AI (at Work) in Finance

PRELIMINARY – DO NOT DISTRIBUTE

Christos A. Makridis * Erik Brynjolfsson† Sophia Kazinnik§

May 7, 2026

Abstract

Using the Gallup Workplace Panel (2023–Q1:2026), we study the diffusion of artificial intelligence in U.S. financial services, focusing on both adoption intensity and task allocation. Frequent AI use in finance rose from 9.9% in 2023 to 46.3% by Q1:2026, compared with 10.8% to 28.8% in the workforce overall. We evaluate three candidate mechanisms behind this acceleration: organizational governance, occupational adaptive capacity, and regulation. Using within-person variation, we find that the strongest overall predictor of frequent AI use is whether workers report that their organization has a clear AI strategy, suggesting governance is a key condition for scaling adoption. Adaptive capacity, especially skill transferability, also predicts more intensive use, but neither governance nor adaptive capacity explain why the share of frequent AI users in finance outpaces other sectors. Regulation does. Higher regulatory restrictions predict lower AI use in finance relative to the rest of the economy, consistent with compliance requirements shaping both when AI can be deployed and which applications firms scale. Finance also shifts AI toward more routine and verifiable tasks and away from uses with higher stakes such as prediction and customer interaction.

Keywords: Artificial Intelligence; Technology Adoption; Financial Services; Regulation; AI Governance; Task Allocation

JEL Codes: G20; J24; O33; D22.

*Corresponding Author: The Gallup Organization, Arizona State University, and Stanford University, Digital Economy Lab (HAI), cmakridi@asu.edu.

†Stanford University, Digital Economy Lab (HAI) and NBER.

‡Stanford University, Digital Economy Lab (HAI).

§These views are our own and do not reflect those of any affiliated institutions.

1 Introduction

AI adoption has varied markedly across the economy. Some sectors have moved relatively quickly, while others remain earlier in the diffusion process, and even among apparent adopters the intensity and organizational depth of use vary substantially (Bick et al., 2024). These differences are important because the returns to AI adoption inevitably vary across sectors: the value of AI depends on the task mix, the quality and standardization of data, the scope for automation or augmentation, and the regulatory and organizational constraints that shape implementation (Bresnahan et al., 2002; Bloom et al., 2012; Makridis, 2025).

Financial services provide a particularly useful setting in this respect. Recent survey evidence places finance and insurance among the more AI active sectors, yet this headline masks substantial heterogeneity in how firms deploy the technology and whether adoption has progressed beyond experimentation. Census data from late 2023 through early 2024 point to relatively high reported AI use in finance and insurance (Census Bureau, 2024), but broader evidence suggests that many organizations remain in the pilot stage rather than full scale deployment. McKinsey's 2025 global survey, for example, finds that although regular AI use is now widespread, most organizations are still experimenting, and only about one third report having begun to scale AI across the business (McKinsey & Company, 2025).

We study this gap between reported adoption and everyday integration using repeated waves of Gallup's Workplace Panel between Q2:2023–Q1:2026. In addition to gathering data on a wide array of topics, the survey measures both how often workers use AI and what they use it for. We use these data to document three main facts. First, finance experiences a sharp rise in regular AI use in 2025: the share of workers using AI every day or a few times a week increases rapidly and moves well above the economy-wide trend. Second, among AI users, finance workers shift toward tasks that are easier to adopt under compliance and oversight, especially routine automation and other applications that are easier to document and audit, and away from higher-stakes uses such as prediction and monitoring. Third, whereas AI use is associated with increases in labor supply in other sectors of the economy, it is associated with a decline in labor supply in financial services, implying that AI helped workers save time in an industry that otherwise focuses much of its time on compliance.

Following Eloundou et al. (2024), we use O*NET task data to construct predicted AI exposure scores for each three digit Standard Occupational Classification (SOC) occupation represented in our finance sample. We merge these scores with the adaptive capacity index of Manning and Aguirre (2026), which captures skill transferability, household wealth, geographic

labor market density, and workforce age composition. This allows us to distinguish between two potential sources of occupational heterogeneity in AI adoption: the extent to which tasks are exposed to AI, and the extent to which workers can adapt when those tasks change. In doing so, we connect the task exposure literature to a broader literature on adjustment costs and skill portability. Within finance, task-level exposure alone does not predict which occupations adopt AI most intensively; most finance occupations are already moderately to highly exposed. Instead, adaptive capacity is the dominant predictor: occupations whose workers have higher skill transferability adopt fastest, while those with lower adaptive capacity lag behind even when their tasks are highly exposed to AI. This occupation-level evidence complements our individual-level findings by showing that, within a high-exposure sector, the binding constraint on AI diffusion is not task suitability alone, but the workforce's capacity to absorb the transition. Governance and compliance help determine where AI can be used in finance, while adaptive capacity helps determine which occupations use it most intensively once it becomes feasible, consistent with Lazear (2009), who argues that workers with broader and more transferable skill portfolios face lower adjustment costs when task demands shift.

These findings fit naturally with what makes finance distinct. AI has been deployed or explored across many functions, including trading, risk management, fraud detection, compliance, and internal operations, but finance is also one of the most heavily regulated parts of the economy. Regulation affects not only whether firms adopt AI, but also where they can use it, how quickly they can deploy it, and what is required to scale it safely. The Financial Stability Board (2024) points out that AI applications in financial services remain concentrated in operational, risk, and compliance functions rather than revenue generating activities, reflecting supervisory expectations and governance constraints. Regulators have also highlighted challenges related to explainability, accountability, and systemic risk.[1] Industry evidence likewise suggests that large firms are using AI extensively in finance and accounting functions for planning, reporting, and internal decision support (KPMG, 2024).

We examine three plausible and potentially overlapping mechanisms that may help explain why finance leads in AI adoption and how that adoption shows up across tasks. First, governance and organizational practices create a governance threshold (Makridis, 2025): firms often need executive oversight and clear internal rules before AI can be used widely, consistent with evidence that CEO-led AI governance strongly predicts impact (McKinsey & Company, 2025). Second, there is an adaptive capacity channel: even when governance

[1] See, for example, Bank for International Settlements, Financial Stability Institute (2025); International Organization of Securities Commissions (2021).

barriers are cleared, adoption may vary across occupations depending on whether workers have the transferable skills and resources to reorganize around AI-augmented workflows. Third, there are regulatory effects: stricter oversight may push AI toward applications that are easier to document and audit, and away from higher-stakes uses, consistent with surveys showing that compliance concerns are a leading barrier to scaling AI in financial services (Census Bureau, 2025; Bank for International Settlements, Financial Stability Institute, 2025).

Exploiting within-person variation between 2023 and 2025, we find support for all three hypotheses, but they operate at different levels and only regulation explains the differential adoption of AI in finance.[2] Governance is the strongest overall predictor of frequent AI use: workers who report that their organization has a clear AI strategy are 11 percentage points more likely to use AI frequently, but this relationship is equally strong inside and outside finance. Adaptive capacity also predicts adoption: workers in occupations with more transferable skills use AI more intensively, but this pattern is not unique to finance. Regulation, by contrast, is the channel that distinguishes finance from the rest of the economy. The interaction between regulatory restriction density and finance is large and negative across all specifications, including models with person fixed effects.

Our paper contributes to a growing literature on AI adoption and the future of work. Recent studies measure AI diffusion using job postings (Babina et al., 2024; Acemoglu et al., 2022), firm disclosures (Cao et al., 2025), and worker surveys (Bick et al., 2024). Experimental and quasi-experimental evidence shows that generative AI can raise productivity in specific tasks such as customer service (Brynjolfsson et al., 2025), writing (Noy and Zhang, 2023), and programming (Peng et al., 2023). Yet we still know relatively little about AI use inside financial organizations at the worker level: how common it is, how intensive it is, what tasks it affects, and how those patterns respond to broader forces such as regulation. Unlike job postings or firm disclosures, the Gallup data capture how employees in finance actually use AI in day-to-day work, including both usage intensity and task allocation. That allows us to observe how AI use takes shape, and how it is governed and constrained in practice.

This perspective also speaks to the literature on AI and financial stability. Financial regulators have increasingly focused on systemic risks from AI adoption, including correlated algorithmic strategies, model monoculture, and concentration in third-party AI providers (Financial Stability Board, 2024; Bank for International Settlements, Financial Stability Institute, 2025). More broadly, Goldstein et al. (2019) argue that fintech can amplify

[2]Table A.1 in the Online Appendix enumerates the various regulatory reforms that have been pursued affecting financial services.

procyclicality and herding, while Boot et al. (2021) emphasize the tension between the efficiency gains from technological automation and the fragility introduced by opacity and interconnectedness in banking. Our evidence on the task composition of AI use speaks directly to these concerns. The finding that finance shifts AI away from predictive and monitoring tasks, and toward more auditable applications, is consistent with firms internalizing, or being required to internalize, the stability risks that regulators emphasize.

Our results also connect to research on organization capital and intangible assets. A large literature argues that firm value increasingly depends on scalable intangible assets such as data, software, business processes, and organizational knowledge (Corrado et al., 2006; Crouzet et al., 2022; Veldkamp, 2023). Organization capital, the productive asset embedded in managerial systems and human coordination, shapes firm value and expected returns (Eisfeldt and Papanikolaou, 2013) and has downstream effects on financing outcomes (Danielova et al., 2023; Hsieh et al., 2025; Falato et al., 2022). A complementary tradition shows that technology adoption depends on managerial and workplace complements (Bresnahan et al., 2002; Bloom et al., 2012; Makridis, 2025). The role of governance in our results is consistent with this literature. AI appears to scale once firms establish internal permissions and clear strategies, and the shift toward auditable workflows suggests that organization capital in finance is evolving to include AI governance infrastructure itself.

2 Data and Measurement

2.1 Gallup Workplace Panel Sampling

We use data from the Gallup Workplace Panel, a nationally representative longitudinal survey of U.S. adults (18 and over) that tracks workplace experiences, management, and well being. Gallup also fields rotating modules on emerging topics, including artificial intelligence. Survey weights ensure the sample reflects the U.S. workforce across key demographics and employment characteristics, and we restrict our sample to individuals between age 20 and 65. Unless noted otherwise, all descriptive statistics and regressions apply these weights.

We link respondents across waves and keep observations where the respondent is employed (full- or part-time) in that survey wave. The resulting panel includes 13,294 unique individuals and 25,261 person-wave observations from 2023-2025.[3] Questions related to AI appear only once per year (in quarter 2) for 2023 and 2024, but starting in Q2:2025, they are asked

[3]In our sample, 81.12% are full-time.

quarterly, which allows us to track AI use at higher frequency during 2025 and on. Figure A.1 in the Online Appendix shows that about half of respondents appear in only one wave, with a mean of 1.90 waves per person. In our main analysis, however, we omit Q1:2026 since we do not have regulatory restrictions data for 2026 yet, which reduces the sample by 6,365 observations. The Gallup Workplace Panel also provides detailed individual covariates, including age, sex, race/ethnicity, education, marital status, children, and job characteristics (full-time status, tenure, income), as well as industry and occupation.

2.2 Industry Classification

To study differences across sectors, we compare finance to non-finance using respondents' self-reported industry. We define finance as both finance and insurance, and define non-finance as all other private sector industries. We exclude public sector respondents (federal and state) to keep the comparison focused on broadly similar institutional settings and incentives.[4]

2.3 Measuring AI Utilization

Gallup measures AI use at work with the following question:

> *"How often do you use artificial intelligence in your role?"*

Before answering, respondents are given the following definition:

> *"Artificial intelligence is when a computer does or enhances a task that is normally performed by a person."*

Responses fall into seven categories: (i) daily, (ii) a few times a week, (iii) a few times a month, (iv) a few times a year, (v) once a year, (vi) less often than once a year, and (vii) never. Following the survey's intended intensity gradient, we group these into three mutually exclusive categories:

- **Frequently:** daily or a few times a week;
- **Occasionally:** a few times a month or a few times a year;
- **Never/rarely:** once a year, less often than once a year, or never.

[4] In industry-by-time descriptive analyses, we also drop industry × year cells with fewer than 200 observations to reduce noise from very small groups.

In our regressions, we include indicators for *Often* and *Sometimes*, with *Never/rarely* as the omitted category. In some specifications, we also use an intensity index that maps the seven categories to a 0–10 scale (e.g., daily=10; a few times a week=8; a few times a month=6; a few times a year=4; once a year=2; less often than once a year=1; never/other=0).

The Gallup Panel also measures *what* AI is used for. AI users are asked whether they use AI for a set of specific activities (e.g., automating basic tasks, consolidating information, generating ideas, interacting with customers, and making predictions). We create activity indicators for each use case and report weighted shares among AI users. In finance versus non-finance comparisons, we report levels (2025) and year-over-year changes (2024–2025), along with 95% confidence intervals computed using survey weights. This design separates *who uses AI* from *how it is used*, which is especially important in regulated settings where governance constraints may shape which applications are feasible even when adoption rises.

2.4 Workplace Practices and AI Strategy

We measure organizational AI governance using an indicator for whether the respondent reports that the organization has communicated a clear AI strategy:

> *"Has your organization communicated a clear plan or strategy for integrating artificial intelligence (AI) technology or tools into current business practices (e.g., to increase productivity, efficiency, and quality)?"*

Similar constructs around "clear AI strategy" and AI governance appear in global surveys such as McKinsey & Company (2025) and Census Bureau (2024), where organizational AI readiness and executive involvement are strongly correlated with reported AI impact. The Gallup measure provides a rare individual level view of these organizational policies.

2.5 Regulatory Restrictions

To measure the regulatory environment faced by different industries, we use the industry-level regulatory restrictions" series from the QuantGov/RegData project (Al-Ubaydli and McLaughlin, 2017). RegData turns the text of the U.S. Code of Federal Regulations (CFR) into annual, industry-specific metrics in two steps. First, it counts the number of legally binding constraint terms (restrictions") in each CFR part, using a standardized dictionary of words and phrases that typically create obligations or prohibitions (e.g., shall," must," may

not," required," "prohibited"). Second, it estimates how relevant each CFR part is to each NAICS industry using supervised machine learning trained on Federal Register notices—the daily publication through which agencies propose and finalize rules—that explicitly reference affected industries. The industry exposure measure combines these components by weighting the restrictions count by predicted industry relevance, yielding an estimate of the number of binding restrictions relevant to each industry in each year.

We use these data to build a sector-by-year panel of regulatory intensity and merge it to respondents by NAICS2 industry and survey year. Our baseline proxy for regulatory intensity is the log of restrictions per 1,000 words. This scaling adjusts for changes in overall regulatory text length and focuses the measure on the density of binding language, which is often more interpretable in proportional terms. Figure 1 plots restrictions per 1,000 words separately for finance/insurance and the rest of the economy across 2020–2025. Both series rise from 2020 to 2024 before declining in 2025. Finance/insurance restrictions increase from roughly 0.321 in 2020 to 0.332 in 2024, then fall to about 0.331 in 2025. The rest of the economy follows a similar trajectory, peaking in 2024 before easing modestly. Notably, the 2025 level in finance remains above its 2020 baseline, indicating that the recent decline partially reverses the prior run-up rather than representing a move to a historically low regulatory environment.

