AI (at Work) in Finance

Christos A. Makridis [*] Erik Brynjolfsson[†] Sophia Kazinnik[§]

February 26, 2026

Abstract

Using the Gallup Workplace Panel (2023–2025), we study the diffusion of artificial intelligence in U.S. financial services, measuring both adoption intensity (how frequently workers use AI) and task allocation (what they use it for). We evaluate three mechanisms driving the acceleration of AI use in finance in 2025: (i) a *governance threshold* channel, where AI use scales only once the firm approves it and has a clear AI plan; (ii) a *customer relationship capital* channel, where AI adoption is concentrated in customer facing work; and (iii) a *regulation induced composition* channel, where compliance rules discourage higher-stakes uses and push firms toward applications that are easier to document and audit. The share of frequent AI users (defined as daily or a few times per week) rises from 9.8% in 2023 to 39.4% by Q4:2025 in finance, versus 10.7% to 25.3% in the workforce overall. Using within-person variation, the best predictor of frequent AI use is whether workers say their organization has a clear AI strategy. This suggests governance acts as a gate for scaling AI use. Regulatory restrictions are especially salient in finance: higher restrictions predict materially lower AI intensity in finance relative to other sectors, consistent with compliance requirements shaping when tools can be deployed and steering AI away from higher-stakes workflows.

Keywords: AI Adoption; Financial Services; Intangible Capital; Regulatory Constraints; Organizational Strategy; Economic Measurement; Productivity; Workplace Technology.

JEL Codes: G20; J24; M54; O33.

[*]Corresponding author. Gallup, Arizona State University, Stanford University, christos_makridis@gallup.com.
[†]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

The financial services industry sits at the center of an AI puzzle. On the one hand, finance was a late, but intensive adopter: by 2025, more than half of U.S. financial firms report using AI, among the highest rates of any sector (see, e.g., reports by FRBNY, 2025; Census Bureau, 2024). On the other hand, adoption often remains cautious and limited in where it shows up. Only about 10% of institutions report integrating generative AI into products or workflows (Stanford HAI, 2025), and only around one third have scaled any AI initiative beyond pilot stages (McKinsey & Company, 2025).

We reconcile these two facts using repeated waves of Gallup's Workplace Panel, which measures worker-level AI use: both *how often* people use AI tools and *what they use them for*. We document three patterns. First, finance sees a sharp rise in *frequent* AI use in 2025. The share of workers using AI daily or a few times a week increases rapidly and pulls away from the economy-wide trend, while occasional ("sometimes") use grows much less. Second, among AI users, finance workers shift toward applications that are easier to deploy under compliance and oversight, especially customer interaction and routine automation, and away from high stakes uses such as prediction and monitoring. Third, these shifts line up with governance and policy frictions: measures of internal approval rules and external regulatory restrictions help explain both the timing of the 2025 surge and the sector specific tilt in which use cases expand.

These findings fit naturally with what makes finance distinct. AI is already used across many functions (e.g., 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

just whether firms adopt AI, but where they can use it, how quickly they can deploy it, and what is required to scale it safely. Recent work by the Financial Stability Board emphasizes 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 (Financial Stability Board, 2024). Regulators have also highlighted challenges around explainability, accountability, and systemic risk (Bank for International Settlements, Financial Stability Institute, 2025; International Organization of Securities Commissions, 2021). Industry evidence similarly suggests that large firms already use AI extensively in finance and accounting for planning, reporting, and internal decision support (KPMG, 2024).

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 touches, and how those patterns respond to broader forces such as regulation. This is a key measurement gap. Unlike job postings (Babina et al., 2024) or firm disclosures (Cao et al., 2025), the Gallup data reflect how actual employees in finance are using AI in workflows, including intensity and task allocation. That allows us to observe how AI adoption is formed, governed, and constrained in practice.

Our interpretation is simple: finance leads in AI uptake not because it escapes regulation, but because it deploys AI where governance allows it. In regulated intermediaries, the relevant margin is often not whether an algorithm exists, but whether it is *used frequently* in frontline workflows and, crucially, *which tasks* it is permitted to touch. Evidence from workers on usage intensity and task assignment helps reconcile two views: AI adoption may be slow because of model risk and governance limits, or fast because finance is information-intensive and the payoff to automation and augmentation is high.

We begin our paper by proposing three plausible and potentially overlapping mechanisms that

could explain why finance leads in AI uptake and how that uptake shows up in frontline task use. First, a governance threshold: 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, a customer facing augmentation channel: Gallup data show that customer service is the most common use of AI in finance, more so than risk or prediction, matching industry accounts in which AI supports frontline communication rather than replacing core decisions (Gallup, Inc., 2025; FRBNY, 2025). Third, regulatory effects: stricter oversight may push AI toward applications that are easier to document and audit, and away from higher-stakes applications, 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).

