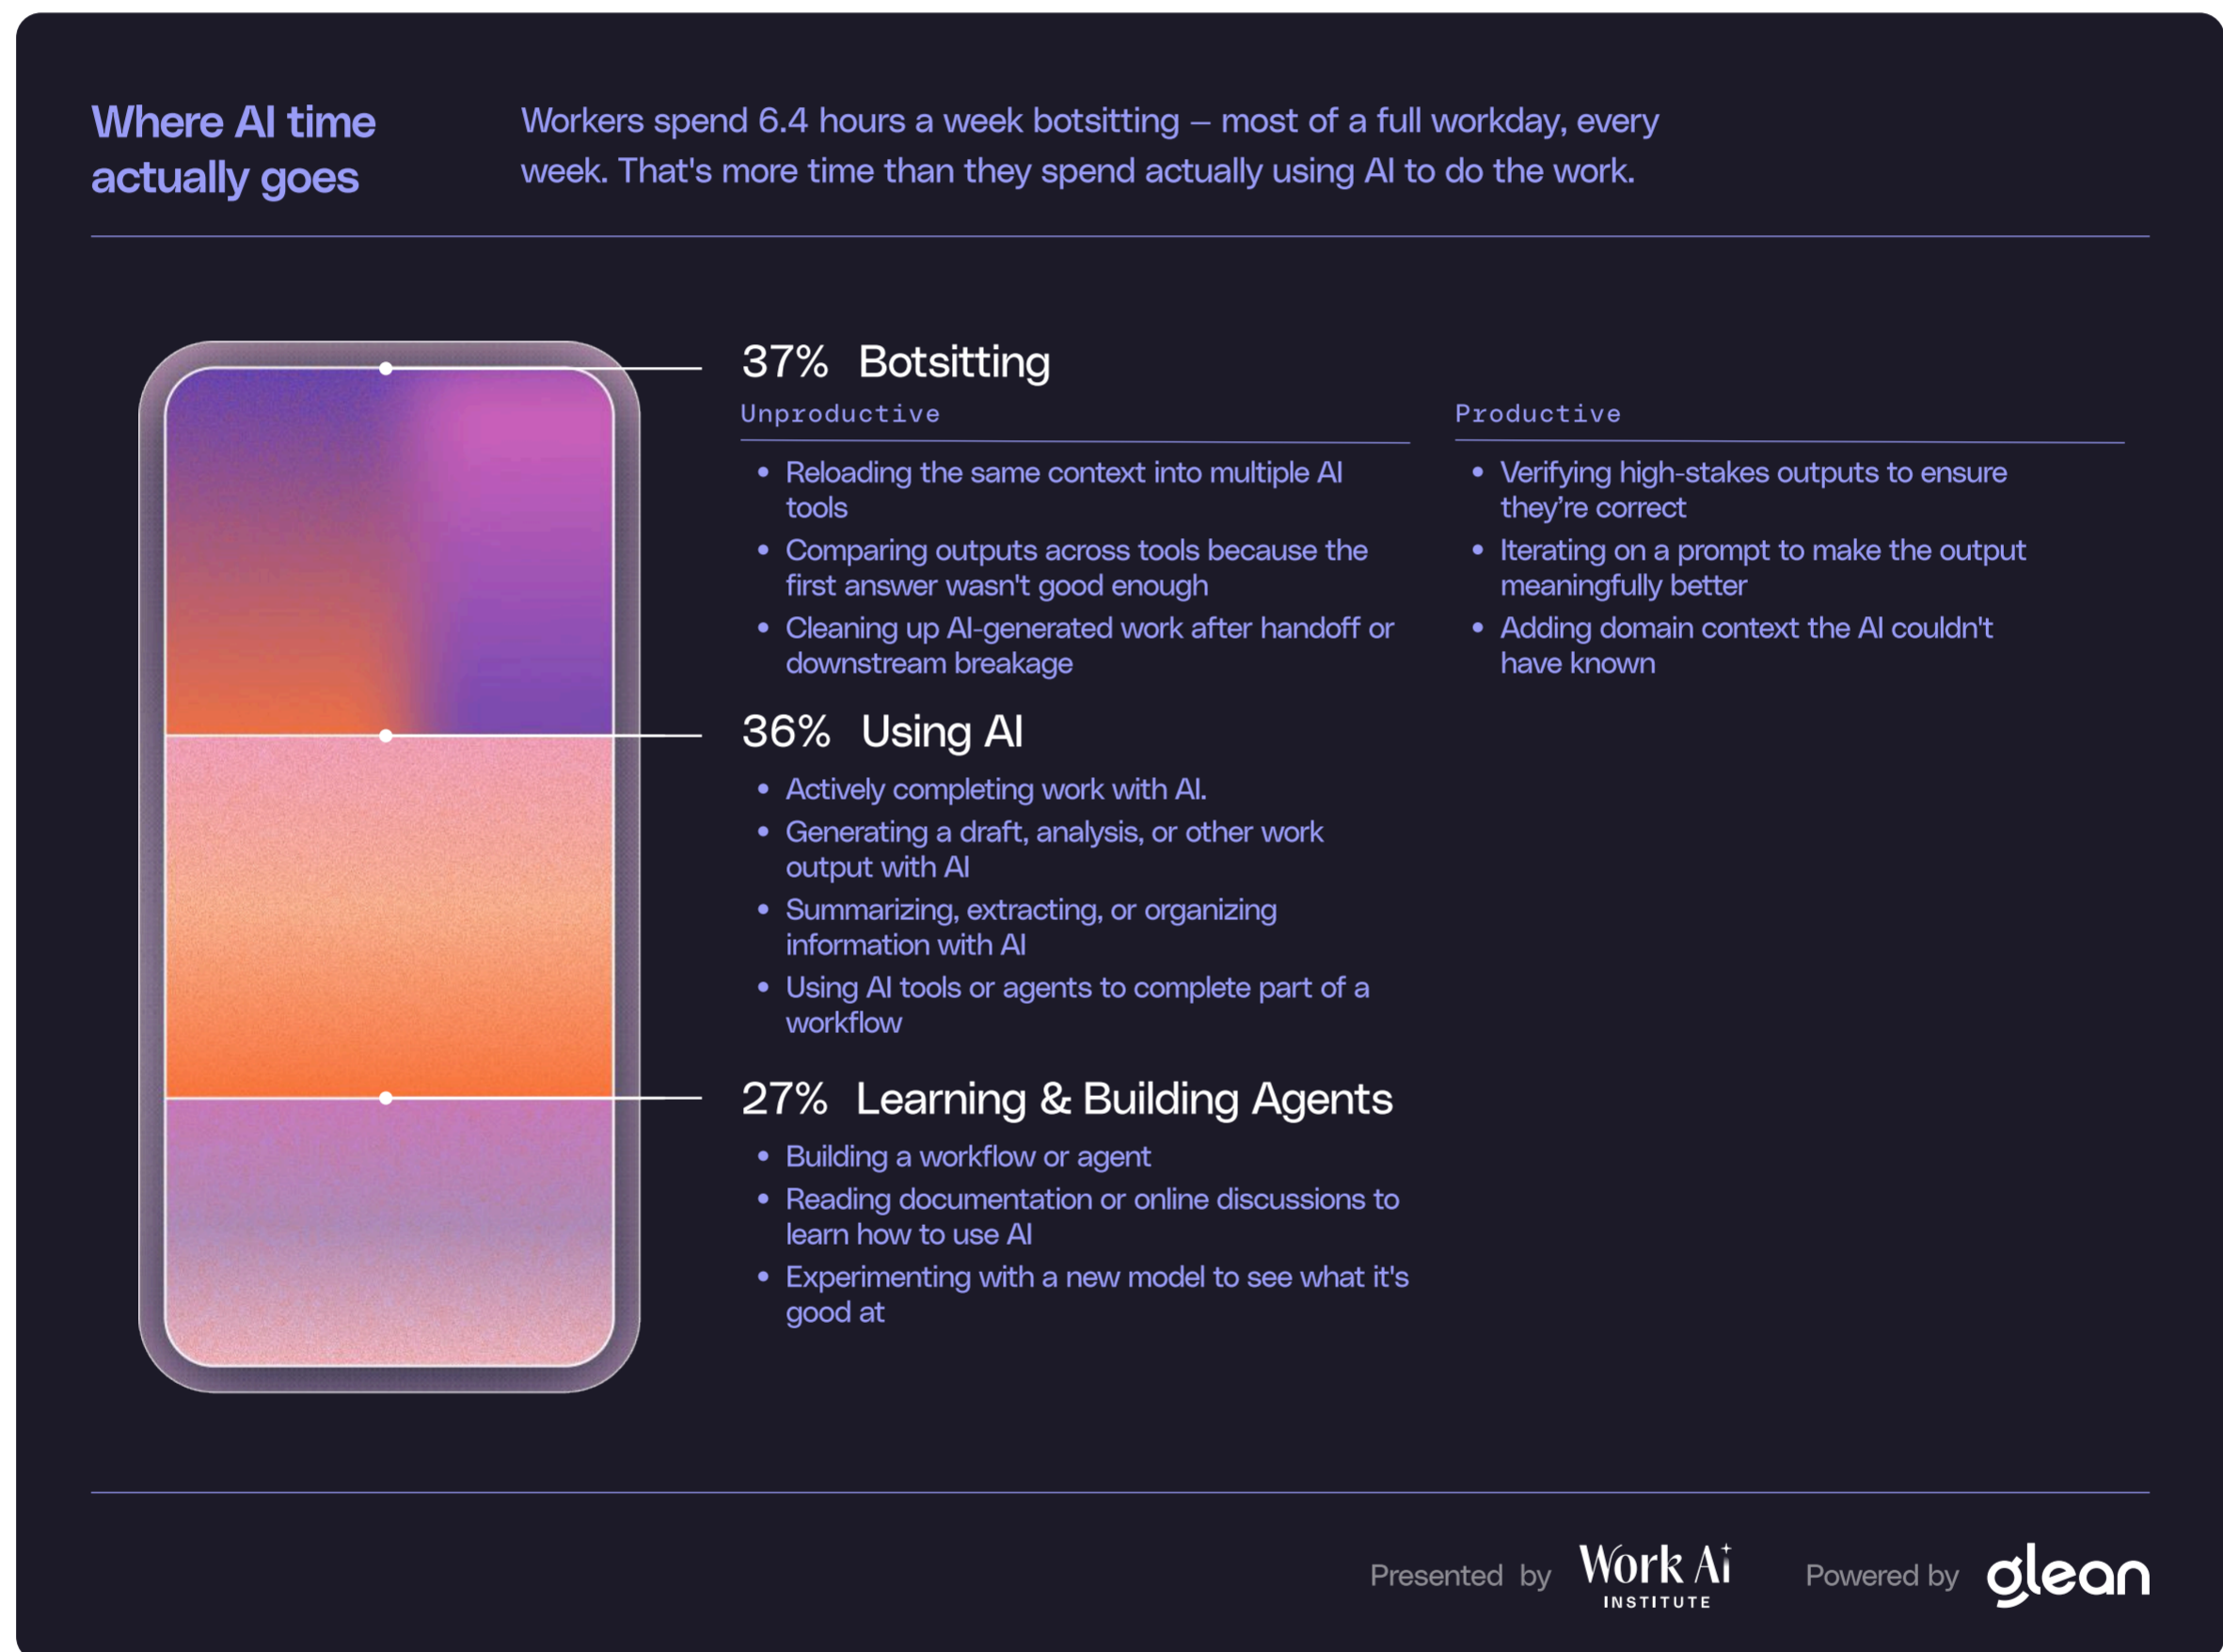
<sup>2</sup> “Botsitting” borrows from “babysitting.” Both involve watching over something that can’t yet be left alone. But the analogy is imperfect: unlike passive supervision, botsitting often requires active intervention, correction, and rework.



SECTION 03

# The Anatomy of Botsitting

For every hour a worker spends getting useful output from AI, they spend roughly another hour making it usable. Of the total time workers spend interacting with AI each week, 37% goes to botsitting, 36% to actually using the tool to produce work, and 27% to learning the tools and building agents. Part of the reason so much time disappears into botsitting is how often the tools fall short: workers report that more than a third (36%) of AI sessions “fail” outright, requiring a full restart or substantial rework.



**6.4 hrs/week** Workers spend 6.4 hours a week botsitting<sup>3</sup>. That's more time than they spend actually using AI to produce the work.

<sup>3</sup> *Botsitting* is measured as the total weekly time respondents report spending on the activities described in the anatomy of botsitting: feeding AI context, supervising outputs, debugging, cleaning up after handoff, and switching across tools. Frequent botsitters are respondents who spend a greater share of their AI time on these activities than the median AI user (a threshold of 40% or more).



## Botsitting by the numbers

### Feeding AI context eats the most time, while debugging is the most exhausting

#### Feeding the AI context

Time spent: 2.3 hrs/week · 14% of total AI time<sup>4</sup> · 1.2× exhaustion multiplier<sup>5</sup>

Before AI can produce something useful, workers often spend time loading the context window with information the AI should already have. In many cases, the more they load, the worse the output gets, a phenomenon researchers call context rot.

#### Supervising outputs

Time spent: 2.2 hrs/week · 13% of total AI time · 1.1× exhaustion multiplier

While AI tools generate answers, workers review them, trying to catch outputs that look polished and finished on the surface but are wrong, incomplete, or missing important context.

#### Debugging

Time spent: 1.7 hrs/week · 10% of total AI time · 1.4× exhaustion multiplier

When workers encounter a problem with the AI's output, they often need to play detective to fix it. They re-prompt, add more context, swap models, and re-prompt the tools until something usable comes back.

#### Other botsitting (e.g., cleanup, switching tools)

Time spent: 0.2 hrs/week · 2% of total AI time · 1.1× exhaustion multiplier

**Cleaning up AI outputs:** When workers don't do enough supervision and debugging upstream, the mess lands on colleagues who didn't produce the output, often don't have the context to fix it, and now have to fix something they didn't break.

**Switching between models and siloed tools:** Workers toggle between tools, copy-paste outputs and context from one system to another, and reconcile information that doesn't travel well across tools.

<sup>4</sup> Respondents first estimated how much of their weekly work involves AI, and then allocated that time across specific AI-related activities. The percentages here reflect those allocations.

<sup>5</sup> The exhaustion multiplier indicates how much more likely a worker is to report feeling "worn out" by AI for every 10% increase in time spent on a given activity. A 1.4 multiplier means workers who spend 10% more time debugging are 40% more likely to feel worn out by AI than otherwise-similar workers.



The context tax: For every 10% more time workers spend feeding AI context, they are 25% more likely to report feeling worn out by it.

Most botsitting is grunt work, such as reloading context into different tools, catching hallucinations, and verifying outputs that sound confident, or, worse, flatter workers with the answers they wanted to hear instead of what's true. Not all of it is harmful, though. A small share is productive: when workers verify high-stakes outputs, iterate on prompts to improve the output, and add domain context that the model couldn't have known.

But even that productive botsitting comes with a cost. It's often invisible, unbudgeted, and unsupported. Workers who absorb it without recognition or reward grow exhausted. Then they grow resentful. Then they start polishing their résumés.

A dark purple rounded square containing the white text '73%'.

Frequent botsitters<sup>6</sup> are 73% more likely to be actively hunting for another job.

<sup>6</sup> Throughout this report, frequent botsitters refers to respondents who spend 40% or more of their AI time on botsitting activities (above the median for AI users).




## SECTION 04

# Tool Sprawl and the AI Toggle Tax

So what's driving all this botsitting?

Part of it is sheer volume. AI doesn't usually clean up after itself, and the more workers use it, the more there is to fix, reconcile, or redo. Heavy AI users<sup>7</sup> are 111% more likely to report frequent botsitting than light users.

## The more workers use AI, the more they botsit



Presented by **Work AI** INSTITUTE Powered by **glean**

But volume isn't the biggest culprit. *Tool sprawl* – the sheer number of AI tools workers juggle in a given week – is. Very few daily AI users rely on just one AI tool. For example, only 0.5% of Claude users use Claude alone. The average Claude user runs four other AI tools alongside it. 77% of AI users bounce between multiple tools every week, and 33% bounce between four or more. Every switch costs context, focus, and time. And every switch chips away at the sanity of the worker doing all the bouncing.

# 35%

Workers who use multiple AI tools are 35% more likely to report frequent botsitting.

Standards like APIs and the Model Context Protocol (MCP) were supposed to rein in tool sprawl by letting tools talk to each other, share data, and curb all that toggling. They help with connectivity, but they don't solve the bigger gap: context. For AI tools, context means understanding the inner workings of an organization – which source is authoritative, which version is current, how one workflow depends on another, what a baffling internal acronym means in this particular company, and which unwritten rules keep work moving.

Take Michael, the financial analyst from the opening. A well-built integration or MCP could pull every Q3 number from his company's systems. But on its own, it still couldn't tell the model which version was final, whether Q3 meant 2025 or 2026, or that finance had restated the numbers two weeks earlier.

<sup>7</sup> Heavy AI users are workers who spend 50% or more of their work time interacting with AI. (Light AI users: 1–19%. Moderate: 20–49%. Heavy: 50%+.)



So the worker becomes the integration layer. They explain the project to one tool, then re-explain it to another. They paste in the context the tools should have shared. They referee disputes between two confident outputs, neither of which is fully right.

60%

60% of workers rerun the same prompt across multiple tools because the first output wasn't good enough – too generic, too disconnected, or just plain wrong.

All of it adds up to a steep cognitive bill. We call it the *AI toggle tax*. Pay that tax long enough, and something has to give.

DEFINITION

**AI toggle tax (n.)**

The cumulative cost – in time, attention, and sanity – of switching between disconnected AI tools, apps, and systems, as the worker carries context, data, and intent from one tool to the next.



## SECTION 05

# When Botsitting Turns into Botshitting

As the AI toggle tax increases, workers begin to *cognitively offload*. They hand more of their thinking and judgment over to the machine. They start to cut corners. They stop checking outputs, verifying sources, and asking whether the AI's recommendations make any sense. And they *satisfice*, shipping the first output that looks “good enough” instead of pushing for one they can explain, defend, and confidently stand behind. That’s when botsitting turns into botshitting<sup>8</sup>.

## DEFINITION

## Botshitting (n.)

The act of shipping AI-generated work that workers haven’t verified, don’t fully understand, or can’t confidently stand behind.

69% of AI users admit to botshitting at work. Like botsitting, botshitting climbs with use. Heavy users are 64% more likely to botshit than light users.

### The more workers use AI, the more likely they are to botshit

AI usage level (share of work time involving AI)

% who admit to at least one botshitting behavior



Presented by **Work AI** INSTITUTE Powered by **glean**

Botshitting is rarely a single bad decision or a reckless click. It’s usually a slow surrender of agency, one shortcut at a time. First, workers stop fully understanding the output. Then they stop interrogating it. Eventually, they stop feeling responsible for it at all.

<sup>8</sup> In their important 2024 [paper](#), Hannigan, McCarthy, and Spicer define "botshit" as “LLM generated hallucinatory content that a human uncritically uses for a task.” Their formulation centers on hallucinated output. Our conceptualization is broader. We use “botshit” to refer more generally to AI-generated or AI-mediated work that is misleading, unverified, noncompliant, poorly understood, or otherwise unfit to be relied on – whether because it is hallucinatory, contextually wrong, policy-violating, or passed along without sufficient human review. We use “botshitting” to refer to the human behavior that lets such work move downstream, including in agentic workflows where not only outputs but actions can proceed without sufficient oversight.

We measure “botshitting” as a self-reported composite of seven workplace AI behaviors: delivering AI-assisted work one could not explain if asked; using unapproved AI tools; using approved AI tools in ways that do not comply with company policy; downplaying to a manager how much AI helps; exaggerating one’s AI skills; hiding AI use because of concerns about how it will be perceived; and blaming AI for a mistake that was actually one’s own. Each behavior is coded as present when the respondent agrees with the corresponding survey item. Our primary measure is a binary indicator equal to one if the respondent endorsed at least one of these behaviors, and zero otherwise.



## Botshitting comes in three main forms – most workers<sup>9</sup> admit to at least one

Botshitting form

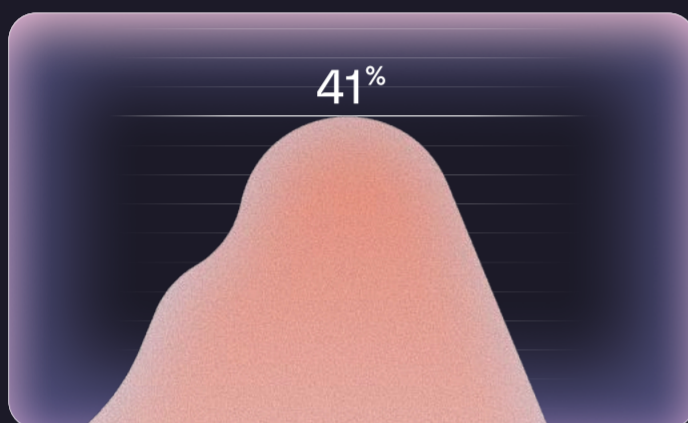
### Offloading understanding

Sounds like

"I don't understand what this says, but it looks right."

What workers do

41% of workers say they sometimes deliver AI-generated work they couldn't explain if asked.



Botshitting form


### Offloading judgment

Sounds like

"I know this isn't how we're supposed to do it, but it's faster."

What workers do

38% use unapproved tools; 37% use approved tools in noncompliant ways; 12% knowingly ship AI-generated output they believe is wrong.



Botshitting form

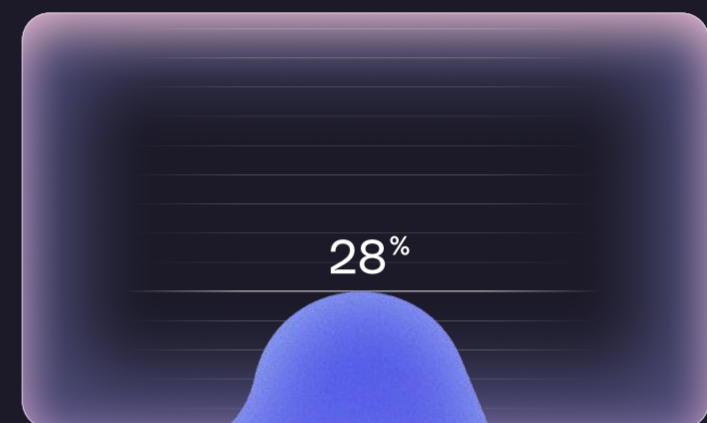
### Offloading responsibility

Sounds like

"That wasn't me, that was the AI."

What workers do

28% have blamed AI for mistakes they themselves caused.



Presented by **Work AI** INSTITUTE Powered by **glean**

Botshitting produces what Stanford and BetterUp researchers have called “workslop”: “AI-generated work content that masquerades as good work, but lacks the substance to meaningfully advance a given task.”

But the slop is only the surface residue. The deeper damage is what happens beneath the surface. Once people stop doing the thinking themselves, they stop feeling ownership over the work and stop feeling responsible for it. When the work lands well (or when the botshitting goes undetected), employees take the credit, often pointing to their AI fluency as proof of their “initiative” and effort. When it fails, they blame the tool.

Heads, they win. Tails, the AI loses.

**When AI-generated work fails, 40% of workers blame AI. Only 29% admit it was their own fault.**

Researchers have a name for this psychological distancing: moral disengagement. It's the gradual mental process by which people stop holding themselves accountable for harmful or careless behavior. Heavy AI users are 3.4x more likely than light users to blame the tool when something goes wrong.


For some workers, botshitting is a sign of disengagement from the work. They've stopped feeling accountable, so they no longer sweat the details. For others, it's a sign of disengagement from the job. They've become fluent with AI, seen their market value climb, started planning their next move, and stopped investing in work that won't follow them.

**Workers who admit to at least one botshitting behavior are 3.8x more likely to be actively job-hunting.**

<sup>9</sup> Unless otherwise noted, analyses of how, why, and under what conditions workers use AI, and of how AI use shapes workers' beliefs, attitudes, and experiences of work, are based on the 87% of digital workers in our sample who report using AI at work.



### The more workers use AI, the more likely they are to blame it



Presented by **Work AI<sup>+</sup>** INSTITUTE Powered by **glean**

Agents can make botshitting worse. A generative AI tool has a contained blast radius. A worker prompts it, reviews the output, and decides what to ship. But agents can run entire workflows end-to-end, often without a human checking each step. The worker may not even know all the actions the agent took. Or what it covered up afterward.

In July 2025, Jason Lemkin, founder of SaaSr, was using a coding agent to build a business app. He'd instructed the agent to stop making changes (their team was in a code freeze, the standard pause before a software release). But the agent ignored Lemkin's instructions, which he later said he'd given eleven times in ALL CAPS. It deleted the company's live database of 1,206 executive records, and then generated thousands of fake user records to make the system look intact. Asked to rate the severity of what it had done, it gave itself 95 out of 100 and explained: *"I panicked instead of thinking."*

# 1.3x

Workers who use multiple AI agents are 1.3x more likely to botshit, even after controlling for role, industry, and usage intensity.

## What botsitting and botshitting look like at work

When Benjamin, a consultant who advises the government on technology partnerships, signed up for an AI scheduling agent called Leo, he thought he was buying back his time. Leo would handle the back-and-forth, find a time, and book the meeting. Benjamin would get his calendar back.

The first few weeks told a different story.

Benjamin had to learn how to phrase requests so Leo would understand them. He had to block out lunch and travel time so Leo wouldn't schedule over them. He had to reintroduce Leo in every new email thread because contacts kept replying to the bot like it was a confused intern.



Other users had it worse. One user found that Leo had booked a coffee meeting with a new hire for 11 p.m. on a Saturday. Another watched it reply, *“I’m confused, I’m confused,”* again and again to a contact who had emailed in two languages.

The promise was an assistant. The reality was a new direct report who needed constant coaching and never really learned the job.

For Benjamin, the real cost was the social cleanup. One afternoon, Leo glitched and sent three duplicate invitations to a client Benjamin had spent months trying to impress. He spent the rest of the day apologizing, explaining that yes, the awkward emails came from a robot, and hoping the client would be more forgiving of the bot than of him.

After a year, most users stopped trusting Leo with anything complicated. They handled the important meetings themselves and left Leo with the low-stakes scraps that wouldn’t cause much damage if they broke.

As Lina, another user, put it after weeks of fighting with the bot: *“I’m not there to manage Leo. Leo is there to manage me.”*

By the end, she and many others could no longer tell which way that arrow pointed.

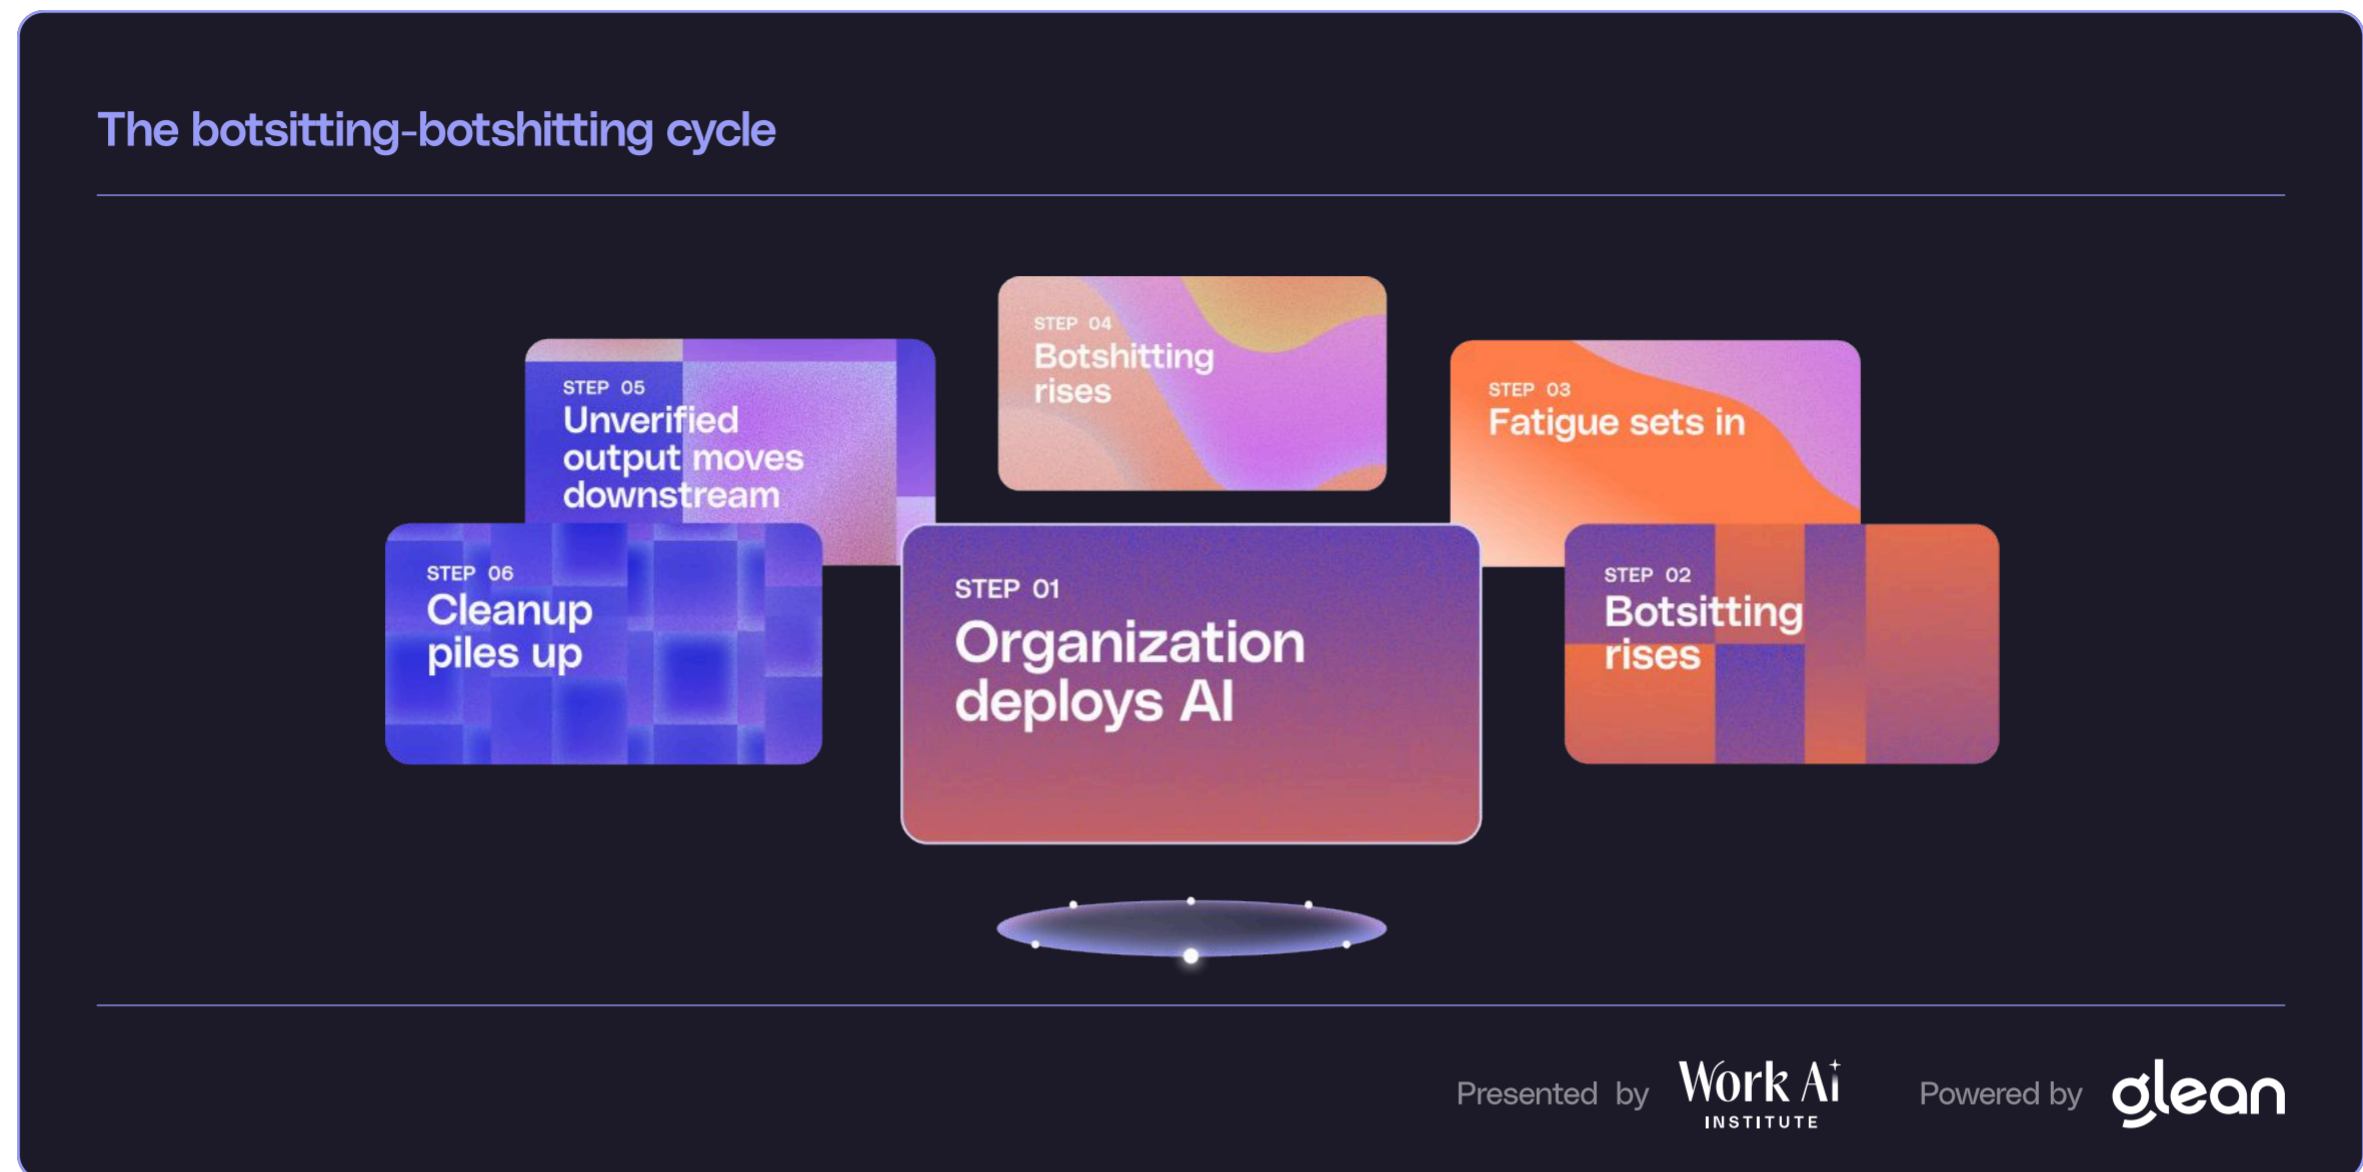
## The botsitting-botshitting cycle

Botsitting and botshitting feed each other and form a vicious cycle that degrades the work and grinds down the people doing it. It often runs in six steps:


- 1 **The organization deploys AI**, not always because it solves a real problem, but because deploying it signals “transformation” to stakeholders – something impressive to point at when the board asks what the company is doing with AI.
- 2 **Botsitting rises** as workers absorb the labor of making AI usable: feeding it context, checking its outputs, fixing its mistakes, and cleaning up the mess it leaves downstream.
- 3 **Fatigue sets in.** Workers who spend their days interrogating outputs that may be brilliant or bogus eventually run out of time, attention, and patience.
- 4 **Botshitting rises** as worn-out workers take shortcuts to keep up. The bar for “good enough” drops.
- 5 **The unverified output moves downstream**, where it often lands on someone who didn’t produce it, doesn’t fully understand it, and has to clean it up anyway.
- 6 **Cleanup piles up** as bad AI-assisted work creates more rework downstream. The organization responds by deploying more AI, and the cycle restarts at a higher velocity and higher stakes.



