The Use and Governance of AI Agents at U.S. firms

Innessa Colaiacovo
University of Oregon

Emma Wiles*†
Boston University

November 19, 2026

Most recent draft here.

Abstract

Autonomous AI agents have been forecast to change work and organizations, but evidence about firms' deployment of these tools is scarce. Drawing on a survey of 1,500 full-time senior managers, we provide some of the first systematic evidence about the diffusion of AI agents, defined as tools able to take actions and complete tasks autonomously. Just over half of all managers surveyed report that their organization has at least one AI agent built into its work processes, but fewer than half of managers who reported this are aware of a formal policy that governs the use of agents at their organization. The correlation between integrating AI agents into work and enacting policies to govern them is positive but small, suggesting that integrating and governing agents might be separate sets of choices made by organizations. These results suggest that while AI agents are a real phenomenon in US workplaces, organizations vary in how they are adjusting work and policies in response, which may affect agents' ultimate impacts.

1 Introduction

The potential for autonomous technological "agents" to reshape work, long anticipated by computer and social scientists (Simon, 1960), is a key distinguishing feature of contemporary advances in artificial intelligence (AI). Since 2024, the volume of commercial publications about agentic AI and the number of companies purporting to offer agentic AI solutions has increased dramatically (Staufer et al., 2026). Prominent tech executives have speculated publicly about the possibilities, with Nvidia's CEO heralding the age of agentic AI in 2025, a year when OpenAI's CEO predicted that the first AI agents would "join the workforce" and "materially change the output of companies" (Altman, 2025; Huang, 2025).

*emma.b.wiles@gmail.com, 595 Commonwealth Ave, Boston, MA 02215.

†This survey received IRB exemption from Boston University's IRB under protocol number 8497X. We would like to thank Julie Bedard, Megan Hsu, and Matthew Kropp for their conversations about AI agents in organizations.

However, it remains unclear whether AI agents are diffusing broadly and if so how organizations are adapting in response. Large-scale official administrative surveys of technology diffusion either do not mention AI agents specifically (McElheran et al., 2024) or bundle them into a broader category of AI technologies (Census Bureau, 2025). Understanding how workers are engaging with AI agents is important, as a large innovation and policy literature suggests that realizing economic gains from a novel technology depends on integrating it effectively into production along with existing workers and assets (Bresnahan et al., 1996; Bloom et al., 2012). In addition, given that AI agents may exercise autonomous judgment, integrating agents into production presents additional challenges of governance, monitoring, and accountability (Goos and Savona, 2024; Wiles et al., 2026). As the CEO of Microsoft counseled other firms in 2026: "You need to give them [agents] identities, you need to give them sandboxes, then you need to set policies to govern them" (Bitter, 2026). Understanding these dimensions of the current state of AI agents at work requires a custom survey of managers designed for that purpose, as proposed by Hunt et al. (2022).

In this paper, we provide some of the first systematic evidence about whether and where agents are deployed in firms, how work has been reorganized around them, and what governance structures organizations have enacted in response. Our data come from a novel survey of 1,500 full-time senior managers representative of U.S. private sector industries.

We first provide evidence that a substantial share of managers are encountering AI agents at work, with over half of managers reporting AI agents incorporated into their firms' workflows across a host of departments. Most managers report that AI agents are purchased or licensed from third-party providers or built on top of foundation models.

We then turn to examining firms that have incorporated AI agents into their workflows and show variance in how employees encounter agents as part of their work and ensuing changes to their roles. Where agents are built into work processes, in most cases managers report that human employees prompt them directly, but in other cases human employees are reviewing work that agents produce autonomously. Nearly 30% of respondents report that agents have not fully taken over tasks at their organization. Where agents have fully taken over tasks previously done by human employees, firms varied in whether they reassign these employees to new roles or tasks, make them responsible for the agent's oversight, or (in a minority of cases) eliminate their position.

We also show that agent-adopting organizations differ in whether they enact formal policies to govern AI agents, who takes accountability for agent errors, and the steps taken to formalize (and humanize) the AI agent and its role. Just under half of managers at agent-adopting firms report that their firm has a formal policy governing the AI agent's use. When AI agents make errors, a plurality of managers (41%) report that the individual employee using the agent is responsible for correcting it. Finally, roughly one quarter to one third of the sample of managers whose organizations use AI agents report that the agent has been given a name or identity, given a job title or

description, listed on a workflow chart, or assigned a human manager. Just under 20% report their organization refers to the agent as an employee or a "digital worker." Given this heterogeneity, as a final step we ask whether the organizations integrating agents into work most aggressively are the same ones enacting policies to govern them, and show that reported work changes and governance changes in response to AI agents are only modestly positively correlated in our sample. We then explore whether organizations cluster meaningfully based on managers' responses by using an unsupervised machine learning approach to identify patterns across all managers' survey responses. This latent class analysis (LCA) divided respondents into four groups which we interpret as a typology suggesting that while two groups of adopters adjust everything or nothing (institutionalized adopters and passive adopters), there are others who are relatively responsive at adjusting either work or policy (policy-led or action-led adopters).

This paper makes two contributions to existing work. First, we provide evidence that agentic AI technologies are a real phenomenon affecting organizations and workers, providing a motivation for future theoretical and empirical work. Second, this survey highlights that not all agentic AI "adoption" is equal. Organizations vary in whether their adoption of agents is accompanied by complementary changes to organizational governance practices and workflow restructuring. Conflating these different types of adopters may obscure important heterogeneity in the impacts of these technologies.

2 Data

Managers' perceptions can shed light on the diffusion and impact of emergent technologies (Hunt et al., 2022) and a growing line of work has designed and deployed custom surveys to track the phenomenon of AI diffusion (Bick et al., 2024), specifically asking managers and executives for their perceptions of AI's impacts on their firms (Karger et al., 2026; Yotzov et al., 2026; Colaiacovo and Koning, 2026). This work is a complement to experimental or quasi-experimental work that seeks to precisely identify the causal impact of deploying AI agents on various outcomes. Participating managers in our survey were asked to respond based on what they knew about their company, and their knowledge may be imperfect. Throughout this analysis, we interpret survey results as suggestive of the ground truth of the deployment of AI agents at firms, particularly with regards to managers' perceptions of overall workforce changes. Perceptions matter even when imperfect because they influence how managers invest their time and attention responding to AI.

The survey described in this paper was conducted in 2026 in partnership with YouGov, which recruited 1,500 U.S. managers and senior managers weighted to match the population of full-time senior managers across private sector industries. A major advantage of partnering with a panel provider is enhanced verification that participants in the survey are real, human respondents. In

contrast to some other online platforms, YouGov relies on a proprietary panel of survey respondents and verifies a respondent's identity in multiple ways, including checking for duplicate or fraudulent accounts, geolocating respondents, and fingerprinting devices. They also automatically check for abnormally fast or inconsistent responses, straight-lining, or copy-paste behavior that would suggest a bot may be taking the survey. Respondents who fail these checks are excluded from the sample.

Respondents were given a definition to clarify the term "AI agent." In the survey, respondents were instructed that these tools "can take actions and complete tasks on their own, such as sending a message, updating a record, or approving a routine request." This definition was chosen to emphasize the autonomous nature of the technology: AI agents as captured in our survey may still be performing a wide range of tasks of varying complexity, reflecting the emerging nature of the technology.

To distinguish between employee-driven use of AI agents and purposeful organizational adoption, respondents were instructed that AI agents can be used in two different ways (the verbatim survey text is reproduced below):

- By individual employees, for their own work. For example, an employee might choose to use an AI assistant that automatically joins their meetings, takes notes, and emails them a summary afterward. Another employee might use Claude Code to automate writing code to clean a dataset.
- By the company, as part of its systems or processes. For example, a company might set up an AI agent that automatically reviews incoming job applications, processes invoices, or routes documents to the right person for approval. The company set this up, and it runs as part of how work gets done — not something one employee is using on their own.

Below, Figure 1 includes both types of agent use while later results focus on managers who reported that agents are part of the firm's systems and processes. All estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

3 Results

3.1 Adoption of AI Agents

A majority of senior managers report that their organizations are using AI agents, and this use extends beyond strictly technical departments. Just over half of managers reported that agents are being built into work processes at their companies (Figure 1). While the most common functional area where respondents report AI use is predictably IT/Engineering/Software Development (26% of respondents), agents are also being used across organizational functions ranging from customer service to human resources to marketing and operations.[1] Sixty-nine percent of all respondents estimate that their organization will either "somewhat" or "significantly" expand their use of AI agents in business processes over the next twelve months (Figure A4).

Companies are obtaining their AI agents by both building them in-house and purchasing or licensing them from third-party vendors. Sixty percent are doing at least some agent development by building on top of a foundation model via API or deploying agents built in-house by a technical team. However, it is still common that these agents are obtained from a third-party (34% of respondents report this is the primary way their firm obtains agentic tools).

Most, but not all, managers perceived this use of AI agents as a deliberate decision by senior leadership. Among managers who report AI agents built into their work processes, a majority reported that it was a deliberate decision by senior leadership, but over a third (37%) reported that the use of AI agents evolved gradually, without a formal decision (Figure A3). This underscores that existing technology providers may play a role in diffusing new technologies to their client firms by embedding agentic tools into their products.

These use rates are high in comparison to official measures of firm-level AI adoption. The Census Bureau's Business Trends and Outlook Survey (BTOS) estimates 32% of businesses on an employment-weighted basis are using AI products in their work processes (37% among firms with more than 250 employees) (Grundy et al., 2026). Our survey is at the level of the manager, weighting the sample toward larger employers that are more likely to be deploying AI. Our estimate is also higher because, as discussed above, the term "agent" likely captures tools of varying sophistication. The two survey approaches are also asking for different perspectives on the same phenomenon. The BTOS asks the firm to report AI used in business functions, while we asked managers about tools employees encounter in their workflows. Because agentic tools are often embedded in third-party products, employees may encounter AI tools that do not register as formal

[1] With a sample size of 1500 we are cautious about statements about agentic AI use across industry and firm size. Figure A1 and Table A1 show weighted shares and sample size for agentic AI use across industries. In line with work showing that larger firms are more likely to adopt other AI technologies (McElheran et al., 2024), we also find that managers at larger firms are more likely to report AI agents integrated into their work processes (Figure A2).

AI adoption at the firm level.

3.2 Variation Among Agent-Adopting Organizations

3.2.1 Integration of agents into work

We then asked managers to report how AI agents were changing work both at an individual level (employees' typical agent interactions) and at the level of the workforce (task takeover or worker displacement by agents) and found variance in their perceptions of whether and how AI agents have changed work.

Where agents are built into work processes, human employees are sometimes prompting them directly and in other cases monitoring and approving the agents' output. Forty-one percent of respondents report that typical use involves employees prompting agents directly, while 15% say it operates in shared channels like Slack and 10% reported that the agent works autonomously with employees reviewing and approving its work. Roughly a third of respondents reported typical use cases were a mix of these options (Figure 2). This is consistent with the amount of authority that managers report agents have within workflows: only 11% of respondents report that agents make some autonomous decisions (Figure 2, Panel b). This suggests that employees are gaining direct exposure through both directly querying agents and in a smaller share of cases monitoring work produced by agents acting autonomously, suggesting adopting AI agents into production has different implications for employees' daily experiences at work.