Our contribution follows directly from these channels. First, we provide direct evidence on AI use inside financial firms, measuring both the intensive margin (how frequently AI is used) and its allocation across tasks. Second, we use these worker-level measures to interpret adoption in a sector that matters for the broader economy. Finance is a core piece of economic infrastructure: it allocates capital, manages risk, and produces information that disciplines firms and reallocates resources. Classic work links financial development to real growth through mobilizing savings, screening projects, and monitoring managers (King and Levine, 1993; Rajan and Zingales, 1998). More recent evidence shows that the modern expansion of finance coincided with large shifts in tasks and human capital, with the sector drawing disproportionately skilled labor and paying a substantial wage premium (Greenwood and Scharfstein, 2013; Philippon and Reshef, 2012; Célérier and Vallée, 2019). Against this backdrop, knowing what finance workers actually do with AI, rather than inferring adoption from disclosures or external text proxies, is key for understanding how AI

reshapes the production of financial services and, ultimately, the financing of real investment and other information-intensive activity.[1]

Second, we link these worker-level patterns to corporate finance research on organizational and intangible capital, where the key idea is that value increasingly resides in scalable, partly non-rival inputs such as software, data, processes, and organization capital (Corrado et al., 2006; Crouzet et al., 2022; Veldkamp, 2023). In this literature, organization capital is a productive asset embedded in managerial systems and human coordination, and it is central to firm value and expected returns (Eisfeldt and Papanikolaou, 2013). A complementary strand in economics shows that the gains from new general-purpose technologies are mediated by management and workplace complements, implying that AI adoption should be read through the lens of permissioning, workflow redesign, and monitoring capacity rather than a frictionless technology choice (Bresnahan et al., 2002; Bloom et al., 2012). Recent corporate finance evidence further suggests that organization capital shapes financing outcomes, such as credit assessments, debt pricing, and payout policy, consistent with intangibles affecting both cash-flow risk and contracting frictions (Danielova et al., 2023; Hsieh et al., 2025; Falato et al., 2022). Our evidence that finance scales AI where it can be governed and audited fits naturally in this view: AI contributes to intangible/organizational capital not simply by raising adoption, but by shifting which workflows can be standardized, monitored, and deployed at scale under regulatory and internal constraints.

The remainder of the paper is organized as follows. Section 2 describes the Gallup Workplace Panel data and how we measure AI use, along with our measures of organizational strategy and regulatory exposure. Section 3 presents the main descriptive patterns in finance relative to the

[1]For example, complementarities between information technology and organizational practices are well documented (Bresnahan et al., 2002; Bloom et al., 2012), and similar complementarities likely matter for AI adoption and impact (Makridis, 2025).

rest of the economy. Section 4 lays out the empirical framework and the testable implications of the mechanisms we study. Section 5 reports the main results and robustness checks. Section 6 concludes.

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. 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 23,399 unique individuals and 43,872 person-wave observations from 2023-2025.[2] Questions related to AI appear intermittently before 2025; starting in Q2:2025 they are asked quarterly, which allows us to track adoption at higher frequency during 2025 and the survey is continuing to date allowing repeated observations.

In our linked sample, 14,029 respondents appear in at least two waves, supporting panel specifications that use within-person changes over time. By year, the samples include 9,326 workers in 2023, 9,992 in 2024, and 30,815 in 2025. The survey also provides detailed individual covariates, including age, sex, race/ethnicity, education, marital status, children, and job characteristics

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

(full-time status, tenure, income brackets), 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.[3]

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:

- **Often:** daily or a few times a week;
- **Sometimes:** a few times a month or a few times a year;

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

- **Never/rarely:** once a year, less often than once a year, or never.

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).

Beyond frequency, the panel 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 overall adoption rises.

2.4 Workplace Practices and AI Strategy

A key advantage of the Gallup panel is its detailed measures of workplace culture and management practices. We capture these practices using six agreement items on a 1–5 scale:

1. I feel connected to my organization's culture;
2. My organization consistently delivers on the promises it makes to customers;
3. I trust the leadership of this organization;
4. My organization cares about my overall wellbeing;

5. I have received meaningful feedback in the past week;

6. I am treated with respect at work.

We focus on answers to the second measure of "customer capital" in mapping our theory of organizational capital to practice, but results are robust to using other dimensions too. In addition, we measure organizational AI governance using an indicator for whether the respondent reports that the organization has communicated a clear AI strategy, using the following question:

> *"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 materials 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 for finance/insurance and for the rest of the economy, and shows a meaningful decline in 2025 relative to 2023–2024.

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. In short, provides a useful time-varying measure of regulatory constraints at the industry-year level.

3 Descriptive Patterns

3.1 AI Adoption Rates

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.[4] Two patterns stand out. First, frequent AI use rises in both finance and the economy between 2023 and late 2025, but the increase is much steeper in finance. Finance begins slightly below the aggregate in 2023 (9.8% versus 10.7%), pulls ahead by 2024 (17.0% versus 12.9%), and the gap widens through 2025 and remains elevated. By Q4:2025, 39.4% of finance respondents are frequent AI users, compared with 25.3% in the full sample, a difference of about 14 percentage points.

Second, the acceleration of AI use in finance is concentrated in 2025. In finance, the share of frequent users rises from 17.0% in 2024 to 31.4% by Q2:2025 and continues to 33.1% in Q3:2025 and 39.4% by Q4:2025. Over the same period, the economy-wide share increases more gradually, from 12.9% in 2024 to 19.8% in Q2:2025, 21.8% in Q3:2025, and 25.3% by Q4:2025.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 further put these trends in perspective, Figure 3 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 Q3:2025. Technology/math stands out as the clear leader, with 54.0% of workers reporting frequent use. Finance/insurance form a second tier: 36.4% of finance workers and 37.3% of insurance workers are frequent AI users, exceeding most other sectors, including consulting (33.4%), real estate (27.2%),

[4]See Figure A.1 in the Online Appendix for workers who use AI only *sometimes*; these shares level off over time.

education (26.7%), and the public sector (25.9%). At the other end of the distribution, frequent use remains below 20% in sectors such as hospitality/travel/food (12.5%), agriculture/forestry (15.6%), and transportation/logistics (16.0%).

