Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The New Digital Divide Inside Organizations: AI Access, Career Opportunities and Wage Inequality

Sarfaraz Abu Abdullah

rajon3430@gmail.com

University of Rajshahi https://orcid.org/0009-0000-3826-1626

Research Article

Keywords: generative AI, digital divide, technology adoption, wage inequality, skill-biased technical change, organizational complementarities, career mobility

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11818566/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

Field experiments on generative artificial intelligence (AI) report their largest gains for novices and lower-skilled workers, and this has encouraged a hopeful reading in which the technology narrows skill gaps. Those experiments, however, give access to everybody by design. Inside real organizations, access is allocated by rank, income, occupation and managerial attention. This paper asks who is actually using AI at work, how that use is linked to career opportunities, and what it implies for wage inequality. I combine a structured synthesis of recent experimental, administrative and survey evidence with new calculations on published figures from the Gallup Workforce Panel (more than 30,000 US employees, 2023 to the first quarter of 2026) and Gallup's February 2026 survey of 23,717 employees. Three findings emerge. First, frequent AI use rose in every group, but the gap between leaders and individual contributors grew from 8 to 25 percentage points between 2023 and the second quarter of 2026, even though the ratio between the two stayed near 1.9. Second, frequent use rises by roughly 10 percentage points for each doubling of household income ($R^2 = 0.81$ across five income brackets). Third, enabling conditions such as manager support and tool integration are associated with differences of 21 to 34 percentage points in frequent use. A simple break-even calculation shows that the equalizing effect found in experiments disappears once the access ratio between lower- and higher-skilled workers falls to the ratio of their gains per user, and the observed gradients sit close to that threshold under plausible assumptions. Posted wage premiums for AI skills (62 percent on average in 2025) and the fall in early-career hiring in exposed occupations suggest that unequal access can turn into unequal careers. The evidence is descriptive and rests on aggregated secondary data; I set out the limits and a research agenda for linked employer-employee data.

1. Introduction

For thirty years the story of workplace technology in rich countries has had a stubborn shape. Computers and software helped people who were already well educated and well paid, and the wage distribution stretched (Katz and Murphy, 1992; Autor, Levy and Murnane, 2003). When generative AI arrived in late 2022, early evidence seemed to promise a different ending. In the best-known field study, customer-support agents given an AI assistant became about 14 percent more productive on average, and the least experienced among them improved by roughly 35 percent, while the most experienced agents gained little (Brynjolfsson, Li and Raymond, 2025). Noy and Zhang (2023) found a similar leveling pattern among college-educated professionals writing mid-level business documents. Brynjolfsson, in a widely quoted remark, called it fascinating that this technology 'goes the other way around'.

There is a catch in that evidence, and it is the starting point of this paper. In an experiment, the researcher decides who gets the tool. In an organization, somebody else decides, usually without announcing the rule. A tool that helps the weakest performer most does not reduce inequality if the weakest performer never receives it, never learns to use it well, or is not allowed to use it for the tasks that carry weight in promotion decisions. The relevant question is therefore not only how large the gain

per user is and whom it favors. It is also how access, depth of use and recognition are distributed across the people who work in the same firm.

Researchers who studied the first internet recognized this problem long ago. The early 'digital divide' concerned who had a connection; the later, second-level divide concerned who had the skills and the purpose to use it well (Hargittai, 2002; van Dijk, 2005). Generative AI raises the question again in a form that sits inside the organization. Almost anybody can open a chatbot in a browser, so the divide between households matters less than it did. What matters is whether the employer supplies licensed tools connected to company data, whether managers encourage experiments, whether training exists, and whether an employee who becomes good at using AI is noticed. The recent data suggest large differences on exactly these points. Gallup's February 2026 survey of 23,717 US employees found that where AI tools were available, 67 percent of leaders used them a few times a week or more, against 46 percent of individual contributors (Gallup, 2026). A survey of 3,600 employees by Perceptyx found that 79 percent of executives but only 58 percent of individual contributors said generative AI was at least somewhat integrated in their workflow, and only 47 percent of contributors said they understood how AI tools were chosen (HCAmag, 2025). Infosys reported that more than a third of senior leaders but 19 percent of junior staff felt always equipped with the AI tools they needed (CIO Dive, 2026).

This paper does three things. It documents the within-organization divide in access and use, using the most recent published data I could assemble, up to the second quarter of 2026. It connects that divide to career opportunities through three channels: skills that employers price, early-career hiring, and promotion into roles that direct AI rather than do the work AI can do. And it asks, with a transparent calibration, under what conditions the leveling effect found in experiments survives unequal access. The contribution is partly empirical, because I bring together figures that have not been analyzed side by side, and partly conceptual, because I argue that the policy-relevant divide is no longer between people with and without a device but between people with and without organizational permission, support and recognition.

A word on what the paper is not. I did not collect primary data, and no survey I use contains wages and AI use for the same worker. All statistics below are computed from published aggregates, and I flag every place where an assumption stands in for a missing variable. The findings are descriptive. They show where inequality is already visible and where theory predicts it will grow, and they identify the data that would settle the causal questions.

The rest of the paper is organized as follows. Section 2 reviews the literature and states five propositions. Section 3 describes the evidence base and the methods. Section 4 reports results. Section 5 discusses mechanisms, implications and limits, and Section 6 concludes.

2. Background and hypotheses

2.1 What experiments say about who gains

Table 1 summarizes the main experimental and quasi-experimental studies. The findings share a pattern, but they differ in what they can tell us about organizations. In the Brynjolfsson et al. study, a staggered roll-out of a GPT-based assistant to roughly 5,200 agents at a Fortune 500 software firm increased resolutions per hour by about 14 percent, with gains near 35 percent for the lowest-skilled and least experienced agents. An agent with two months of tenure using the tool performed about as well as an agent with six months of tenure without it. The authors argue that the model captures and spreads the behavior of the best agents. In Noy and Zhang's preregistered experiment with 453 professionals, ChatGPT cut the time taken by 40 percent and raised graded quality by 18 percent, and inequality between workers fell.

Dell'Acqua and colleagues (2026) ran a preregistered field experiment with 758 consultants at the Boston Consulting Group. On tasks inside what they call the 'jagged technological frontier', consultants with GPT-4 completed 12.2 percent more tasks and finished 25.1 percent faster, with large quality gains. The consultants who scored worst at baseline gained most (Mollick reports 43 percent for the bottom half). But on a task outside the frontier, accuracy fell from 84.5 percent for the control group to 70.6 percent for consultants with GPT-4 and 60 percent for those who also received a prompting overview. The lesson for this paper is that gains depend on judgment about where the tool is reliable, and judgment is itself unequally distributed and trained.