RegData is useful for our setting for two reasons. First, it provides a transparent, replicable, annually updated panel that matches standard economic industry codes, making it straightforward to link to worker outcomes. Second, by separating text volume (word counts), binding language density, and predicted industry relevance, it offers a clear way to compare levels and changes in written regulatory constraints across sectors. A key caveat is that this is a proxy for regulation, not enforcement: it captures the prevalence of binding language and predicted applicability in the CFR, not supervision intensity, litigation risk, or agency enforcement resources. We are, therefore, likely to underestimate the full effect of regulation.

2.6 AI Exposure and Adaptive Capacity

To study heterogeneity in AI adoption across occupations within finance, we construct two additional measures at the three digit SOC level: predicted AI exposure and adaptive capacity. Our finance sample includes approximately 3,900 respondents employed in finance and insurance across 2023–2025, spanning 21 distinct three digit SOC occupation groups.

To measure exposure to generative AI for each occupation, we follow the task level classification framework of Eloundou et al. (2024). The O*NET database catalogs detailed

task statements for each occupation. We use this data to classify each task into one of four exposure categories based on whether an LLM-based tool can reduce the time required to complete the task by at least half without compromising quality:

- **No Exposure (E0)**: An LLM-based tool cannot reduce task completion time by half, or its use diminishes output quality.
- **Direct Exposure (E1)**: Direct access to an LLM alone can halve the time needed.
- **Indirect Exposure (E2)**: An LLM alone does not suffice, but additional software built on top of the LLM could halve the time.
- **Image Exposure (E3)**: An LLM with image-processing capabilities substantially reduces completion time.

We assign numeric scores to each category (E0 = 0, E1 = 0.5, E2 = 0.75, E3 = 1) and aggregate to the occupation level by averaging across all tasks within each eight digit SOC code.[5] We then aggregate from the eight digit to the three digit SOC level using employment weighted averages and merge to the occupations represented in our finance sample, achieving 100% coverage of finance respondents.

We also merge our data with the occupation level adaptive capacity index constructed by Manning and Aguirre (2026), which measures workers' ability to adapt if displaced by AI. The index combines four components:

1. **Skill transferability**: a growth-weighted measure of how portable an occupation's O*NET skills are to other occupations, capturing redeployment options.
2. **Median net liquid wealth**: the median household liquid assets minus unsecured debt from the Survey of Income and Program Participation (SIPP, 2022–2024), capturing the financial cushion available to absorb transitions.
3. **Geographic density**: the expected log population density of an occupation's workforce, capturing access to thick labor markets with more alternative employment.
4. **Older worker fraction**: the share of workers aged 55 and over from the American Community Survey (ACS, 2024), entered with a negative sign so that a younger workforce implies higher adaptive capacity.

[5]Following Eloundou et al. (2024), direct LLM exposure (E1) receives a score of 0.5, indirect exposure via LLM powered applications (E2) receives 0.75, and exposure through image capabilities (E3) receives 1.0.

Building on Eloundou et al. (2024), we also convert the Manning and Aguirre (2026) data, which is defined at the modified SIPP occupation level, to our three digit SOC codes using their SOC mapping and employment weighted aggregation. The merge yields adaptive capacity scores for 20 of 21 finance occupation groups, covering 99.7% of finance respondents.

3 Descriptive Patterns

3.1 Summary Statistics and AI Usage Rates

Table 1 reports descriptive statistics. Finance and insurance workers are more AI-active, better governed, and more highly educated than the rest of the private-sector workforce. The share of frequent AI users in finance (0.38) is roughly 46% higher than in the rest of the economy (0.26), and occasional use follows a similar pattern (0.25 vs. 0.19). Organizational readiness mirrors this gap: 46% of finance workers report that their firm has a clear AI strategy, compared with 24% elsewhere. Finance workers also score higher on both GenAI exposure (0.50 vs. 0.39) and adaptive capacity (0.66 vs. 0.58), meaning their tasks are more amenable to AI augmentation and their occupations are better positioned to absorb the transition. On demographics, finance skews slightly more male (0.54 vs. 0.51), works modestly longer hours (41.67 vs. 39.26 per week), and is substantially more college-educated: 58% hold at least a four-year degree, versus 44% in the rest of the economy, with correspondingly lower shares at the high school or less level (0.17 vs. 0.28). Racial composition and family structure are broadly similar across the two groups. The finance sample is not simply a high-adoption outlier—it is a sector with a distinctive combination of task exposure, organizational governance, and human capital that collectively predisposes it toward intensive AI use.

Figure 2 plots the share of *frequent* AI users, i.e., respondents who report using AI *daily* or *a few times a week*, both in the financial sector and in the overall economy.[6] Two patterns stand out. First, frequent AI use rises in both finance and the economy between 2023 and Q1:2026, but the increase is much steeper in finance. Finance begins slightly below the aggregate in 2023 (9.9% versus 10.8%), pulls ahead by 2024 (17.0% versus 13.0%), and the gap widens steadily thereafter. By Q1:2026, 46.3% of finance respondents are frequent AI users, compared with 28.8% in the full sample, a difference of roughly 17 percentage points.

Second, the acceleration of AI use in finance is concentrated in 2025 and continues into early 2026. In finance, the share of frequent users rises from 17.0% in 2024 to 31.8% by

[6] Figure A.2 in the Online Appendix shows that the shares of workers who use AI only *sometimes* level off.

Q2:2025, and continues to 34.2% in Q3:2025, 40.3% by Q4:2025, and 46.3% by Q1:2026. Over the same period, the economy-wide share increases more gradually, from 13.0% in 2024 to 20.5% in Q2:2025, 22.7% in Q3:2025, 26.2% by Q4:2025, and 28.8% by Q1:2026. Notably, the gap between finance and the broader economy continues to widen into 2026, with no sign of convergence. The fact that growth is concentrated in frequent use, rather than occasional experimentation, points to scaled deployment in everyday workflows, consistent with external survey evidence that finance is an intensive AI adopter (FRBNY, 2025; Census Bureau, 2024).

To put these trends in perspective, Figure 4 subsequently documents the share of frequent AI users across industries (Panel A) and the growth over the past year for each industry (Panel B). Panel A shows substantial cross-sectional dispersion in frequent AI use by Q1:2026. Technology/math stands out as the clear leader, with 61.3% of workers reporting frequent use. Finance and insurance form a strong second tier: 46.3% of finance workers and 45.1% of insurance workers are frequent AI users, well ahead of consulting (41.9%), education (31.0%), arts/entertainment/media (28.4%), legal (27.1%), and health care (27.0%). On the other side, frequent use remains below 20% in sectors such as hospitality/travel/food (14.5%), transportation/logistics (18.2%), and industrial/extraction (20.9%).

Panel B reports the change in frequent AI use between Q2:2025 and Q1:2026. Growth is broad but uneven: finance records the largest gain of any sector (14.5 percentage points), followed by arts/entertainment/media (13.9 pp), insurance (13.7 pp), and health care (12.9 pp). Technology/math, despite its high baseline, also exhibits a sizeable gain (9.8 pp), while consulting (7.3 pp), education (4.9 pp), and community/social services (4.5 pp) expand more modestly, and hospitality/travel/food records the smallest increase (1.8 pp). In this sense, Figure 4 points to two key results: while finance is not the economy-wide leader in frequent AI use, as technology/math remains far ahead at 61.3%, finance and insurance rank second and third across all sectors, and finance records the steepest growth of any industry over the past three quarters, consistent with the view that AI deployment in financial services has moved well beyond early experimentation into routine, intensive use.

A related question is how AI use intensity relates to labor supply. Figure 5 plots mean hours worked per week by AI use category (frequent, occasional, and never/rarely) separately for finance and insurance (Panel A) and the rest of the economy (Panel B). Several patterns stand out. In finance, frequent AI users work substantially more hours than non-users throughout the sample period, starting at roughly 47 hours per week in 2023 and converging toward 43 hours by 2025, while never/rarely users hold steady around 40.5 hours. Occasional users in finance fall between the two, tracking closer to never/rarely users over time. The gap between

frequent and never/rarely users narrows across the sample period but remains meaningful through Q1:2026, consistent with AI being concentrated among workers in high-effort roles rather than substituting for labor input.

Outside finance (Panel B), the pattern is qualitatively similar but compressed: frequent users work modestly more hours than never/rarely users, and the gap is smaller and less stable across periods. Occasional users outside finance actually report higher hours than frequent users in 2023 before converging by 2025. The cross-sector contrast suggests that the association between intensive AI use and longer hours is strongest in finance, where AI appears to complement rather than substitute for work effort, while the relationship is weaker and noisier in the broader economy. In short, frequent AI users in finance worked roughly 9% fewer hours by Q1:2026 than in 2023 (47 to 43 hours per week), while frequent users outside finance worked modestly more (+5%), suggesting that intensive AI adoption is associated with time savings in a high-intensity regulated sector but complements work effort elsewhere.

3.2 Composition of AI Activities

Most measures of AI adoption tell us *whether* firms are using AI, but not what workers actually use it for. Gallup's task-level questions allow us to track the mix of activities where AI shows up in day-to-day work, and how that mix differs across industries.

Panel A of Figure 6 compares activity shares among AI users in 2025. Relative to non-finance, finance users are substantially more likely to report using AI to *automate basic tasks* (47.9% vs. 35.8%) and *consolidate information* (46.0% vs. 44.3%). Finance also reports higher use for *making predictions* (10.4% vs. 8.8%). By contrast, non-finance users report higher shares in several other categories, including *learning new things* (37.3% vs. 34.7%), *generating ideas* (41.3% vs. 39.4%), *collaborating with others* (11.8% vs. 9.6%), and *setup/operate/monitor devices* (8.0% vs. 5.5%). Customer-facing use is moderately lower in finance: *interacting with customers* is 11.2% in finance versus 13.5% outside finance. *Identifying problems* is also somewhat lower in finance (18.7% vs. 20.7%), while *other* use cases are nearly identical across the two groups (10.9% vs. 11.0%).

Panel B of Figure 6 reports percent changes in activity shares from 2024 to 2025 among AI users. Finance exhibits a pronounced reallocation across activities. The share using AI for *automating basic tasks* increases by 15.4%, while *interacting with customers* falls sharply (−58.1%) and *making predictions* declines substantially (−47.0%). Finance also shows large declines in *other* (−31.4%), and more modest reductions in *generate ideas* (−10.3%) and

consolidate information (−7.1%). By contrast, *setup/operate/monitor devices* rises notably in finance (+17.5%), though from a low base. Non-finance displays a markedly different pattern, with large increases in *learning new things* (+42.6%), *identifying problems* (+18.1%), *interacting with customers* (+12.4%), and *automating basic tasks* (+8.5%), alongside declines in *consolidating information* (−24.8%), *generating ideas* (−16.9%), and *making predictions* (−12.9%). This shows that finance's 2024–2025 shift is driven by a strong expansion in routine automation and a sharp contraction in customer interaction and predictive use cases.

These shifts look more like a reallocation in where AI is used in finance than a uniform change in overall adoption. In particular, the steep decline in *interacting with customers* (−58.1%) and *making predictions* (−47.0%) is consistent with firms pulling back from, or centralizing, applications that are more likely to trigger internal controls—because they resemble decision support, client-facing communication, or model-driven outputs—and therefore face greater scrutiny from model risk, audit, and compliance functions. Importantly, these constraints can remain binding even when the broader policy climate becomes more deregulatory, since supervisory expectations in banking and related financial services often rest on durable examination and risk management frameworks rather than frequent statutory change. Thus, finance can scale AI while still narrowing the set of tasks it allows AI to touch.

Nevertheless, automation of basic tasks remains the dominant and growing use case in finance. A natural interpretation is that firms have scaled AI most aggressively in back-office and workflow automation—drafting, retrieval, summarization, and task routing—where controls such as approved workflows, logging, and human review are easier to enforce. This would produce exactly the pattern in Figure 6: less emphasis on prediction and customer interaction, and more on routine internal automation.

These patterns are consistent with survey evidence that governance remains a practical bottleneck for many firms. For example, ABS and McKinsey report that many organizations still lack a clear AI strategy and that compliance and output-quality concerns often limit scaling after deployment (Census Bureau, 2025; McKinsey & Company, 2025). The retrenchment in customer-facing and predictive use cases in finance is consistent with firms responding to these frictions by concentrating AI in lower-scrutiny, back-office workflows, even as overall adoption expands. This dynamic is especially relevant in financial services, where model risk management expectations and supervisory guidance have historically constrained deployment in decision-relevant and client-facing applications, and where that guidance remains unsettled.

3.3 Occupation Level Variation in AI Adoption

We now shift from individual-level patterns to occupation-level variation in AI adoption, asking whether predicted task exposure and adaptive capacity explain how intensively workers use AI across occupations. This analysis aggregates respondents to three-digit SOC occupation codes and merges occupation-level measures of predicted AI exposure and adaptive capacity described in Section 2.6.

Figure 3 compares these relationships for finance and insurance versus the rest of the economy. The top panels plot predicted generative AI exposure against the share of workers using AI frequently, while the bottom panels plot adaptive capacity against the same outcome. Two patterns stand out. First, across the rest of the economy, higher predicted exposure is associated with more frequent AI use. Second, within finance and insurance, the relationship between exposure and frequent AI use is much flatter, while adaptive capacity is a stronger predictor of adoption intensity. These patterns suggest that, within finance, the relevant margin is less whether tasks are technically exposed to AI (as many evidently are) and more whether workers can adapt to AI-augmented workflows.

Most occupations represented in finance are clustered in a moderately to highly AI-exposed part of the occupational distribution, so variation in exposure within the sector is relatively compressed. Within this range, exposure alone does not clearly differentiate adoption. By contrast, the occupation-level evidence points to adaptive capacity—and especially skill transferability—as the more important correlate of adoption within finance, shaping which occupations are better able to absorb the transition to AI-augmented workflows. Despite the weak exposure–adoption link, substantial heterogeneity in AI adoption remains across finance occupations. Computer occupations report high rates of frequent AI use, with 55% of workers using AI frequently (exposure: 0.54, adaptive capacity: 0.92), followed by art and design workers in finance (58%, exposure: 0.36, adaptive capacity: 0.84) and healthcare practitioners in finance (52%, exposure: 0.34, adaptive capacity: 0.83).[7] At the lower end, personal care (11%), financial clerks (12%), and retail sales (13%) report much lower rates of frequent AI use. Notably, financial clerks and secretaries/admin have among the highest exposure scores in our sample (0.37 and 0.63, respectively) but among the lowest adoption rates, consistent with adaptive capacity acting as a binding constraint even when task suitability is high.