It is not common for agents to have fully taken over human workers' tasks, and where this does happen it is more common for workers' jobs to change rather than be eliminated. Twenty-nine percent of managers at firms where agents are embedded in workflows reported that no tasks were fully taken over by the agents at their firms (Figure 3). Thirty-three percent of firms report employees being assigned to new positions, given new tasks in the same role (32%), or taking on oversight of the AI agents (31%) following agentic AI deployment. Fifteen percent of respondents report positions being eliminated. It is unclear whether these changes are transient or permanent, but this suggests that some firms are changing employees' individual roles given the integration of agents.

Managers also report varied perceptions of overall workforce changes at their firm since the deployment of AI agents. Notably, equal shares of managers report that employees have been reassigned to different roles as report "no significant workforce changes" have taken place (36%). Equal shares (27%) also report changes in job descriptions and hiring employees to oversee AI agents. While this question may reflect workforce changes related to factors other than AI agents, the variance in managers' perceptions is striking and suggests that agent-adopting firms adjust their personnel in different ways along different margins.

Figure 1: Prevalence of adoption of AI agents in organizations

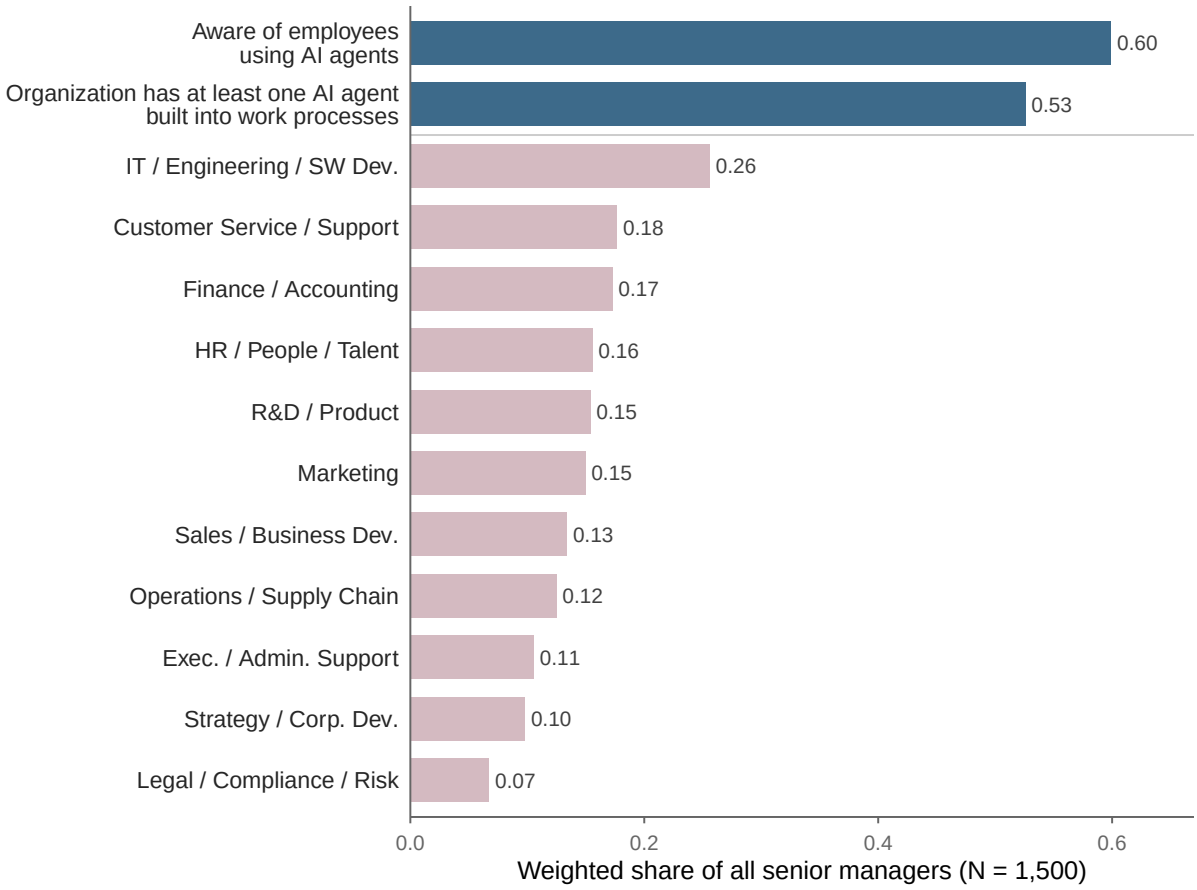

Notes: This plot shows the share of all respondents who responded yes to questions about the prevalence of AI agents in their organization. The first question asked "Are you aware of employees at your company choosing to use AI agents on their own to automate their personal work tasks? (e.g., an employee using an AI assistant that automatically joins their meetings, takes notes, and emails them a summary afterward, or using Claude Code to automate writing code to clean a dataset— not integrated into formal company workflows.)" The second asked "Are you aware of your company having officially integrated AI agents into its systems or workflows to complete tasks automatically? (e.g., use of an AI agent to automatically screen job applications, processing invoices, or route documents for approval)". The second panel asks respondents if they have AI agents formally integrated into the workflow of each department. N = 1,500.

Figure 2: AI Agent Employee Interaction and Authority

(a) Which best describes the typical way employees interact with AI agents at your company?

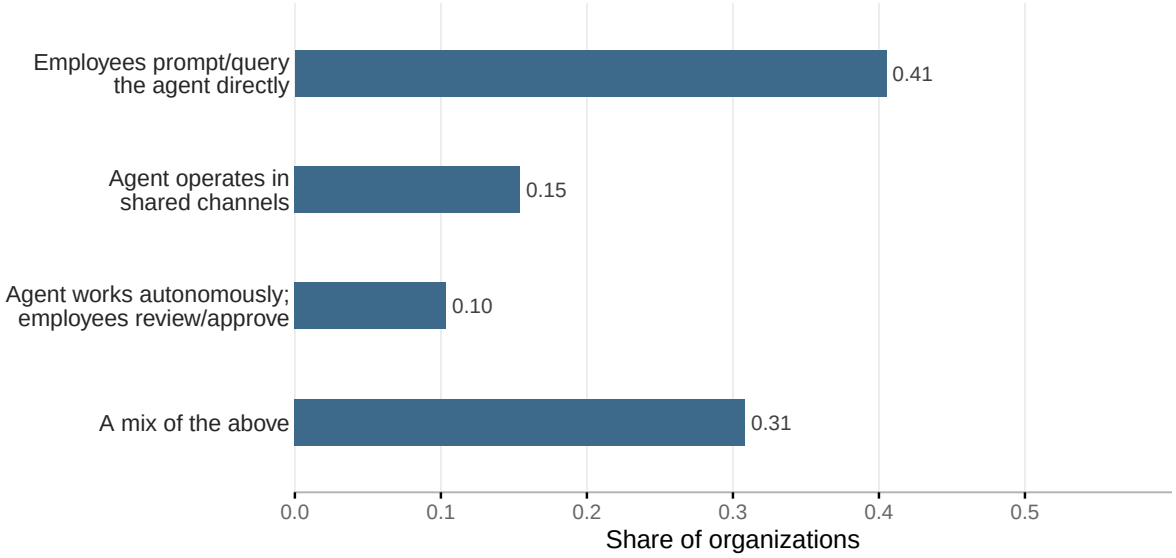

(b) Which best describes the level of authority [AI agents used in business processes] typically have within workflows?

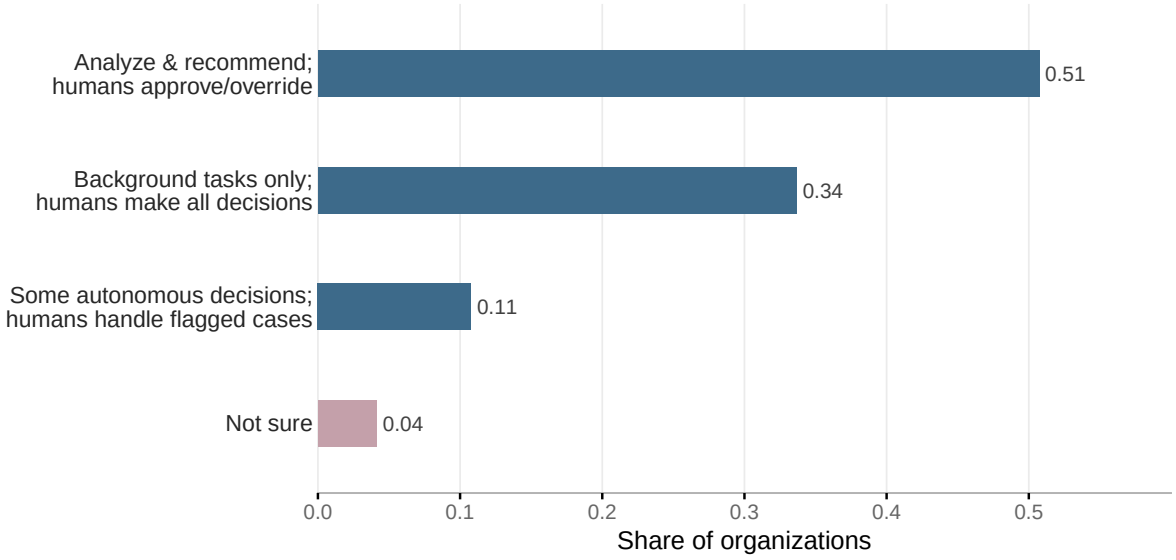

Notes: This figure shows, conditional on managers reporting that their organization has AI agents built into their work processes, how employees interact with the agents (N = 764). Panel (a) reports the answer to the question "Which best describes the typical way employees interact with AI agents at your company? (Select one)?" Panel (b) reports the answer to "Which best describes the level of authority they [AI agents used in business processes] typically have within workflows?" Estimates are weighted using YouGov respondent weights (see Section 2 for details).

Figure 3: Reorganization of workforce

(a) To the best of your knowledge, what happened to employees whose tasks were taken over by AI agents?

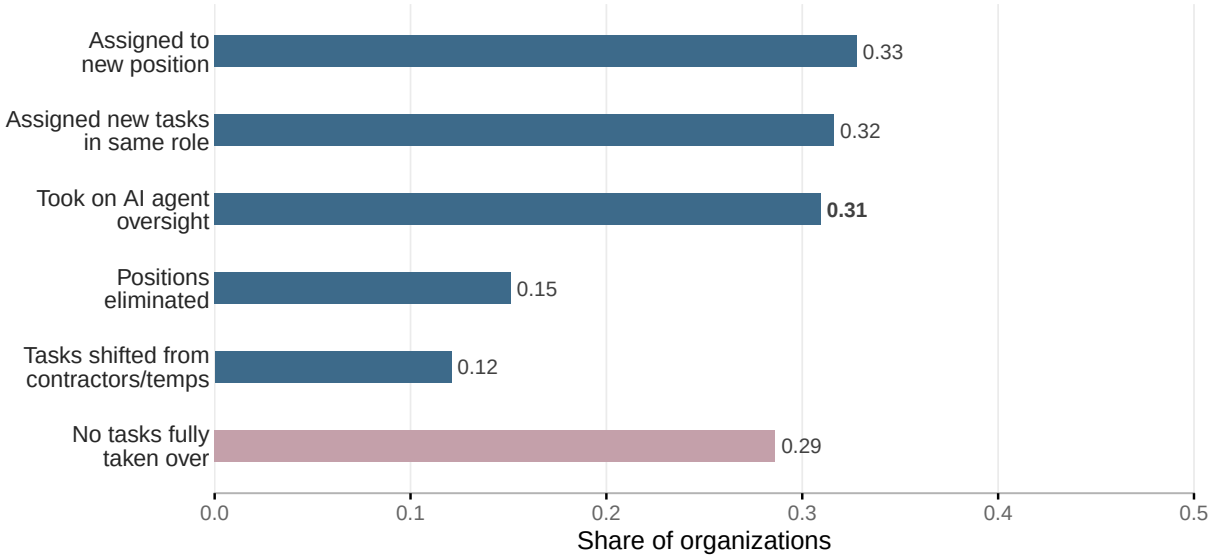

(b) Since deploying AI agents in business processes, has your company...?