Table 1
Selected experimental and quasi-experimental evidence on AI and worker productivity

Study	Setting / sample	Average effect	Distributional finding
Brynjolfsson, Li and Raymond (2025)	Customer support; about 5,200 agents; staggered roll-out	About + 14% resolutions per hour	About + 35% for least experienced / lowest skill; little or none for the most experienced
Noy and Zhang (2023)	Online experiment; 453 college-educated professionals	−40% time; +18% graded quality	Inequality between workers fell
Dell'Acqua et al. (2026)	Field experiment; 758 BCG consultants	+ 12.2% tasks; +25.1% speed (inside frontier)	Weakest at baseline gained most; outside the frontier, accuracy fell from 84.5% to 70.6% (60% with prompting overview)
Humlum and Vestergaard (2025, 2026)	Danish registers linked to surveys; 11 exposed occupations; 25,000 workers; 7,000 workplaces	Precise null on earnings and hours; effects larger than 2% ruled out	Task reorganization and movement into higher-paying occupations, but too few to move average earnings
Brynjolfsson, Chandar and Chen (2025; rev. 2026)	ADP payroll data; millions of US workers through June 2026	No economy-wide displacement	Employment of workers aged 22 to 25 in AI-exposed jobs about 19% below the path of less-exposed peers; no comparable gap for experienced workers

Note: Figures are as reported by the authors. Brynjolfsson et al. earlier headlined regression estimates of a 13 percent (and in a later draft 16 percent) relative decline after adjusting for firm-level shocks; the 19 percent figure is the unadjusted gap in the latest version and the authors describe all of these as descriptive.

2.2 Exposure is not access

A second body of work measures exposure, meaning how closely the tasks in an occupation match what AI can do. Eloundou et al. (2023) estimated that about 80 percent of the US workforce could have at least 10 percent of their tasks affected by large language models and about 19 percent could have at least half affected, and, unlike earlier waves of automation, exposure rose with wages. The International Labour Organization's refined index finds that about one in four workers worldwide is in an occupation with some exposure, that 3.3 percent of global employment sits in the highest exposure gradient, and that women are more concentrated there (4.7 percent of female employment against 2.4 percent of male employment; 9.6 against 3.5 percent in high-income countries) (Gmyrek et al., 2025). Exposure also tracks national income: 34 percent of employment in high-income countries is exposed against 11 percent in low-income countries. A companion paper finds that conventional exposure measures overstate the impact in developing economies, where workers perform fewer non-routine analytical tasks and digital infrastructure is thinner (Gmyrek, Viollaz and Winkler, 2025). A recent review for the ILO

concludes that exposure measures converge on higher-wage jobs and that the effect on inequality remains inconclusive (del Rio-Chanona et al., 2025).

Exposure describes the technical potential of a job. It says nothing about whether the incumbent has the tool. A well-paid lawyer in a firm that prohibits chatbots is highly exposed and entirely unaided. Conversely, an analyst in a less-exposed occupation may use an embedded AI function every day. The gap between potential and use is where organizations enter.

2.3 Inside the firm: managers, strategy and trust

Organizational economics has long held that information technology pays off only when complementary practices and skills are in place (Bresnahan, Brynjolfsson and Hitt, 2002; Brynjolfsson, Rock and Syverson, 2021). Makridis (2026) applies this to generative AI using the Gallup Workforce Panel. Frequent use (daily or a few times a week) rose from under 10 percent of workers in 2023 to 27.4 percent in the first quarter of 2026. Senior leaders reached the low 40s by late 2025 while individual contributors were in the low 20s, and the income gradient on frequent use widened sharply. Employees who say their organization has communicated a clear AI strategy are about 27 percentage points more likely to use AI frequently (coefficient 0.268, s.e. 0.012) after controls. Frequent use without a clear strategy is associated with a slightly negative engagement effect (− 0.014 SD) and higher burnout (0.115 SD), while the interaction with a clear strategy is positive for engagement (0.213 SD). With person fixed effects, the engagement interaction shrinks to 0.064 SD, which tells us that part of the cross-sectional pattern is selection. Among frequent users, those in high-trust workplaces report 0.61 to 0.72 SD larger perceived productivity gains and tilt their use toward learning.

Gallup's February 2026 survey makes the same point with simple contrasts. Among employees in organizations that make AI tools available, 78 percent of those who strongly agree that their manager actively supports AI use are frequent users, against 44 percent of the rest; the corresponding contrasts are 88 against 55 percent for tools that integrate with existing systems, 72 against 44 percent for organizational support of experimentation, and 68 against 47 percent for clear guidelines. Boston Consulting Group's 2026 survey (n = 11,749) reports that only 36 percent of employees feel they received sufficient training and that only 33 percent of frontline employees say leadership communicates clearly about AI. A Salesforce and YouGov survey found that 76 percent of workers say their preferred AI tools lack access to company data and context. The picture is of access that is nominal for many and substantive for few.

2.4 From use to careers and wages

Three links connect unequal use to unequal careers. The first is skill pricing. PwC's 2026 Global AI Jobs Barometer, based on more than one billion job advertisements, reports that advertised jobs requiring AI skills carry an average wage premium of 62 percent, up from 57 percent a year earlier, and that jobs requiring AI skills grew 69 percent against 9 percent for the whole market. The authors caution that this is a snapshot that does not control for education, experience or location. The second link is entry-level hiring. Brynjolfsson et al. find that the adjustment is occurring mainly through fewer hires of young

workers in occupations where AI automates rather than augments; PwC finds that in the United States entry-level roles that have been 'seniorized', meaning that they now ask for skills traditionally associated with senior staff, grew 35 percent since 2019 while other entry-level roles fell 10 percent. The third link is task design. PwC's classification, adapted from Teeselink and Carey, labels jobs as 'professionalised' when AI takes the basic work and leaves expert tasks to people, and 'democratised' when AI takes the complex work and leaves less expert tasks. About 22 percent of advertised jobs are professionalised and 52 percent democratised, and advertised wages in the former have grown faster since 2021 (37 against 26 percent).

Against this stands a sobering counterpoint. Humlum and Vestergaard (2025) find precise zeros: with administrative data from Denmark, chatbot adoption has had no measurable effect on earnings or recorded hours, ruling out average effects larger than 2 percent two years after the launch of ChatGPT. They find instead that work is being reorganized and that adopters move into higher-paying occupations where chatbots are more relevant. Any claim that unequal AI access already produces unequal pay must therefore be held loosely. The more defensible claim is that the pathways through which it could do so are open and that early indicators point in that direction.

2.5 Propositions

The review suggests five propositions, which structure the empirical work.

P1 (hierarchy). Frequent AI use is higher at higher organizational levels, and the absolute gap between levels has widened as adoption has spread.

P2 (income). Frequent use increases with household income, and the gradient is steeper at the top of the distribution than at the bottom.

P3 (enablement). Organizational conditions that reduce the cost of experimenting (manager support, tool integration, clear rules, experimentation) are associated with large differences in frequent use, and these conditions are themselves unevenly distributed.

P4 (leveling threshold). The equalizing effect observed in experiments survives only if the ratio of access between lower- and higher-skilled workers exceeds the ratio of higher- to lower-skilled gains per user.