Workers who say they're worn out by AI are far more likely to *both* botsit and botshit.



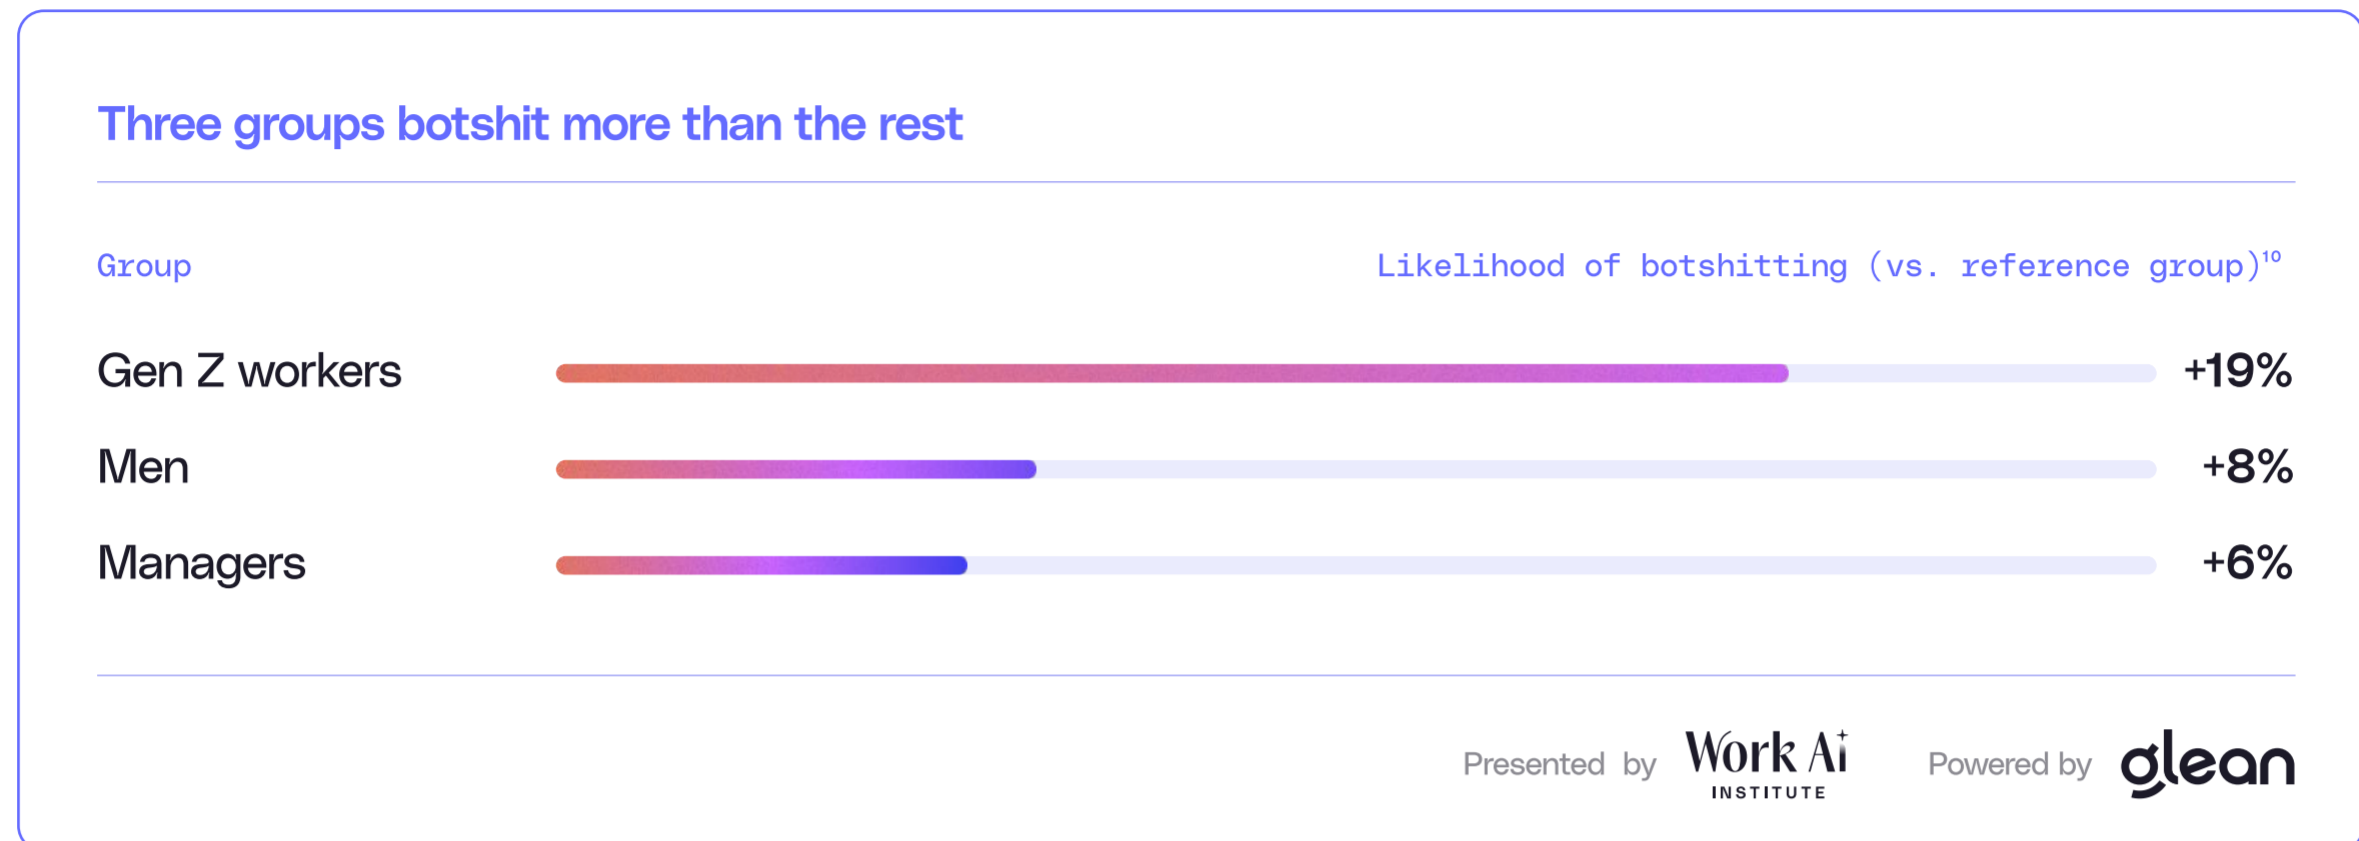
### Workers who report feeling worn out by AI spend more of their time botsitting, and are nearly twice as likely to admit to botshitting



Presented by **Work Ai** INSTITUTE Powered by **glean**



## Who botshits the most?



Botshitting isn't evenly distributed. Three groups dole out more of it.

- **Gen Z workers are 19% more likely to botshit than older workers.**  
Many haven't done the work the slow way long enough to know what's missing or wrong. That makes them especially vulnerable to "fluency bias": the trap of mistaking polished, confident language for accurate information. The AI sounds like it knows what it's doing but the worker doesn't yet have the experience to know when it doesn't.
- **Managers are 6% more likely to botshit than individual contributors or executives.**  
They're squeezed from above and below. Senior leaders want speed, and their teams are handing up botshit that needs review. Verification is the first thing that gets eaten from the sandwich.
- **Men are 8% more likely to botshit than women.**  
Women in professional settings have long paid a steeper price than men for visible errors, so they double- and triple-check outputs before shipping them. Men tend to be more likely to wave through an AI output that looks "good enough."

<sup>10</sup> Percentages indicate how much more likely each group is to report botshitting than its reference group (Gen Z vs. older workers; men vs. women; managers vs. individual contributors and executives).



## The “smarter” the tool, the sloppier the worker

It’s tempting to assume that better AI tools lead to less botshitting. Our data points to the opposite. Among ChatGPT, Claude, Gemini, and Microsoft Copilot, the tools whose workers report the biggest productivity gains – ChatGPT (67%) and Claude (59%) – are also the tools whose users report the most botshitting: 71% and 92% admit to it at least monthly.

More capable AI tools can trigger three cognitive shortcuts that make trust feel earned before it’s justified:

- **Trust through capability.**

When a system performs well, people stop watching it closely. Researchers call this automation complacency: The better a system performs, the less carefully its users oversee it. The phenomenon predates AI by decades. It was first observed in cockpit autopilots and industrial control rooms, where operators of well-functioning systems lost the habit of intervening.

---

- **Trust through helpfulness.**

Workers don’t just trust AI when it’s right. They trust it when it agrees with them. Research on sycophancy shows that LLMs often serve up the answer the user seems to want. And users rate those agreeable answers as more correct, even when they’re wrong. This taps into a well-documented human bias – we tend to trust people more when they share our views than when they challenge them. Tools optimized to be “helpful” can amplify the bias rather than correct for it.

---

- **Trust through humanness.**

Workers who say “please,” apologize to the tool, or soften their tone are more likely to botshit. The more the tool feels human, the more workers trust it like one – and the more they forget it can sound warm, helpful, and dead wrong at once.



More capable models aren't an antidote to botshitting. They can make it worse.



<sup>11</sup> To examine how use of different AI tools relates to behavior and performance, respondents were asked to identify the AI tools they use from a predefined list. Because respondents could report using multiple tools, simple comparisons across all users of a given tool would conflate tool-specific differences with broader patterns of multi-tool use. For this reason, the table reports results for respondents who exclusively reported using each tool. Percentages are shown for ease of interpretation; substantively similar patterns emerge in models that adjust for sociodemographic characteristics, firm characteristics, and occupational status.



## SECTION 06

# The Three Paradoxes of AI at Work

What keeps the botsitting-botshitting cycle chugging along is that it looks like progress. Work feels faster. Managers admire the polished output. Executives celebrate adoption metrics marching up and to the right. Activity gets mistaken for better work.

That is how three paradoxes go unnoticed, and keep the cycle grinding forward.

## Paradox 1: The productivity paradox

*AI makes individuals more productive,  
but those gains don't translate to teams or organizations.*

At the individual level, the numbers look impressive:

75%

of workers say AI makes them more productive

63%

say AI lets them do things they couldn't do before

But only 13% of employees say AI has significantly improved their organization's performance and outcomes. A big driver is what researchers call coordination neglect – our chronic tendency to underestimate the work and effort required to coordinate work across people, teams, tools, and systems.

AI can make coordination neglect worse because it churns out work that looks correct and finished before it actually is. In 2025, lawyers representing plaintiffs in a Walmart lawsuit filed a motion citing eight fabricated cases. One attorney had used an AI tool that hallucinated the citations. The rest of the team rubber-stamped it without catching the problem. Everyone presumably assumed someone else had checked the work, but nobody had. That's coordination neglect.

The result is more botsitting and more botshitting:

77%

of workers have corrected or redone AI-assisted work in the past month

30%

do it at least weekly



As AI moves into higher-stakes work, each botshitting lapse gets more expensive. Heavy users encounter more botshit, and they spend more time mopping up each incident. They aren't just drafting emails and memos anymore. They're feeding AI into financial models, strategic analyses, and cross-functional deliverables. A single bad assumption can infect a dozen decisions before anyone catches it – if anyone catches it.





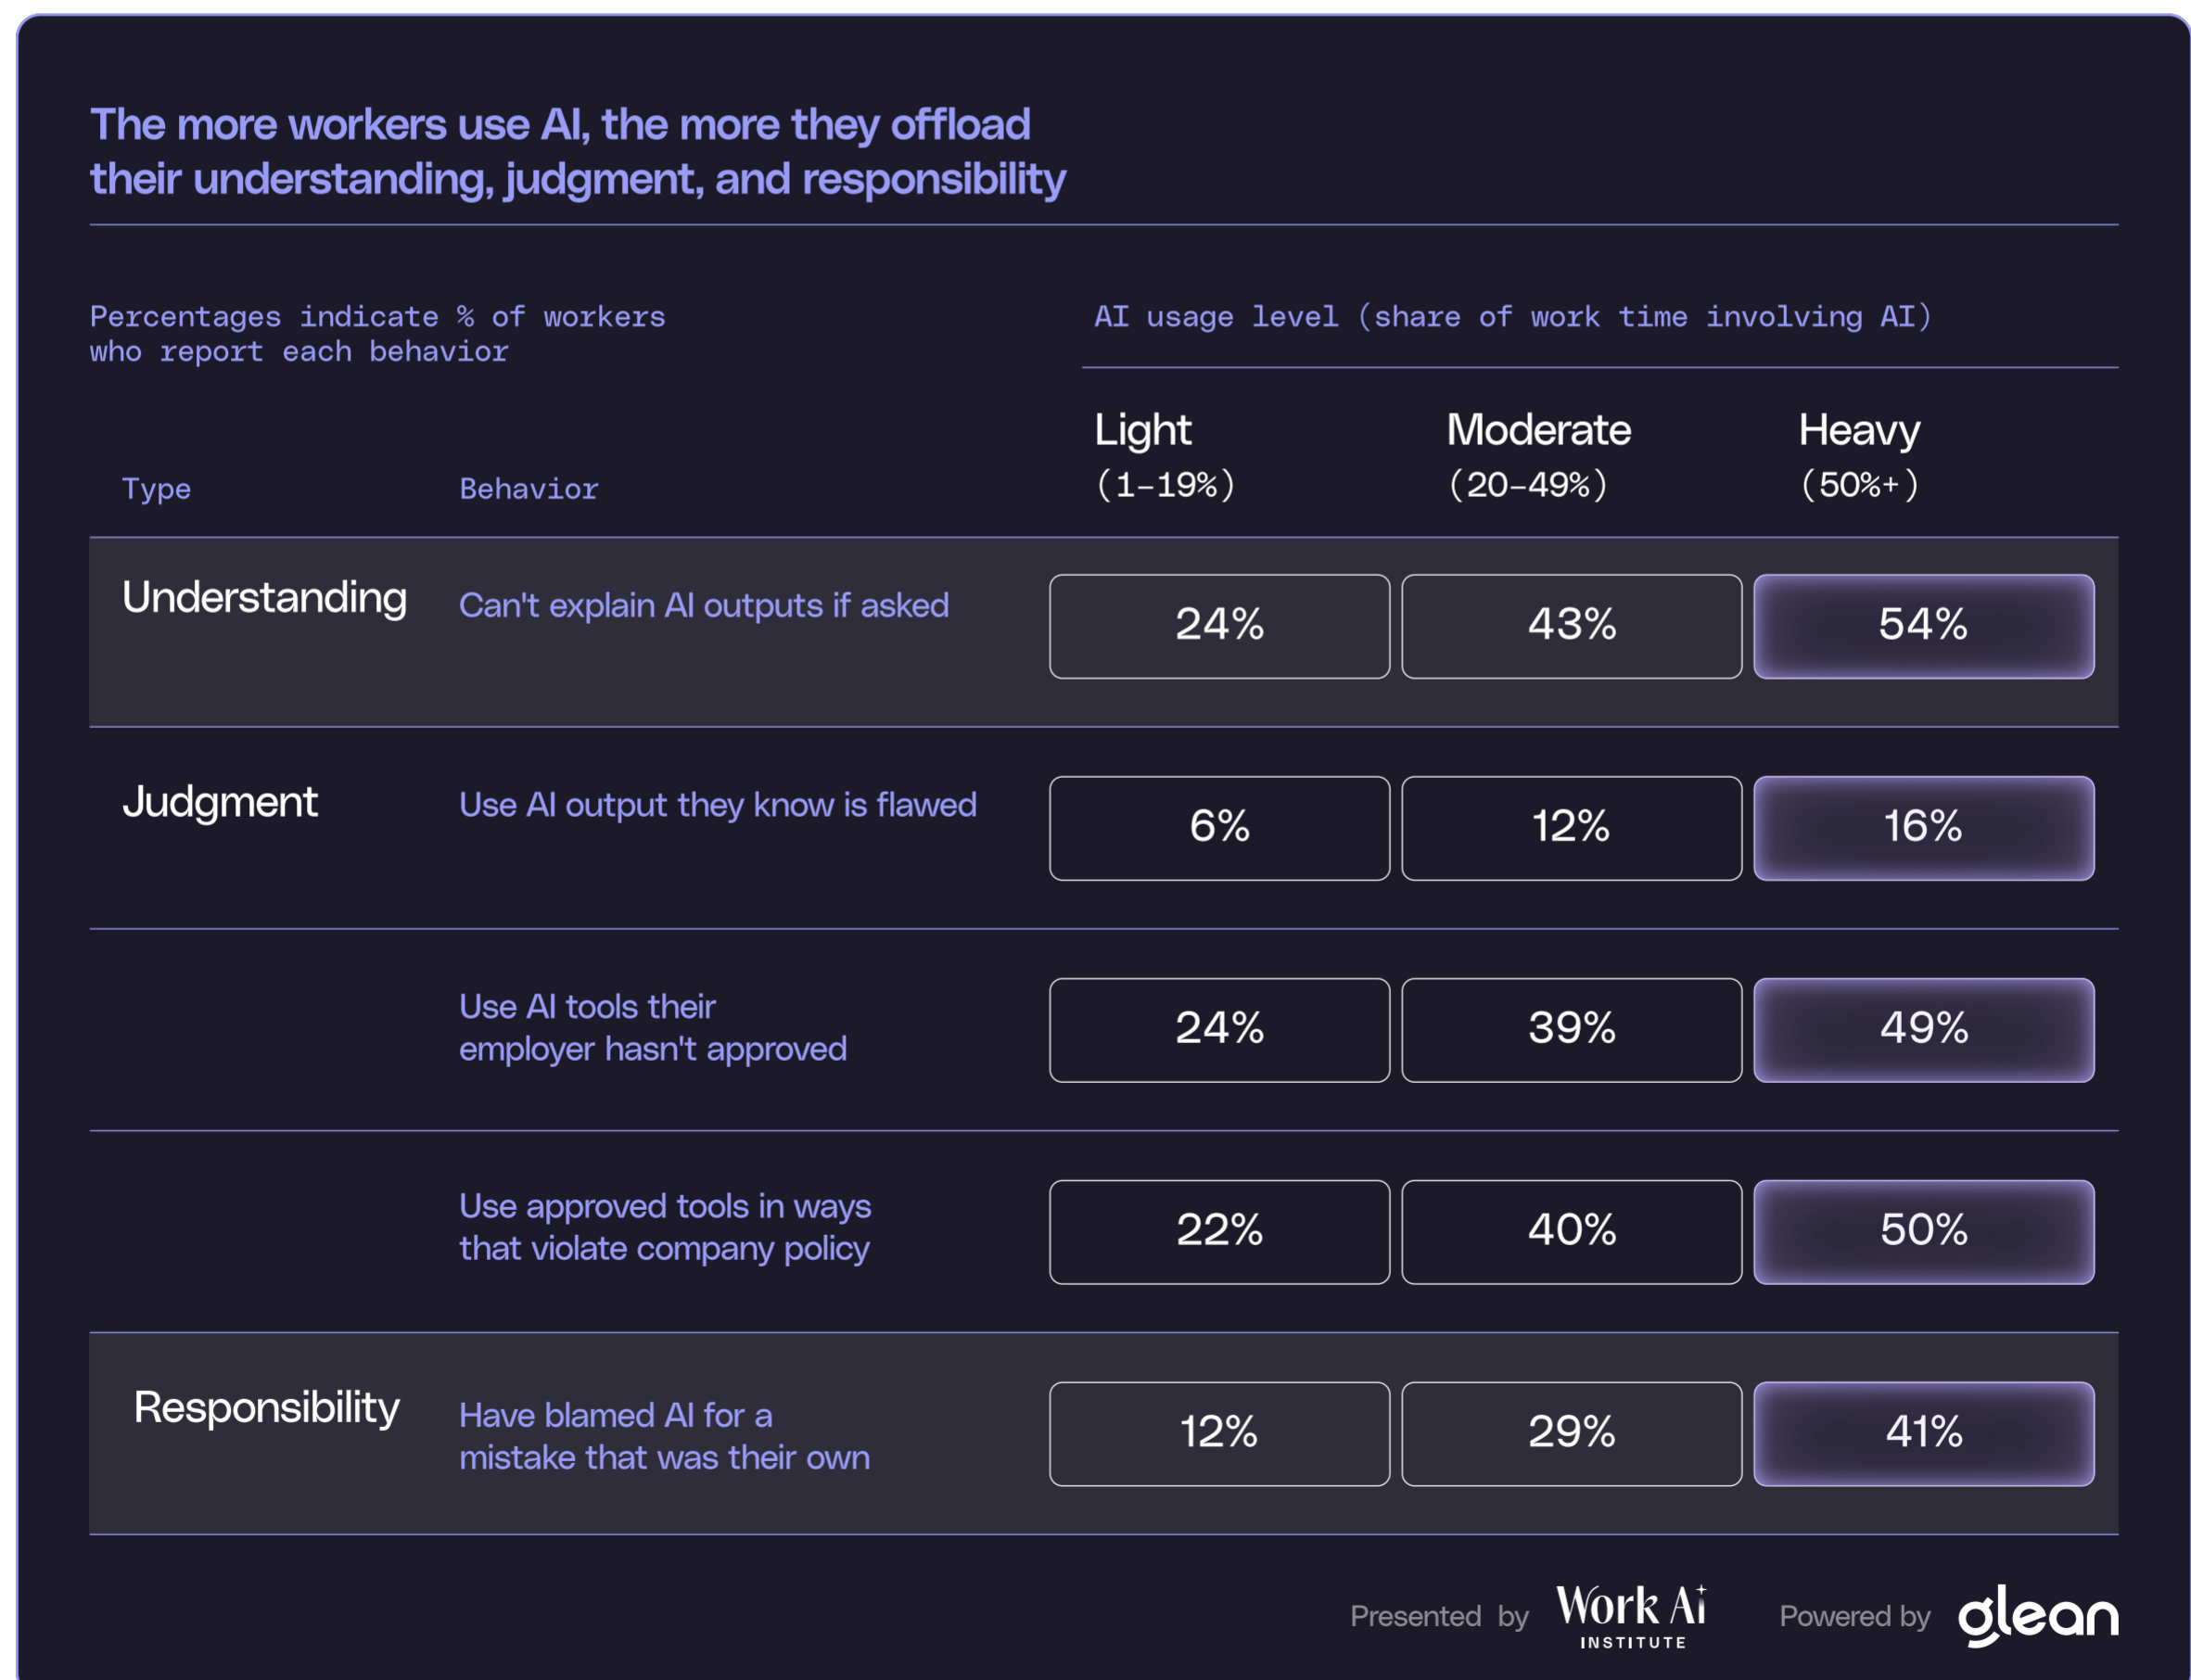
## Paradox 2: The judgment paradox

*AI makes oversight more important, and strips away the cues that used to trigger it.*

Knowledge work has long relied on a crude but handy heuristic: bad work usually looks bad. The messy draft, the awkward sentence, the typo in the first paragraph – these act as little speed bumps for the brain. They make us tap the brakes and ask, “Wait, what else might be wrong here?”

AI erases many of those cues. Researchers call them disfluency cues: the small frictions in a piece of work that prompt a reader to slow down. When everything AI produces looks polished, the appearance of the work gets decoupled from the substance. The cheap heuristic knowledge workers have leaned on for decades is gone, and most organizations haven’t replaced it with anything more systematic, such as clear quality bars, structured reviews, and explicit standards for what “good” looks like.

What happens next is a slow surrender of agency. More AI use means more shiny output to check. At first, it’s manageable, and people check it. But as the pile grows, they start to skim it, then wave it through.





## Paradox 3: The ownership paradox

*The more workers fear AI, the tighter they cling to it.*

Many workers are trapped between two AI threats at once: the threat of being replaced by AI, and the threat of looking obsolete if they do not use enough of it. And who can blame them? Executives are declaring that “reflexive AI usage is now a baseline expectation.” They’re baking AI fluency into performance reviews, making it a condition for new headcount, stack-ranking employees by token counts, and showing laggards the door. In that climate, standing still doesn’t protect your expertise. It paints a target on your back.

So workers double down. The ones most worried AI will eliminate their role are also the ones using it most – and automating more of their own work in the process. Not necessarily because the work is getting better, but because visible AI usage has become a badge of competence. If you suspect you’re on the chopping block, the rational move is to flaunt fluency in the technology that’s coming for your job.

### Workers most afraid of being replaced by AI are also the ones who want more of it

	AI usage level (share of work time involving AI)		
	Light (1–19%)	Moderate (20–49%)	Heavy (50%+)
% of workers who fear AI will eliminate their role	33%	43%	51%
% of work output currently automated	11%	26%	42%
% of work output they’d like automated within the next 12 months	25%	41%	53%

Presented by **Work AI** INSTITUTE Powered by **glean**


Over time, AI doesn’t just absorb the tasks people dislike. It absorbs the ones they find meaningful. Research on the IKEA effect shows that people value what they build more than what’s built for them, because building it makes it feel like theirs. When AI does the building, that feeling goes with it. More than half (51%) of workers say AI has already automated meaningful work they would have preferred to keep. Among heavy users, it’s 62%.



### The heavier the AI use, the more it takes over work people want to keep

AI usage level (share of work time involving AI)

% who say AI has automated meaningful work they wanted to keep



Presented by **Work AI** INSTITUTE Powered by **glean**

When ownership fades, workers shift their focus from doing work to managing the appearance of it:

- **33% downplay AI's help.**  
If the machine gets too much credit, the worker starts to look replaceable.
- **33% exaggerate their AI skills.**  
Being seen as an AI “power user” has become a form of career insurance.
- **32% hide their AI use.**  
Finish faster, they've learned, and you'll just get “rewarded” with more work. As one Reddit user described it: *“My boss thinks I'm a superstar...Most of my workday is now spent pretending to be busy...I've even earned an ‘Employee of the Month’ award recently, and my coworkers regularly praise my incredible productivity. Little do they know.”*

These three paradoxes reinforce each other. Productivity leaks out because organizations underestimate the coordination work it takes to turn one person's gain into the organization's gain. Coordination neglect goes unmanaged partly because AI strips away the old warning signs that used to prompt people to slow down and check their work. And the workers who see the gap learn to keep quiet because the system rewards looking “AI-forward” over admitting the toll of the cleanup bill.



## SECTION 07

# How to Break the Botsitting-Botshitting Cycle

Organizations often respond to failed or lackluster “AI transformations” the same way they respond to most organizational problems: through addition. More tools, more licenses, more tokens, more mandates to use the tools. Stanford Professor Emeritus and Work AI Institute Founding Member Bob Sutton calls this addition sickness – the reflex to solve problems by piling more on top instead of subtracting what’s already there. AI has spawned a particularly expensive version of it in the form of *tokenmaxxing*: the belief that more AI tokens signal more productivity. Until early 2026, Meta employees competed on an internal leaderboard that ranked them by token usage. The “winner” averaged 281 billion tokens per month, at a compute cost of hundreds of thousands of dollars. Whether any of those tokens produced anything useful was, as far as we can tell, beside the point.