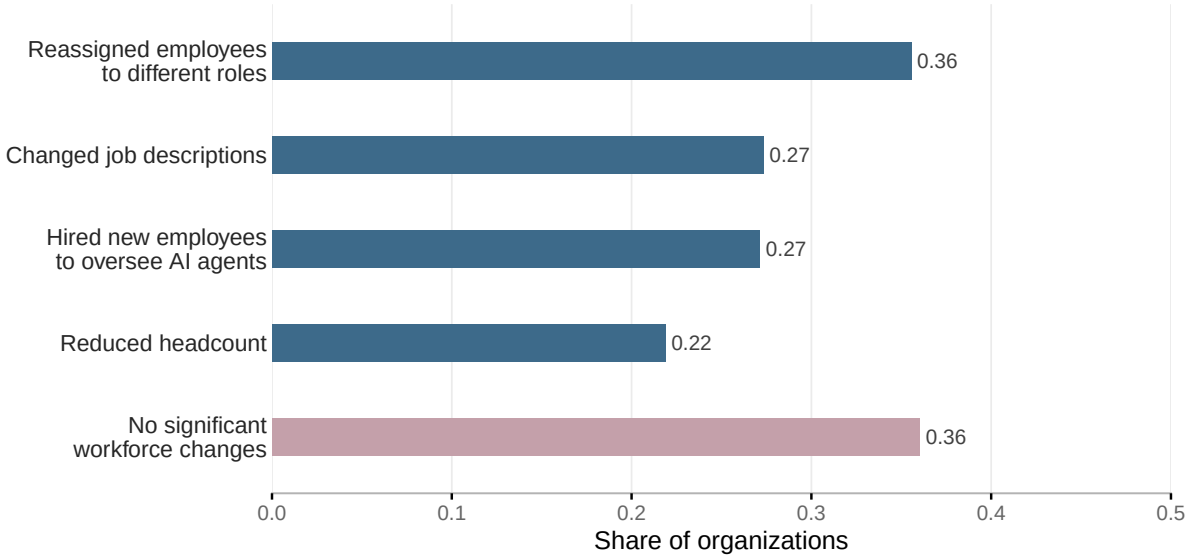

Notes: Panel a) reports the answers to the question "To the best of your knowledge, what happened to employees whose tasks were taken over by AI agents? Select all that apply." Panel b) reports the answers to the question "Since deploying AI agents in business processes, are you aware of your company making any of the following changes? Select all that apply." These questions were asked to all respondents who indicated their organization has adopted AI agents into their work processes, N = 764. Estimates are weighted using YouGov respondent weights (see Section 2 for details).

3.2.2 Governance and formalization of agents

Firms with AI agents as part of their workflows varied in whether they had enacted formal policies to govern the use of AI agents, who takes accountability for errors, and steps taken to formalize (and humanize) the AI agent and its role.

In terms of governance, fewer than half of respondents (48%) report that their organization has a formal policy to govern the use of AI agents (Figure 4). Even fewer report that they have auditing or monitoring of its outputs (41%), formal human managers (38%), or have employee roles focused on AI agent oversight (36%). Ten percent of managers at firms where agents have been integrated into workflows indicated that "none of the above" practices had been implemented.

When an AI agent makes an error, firms varied in whether individuals or teams owned the agents' mistakes. Most commonly, the employee using the agent is primarily responsible for addressing that error (41% of cases). However, in 25% of cases accountability fell to the team or department who owns the process. A smaller share of respondents indicated that errors were treated as an issue to be addressed by the team that implemented the system (15%) or a systems issue overall (10%). This represents another way that AI agents could change work without displacing employees: even if agents are very effective, individual humans may be responsible for dealing with residual cases where the agent makes mistakes, without necessarily being assigned a formal oversight role.

Finally, agents received different types of formal recognition at the firms using them. Twenty-nine percent of managers at firms where AI agents are used report the agent has a name, and a fifth of respondents report that the agent is referred to as a digital worker or employee. At a quarter of firms the agent is listed on an organization chart, given a job title, and assigned human managers (shown in Figure 5). Most firms (67%) had done at least one of these things while only 2% had done all five (Figure A5). This is important because whether agents are formally recognized (and sometimes anthropomorphized) may impact how humans perceive and react to their output (Fumagalli et al., 2022; Wiles et al., 2026).

Figure 4: Governance and Accountability

(a) Which governance practices, if any, does your company currently have in place for AI agents?

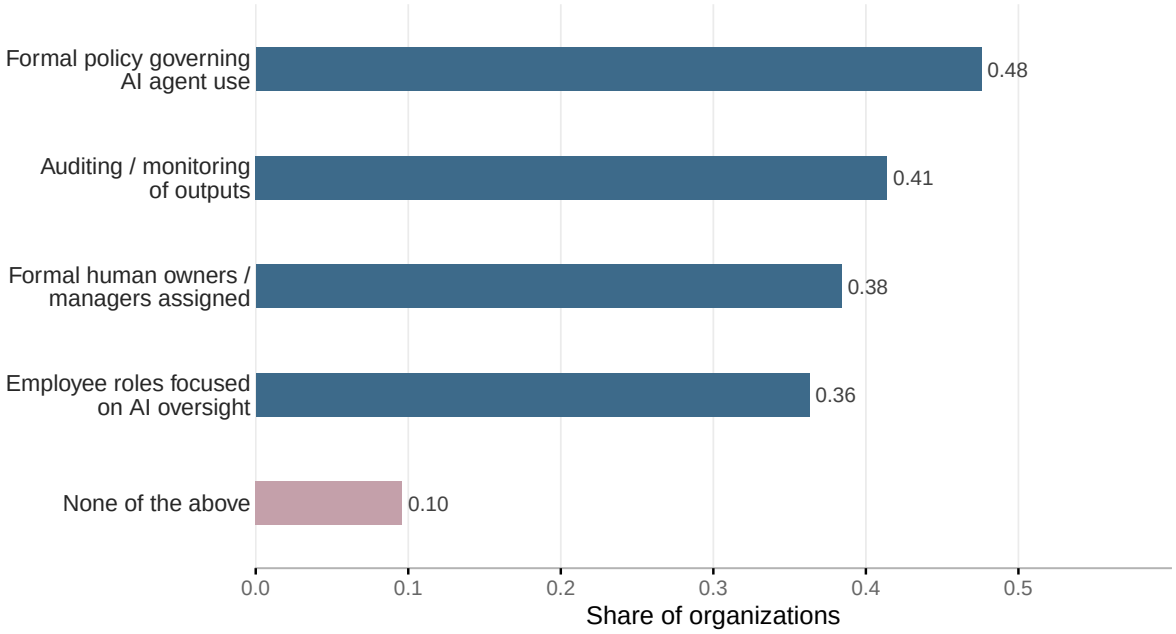

(b) When an AI agent makes an error, who is responsible for addressing it?

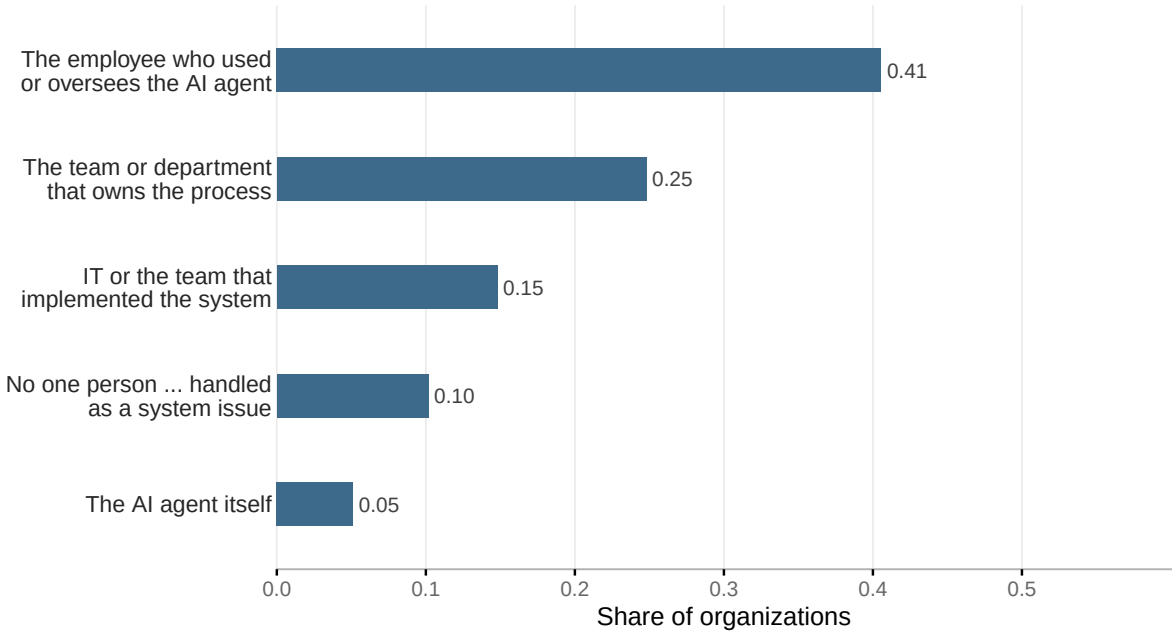

Notes: This plot reports the answers to the question "When an AI agent built into your company's workflows makes an error, who is primarily responsible for addressing it?" Here we plot the results for only those senior managers who report having AI agents embedded in their organization, N = 764. Responses "not sure" or "other" are not plotted. Estimates are weighted using YouGov respondent weights (see Section 2 for details).