P5 (career pathway). Skills that employers price and entry-level opportunities are concentrated where AI use is concentrated, so that unequal use is a plausible precursor of unequal pay and mobility.

3. Data and methods

3.1 Research design

The design is a structured evidence synthesis combined with secondary analysis of published aggregates and a transparent calibration exercise. I did not run a survey, and I do not have microdata.

This is a limitation I accept because the question is time-sensitive and the best available data are held by survey firms and payroll providers that publish aggregates. Where the published information is enough to compute a statistic, I compute it from the published numbers and show the arithmetic. Where it is not, I say so.

3.2 Sources and selection

I searched for peer-reviewed articles, working papers from established economics series (NBER, CESifo, ILO, Stanford Digital Economy Lab, Becker Friedman Institute) and large-sample surveys from organizations that publish sample sizes and survey dates. Inclusion required a sample or dataset described with enough detail to judge its coverage, a date no earlier than 2023, and a result concerning productivity, adoption, employment or wages related to generative AI. Industry sources were used for adoption and access measures where academic data do not exist, and I treat them as lower in evidential weight, particularly when sampling frames are not probability-based. Table 2 lists the quantitative sources used for the new calculations.

Table 2
Quantitative sources used in the calculations

Source	Population / sample	Date	Measure used
Gallup Workforce Panel, Makridis (2026)	US employees aged 20–65; 33,743 workers, at least 78,939 person-waves; probability-based panel	2023 to Q1 2026	Frequent AI use by income bracket, role, industry; regression coefficients
Gallup (2026), workplace survey	23,717 employed US adults; margin of error ± 0.9 points	4–19 Feb 2026	Frequent use by role in organizations that provide AI; enabling conditions
Gallup (2026), AI and productivity	US employees	Q2 2026	Frequent use by leaders (51%), managers (36%), individual contributors (26%); 2023 values 17%, 15%, 9%
BCG, AI at Work (2025; 2026)	10,600+ (2025); 11,749 (2026) leaders, managers and white-collar frontline employees	2025; 2026	Regular use by level; training; leadership communication
Perceptyx via HCAmag (2025); Infosys via CIO Dive (2026)	3,600 + employees (Perceptyx); senior leaders and junior staff (Infosys)	2025; 2026	Workflow integration; feeling equipped
PwC Global AI Jobs Barometer (2026)	More than one billion job advertisements; 27 countries and territories	2018–2025	AI-skill wage premiums, sector dispersion, entry-level skill change
ILO working papers 140 and 166	Task-level scores for 436 ISCO occupations; labor force surveys in 135 countries	2025	Exposure gradients by sex and income group

Note: Percentages are those published by the sources. Gallup and BCG define 'frequent' or 'regular' use differently (daily or a few times a week in both, but with different samples and question wording), so levels are not strictly comparable across them.

3.3 Measures and calculations

Hierarchy gradient (P1). I use Gallup's published frequent-use shares for leaders, managers and individual contributors in 2023 and Q2 2026. I report percentage-point gaps and ratios, because absolute and relative gaps can tell different stories when a technology is spreading.

Cross-survey access ratio (P1, P3). For five surveys I compute the ratio of a frontline or contributor measure to the corresponding leader or executive measure. A ratio of 1 means equality. I report the mean, standard deviation and a t-based 95 percent interval with four degrees of freedom. The five ratios are not independent draws from a population of surveys, they use different outcome definitions, and two (Infosys, BCG 2025) are approximate because one side of the ratio is described only as 'more than' a value. The interval is therefore a summary of dispersion, not a formal estimate of a population parameter.

Income gradient (P2). Gallup reports frequent use for five household-income brackets in the first quarter of 2026: 19 percent (below $48K), about 20 percent ($48K to $89.9K), about 26 percent ($90K to $119.9K), about 33 percent ($120K to $179.9K) and nearly 46 percent ($180K and above). I assign each bracket a midpoint (36, 69, 105, 150 and 210 thousand dollars; the first and last are my assumptions because they are open-ended) and regress frequent use on the base-2 logarithm of the midpoint, so that the slope is the change in percentage points per doubling of income. Because the midpoints are assumed, I report a sensitivity grid (lower midpoint of 30, 36 or 42 and upper midpoint of 190, 210, 250 or 300) and leave-one-out slopes. With five observations, inference is indicative only. Note that the variable is household income, not individual pay, and that income proxies education, occupation and seniority, which also predict use (Makridis reports a college-degree coefficient of about 0.09 in the saturated model).

Enablement contrasts (P3). I use the four published contrasts (strongly agree versus not) and report differences and ratios.

Leveling threshold (P4). Let a_L and a_H be the shares of lower- and higher-skilled workers who use AI regularly, and g_L and g_H the percentage productivity gains per user. Expected gains per worker in each group are a_L·g_L and a_H·g_H. The leveling effect, meaning that the expected gain is at least as large for the lower-skilled group, holds if a_L/a_H ≥ g_H/g_L. I set g_L = 35 percent, the figure reported for the least experienced agents in the customer-support study, and vary g_H over 5, 7, 14 and 21 percent. Brynjolfsson et al. report little or no gain for the most experienced agents and Dell'Acqua et al. report smaller gains for top performers, so g_H is below the average gain of 14 percent in most settings, but I include higher values because gains are task-specific and the values for professional work are not established. I compare the break-even ratio with the observed ratios of access in Section 4. The mapping from 'lower-skilled' to 'individual contributor' or 'lower income' is an assumption, not an observation, and the calculation is illustrative. It is a way of asking how much unequal access the equalizing finding can bear.

Wage pathway (P5). I summarize the dispersion of sector wage premiums reported by PwC (coefficient of variation across eight sectors) and report the published contrasts for entry-level and task-design categories. No inferential test is attempted because the published data are not at the worker level.

All calculations were run in Python (NumPy and SciPy), and the code and outputs are available from the author on request. Figures use the same published values.

4. Results

4.1 The hierarchy gap has widened in absolute terms

Frequent AI use rose for every group (Fig. 1). Among leaders it moved from 17 percent in 2023 to 51 percent in the second quarter of 2026, a threefold increase; among managers from 15 to 36 percent; and among individual contributors from 9 to 26 percent, a 2.9-fold increase. The relative gap barely moved:

leaders were 1.89 times as likely as contributors to be frequent users in 2023 and 1.96 times as likely in 2026. The absolute gap, however, grew from 8 to 25 percentage points, and the manager-contributor gap from 6 to 10 points. This distinction matters. If frequent use is understood as the upper part of an S-shaped diffusion curve, a constant proportional advantage will mechanically produce a rapidly growing absolute gap in the middle of the curve. Whether that gap narrows later depends on whether contributors catch up in the saturation phase, as occurred with earlier office technologies, or whether they remain a fixed fraction behind.