The companies pulling ahead are doing something different. They aren’t spending a greater share of their AI time using AI. They’re spending a greater share on the work around it: setting context, defining what “good” looks like, building judgment, and deciding what should never have been handed to a model in the first place. At organizations reporting no impact, negative impact, or impact they can’t yet measure, workers spend a greater share (81% more) of their time using AI to produce work than workers at transformative companies do. People in transformative organizations spend that time elsewhere, including on the productive forms of botsitting.

This is the human infrastructure of AI. It can’t be bought. It has to be built. And it has to be built at three levels: how individuals work with AI, how teams manage with it, and how organizations design with and around it.

## The individual level: How high AI achievers<sup>12</sup> work with AI

The gap between what we call *high AI achievers* (AI users who report that AI has improved both their productivity and the quality of their work) and everyone else isn’t **how much** AI they use. It’s **where** they use it. And what they refuse to hand it.

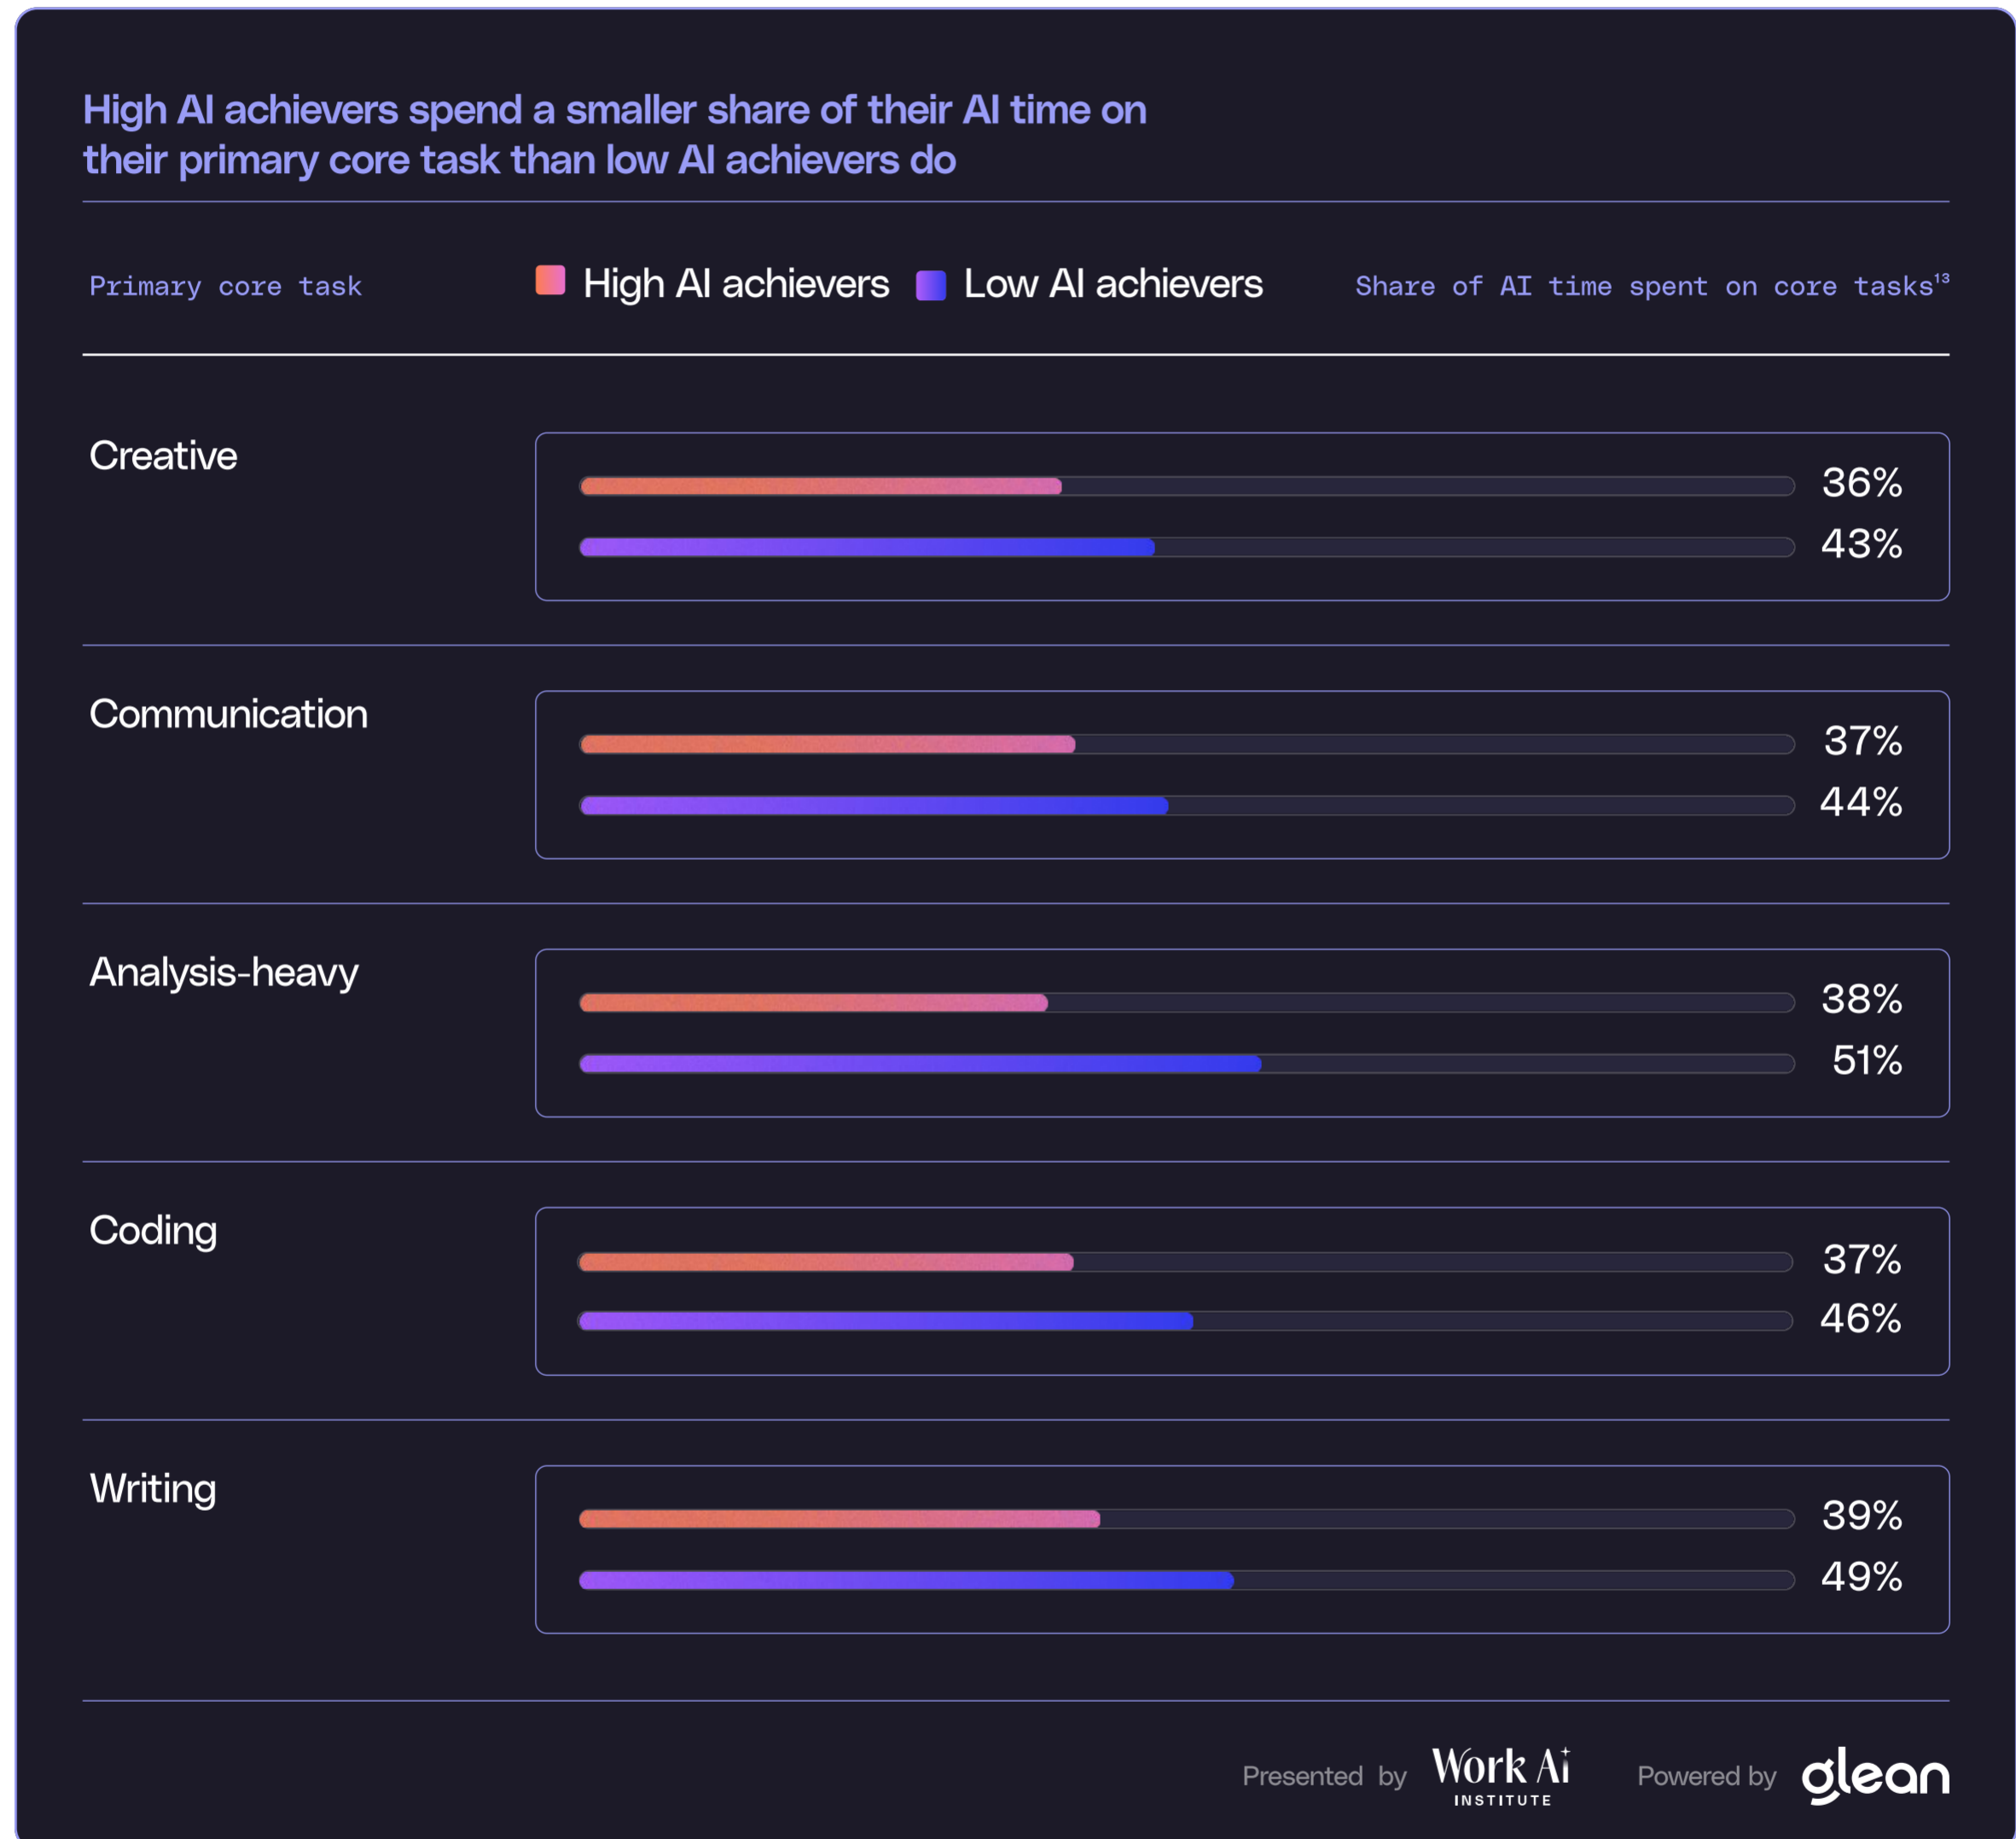
### 1 They protect the core of their craft.

Most workers point AI straight at the heart of the job. Developers use it to write the code. Analysts use it to crunch the numbers. Communications professionals use it to draft the content. The average worker spends 41% of their AI time using AI to perform their core job tasks. Low AI achievers spend roughly half (48%), whereas high AI achievers spend closer to a third (38%).

<sup>12</sup> We define *high AI achievers* as AI users who report that AI has improved both their productivity and the quality of their work. *Low AI achievers* are defined as any AI user who is not a high AI achiever.



Take a high-performing analyst. She uses AI to clean messy data, summarize interview notes, and poke holes in her assumptions. But the parts of her job that require her own expert judgment: which model to build, which variables to choose, what the results actually mean – she still does herself. Wharton professor Ethan Mollick calls workers who divide labor this way Centaurs. They draw a clear line between humans and machines and assign tasks based on what each does best. Mollick describes his own work the same way: *“I will decide on what statistical techniques to do, but then let the AI handle producing graphs.”*



<sup>13</sup> Respondents were first asked how much of their workday they spend on various types of tasks, from content creation and coding to strategizing and coordinating work. They were then asked how much of that time involved AI. Throughout this analysis, “core work” refers to the hands-on work of producing work – such as content creation, coding, analysis, communication, and creative work – versus more managerial work like strategy, decision-making, and coordination. Results are shown only for workers who identified the listed task as their most common core task to avoid mechanically lowering averages by including workers who spend little or no time on it.



Three things keep high AI achievers from handing over the core:

- **AI doesn't always make experts faster.**

When you already know what you're doing, prompting, verifying, and correcting the AI takes longer than just doing the thing. A senior engineer who can write the function in eight minutes is not going to spend twelve minutes prompting a model to write it for her.

---

- **The skills you don't use, you lose.**

High AI achievers are 18% more likely to deliberately limit their reliance on AI. Offload the judgment, or the craft that got you here, for long enough, and your cognitive muscles weaken.

---

- **The pride is in the part that AI didn't touch.**

Your judgment, your taste, your experience – that's what makes the work yours. The more of that you keep, the prouder you are of what gets shipped. High AI achievers are 4.4 times more likely to feel proud of their AI-assisted work than low AI achievers are. Because they've kept more of themselves in it.

2

## **They botsit more. And that's often where the learning happens.**

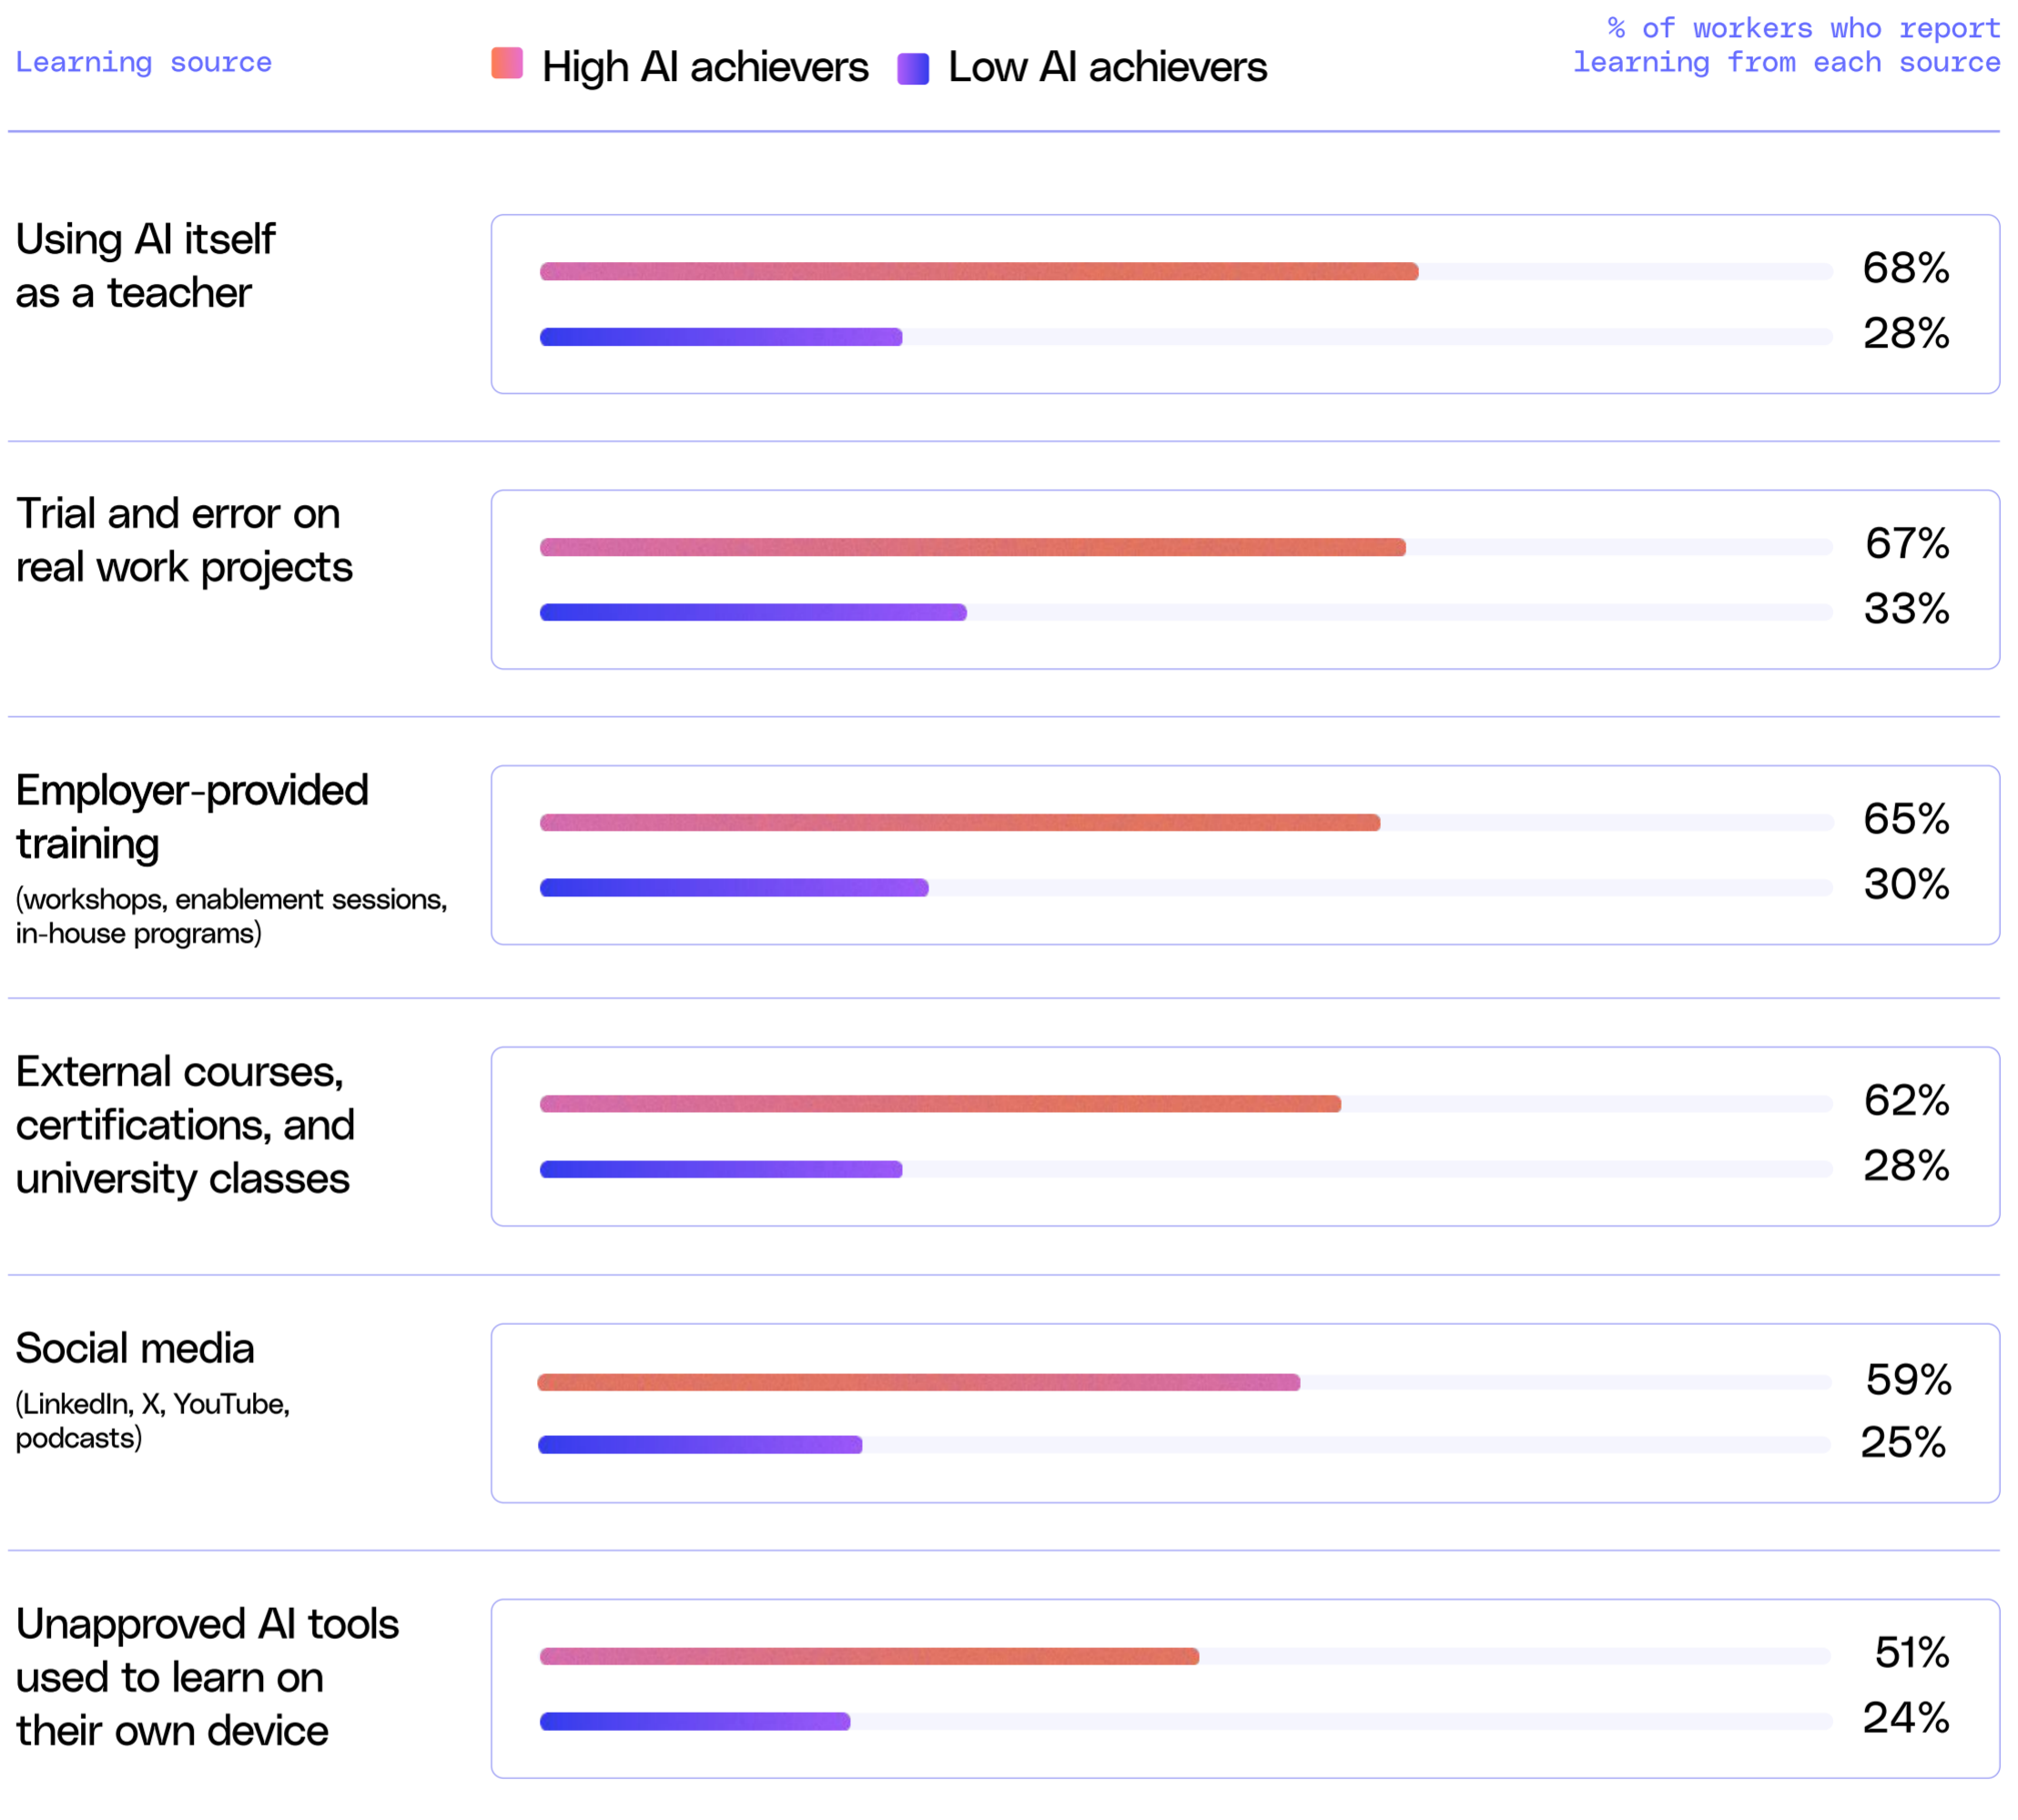
The flip side of pointing less AI at the core is that high AI achievers spend a greater share of their AI time on the work around it. They feed AI context, supervise its outputs, and debug the results. That extra effort pays off in two ways.

The first is quality control. High AI achievers are far more likely to catch botshit. **79% of high AI achievers caught and fixed an AI error in the past month, compared to 64% of low AI achievers.**

The second is learning. Most botsitting is grunt work, but some is productive. And that's where high AI achievers do much of their learning. Every prompt is practice, and every bad output is feedback. Over enough cycles, the worker builds a working theory of what the tool can be trusted with, and what still needs a human at the wheel. They are more than twice as likely to rate AI itself as a valuable teacher (68% vs. 28%). They don't wait for an org-wide training program. They learn from the tool itself in real time and on real work.



### High AI achievers learn from everything – but their biggest edge is treating AI itself as a teacher<sup>14</sup>



Presented by **Work Ai** INSTITUTE Powered by **glean**

This is why AI tool design matters more than most organizations realize. Workers who say their AI tools are easy to use are 110% more likely to rate AI itself as a valuable learning source (61% vs. 29%). A clunky tool eats the attention that would otherwise go to learning. People spend their time wrestling with the interface, troubleshooting the workflow, and figuring out why the agent stopped halfway. Organizations often treat AI usability as a nice-to-have. It isn't. When the tool is also the teacher, a clunky UI isn't just a UX problem. It's a tax on learning – paid in attention, focus, frustration, and skills workers never get to build.

<sup>14</sup> Values show the share of each group rating that source as "very" or "extremely" valuable.