Figure 5: **Which, if any, formalization practices has your organization taken?**

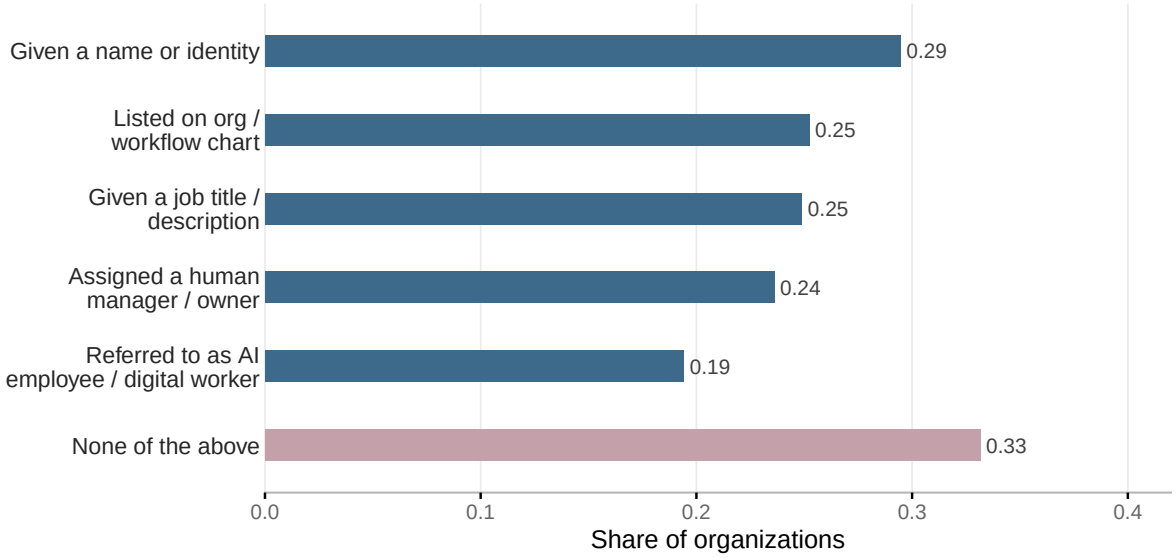

Notes: This plot reports the answers to the question "Are you aware of your company taking any of the following steps to formally recognize these AI Agents as part of your organization? Select all that apply." Here we plot the results for only those senior managers who report having AI agents embedded in their organization, N = 764. Responses "not sure" or "other" are not plotted. Estimates are weighted using YouGov respondent weights (see Section 2 for details).

3.2.3 Typology of AI-adopting organizations

This variation in how managers report their organization has integrated AI agents and enacted policies to govern them raises a natural next question: are organizations that integrate AI agents most completely those that also enact the most governance? These two margins of adjustment could reflect the same underlying organizational capability or absorptive capacity to adjust to a novel technology. Conversely, if these reflect different aspects of organizational implementation of AI agents, this opens up the possibility for more variance in the implementation of AI agents and hence differences in their impacts on firms, work, and employees.

We find a positive but low correlation between reported work and governance changes. For example, the correlation between a respondent indicating that they had reassigned employees to new tasks when tasks were taken over by an agent and indicating that their organization had adopted a formal policy governing the use of AI agents is .283 (Figure A6 presents a heatmap showing the correlation between all levels of these questions). This suggests that the between-respondent variation in integration of AI agents and agentic AI governance practices is not solely driven by some respondents indicating their organization has taken a comprehensive approach to both dimensions of agentic AI adoption.

To further explore the relationship between the integration and governance of AI agents, we used all survey questions to group respondents into distinct categories using latent class analysis,

a form of unsupervised machine learning designed to uncover subgroupings. We group the 764 respondents who indicated AI agents are part of their company's system or processes based on their unweighted responses to all questions shown or discussed in-text (Q3-Q14 in the Appendix Section B). Four distinct groups emerged as the preferred model that minimized Bayesian information criteria (Δ BIC relative to K=3: 293; relative to K=5: 57). The four groups were distinctly separated from each other: the average respondent could be assigned to their group with 95% confidence (a mean maximum posterior class probability of 0.952) and only about 1 in 20 respondents (5.4%) were "ambiguous" cases (a maximum posterior probability below 0.70, conventionally acceptable levels).[2] Taken together, this suggests that these clusters are well-separated and the four profiles reflect managers reporting distinct approaches to integrating and governing AI agents at their firms. Figure 6 illustrates the groups (with labels assigned by the authors) and Table 1 presents the variables that deviate the most between that group and their overall prevalence in the sample.

There are two groups whose integration and governance are both high or low relative to the rest of the sample. The group we label "passive" adopters (about a third of the sample) are both more likely to report no workforce changes or task takeover due to AI agents in their workflows, and to report no future governance plans. On the other hand, the group we label "institutionalized" adopters (almost a fifth of the sample) report relatively intensive work changes and policy responses to the introduction of AI agents. At these firms, respondents report that job descriptions have been changed and employees assigned to new tasks, AI oversight roles have been created and human managers assigned to AI agents, new employees have been hired to manage AI agents, and headcount has been reduced.

In contrast, the two other groups score high on integration or governance and low on the other dimension relative to the rest of the sample. A first group, which we label "policy-led adopters" (13.5% of the sample), scores relatively high on reported governance changes but low on integration. These respondents are more likely relative to the sample to report having AI agents deployed across multiple departments and implemented as a result of a deliberate decision by senior leadership, but not to have made many major workforce changes despite enacting organizational policies to govern them. The final third of the sample, which we label "action-led adopters," report reassigning employees to new tasks and roles and hiring new employees to manage AI. However, these major adjustments are not paired with formalization or governance practices. The existence of these two groups suggests that agentic AI integration and governance need not co-occur, and that organizations may differ in how they experiment with implementing AI.

[2] Appendix Table A2 contains additional details on model selection and Appendix Section B.1.1 additional robustness checks. In the robustness specification we redefine multi-select variables and find that in this case BIC is still minimized when K =4.

Figure 6: Typology of agentic AI adoption

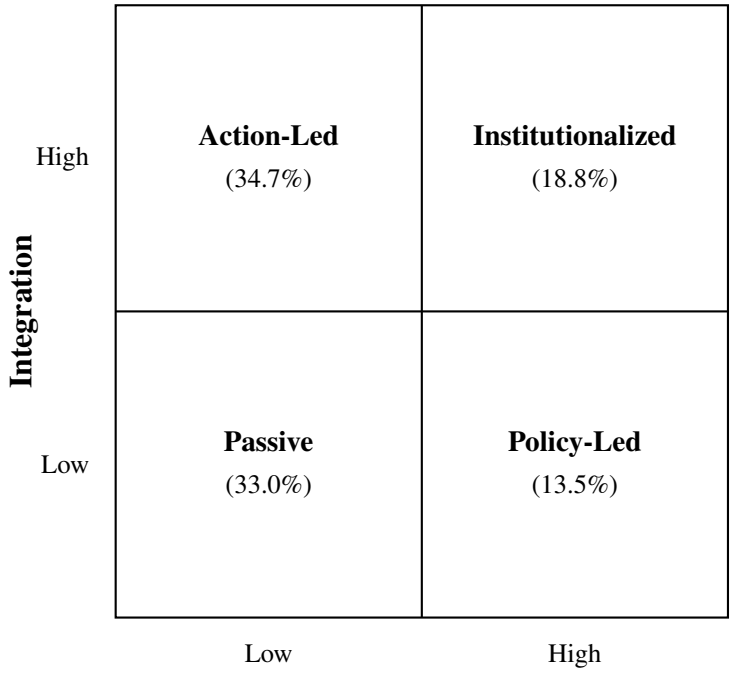

Notes: Figure shows a typology of the adoption of AI agents by organizations based on the groupings that emerged from latent class analysis. Organizations can score high or low on agentic integration and governance relative to the rest of the sample. Institutionalized adoption reflects firms making both major workforce changes and enacting policies to formalize and govern their AI agents. Passive adoption reflects firms using AI agents but making neither major workforce changes nor enacting such policies. The off-diagonals represent firms where adoption of AI agents has produced labor adjustments but not policy adjustments (action-led adoption) and those firms that have made policy adjustments but no major workforce changes (policy-led adoption).

Table 1: Latent class typology of AI agent adoption (K=4)

Class	Share	Most distinctive variables
Policy-Led Adoption	13.5%	Develop/update AI policy (+20%); Deliberate leadership decision (+19%); Wide departmental reach (+35–41% across depts); No workforce changes (+26%)
Passive Adoption	33.0%	No workforce changes (+47%); No task takeover (+38%); No governance plans (+23%)
Action-Led Adoption	34.7%	Reassigned employees affected by automation to new tasks (+18%); Moved affected employees to new roles (+12%); Hired new employees to manage AI (+11%)
Institutionalized Adoption	18.8%	Changed job descriptions (+54%); Assigned human manager (+43%); Created AI oversight roles (+43%); Reassigned employees affected by automation to new tasks (+42%); Hired new employees to manage AI (+39%); Reduced headcount (+32%)

Notes: The K=4 latent class solution from a poLCA model estimated on N=764 respondents using a 47-variable manifest. Share is the proportion of respondents in each class. Most distinctive indicators are the binary variables with the largest positive deviation between the estimated class conditional probability and prevalence in the whole sample; deviations shown in parentheses.

4 Discussion

4.1 Limitations

There are five important caveats to these results. First, we report on senior managers' perceptions of their company's use of AI agents, which may be imperfect. Managers may mis-perceive both the depth of integration and the extent of firm governance. Second, these data are a single cross-sectional snapshot of AI adoption. We document a remarkable breadth of AI agent adoption by organizations that notably mirrors the breadth of consumer use of other generative AI technologies (Chatterji et al., 2025). However, this use may increase or decrease over time as organizations discover whether these technologies actually benefit their firm. Similarly, while we find that at the moment of the survey some managers report observing action-led vs. policy-led adoption we cannot tell if their organizations later catch up on their other dimension of adoption. For example, the organizations we call action-led adopters may lead with workforce adjustments and later enact governance policies, or they may never do so. A single survey cannot distinguish between these cases, though they may have different outcomes for the firm. Third, as previously

discussed, AI "agents" may conflate AI products of varying sophistication, ranging from third-party produced AI-powered assistants like Microsoft Copilot to purpose-built agents for customer support, to agentic diagnostic tools. While some of this variation is captured by asking about how agents are obtained, managers' understanding of "agent" may still capture a broad range of functionality and sophistication, which also reflects the breadth of this emerging technology category. Fourth, our survey is limited to managers at U.S. firms. Finally, we have no data on whether AI agents have enhanced firm performance or productivity at these adopting firms: an open question for future work.

4.2 Future Work

This survey has suggested that agentic AI technologies are diffusing across the U.S. workforce and in multiple functional areas of organizations, but that there are differences in concurrent workforce and policy adjustments. This suggests three areas for future work.

First, what drives an organization to pursue deep integration of agentic AI, extensive governance, or both? Just as AI users differ systematically in their characteristics and how they acquire AI capabilities (Calvino and Fontanelli, 2026), we show that adopters of agentic AI differ in how they integrate and govern these tools once acquired. As with previous types of AI, firm characteristics might predict agentic AI adoption (McElheran et al., 2024) and prior capabilities or competencies might impact whether managers and firms can capitalize on investments in agentic AI (Teece, 1986; Cohen and Levinthal, 1990; Kaplan and Tripsas, 2008). Adapting an organization in response to a promising new technology is notoriously difficult, and future work might explore which firms and managers have the foresight and capability to recognize the need for certain types of adjustments in response.

Second, is this widespread diffusion of agents changing employees' experiences at work? For example, we found that in a plurality of cases individual workers are assigned accountability for errors committed by agents. However, humans may apply different standards when judging AI outputs as compared to human outputs, including by reducing monitoring intensity (Fumagalli et al., 2022; Wiles et al., 2026). While some organizations in our sample treat agents as "digital employees," this approach may have unintended consequences for organizations. There is also mounting discussion of whether reliance on AI tools may "de-skill" professionals who delegate work to them (Lenharo, 2026), including in ways that may diminish core human capacities (Ferdman, 2026). While we find that only a minority of managers report that agents exercise decision authority, if agents begin to exercise more autonomy this raises fundamental questions about human-machine interaction, co-production, and the value of expertise (Agrawal et al., 2018; Murray et al., 2021; Autor and Thompson, 2025).