Conditioning on access narrows but does not remove the gap. In Gallup's February 2026 survey, restricted to organizations that make AI tools available, 67 percent of leaders, 52 percent of managers, 50 percent of project managers and 46 percent of individual contributors used AI frequently, a 21-point difference and a contributor-to-leader ratio of 0.69. Non-use is also graded: 8 percent of leaders but 24 percent of individual contributors do not use AI at all, even where it is provided.

The cross-survey comparison in Table 3 shows the same direction in every source I could find, with ratios from 0.51 to 0.73 (mean 0.64, standard deviation 0.09, 95 percent interval 0.52 to 0.75). The ratio is lowest when the outcome is use across all workers (Gallup, Q2 2026: 0.51) and highest when the outcome is perceived integration in workflow (Perceptyx: 0.73).

Table 3
Ratio of frontline or contributor measure to leader or executive measure across surveys

Survey and measure	Frontline / contributor	Leader / executive	Ratio
Gallup, Q2 2026: frequent use, all workers	26%	51%	0.51
Gallup, Feb 2026: frequent use, tools available	46%	67%	0.69
Perceptyx: AI integrated in workflow	58%	79%	0.73
BCG 2025: regular use (leaders 'more than three-quarters', taken as 75%)	51%	75%	0.68
Infosys: 'always equipped' (senior 'more than one-third', taken as 33.3%)	19%	33.3%	0.57
Mean (SD)			0.64 (0.09)

Note: Ratios computed by the author from published percentages. For BCG 2025 and Infosys the leader value is a published lower bound, so the ratio is an upper bound for those rows. The five ratios are not independent and the outcome definitions differ.

The most recent BCG survey complicates the picture and I report it deliberately. BCG's 2026 edition finds that 74 percent of frontline white-collar employees now describe themselves as regular users, a rise of 23 points on the 2025 figure of 51 percent, and speaks of the end of a 'silicon ceiling'. The result is not necessarily inconsistent with the Gallup data, since BCG samples white-collar respondents in 11

countries and regions, defines the frontline as employees with no managerial responsibilities, and may over-represent organizations that have invested in AI. Two readings are possible. Either the gap in simple regular use is closing in the white-collar workforce while it persists elsewhere in the economy and in measures of depth, or survey design explains much of the variation. The same BCG survey reports that only 36 percent of employees felt they had enough training and that only 33 percent of frontline employees said leadership communicated clearly, which suggests that the divide may be moving from whether people use AI to how well they are supported in using it.

4.2 Frequent use climbs steeply with income

Figure 2 shows frequent use by household income. In the first quarter of 2026, workers in households earning $180,000 or more were 2.4 times as likely to be frequent users as workers in households under $48,000 (46 against 19 percent), a gap of 27 points. The log-linear fit implies an increase of 10.1 percentage points for each doubling of income (standard error 2.8, $R^2 = 0.81$, $p = 0.036$ with five points). The rank correlation is perfect (Spearman $\rho = 1$; exact two-sided permutation $p = 0.017$). Sensitivity to the assumed midpoints is modest: across the twelve combinations of lower and upper bounds the slope lies between 8.3 and 11.5 points per doubling. Leave-one-out estimates vary more, from 6.6 (dropping the top bracket) to 15.8 (dropping the lowest), reflecting the convexity visible in the figure: the bottom two brackets are nearly identical (19 and 20 percent) and the rise accelerates at higher incomes. Makridis notes that the occasional-use gradient is much flatter, so higher earners are not simply trying AI more casually but converting to heavy, regular use.

Two cautions apply. Household income is correlated with education, occupation and seniority, so this gradient is not a pure income effect. And the statistic describes a gap in intensity of use, not in access to tools. It is nonetheless informative about who is turning access into routine practice.

Industry and occupation tell a related story. In Q1 2026, about 60 percent of workers in technology and mathematics report frequent use, compared with 43 to 47 percent in finance, insurance, real estate and consulting, about 30 percent in education, 20 to 21 percent in retail and industrial sectors, and about 15 percent in hospitality. The occupational profile is similar: 57 percent frequent use among computer and mathematical occupations, 44 percent in business and finance, and under 15 percent in personal care and transportation (Makridis, 2026). These sectors are also those whose pay and bargaining positions differ most, which is why the divide in use risks reproducing existing inequality rather than creating a new axis.

4.3 Enabling conditions are large and unevenly supplied

Table 4 reports the four contrasts. The differences between employees who strongly agree and those who do not range from 21 to 34 percentage points (mean 29), and the ratios of frequent-use rates from 1.45 to 1.77 (mean 1.62). Manager support has the largest difference (78 against 44 percent). These are bivariate contrasts from a survey, and they are plausibly inflated by reverse causation: people who use AI may be more likely to perceive support. Makridis's panel estimates with person fixed effects reduce the engagement complementarity to about 30 percent of the pooled estimate (0.064 against 0.213 SD), a

rough indication of how much of the cross-sectional association might be sorting. Even so, the survey evidence on training, communication and integration shows that these conditions are scarce in the very group that experiments suggest would benefit most.

Table 4
Frequent AI use by organizational condition (percent), US employees, February 2026

Condition (strongly agree vs. not)	Strongly agree	Not strongly agree	Difference (pp)	Ratio
Manager actively supports team's AI use	78	44	34	1.77
Tools integrate with existing systems	88	55	33	1.60
Organization supports experimenting	72	44	28	1.64
Clear AI guidelines and policies	68	47	21	1.45
Mean			29.0	1.62

Source: Gallup (2026), as reported; differences and ratios computed by the author. Frequent use means a few times a week or more.

4.4 How much unequal access can the leveling result bear?

Table 5 and Fig. 3 implement the threshold calculation. If lower-skilled workers gain 35 percent per user and higher-skilled workers gain 14 percent, the break-even access ratio is 0.40. If higher-skilled workers gain only 7 percent, the break-even ratio is 0.20 and leveling is robust to very unequal access. If they gain 21 percent, the break-even ratio rises to 0.60 and any access ratio below that reverses the ordering.

The observed ratios are 0.41 across income brackets (19 against 46 percent) and 0.51 between contributors and leaders in the second quarter of 2026. They lie on either side of the middle break-even value of 0.40 and below the 0.60 value. In other words, when the gain for higher-skilled workers is roughly the average gain in the customer-support study, the observed income gradient in access is almost exactly enough to offset the leveling effect, and the hierarchy gradient leaves lower-ranked workers modestly ahead (an expected-gain ratio of 1.25 at an access ratio of 0.5). When the experimental pattern of near-zero gains at the top holds, leveling survives comfortably. The calculation does not say which of these worlds we are in. It says that the conclusion is hostage to the very quantity, the distribution of access, that experiments hold fixed and that surveys now show to be unequal.

Table 5
Expected gain ratio (lower- over higher-skilled) by access ratio and gain of higher-skilled workers, with g_L = 35%

Gain of higher-skilled workers (g_H)	a ratio 0.3	0.4	0.5	0.6	0.7
5%	2.10	2.80	3.50	4.20	4.90
7%	1.50	2.00	2.50	3.00	3.50
14%	0.75	1.00	1.25	1.50	1.75
21%	0.50	0.67	0.83	1.00	1.17

Note: Entries are (a_L/a_H)·(g_L/g_H). Values above 1 mean that expected gains per worker favor the lower-skilled group. Observed access ratios: 0.41 (income, Q1 2026) and 0.51 (contributors to leaders, Q2 2026).