The four panels reveal an asymmetry between the rest of the economy and finance in

[7] If one uses the broader healthcare SOC grouping rather than healthcare practitioners, the corresponding exposure is 0.33 rather than 0.34.

how task exposure and adaptive capacity relate to AI adoption. Across the broader economy, generative AI exposure is a strong positive predictor of frequent AI use (Panel A): occupations with higher predicted task exposure adopt more intensively, consistent with the view that technical suitability is the binding constraint when governance barriers are not yet fully cleared. Within finance and insurance, however, this relationship is essentially flat (Panel B)—exposure varies little across occupations in the sector and provides no discriminating power over who actually uses AI. The adaptive capacity panels tell the complementary story. Outside finance, the positive relationship between adaptive capacity and frequent use is present but noisier (Panel C), reflecting the fact that exposure is doing most of the work economy-wide. Within finance, adaptive capacity is a considerably stronger and cleaner predictor (Panel D): occupations whose workers possess more portable skill sets adopt AI more intensively, even though their tasks are similarly exposed to the technology. The contrast across the two columns points to a hierarchy of constraints. In the broader labor market, the primary barrier to AI adoption is task suitability, or whether the technology technically applies to what a worker does. In finance, where many jobs are already exposed to AI, the main difference is not exposure itself but whether workers can adapt to AI-based workflows.

4 Empirical Strategy

The descriptive evidence establishes two facts that require explanation. First, frequent AI use in finance rises sharply in 2025 and continues into Q1:2026, outpacing the rest of the economy by a widening margin. Second, adoption is not uniform across tasks: it expands in routine automation while contracting in prediction and monitoring. These patterns suggest that firms are reallocating AI across workflows, not simply adopting more of it. We now lay out an empirical framework that quantifies the finance-specific acceleration and then unpacks the mechanisms behind it, running a horse-race among different theories.

We begin by estimating an event-style specification to establish the timing and magnitude of the finance–economy divergence in AI adoption. Then, we test three candidate channels (governance, regulation, and occupational adaptive capacity) that may explain both the aggregate acceleration and the reallocation across tasks.

We use the panel structure of the Gallup data in two ways. First, we estimate an event-style specification to document when finance begins to diverge from the rest of the economy in frequent AI use, conditioning on common time effects and observable worker characteristics. This specification is intended to establish the timing and magnitude of the finance-specific

acceleration rather than to isolate within-worker changes. Second, in the mechanism analysis, we exploit the longitudinal structure more directly by estimating models with person fixed effects, so that identification comes from within-worker changes over time rather than from cross-sectional differences in who works in finance.

4.1 Frequency of AI in Finance and the Aggregate

We begin by documenting the relative increase in frequent AI use in finance. To trace the timing and size of this divergence, controlling for demographic factors unlike in Figure 2, we estimate an event-style difference-in-differences specification:

$$\text{FreqAI}_{it} \;=\; \alpha \;+\; \sum_{\tau \neq \tau_0} \beta_\tau (\text{Finance}_i \times \mathbb{1}\{t = \tau\}) + \delta_t + \Gamma X_{it} + \varepsilon_{it}, \tag{1}$$

where FreqAI_{it} equals one if respondent i reports using AI daily or a few times a week in period t, and zero otherwise. Finance_i indicates employment in finance or insurance. τ_0 is the pre-surge baseline period (2023), δ_t absorbs common time shocks, and X_{it} includes demographics and job characteristics (age, sex, race, education, marital status, children, full-time status, tenure, and income). Standard errors are clustered at the respondent level.

The coefficients $\{\beta_\tau\}$ trace out the finance-specific deviation from the economy-wide trend in each period. Under the null of parallel trends, these coefficients should be near zero before the acceleration begins. If the 2025 surge reflects a real shift in deployment rather than differential composition or measurement, β_τ should rise sharply from Q2:2025 onward and continue through Q1:2026. This specification establishes the object to be explained: a finance-specific, intensive-margin acceleration concentrated in a narrow window. Because it conditions on common time effects and observable worker characteristics, it is useful for documenting the timing of the divergence. The mechanism analysis that follows then asks whether governance, regulation, and adaptive capacity can account for that pattern, including in specifications with person fixed effects.

4.2 Disentangling the Mechanisms

Three forces may shape both the pace and type of AI adoption in finance. Each implies distinct testable predictions, and each operates at a different level of variation: governance at the firm–worker level, regulation at the industry–year level, and adaptive capacity at the occupation level. We test them jointly to assess their relative explanatory power.

4.2.1 Governance

In regulated industries, deploying AI in day-to-day workflows typically requires a set of upfront investments: policy review, model risk sign-off, vendor due diligence, logging infrastructure, and workflow redesign. These fixed costs create a governance threshold. Before the threshold is crossed, AI use remains experimental; once it is crossed, the marginal cost of extending AI to additional use cases falls, and frequent use can scale rapidly. This mechanism builds on a long line of evidence that technology delivers its largest gains when paired with complementary organizational investments (Bresnahan et al., 2002; Bloom et al., 2012; Makridis, 2025).

We measure governance using the Gallup item that asks whether the respondent's organization has communicated a clear plan or strategy for integrating AI into business practices. To test whether this organizational gate predicts intensive use, we estimate:

$$\text{FreqAI}_{it} = \alpha + \delta_t + \theta_1 \,\text{Finance}_i + \theta_2 \,\text{Strategy}_{it} + \theta_3(\text{Finance}_i \times \text{Strategy}_{it}) + \Gamma X_{it} + \varepsilon_{it}, \quad (2)$$

where Strategy_{it} is an indicator for whether the respondent reports a clear AI strategy at their firm. A positive θ_2 supports the governance-threshold view: workers at firms that have crossed the governance threshold are more likely to use AI frequently. The interaction θ_3 tests whether governance matters differentially in finance. If governance is a general prerequisite for scaling AI rather than a finance-specific channel, θ_3 should be small.

We estimate (2) with both NAICS2-by-time fixed effects (absorbing sector-level trends) and person-by-time fixed effects (absorbing all time-invariant individual heterogeneity). The person fixed effects specification is demanding: it asks whether, within a given worker, a change in perceived strategy clarity is associated with a transition into frequent AI use.

4.2.2 Regulation

A second mechanism operates at the sector level. Compliance requirements raise the cost of deploying AI in applications that resemble decision-making, surveillance, or high-stakes prediction, because these uses require heavier documentation, validation, and auditability. By contrast, drafting, summarization, and routine task automation are easier to oversee and control. If regulation shapes the feasibility frontier for AI deployment, tighter regulatory environments should be associated with lower overall AI intensity and, within finance specifically, a stronger negative relationship because financial regulation targets precisely the high-stakes applications where AI is hardest to govern. For example, there is emerging

evidence that AI adoption can increase operational risk—fraud, business disruption, and process failures—particularly in firms with weaker internal controls (McLemore and Mihov, 2026). If regulators anticipate these vulnerabilities, the compliance burden for high-stakes AI uses rises further, reinforcing a tilt toward lower-risk deployments.

We measure regulatory intensity using the QuantGov/RegData series (Al-Ubaydli and McLaughlin, 2017), which counts the number of legally binding restriction terms in each part of the Code of Federal Regulations and weights them by predicted relevance to each NAICS industry. Our baseline proxy is the log of restrictions per 1,000 words at the NAICS2-by-year level (see Makridis and McLaughlin (2026) for an explanation). We estimate:

$$\text{FreqAI}_{it} = \alpha + \delta_t + \beta_1 \,\text{Finance}_i + \beta_2 \,\text{Reg}_{s(i),t} + \beta_3 \Big(\text{Finance}_i \times \text{Reg}_{s(i),t}\Big) + \Gamma X_{it} + \varepsilon_{it}, \quad (3)$$

where $\text{Reg}_{s(i),t}$ is the regulatory restrictions measure for respondent i's sector s in year t. A negative β_2 indicates that tighter regulation is associated with lower frequent AI use overall. A negative β_3 indicates that this relationship is amplified in finance, consistent with compliance requirements binding more tightly in a sector where many potential AI applications touch on model risk, client-facing communication, or supervisory expectations.

QuantGov measures the prevalence of binding language in the CFR, not enforcement intensity or litigation risk. It is a proxy for the written regulatory environment, not the full cost of compliance. We treat it as a useful time-varying measure of regulatory constraints at the industry-year level, not as a sufficient statistic for the regulatory burden.

4.2.3 Adaptive Capacity

A third mechanism operates at the occupation level. Even within a sector where governance thresholds have been crossed and regulatory constraints are held constant, AI adoption may vary across occupations depending on the workforce's capacity to absorb the transition. Most finance occupations are already moderately to highly exposed to generative AI (mean exposure score: 0.50 vs. 0.39 in the rest of the economy), so task-level suitability alone provides limited discriminating power within the sector. The binding constraint may instead be whether workers can reorganize around AI-augmented workflows—whether their skills transfer to new task configurations, whether they have the financial resources to absorb adjustment costs, and whether they operate in labor markets thick enough to support reallocation. This builds on Lazear (2009) who argues that workers with broader and more transferable skill portfolios

face lower adjustment costs when task demands shift.

We measure adaptive capacity using the occupation-level index constructed by Manning and Aguirre (2026), which combines four components: skill transferability (a growth-weighted measure of how portable an occupation's O*NET skills are to other occupations), median net liquid wealth (from the Survey of Income and Program Participation), geographic labor market density (expected log population density of an occupation's workforce), and the share of workers aged 55 and over (entered negatively, so that a younger workforce implies higher capacity). We merge this index to respondents at the three digit SOC level. To separate adaptive capacity from raw task exposure, we also construct predicted AI exposure scores following the task-level classification framework of Eloundou et al. (2024), which classifies each O*NET task by whether an LLM-based tool can reduce completion time by at least half without compromising quality. We estimate:

$$\text{FreqAI}_{it} \;=\; \alpha + \delta_t + \lambda_1 \,\text{Finance}_i + \lambda_2 \,\text{AC}_{o(i)} + \lambda_3 \Big(\text{Finance}_i \times \text{AC}_{o(i)}\Big) + \Gamma X_{it} + \varepsilon_{it}, \qquad (4)$$

where $\text{AC}_{o(i)}$ is the adaptive capacity index for respondent i's occupation o. A positive λ_2 indicates that workers in occupations with higher adaptive capacity adopt AI more intensively, conditional on demographics and job characteristics. A positive λ_3 indicates that this relationship is stronger in finance—consistent with adaptive capacity acting as the binding constraint in a sector where task exposure is uniformly high and the relevant margin is whether the workforce can adjust to changing task demands.

To identify which dimension of adaptive capacity drives the result within finance, we also estimate a decomposed specification at the occupation level:

$$\text{AIuse}_o \;=\; \alpha + \gamma_1 \,\text{Exposure}_o + \gamma_2 \,\text{Transferability}_o + \gamma_3 \,\text{Wealth}_o + \gamma_4 \,\text{Density}_o + \gamma_5 \,\text{Age}_o + \varepsilon_o, \quad (5)$$

where AIuse_o is the share of frequent AI users in three digit SOC occupation o among finance and insurance workers in 2025, and the right-hand-side variables are the individual components of the adaptive capacity index. We estimate (5) by weighted least squares, using the number of finance respondents in each occupation as weights. This decomposition distinguishes whether the relevant margin within finance is skill portability, financial resources, labor market thickness, or workforce demographics.

Table A.2 decomposes the adaptive capacity index into its four components. Within finance, skill transferability is the only individually significant component ($\beta = 0.081$, $p = 0.04$):

occupations whose O*NET skill profiles are more portable to other occupations adopt AI more intensively, even conditional on task exposure. By contrast, median net liquid wealth ($p = 0.21$), geographic density ($p = 0.24$), and the older worker fraction ($p = 0.92$) are not statistically significant within finance. This pattern differs from what one would find across the full economy, and the reason is intuitive: within finance, variation in household wealth and workforce age is much smaller than across all occupations. The binding margin within the sector is skill portability, whether workers can redeploy their expertise to AI-augmented workflows, rather than financial resources or demographics.[8]

Why not use exposure to generative AI instead? Table A.3 in the Online Appendix reports weighted least squares cross-sectional regressions for 2025. Columns 1 and 4 include only predicted exposure; columns 2 and 5 add the composite adaptive capacity index; columns 3 and 6 add the exposure $\times$ AC interaction. First, within finance, predicted AI exposure alone has essentially no explanatory power ($R^2 < 0.01$ for the share of frequent users). This is consistent with the descriptive evidence in Section 3.3: finance occupations occupy a relatively compressed range of exposure scores, so task suitability alone does not differentiate adoption. Second, adaptive capacity is the dominant predictor. Adding AC to the model raises R^2 from near zero to 0.66 for the share of frequent users and 0.58 for usage frequency. A one standard deviation increase in adaptive capacity is associated with a 9.8 percentage point higher share of frequent AI users ($p < 0.001$) and 0.96 points higher usage frequency ($p < 0.001$). By contrast, we do not find evidence that the association between exposure and adoption varies systematically with adaptive capacity, although the small sample limits power to detect interaction effects, so we focus on results using adaptive capacity instead of exposure in our study of the surge in AI use in finance in the analysis that follows.

[8]We probe the robustness of these results in Appendix C.1 using two alternative predictors. First, BLS median wages correlate highly with adaptive capacity within finance ($r = 0.82$) and predict adoption equally well when used alone. However, when both are included, neither is individually significant, reflecting the collinearity between wages and the components of AC within this narrow cross-section. Second, an index of AI-complementary skills from O*NET (programming, critical thinking, complex problem solving, mathematics) also predicts adoption strongly in isolation, but becomes insignificant when AC is added, while AC remains significant ($p = 0.02$). This suggests that adaptive capacity captures variation beyond narrowly technical skills, consistent with skill *transferability* mattering more than any single skill dimension.

4.2.4 Joint Specification

To assess the relative explanatory power of all three channels simultaneously, we estimate a saturated model:

$$\begin{aligned}\text{FreqAI}_{it} &= \alpha + \delta_t + \phi_1 \,\text{Finance}_i + \phi_2 \,\text{Strategy}_{it} + \phi_3(\text{Finance}_i \times \text{Strategy}_{it}) \\ &\quad + \phi_4 \,\text{Reg}_{s(i),t} + \phi_5\Big(\text{Finance}_i \times \text{Reg}_{s(i),t}\Big) \\ &\quad + \phi_6 \,\text{AC}_{o(i)} + \phi_7\Big(\text{Finance}_i \times \text{AC}_{o(i)}\Big) + \Gamma X_{it} + \varepsilon_{it},\end{aligned} \tag{6}$$

where $\text{AC}_{o(i)}$ is the adaptive capacity index for respondent i's three digit SOC occupation, constructed by Manning and Aguirre (2026) from skill transferability, median net liquid wealth, geographic labor market density, and workforce age composition. This specification nests all three channels and asks whether each explains a distinct component of the variation in frequent AI use. We estimate (6) with both NAICS2 and person fixed effects.