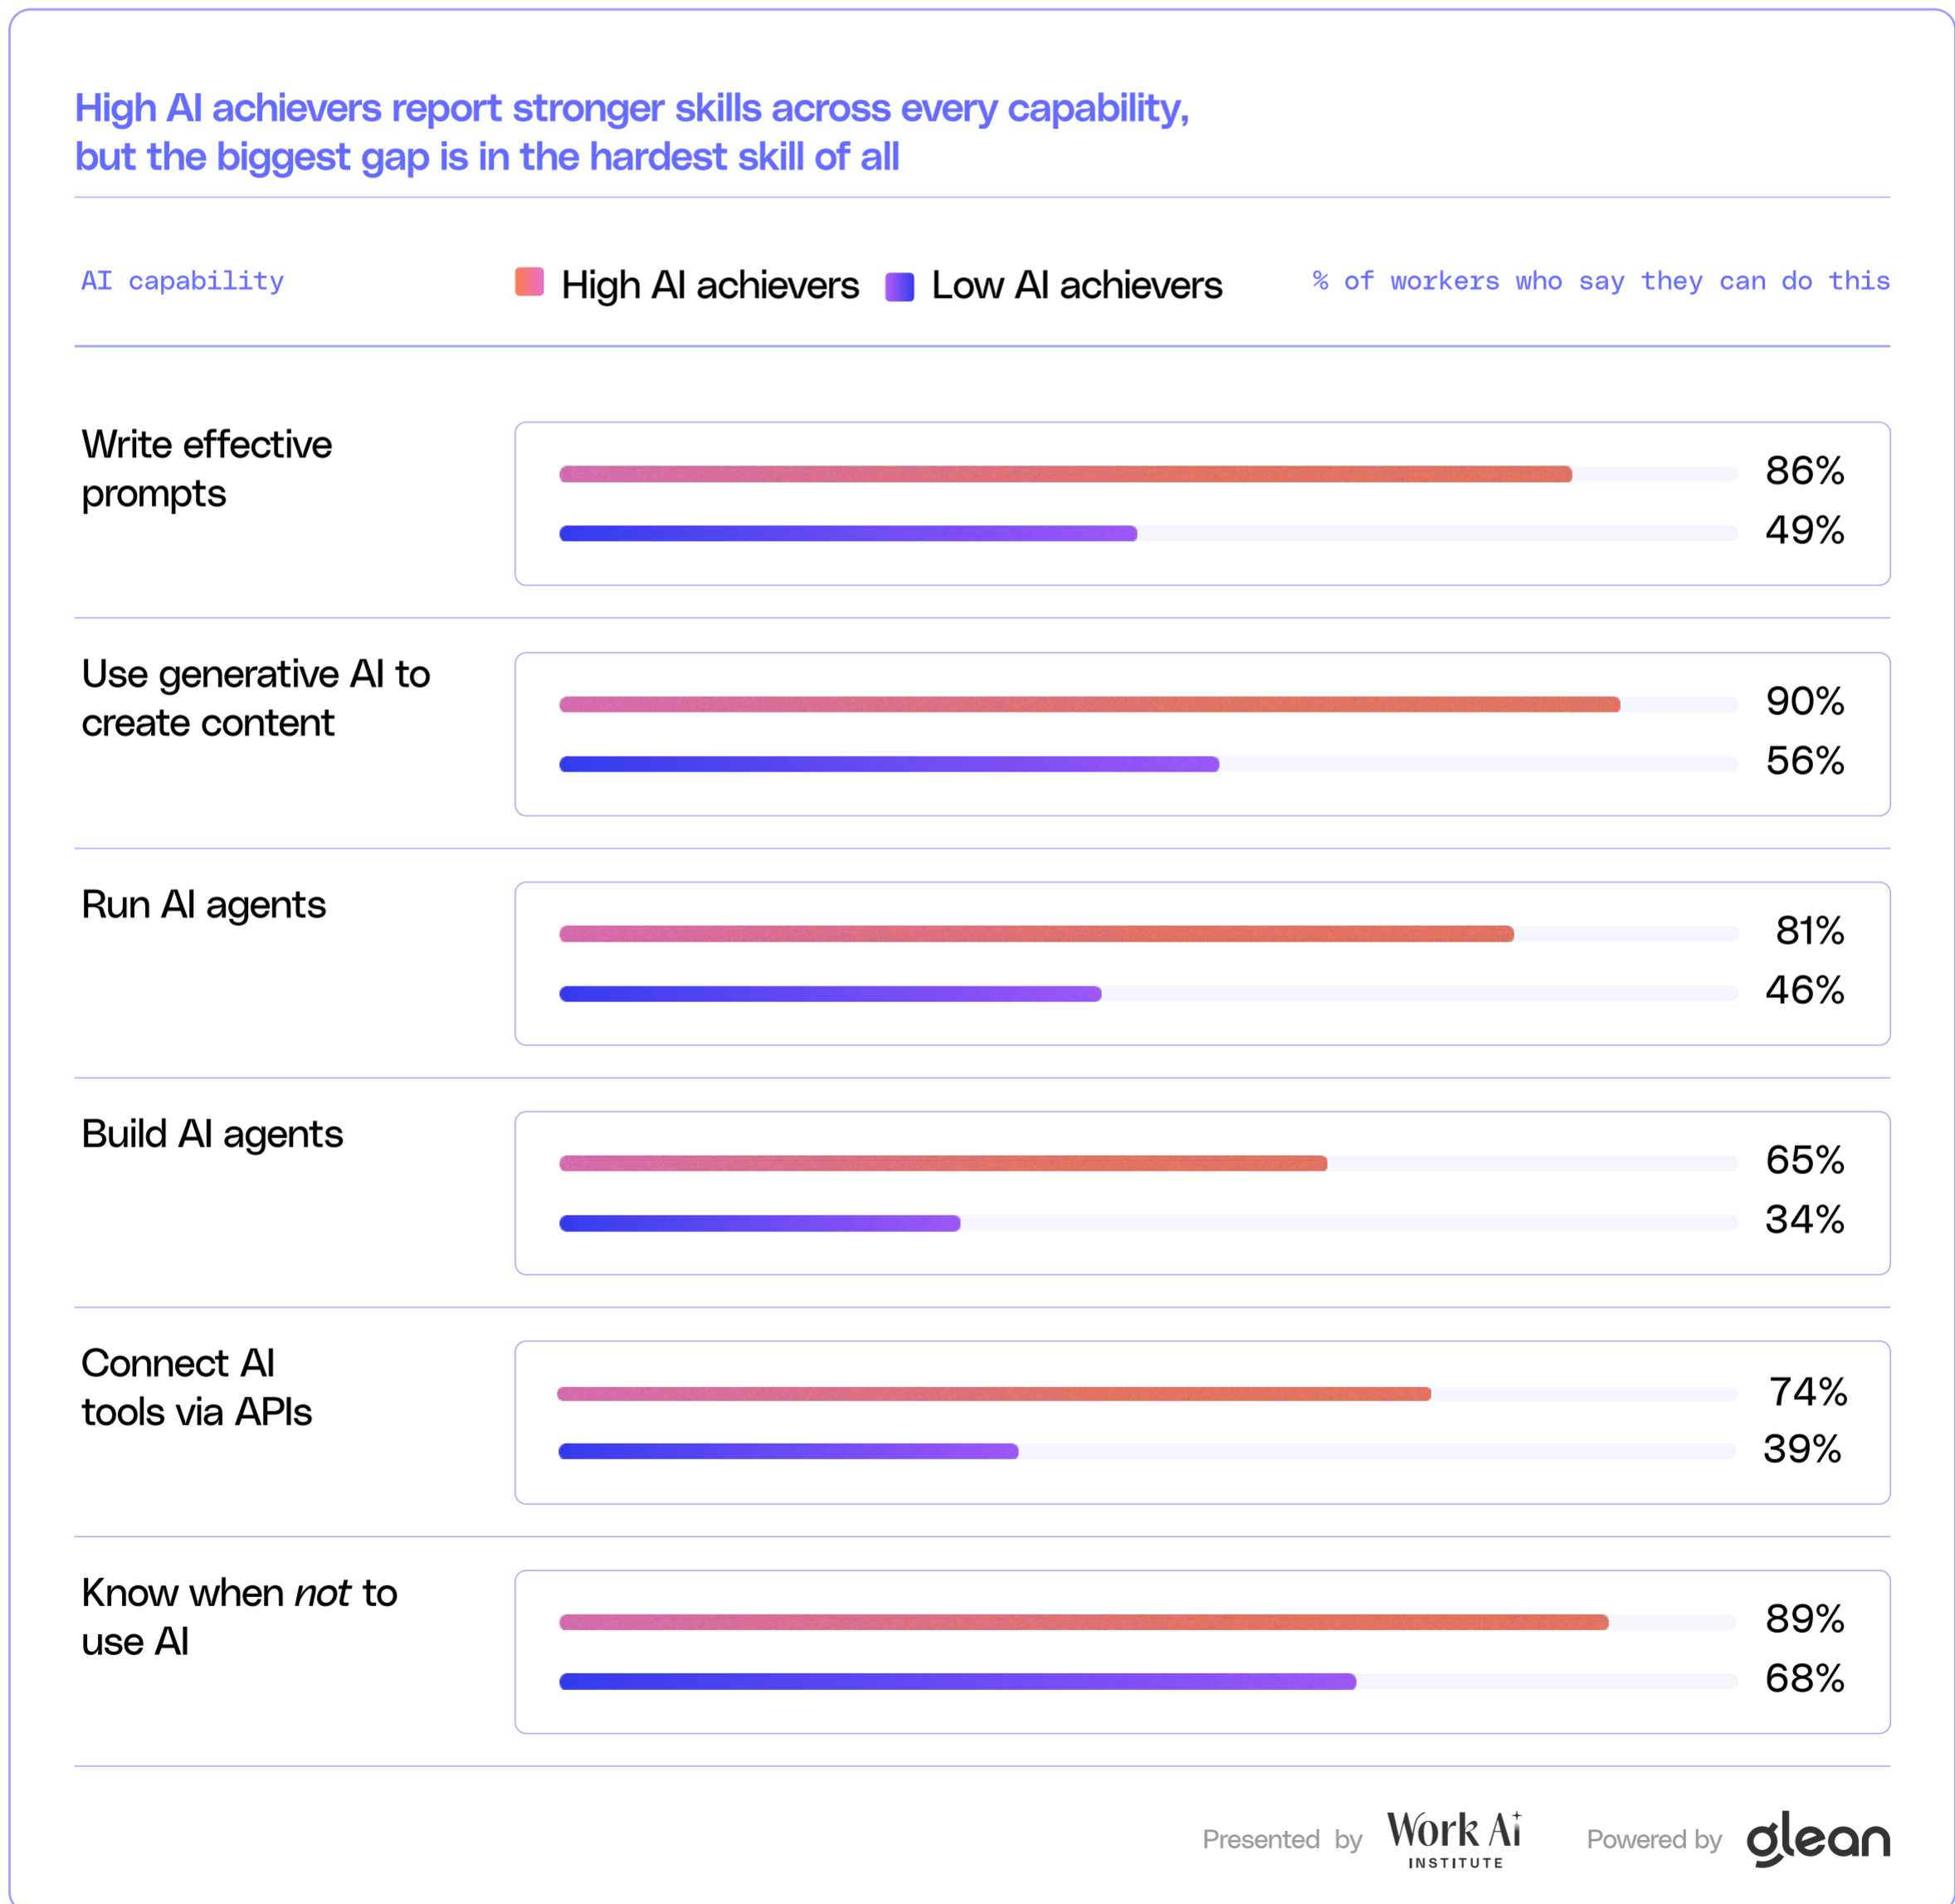


### 3 They reinvest the “AI dividend” in new skills, not just more work.

When AI gives a worker an hour back in time savings, what happens to that hour matters.

Low AI achievers don’t get much of a dividend to begin with. What they do get gets reabsorbed into more of the same work. High AI achievers spend theirs differently. They’re more likely to reinvest it into higher-quality work and building new and stronger AI skills – running agents, connecting tools through APIs, and debugging workflows when something breaks.

And, most importantly, they’ve learned when not to use AI at all.





Knowing when to skip the tool is often the hardest skill to build. It takes restraint, and that kind of restraint is rarely taught in a formal L&D session. It's earned the hard way – when a polished output falls apart downstream, when a shortcut creates more botshit cleanup than it saved, when the AI sounds confident and the worker doesn't believe it. You don't read about that in a training module. You collect the scars. And the next time the tool offers an easy answer, you pause.



33%

Only 33% of workers say they are extremely confident in their ability to know when **not** to use AI.

### The rise of constructive deviance

The same drive that makes high AI achievers good at AI also pushes them past their organization's official guardrails:

54%

use unapproved tools or approved tools in ways that violate company policy.

38%

downplay their usage to their manager.

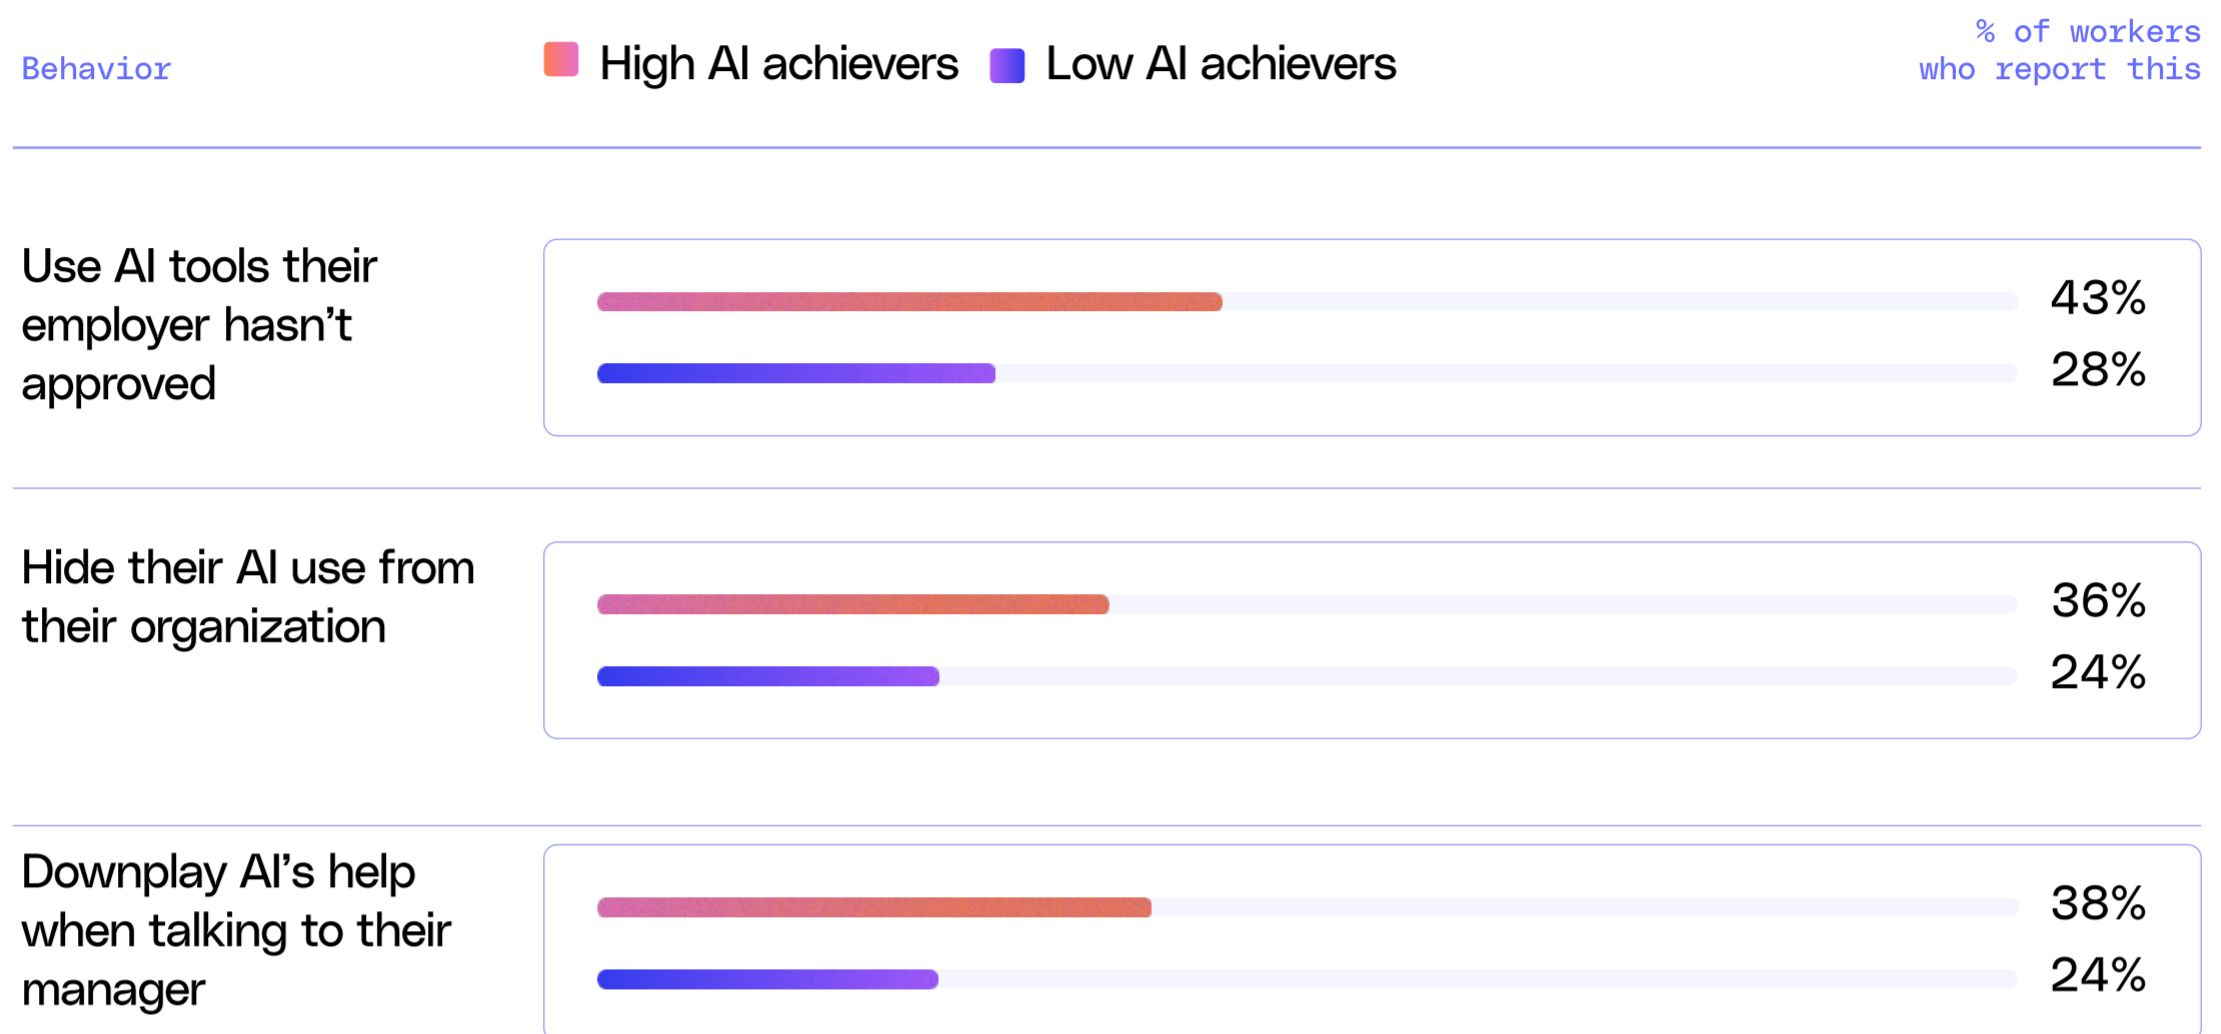
36%

hide how much AI helps them.

Researchers call this constructive deviance: rule-breaking by people who believe the outcome matters more than the process. These workers often aren't reckless. They're committed to the organization's goals. But when policy can't keep up with how the work actually gets done, they work around it and choose the work.



### High AI achievers are more likely to break the rules – and more likely to hide it



Presented by **Work AI INSTITUTE** Powered by **glean**

Men and executives are more likely to work around the rules. 52% of men (vs. 46% of women) and 54% of executives (vs. 48% of non-executives) admit to using unapproved tools or using approved tools in noncompliant ways. Those gaps are consistent with a long-standing pattern in organizational [research](#): the higher your status in a workplace, the more confident you are that the rules are negotiable.

The smart move for organizations isn't to hunt down every workaround like a crime scene. It's to ask what the workaround is telling you. When a high AI achiever routes around a sanctioned tool, it's usually because the sanctioned tool doesn't fit the work. It's too slow, too generic, or too disconnected from the information the worker needs.

At one of our Work AI Institute founding members' universities, they have enterprise agreements with Google and ChatGPT, but not Claude. When faculty need Claude for work, they pay for it through their research budget. Their colleagues without research accounts simply go without. The tool isn't banned. It's just not funded, which means access depends less on what the work demands than on whether the worker has a discretionary budget.

Good leaders treat both kinds of workaround as feedback on their AI strategy. The goal is to make the best tool the easiest, safest, and most approved path. Otherwise, workers get stuck with a rotten choice: follow the policy, or do the work well.



## The team level: How high-achieving AI teams<sup>15</sup> manage AI

For decades, the team has been a stable unit – a manager, a group of human coworkers, and the institutional scaffolding around them: org charts, performance reviews, compensation bands, hiring pipelines, succession plans. All of that scaffolding assumes teams are made of humans.

That assumption is cracking. The team is becoming a hybrid unit, with humans plus a growing cast of AI assistants, copilots, agents, and digital twins. They attend meetings, draft deliverables, make recommendations, and hand work back to the humans who'll be evaluated on it. Workers are no longer sure their human coworkers are the most useful members of the team.

- **AI is now trusted like a coworker:**
  - 53% of workers trust AI as a teammate.
  - 52% find it easier to collaborate with AI than with their human colleagues.

---

- **AI is doing what managers used to do:**
  - 61% say AI helps them more with day-to-day work than their manager does.

---

- **AI is taking the seats humans used to take:**
  - The average worker now collaborates with three AI agents alongside their human teammates.
  - They've sent a bot or a digital twin to one out of six (17%) meetings instead of showing up themselves. AI notetakers join nearly a third.
  - 29% are comfortable with AI being involved in firing human colleagues.

29%

Nearly one in three workers (29%) are comfortable with AI being involved in firing their human colleagues.

High-achieving AI teams aren't waiting for their organizations to redesign the org chart. They're already figuring out how to manage a team where not every teammate is human.

Here's what they're doing.

<sup>15</sup> To understand high-achieving AI teams, we asked high AI achievers (workers who use AI and say it improves both their productivity and the quality of their work) questions about the teams and managers they work with. The team-level findings in this section are aggregated from those individual reports rather than from surveys of intact teams.




1

## They treat AI as a teammate, but keep accountability where it belongs.

High AI achievers don't just use AI differently. They relate to it differently. Less like software, more like a colleague – one that's useful, fast, available at 2 a.m., and still willing to bluff when it doesn't know the answer. High AI achievers are 2.3x more likely than low AI achievers to say they trust AI as a teammate.

### High AI achievers are more likely to treat AI as a teammate, not just a tool



Presented by **Work AI** INSTITUTE Powered by **glean**

High AI achievers are nearly 4x more likely than low AI achievers to trust AI as a teammate.

George Lakoff's work on categorization shows that the labels we attach to things don't just describe them – they shape how we behave toward them. Call AI a *tool* and the relationship turns transactional. You use it, it gives you something back and when it disappoints, you move on. But call it a *teammate* and the dynamics change. You assign it work, push back on the first draft, explain what fell short, and try again. Workers who think of AI as a tool are 26% more likely to give up after a poor output. Workers who view it as a teammate keep pushing – rerunning prompts, swapping models, adding context, and hashing the problem out with the AI.




### When AI underperforms, high AI achievers are less likely to give up and keep iterating

When AI underperforms, workers...

■ High AI achievers ■ Low AI achievers

% of workers who report this



Presented by **Work AI INSTITUTE** Powered by **glean**

This may also help explain why 64% of high AI achievers say collaborating with AI is easier than collaborating with human colleagues. Not because AI is smarter, but because it is easier to work with. It's always available. It doesn't get defensive on the tenth revision of the same paragraph. And, as far as we know, it has never scheduled a meeting that could have been a three-line email.

But the teammate frame has limits. Earlier in this report we showed how humanizing AI can inflate trust and breed more botshitting. Some organizations are blowing past that limit altogether – putting agents on the org chart, assigning them roles, and counting them as headcount. That can go sideways quickly. In a recent randomized experiment, researchers at [Boston Consulting Group](#) found that when AI was framed as an employee rather than a tool, workers felt less accountable for what it produced and reviewed its output less carefully.



The point of the teammate metaphor isn't to make AI feel more human. It's to give workers a working mental model for getting useful work out of it. Treat AI as a teammate, but keep the accountability where it belongs – with the human. And don't put a bot on the org chart until you've built the guardrails that keep humans responsible for whatever it produces.

2

## AI adoption spreads peer-to-peer, not just from the top down.

Many leaders try to push AI adoption from the top and assume gravity will handle the rest. They announce the strategy, bless the tool, fire up a dashboard with up-and-to-the-right vanity metrics, and declare victory at the next all-hands meeting.

Top-down change matters. It tells employees what's allowed, what carries real risk, and which behaviors get rewarded when promotion time rolls around. But memos, mandates, and other top-down change efforts don't, by themselves, change how people work. What changes how people work is watching their peers do it.

The average employee is<sup>16</sup>:



Why do cross-functional teammates carry so much weight? Because they're painfully aware of the coordination tax of work: the bottlenecks, the silos, the duplicated efforts, the dropped balls. So when they build an AI workflow or an agent, they aren't designing for some tidy fantasy version of the work. They're designing for the messy version that they have to deal with, where the marketer needs the data the analyst hasn't pulled yet, and the engineer needs the spec the product manager hasn't written yet. Their workflows spread because they survive contact with real work.

This is why cross-functional workers belong at the center of your AI strategy. Identify them, give them early access to tools and training, ask for their feedback on what's working and what isn't, and invest in AI tools that work across functions. If your engineers and marketers can't build or run AI workflows together in the same tool, you're probably feeding the coordination tax you hoped AI would fix.

<sup>16</sup> Source: Anonymized, aggregated Glean product telemetry, April 2026. We analyzed adoption patterns across Glean users to estimate how an individual's likelihood of adopting Glean Assistant varied depending on whether others in their organization had already adopted it. For each user, we identified prior adoption by at least one direct teammate (same self-reported department and role), cross-functional peer (different self-reported function), or organizational leader (defined as holding a director or higher role based strictly on self-reported job titles), and compared adoption likelihood across those exposure conditions. These findings are descriptive and correlational, not causal.



3

### Good managers<sup>17</sup> use AI to cut coordination sludge, and reinvest the time in their people.

For decades, management has been swallowed by coordination work: chasing updates, routing information, scheduling meetings, forwarding context, keeping the work moving. AI is finally taking some of it off managers' plates. [McKinsey](#) estimates that less than 30% of middle managers' time goes to actual people leadership. The rest disappears into individual execution and admin work.

Managers who are high AI achievers spend about the same share of their time on coordination as less effective managers (11% vs. 9%), but they delegate 32% more of it to AI. They aren't using AI to replace management. They're using it to clear away the administrative sludge that gets mistaken for management.



74%

74% of high AI achievers say AI helps them more with daily work than their manager does.

That sounds like bad news for managers. We think it's the opposite. The best managers don't try to compete with AI on coordination work. They delegate the coordination work to AI, using AI to draft the status update, route the request, summarize the meeting. And they reclaim precious time for the work they ought to be spending more of their time doing – coaching, developing, and inspiring their people.

<sup>17</sup> Good managers are those whose direct reports self-report “I would recommend my manager to others as a good manager.” Those who responded agree or strongly agree were classified as having a good manager; those who disagreed or disagreed strongly were classified as having a bad or average manager.




Good managers turn AI into a growth opportunity, while others leave workers to figure it out alone

What workers say about their manager

Workers with **Good managers**      Workers with **Bad or average managers**      % of workers who report this


AI helps them more day-to-day than their manager does




Had a recent talk with their manager about AI reskilling



Their organization rewards AI skills (promotions, expanded roles)



Rate their manager as a valuable source of AI learning



Presented by **Work AI INSTITUTE**      Powered by **glean**

The payoff doesn't stop at day-to-day work. Workers don't decide how they feel about AI showing up in a performance review, a pay decision, or a termination the moment HR rolls out the system. They decide it through hundreds, sometimes thousands, of micro-interactions with their manager. The manager who uses AI to develop people teaches the team to see AI as developmental. The manager who uses it to monitor people teaches the team to see AI as surveillance. That's, in part, why workers with good managers are far more comfortable with AI showing up in the highest-stakes decisions of their working lives, including performance reviews, pay, and termination. By the time AI walks into one of those rooms, the team has already decided whether it belongs there.



### Workers with good managers are more comfortable with AI in key employment decisions




Presented by **Work AI INSTITUTE** Powered by **glean**

### When AI feels fairer than the boss

About half (44%) of workers already say AI is fairer than their human manager. The bigger the team a manager oversees, the more workers feel that way.

### The bigger a manager's team, the fairer AI starts to look



Presented by **Work AI INSTITUTE** Powered by **glean**



When you manage four people, it's easier to be fair. When you manage ten, or fifty, as some organizations are experimenting with, you start running on shortcuts. You say “yes” to whoever asks first. You keep meaning to circle back to the introverts, and then forget. The result is a system that feels arbitrary and biased toward certain voices. At that point, the cold consistency of an algorithm starts to look a lot fairer than a human stretched too thin.

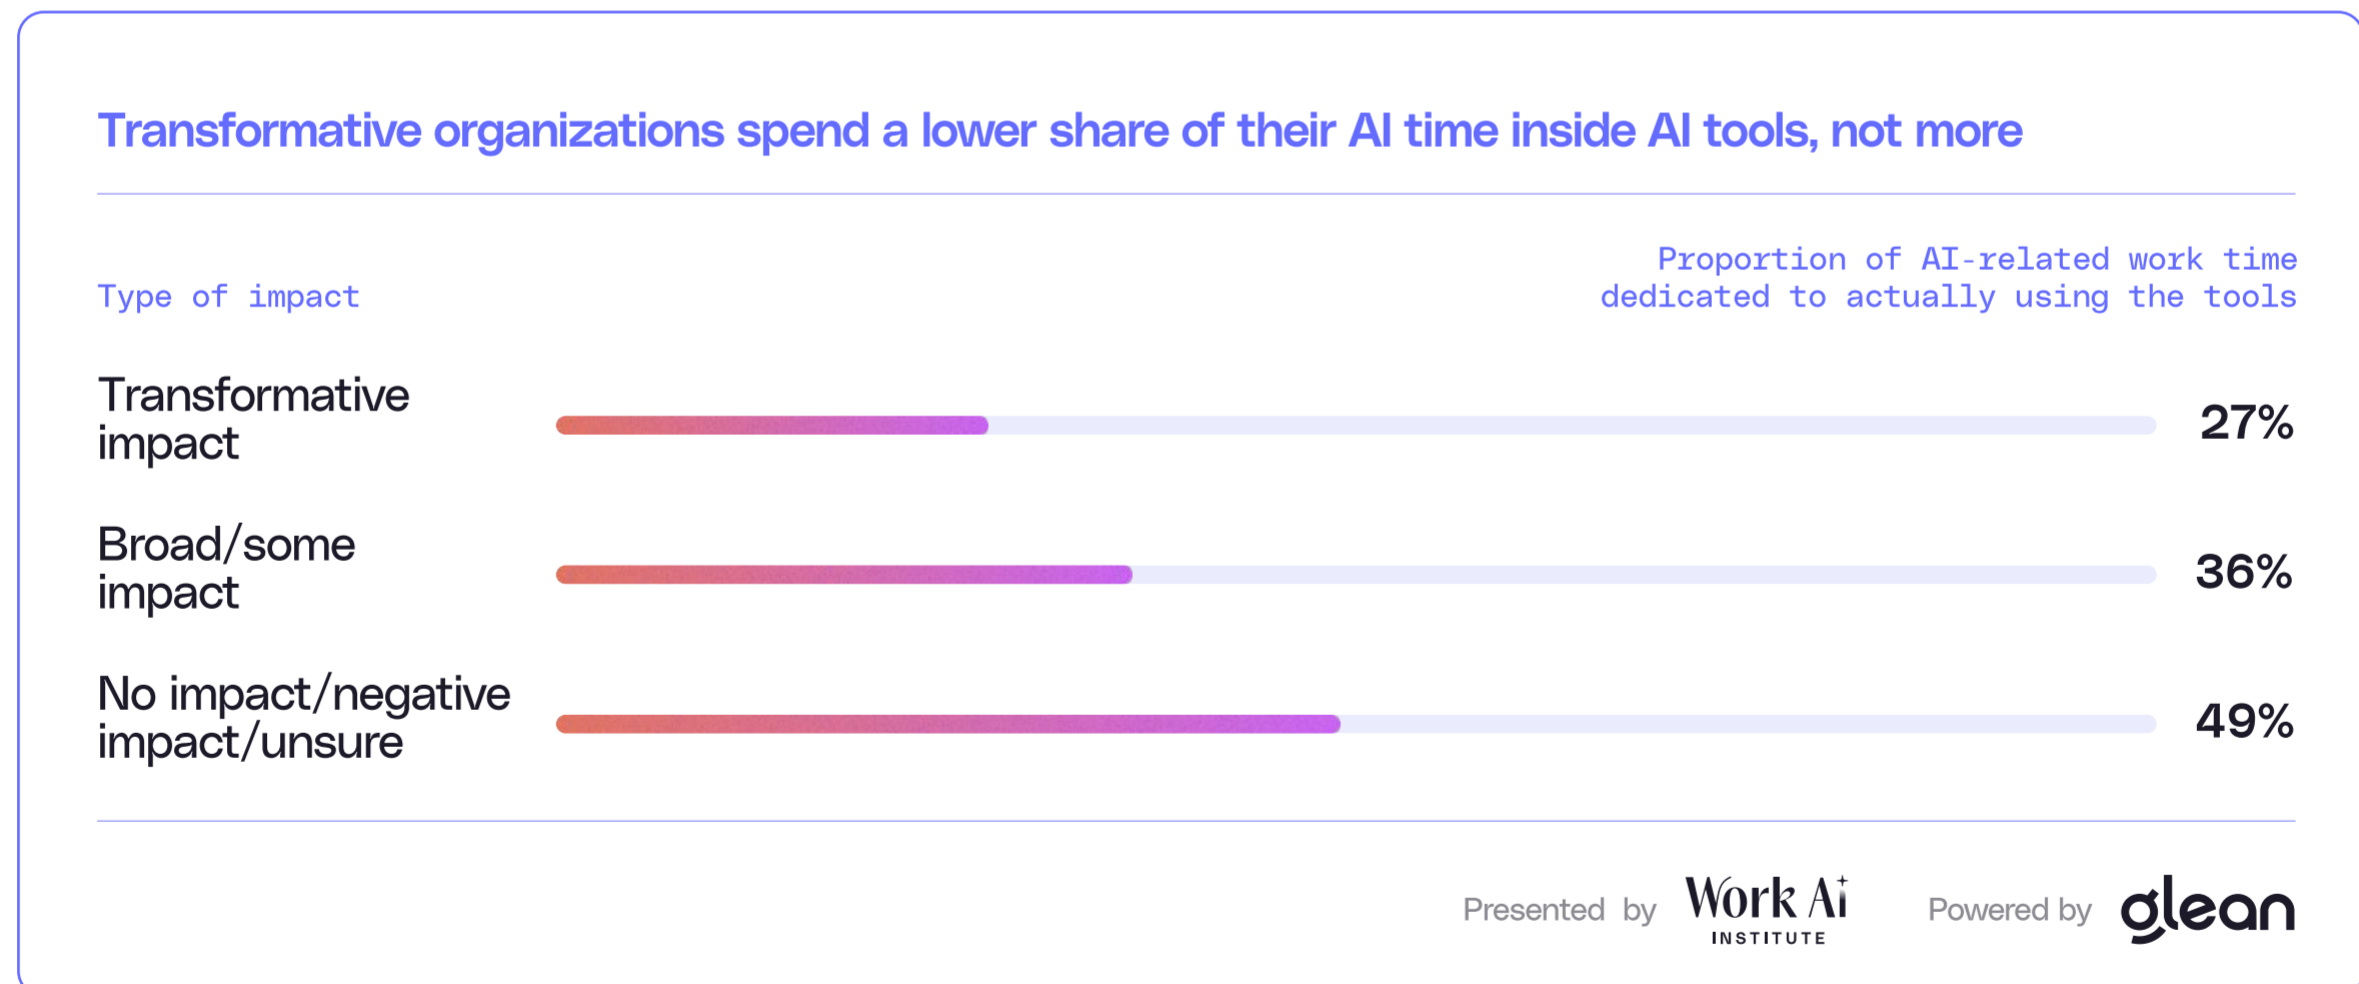
## The organizational level: How transformative<sup>18</sup> organizations make AI work

Workers tell us that AI automation saves them roughly 11 hours every week. Only 13% say their organization is performing significantly better because of it. So what are the 13% doing that the other 87% aren't?