4.5 From access to pay and careers: the pathway evidence

The published wage evidence concerns premiums for AI skills, not for AI access, and the two are linked only through learning by doing. Still, the dispersion of the premium is informative about the stakes. PwC reports an average AI-skills premium of 62 percent in 2025, up from 57 percent, but sector premiums range from 16 percent in government to 118 percent in consumer markets (Fig. 4). Across the eight sectors I could document, the mean is 65 percent, the standard deviation is 31 points and the coefficient of variation is 0.47; the highest premium is 7.4 times the lowest. In the United Kingdom the average premium rose from 11 percent in 2024 to 34.2 percent in 2025 according to PwC's UK edition. PwC cautions that these are comparisons between similar advertised roles that do not control for education, experience or location, and premiums on advertised pay can differ from realized pay.

On the career side, the evidence points to the bottom rung of the ladder. In the Brynjolfsson et al. payroll data, young workers (22 to 25) in the most AI-exposed occupations had employment 19 percent below the level implied by less-exposed peers by mid-2026, with no comparable gap for experienced workers, and adjustment occurring mainly through reduced hiring rather than separations. Earlier versions with firm-level controls reported 13 and 16 percent relative declines. The authors treat these as descriptive and note that the pattern attenuates when controlling for education and shows some pre-trends. PwC finds that in the top exposure quartile 52 percent of newly demanded skills are traditionally senior, against 7 percent in the lowest quartile, and that the most exposed entry-level jobs are seven times more likely to demand senior skills. If entry-level jobs disappear or are redefined as senior ones, the apprenticeship through which novices benefit from AI-assisted learning, as in the customer-support study, is also at risk.

Task design provides the third link. If the majority of advertised jobs are of the 'democratised' type, in which AI absorbs the complex work, then the people employed in them are likely to hold the jobs with slower wage growth (26 against 37 percent since 2021 in PwC's data) regardless of how much they use

AI. The question of who gets to do the work that AI cannot is, in the end, a question about organizational design, and about who designs.

Finally, the Danish null deserves emphasis. Humlum and Vestergaard find no aggregate wage effect two years in, and their interpretation is that technological change reshapes work well before it shows up in earnings. This is consistent with the present argument, which is about distributions and pathways rather than averages, but it also cautions against reading the 62 percent premium as a returns-to-use estimate. A premium posted by employers for AI skills may reflect scarcity, selection into technical roles and signaling rather than productivity from use.

5. Discussion

5.1 Three layers of the divide

Putting the evidence together, I find it useful to separate three layers. The first is *permission and provision*: whether a worker is allowed and equipped to use AI at work. Gallup's finding that only 2 to 9 percent of infrequent users and non-users cite lack of access or permission suggests that this layer is not the main barrier among employees whose organizations provide tools, although it may be the main barrier for workers in organizations that do not provide them, who are outside the sample. The second layer is *support and integration*: whether managers encourage use, tools connect to company data, rules are clear and training exists. The 21 to 34 point differences, the 36 percent training figure, and the 76 percent of workers who say their preferred tools lack company data all point to this layer. The third layer is *recognition*: whether skillful use is seen, rewarded and turned into advancement. Here the evidence is thinnest. The AI-skill premium and the seniorization of entry-level jobs are signals, but no data set I could find links an individual's AI practice to promotion or pay in the same firm.

The term 'digital divide' therefore may need a third generation. The first concerned connection, the second skills. A third concerns organizational standing: the position from which a person's use of technology is supported, observed and rewarded. This is close in spirit to complementarity arguments in organizational economics, but it adds a distributional point: complementary assets such as trust, feedback and strategic clarity are themselves unevenly allocated across ranks, so technology can widen gaps even when its direct effect is leveling.

5.2 Reconciling leveling and widening

The apparent conflict between experiments and surveys dissolves once one asks what is held fixed. Experiments randomize access and often supply a tool configured by the researcher or firm, so they measure the effect of use on users. Surveys measure who uses. Production inequality in the economy depends on the product of the two, and the break-even calculation in Section 4.4 shows that the product can go either way. The evidence on novices is encouraging and I take it at face value. What I doubt is that it can be extrapolated to organizations without asking who the novices are, whether they have the tool and whether anybody looks at their output.

There is a second reconciliation. Brynjolfsson et al. explain the leveling effect by the model's capture of best practice. The Dell'Acqua results add that the tool is reliable only inside a frontier that users must learn to see. Learning to see it is a form of expertise that probably accumulates with practice and feedback, which are the same resources that manager support and trust supply in Makridis's analysis (high-trust adopters tilt toward learning use). If so, access, support and judgment are complements, and workers who lack the first two will also lack the third.

5.3 Implications for practice and policy

For employers, the data suggest that the cost of a divide in use is not only a fairness concern. Gallup's contrasts and BCG's training and communication figures imply that firms may be getting less from their AI investment than they think because the people who do routine work do not use the tools well. Concretely, organizations might audit AI use by rank, function and demographic group as they already do for pay; provide licensed tools with company data access to the frontline rather than to leadership first; make AI-related tasks visible in performance reviews; and preserve entry-level roles in which novices can learn alongside the tools. Managers matter: the 34-point difference associated with active manager support is the largest of those in Table 4.

For policy, the pattern suggests three directions. Statistical agencies could add employer-provision and use questions to labor force and establishment surveys so that the divide can be measured in probability samples rather than industry polls; Makridis cautions that online panels are vulnerable to synthetic respondents, which strengthens the case for probability-based frames. Public-sector employers, who in the Gallup data are close to the private sector by 2026, could set standards for equitable provision. And training subsidies could be directed to workers in occupations that are exposed but under-supported rather than to those already in technical roles. The ILO evidence on gender is a reminder that the distribution of exposure can be uneven in ways that interact with access: women are more concentrated in the clerical occupations with the highest exposure, which are also the occupations where the displacement route of AI is most plausible.

5.4 Limitations

The limitations are substantial and I want to state them plainly. First, I use aggregated published figures and no microdata. I cannot control for confounders, and the cross-survey comparison mixes different populations, wording and dates. Second, several inputs are approximate: Gallup's bracket values are read from published text and rounded, the leader values in BCG 2025 and Infosys are lower bounds, and the bracket midpoints are assumptions. Third, the income regression uses five observations and I use it descriptively. Fourth, the break-even calculation depends on gains per user taken from specific tasks in specific firms, and on mapping skill groups to rank and income groups; it is a thought experiment disciplined by numbers, not an estimate. Fifth, the wage evidence concerns advertised premiums from job postings, which can overstate realized pay, and the one study with administrative earnings finds nulls. Sixth, many sources are working papers or industry reports that have not been through peer review, and industry surveys have commercial interests, though the large-sample Gallup panel and the academic

studies reduce this concern. Seventh, the field moves quickly: the BCG 2026 finding that frontline regular use has risen to 74 percent shows how fast the numbers can change, and conclusions about the gap in 2026 may not hold in 2027. Finally, much of the evidence is from the United States; ILO evidence suggests that the pattern in lower-income countries is different, with lower exposure but infrastructure gaps that may prevent gains.