The predictions are as follows. If governance is a broad prerequisite for scaling AI, ϕ_2 should be large and positive regardless of sector. If regulation acts as a finance-specific constraint, ϕ_5 should be negative and significant: tighter regulatory environments should depress frequent AI use disproportionately in finance. If adaptive capacity determines which workers can absorb the transition to AI-augmented workflows, ϕ_6 should be positive—workers in occupations with more portable skill sets and stronger financial cushions adopt more intensively. The interaction ϕ_7 tests whether workforce-readiness matters differently in finance where most occupations are already moderately to highly exposed and the relevant constraint may be the capacity to reorganize around AI rather than the technical applicability.

5 Results

The event-study estimates confirm a sharp, finance-specific acceleration in frequent AI use that begins in 2025 and widens through Q1:2026. The $\hat{\beta}_\tau$ coefficients from Equation (1) are plotted in Figure 7. The 2024 coefficient is close to zero and statistically indistinguishable from the 2023 baseline, indicating that finance and the rest of the economy were on similar trajectories through 2024 after conditioning on demographics.

The divergence begins in Q2:2025, when the finance-specific premium reaches roughly 8 percentage points, and continues to widen: by Q1:2026, finance workers are approximately 14 percentage points more likely to use AI frequently than observably similar workers in other sectors, relative to the 2023 baseline. The monotonic increase across the four quarters of

2025 and into early 2026 is consistent with scaled deployment, and the tightening confidence intervals in later periods reflect the larger quarterly samples available from Q2:2025 onward. These estimates control for age, sex, race, education, marital status, children, and full-time status, so the acceleration is not driven by compositional differences between finance and the rest of the economy along observable dimensions. The pattern matches the raw trends in Figure 4 but puts them on a more rigorous footing: the finance surge survives conditioning on worker characteristics and common time effects.

Table 2 reports our main results. Columns 1–3 (and 5–7) estimate each channel separately, while columns 4 (and 8) include all channels together in a fully specified model. All columns estimate on the same sample, so coefficient movements across specifications reflect the addition of controls rather than compositional changes in the estimation sample.

The most pronounced finding is that internal governance is strongly related to frequent AI use. Across specifications, workers who report that their organization has a clear AI strategy are approximately 29 percentage points more likely to use AI frequently in the NAICS2/time models (columns 1 and 4) and 11 percentage points more likely in the person fixed-effects models (columns 5 and 8). The fact that the estimate survives person fixed effects implies that, within a given worker, changes in perceived strategy clarity are closely tied to transitions into frequent AI use, consistent with the evidence of complementarities in Makridis (2025). However, the interaction between strategy clarity and finance/insurance is small and statistically indistinguishable from zero across all specifications, suggesting that governance acts as a general gate for scaling AI rather than a finance-specific channel.

The second result is that adaptive capacity independently predicts frequent AI use. Workers in occupations with higher adaptive capacity—more portable skill sets, greater financial resources, thicker local labor markets, and younger workforces—are significantly more likely to use AI frequently. The coefficient is 0.065 in the NAICS2/time specification (column 2, $p < 0.01$) and 0.024 in the person fixed-effects model (column 6, $p < 0.05$), and it remains stable in the saturated models (0.054 in column 4; 0.025 in column 8). The interaction with finance/insurance is close to zero and insignificant in all specifications, indicating that adaptive capacity predicts AI adoption broadly across the economy rather than differentially in finance. This is consistent with the view that, within any given sector, the workforce's capacity to reorganize around AI-augmented tasks matters for who adopts, even after conditioning on governance and regulatory conditions. One reason for less economically and statistically significant coefficient estimates stems from only having occupational granularity at the three digit SOC level, bundling many sub-occupations.

Third, the regulatory environment is strongly related to intensive AI adoption in finance. Regulatory restrictions are negatively associated with frequent AI use overall, and the interaction with finance/insurance is large and negative. In the NAICS2/time specifications, higher restriction intensity predicts substantially lower frequent AI use, with a much stronger negative relationship in finance/insurance: the interaction coefficient is -54.8 in column 3 ($p < 0.01$) and -44.8 in the saturated model of column 4 ($p < 0.05$). This finance-specific regulatory penalty persists in the person fixed-effects models, where the interaction reaches -46.0 in column 7 ($p < 0.10$) and -47.0 in column 8 ($p < 0.10$), indicating that the result is not driven by time-invariant differences across workers who sort into finance.[9]

Are these estimates plausible? Figure 2 shows that the share of frequent AI users in finance rose from roughly 17% in 2024 to 46% by Q1:2026, a gain of approximately 29 percentage points. Furthermore, Figure 1 shows that regulatory restrictions per 1,000 words in finance fell from approximately 9.65 in 2024 to 9.55 in 2025, a decline of about 1.0%.[10] Using the saturated person fixed-effects estimate from column 8—where the finance $\times$ regulation interaction is -47.0—a 1.0% decline in restriction density predicts an increase in frequent AI use in finance of $-47.0 \times (-0.01) \approx 0.47$ percentage points. That is, the observed easing in regulatory text density can account for roughly 1.6% of the 29 percentage point increase in frequent AI use in finance over this period ($0.47/29 \approx 0.016$). The regulation channel is statistically significant and directionally consistent with the hypothesis that deregulation facilitated AI deployment, but the magnitude implied by the 2024–2025 change in restriction density alone is modest. In particular, the decline in regulatory restricts per 1,000 words was small and may overlook declines in enforcement actions too. As a result, this calculation should be interpreted as a simple diagnostic for a lower bound and nothing more.

The results support a hierarchy-of-constraints interpretation. Internal governance is a broad prerequisite for scaling AI use: it predicts frequent adoption across all sectors with roughly equal force. Occupational adaptive capacity determines which workers can absorb the transition to AI-augmented workflows regardless of sector. Regulation, by contrast, is the distinctly finance-specific channel: cross-industry differences in regulatory stringency help explain why finance is especially sensitive to the external environment, even with worker fixed effects, common time shocks, and controls for governance and workforce adaptability.

Finally, as a placebo check, we create an alternative indicator for finance using occupational classifications. If our results reflect heterogeneity in exposure to generative AI (e.g., based

[9] Table 3 replicates these results using a more continuous AI utilization index that uses the full set of potential responses in the survey.

[10] $\ln(9.65) - \ln(9.55) \approx 0.0104$, or roughly 1.0%.

on (Eloundou et al., 2024)), rather than industry-specific regulation, then we should see effects using this classification too. However, Table 4 shows that no such interaction effects between occupation-level finance respondents and regulatory restrictions exist and, in the fully saturated specification, no association exists on the direct effects of regulatory restrictions either.[11]

6 Conclusion

AI is spreading rapidly in financial services, but its diffusion depends not only on technical feasibility, but also on organizational readiness, workforce adaptability, and regulation. Using nationally representative panel data from the Gallup Workplace Panel, we show that frequent AI use in finance rose sharply in 2025 and continued into the first quarter of 2026, indicating that AI adoption in the sector has moved beyond experimentation and into routine use.

We highlight three mechanisms behind this pattern. First, governance acts as a threshold for adoption. Workers are much more likely to use AI frequently when their organization has communicated a clear AI strategy, but this relationship is not unique to finance. Governance appears to be a general condition for scaling AI from experimentation to everyday use.

Second, adaptive capacity helps explain which occupations adopt AI most intensively once AI becomes feasible. Within finance, task exposure alone does not explain much of the variation in adoption, likely because most occupations are already moderately to highly exposed. What matters more is whether workers can adjust to AI-supported workflows. Occupations with more transferable skill portfolios adopt AI more intensively, both within and outside finance.

Third, regulation shapes both the pace and composition of AI use in finance. Regulatory restrictions are associated with lower frequent AI adoption in finance relative to other sectors, even after accounting for governance and adaptive capacity. The descriptive evidence also suggests that finance concentrates AI in applications that are easier to document, monitor, and audit, while use in more predictive or customer-facing settings remains more limited.

These results suggest a hierarchy of constraints on AI diffusion in finance. Governance helps determine whether organizations can scale AI at all. Adaptive capacity helps determine which workers and occupations can adjust once deployment becomes feasible. Regulation then influences both how intensively AI is used and where firms are willing to deploy it. But

[11]The occupation-based finance indicator is defined at the SOC3 level and flags respondents in (i) business and financial operations occupations (SOC3 13-1xx, 13-2xx, 13-3xx), (ii) financial services sales occupations (SOC3 41-3xx), and (iii) financial clerical and records occupations (SOC3 43-3xx).

the key barriers to AI adoption in high-stakes sectors may lie less in technical feasibility alone than in the organizational and institutional conditions required for safe deployment.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. Working Paper 28257, NBER.

Al-Ubaydli, O. and McLaughlin, P. A. (2017). RegData: A Numerical Database on Industry-Specific Regulations for All U.S. Industries and Federal Regulations, 1997–2012. *Regulation & Governance*, 11(1):109–123.

Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2024). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics*, 151:103745.

Bank for International Settlements, Financial Stability Institute (2025). Managing explanations: How regulators can address ai explainability. Fsi paper, Bank for International Settlements, Basel.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. Working Paper 32966, National Bureau of Economic Research.

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.

Boot, A., Hoffmann, P., Laeven, L., and Ratnovski, L. (2021). Fintech: What's old, what's new? *Journal of Financial Stability*, 53:100836.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1):339–376.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2):1–49. Originally NBER WP 31161, 2023.

Cao, Y., Liu, M., Qiu, J., and Zhao, R. (2025). Determinants and Information Content of Voluntary Artificial Intelligence Disclosure. Working paper.

Census Bureau (2024). How Many U.S. Businesses Use AI? https://www.census.gov/library/stories/2024/11/how-many-us-businesses-use-ai.html. Business Trends and Outlook Survey, America Counts Brief, November 2024.

Census Bureau (2025). 2023 Annual Business Survey: Tech Adoption Supplement. https://www.census.gov/programs-surveys/abs.html. America Counts Brief, September 2025.

Corrado, C., Hulten, C., and Sichel, D. (2006). Intangible capital and economic growth. Working Paper 11948, National Bureau of Economic Research.

Crouzet, N., Eberly, J. C., Eisfeldt, A. L., and Papanikolaou, D. (2022). The economics of intangible capital. *Journal of Economic Perspectives*, 36(3):29–52.

Danielova, A., Francis, B. B., Teng, H., and Wu, Q. (2023). The effect of organization capital on the cost of bank loans. *Journal of Financial and Quantitative Analysis*, 58(6):2579–2616.

Eisfeldt, A. L. and Papanikolaou, D. (2013). Organization capital and the cross-section of expected returns. *The Journal of Finance*, 68(4):1365–1406.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Falato, A., Kadyrzhanova, D., Sim, J. W., and Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the u.s. corporate savings glut. *The Journal of Finance*, 77(5):2799–2852.

Financial Stability Board (2024). The Financial Stability Implications of Artificial Intelligence. Report, Financial Stability Board. Nov. 14, 2024.

FRBNY (2025). New York Fed Regional Business Survey. https://libertystreeteconomics.newyorkfed.org/2025/08/ai-in-business-survey/. Liberty Street Economics Blog, August 2025.

Goldstein, I., Jiang, W., and Karolyi, G. A. (2019). To FinTech and beyond. *Review of Financial Studies*, 32(5):1647–1661.

Hsieh, A., Lee, C.-F., and Yu, H.-C. (2025). Organizational capital and debt pricing: evidence from bank loans and public bonds. *Review of Quantitative Finance and Accounting*, 65:1671–1708.

International Organization of Securities Commissions (2021). The use of artificial intelligence and machine learning by market intermediaries and asset managers. Report, International Organization of Securities Commissions, Madrid.

KPMG (2024). AI in Finance: From Automation to Transformation. Industry report, KPMG International, London.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of Political Economy*, 117(5):914–940.

Makridis, C. A. (2025). Organizational transmission of ai: Managerial influence on generative ai adoption. *CESifo working paper*.

Makridis, C. A. and McLaughlin, P. (2026). Regulatory stringency and industry performance: Evidence from regdata u.s. 6.0. *Journal of Regulatory Economics, invited special issue*.

Manning, S. J. and Aguirre, T. (2026). How adaptable are american workers to ai-induced job displacement? Technical report, National Bureau of Economic Research.

McKinsey & Company (2025). The State of AI in 2025. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2025. McKinsey Global AI Survey.

McLemore, P. and Mihov, A. (2026). Ai and operational losses: Evidence from us bank holding companies. *Review of Corporate Finance Studies*, page cfag003.

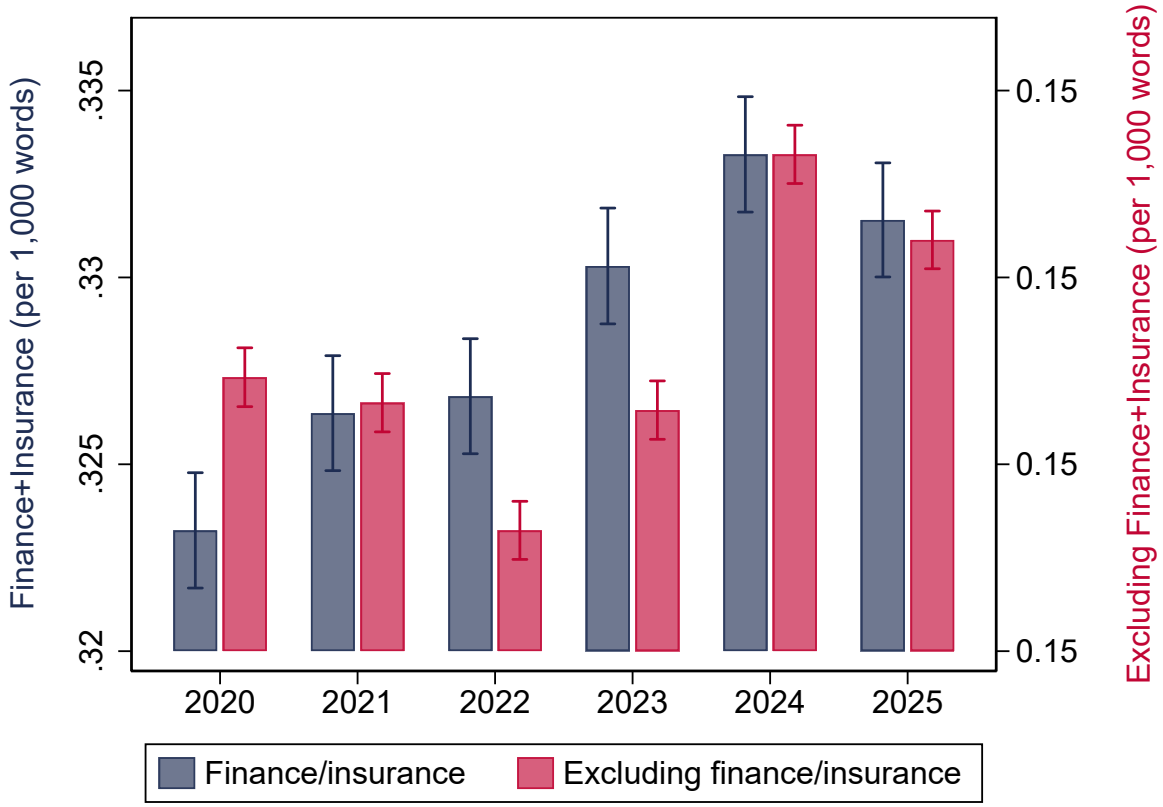
Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. Technical Report 2302.06590, arXiv.