Third, can agentic AI adoption enhance productivity, and if so what type of adoption does that involve? This paper has suggested that adopting AI agents is not a single choice, but rather a series of choices about reorganizing workflows, reassigning employees, and establishing policies. As in prior waves of potentially transformative technologies, conflating all firms where agents are present as "adopters" might mask important heterogeneity in the adoption of AI agents and hence its impacts.

5 Statement on AI

The authors used Generative AI tools to assist with code generation. The authors directed its use, reviewed all generated code, and take full responsibility for all aspects of the paper. Generative AI tools were not used to write any part of the paper.

References

Agrawal, A., Gans, J., and Goldfarb, A. (2018). Prediction, judgment, and complexity: a theory of decision-making and artificial intelligence. In *The economics of artificial intelligence: An agenda*, pages 89–110. University of Chicago Press.

Altman, S. (2025). Reflections.

Autor, D. and Thompson, N. (2025). Expertise. *Journal of the European Economic Association*, 23(4):1203–1271.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative ai. Technical report, National Bureau of Economic Research.

Bitter, A. (2026). Satya Nadella says AI agents should be treated like employees with identities, permissions, and audits.

Bloom, N., Sadun, R., and Reenen, J. V. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review*, 102(1):167–201.

Bresnahan, T., Greenstein, S., Brownstone, D., and Flamm, K. (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity. Microeconomics*, 1996:1–83.

Calvino, F. and Fontanelli, L. (2026). AI users are not all alike: The characteristics of French firms buying and developing AI. *Research Policy*, 55(5):105473.

Census Bureau (2025). Business trends and outlook survey content v2.

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Yan Shan, C., and Wadman, K. (2025). How People Use ChatGPT. Technical report.

Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152.

Colaiacovo, I. and Koning, R. (2026). Delegating to AI: Beliefs and the Organization of Work at New Ventures. Working paper. Available at SSRN: `https://ssrn.com/abstract=6962978`.

Ferdman, A. (2026). AI deskilling is a structural problem. *AI & SOCIETY*, 41(4):3001–3013.

Fumagalli, E., Rezaei, S., and Salomons, A. (2022). Ok computer: Worker perceptions of algorithmic recruitment. *Research Policy*, 51(2):104420.

Goos, M. and Savona, M. (2024). The governance of artificial intelligence: Harnessing opportunities and mitigating challenges. *Research Policy*, 53(3):104928.

Grundy, A., Breaux, C., and Khatiwoda, D. (2026). Large Firms With at Least 20 Employees Biggest AI Users.

Huang, J. (2025). NVIDIA GTC Taipei 2026 Keynote Replay.

Hunt, W., Sarkar, S., and Warhurst, C. (2022). Measuring the impact of AI on jobs at the organization level: Lessons from a survey of UK business leaders. *Research Policy*, 51(2):104425.

Kaplan, S. and Tripsas, M. (2008). Thinking about technology: Applying a cognitive lens to technical change. *Research policy*, 37(5):790–805.

Karger, E., Kuusela, O., Abaluck, J., Bryan, K., Halperin, B., Jones, T., Murphy, C., Trammell, P., Reynolds, M., Mayland, D., Viswanathan, R., Mittal, A., Ceppas de Castro, R., Rosenberg, J., and Tetlock, P. E. (2026). Forecasting the economic effects of AI. Working paper.

Lenharo, M. (2026). Is AI ruining our skills? Early results are in — and they're not good.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., and Zolas, N. (2024). Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415.

Murray, A., Rhymer, J., and Sirmon, D. G. (2021). Humans and technology: Forms of conjoined agency in organizations. *Academy of Management Review*, 46(3):552–571.

Simon, H. A. (1960). The corporation: Will it be managed by machines? In Anshen, M. L. and Bach, G. L., editors, *Management and the Corporations, 1985*, pages 17–55. McGraw-Hill, New York.

Staufer, L., Feng, K., Wei, K., Bailey, L., Duan, Y., Yang, M., Ozisik, A. P., Casper, S., and Kolt, N. (2026). The 2025 ai agent index: Documenting technical and safety features of deployed agentic ai systems. In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '26, page 1536–1576. ACM.

Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6):285–305.

Wiles, E., Hsu, M., Bedard, J., and Kropp, M. (2026). Putting AI on the Org Chart: Evidence on Delegation and Oversight. *Working Paper*.

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G., and Wang, B. Z. (2026). Firm Data on AI. *NBER Working Paper 34836*.

A Additional Tables and Figures

Figure A1: Adoption of AI Agents in Organizations, by NAICS Sector

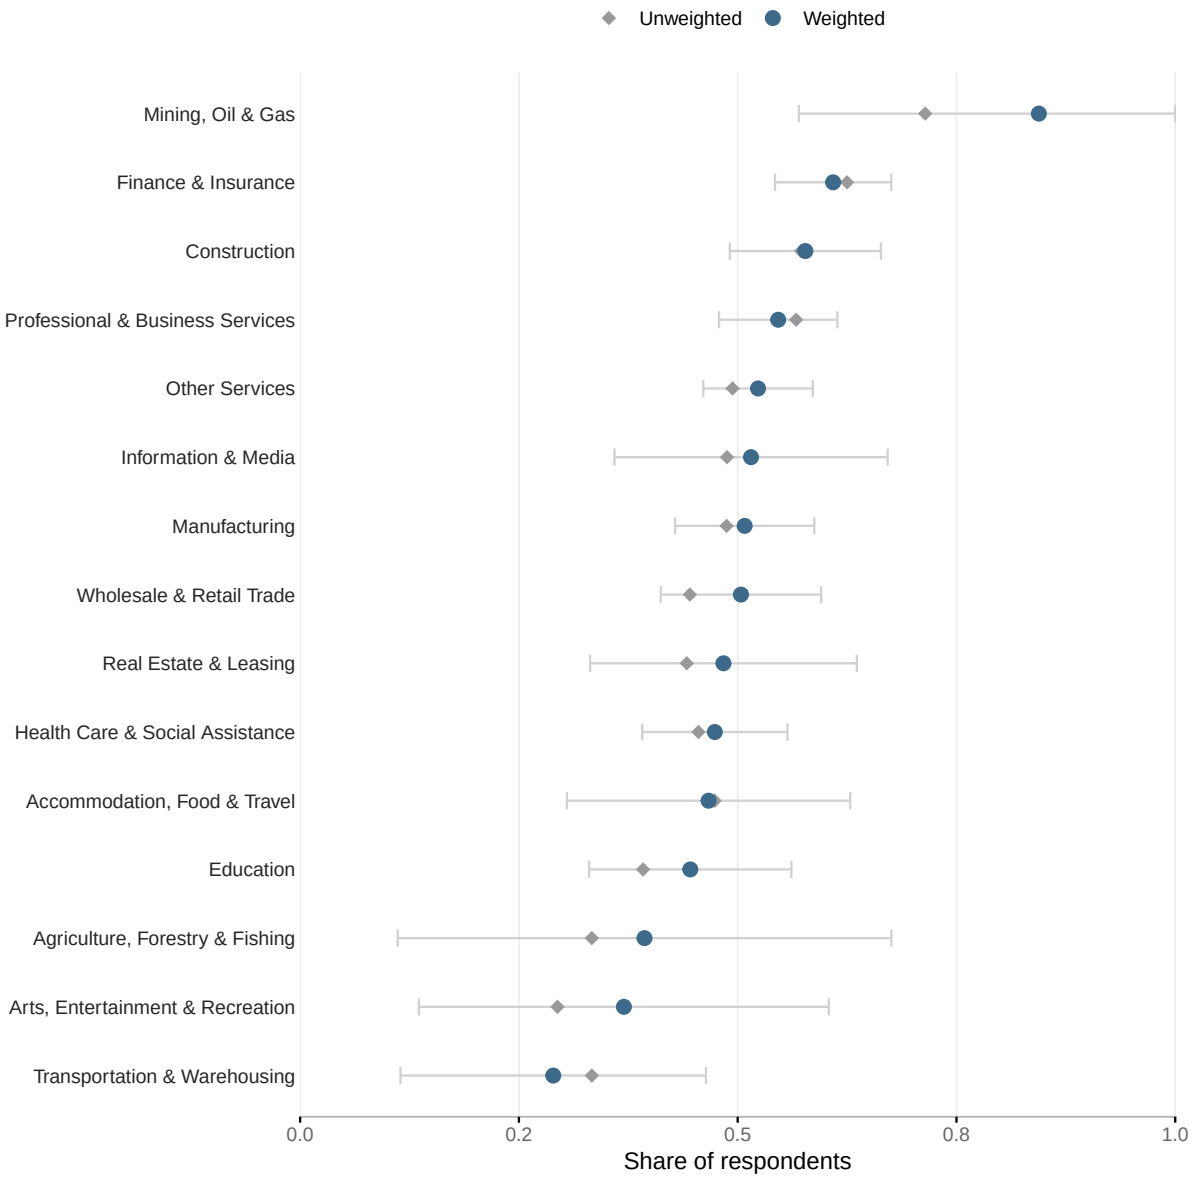

Notes: This plot shows the share of all respondents who responded yes to "Are you aware of your company having officially integrated AI agents into its systems or workflows to complete tasks automatically? (e.g., use of an AI agent to automatically screen job applications, processing invoices, or route documents for approval)" by 2 digit NAICS sector. N = 1,500. Estimates are weighted using YouGov respondent weights (see Section 2 for details)

Table A1: Share of organizations reporting AI agents built into their work processes, by NAICS sector

Sector code	Sector name	N	Wtd. share
11	Agriculture, Forestry, Fishing and Hunting	12	0.394
21/22	Mining, Oil and Gas Extraction, and Utilities	7	0.844
23	Construction	131	0.577
31–33	Manufacturing	158	0.508
42	Wholesale Trade	36	0.563
44–45	Retail Trade	83	0.479
48–49	Transportation and Warehousing	27	0.289
51	Media, Communications, and Digital Entertainment	41	0.515
52	Finance and Insurance	216	0.609
53	Real Estate, Rental, and Leasing	43	0.484
54	Professional, Scientific, and Technical Services	217	0.546
61	Education	74	0.446
62	Health Care and Social Assistance	145	0.474
71	Arts, Entertainment, and Recreation	17	0.370
72	Restaurant, Travel and Lodging	38	0.467
81	Other Services	255	0.523

Notes: shows the share of all respondents who responded yes to "Are you aware of your company having officially integrated AI agents into its systems or workflows to complete tasks automatically? (e.g., use of an AI agent to automatically screen job applications, processing invoices, or route documents for approval)" by 2 digit NAICS sector. N = 1,500. Estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

Figure A2: Share of organizations reporting AI agents built into their work processes, by firm size