Panel B reports the change in frequent AI use between Q3:2025 and Q4:2025. Growth is broad but uneven: technology/math again shows the largest increase (25.4 percentage points), followed by insurance (21.4 pp) and finance (19.4 pp). Several white-collar sectors also exhibit sizeable gains over this short window, including arts/entertainment/media (18.6 pp), education (16.4 pp), and consulting (14.6 pp), while retail/sales (5.5 pp) and industrial/extraction (7.5 pp) expand more modestly. In this sense, Figure 3 points out two key results: while finance is not the economy-wide leader in frequent AI use (technology/math/professional services remain far ahead), but finance and insurance are among the strongest adopters outside of tech, and their late 2025 acceleration is among the most pronounced across sectors.

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 between finance and other industries.

Panel A of Figure 4 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* (47.6% vs. 35.1%). Finance also reports higher use for *consolidating information* (46.3% vs. 44.1%) and *making predictions* (10.9% vs. 8.6%). By contrast, non-finance users report higher shares in several other categories, including *learning new things* (37.2% vs. 35.3%), *collaborating with others* (11.9%

vs. 9.3%), and *setup/operate/monitor devices* (8.1% vs. 5.4%). Customer-facing use is similar but slightly lower in finance: *interacting with customers* is 11.8% in finance versus 13.8% outside finance. Several remaining categories are broadly comparable across groups, including *generating ideas* (39.3% in finance vs. 40.7% in non-finance) and *identify problems* (19.0% vs. 20.9%).

Panel B of Figure 4 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 14.8%, while *setup/operate/monitor devices* falls sharply (−58.7%) and *making predictions* declines (−46.1%). Finance also shows declines in *other* (−29.3%) and a modest reduction in *interacting with customers* (−11.7%). Non-finance displays a different pattern, with large increases in *collaboration* (+36.7%) and smaller increases in *consolidating information* (+10.0%) and *learning new things* (+11.1%), alongside declines in *automation* (−16.8%), *setup/operate/monitor devices* (−25.0%), and *making predictions* (−17.4%). The contrast highlights that finance's 2024–2025 shift is driven by a strong expansion in routine automation and a contraction in operational monitoring 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 declines in *making predictions* and *setup/operation/monitoring* are consistent with firms limiting or centralizing applications that are more likely to trigger internal controls, because they resemble decision support, surveillance, or model-driven monitoring, 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, customer interaction remains a prominent use case. A natural interpretation is that firms scaled AI in communication-heavy workflows (agent-assist, drafting, retrieval, summarization, and customer support), where controls, such as approved workflows, logging, standardized disclosures, and human review, are easier to enforce. This would produce exactly the pattern in Figure 4: less emphasis on prediction and monitoring and more on customer facing use.

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). A shift away from tightly controlled, higher-stakes use cases in finance is consistent with firms responding to these frictions, especially when guidance remains unsettled; it is only beginning to change now.

Finally, these descriptive shifts are consistent with (though not definitive evidence of) heightened regulatory uncertainty in 2025, even as policy rhetoric emphasized lowering barriers to innovation. In financial market regulation, the rulemaking agenda saw notable resets, including the SEC's withdrawal of proposals related to predictive analytics and conflicts of interest. Such withdrawals can ease near-term compliance pressure, while still encouraging a "wait-and-see" stance for higher risk AI deployments, especially predictive tools, if firms expect revised rules or shifting enforcement priorities. Cross-border dynamics may reinforce this caution: as the EU AI Act moves into implementation, multinational firms may harmonize controls upward, with particularly large effects in heavily regulated sectors such as finance.

3.3 Labor Supply

A natural place to look for economic effects of workplace AI is labor supply (especially hours worked) and this margin is particularly relevant in finance. Finance has long been associated with high-skill, high-intensity work. Over the twentieth century and into the 2000s, compensation in finance rose sharply relative to other sectors alongside shifts in the tasks and human capital the industry demands (Philippon and Reshef, 2012). Using longitudinal data on elite MBA graduates in corporate and financial occupations, Bertrand et al. (2010) show that long hours and limited flexibility are closely tied to career advancement, consistent with client driven demands and co-ordination frictions that make time requirements central. Finance also exhibits high returns to talent and strong complementarities between talent and scale (Célérier and Vallée, 2019).

Technologies that raise the productivity of information processing and communication could therefore affect hours in either direction: they may reduce time spent on routine work (lowering hours for a given output) or increase the marginal return to effort in scalable roles (raising hours if effort and scale are complements). Consistent with the time-savings channel, Dillon et al. (2025) conduct a large-scale deployment of generative AI for workplace communication, finding that workers spend less time on email and reduce after-hours work, suggesting that AI can compress routine communication tasks and relax time constraints in day-to-day workflows. Bick et al. (2024) similarly document the rise of generative AI, showing that users report time savings. Our results add to this debate about the effects of genAI adoption at the individual-level on labor supply.

Figure 5 relates AI use to weekly hours worked in finance (Panel A) and in other sectors (Panel B), conditional on demographic controls. First, in finance, both frequent AI use ("often") and occasional use ("sometimes") are generally associated with fewer weekly hours relative to

the "never" group. The largest negative association for frequent users appears around Q3:2025, though the confidence intervals are wide. This pattern does not fit a simple "AI increases work intensity" story within finance. Instead, it is consistent with AI being used in ways that save time in work processes with strict controls around how work is done and documented, or with selection into frequent AI use among finance roles that already have lower baseline hours.