Veldkamp, L. (2023). Valuing data as an asset. *Review of Finance*, 27(5):1545–1562.

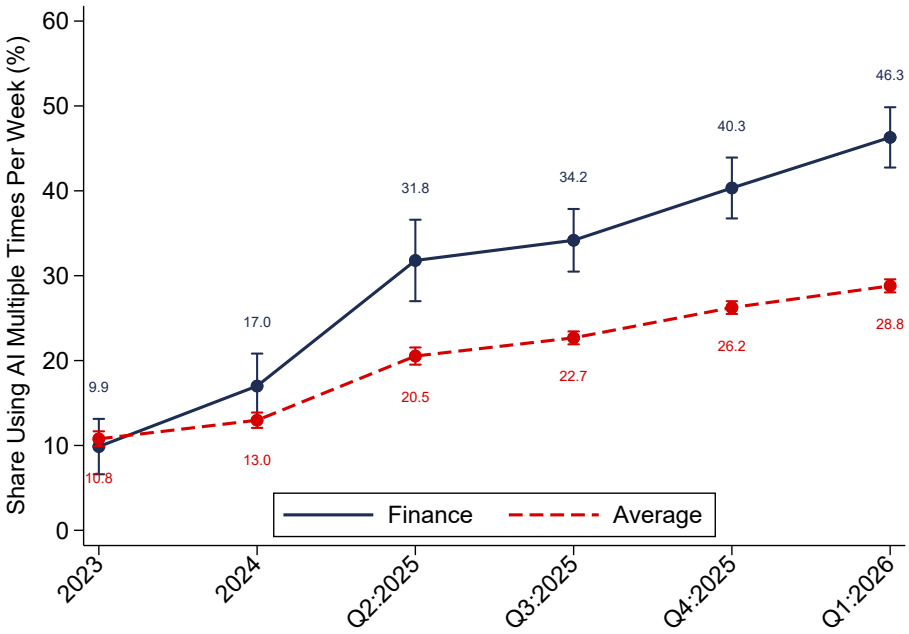
Tables and Figures

Figure 1: Time Series of Regulatory Restrictions



Notes.—Sources: QuantGov from Al-Ubaydli and McLaughlin (2017). The figure plots the regulatory restrictions per 1,000 words from 2020 to 2025, for finance/insurance and the rest of the economy (excluding finance/insurance).

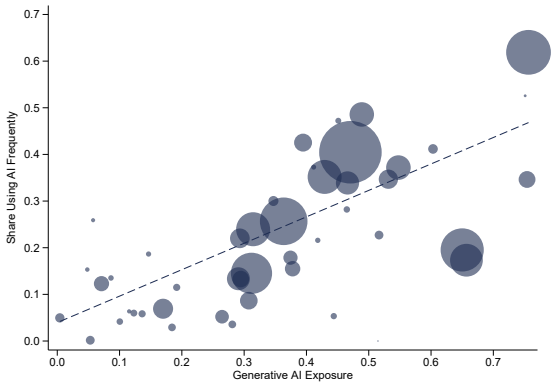
Figure 2: Frequent AI Users in Finance and the Aggregate Economy



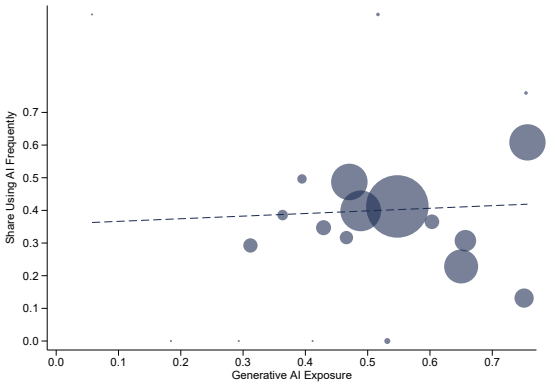
Notes.—Source: Gallup Workplace Panel, 2023–Q1:2026. The sample is restricted to respondents aged 20–65. The figure plots the share of frequent AI users, defined as those who report using AI "multiple times a week" or "daily" to the question "How often do you use artificial intelligence in your role?", separately for financial services workers and the full sample. Other potential answers include: a few times a month, a few times a year, once a year, less often than once a year, and never. Before asking this question, respondents are given the following definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." All series are aggregated using sample weights. Bands denote 95% confidence intervals.

Figure 3: Generative AI Exposure, Adaptive Capacity, and Frequent AI Use

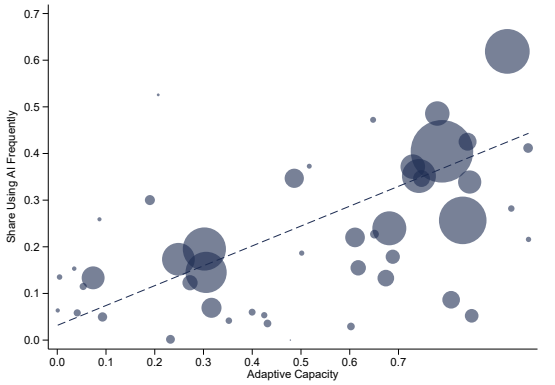
(a) Panel A: Exposure, Rest of Economy (RoE)



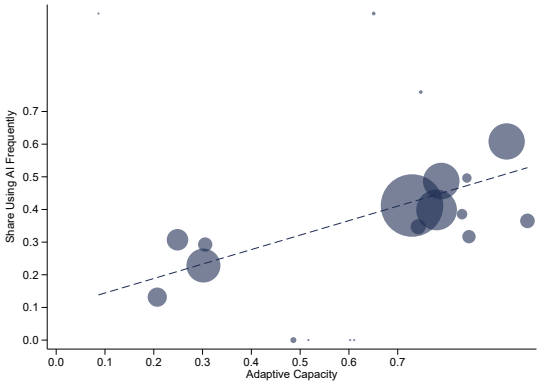
(b) Panel B: Exposure, Finance (FIN)



(c) Panel C: Adaptive Capacity, RoE

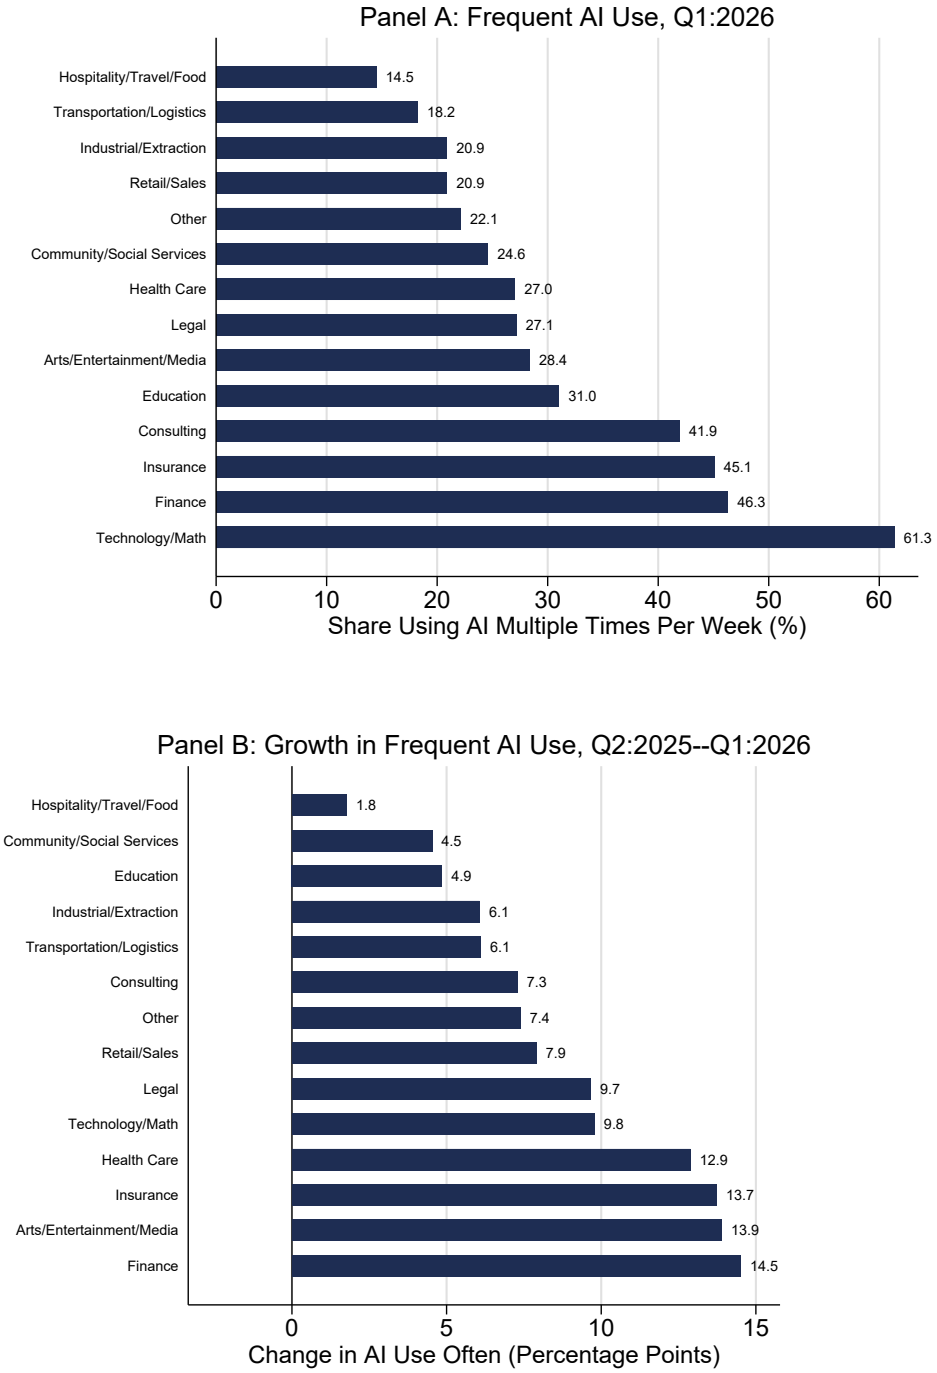


(d) Panel D: Adaptive Capacity, FIN



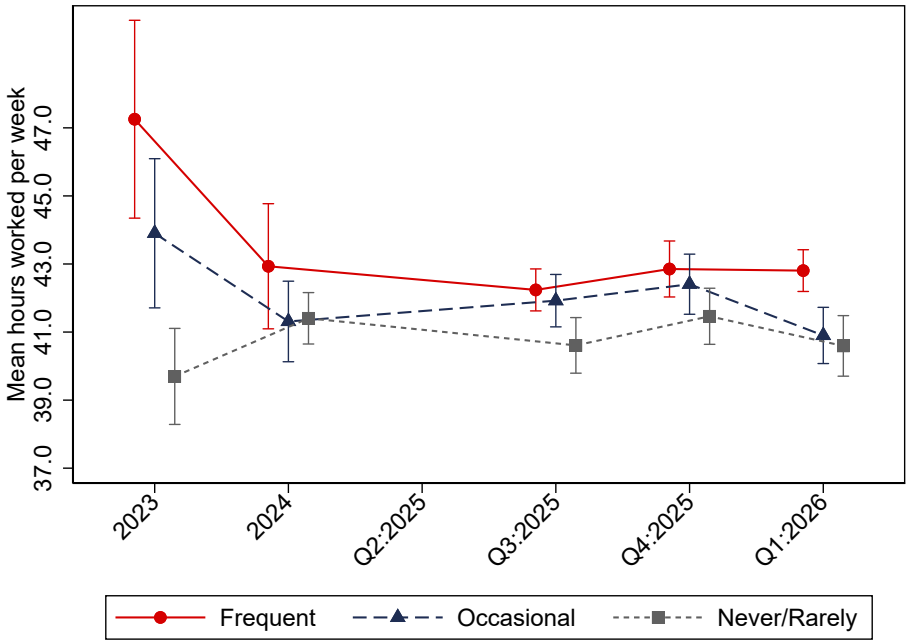
Notes.—Source: Gallup Workplace Panel (2025–Q1:2026). Each bubble represents a three digit SOC occupation group; bubble size is proportional to the number of survey respondents. Lines are weighted least squares fits. Frequent AI use is defined as using AI daily or a few times a week, based on respondents' answer to how often they use artificial intelligence in their role, where artificial intelligence is defined as when a computer does or enhances a task that is normally performed by a person. GenAI Exposure is the task-based predicted exposure score constructed following Eloundou et al. (2024), aggregated to the three digit SOC level. Adaptive Capacity is the occupation-level index from Manning and Aguirre (2026), combining skill transferability, median net liquid wealth, geographic labor market density, and workforce age composition. The sample is restricted to employed respondents ages 20–65 with non-missing AI use frequency and industry and occupation codes. Public-sector respondents are excluded. Observations are weighted by the number of respondents in the occupation-industry cell.

Figure 4: Frequent AI Utilization Across Sectors, 2025 and 2025-26 Growth

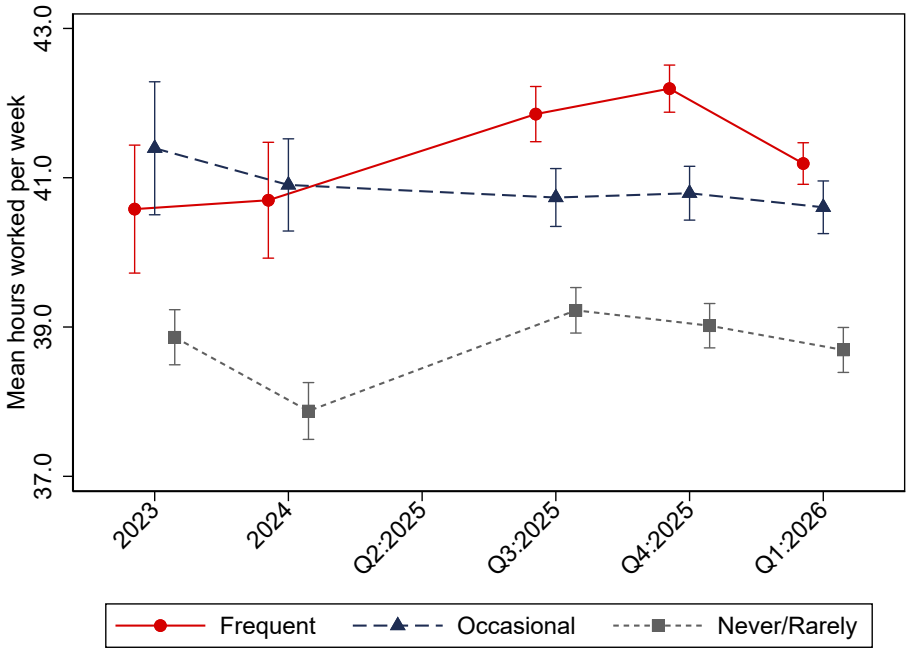


Notes.—Source: Gallup Workplace Panel, 2023–Q1:2026. Panel A plots the share of frequent AI users across industries as of Q1:2026. Panel B plots the change in that share from Q2:2025 to Q1:2026. Frequent AI use is defined as reporting using AI "daily" or "a few times a week" in response to the question: "How often do you use artificial intelligence in your role?" Other response options are: a few times a month, a few times a year, once a year, less often than once a year, and never. Before this question is asked, respondents are given the following definition: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." The sample is restricted to full- and part-time workers aged 20–65. All observations are aggregated using sample weights.