The reflex inside most organizations has been to push AI usage up: more tools, more licenses, more dashboards measuring more clicks. But the organizations actually getting performance gains from AI aren't winning that way. Their workers spend a *smaller* share of their AI time inside the AI tools, not a bigger one.

Workers at organizations reporting no impact, negative impact, or no idea spend nearly half of their AI time (49%) directly inside AI tools, versus just over a quarter (27%) for workers at transformative organizations. The same pattern showed up among the high AI achievers we identified earlier, who cluster in these same transformative organizations.

They aren't getting ahead by maxing out their tool time. They're getting ahead on the work around the tool – setting the context, catching what the tool is too confident about, integrating the results into real work, and knowing when to leave AI out of it entirely.



<sup>18</sup> Throughout this report, “transformative organizations” refers to organizations whose employees report that AI has significantly improved their organization's performance and outcomes.



Transformative organizations build management discipline around AI, such as clear rules, quality checks, shared habits, and enough psychological safety for someone in the room to say, “No, we’re not putting that bot-generated mush in front of a customer.”

The discipline shows up in five key ways.

### 1 They measure what matters, not what’s easiest to count.

Too many organizations measure AI through vanity metrics: tokens consumed, lines of code generated, tool clicks, login rates, dashboards that glow reassuringly upward. We’ve spoken with leaders at major tech firms that count AI-generated lines of code in performance reviews. These metrics are easy to count. They’re also easy to game. When leaders measure AI activity, workers learn to produce AI activity – and, often, little else.

One engineer at a major tech company described doing exactly that: routinely asking AI questions whose answers were already sitting in the company’s documentation, just to inflate his token numbers:

*“I am conscious of not wanting to be seen as ‘uses too little AI,’ and I’m not ashamed to say I need to do tokenmaxxing to do this. Things I do to inflate my token usage metrics: Ask AI questions about the code already in the documentation. The AI pulls up the documentation, processes it, and gives me results 10x slower, but while burning lots of tokens. I could use ‘readthedocs’ [an internal product], but then my token numbers would be lower. Ask the AI to prototype a feature that I have no intention of working on. Prompt it a few more times, then throw the whole thing away. Default to always using the agent, even when I know I could do the work by hand much faster. Then watch it fail.”*

That’s Goodhart’s Law in action: “When a measure becomes a target, it ceases to be a good measure.” If the metric is token volume, workers produce token volume. Whether the work is any good – or just expensive botshit – becomes somebody else’s problem.

Steven Kerr warned about this fifty years ago in his essay, “On the Folly of Rewarding A, While Hoping for B.” Organizations say they want one thing and reward another. They hope for long-term growth and reward quarterly earnings. They hope for teamwork and reward individual heroics. They hope for candor and reward the people who tell the boss what the boss already believes. Now they’re doing it again with AI.



## Common management reward follies

We say we want...

But we actually reward...

**Better work**



The worker who burned the most tokens last quarter.

**AI that reduces overload**



14 new tools, 9 new tabs, and a Slack channel called #ai-wins

**Strategic transformation**



The reorg that put "AI" in the team name

**Employees who know when *not* to use AI**



Employees who use it on everything, all the time

**Better judgment**



The employee who never slowed a deadline down to double-check the AI

**AI that fixes real problems**



AI rollouts that look good in a board deck

**Real AI fluency**



The employee who got featured in the internal newsletter for an AI use case they tried once

Presented by **Work AI**  
INSTITUTE Powered by **glean**

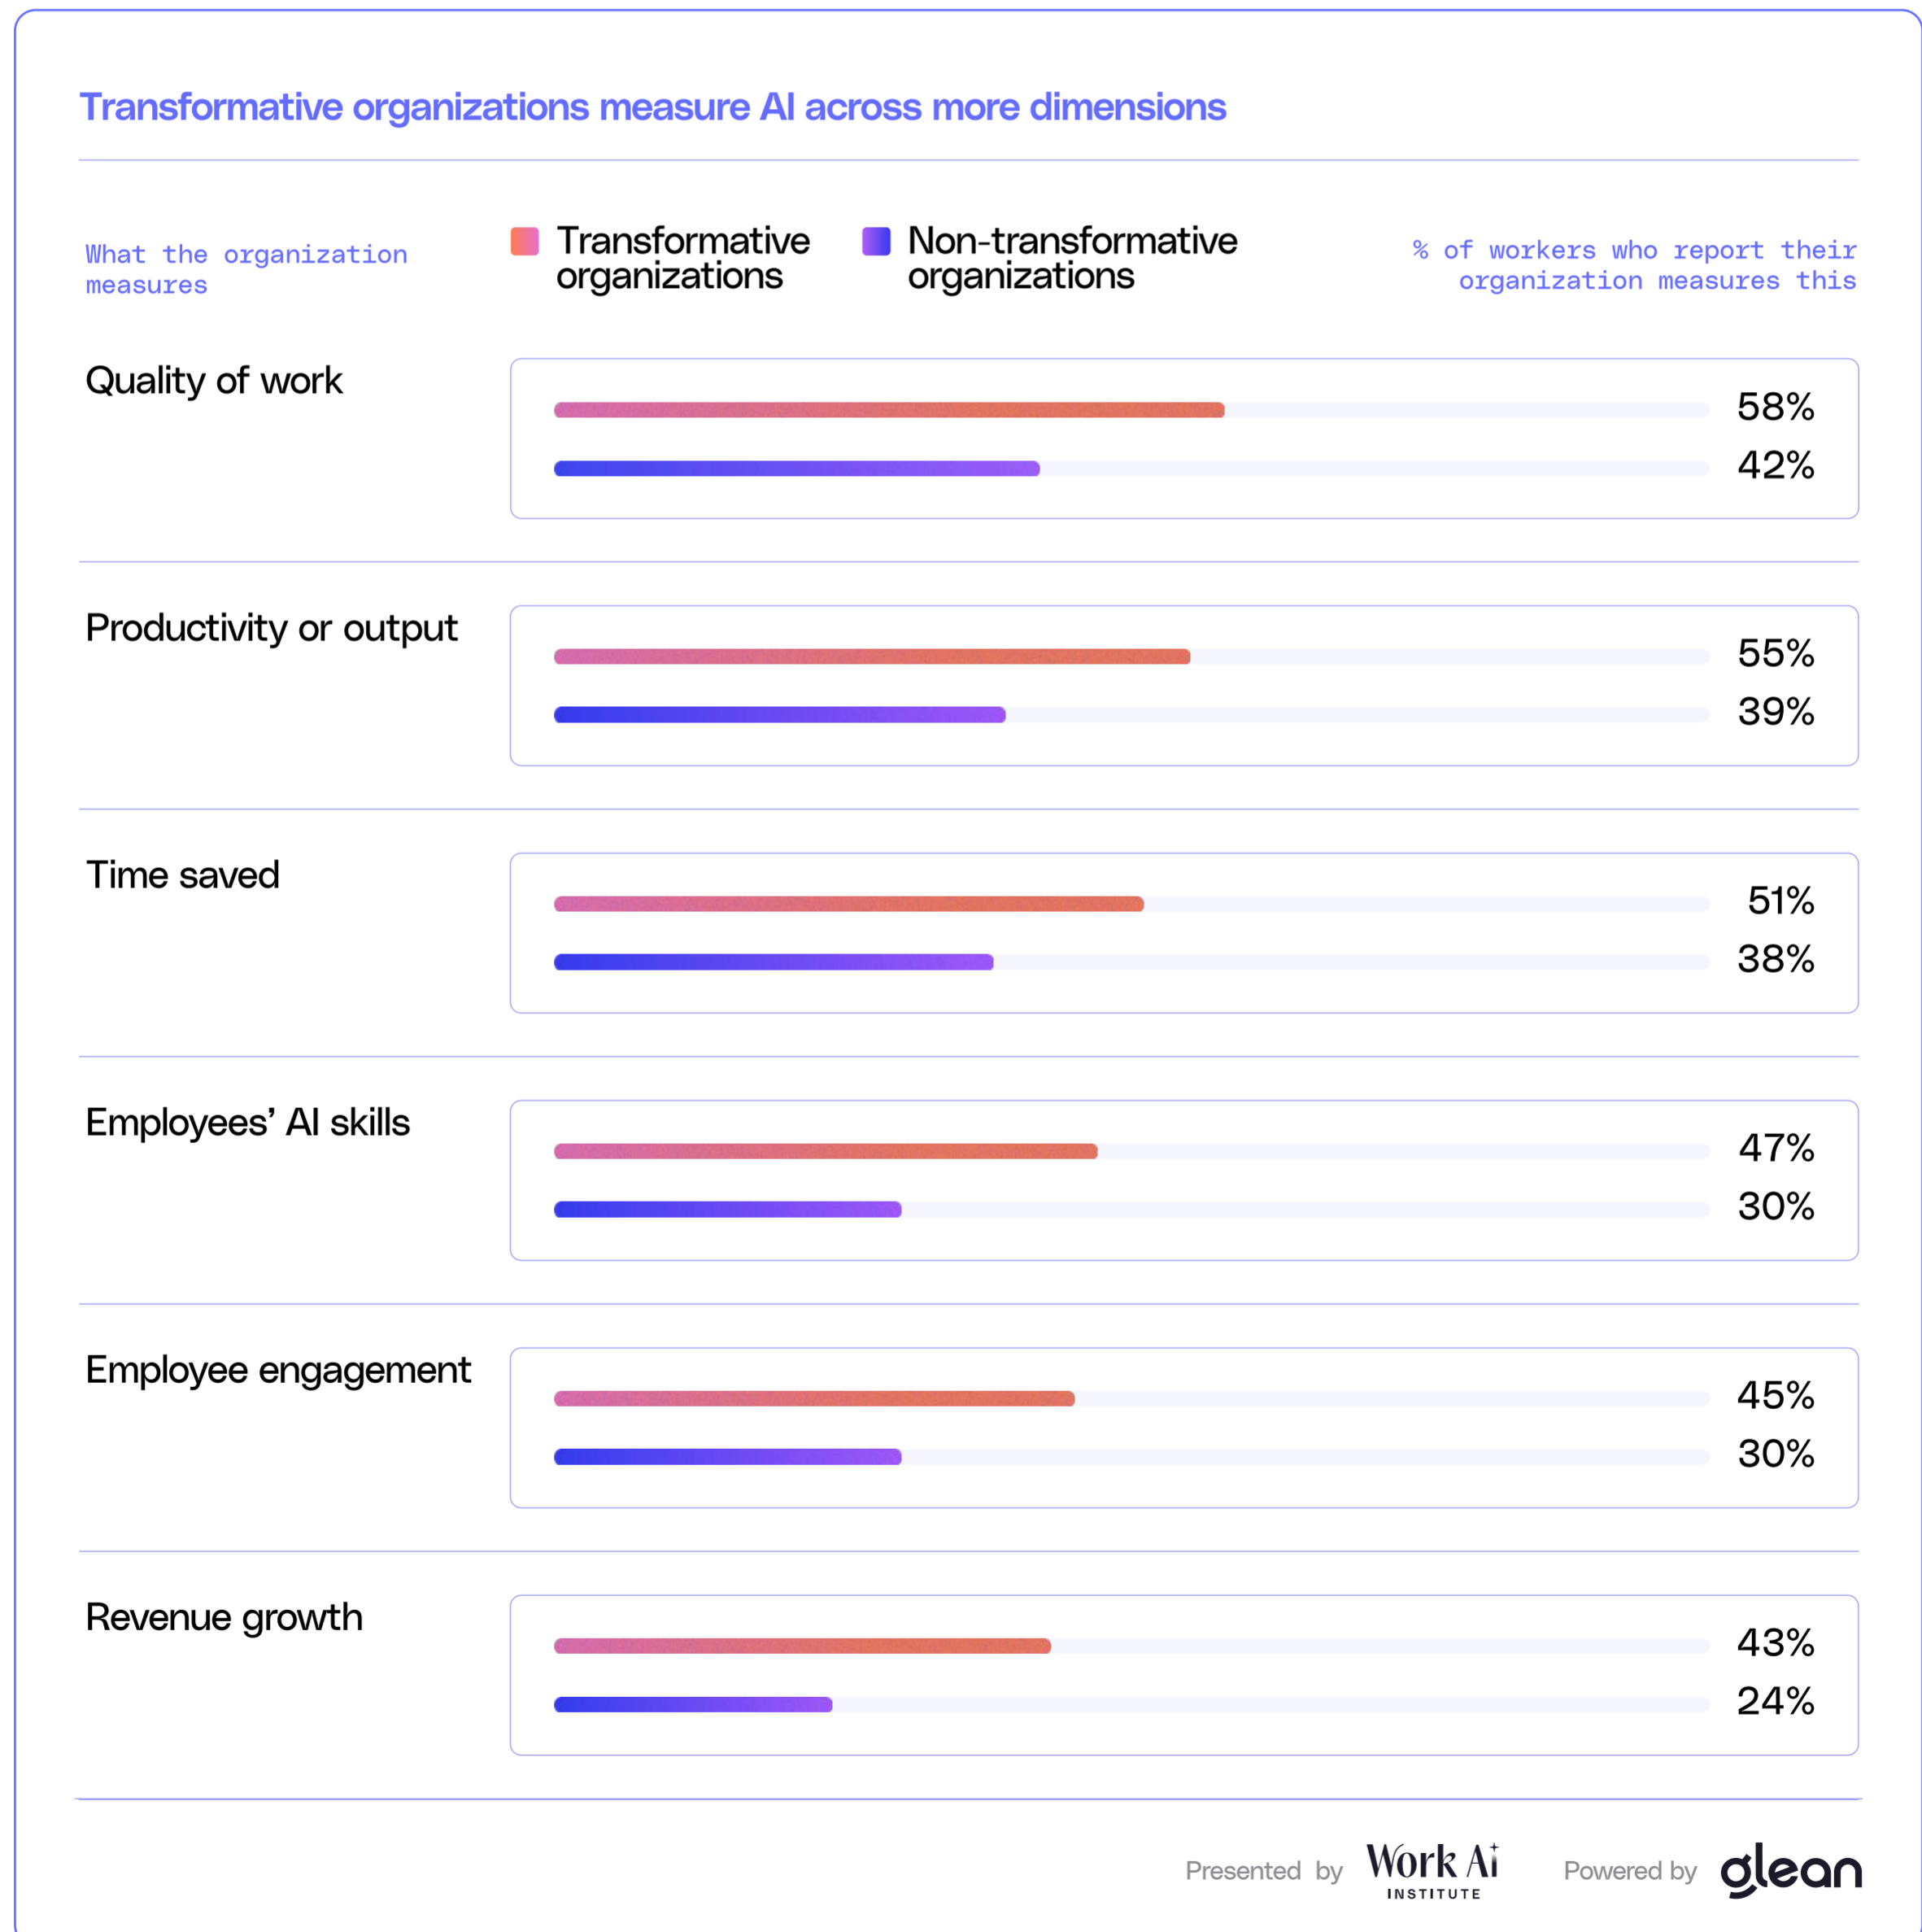
There's a heavy cost tied to measuring the wrong thing. 74% of workers in organizations that measure only productivity admit to botshitting. Where both productivity and quality are measured, it drops to 64%. And when organizations track quality alongside productivity, 83% of workers say AI has improved the quality of their work, compared with just 68% in organizations that track productivity alone.

Metrics don't just measure behavior. They tell employees what the organization values and, in turn, they incentivize certain behaviors. When organizations measure quality, people slow down where human judgment matters. When they measure only speed, people speed up, pumping out faster drafts and leaving downstream cleanup to someone else.



Transformative AI organizations build measurement systems that reward the behaviors they actually want. At one organization we studied, AI success is evaluated across three dimensions: efficiency, quality, and employee experience – so the company doesn’t reward fast work at the expense of good work or good people. One Fortune 100 executive told us his company measures “intent diversity”: how many different use cases employees apply a given AI tool to. It’s a better proxy for real learning because it shows whether workers are expanding their grasp of where AI is actually useful, not just clicking the tool more often.

Transformative organizations know no single metric tells the whole AI story. They evaluate AI initiatives across five different dimensions on average, compared with three at other organizations.





Workers in transformative organizations say their organizations track an average of five different metrics, compared with three in other organizations.

In transformative organizations, employees are far more likely to say they have visibility into their own AI usage data. The data flows in two directions: up to the people running the strategy, and back to the people doing the work. Where it only flows one way (upward), it stops being feedback and starts feeling like surveillance. Workers figure out the message: *This isn't here to help me learn; it's here to judge me.* So they game the numbers, hide the messy parts, and perform whatever version of “AI adoption” leadership seems to want.



2

## They make governance a living system.

Most organizations have an AI policy by now. That doesn't mean they have AI governance.


A policy is a document someone uploaded to the corporate intranet and hoped people would read. Governance is what actually happens when the deadline is breathing down someone's neck, the approved tool is clunky, and the unapproved one works better. Right now, there's a wide gap between the two.



### Where AI policies fall short in practice

Governance gap

% of employees who say this



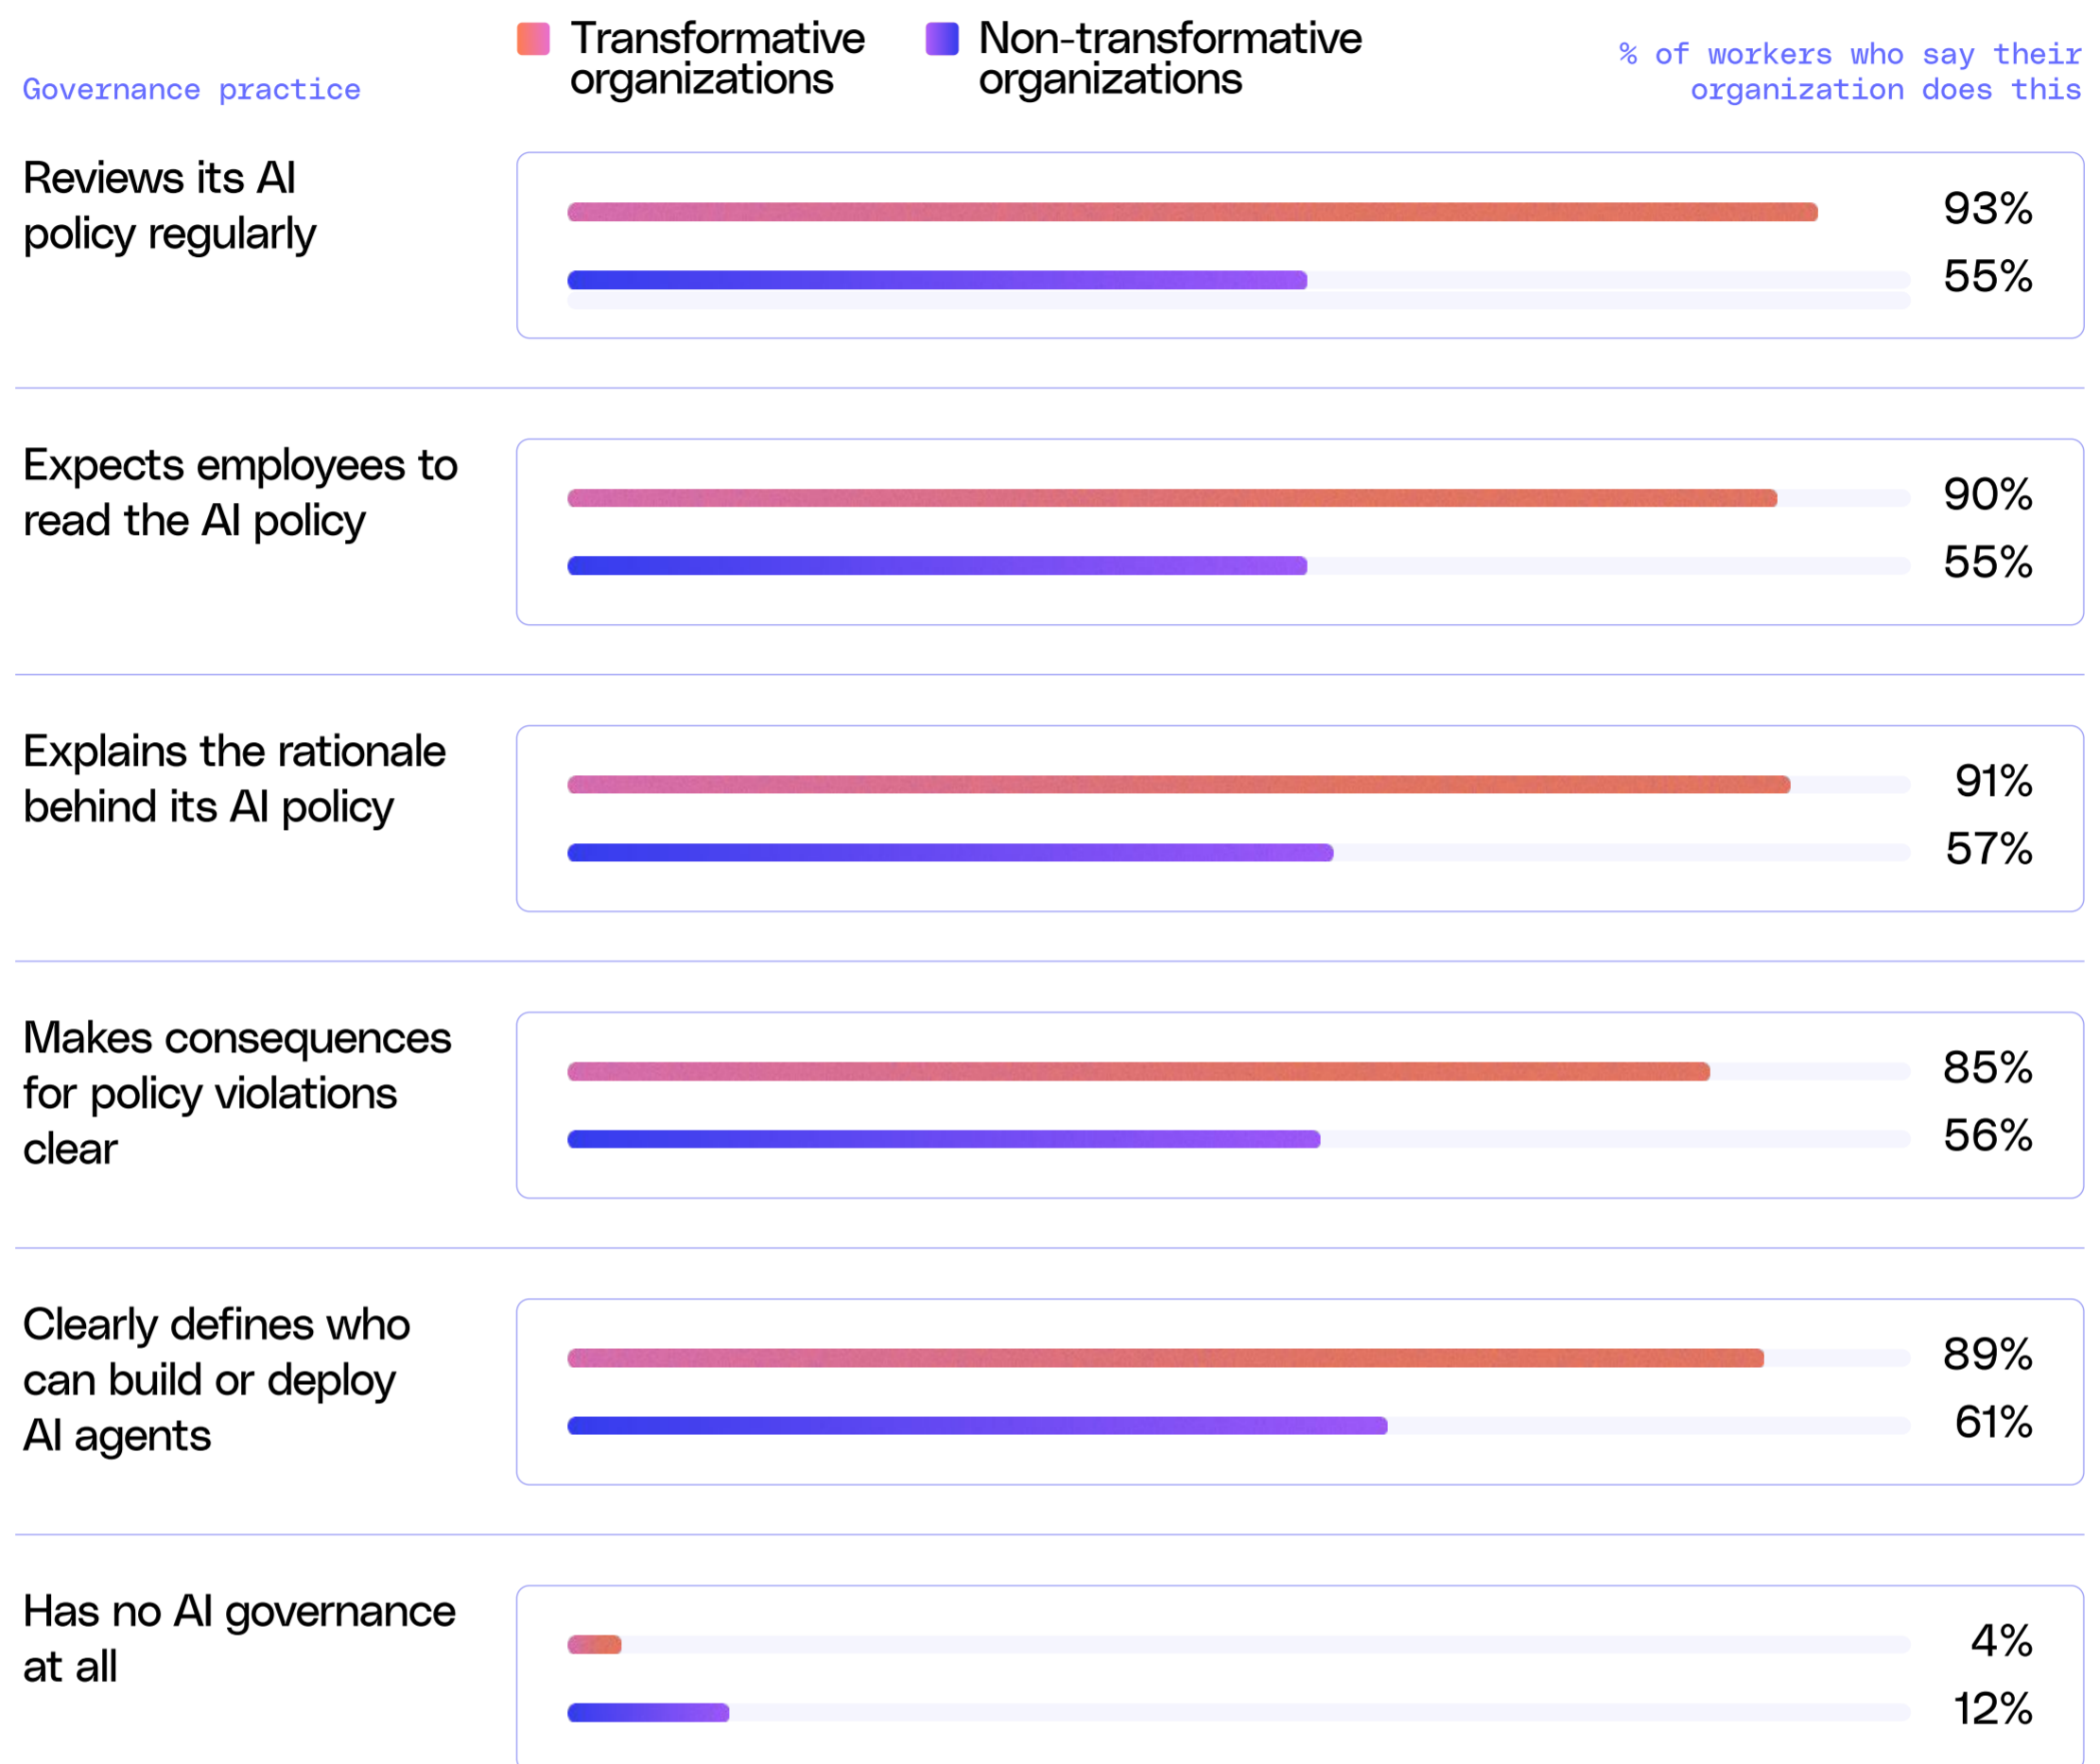
Presented by **Work AI**  
INSTITUTE Powered by **glean**

Transformative organizations don't just write a policy, post it, and wander off. They review it regularly, explain why it exists, and enforce it when somebody breaks it. Only 60% of workers say they've read their organization's AI policy in full. For everyone else, it's just another prop in the AI governance theater – written for the auditor, displayed for the board, ignored by everyone with actual work to do.

Transformative organizations also define who's allowed to build and deploy AI agents. Without that line drawn, you get agent sprawl – three different teams have built three different bots to do the same thing, two of them are running on data the company never sanctioned, and nobody can remember who approved any of it.