Second, outside finance, the association between AI use and weekly hours varies sharply with usage intensity. Frequent AI use is positively associated with weekly hours, while occasional use is negatively associated with hours, creating a stable gap between "often" and "sometimes" users across waves. One interpretation is that, in non-finance sectors, intensive AI use is more common in high-pressure settings where workloads and time demands are already high (so AI complements long hours), whereas occasional use reflects lighter-touch experimentation in roles with lower time pressure. In this sense, the finance surge in AI utilization documented earlier is not accompanied by a parallel rise in hours worked. Instead, the contrast suggests different modes of integration: in finance, AI is used as a time-saving, compliance-friendly layer; elsewhere, AI use appears more like a complement to high-hours environments.

This interpretation matches business surveys showing that finance tends to use AI to make internal processes faster, especially reporting, compliance, and decision support, rather than to boost production or cut headcount (KPMG, 2024; FRBNY, 2025). The fact that the hours–AI relationship looks different outside finance supports the same point: finance adoption centers on time-saving uses with clearer guardrails, while heavy AI use in other sectors shows up more in high-demand work environments.

4 Empirical Strategy and Testable Mechanisms

The descriptive evidence points to a clear pattern: finance sees a sharp rise in *frequent* AI use in 2025, but growth is not evenly distributed across tasks. Adoption expands in customer-facing work and routine automation, while predictive and monitoring uses grow less, or even decline. We interpret this as evidence that firms are not simply adopting more AI; they are reallocating AI toward applications that are easier to deploy and scale. We formalize three mechanisms that can generate this pattern:

1. **Governance threshold.** In regulated firms, using AI routinely usually requires upfront work to get approvals and put controls in place (e.g., policy review, model risk sign off, vendor due diligence, logging and monitoring, and workflow redesign). Once these guardrails are built, frequent use can scale quickly.

2. **Customer-relationship capital.** In finance, the biggest payoffs from AI may come from improving how firms communicate with and serve customers. If AI boosts productivity in customer interactions (e.g., drafting, personalization, agent-assist), firms have strong incentives to roll it out broadly, shifting adoption toward customer-facing uses.

3. **Regulation-induced composition.** Compliance rules make some AI uses much harder to scale than others. Applications that look like decision making, surveillance, or high-stakes prediction typically require heavier documentation, validation, and auditability, which raises the cost of deployment. By contrast, drafting, summarization, and support triage are easier to oversee and control, so adoption may tilt toward these lower-risk activities.

These mechanisms are hard to study with firm-level proxies such as disclosures or job postings. The Gallup panel allows more direct tests by linking worker-level AI intensity and task use to both reported organizational governance and industry-year regulatory exposure. Our governance threshold mechanism also builds on classic evidence that technology delivers the largest gains when paired with complementary organizational investments and management practices (Bresnahan et al., 2002; Bloom et al., 2012), including in the context of AI (Makridis, 2025).

A central implication is that regulation shapes not only *how much* AI is used, but *where* it is used. When higher-stakes applications require substantial validation and oversight, firms may rationally scale AI first in workflows where guardrails are easier to implement, while delaying or avoiding deployments that invite greater model risk or supervisory scrutiny. This view is consistent with emerging evidence that AI adoption can increase operational risk, such as fraud, business disruption, and process failures, particularly in firms with weaker internal controls (McLemore and Mihov, 2025). If regulators anticipate these vulnerabilities, the compliance burden for high-stakes AI uses may rise further, reinforcing a shift toward lower-risk deployments.

Each mechanism implies testable predictions about (a) how often workers use AI and (b) what they use it for. We test these predictions using three sets of models:

1. Event-style difference-in-differences. To quantify the 2025 surge in finance-specific adoption, we estimate:

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

where OftenAI_{it} indicates whether respondent i uses AI *daily* or *a few times a week*, τ_0 is a pre-surge baseline, and X_{it} includes demographics and job characteristics. This model tracks the timing and size of the finance-specific acceleration. If adoption is concentrated on the intensive margin, β_τ should rise sharply in 2025 but stay flat before that.

2. Governance-threshold mechanism. To test whether internal permissioning drives intensive use, we estimate:

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

where Gov_{it} captures whether a firm has a clear AI policy or integration strategy. A positive θ_2 supports the governance-threshold view. Conditional on governance, the finance premium should shrink. We also test whether governance predicts which types of AI use expand, expecting it to matter more for frequent, auditable applications like customer interaction and automation, rather than uniformly.

3. Regulation-induced composition. To test whether regulation affects the mix of use cases, we use QuantGov's measure of regulatory restrictions by sector-year:

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

and among AI users, we estimate:

$$\text{UseCase}^k_{it} = \alpha^k + \delta^k_t + \pi^k_1 \text{Finance}_i + \pi^k_2 \text{Reg}_{s(i),t} + \pi^k_3 (\text{Finance}_i \times \text{Reg}_{s(i),t}) + \Gamma^k X_{it} + u^k_{it}. \tag{4}$$

We expect $\text{Reg}_{s,t}$ to be associated with *less* use of high-stakes applications (e.g., prediction, monitoring) and *more* use of low-risk applications (e.g., automation, customer support) in finance.

These predictions match the descriptive patterns. In 2025, finance shifts AI use toward customer interaction and routine automation, while scaling back prediction and monitoring, the applications most likely to face stricter internal controls or external scrutiny.

When interpreting the results, we focus on two features of our empirical design. First, the main identifying variation comes from within-sector changes over time, not from fixed differences across sectors. Second, we think regulation primarily shapes which AI applications are feasible, rather than the overall level of adoption, which helps guide how we interpret the reduced-form coefficients.