Figure 5: Mean Hours Worked per Week by AI Use Intensity



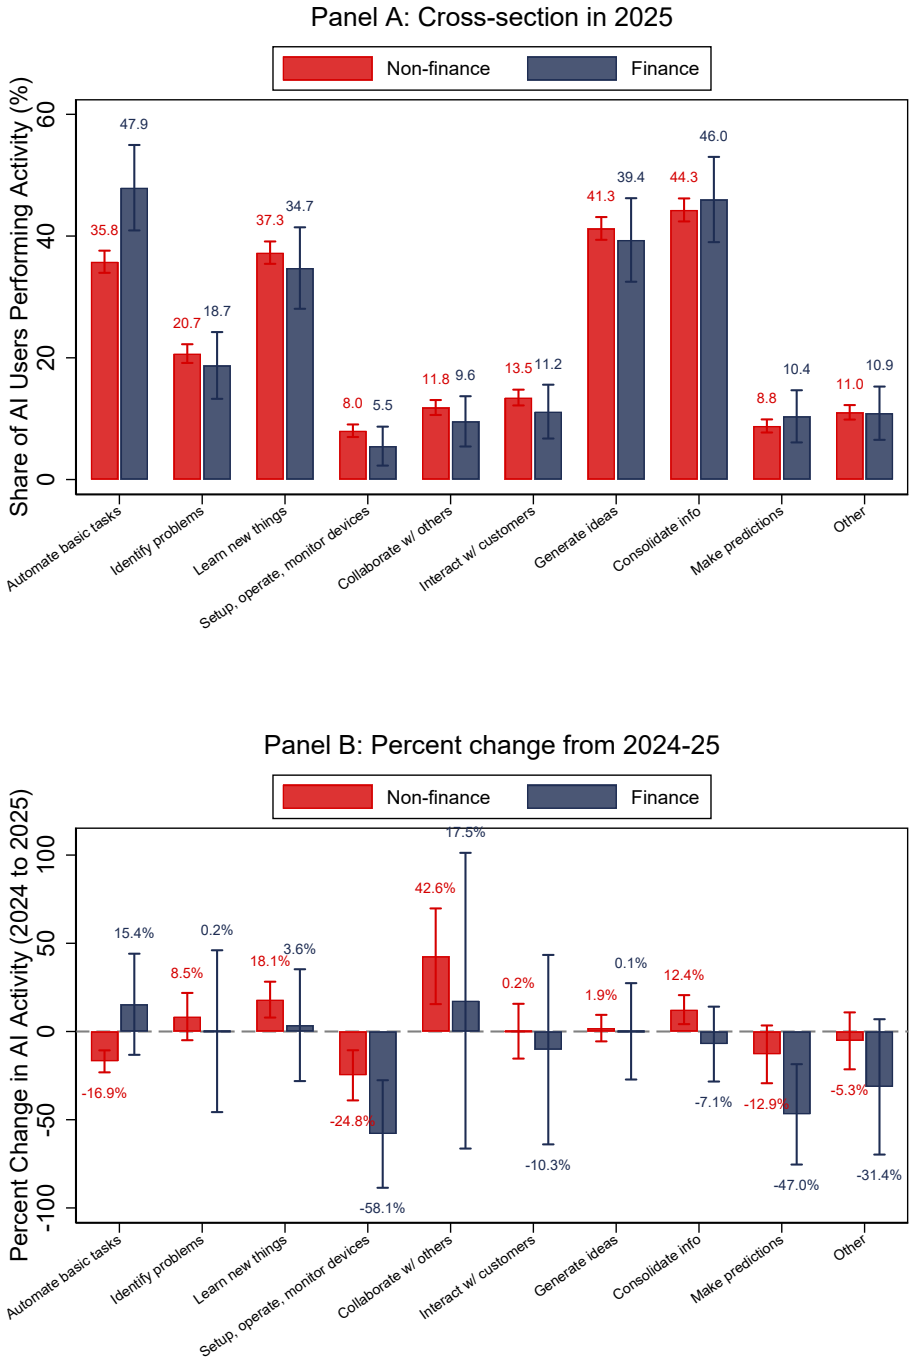
(a) Finance/Insurance



(b) Rest of Economy

Notes.—Source: Gallup Workplace Panel (2023–Q1:2026). The figure plots mean hours worked per week by AI use intensity separately for finance and insurance workers (Panel A) and the rest of the private-sector economy (Panel B). Frequent AI use is defined as using AI daily or a few times a week; Occasional AI use is defined as using AI a few times a month or a few times a year; Never/Rarely is defined as using AI once a year or less. Bands denote 95% confidence intervals. Points are jittered horizontally to reduce overlap. The sample is restricted to employed respondents ages 20–65 working at least 10 hours per week. Public-sector respondents are excluded. Observations are weighted by sample weights.

Figure 6: Composition of AI Activities Among Finance and Non-Finance



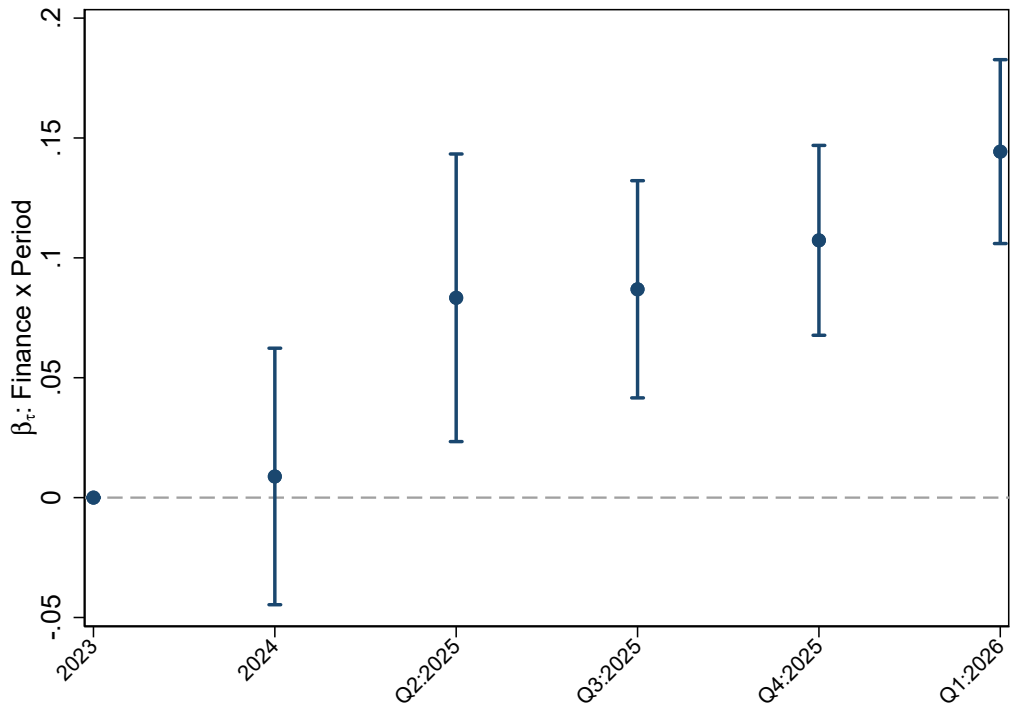
Notes.—Sources: Gallup Panel, 2024–2025. Panel A plots the share of respondents who report different AI use-cases and activities, conditional on using AI at least somewhat during the year. Panel B plots the percent change from 2024-25. The sample is restricted to full and part time workers ages 20-65, and the observations are weighted by sample weights.

Table 1: Summary Statistics on the Gallup Workplace Panel

	Pooled		Finance/Insurance		Rest of Economy	
	(1)		(2)		(3)	
	mean	sd	mean	sd	mean	sd
Frequent AI Use	0.27	0.44	0.38	0.49	0.26	0.44
Occasional AI Use	0.19	0.40	0.25	0.43	0.19	0.39
Clear AI Strategy	0.27	0.44	0.46	0.50	0.24	0.43
GenAI Exposure	0.40	0.15	0.50	0.08	0.39	0.15
GenAI Adaptive Capacity	0.59	0.26	0.66	0.23	0.58	0.27
Hours Worked/Week	39.49	10.58	41.67	8.14	39.26	10.78
Age	41.86	12.20	41.14	11.81	41.94	12.24
Male	0.51	0.50	0.54	0.50	0.51	0.50
White	0.68	0.47	0.65	0.48	0.68	0.47
Black	0.10	0.29	0.10	0.31	0.09	0.29
Has Children	0.36	0.48	0.37	0.48	0.35	0.48
Married	0.54	0.50	0.55	0.50	0.54	0.50
HS or Less	0.27	0.45	0.17	0.38	0.28	0.45
Some College	0.27	0.45	0.25	0.43	0.28	0.45
College or More	0.45	0.50	0.58	0.49	0.44	0.50
Observations	25261		2700		22561	

Notes.—Source: Gallup Workplace Panel (2023–Q1:2026). The table reports summary statistics on key variables for the pooled sample and separately for workers in finance or insurance versus all other private-sector industries. Frequent AI use is defined as using AI daily or a few times a week; Occasional AI use is defined as using AI a few times a month or a few times a year. AI use is measured by asking respondents how often they use artificial intelligence in their role, where artificial intelligence is defined as when a computer does or enhances a task that is normally performed by a person. GenAI Exposure is the task-based predicted exposure score constructed following Eloundou et al. (2024), aggregated to the three digit SOC level. GenAI Adaptive Capacity is the occupation-level index from Manning and Aguirre (2026), combining skill transferability, median net liquid wealth, geographic labor market density, and workforce age composition. Demographics include age, sex, race (White, Black), full-time status, children, marital status, and education. The sample is restricted to employed respondents ages 20–65 with non-missing AI use frequency, demographics, and industry and occupation codes. Public-sector respondents are excluded. Observations are weighted by sample weights.

Figure 7: Finance-Specific Acceleration in Frequent AI Use



Notes.—Source: Gallup Workplace Panel, 2023–Q1:2026. The sample is restricted to employed respondents aged 20–65. The figure plots $\hat{\beta}_\tau$ coefficients from Equation (1), which regresses an indicator for frequent AI use (daily or a few times a week) on interactions between a finance/insurance indicator and period dummies, with 2023 as the omitted baseline. Controls include age, sex, race (White, Black), full-time status, children, marital status, and education (some college, college or more). Frequent AI use is defined as reporting "daily" or "a few times a week" to the question "How often do you use artificial intelligence in your role?" Before this question, respondents are given the following definition: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." Bands denote 95% confidence intervals. Standard errors are clustered at the respondent level. Observations are weighted by sample weights.

Table 2: Examining the Determinants of Frequent AI Utilization

Dep. var. =	Frequent AI Use (1/0 Indicator)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	.292***			.283***	.112***			.113***
	[.011]			[.011]	[.015]			[.015]
× Finance/Insurance	-.036			-.040	-.052			-.065
	[.030]			[.030]	[.039]			[.039]
Adaptive Capacity		.065***		.054***		.024**		.025**
		[.005]		[.005]		[.012]		[.011]
× Finance/Insurance		.010		.003		-.044		-.040
		[.019]		[.019]		[.030]		[.030]
log(Restrictions/1000 Words)			-5.307***	-5.318***			.047	.038
			[.931]	[.875]			[.031]	[.031]
× Finance/Insurance			-54.769***	-44.756**			-46.012*	-46.958*
			[19.064]	[18.363]			[25.657]	[25.460]
Adj. R2	.17	.11	.10	.18	.60	.59	.59	.60
N	28548	28548	28548	28548	20132	20132	20132	20132
Controls	Yes	Yes	Yes	Yes	No	No	No	No
NAICS2 FE	Yes	Yes	Yes	Yes	No	No	No	No
Person FE	No	No	No	No	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Sources: Gallup Panel, 2023–2025. The table reports coefficients from regressions of an indicator for whether the respondent uses AI daily or a few times a week (zero otherwise) on measures capturing three hypothesized channels for the rise of AI in finance: (i) *governance*, proxied by an indicator for whether the respondent reports that their organization has a clear AI strategy; (ii) *adaptive capacity*, measured by the occupation-level index from Manning and Aguirre (2026), which combines skill transferability, median net liquid wealth, geographic labor market density, and workforce age composition, merged at the three digit SOC level; and (iii) *regulatory constraints*, proxied by the log of regulatory restrictions per 1,000 words matched at the NAICS2-by-year level. Each channel is interacted with an indicator for employment in finance or insurance. Columns (1)–(4) include NAICS2 and year fixed effects and demographic controls; columns (5)–(8) include person and year fixed effects without demographic controls. Demographics include: age, race (White, Black), full-time status, having children, being married, and education (some college, and college or more). The sample is restricted to full and part time workers with non-missing covariates, and observations are weighted. Standard errors are clustered at the respondent level.

Table 3: Examining the Determinants of Frequent AI Utilization (Robustness Using Index)

Dep. var. =	AI Utilization Index (0-10 Index)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	3.031***			2.930***	1.161***			1.167***
	[.092]			[.093]	[.125]			[.124]
× Finance/Insurance	-.401*			-.413*	-.295			-.431
	[.229]			[.232]	[.301]			[.292]
Adaptive Capacity		.725***		.610***		.088		.095
		[.048]		[.045]		[.097]		[.093]
× Finance/Insurance		-.018		-.090		-.544*		-.518*
		[.167]		[.162]		[.301]		[.289]
log(Restrictions/1000 Words)			-33.344***	-33.720***			.348*	.255
			[7.963]	[7.395]			[.206]	[.198]
× Finance/Insurance			-492.501***	-387.800***			-511.088***	-491.495***
			[157.810]	[147.679]			[188.578]	[182.050]
Adj. R2	.26	.18	.16	.28	.71	.70	.70	.71
N	28548	28548	28548	28548	20132	20132	20132	20132
Controls	Yes	Yes	Yes	Yes	No	No	No	No
NAICS2 FE	Yes	Yes	Yes	Yes	No	No	No	No
Person FE	No	No	No	No	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Source: Gallup Workplace Panel, 2023–2025. The table reports coefficients from regressions of an AI utilization index on measures capturing three hypothesized channels for the rise of AI in finance: (i) *governance*, proxied by an indicator for whether the respondent reports that their organization has a clear AI strategy; (ii) *adaptive capacity*, measured by the occupation-level index from Manning and Aguirre (2026), which combines skill transferability, median net liquid wealth, geographic labor market density, and workforce age composition, merged at the three digit SOC level; and (iii) *regulatory constraints*, proxied by the log of regulatory restrictions per 1,000 words matched at the NAICS2-by-year level. Each channel is interacted with an indicator for employment in finance or insurance. Columns (1)–(4) include NAICS2 and year fixed effects and demographic controls; columns (5)–(8) include person and year fixed effects without demographic controls. The AI utilization index is defined as follows: daily = 10, a few times a week = 8, a few times a month = 6, a few times a year = 4, once a year = 2, less often than once a year = 1, and never = 0. Demographics include age, race (White, Black), full-time status, having children, being married, and education (some college, and college or more). The sample is restricted to full- and part-time workers with non-missing covariates, and observations are weighted. Standard errors are clustered at the respondent level.