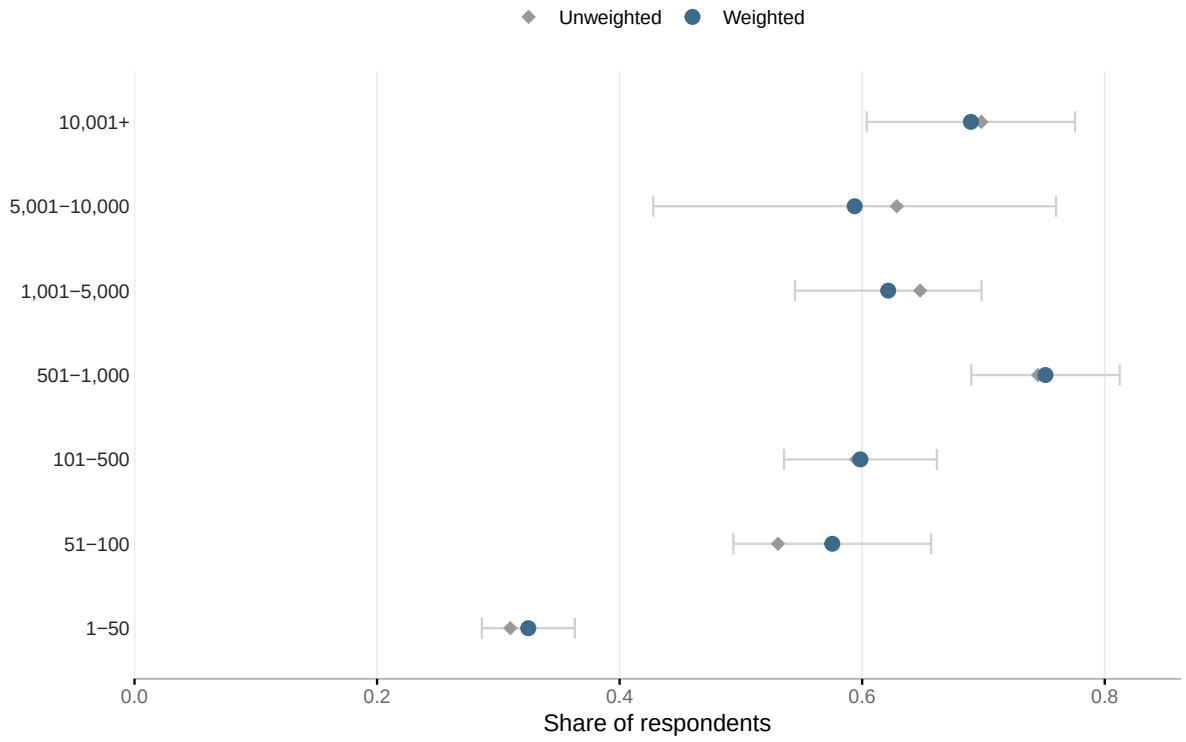

Notes: This plot shows the share of all respondents who responded yes to "Are you aware of your company having officially integrated AI agents into its systems or workflows to complete tasks automatically? (e.g., use of an AI agent to automatically screen job applications, processing invoices, or route documents for approval)" bucketed by firm size. N = 1,500. Estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

Figure A3: **What drove the decision to adopt AI agents?**

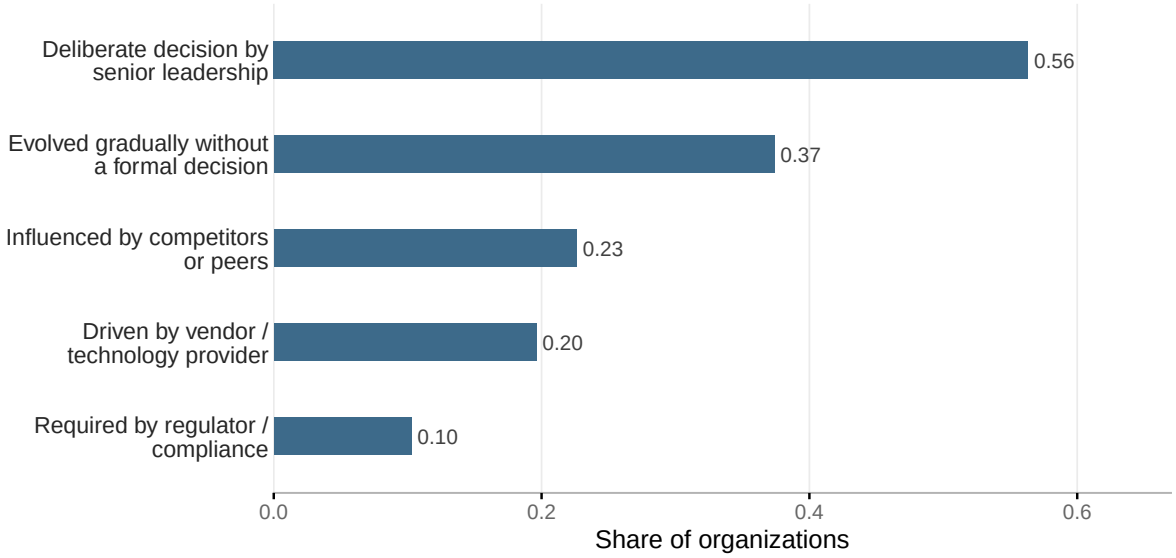

Notes: This figure reports how organizations came to have AI agents built into their work processes, conditional on managers reporting that their organization has AI agents built into their work processes. The figure reports the answer to the question "How did your company arrive at the current way it uses and describes AI agents?" Responses "not sure" or "other" are not plotted. Estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

Figure A4: AI agent use and governance plans for the next 12 months

(a) Future plans for use of AI agents in business processes

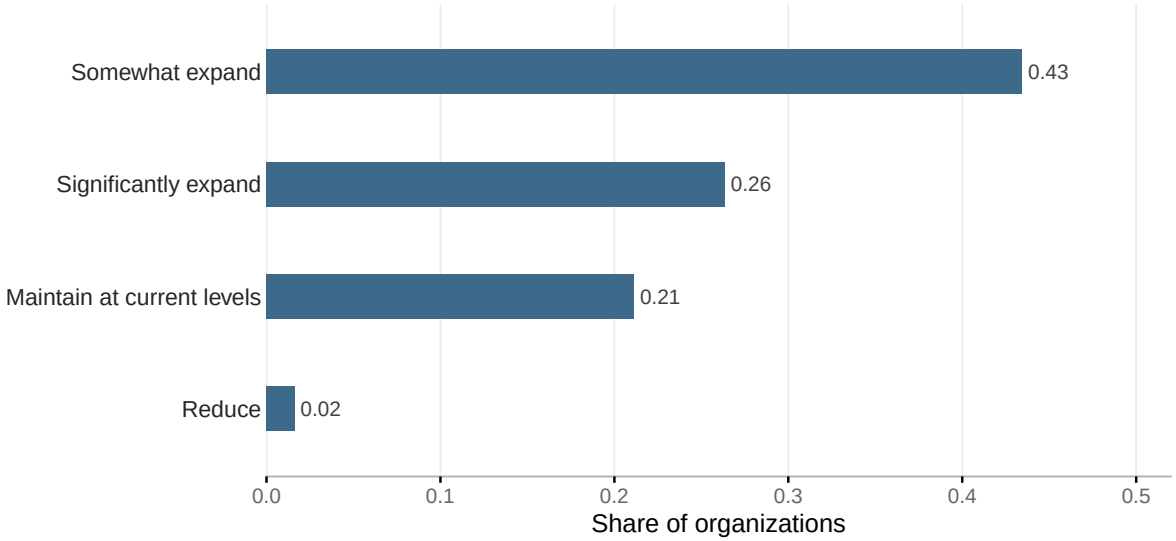

(b) Future governance plans

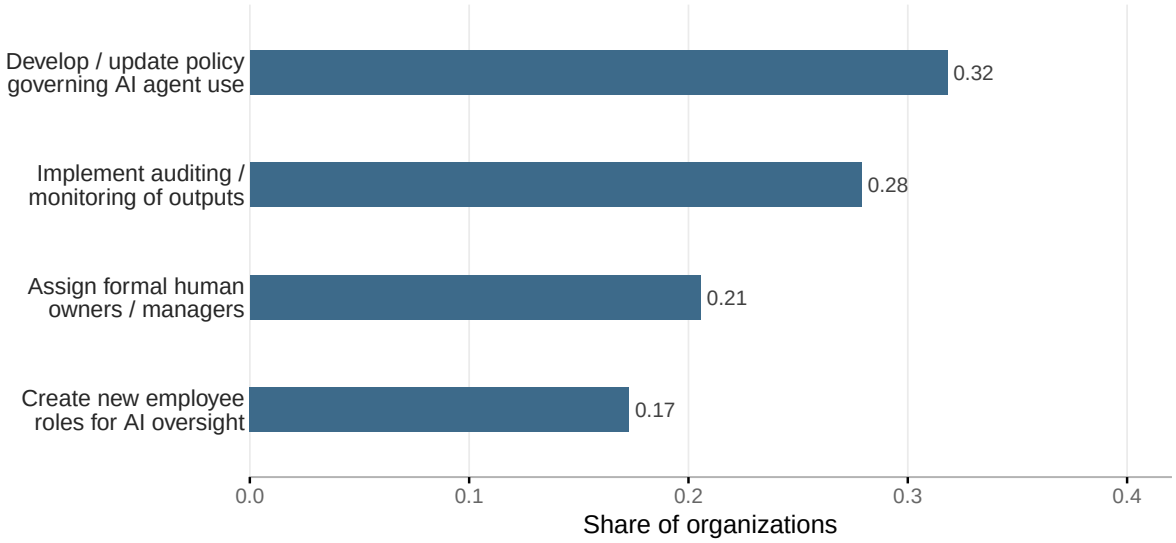

Notes: Panel a) reports the answers to the question "Over the next 12 months, does your company plan to expand, maintain, or reduce its use of AI agents in business processes?" Panel b) reports the answers to the question "Which of the following governance changes, if any, is your company planning to make regarding AI agents in the next 12 months?" These questions were asked to all respondents who indicated their organization has adopted AI agents into their work processes, N = 1500. Estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