Finally, we treat governance and regulation as related but distinct constraints. Governance reflects firm-level readiness and risk tolerance and may be endogenous. Regulation is a sector-level constraint that varies across industries and over time. This distinction lets us test two ideas: (i) whether frequent AI use rises as internal constraints loosen ((2)), and (ii) whether higher-stakes applications shrink more in sectors facing tighter restrictions ((4)).

5 Main Results

Table 1 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. 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 much more likely to use AI often. This relationship appears both in models with NAICS2 and time fixed effects (columns

1 and 4) and in models with person and time fixed effects (columns 5 and 8). The fact that the estimate remains in the person fixed-effects models implies that, within a given worker, changes in perceived strategy clarity are closely tied to moves into frequent AI use. By contrast, the interaction between strategy clarity and finance/insurance is small and statistically indistinguishable from zero, suggesting that governance acts as a general gate for scaling AI rather than a finance-specific channel.

The second result is that our proxy for customer-related organizational capital does not reliably predict frequent AI use once we control for fixed effects and the other channels. The coefficients are small and statistically insignificant, and the interaction with finance/insurance is also close to zero in both the NAICS2/time and person/time specifications (columns 2 and 6), including in the full models (columns 4 and 8). This aligns with the descriptive evidence: the 2025 surge in finance is not well explained by differences in broad workplace practices or customer-delivery perceptions as captured by this measure.

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 lower frequent AI use, with a much stronger negative relationship in finance/insurance (columns 3 and 7). Although the overall (level) effect is sensitive to the full set of controls (column 4), the finance/insurance interaction remains large and statistically significant in the person fixed-effects model (column 8). Taken together, the results support a constrained-adoption interpretation. Internal governance is a broad prerequisite for scaling AI use, while cross-industry differences in regulatory stringency help explain why finance is especially sensitive to the external environment, even after absorbing worker fixed effects and common time

shocks.

Table 2 reports robustness results using a broader AI intensity index that ranges from 0 to 10, constructed by mapping AI use frequency into a cardinal scale (not just the indicator). The qualitative conclusions from Table 1 remain intact. Clear AI strategy is strongly and precisely associated with higher AI intensity across specifications (columns 1 and 4; and columns 5 and 8 with person fixed effects), while its interaction with finance/insurance remains small and statistically insignificant. Customer capital continues to exhibit limited explanatory power, with small coefficients and no evidence of a differential association in finance/insurance once fixed effects are included. By contrast, the regulatory restrictions measure enters negatively and, more importantly, the interaction with finance/insurance is large and statistically significant in both the baseline and person fixed effects specifications (columns 3 and 7; and in the saturated model in column 8), indicating that tighter regulatory environments are associated with materially lower AI intensity in finance relative to other sectors. Overall, the index-based results show that our main findings do not depend on how we define "frequent" AI use.

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 on (Eloundou et al., 2024)), then we should see effects using this classification too. However, Table A.1 in the Online Appendix shows that no such interaction effects between occupation-level finance respondents and regulatory restrictions exist, although the negative direct effects of regulatory restrictions on AI use are still present.[5] This is consistent with a causal channel coming from regulation.

[5]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).

6 Conclusion

AI is spreading quickly across financial services, but its deployment is shaped as much by governance and regulation as by the technology itself. Using nationally representative panel data from the Gallup Workforce Panel, we show that frequent AI use in finance rose sharply in 2025. This increase reflects regular, embedded use, not just experimentation. At the same time, adoption shifted across tasks: AI expanded in customer service and routine automation, while predictive and monitoring uses grew more slowly and in some cases declined.

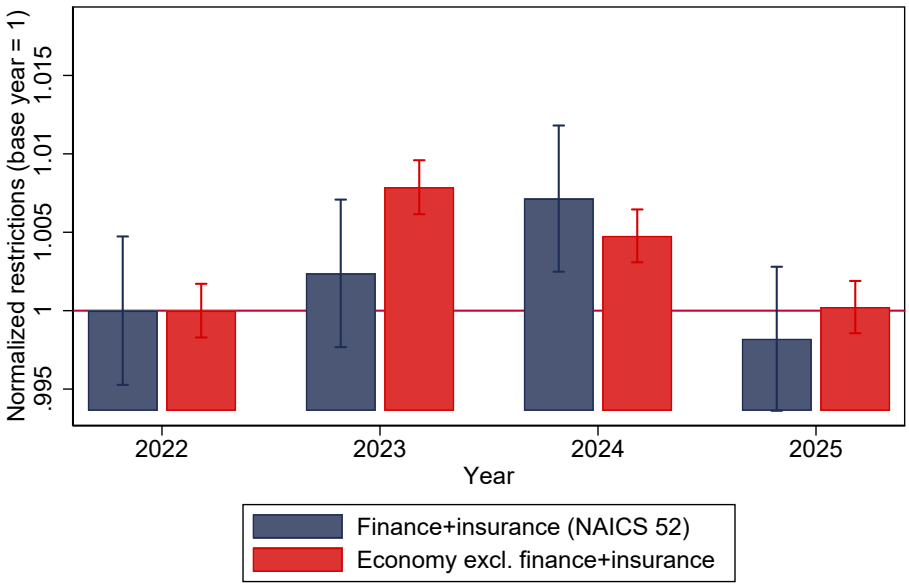
We interpret these patterns through three mechanisms. First, a *governance threshold*: scaling AI often requires upfront investments in oversight, approvals, and workflow integration, so use can accelerate once these prerequisites are in place. Second, a *customer-relationship channel*: firms may prioritize AI where it improves client communication and service. Third, *regulation-induced composition*: compliance pressures raise the cost of higher-stakes, harder-to-audit applications and tilt deployment toward uses that are easier to document and control. In the data, internal governance predicts frequent AI use across sectors, but does not appear uniquely stronger in finance. By contrast, regulatory restrictions are disproportionately related to both adoption and the mix of applications in finance.