Table 4: Placebo Determinants of Frequent AI Utilization Using Occupation

Dep. var. =	Frequent AI Use (1/0 Indicator)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	.312***			.287***	.107***			.106***
	[.011]			[.011]	[.015]			[.014]
× Finance Role	.017			.028	-.020			-.016
	[.032]			[.032]	[.045]			[.045]
Adaptive Capacity		.066***		.043***		.021*		.021*
		[.005]		[.005]		[.011]		[.011]
× Finance Role		.010		.016		.023		.029
		[.039]		[.031]		[.028]		[.028]
log(Restrictions/1000 Words)			.205***	.111***			.052	.045
			[.014]	[.013]			[.032]	[.032]
× Finance Role			.008	.003			-.067	-.062
			[.056]	[.051]			[.066]	[.064]
Adj. R2	.15	.08	.08	.17	.60	.59	.59	.60
N	28548	28548	28548	28548	20132	20132	20132	20132
Controls	Yes	Yes	Yes	Yes	No	No	No	No
Person FE	No	No	No	No	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Source: Gallup Workplace Panel, 2023–2025. The table reports placebo regressions that replace the industry-based finance indicator used in the main text with an occupation-based finance indicator defined at the SOC3 level. The specification is included as a falsification exercise for whether the main finance result is driven purely by occupation mix rather than by the financial-services industry environment. Each row corresponds to one explanatory variable included in the placebo specification. Columns (1)–(4) include SOC3 and year fixed effects and demographic controls; columns (5)–(8) include person and year fixed effects without demographic controls. Standard errors are clustered at the respondent level.

Online Appendix for

"AI (at Work) in Finance"

Abstract

This Internet Appendix provides additional material on the data, methodology, and some robustness exercises related to the main text.

A U.S. Federal Supervisory Guidance

This section summarizes major federal supervisory guidance, rulemakings, and executive actions relevant to artificial intelligence and machine learning in financial services. The timeline is intended to contextualize the regulatory environment discussed in the main text during our sample period. It is descriptive and does not, by itself, identify the causal effect of any single policy change.

Table A.1 presents a chronological summary of major U.S. federal banking supervisory guidance, rulemakings, and executive actions related to artificial intelligence and machine learning in financial services. Several features are notable. First, SR 11-7 (2011) remains the foundational model risk management framework; regulators have confirmed that it applies to AI/ML models, but no comprehensive replacement guidance has been issued. Second, the CFPB issued some of the most operationally specific guidance through circulars on adverse action notice requirements when lenders use algorithmic credit models. Third, policy direction shifted significantly with the change in administration in January 2025, moving from a risk-management orientation (Executive Order 14110) to a more pro-innovation stance (Executive Order 14179 and the December 2025 national framework order). The SEC's withdrawal of its predictive data analytics rule in June 2025 further illustrates this shift.

B Supplementary Descriptive Evidence

This section reports additional descriptive evidence that complements the figures in the main text. Unless otherwise noted, descriptive appendix figures use the same sample conventions as the corresponding figures in the paper and extend through Q1:2026 when those data are available.

Table A.1: U.S. Federal Supervisory Actions Related to AI in Financial Services

Date	Agency	Document	Description
Apr 2011	Fed, OCC	SR 11-7 / OCC 2011-12: Model Risk Management	Foundational interagency guidance on model development, validation, and governance; now broadly applied to AI/ML models.
Mar 2021	Fed, OCC, FDIC, CFPB, NCUA	Interagency RFI on AI/ML Use (86 FR 16837)	Five-agency request for information on AI governance and risk management in financial institutions.
May 2022	CFPB	Circular 2022-03: Adverse Action Notices and Complex Algorithms	Creditors using AI for credit decisions must provide specific adverse action reasons; lack of model understanding is not a defense.
Jul 2023	SEC	Proposed Rule on Predictive Data Analytics (Rel. 34-97990)	Proposed rules requiring broker-dealers and advisers to eliminate conflicts from "predictive data analytics" technologies.
Sep 2023	CFPB	Circular 2023-03: Adverse Action and Sample Forms	Standard checkbox forms are insufficient if they do not reflect actual AI-driven denial reasons.
Oct 2023	White House	Executive Order 14110: Safe, Secure, and Trustworthy AI	Directed Treasury to report on AI cybersecurity risks and FSOC to monitor AI-related financial stability risks.
Mar 2024	Treasury	Report: AI-Specific Cybersecurity Risks in Finance	Based on 42 interviews; 10 recommendations including a common AI lexicon and expansion of the NIST AI RMF for finance.
Jun 2024	OCC, Fed, FDIC, NCUA, CFPB, FHFA	Final Rule: AVM Quality Control Standards	Requires AI-driven automated valuation models to meet standards for accuracy, nondiscrimination, and random testing.
Dec 2024	Treasury	Report: AI in Financial Services	Comprehensive report on AI/GenAI adoption; highlights data privacy, bias, third-party concentration, and model opacity risks.
Jan 2025	White House	Revocation of EO 14110; EO 14179 (Removing Barriers to AI Leadership)	Rescinded Biden AI order; mandates review of policies that could impede AI innovation.
May 2025	GAO	GAO-25-107197: AI Use and Oversight in Financial Services	Reviews regulators' AI oversight; finds gaps in NCUA guidance; recommends expanded examination authority.
Jun 2025	SEC	Withdrawal of Predictive Data Analytics Rule	Formally withdrew the 2023 proposed rule, citing industry criticism of overly broad scope.
Dec 2025	White House	EO: National Policy Framework for AI	Establishes a "minimally burdensome" AI framework; directs an AI Litigation Task Force to challenge inconsistent state laws.

Notes.— Compiled from Federal Register notices, agency press releases, and executive orders. See text for full citations and URLs.

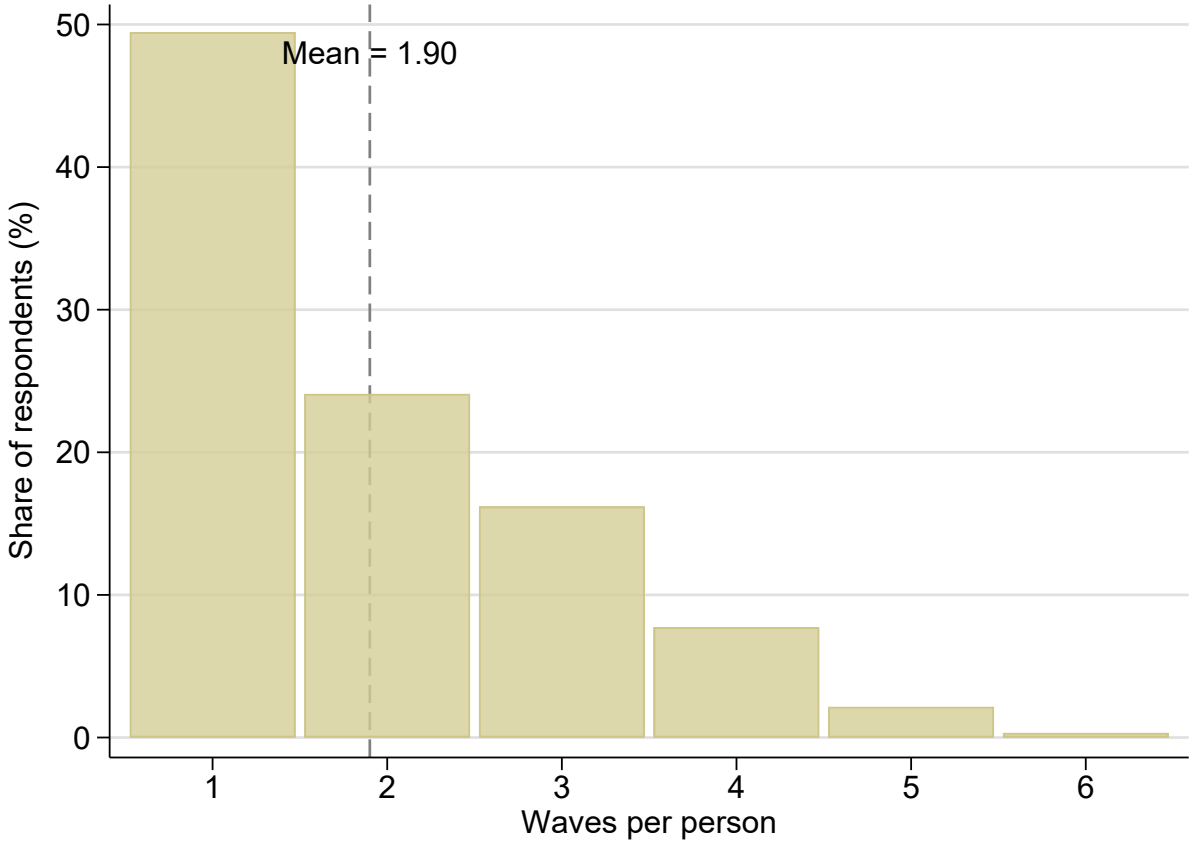


Figure A.1: Descriptive Evidence on Longitudinal Variation

Notes.—Source: Gallup Workplace Panel, 2023–Q1:2026. The figure plots the share of respondents across the potential number of waves in the data. The sample is restricted to respondents aged 20–65.

C Occupation-Level Supplementary Evidence

This section provides supplementary occupation-level evidence for the adaptive capacity interpretation in the main text. We first report cross-sectional finance-only regressions, then show additional scatterplots and descriptive figures that illustrate how exposure and adaptive capacity relate to AI adoption across occupations.

Table A.2: Decomposing Adaptive Capacity Within Finance (2025)

	Share Using AI Frequently	**AI Usage Frequency**
	(1)	(2)
Constant	0.318***	4.362***
	(0.018)	(0.191)
AI Exposure (z)	0.053	0.549
	(0.030)	(0.318)
Skill Transferability (z)	0.081**	0.779**
	(0.036)	(0.379)
Net Liquid Wealth (z, log)	0.047	0.519
	(0.037)	(0.385)
Geographic Density (z)	0.054	0.550
	(0.044)	(0.461)
Older Worker Fraction (z)	−0.003	−0.037
	(0.030)	(0.312)
Observations	20	20
R^2	0.746	0.702

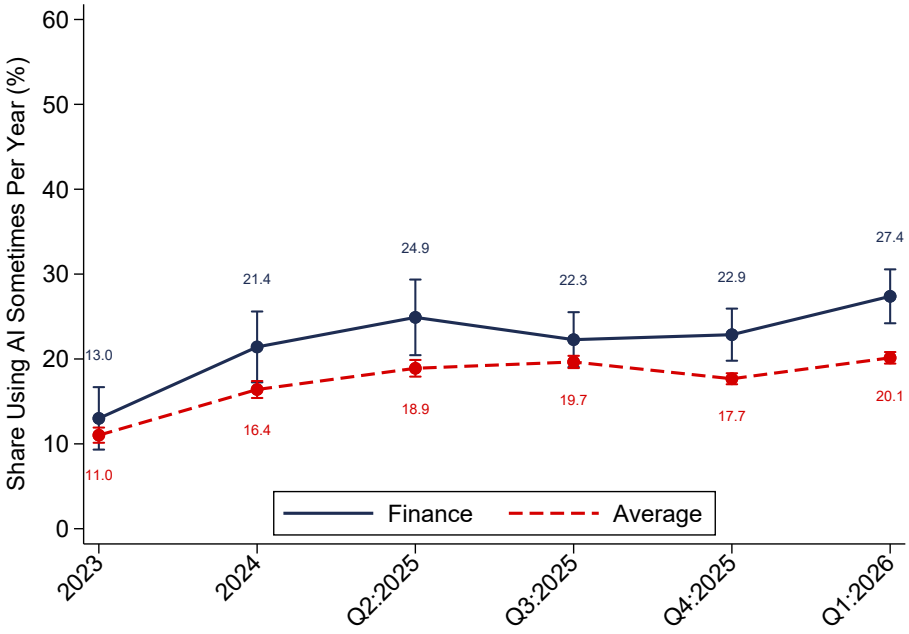
Notes.— Cross-sectional WLS regressions for 2025, restricted to finance and insurance respondents. All regressors are standardized. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3: AI Adoption, Exposure, and Adaptive Capacity: Finance Occupations (2025)

	Share Using AI Frequently			**AI Usage Frequency**		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.318***	0.318***	0.318***	4.362***	4.362***	4.362***
	(0.028)	(0.016)	(0.016)	(0.271)	(0.175)	(0.182)
AI Exposure (z)	0.012	0.031	0.032	0.026	0.208	0.208
	(0.034)	(0.021)	(0.021)	(0.350)	(0.243)	(0.253)
Adaptive Capacity (z)		0.098***	0.121***		0.963***	0.962***
		(0.020)	(0.025)		(0.202)	(0.338)
Exposure × AC			−0.026			0.001
			(0.023)			(0.238)
Observations	20	20	20	20	20	20
R^2	0.007	0.662	0.687	0.000	0.577	0.577

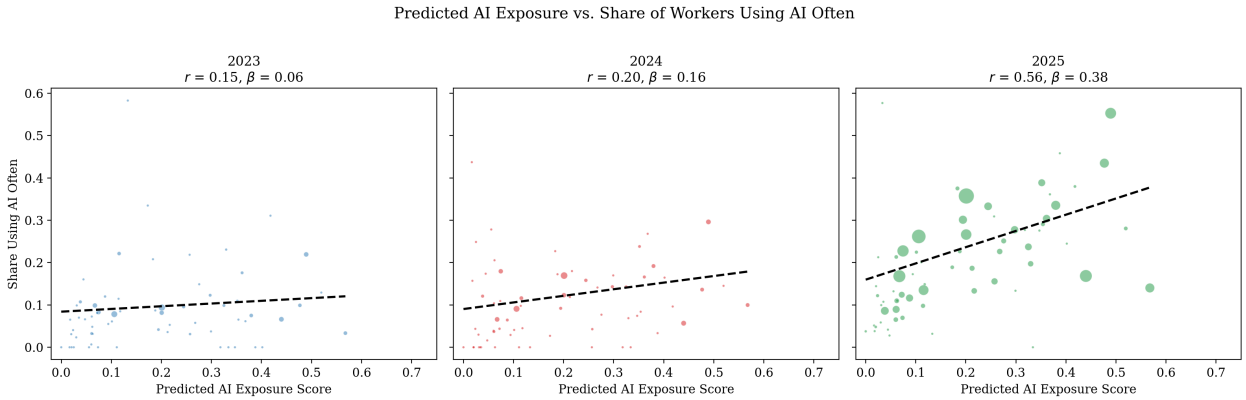
Notes.— Cross-sectional WLS regressions for 2025, restricted to respondents in finance and insurance. Weights are survey respondent counts. All regressors are standardized (z-scores). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A.2: "Sometimes" AI Users in Finance and the Aggregate



Notes.—Source: Gallup Workplace Panel, 2023–Q1:2026. The sample is restricted to respondents aged 20–65. The figure plots the share of respondents who answer "a few times a month" or "a few times a year" to the question "How often do you use artificial intelligence in your role?" Before answering, respondents are given the following definition: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." All series are aggregated using sample weights. Bands denote 95% confidence intervals.