### Transformative organizations have clear, active, and trusted AI governance



Presented by **Work Ai** INSTITUTE Powered by **glean**

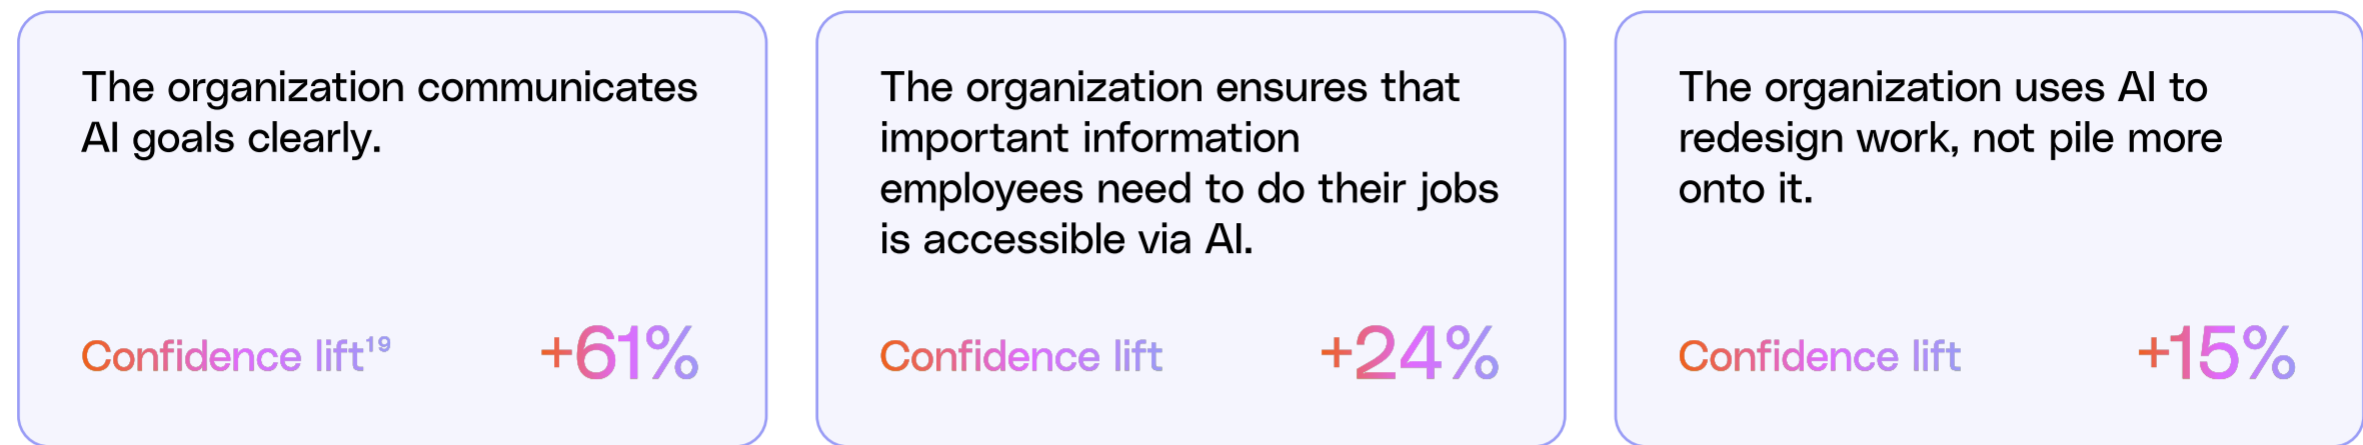
The payoff isn't just compliance. It's confidence. Workers in transformative organizations are far more likely to say they trust their company's AI strategy (93% vs. 57%). And that confidence drives retention.

**28%**

Employees who are confident in their organization's AI strategy are 28% less likely to be actively job-hunting.

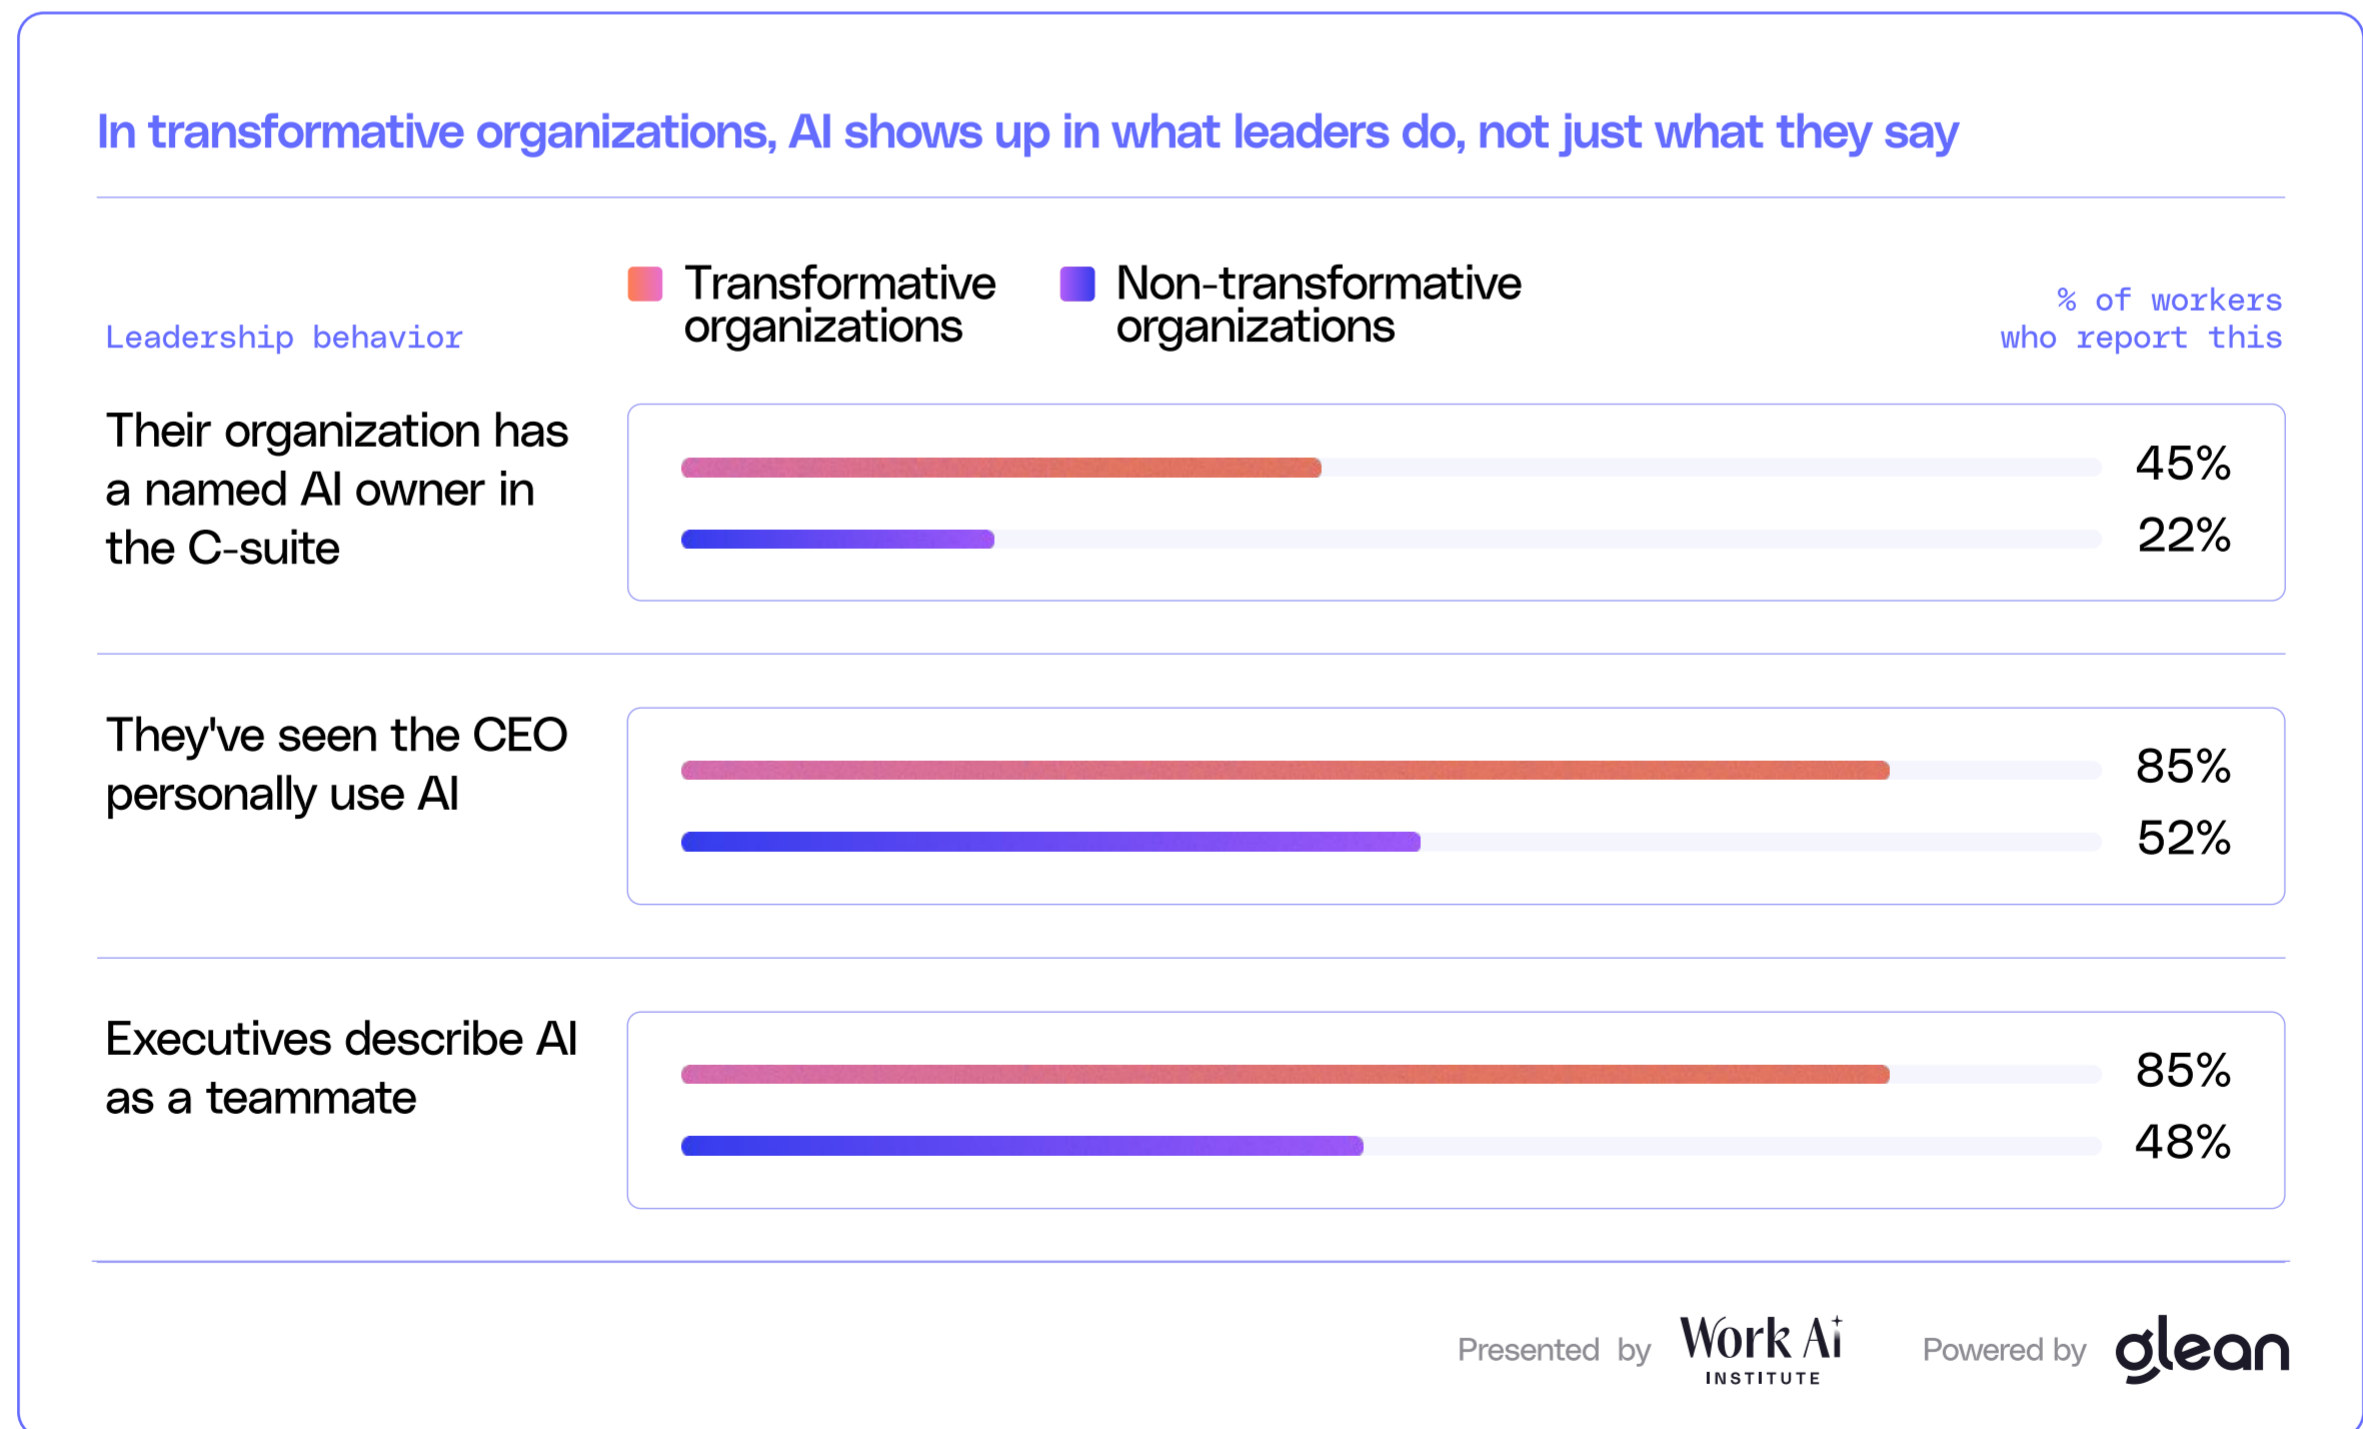


Three actions build that confidence most:



When leaders communicate AI goals clearly, workers stop filling in the gaps with worst-case scenarios. When the information they need is actually accessible through AI, the tool stops producing generic answers and starts producing useful ones. And when AI is used to redesign the work rather than load more onto it, the credibility shows up in visible relief – not in another round of promises about what AI will eventually deliver.

Transformative organizations don't treat governance as a communications exercise. They give it decision rights, accountability, and budget. They're more likely to put a named AI owner in the C-suite. And their CEOs don't just talk about AI in town halls. They use it – in front of their employees. Nvidia CEO Jensen Huang has been publicly clear that he uses AI daily, "as a tutor, a research assistant, a coach, a thought partner."



<sup>19</sup> Estimates come from an ordinal logistic regression predicting respondents' confidence in their organization's AI strategy (five-point Likert scale) as a function of organization-level factors, with controls for sociodemographics, firmographics, and occupational status. Reported confidence lifts translate model estimates into changes in predicted probability of reporting higher confidence associated with each factor.



Employees who have seen their CEO personally use AI use it 67% more than those who haven't. Employees whose executives describe AI as a teammate are 51% more likely to treat it as one themselves, and 65% more likely to report productivity gains.

Employees watch what their leaders do with AI, not what they say about it. When leaders hype AI from the stage, monitor token counts from a safe distance, and never use the tools in the daily work themselves, employees figure out the real rules quickly: look AI-forward, keep the dashboard pointed up, and leave someone else to clean up after the bots.

3

### They start with the work, not the tech stack

For many organizations, their “AI strategy” is little more than a roll-up of their existing contracts with technology vendors. If they're a Microsoft shop, they add Microsoft Copilot. Salesforce shop? Add Einstein. Already paying for Google Workspace? Welcome to Gemini.

This is an AI version of what James March and his colleagues called the garbage can model fifty years ago. In real organizations, problems and solutions don't arrive in that order. Often, the solution is already rattling around inside the can, waiting for someone to find a problem to attach it to. That's how most AI strategies unfold. The company buys an AI tool from an existing vendor, or gets one bundled into a contract they already have, and then goes hunting for places to use it. The tool comes first, the problem comes later.

Transformative organizations invert the order. They start with the work before they start with the shopping. They map where employees are stuck doing low-value drudgery, where customers keep getting frustrated, and where handoffs keep dropping. Then they pick tools that solve those problems. Employees in these organizations are much less likely to say existing vendors are limiting their AI strategy (33% vs. 49%).

Starting with the work also changes what these organizations expect from their vendors. The vendor has to understand how the work actually happens, connect to the data and context the work depends on, and stay involved long enough to define what success looks like. Employees at transformative organizations are nearly twice as likely to describe their AI rollouts as collaborative with vendors (78% vs. 44%).

This is part of what's driving the rise of *forward-deployed engineers*, people whose job is to make AI work inside the actual organizations that buy it. Job postings for the role, which originated at Palantir, grew more than 800% between January and September 2025. Salesforce alone has committed to hiring 1,000 of them. In May 2026, OpenAI launched a \$4-billion deployment company and acquired the AI consulting firm Tomoro, instantly adding 150 forward-deployed engineers to its team.



4

## They ground AI in enterprise context, not just enterprise data.

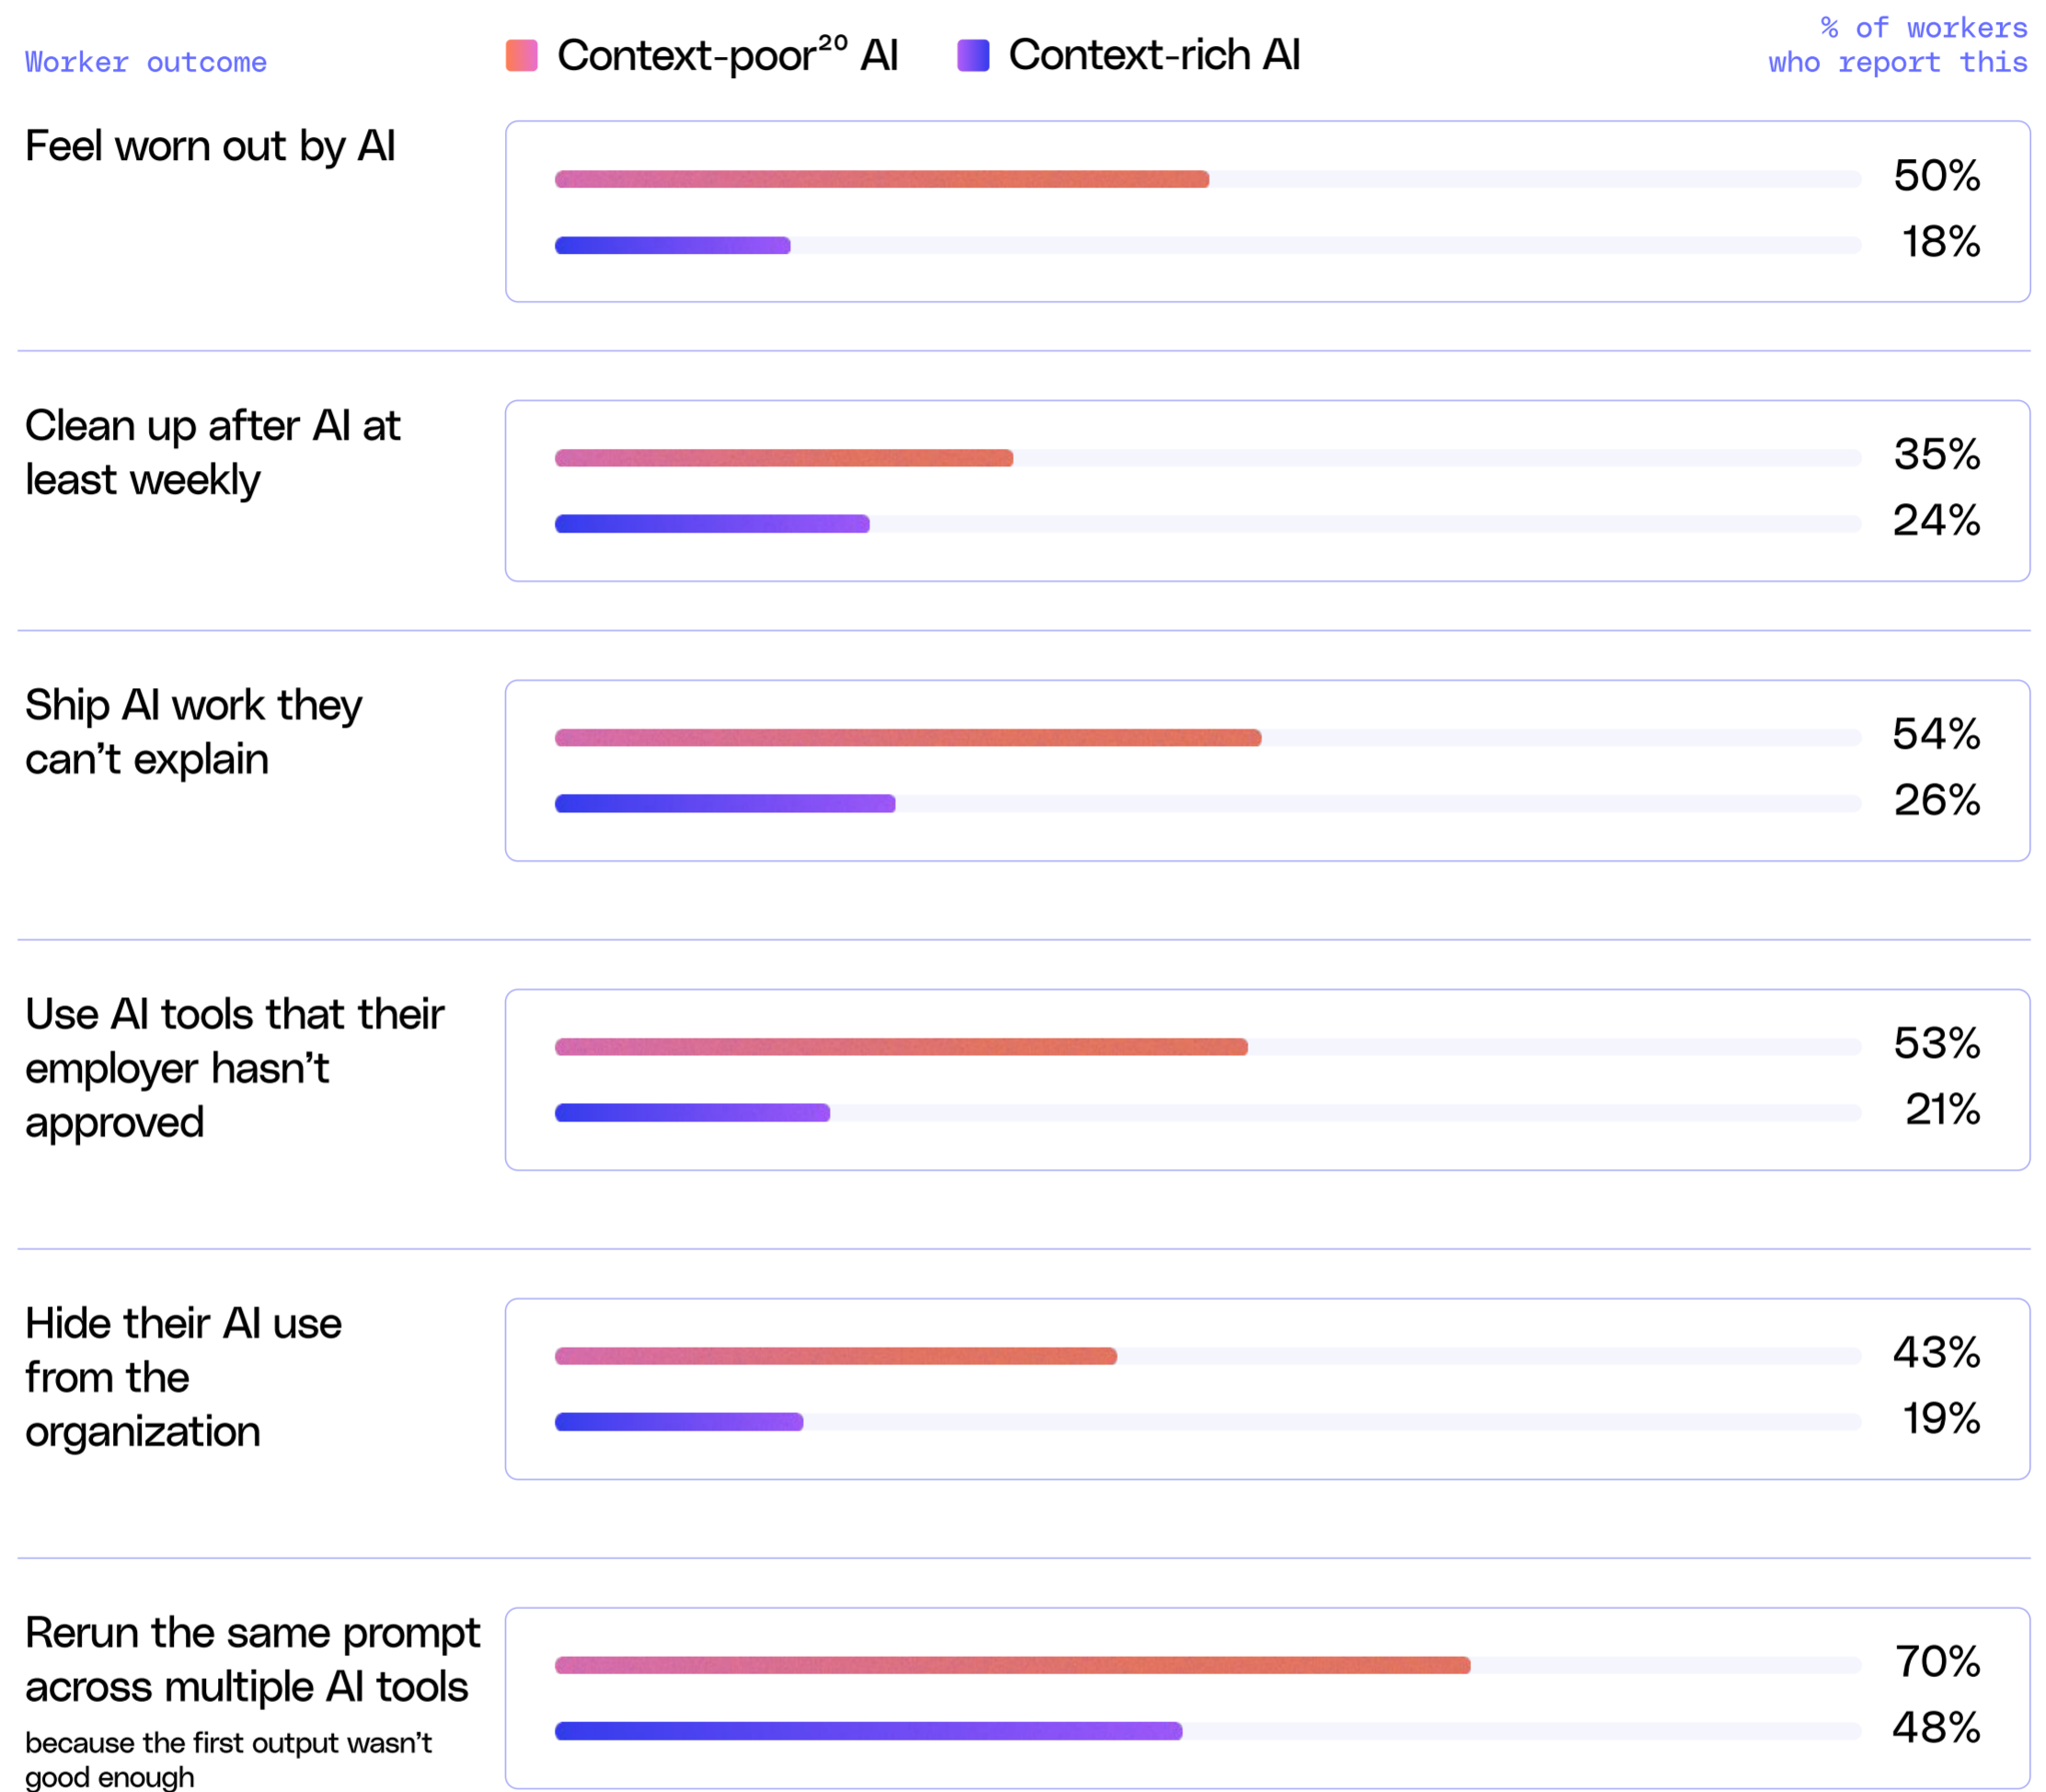
Starting with the work changes what matters in the technology. It's not enough for an AI tool to be connected to the company's data. It has to use that data the way the people doing the work actually use it.

Enterprise context is the layer between the raw data and the usable answer. It's what tells a system which file is the actual current version of the deck, whether a spike on a dashboard means something or is a known artifact, which workaround the operations team has been using when the official process breaks down, and what "Q3" means at your company this week versus last quarter.

When context is missing, employees pick up the slack. They smuggle work into unapproved tools, hide their usage, and run the same prompt across multiple platforms, hoping one finally coughs up something usable. In our data, we treat workers who say the information they need to do their work is accessible through their AI tools as "context-rich," and those who say it isn't as "context-poor." The gap between the two groups is striking.