5.5 A research agenda

The decisive evidence would come from linked employer-employee data that record AI use (or licensing) at the worker level alongside pay, promotion and exit. Danish registers linked to adoption surveys already show how this can be done for earnings and hours; the next step is to add rank and within-firm promotion. Randomized roll-outs in which organizations vary who gets access first (frontline against leaders) and what support accompanies it would identify the access channel directly. Diary or log data could measure the depth of use, which is not captured by a frequency question. And studies of recognition, for instance whether managers rate AI-assisted work differently depending on who produced it, would fill the thinnest layer of the divide.

6. Conclusion

Generative AI helps novices more than experts in the settings where it has been tested properly. That finding is real and good news. But organizations are not laboratories. In the surveys I examined, frequent use is higher at the top of the hierarchy and the income distribution, and the absolute gap between leaders and contributors has tripled since 2023 even though the proportional gap is stable. The conditions that turn access into benefit, namely manager support, integrated tools, clear rules and trust, are strongly associated with use and appear to be unevenly supplied. A simple calculation shows that the observed gradients are close to the point at which unequal access cancels the leveling effect, so that whether AI narrows or widens inside-firm inequality depends on parameters that current data cannot pin down. Posted premiums for AI skills and the weakening of entry-level hiring in exposed occupations suggest that the pathway from unequal use to unequal careers is open, even though administrative data have not yet shown a pay effect.

What I take from this is a change of question. Asking whether AI is skill-biased or skill-leveling treats the technology as the actor. The evidence says that organizations are the actors, because they decide who is equipped, supported and recognized. The new divide inside organizations is a managerial and institutional choice, and so it can be measured, audited and changed.

Declarations

Data availability.

All data are drawn from published sources listed in the references. Derived calculations (analysis script and output) are available from the author on request. *Funding.* None declared. *Conflict of interest.* None

declared. *Use of AI tools.* [To be completed by the authors according to the target journal's policy.]

References

1. Autor DH (2015) Why are there still so many jobs? The history and future of workplace automation. J Economic Perspect 29(3):3–30
2. Autor DH, Levy F, Murnane RJ (2003) The skill content of recent technological change: An empirical exploration. Quart J Econ 118(4):1279–1333
3. Boston Consulting Group (2025) *AI at Work: Momentum Builds, but Gaps Remain.* https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain
4. Boston Consulting Group (2026) *AI at Work: Strategy Matters More Than Tools*(fourth edition, June 2026; n = 11,749). https://www.bcg.com/publications/2026/ai-at-work-why-strategy-matters-more-than-tools
5. Bresnahan TF, Brynjolfsson E, Hitt LM (2002) Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. Quart J Econ 117(1):339–376
6. Brynjolfsson E, Chandar B, Chen R (2025) revised August 2026). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.*Stanford Digital Economy Lab. https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine/
7. Brynjolfsson E, Li D, Raymond LR (2025) Generative AI at work. NBER Working Paper 31161
8. Brynjolfsson E, Rock D, Syverson C (2021) The productivity J-curve: How intangibles complement general purpose technologies. Am Economic Journal: Macroeconomics 13(1):333–372
9. CIO Dive (2026) Report on the Infosys study of AI access gaps between senior leaders and junior staff. https://www.ciodive.com/news/AI-access-gap-infosys/826888/
10. Dell'Acqua F, McFowland E, Mollick E, Lifshitz-Assaf H, Kellogg KC, Rajendran S, Krayer L, Candelon F, Lakhani KR (2026) Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. Organ Sci 37(2):403–423
11. del Rio-Chanona M, Ernst E, Merola R, Samaan D et al (2025) Generative AI and the labour market: Review of exposure measures, empirical evidence and research gaps. ILO discussion paper; arXiv:2509.15265
12. Eloundou T, Manning S, Mishkin P, Rock D (2023) GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv:2303.10130.
13. Gallup (2026) What separates AI adopters from holdouts in the workplace. Survey fielded 4–19 February 2026, n = 23,717. https://www.gallup.com/workplace/704252/workplace-separates-adopters-and-holdouts.aspx
14. Gallup (2026) AI in the workplace and productivity (Q2 2026 findings). https://www.gallup.com/workplace/713063/ai-workplace-productivity.aspx

15. Gmyrek P, Berg J, Kamiński K, Konopczyński F, Ładna A (2025) Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140. https://doi.org/10.54394/HETP0387
16. Gmyrek P, Viollaz M, Winkler H (2025) Disruption without dividend? How the digital divide and task differences split GenAI's global impact. ILO Working Paper 166
17. Hargittai E (2002) Second-level digital divide: Differences in people's online skills. First Monday, 7(4)
18. HCAmag (2025) Executives, managers far more likely to use AI than individual contributors: Report (Perceptyx survey of 3,600 + employees). https://www.hcamag.com/ca/specialization/transformation/executives-managers-far-more-likely-to-use-ai-than-individual-contributors-report/553602
19. Humlum A, Vestergaard E (2025) Large language models, small labor market effects. NBER Working Paper 33777; revised (2026) as Still waters, rapid currents: Early labor market transformation under generative AI
20. Katz LF, Murphy KM (1992) Changes in relative wages, 1963–1987: Supply and demand factors. Quart J Econ 107(1):35–78
21. Makridis CA (2026) The organizational transmission of AI: The role of managers on AI adoption and impact. CESifo Working Paper 12373
22. Noy S, Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. Science 381(6654):187–192
23. PwC (2026) *2026 Global AI Jobs Barometer: Two Futures for Jobs in an AI Era*. https://www.pwc.com/gx/en/issues/artificial-intelligence/job-barometer/
24. Salesforce and YouGov (2026) Survey on workers' AI tools and access to company data. https://www.salesforce.com/au/news/?p=13031
25. SHRM (2026) *Navigating AI in the Workplace*(survey of 5,875 US workers, March–April 2026). https://www.shrm.org/topics-tools/research/navigating-ai-in-the-workplace/full-report
26. van Dijk J (2005) The Deepening Divide: Inequality in the Information Society. Sage