These results help explain an apparent contradiction in discussions of AI in finance: the sector can adopt quickly while remaining selective about where AI is allowed to operate. More generally, governance and regulation may shape not only the pace of diffusion but also which applications are feasible to deploy at scale. By measuring AI use directly at the worker level our evidence complements prior work based on job postings or firm disclosures.

More broadly, the main barriers to AI adoption in high-stakes sectors may be less about technical feasibility and more about organizational readiness and institutional constraints. Understanding how internal controls and external rules shape deployment will matter for researchers, practitioners, and policymakers as AI continues to diffuse through regulated parts of the economy.

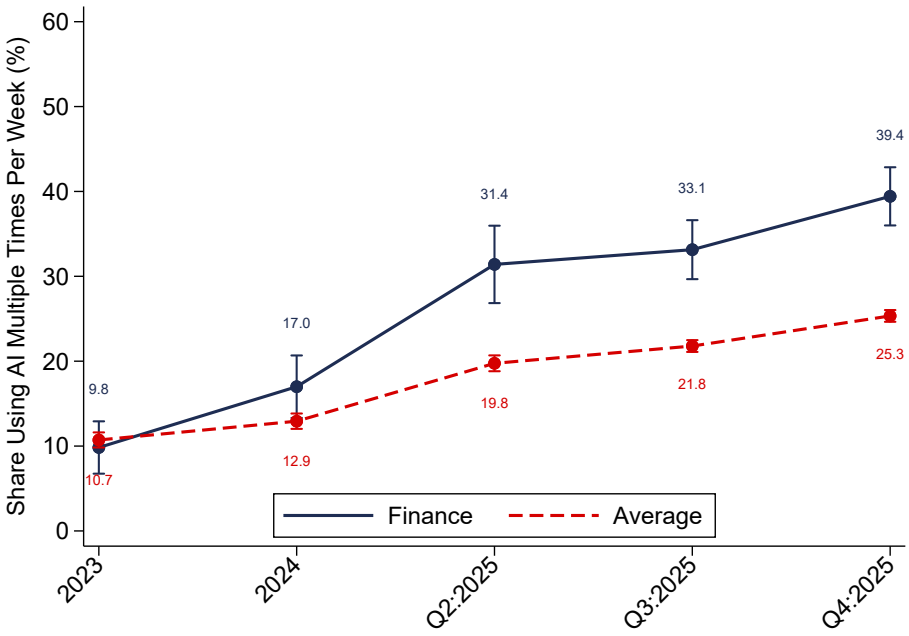
Tables and Figures

Figure 1: Time Series of Regulatory Restrictions



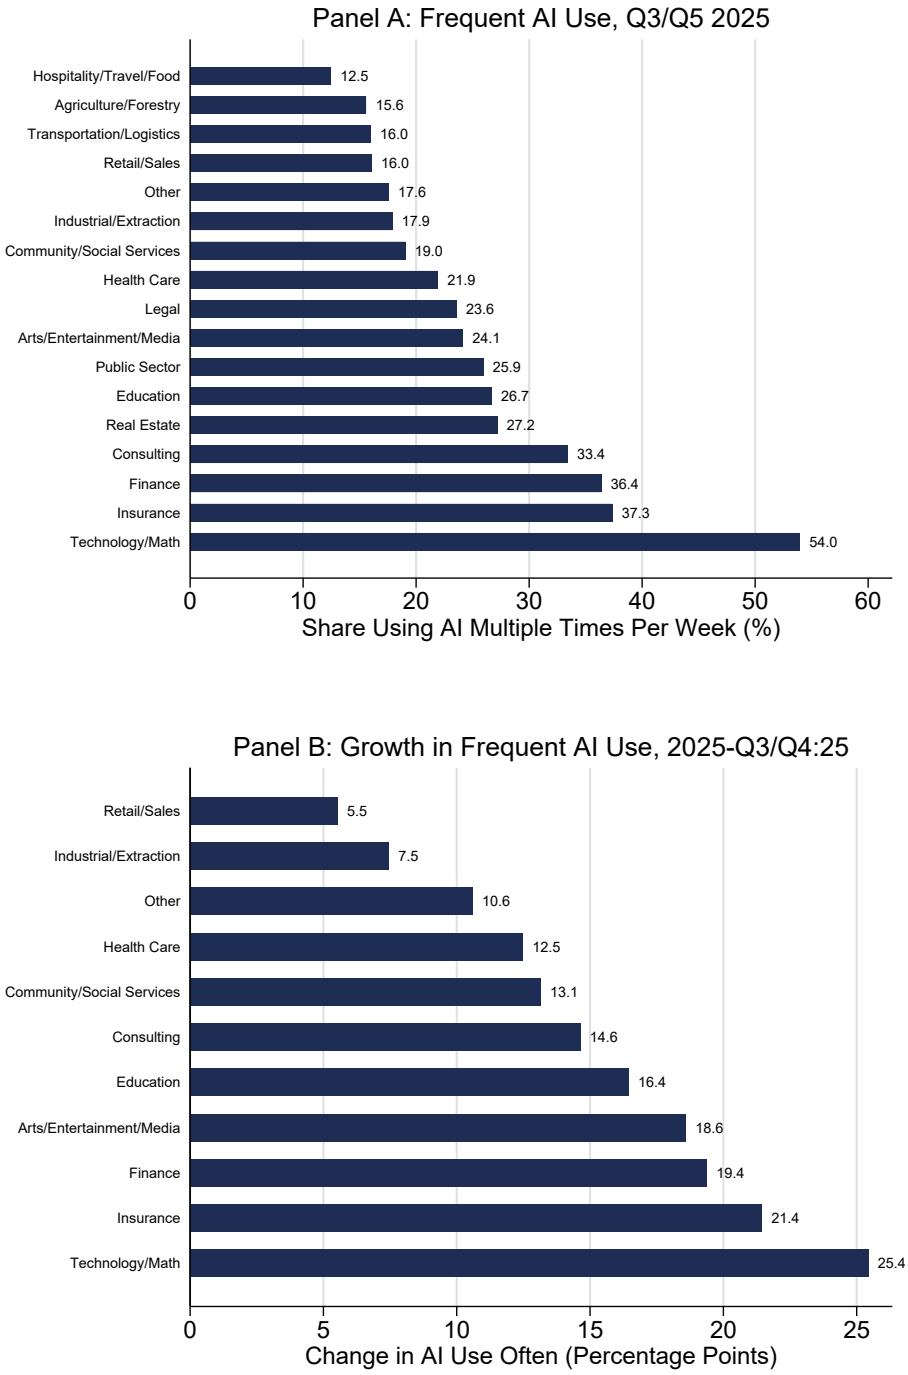
Notes.—Sources: QuantGov from Al-Ubaydli and McLaughlin (2017). The figure plots the regulatory restrictions per 1,000 words from 2022 to 2025, normalized to 2022, for finance/insurance and the rest of the economy.