Figure A5: Counts of formalization practices

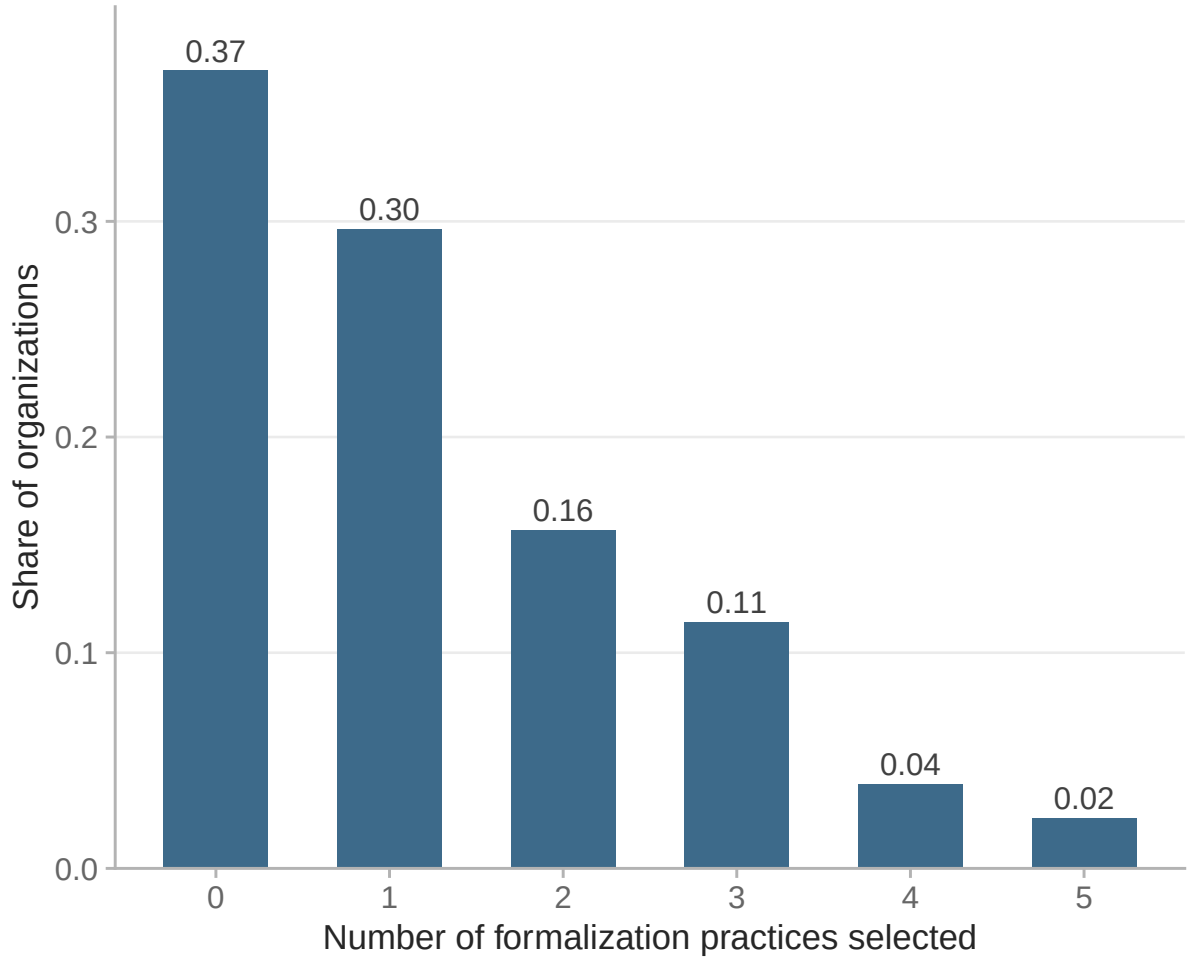

Notes: This figure reports how many different ways organizations formalized agents in their organizations, conditional on managers reporting that their organization has AI agents built into their work processes. It reports the number of possible decision drivers the senior manager selected to answer the question "Are you aware of your company taking any of the following steps to formally recognize these AI Agents as part of your organization? Select all that apply." N = 764. Estimates are weighted using YouGov respondent weights, which correct for within-sector differences in respondent demographics, including age, gender, race/ethnicity, education, and employment characteristics, relative to the US manager population.

Figure A6: Correlation between governance practices and workforce integration outcomes

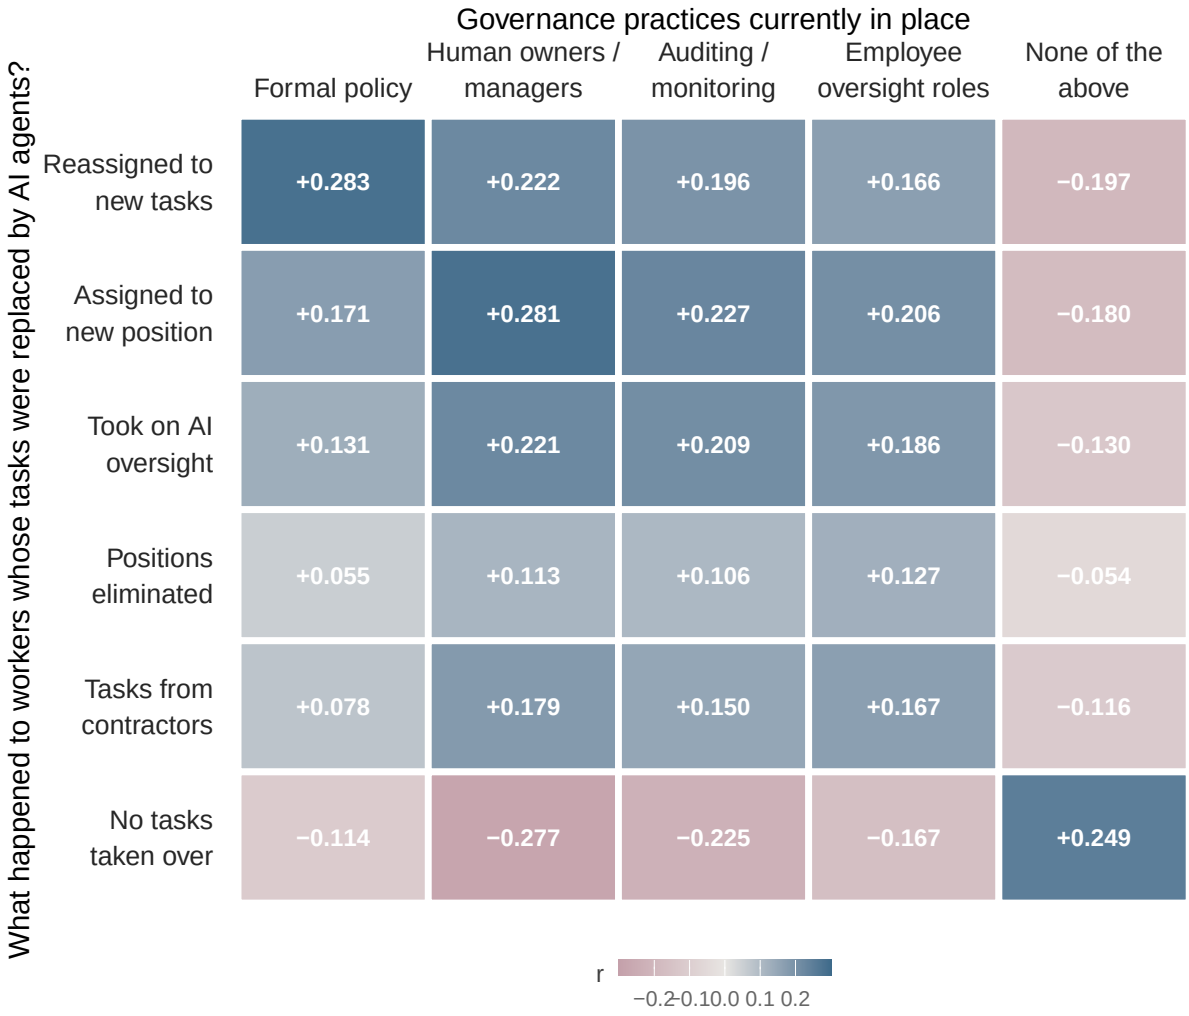

Notes: This figure shows pairwise Pearson correlations between Q9 governance mechanisms on the x axis and Q12 workforce integration outcomes on the y axis, among organizations who use AI agents. N = 764. All variables are binary.

B Survey Questions

Q1: Are you aware of employees at your company choosing to use AI agents on their own to automate their personal work tasks? (e.g., an employee using an AI assistant that automatically joins their meetings, takes notes, and emails them a summary afterward, or using Claude Code to automate writing code to clean a dataset— not integrated into formal company workflows.)

Q2: Are you aware of your company having officially integrated AI agents into its systems or workflows to complete tasks automatically? (e.g., use of an AI agent to automatically screen job applications, processing invoices, or route documents for approval)

Q3: [Ask if Q2 == “Yes”] In which functions or departments has your company deployed AI agents into workflows? Select all that apply.

Q4: [Ask if Q2 == “Yes”] Which best describes the typical way employees interact with AI agents at your company? (Select one)

Q5: Organizational formalization*[Ask if Q1 or Q2 == Yes]* Are you aware of your company taking any of the following steps to formally recognize these AI Agents as part of your organization? Select all that apply.

Q6: Decision Authority*[Ask if Q2 == Yes]*
****For AI agents used in your company’s business processes, which best describes the level of authority they typically have within workflows?

Q7: Accountability*[Ask if Q2 = = Yes]*
When an AI agent built into your company’s workflows makes an error, who is primarily responsible for addressing it?

Q8: How did your company arrive at the current way it uses and describes AI agents? Select all that apply.

Q9: Which of the following governance practices, if any, does your company currently have in place for AI agents? Select all that apply.

Q10: How were the AI agents used in your company’s business processes primarily sourced?

Q11: Since deploying AI agents in business processes, are you aware of your company making any of the following changes? Select all that apply.

Q12: To the best of your knowledge, what happened to employees whose tasks were taken over by AI agents? Select all that apply.

Q13: Over the next 12 months, does your company plan to expand, maintain, or reduce its use of AI agents in business processes?

Q14: Which of the following governance changes, if any, is your company planning to make regarding AI agents in the next 12 months?

B.1 Cluster analysis

Figure A7: BIC by number of latent classes

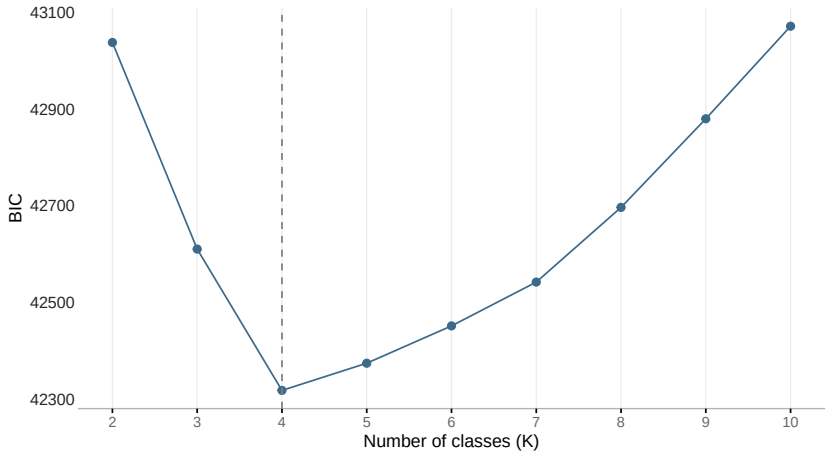

Notes: BIC by number of latent classes (K=2–10), estimated via poLCA on N=764 managers whose organizations use AI agents. Each model was fit with 200+ random starts; the plotted value is the best (highest) log-likelihood across starts. The dashed vertical line marks the minimum BIC at K=4, the preferred number of classes.