Figures

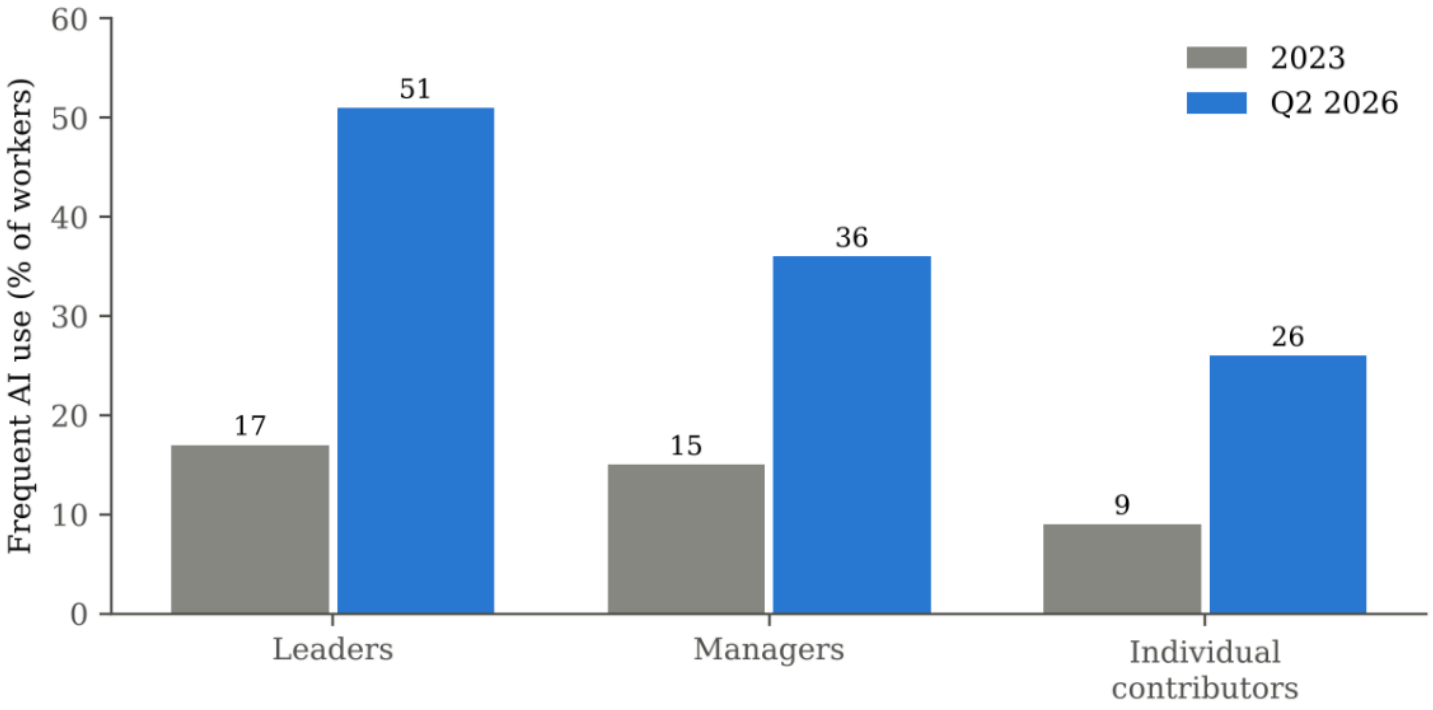

Figure 1

Frequent AI use at work by organizational level, 2023 and Q2 2026 (percent of US workers). Source: Gallup (2026), as reported; author's presentation.

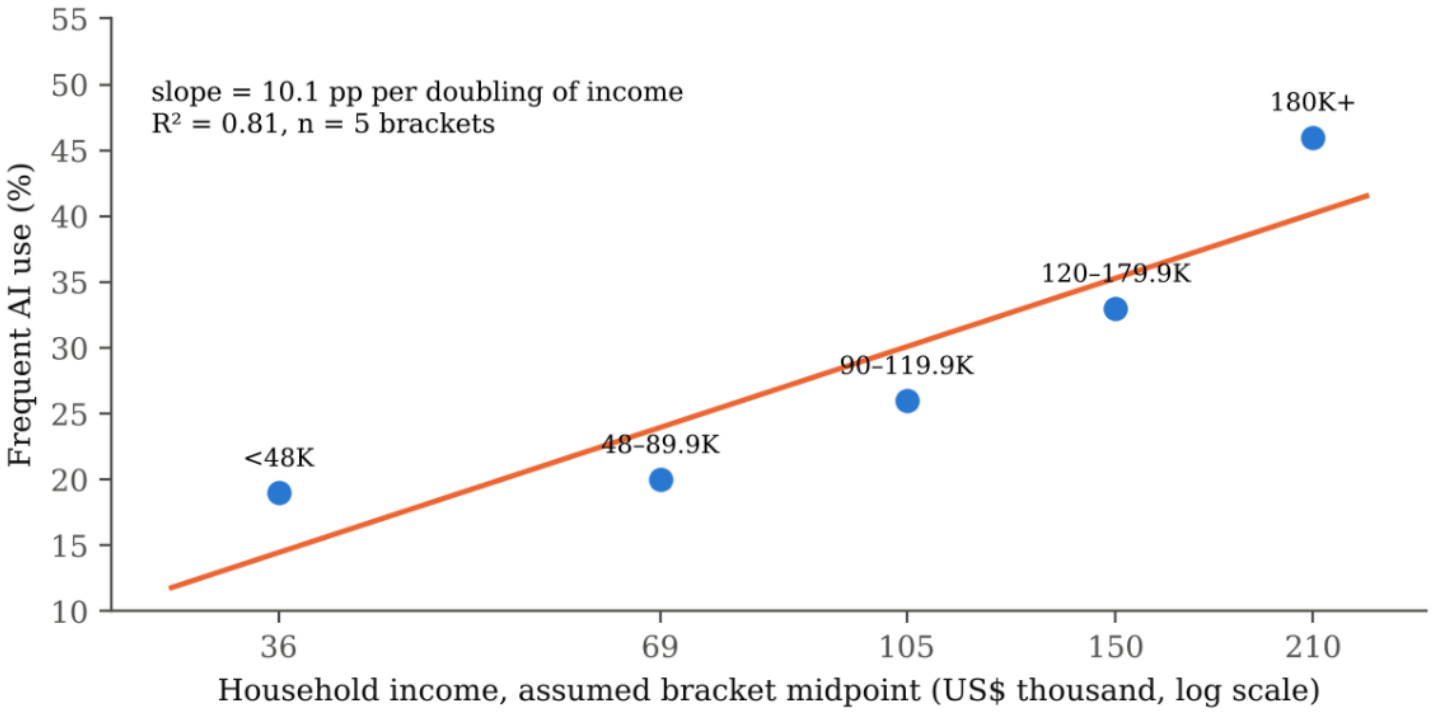

Figure 2

Frequent AI use by household income, Q1 2026. Points are Gallup's bracket values (rounded as reported); the line is an OLS fit on log2 of assumed bracket midpoints. Source: Makridis (2026); author's calculations.