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



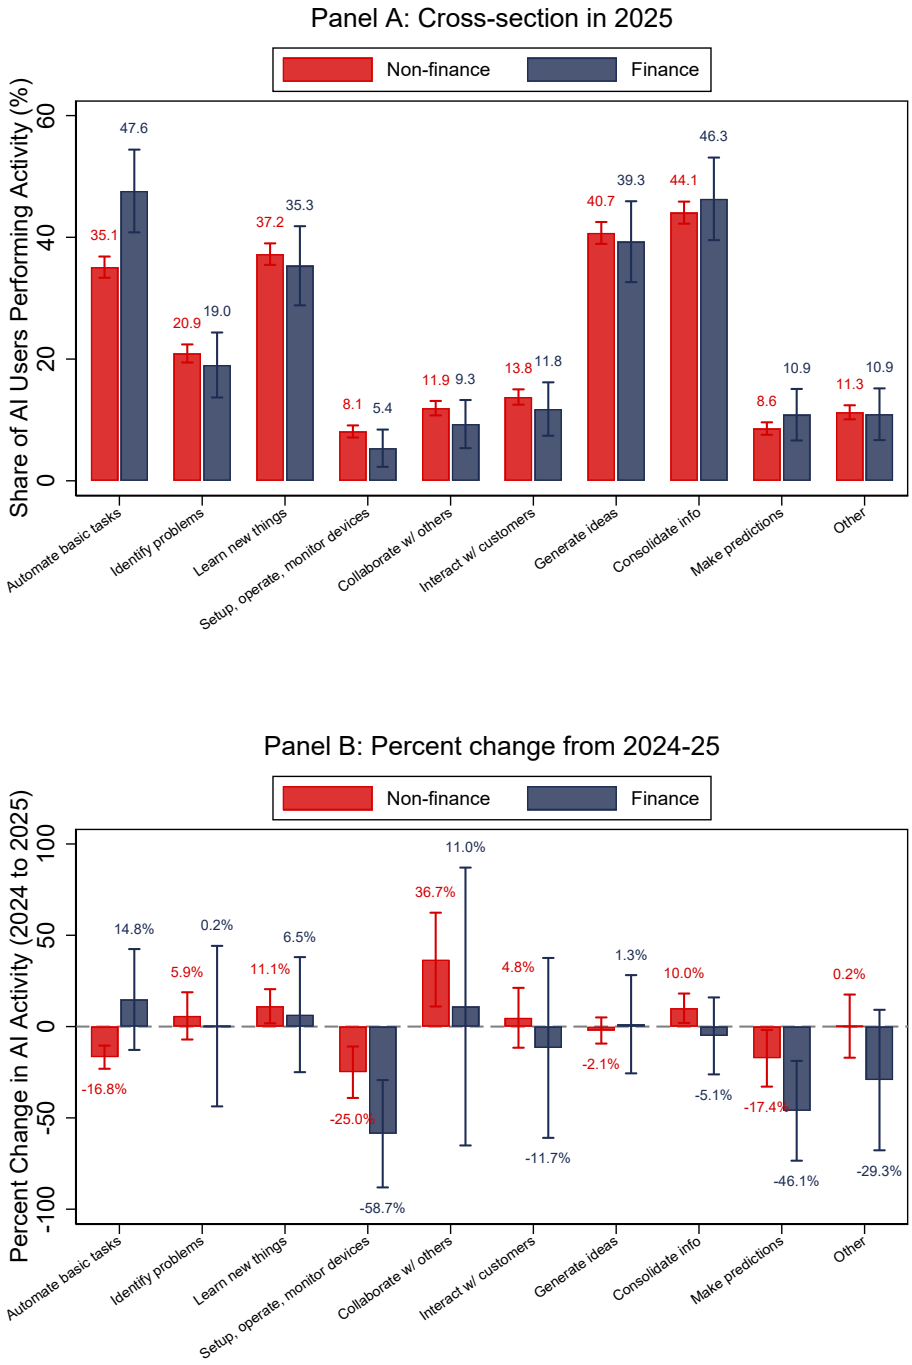
Notes.—Sources: Gallup Workforce Panel, 2023-2025. The figure plots the share of frequent AI users, defined as those who report using AI "multiple times a week" or "daily" to the question "How often do you use artificial intelligence in your role, separately for financial services and the economy as a whole. Other potential answers include: a few times a month, few times a year, once a year, less often than once a year, and never. Before asking this question, respondents are given a definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." The series is aggregated using sample weights.