Figure A8: Two-dimensional visualization of the four classes of organizations

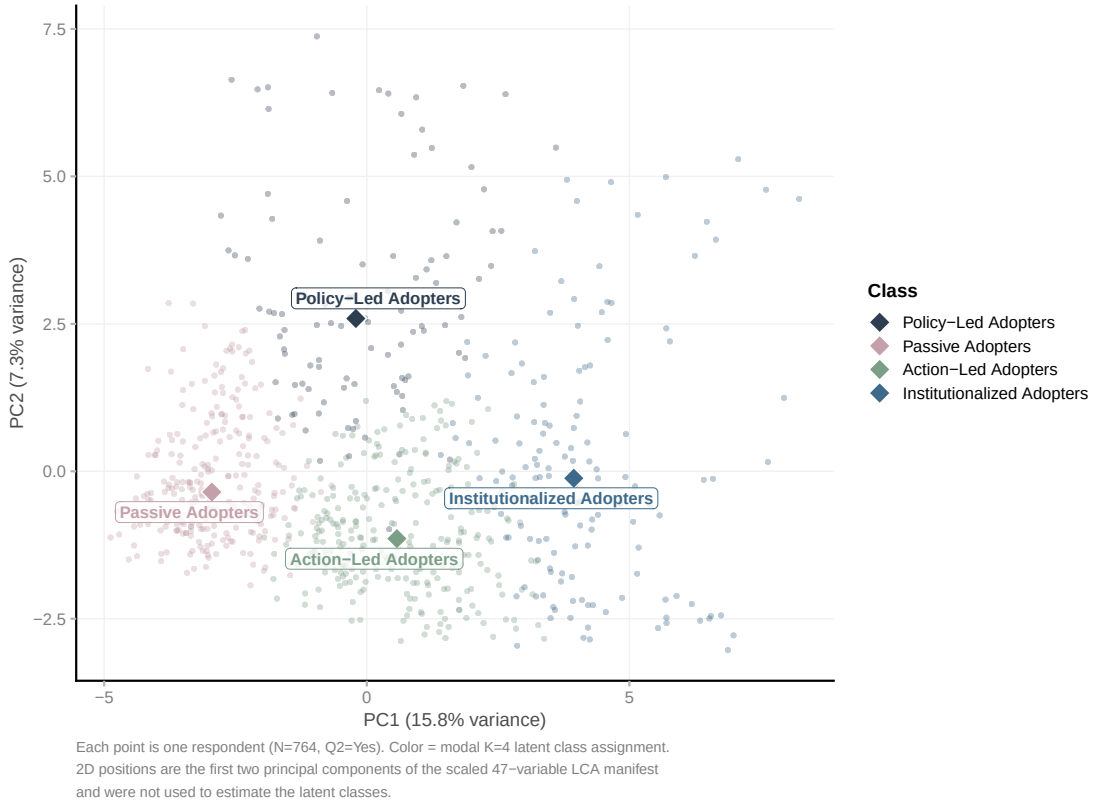

Notes: Principal component projection of the K=4 latent class solution. Each point is one respondent (N=764); color indicates modal class assignment from the K=4 LCA. Axes are the first two principal components of the scaled 47-variable LCA manifest (PC1 = 15.8% of variance, PC2 = 7.3%). PC1 is a workforce-restructuring gradient: respondents scoring higher were more likely to report employee reassignments, job-description changes, and governance activity; those scoring lower reported no significant workforce changes or task displacement. Diamond markers show class medoids. The two-dimensional positions are a linear projection of the high-dimensional survey responses and were not used to estimate the latent classes.

Figure A9: Most distinctive variables by for each class of organization

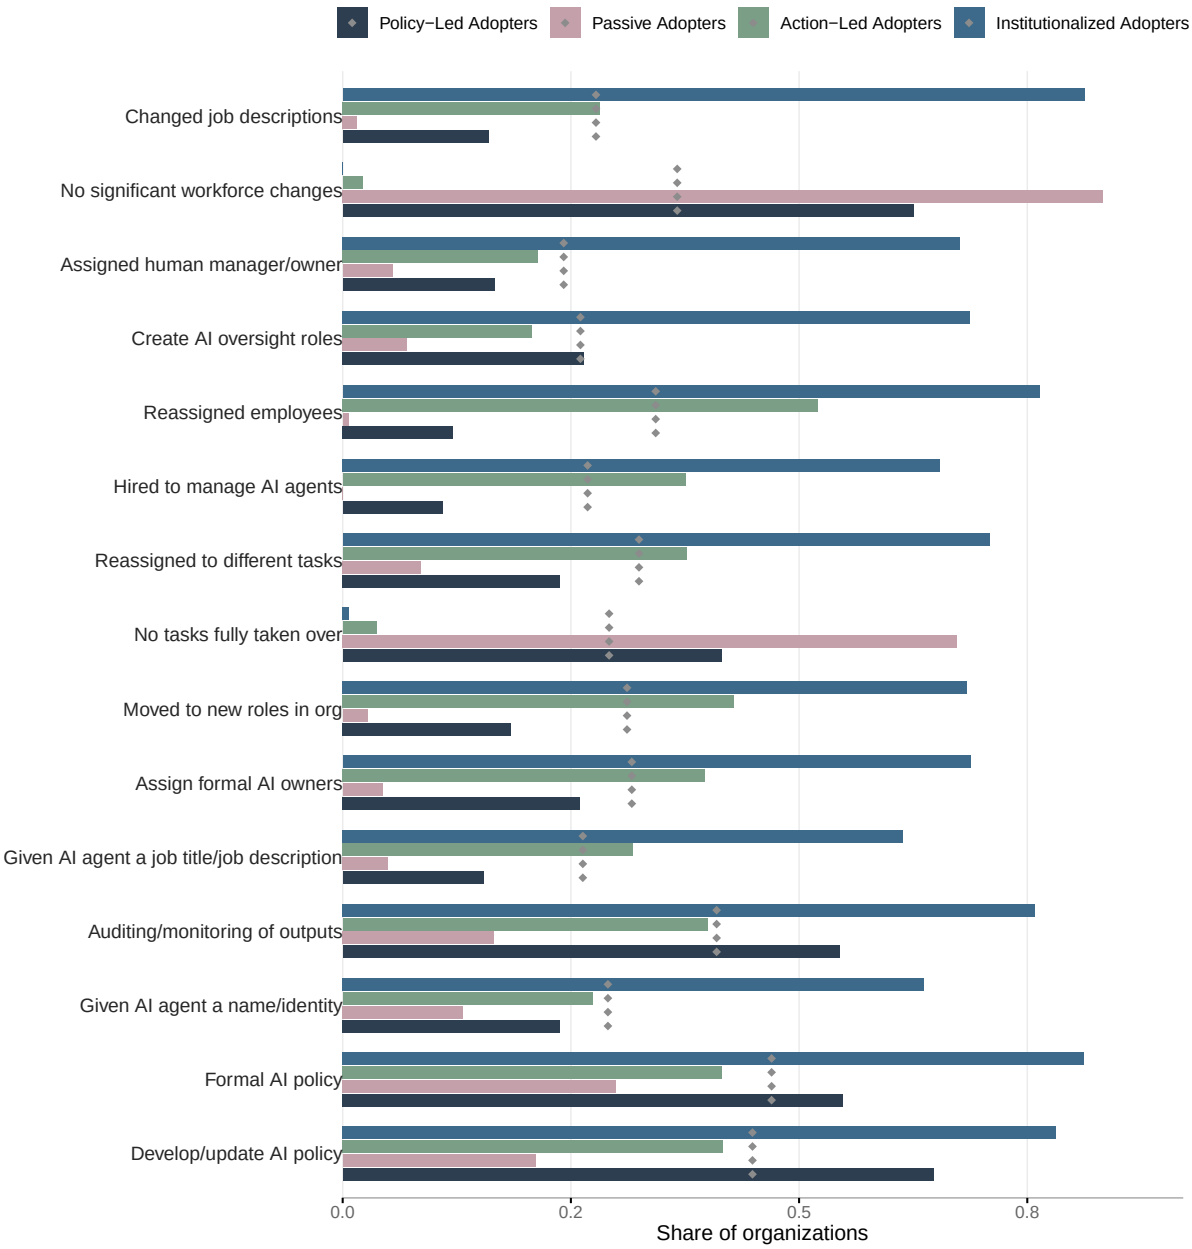

Notes: Top 15 items differentiating the four latent classes. Bars show the estimated probability of selecting each item given class membership, P(selected | class); diamond markers show the overall prevalence among all AI-agent-using organizations (N=764). Items are ranked by maximum absolute deviation from overall prevalence across the four classes. Department indicators are excluded from the plot. Classes: *Policy-Led Adopters* (13.5%) — broad departmental deployment with governance infrastructure but no workforce restructuring; *Passive Adopters* (32.9%) — AI use with minimal organizational impact; *Action-Led Adopters* (34.7%) — active workforce adjustment and emerging governance; *Institutionalized Adopters* (18.8%) — formal governance, changed job descriptions, and extensive workforce restructuring.

Table A2: Comparison of latent class models with 2 to 10 classes

K	Log likelihood	BIC	Δ BIC	Min class size	Mean max posterior	% below 0.70
2	-21124	43038	721	0.468	0.972	3.80
3	-20711	42610	293	0.222	0.953	4.97
4	-20365	42317	0	0.135	0.952	5.37
5	-20194	42374	56	0.077	0.956	4.84
6	-20034	42451	134	0.086	0.933	7.20
7	-19880	42542	224	0.087	0.935	6.68
8	-19758	42696	379	0.073	0.933	7.20
9	-19651	42880	563	0.029	0.941	6.02
10	-19548	43072	755	0.039	0.938	6.54

Notes: This table compares latent class models with 2 to 10 classes. K is the number of classes. Log likelihood is the maximized log likelihood of the fitted model. BIC is the Bayesian information criterion. ΔBIC is the difference between a model's BIC and the minimum BIC across the models considered, so that the preferred model has $\Delta\text{BIC} = 0$. Min class size is the minimum estimated class share for each K. Mean max posterior is the mean across respondents of their maximum posterior class-membership probability, and % below 0.70 is the percentage of respondents whose maximum posterior class-membership probability is below 0.70. $N = 764$, with 47 class-defining variables.

B.1.1 LCA Robustness

The primary models use all 47 variables created by making binary versions of each question. When a question allows for multiple selections we make each option its own binary variable. This induces correlation across variables within each question. Therefore in this robustness specification we collapse multi-select questions into composites variables.

The twelve department items in Q3 were grouped into four functional domain composites — technical (IT/Engineering, R&D, Strategy), customer-facing (Marketing, Sales, Customer Service), administrative (Finance, HR, Executive/Admin), and operations/risk (Operations, Legal, Procurement) — each coded 1 if any item in the domain was selected. The five AI agent formalization practices in Q5 and the four current governance mechanisms in Q9 were each collapsed into a single binary indicator of whether the organization had adopted any governance practice. For Q11 (workforce changes), Q12 (task takeover outcomes), and Q14 (planned governance changes), we constructed a single three-level ordinal variable for each question: 1 if no active options were selected, 2 if exactly one was selected, and 3 if two or more were selected. The five adoption driver items in Q8 were kept as individual outcomes, as each option represents a substantively distinct pathway to adoption. Single-choice questions (Q4, Q6, Q7, Q13) and the sourcing method question (Q10, recoded as a purchased/licensed binary) were included without modification. The resulting data contains 22 variables — 18 binary and 4 original — compared to 47 variables in the main analysis.

Table A3 shows that BIC is still minimized when K = 4.

Table A3: Comparison of LCA models with 2 to 10 classes (robustness: composite variables)

K	Log likelihood	BIC	Δ BIC	Min class size	Mean max posterior	% below 0.70
2	-10990	22439	231	0.469	0.988	1.44
3	-10762	22214	6	0.224	0.977	2.09
4	-10642	22208	0	0.121	0.970	2.49
5	-10532	22220	12	0.061	0.965	3.40
6	-10434	22256	49	0.061	0.965	2.88
7	-10359	22337	130	0.025	0.968	2.75
8	-10289	22430	222	0.029	0.943	5.76
9	-10217	22518	311	0.023	0.922	9.03
10	-10154	22625	417	0.025	0.922	9.16

Notes: This table compares latent class models with 2 to 10 classes. K is the number of classes. Log likelihood is the maximized log likelihood of the fitted model. BIC is the Bayesian information criterion. ΔBIC is the difference between a model's BIC and the minimum BIC across the models considered, so that the preferred model has $\Delta\text{BIC} = 0$. Min class size is the minimum estimated class share for each K. Mean max posterior is the mean across respondents of their maximum posterior class-membership probability, and % below 0.70 is the percentage of respondents whose maximum posterior class-membership probability is below 0.70. $N = 764$, with 22 class-defining variables constructed by collapsing multi-select questions into binary or composite variables.