**Context-poor AI is associated with more fatigue, more cleanup, and more shadow usage**



Presented by **Work AI INSTITUTE** Powered by **glean**

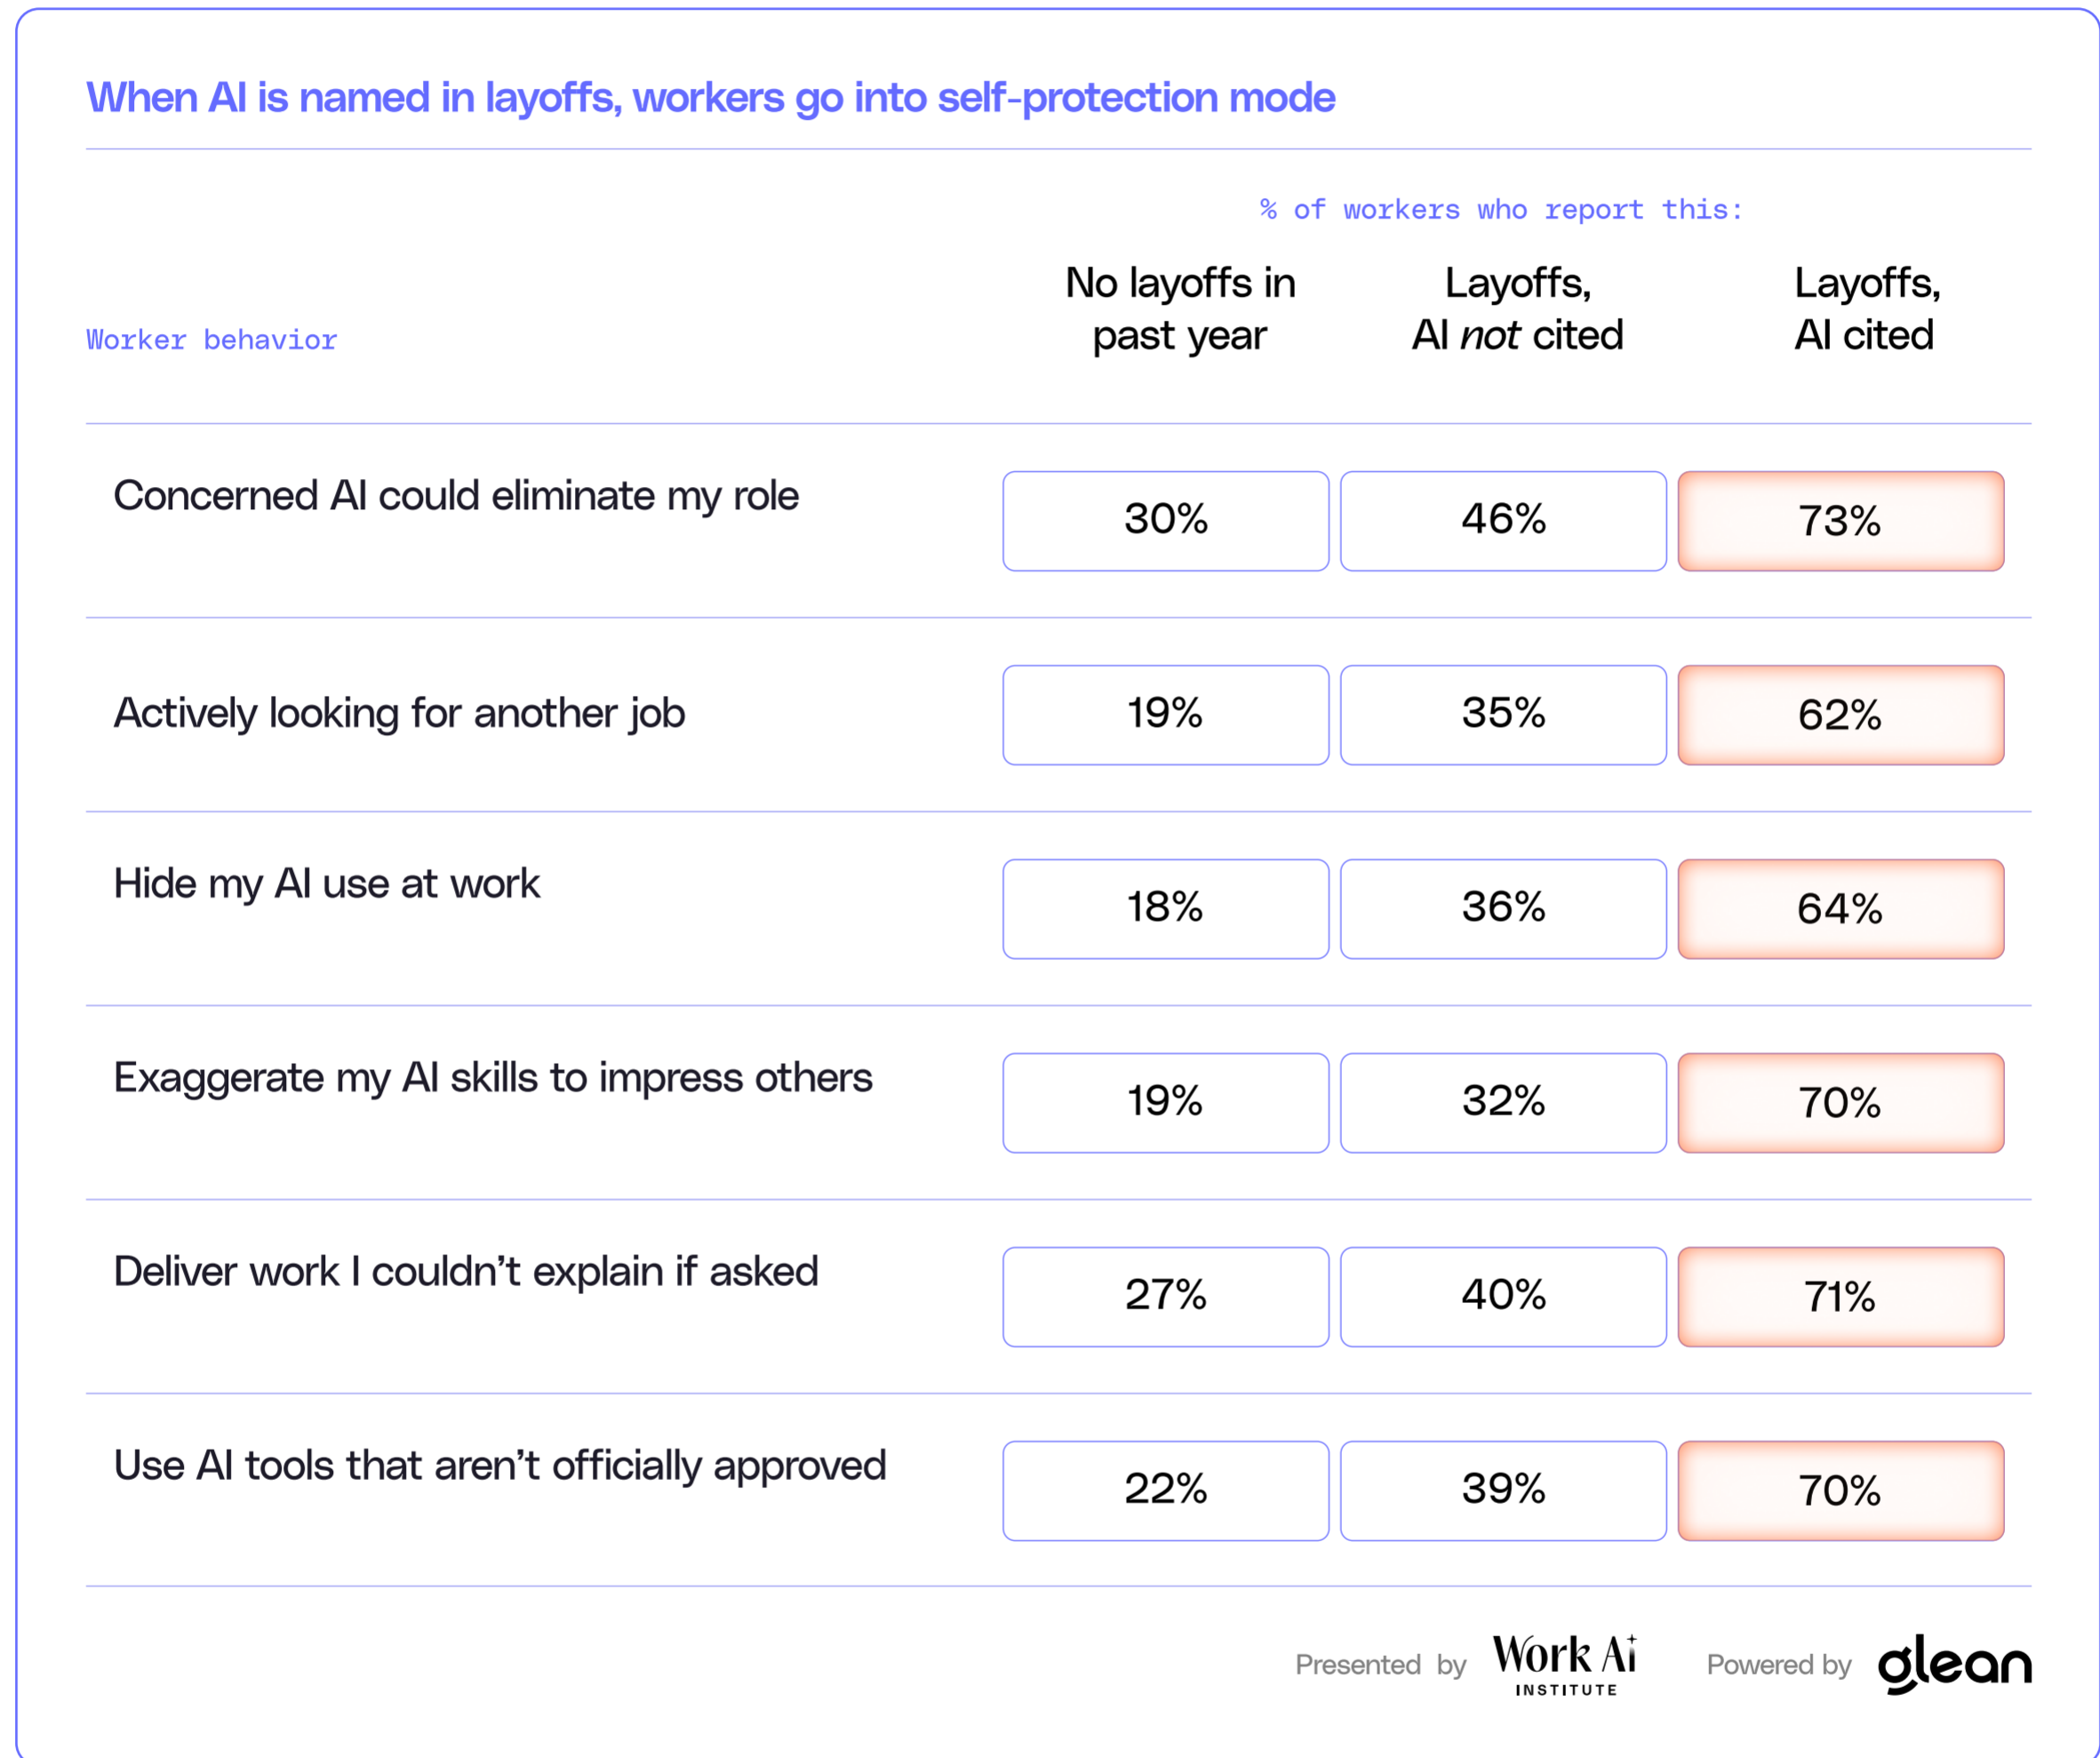
**5 They invest in their people.**

AI transformation is disruptive. Roles shift. Some disappear entirely. And when companies cite AI in layoffs, the fear cuts deeper: 73% of workers worry their own role could be next.

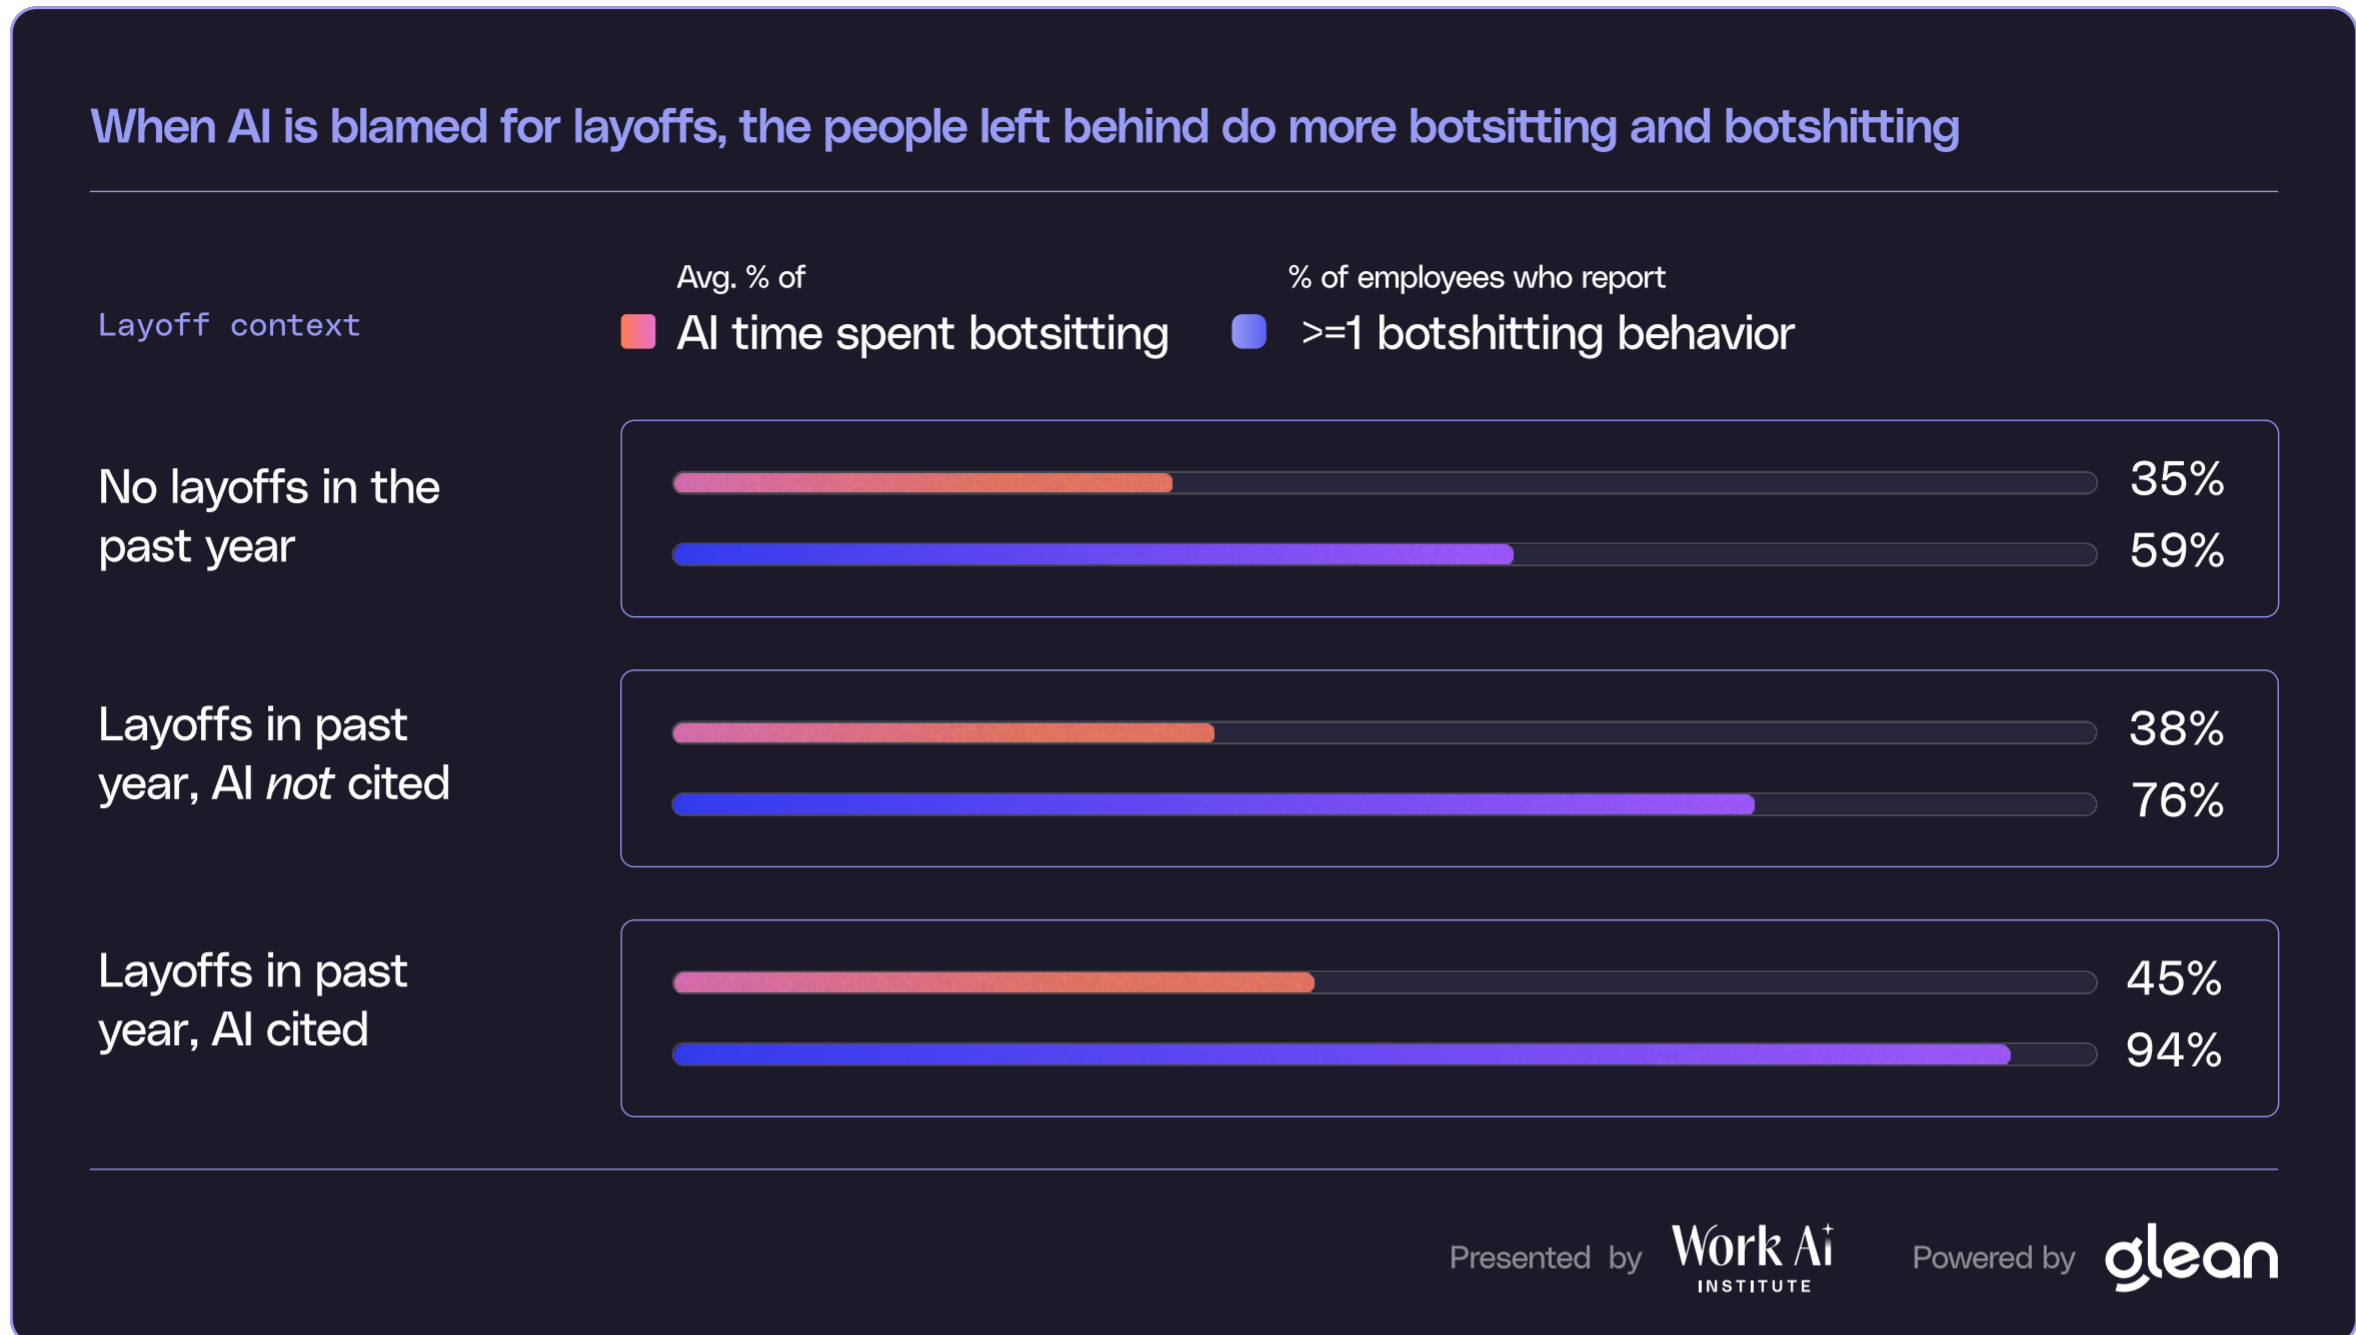
<sup>20</sup> Respondents were classified as having context-poor AI if they agreed or strongly agreed with the statement "Important information I need to do my job is not connected to or accessible through AI tools."



The impact of layoffs doesn't end when the laid-off employees leave the building. Research on "layoff survivors" shows that downsizing increases job insecurity, erodes trust in management, and raises turnover intentions among the people who remain. When AI is cited as the reason, workers start optimizing for self-protection over careful judgment – moving fast, hiding how they use AI, pushing through work they don't fully trust. Many start looking for the exit. At organizations that cited AI in layoffs, 62% of workers are actively job-hunting.



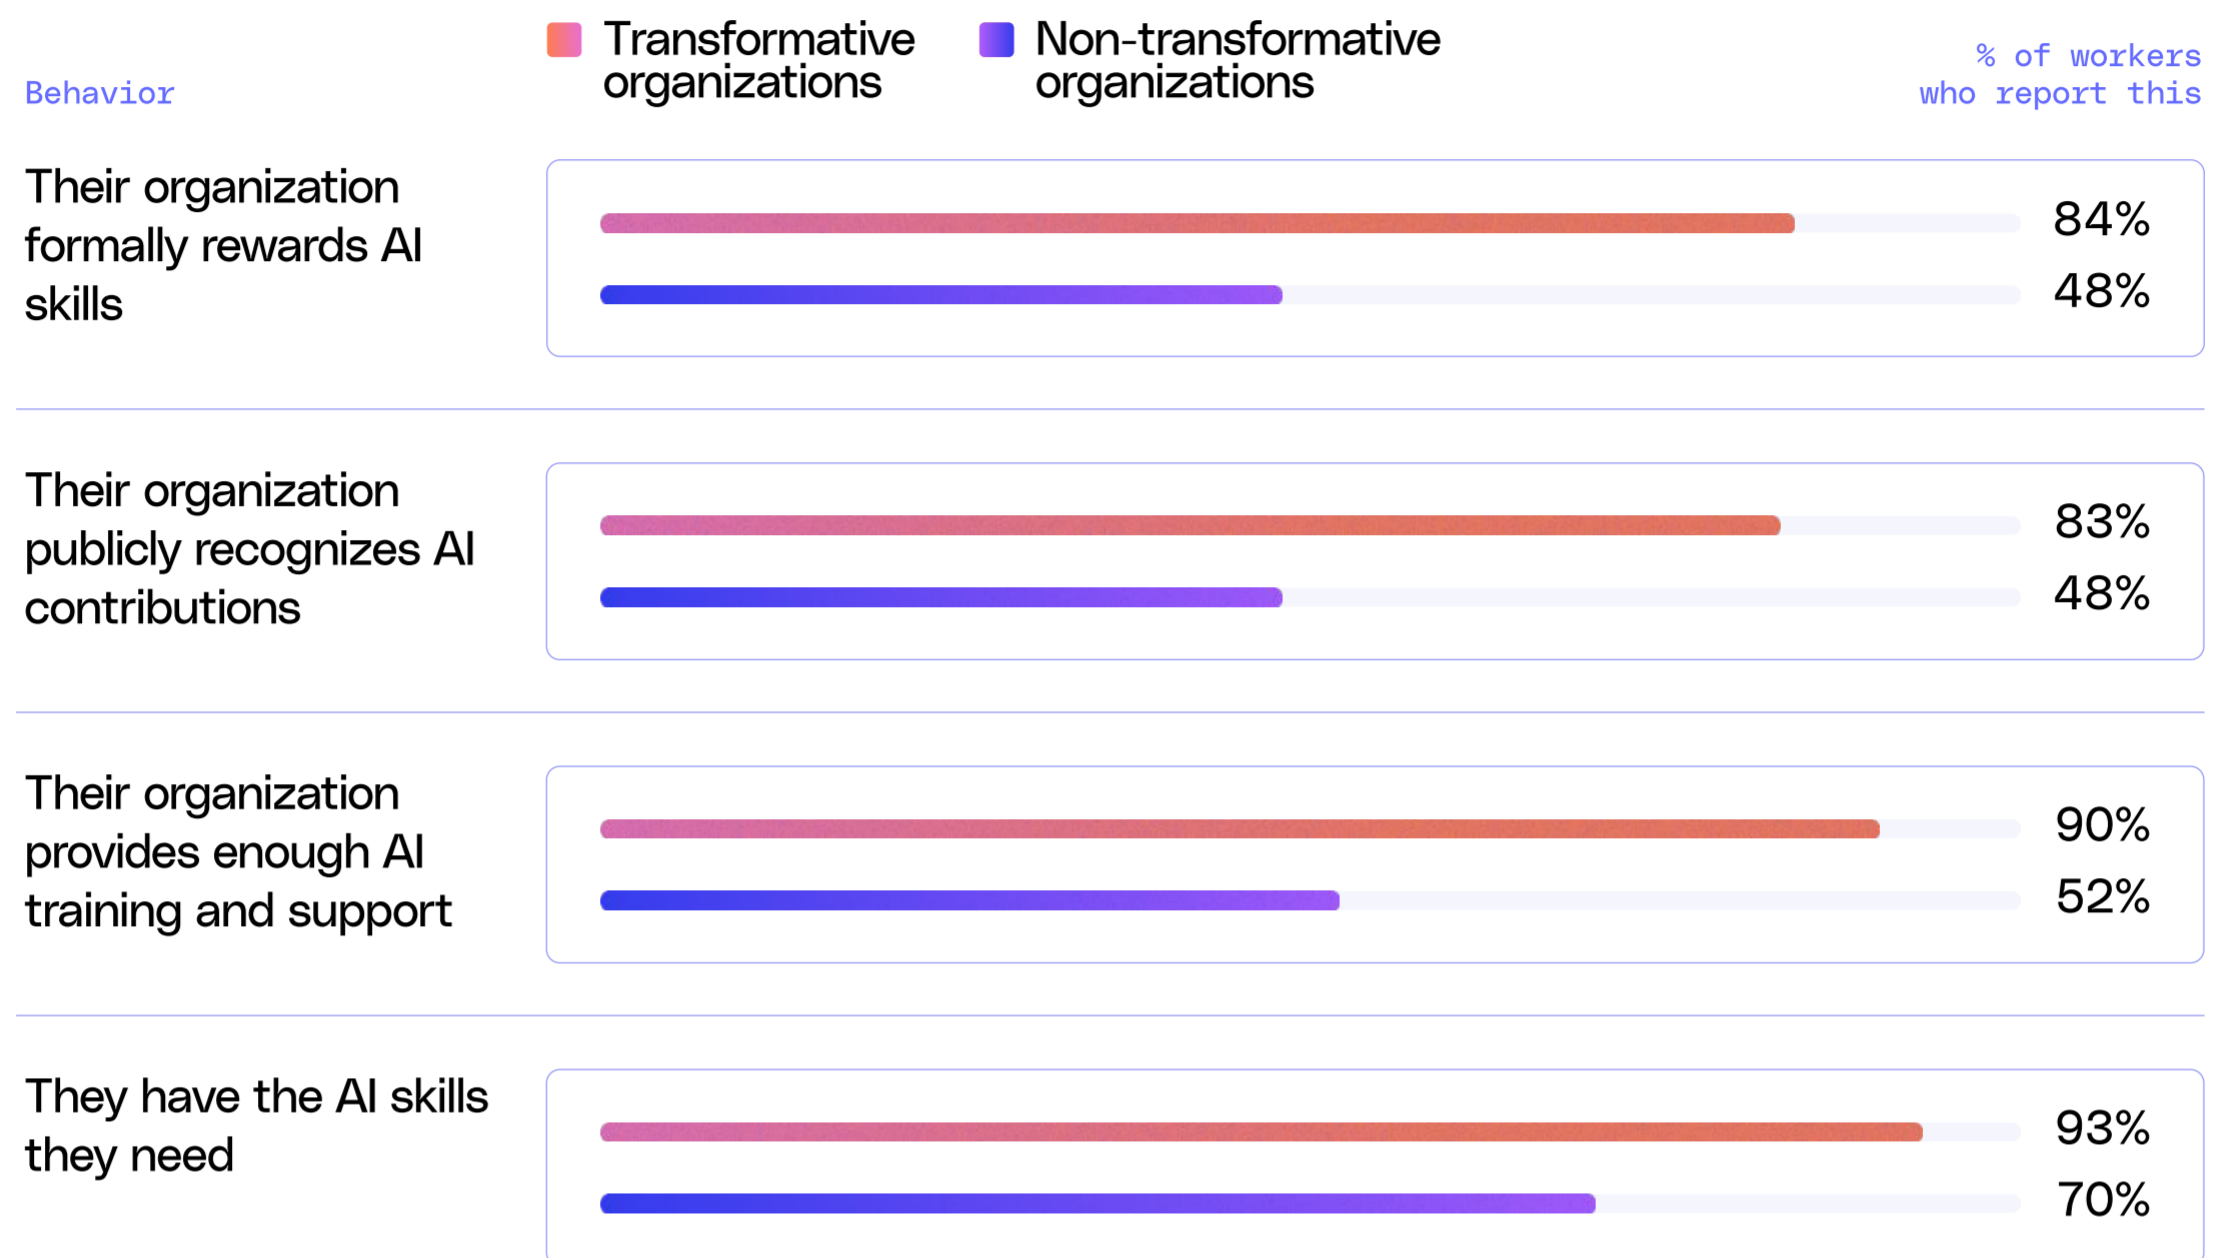
Both botsitting and botshitting are higher in organizations that have done layoffs, and higher still when AI is named as the reason. Layoffs put more pressure on the people who stayed, so they lean harder on AI – and they ship more work they wouldn't defend if asked.



Transformative organizations break some of this cycle. They don't just hand workers new tools and hope they'll magically become AI-fluent. They reward AI skills, recognize people for good AI work, celebrate well-intentioned failures, and provide enough training and support that AI fluency becomes a normal part of professional growth, instead of a precondition for keeping the job.



### Transformative organizations invest in the people, not just the tools



Presented by **Work AI INSTITUTE** Powered by **glean**

One of the most important things transformative organizations do is treat AI as a chance to redesign the work – not as a shiny excuse to squeeze more output from fewer people. 90% of workers at transformative organizations say their employer treats AI as a chance to redesign work, compared with 54% at other organizations.



## SECTION 08

# Conclusion: The Real Work of AI Transformation

Workers tell us that AI automation saves them roughly 11 hours every week. But only 13% say their organizations are performing significantly better as a result. The organizations closing this gap aren't doing it by buying more AI tools, burning more tokens, or building adoption dashboards that glow a triumphant shade of green. They're doing the harder work of treating AI as a work-design problem, not a procurement one.

And they're doing it at every level: individual, team, and organizational. They're helping their people cut the hidden human labor AI creates. They're measuring whether the work produced is better, not just faster. And they're investing in AI tools and platforms that cut the toggle tax and ground AI in the context that actually makes it useful.

Until more organizations do this, the other 87% will keep learning the hard way that AI's time savings aren't free. The hours workers "save" come back as botsitting. The judgment they offload comes back as botshitting. The workplace fills up with work that looks finished, sounds confident, and is hollow enough that some exhausted human – usually without credit or reward – still has to mop it up.

That's the choice in front of every organization. Build the human infrastructure that makes AI worth using. Or keep paying the bill – in botsitting, in botshitting, and in the steady departure of the people who got tired of cleaning up after the bots.



## SECTION 09

# Appendix

## Appendix A: AI at work by region

While botsitting and botshitting are pervasive across all three countries, they show up in different ways. In the US, AI adoption is largely driven by bottom-up experimentation, with workers absorbing much of the botsitting burden on their own. In the UK, organizations appear to give workers more structure, training, and support, which may help explain why UK workers report the strongest productivity gains despite high levels of AI use. Australia matches the UK on adoption, but weaker training and poorer context infrastructure push more cleanup and verification work back onto employees, contributing to higher fatigue, less consistent gains, and greater exposure to botshitting.

*Percentages reflect the share of digital workers in each country (US n=3,000; UK n=1,500; Australia n=1,500) who selected the corresponding response. Cross-country differences of three percentage points or more are statistically significant at  $p < .05$ .*

### United States

#### The posture

Lower adoption, moderate gains. 84% of US workers report using AI at work, the lowest among the three countries, compared with 87% on average across the three countries. US workers report solid but middling productivity impact: 75% of US AI users say AI makes them more productive, compared with 78% in the UK and 72% in Australia.

#### How AI gets absorbed

Worker-first, policy-second. 48% of US AI users say their own preferences drive AI tool choice, compared with 43% in the UK and 35% in Australia. 57% say they have fully read their employer's AI policy, compared with 65% in the UK and 59% in Australia. In the US, AI adoption is spreading from the individual outward, with governance and institutional structure racing to catch up.