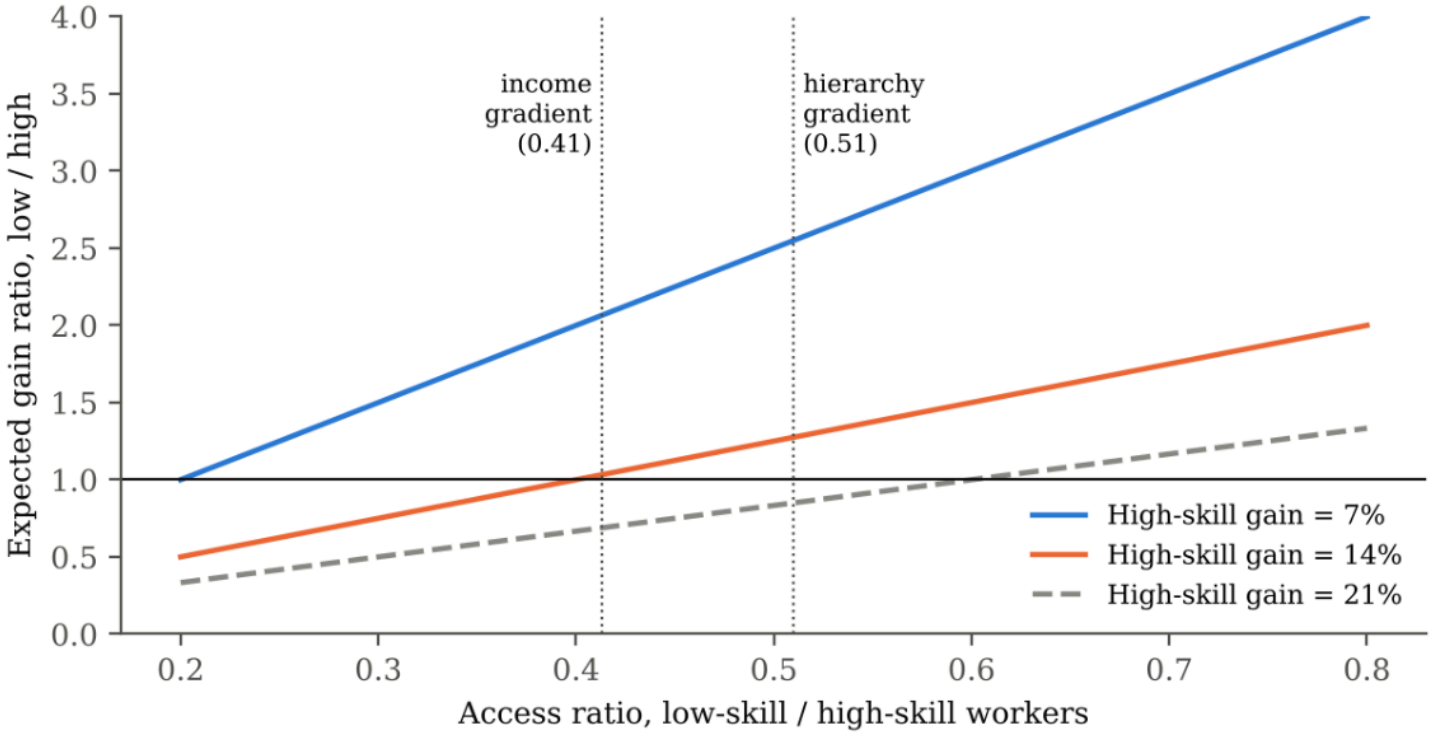

Figure 3

Leveling threshold. The expected gain ratio equals 1 where the access ratio equals g_H/g_L. Dotted lines mark the observed access ratios. Author's calculations.

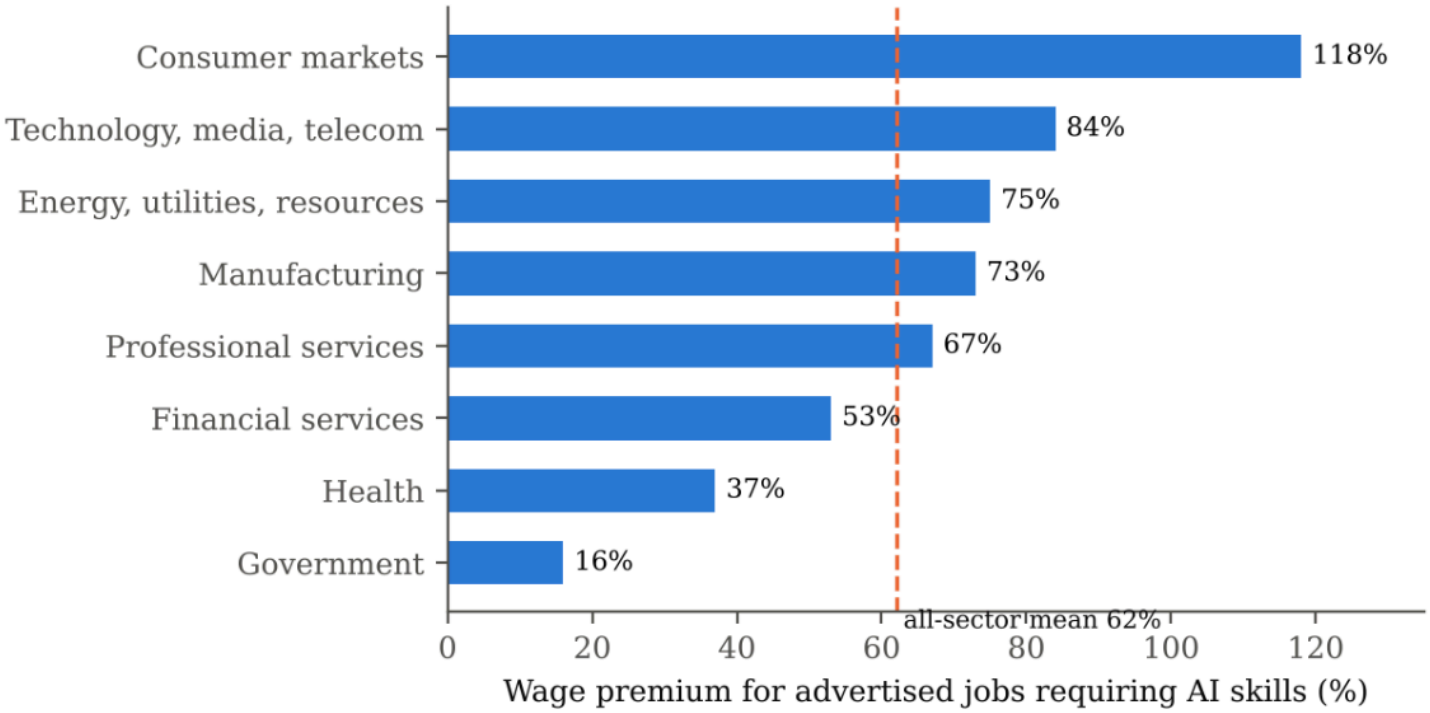

Figure 4

Wage premium for jobs requiring AI skills by sector, 2025. Source: PwC (2026) global findings, as reported; the dashed line marks the 62 percent average. Sector values are those in the global report and may differ from country-specific editions.