Figure A.3: Predicted AI Exposure vs. Share of Workers Using AI Frequently



Notes.—Each panel plots predicted AI exposure (horizontal axis) against the share of workers using AI frequently (vertical axis) for three-digit SOC occupations. Bubble size is proportional to the number of survey respondents. Lines are weighted least squares fits.

Figure A.4: High-Exposure Occupations Are Adopting AI Faster



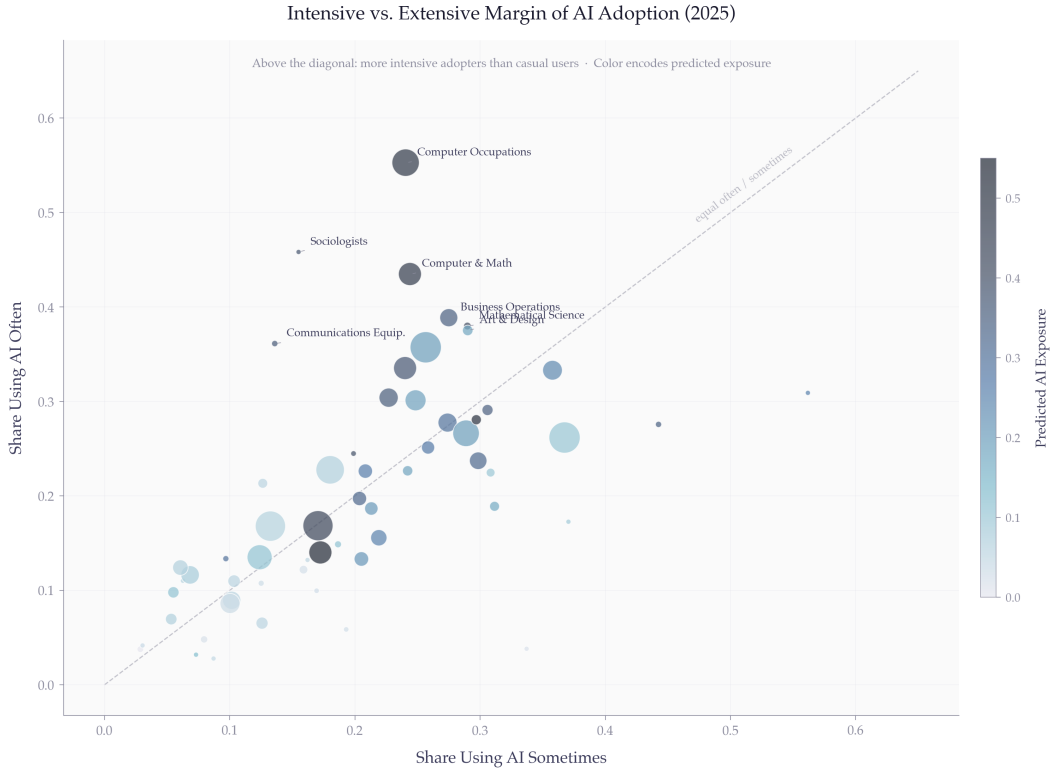
Notes.—Each bubble represents a three-digit SOC occupation present in both 2023 and 2025. The horizontal axis plots predicted AI exposure. The left panel plots the change in AI usage frequency; the right panel plots the change in the share of workers using AI frequently. Lines are weighted least squares fits. Bubble size is proportional to the average number of respondents across the two years.

Table A.4: Task-Level Exposure Classification Distribution (Newly Classified Tasks)

Label	Description	Tasks	Share
E0	No Exposure	210	81.1%
E1	Direct Exposure	43	16.6%
Unclassified	Label not parsed	6	2.3%
Total		**259**	**100%**

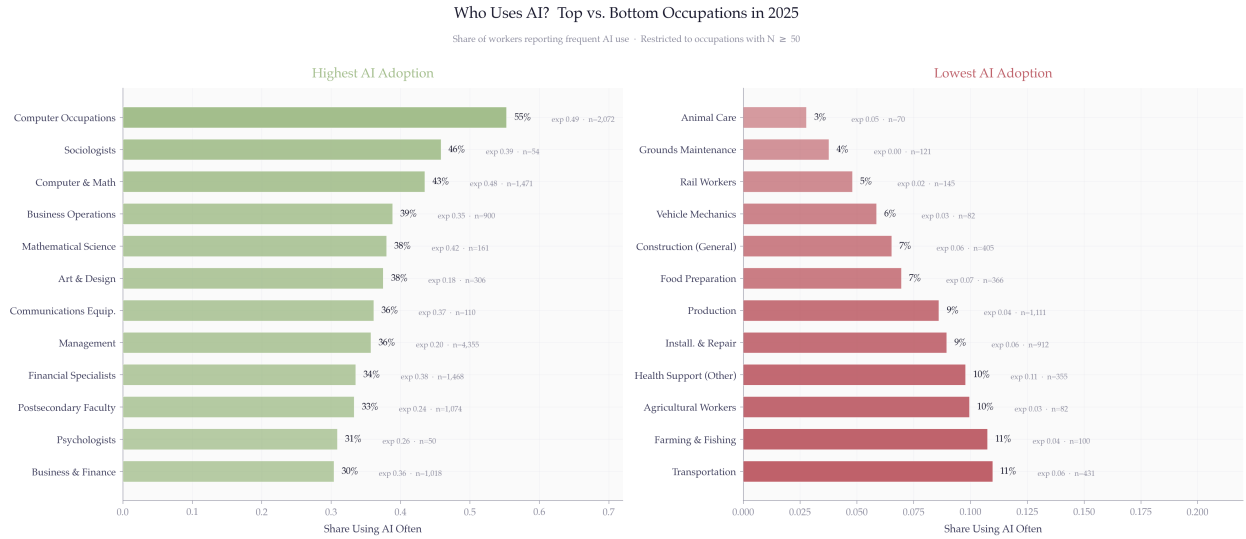
Notes.— Classification of 259 O*NET tasks across four three-digit SOC occupation groups not covered by the original analysis. Tasks were classified using GPT-4o-mini following the exposure rubric in Eloundou et al. (2024). These classifications are used only in supplementary appendix analyses.

Figure A.5: Intensive vs. Extensive Margin of AI Adoption (2025)



Notes.—Each bubble represents a three-digit SOC occupation in 2025. The horizontal axis plots the share of workers using AI occasionally; the vertical axis plots the share using AI frequently. Color encodes predicted AI exposure (darker = higher exposure). Points above the diagonal have more intensive users than casual users. Bubble size is proportional to the number of survey respondents.

Figure A.6: Who Uses AI? Top vs. Bottom Occupations in 2025



Notes.—The left panel shows the twelve occupations with the highest share of workers using AI frequently in 2025; the right panel shows the twelve occupations with the lowest share. Bars report the share of frequent AI users. Exposure scores and sample sizes are reported alongside each bar. The sample is restricted to three-digit SOC groups with at least 50 survey respondents.

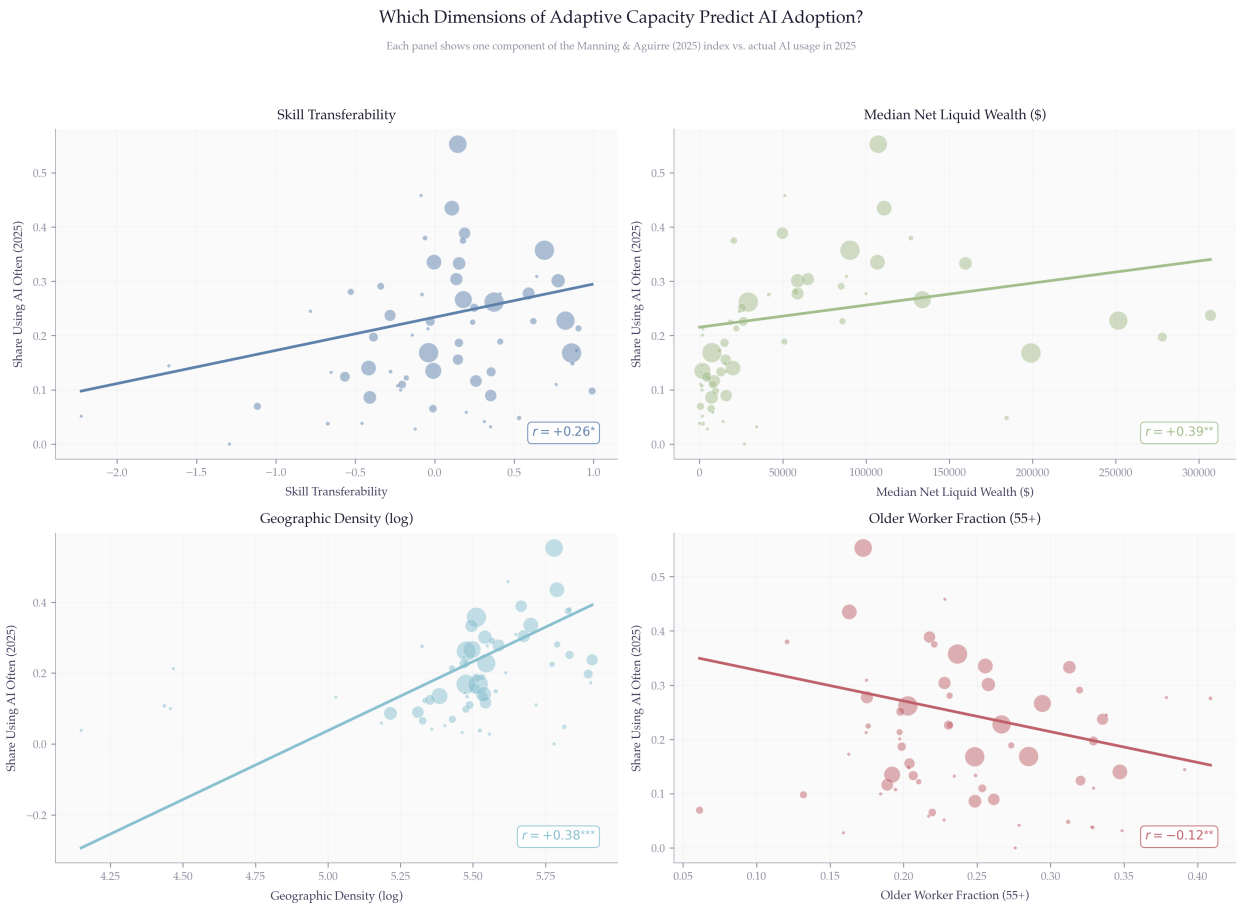
C.1 Robustness: Wages and Skill Composition

This subsection probes whether the adaptive capacity result in the main text is driven by simpler occupation-level correlates such as wages or AI-complementary skill intensity.

Table A.5 reports robustness checks that examine whether the adaptive capacity result in the main text (Table A.3) is driven by occupation-level wages or skill composition.

Using median annual wages from the Bureau of Labor Statistics Occupational Employment and Wage Statistics (OEWS, May 2024), we find that log wages are strongly correlated with adaptive capacity within finance ($r = 0.82$). Replacing AC with log wages yields a comparable fit ($R^2 = 0.67$ vs. 0.66 with AC alone; columns 1 and 2). When both are included simultaneously, neither is individually significant ($p = 0.25$ for AC, $p = 0.20$ for wages; column 3), and the R^2 rises only modestly to 0.70. This collinearity is expected: within finance, higher-paid occupations tend to have more portable skill sets, greater household wealth, and younger workforces—exactly the components of adaptive capacity. The wage result confirms

Figure A.7: Which Dimensions of Adaptive Capacity Predict AI Adoption?



Notes.—Each panel plots one component of the Manning and Aguirre (2026) adaptive capacity index (horizontal axis) against the share of workers using AI frequently in 2025 (vertical axis). Bubble size is proportional to the number of survey respondents. Lines are weighted least squares fits. Correlation coefficients and significance levels are reported in the figure.

that the adaptive capacity finding is robust to a simpler and more transparent proxy, but the two measures are too correlated to be cleanly separated within a 20-occupation cross-section.

We construct an AI-complementary skill index by averaging O*NET importance scores for programming, critical thinking, complex problem solving, and mathematics at the SOC-3 level. AI-complementary skills alone predict adoption as strongly as adaptive capacity. However, when both adaptive capacity and the skill index are included, adaptive capacity remains significant while the skill index becomes insignificant. This suggests that adaptive capacity captures variation beyond narrowly technical skills, consistent with skill *transferability* mattering more than any specific skill dimension.

Table A.5: Robustness: Wages and O*NET Skills as Alternative Predictors (2025)

	Share Using AI Frequently				
	(1)	(2)	(3)	(4)	(5)
Constant	0.318***	0.321***	0.322***	0.324***	0.334***
	(0.016)	(0.017)	(0.017)	(0.023)	(0.018)
AI Exposure (z)	0.031	0.032	0.032	0.025	0.048*
	(0.021)	(0.021)	(0.020)	(0.027)	(0.022)
Adaptive Capacity (z)	0.098***		0.049		0.117**
	(0.017)		(0.041)		(0.041)
Log Median Wage (z)		0.121***	0.066		
		(0.021)	(0.050)		
AI-Complementary Skills (z)				0.112***	−0.013
				(0.025)	(0.047)
Observations	20	20	20	13	13
R^2	0.662	0.668	0.695	0.677	0.831

Notes.— Cross-sectional WLS regressions for 2025, restricted to finance and insurance respondents. Weights are survey respondent counts. All regressors are standardized. Wages are from BLS OEWS (May 2024). The AI-complementary skills index is the average O*NET importance score for programming, critical thinking, complex problem solving, and mathematics (O*NET 29.1). The skills sample is smaller because O*NET data are available for 13 of 20 finance occupations at the SOC-3 level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.