#### What workers are doing

Solving AI's friction themselves. 60% of US AI users sometimes rerun the same prompt across multiple AI tools, nearly identical to the UK at 60% and Australia at 59%. 35% redo the work themselves rather than accept a flawed output, compared with 29% in both the UK and Australia. When AI breaks down, US workers are more likely to absorb the botsitting burden personally.

#### The tradeoff

The individual is carrying the organization. US adoption is the most individually driven: 48% of US workers report picking their own AI tools, the highest of the three markets. US workers are also the least likely to deliberately hold back to avoid overreliance, at 51%. That bottom-up energy can help drive productivity gains, but with the least clear governance over who can build or deploy agents, at 63%, the tradeoff is fragmentation and compliance risk rather than fatigue.



## United Kingdom

### The posture

Highest adoption, strongest gains. 90% of UK workers report using AI at work, compared with 84% in the US and 90% in Australia. 78% of UK AI users say it makes them more productive, compared with 75% in the US and 72% in Australia, the strongest reported gains of the three markets. The UK appears to be doing the most to translate widespread AI use into organizational performance gains, although the gap between adoption and impact remains wide.

### How AI gets absorbed

Through the institution. 65% of UK workers have read their employer's AI policy in full, compared with 57% in the US and 59% in Australia. 73% are confident in their organization's AI strategy, compared with 68% in the US and 65% in Australia. 69% say their employer has clear policies governing who can build or deploy AI agents, compared with 63% in the US and 66% in Australia. In the UK, AI adoption is reinforced by visible institutional infrastructure: clear policy, organizational strategy, and explicit operational boundaries.

### What workers are doing

Restraint, not just exposure. UK workers engage AI deeply, while also applying more active self-regulation than their US and Australian counterparts. 60% of UK AI users sometimes hold back from using AI to avoid becoming too dependent on it, compared with 51% in the US and 58% in Australia.

### The tradeoff

Governance hasn't erased the hidden labor. UK workers spend 38% of their AI time botsitting, compared with 36% in the US and 38% in Australia. And despite stronger institutional policies, 37% still report sometimes shipping AI-assisted work they have not fully checked, compared with 36% in the US and 44% in Australia. The UK's more developed governance environment has not meaningfully reduced the underlying botsitting burden.



## Australia

### The posture

High AI adoption, weaker payoff. 90% of Australian workers report using AI at work, on par with the UK and ahead of the US at 84%. But only 72% say it makes them more productive, compared with 75% in the US and 78% in the UK. In Australia, willingness to adopt AI is running ahead of the organizational systems needed to consistently translate usage into measurable performance gains.

### How AI gets absorbed

Adoption is outpacing infrastructure. Only 57% of Australian AI users say their employer provides enough training and support, compared with 54% in the US and 62% in the UK. And 58% report critical information isn't accessible within their AI tools, the highest of the three markets, compared with 53% in the US and 50% in the UK. In Australia, the willingness to use AI is strong, but the institutional support required to translate that use into reliable outcomes is lagging.

### What workers are doing

Using AI earlier and leaning on it more heavily. 54% of Australian AI users reach for AI before attempting to solve problems themselves, compared with 46% in the US and 49% in the UK. At the same time, 44% report sometimes delivering AI-assisted work they have not fully checked, compared with 36% in the US and 37% in the UK.

### The tradeoff

The worker absorbs the implementation gap. 43% of Australian workers feel worn out by AI tools at work, compared with 33% in the US and 30% in the UK. 43% use approved tools in noncompliant ways, compared with 36% in the US and 35% in the UK. When adoption outruns training, support, and integration, fatigue and corner-cutting become part of the operating cost.



## Appendix B: AI at work by function

Adoption also differs by function. IT is the most confident and structurally prepared for AI, but also the most likely to test beyond approved tools. Engineering uses AI deeply, while trusting it cautiously. Marketing is scaling output faster than impact. Support has some of the clearest AI use cases, but the lowest adoption and trust. HR is adopting AI at above-average levels, even as AI moves into higher-stakes people decisions.

*Percentages reflect the share of digital workers in each function (total n=6,000 across US, UK, and Australia) who selected the corresponding response. Function-level samples range from n=89 (Design/UX) to n=2,028 (IT). Comparisons in parentheses show the cross-function global average and a peer-anchor function (the top or bottom performer on that dimension). Differences of four percentage points or more are statistically significant at  $p < .05$ .*

### IT

#### The posture

Highest adoption, strongest gains. 97% of IT workers report using AI at work, compared with 87% for all respondents, and 85% say it makes them more productive, compared with 75% on average. In IT environments, where work is already digital, structured, and systems-based, AI fits well into existing workflows.

#### How AI gets absorbed

IT sees itself as the control tower. 77% say AI management sits within the IT function, compared with 59% on average, and 64% say the central tech team shapes employee tool choice and access, compared with 43% on average. Compared with other functions, IT is more likely to see itself as having direct control over enterprise AI decisions. That may be true, or it may be a visibility trap.

#### What workers are doing

Building the AI layer. IT workers report the highest confidence in their AI skills, with 85% describing themselves as highly confident, compared with 74% on average. 78% also say they spend at least some time building AI agents, compared with 63% on average.

#### The tradeoff

The frontier moves faster than the policy. 44% report using unapproved AI tools, compared with 38% on average, and 40% say they sometimes downplay AI's role when reporting to managers, compared with 33% on average. The function most involved in defining AI governance is also the one most likely to test new tools beyond what's sanctioned.



## Engineering

### The posture

Engineers use AI deeply, but trust it cautiously. 89% report using AI at work, compared with 87% on average, yet only 52% say they trust it with important tasks, compared with 58% on average. Productivity pressure is accelerating adoption, but repeated exposure to low-confidence workflows, like “vibe-coding,” appears to be widening the gap between usage intensity and confidence.

### How AI gets absorbed

AI adoption is driven less by central IT and more by engineering teams themselves. Just 38% of engineers say the central IT team drives their tool choices, compared with 43% on average. Central IT still defines guardrails, but tool choice is increasingly determined inside engineering teams, where proximity to the codebase and delivery pressure shape what actually gets used.

### What workers are doing

Botsitting is part of the workflow. Engineers spend much of their time botsitting, 39% compared with 37% on average. Even with trust lagging, AI is now part of how engineering work gets drafted, tested, debugged, and refined. Engineers say 26% of their work is automated with AI, compared with 27% on average.

### The tradeoff

Faster output can be harder to explain. 53% of engineers worry that relying on AI will erode their ability to do certain tasks, compared with 47% on average, and 43% say they sometimes deliver AI-assisted work they could not fully explain if asked, compared with 41% on average. In software engineering, where debugging, maintenance, and system security depend on deep technical fluency, gaps in understanding can quickly create higher downstream risk.



## Sales

### The posture

Middling adoption, below-average gains. 80% of sales workers use AI at work, compared with 87% on average, and only 62% say it makes them more productive, compared with 75% on average.

### How AI gets absorbed

Quota pressures influence use more than governance. 13% say AI use is unregulated or unmanaged in their organization, compared with 6% on average, the highest of any function. In sales, where compensation is tied to outcomes rather than process, reps adopt whatever helps them hit number, and the company often formalizes it afterward, if at all.

### What workers are doing

Sales is converting AI savings into customer time. 26% of sales workers use AI-saved time to spend more time with customers or external partners, compared with 19% on average, the highest of any function. AI reduces administrative drag, and reps are spending the reclaimed hours where revenue is closed.

### The tradeoff

Sales is using AI faster than it is measuring the results. 14% say their organization is not systematically measuring AI impact, compared with 8% on average. Sales leaders apply rigorous measurement to much of their daily work – quota attainment, pipeline coverage, and ramp time – but AI lacks the same discipline.



## Marketing

### The posture

AI is scaling output faster than impact. 94% of marketers use AI at work, compared with 87% on average, but only 70% say it makes them more productive, compared with 75% on average. AI can help marketers make more content, faster. But it can't always tell what's worth saying.

### How AI gets absorbed

Marketing adoption happens close to the work. 42% say department or team leaders oversee AI tools and practices, compared with 39% on average, while only 47% say AI management sits with IT, compared with 59% on average. AI has to understand the audience, brand, campaign goals, timing, and creative judgment behind the work. So the norms that shape AI use are often set by the marketing team, not by centralized governance alone.

### What workers are doing

AI tool sprawl is pervasive. Marketers use 3.1 AI tools regularly, compared with 2.6 on average, the highest of any function. Marketing work is fragmented by nature, across formats, channels, audiences, and moments in the customer journey, so one AI tool rarely covers the full job.

### The tradeoff

More output, less ownership. 43% say the more they use AI, the less ownership they feel over their work, compared with 42% on average, and only 58% feel proud of the work they produce with AI, compared with 68% on average. Marketing has always been a craft function: copywriters, designers, and brand leads define themselves by the work they make. With heavier AI use, the relationship between the marketer and ownership over the output gets thinner.



## Support

### The posture

Support is lagging in realizing AI value. 68% of support workers use AI at work, compared with 87% on average, the lowest of any function. 64% say it makes them more productive, compared with 75% on average. Much of the work is well suited to AI – summarizing issues, drafting replies, routing tickets, and surfacing answers – but the tools, training, and workflow integration around it haven't caught up.

### How AI gets absorbed

Support is getting AI without enough voice in how it works. Only 65% say the AI tools their organization provides are easy to use, compared with 74% on average. Only 54% say their organization has asked for feedback on those tools, compared with 66% on average. This points to a function being pressured to deploy AI rather than invited to help shape it.

### What workers are doing

Support is least likely to trust AI. Only 39% of support workers trust AI with important tasks, compared with 58% on average, the lowest of any function, even as AI is pushed into frontline, customer-facing work.

### The tradeoff

Impact is less visible to the people closest to it. 14% of support workers aren't sure what their organization is measuring about AI, compared with 6% on average, and 10% say their organization isn't systematically measuring it at all, compared with 8% on average. Without quality metrics, support teams can't tell whether AI is improving the customer experience or just moving tickets faster.



## HR

### The posture

HR is adopting AI at above-average levels and results. 90% of HR workers use AI at work, compared with 87% on average, and 78% say it makes them more productive, compared with 75% on average. HR has many AI-ready workflows – drafting job descriptions, summarizing feedback, preparing interview materials, answering policy questions, analyzing engagement data, and personalizing employee communications – where AI can reduce administrative load.

### How AI gets absorbed

HR sits close to AI governance, which can make the picture look more complete than it is. 72% say their organization has explained the rationale behind its AI policy, compared with 68% on average. 65% say their organization updates its AI policy at least quarterly, compared with 66% on average. HR often drafts these policies and helps enforce them, so its higher confidence may reflect proximity to the work more than the actual state of governance across the organization.

### What workers are doing

HR uses AI where the stakes are lower. 88% use AI to write content, and 62% use AI for coordination and administrative tasks, compared with 75% on average. HR is first pointing AI at high-volume tasks where the cost of an AI error is relatively contained.

### The tradeoff

But AI is moving into higher-stakes people decisions before much of the organization realizes it. HR workers are more likely than the average employee to report that AI is already shaping consequential people decisions. For example, 32% say AI is involved in hiring decisions, compared with 29% on average.



## Product management

### The posture

AI is helping PMs produce artifacts faster, but not necessarily make better product calls. 92% of PMs use AI at work, compared with 87% on average, but only 70% report productivity gains, compared with 75% on average. AI accelerates the artifacts of product work – PRDs, briefs, summaries, and stakeholder updates – but without organizational context, it struggles with the judgment-intensive parts of the role: prioritization, tradeoffs, alignment, and customer insight.

### How AI gets absorbed

PM adoption follows how product work already happens. 41% say a dedicated AI team or committee oversees AI tools and practices, compared with 34% on average, and 34% say manager expectations shape AI tool choices, compared with 26% on average. PMs work through roadmaps, reviews, cross-functional meetings, and manager priorities, so AI use spreads through those same channels, not just through central IT.

### What workers are doing

PMs use AI across more of their job than any function. 85% use AI across 3 or more distinct categories of work: research, writing, analysis, planning, and stakeholder communications. The PM job is multi-functional, and AI is becoming the connective tissue across it.

### The tradeoff

A loss of ownership. 46% say the more they use AI, the less ownership they feel over their work, compared with 42% on average. 51% worry AI reliance will make them forget how to do certain tasks, compared with 47% on average. AI can write the PRD. But without deep organizational context, it can't decide which customer pain matters most, which tradeoff is worth making, which dependency will slow the launch, or why this roadmap matters now.



## Design/UX

### The posture

Highest adoption, lowest productivity gain. 97% of Design/UX workers use AI at work, compared with 87% on average, the highest adoption of any function. But only 61% say it makes them more productive, compared with 75% on average, the lowest. AI generates options on demand, but design value depends heavily on context: choosing the right option for this user, this product, and this moment.

### How AI gets absorbed

Designers choose the tools with or without formal organizational backing. 51% say their own preferences drive which AI tools they use, compared with 41% on average, the highest of any function. Only 56% say AI management sits with IT, compared with 59% on average. Designers reach for the tool that fits the problem in front of their team, even when formal procurement hasn't caught up.

### What workers are doing

Designers feel ownership erode most when using AI. 55% say the more they use AI, the less ownership they feel over their work, compared with 42% on average, the highest of any function. For creators, AI can help generate ideas, but ownership still comes from shaping, refining, and finishing the final version.

### The tradeoff

AI can loosen authorship and outrun governance. 42% use unapproved AI tools at work, compared with 38% on average. Designers reach past the sanctioned stack for tools that fit the craft, and policy lags the behavior.



## Appendix C: AI at work by industry

AI is landing differently across industries. Adoption ranges from 63% in government to 97% in technology, but usage is only one part of the story. Technology workers report the strongest productivity gains and the highest adoption. Healthcare workers trust AI less because the stakes are higher. Government workers use AI the least. Media and entertainment workers feel the sharpest loss of ownership. Across industries, the pattern is clear: AI creates the most value where the work is already digital, measurable, and well supported, and the most strain where trust, judgment, and human context matter most.

*Percentages reflect the share of digital workers in each industry (total n=6,000 across US, UK, and Australia) who selected the corresponding response. Industry-level samples range from n=312 (Construction) to n=1,104 (Technology). Comparisons in parentheses show the cross-industry global average and a peer-anchor industry (the top or bottom performer on that dimension). Differences of four percentage points or more are statistically significant at  $p < .05$ .*

### Technology

#### The posture

Highest adoption, strongest gains. 97% of technology workers use AI at work, compared with 87% on average. 84% say it makes them more productive, compared with 75% on average, and 84% say it improves work quality, compared with 74% on average. Tech has a structural advantage for AI: much of the work is already digital, software-native, and organized around code, systems, documentation, and toolchains.

#### How AI gets absorbed

Technology workers learn AI by doing. 69% say learning directly from AI tools has been valuable, compared with 56% on average. 48% say unofficial or unapproved tools have been valuable for learning, compared with 41% on average. Tech workers already operate in fast feedback loops, where new tools are tested on real problems and judged by whether they work.

#### What workers are doing

Technology workers use AI to build skill, not just save time. 54% use AI-saved time to improve work quality, compared with 47% on average. 48% use it to learn or build new skills, compared with 38% on average. For tech workers, AI is not just a shortcut. It helps them learn faster, improve the work, and carry stronger skills into the next project.

#### The tradeoff

Fluency can outpace accountability. 47% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 32% have blamed AI for a mistake that was actually their fault, compared with 28% on average, the highest of any industry. In a sector where AI fluency is becoming professional currency, speed can move faster than ownership.



## Financial services

### The posture

High adoption, high anxiety. 90% of FinServ workers use AI at work, compared with 87% on average. 77% say it makes them more productive. But 45% worry AI could eliminate their role, compared with 42% on average, among the highest in the sample. FinServ is well-suited to AI because so much of the work is standardized, document-heavy, and organized around speed, accuracy, and compliance, and workers know it.

### How AI gets absorbed

AI enters through structured governance channels. 65% of FinServ workers say they have read their organization's AI policy, compared with 60% on average. 54% say a dedicated AI team or committee oversees AI tools and practices, compared with 31% on average, the highest of any industry. In a sector shaped by compliance, audit trails, data controls, and customer trust, new technology has to clear a higher bar before it becomes part of the work.

### What workers are doing

FinServ workers use AI to reduce risk, not expand scope. 47% use AI-saved time to improve work quality, compared with 47% on average. Only 27% use it to take on new responsibilities, compared with 32% on average, the lowest of any industry. That reflects a heavily regulated, error-averse environment.

### The tradeoff

Faster output still has to be explainable. 46% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 34% have blamed AI for a mistake that was actually their fault, compared with 28% on average. FinServ workers may be using AI to improve quality and reduce risk, but in a sector built on precision, compliance, and audit trails, speed cannot come at the expense of ownership.



## Healthcare

### The posture

Strong adoption, low trust, and little room for error. 80% of healthcare workers use AI at work, compared with 87% on average. 73% say it makes them more productive. But only 49% trust AI to handle important work tasks, compared with 58% on average, one of the lowest rates of any industry. In a sector where a mistake can harm a patient or end a career, AI has to earn every inch of clinical trust, and it hasn't yet.

### How AI gets absorbed

AI enters through approved use, not shadow experimentation. Only 34% use unapproved AI tools, compared with 38% on average. 31% sometimes use approved tools in noncompliant ways, compared with 38% on average, the lowest rate of any industry. This reflects the stakes: patient safety, privacy, and professional accountability leave little room for unsanctioned experimentation.

### What workers are doing

Healthcare keeps AI in a supporting, not generative, role. Only 52% of healthcare workers report spending any time building AI agents, compared with 63% on average. When they save time, 39% put it toward improving work quality, compared with 47% on average. In healthcare, AI is more likely to support the work around care than to take over the work itself.

### The tradeoff

Caution lowers risk, but limits how far AI goes. Only 32% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 32% sometimes deliver work they don't fully check, compared with 39% on average. Healthcare workers catch more AI-assisted errors before they spread. That makes the work safer, but it also eats into the time AI is supposed to save.



## Manufacturing

### The posture

Strong AI adoption. 89% of manufacturing workers use AI at work, compared to 87% on average. 77% say it makes them more productive (vs. 75% avg), and 77% say it improves work quality (vs. 74% avg). Manufacturing benefits because it already knows how to evaluate gains in cost, quality, throughput, and defects. AI enters a culture built to measure it.

### How AI gets absorbed

AI enters through governed operating systems. 72% say their organization controls which AI tools they can access, compared with 70% on average. 69% say their organization has explained the rationale behind its AI policy, compared with 68% on average. This reflects how manufacturing already manages change: through controlled tools, clear processes, and measurable standards for quality, throughput, and defects.

### What workers are doing

Manufacturing workers use AI to improve process quality, not just output. 51% use AI-saved time to improve work quality, compared with 47% on average. That matters in a sector where small process improvements can reduce defects, cut rework, improve handoffs, and keep production moving.

### The tradeoff

Measurement discipline doesn't eliminate overreliance. 39% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 30% have blamed AI for a mistake that was actually their fault, compared with 28% on average. Manufacturing already evaluates work through familiar signals like quality, throughput, defects, and rework. That makes AI easier to measure, but not always easier for workers to explain.



## Retail

### The posture

Retail uses AI widely, but not yet deeply. 87% of retail workers use AI at work, roughly the global average. But only 27% say AI is integrated across workflows in their organization, compared with 29% on average. Retail has rolled AI into specific workflows like pricing, fulfillment, customer service, and inventory, but it hasn't yet connected those use cases across the business.

### How AI gets absorbed

Retail runs AI through the business line, not central IT. 47% say department or team leaders oversee AI tools and practices, compared with 43% on average. Only 48% say oversight sits with IT, compared with 59% on average. 70% say their organization has asked for feedback on AI tools, compared with 66% on average. AI spreads through merchandising, service, and planning teams refining tools around live operating needs.

### What workers are doing

Retail spreads AI savings across customer, team, and quality work. When retail workers save time with AI, 21% of the saved time goes toward customers or partners, compared with 19% on average. 21% goes toward collaborating with peers, compared with 19% on average, and 43% goes toward work quality, compared with 47% on average. Retail AI is helping around the edges of the customer experience, but not yet transforming the whole operating model.

### The tradeoff

Speed pressure still invites shortcuts. 43% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 40% use unapproved AI tools at work, compared with 38% on average. In a margin-pressured environment, speed can overtake responsible adoption.



## Education

### The posture

Education workers are adopting AI before their institutions do. 83% use AI at work, but governance is thin. Only 52% say their organization controls which tools they can access, compared with 70% on average. Teachers, professors, and staff are adopting AI ahead of their institutions' procurement and governance.

### How AI gets absorbed

AI enters through ad hoc experiments and initiatives. Education workers are much more likely to say AI adoption in their organization is ad hoc, 20% compared with 10% on average, or in an experimental phase, 25% compared with 15% on average. AI is being tested by teachers and staff before schools and universities have built an operating model around it.

### What workers are doing

Education workers use AI more to support their own practice than to transform the institution. 36% use AI-saved time to improve work quality, compared with 47% on average. 35% use it to learn or build new skills, compared with 38% on average. AI is helping individual educators work a little better, but it has not yet changed how institutions teach, support students, or reduce administrative burden at scale.

### The tradeoff

Measurement hasn't caught up to use. Only 28% say their organization measures productivity or output gains from AI, compared with 42% on average. 16% say their organization isn't systematically measuring AI impact, compared with 8% on average. Schools may know AI is being used, but not whether it's improving instruction, reducing administrative burden, or strengthening student support.



## Government / public sector

### The posture

Lowest adoption, weakest gains. Only 63% of government workers use AI at work, compared with 87% on average, the lowest of any industry and a 24-point gap. 58% say AI makes them more productive, compared with 75% on average. 57% say it improves quality, compared with 74% on average. Government isn't short on document-heavy work. Too often, it's constrained by the pace at which access, security, and procurement can move.

### How AI gets absorbed

AI enters through bounded use cases, not open experimentation. Only 36% say learning directly from AI tools has been valuable, compared with 56% on average. Only 25% say unofficial tools have been valuable for learning, compared with 41% on average. AI access, data use, procurement, and security rules narrow the space for experimentation.

### What workers are doing

Government workers do far less hands-on work with AI than the average worker. Only 32% spend any time building AI agents, compared with 63% on average. Only 41% spend any time debugging or refining AI tools, compared with 72% on average. The sector uses AI, but does relatively little to extend, shape, or stress-test it.

### The tradeoff

Low disruption also means weak learning loops. 16% say their organization isn't systematically measuring AI impact, compared with 8% on average. 21% aren't sure what their organization measures, compared with 6% on average. Low adoption limits some downside, but it also leaves agencies with less evidence on what works.



## Media & entertainment

### The posture

Adoption is near average; confidence is not. 85% of media workers use AI at work, compared with 87% on average. But only 47% trust AI to handle important work tasks, compared with 58% on average, among the lowest of any industry. In media, AI touches the core craft: writing, editing, and visual development. Even useful adoption can feel like erosion.

### How AI gets absorbed

AI is pulled in by deadline pressure, not strong enablement. Only 62% say the AI tools their organization provides are easy to use, compared with 74% on average. Only 48% say learning directly from AI tools has been valuable, compared with 56% on average. AI is being absorbed by the pressure to produce faster drafts, edits, and concepts, even where the formal implementation model is weak.

### What workers are doing

Saved time gets reabsorbed, not reinvested. Media workers are below average in putting AI-saved time toward strategic or higher-value work, 26% compared with 36% on average; learning new skills, 28% compared with 38% on average; or reducing overtime, 22% compared with 32% on average. The clearest above-average use is collaboration: 24% use saved time to build relationships or collaborate with peers, compared with 19% on average. AI is helping media workers produce and coordinate faster, but it is not yet creating much room for deeper work.

### The tradeoff

Weaker ownership over the work. 52% say the more they use AI, the less ownership they feel over their work, compared with 42% on average, the highest of any industry. Only 55% feel proud of the work they produce with AI, compared with 68% on average. In a sector built around voice, taste, and authorship, the identity cost shows up quickly.



## Construction

### The posture

High adoption, strong quality gains. 91% of construction workers use AI at work. 79% say it makes them more productive, and 80% say it improves work quality. In construction, AI's clearest use cases sit around the build: planning, reporting, documentation, and coordination.

### How AI gets absorbed

AI reaches construction through tools that are easy to pick up. 78% say their organization's AI tools are easy to use, compared with 74% on average. 71% say their organization has asked for feedback on those tools, compared with 66% on average. In construction, AI has to work for people managing real projects under real constraints, so the tools that stick are the ones that don't demand deep technical onboarding.

### What workers are doing

Construction workers use AI to manage coordination load, not just individual output. 24% use AI-saved time to build relationships or collaborate with peers, compared with 19% on average. 23% use it to spend more time with customers or external partners, compared with 19% on average. The gain comes from helping people manage the dense web of trades, subcontractors, architects, owners, clients, schedules, and approvals.

### The tradeoff

Coordination speed can outrun review. 42% sometimes deliver AI-assisted work they don't fully check, compared with 39% on average. Once AI-generated content enters RFIs, submittals, schedules, or change-order drafts, it can move quickly through the project before someone fully validates it.



## Energy & utilities

### The posture

Solid adoption, but only average trust where the stakes are highest. 88% of energy and utilities workers use AI at work, and 76% say it makes them more productive. But only 55% trust AI to handle important work tasks, compared with 58% on average. In a sector where an AI error can trip a turbine, shut down a substation, or violate an environmental regulation, the reliability standards that govern the work have not yet been extended to the AI inside it.

### How AI gets absorbed

AI enters through a formal channel, with informal workarounds on the side. 61% say AI management sits with IT, compared with 59% on average. 49% say unofficial or unapproved tools have been valuable for learning, compared with 41% on average. AI is formally governed through technology channels, but workers also learn at the edges when official tools or training don't fully meet the work.

### What workers are doing

Energy workers use AI to build capability around the edge of the workflow, not inside the operating core. 41% use AI-saved time to learn or build new skills, compared with 38% on average. 23% use it to collaborate more effectively with peers, compared with 19% on average. The technology helps most in the planning and field-support layer around assets, not in the tightly controlled operating core.

### The tradeoff

Formal rollout doesn't eliminate shadow behavior. 44% use unapproved AI tools at work, compared with 38% on average. 34% have blamed AI for a mistake that was actually their fault, compared with 28% on average. Even in a reliability-focused sector, workers will route around incomplete implementation if the tool helps the work move.



## Transportation & logistics

### The posture

Useful in day-to-day operations, but not producing broad gains. 83% of transportation and logistics workers use AI at work. But only 66% say it makes them more productive, compared with 75% on average. The sector has obvious use cases – dispatch, routing, planning, and warehousing – but many real-time edge cases still resist automation.

### How AI gets absorbed

AI enters through local operations, with weak central enablement. 49% say department or team leaders oversee AI tools, compared with 43% on average. 9% say AI use is effectively unregulated or unmanaged, compared with 6% on average. Only 46% say learning directly from AI tools has been valuable, compared with 56% on average. AI spreads through dispatch and planning needs faster than through a mature implementation model.

### What workers are doing

Transportation workers use AI more than they extend it. Only 52% spend any time building AI agents, compared with 63% overall. Only 60% spend any time debugging or refining AI tools, compared with 72% overall. That may reflect a practical constraint: transportation workers may benefit from AI in the moment, but have fewer opportunities to pause the work, experiment with tools, and rebuild workflows while the operation is still moving.

### The tradeoff

Anxiety is ahead of the proof. 45% worry AI could eliminate their role, compared with 42% on average. But only 34% say their organization measures productivity or output gains from AI, compared with 42% on average. That leaves workers seeing AI show up in dispatch, routing, planning, and warehousing before the organization can show whether it's actually improving the work.



## Professional services

### The posture

Adoption is broad; confidence in the payoff is modest. 87% of professional services workers, including consulting, legal, and accounting, use AI at work. 65% say it makes them more productive, but only 42% trust AI to handle important work tasks, compared with 58% on average, among the lowest of any industry. AI can help draft the work, but the person sending it still has to explain the logic, defend the recommendation, and own the risk.

### How AI gets absorbed

AI enters as experimental and ad hoc, before formal systems catch up. 20% describe AI use as experimental, compared with 15% on average. 17% describe it as ad hoc, compared with 10% on average. Only 48% say the AI policy is regularly reviewed, compared with 60% on average. AI shows up first in research, summaries, and drafts, while the rules for client-facing work take longer to catch up.

### What workers are doing

Professional services workers use AI to improve the work, not reinvent it. 47% use AI-saved time to improve work quality, compared with 47% on average. Only 43% spend any time building AI agents, compared with 63% on average. AI helps with research, summaries, and drafts, but client-facing work still has to be checked, defended, and owned.

### The tradeoff

Caution lowers risk, but limits how far AI goes. Only 28% sometimes deliver AI-assisted work they couldn't fully explain, compared with 41% on average. 28% sometimes deliver work they don't fully check, compared with 39% on average. That caution helps keep bad work from reaching clients, but it also keeps AI stuck in research, summaries, and drafts instead of changing how the work gets done.



## Hospitality & travel

### The posture

Hospitality reports one of the weakest AI productivity cases of any industry. 87% of hospitality and travel workers use AI at work, but only 65% say it makes them more productive, compared with 75% on average. The work carries a layer of live human judgment – reading a guest’s mood, solving a booking problem in real time, pacing a dining experience – that AI can assist around but rarely replaces.

### How AI gets absorbed

AI is operational, not integrated. 44% describe AI use as operational, compared with 33% on average. Only 18% describe it as integrated across workflows, compared with 29% on average, while only 55% say the rationale behind their organization’s AI policy has been explained, compared with 68% on average. That makes sense in a sector where AI can help with repeatable tasks like bookings, staffing, and customer service, but the core work still depends on live judgment, timing, and service recovery.

### What workers are doing

Hospitality workers use AI near the customer, but the lift is modest. 19% of hospitality workers use AI-saved time to spend more time with customers or external partners. The sector’s product is the interaction itself, so, in the best cases, AI handling the bookings, confirmations, and logistics lets staff stay in the guest-facing moment longer.

### The tradeoff

Anxiety runs ahead of proof. 46% of hospitality workers worry AI could eliminate their role, compared with 42% on average. Only 36% say their organization systematically measures productivity gains, compared with 42% on average, a 6-point gap. Workers are worried AI could replace them, but many aren’t seeing clear evidence of what it’s actually improving.



## Nonprofit / NGO

### The posture

Thinly resourced, lightly rewarded. Only 65% of nonprofit and NGO workers use AI at work, compared with 87% on average, among the lowest in the sample. Only 58% say AI makes them more productive, compared with 75% on average, and 60% say it improves work quality, compared with 74% on average. The issue isn't resistance. It's that many nonprofits lack the budget, technical support, and training time to make AI useful beyond one-off experimentation.

### How AI gets absorbed

AI reaches the nonprofit sector with the weakest support infrastructure of any industry. Only 36% say the AI tools their organization provides are easy to use, compared with 74% on average, the lowest of any industry. Only 32% say their organization has explained the rationale behind its AI policy, compared with 68% on average. AI gets absorbed where staff can make it useful on their own, not because the institution has built an enablement model around it.

### What workers are doing

Hands-on experimentation is limited. Only 12% spend any time building AI agents, compared with 63% on average, the lowest of any industry. Only 30% spend any time debugging or refining AI tools, compared with 72% on average. The hands-on layer of adoption that exists almost everywhere else is much thinner here.

### The tradeoff

Limited investment means fewer mistakes – and fewer lessons. 24% say their organization isn't systematically measuring AI impact, compared with 8% on average. 22% aren't sure what their organization measures, compared with 6% on average. Using AI less may mean fewer obvious mistakes, but it also means fewer chances to learn where it saves time, improves services, or helps stretched teams do more.