Figure 3: Frequent AI Utilization Across Sectors, 2025 and 2024-25 Growth



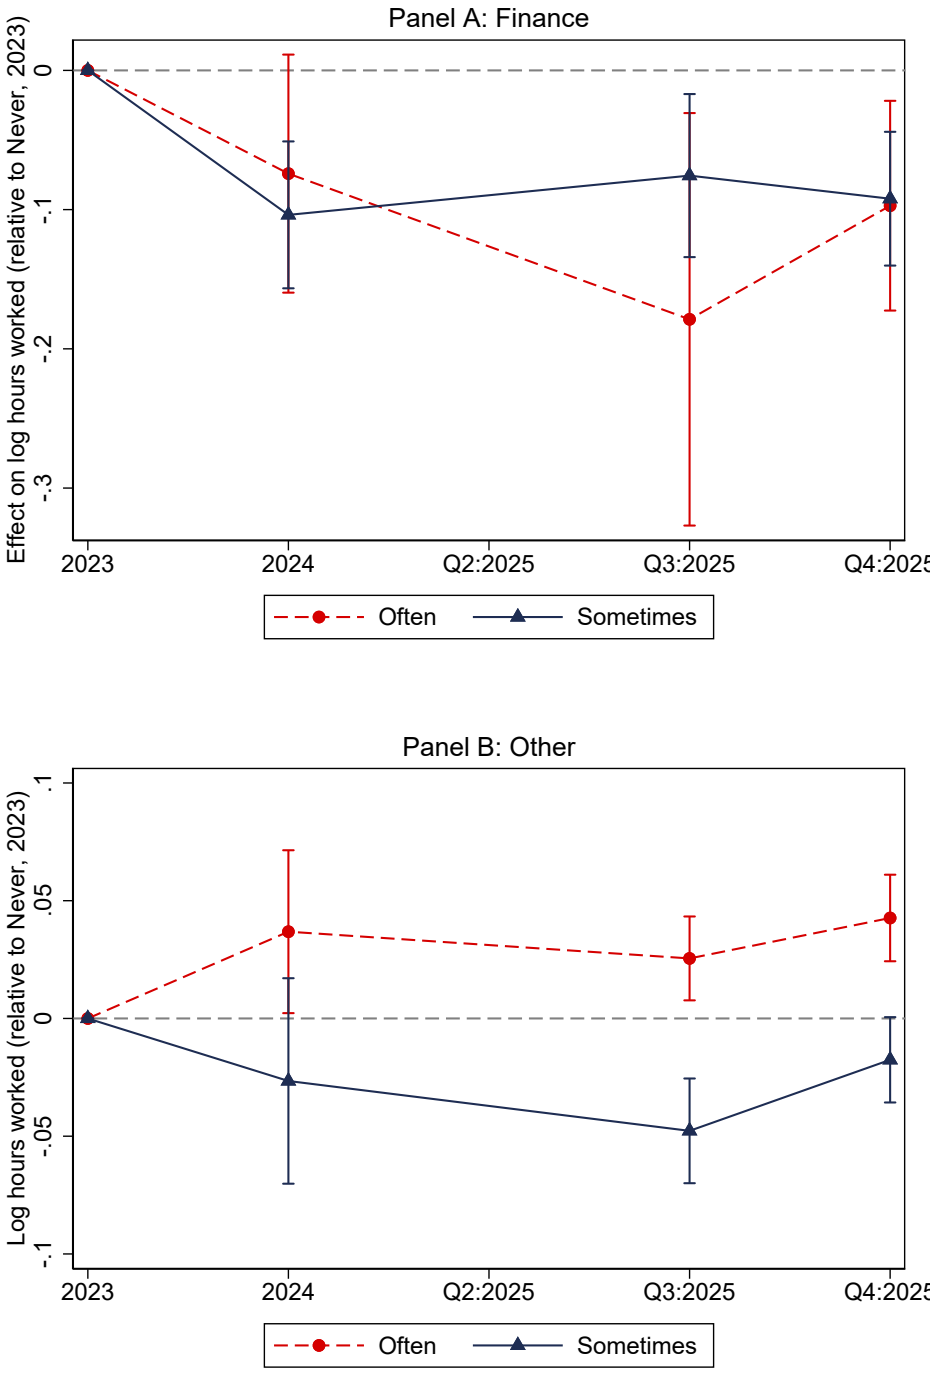
Notes.—Sources: Gallup Panel, 2025. Panel A plots the share of frequent AI users across all industries for Q3/Q4:2025. Panel B plots the growth in the share across 2024 to Q3/Q4:2025. Respondents answer the question: "How often do you use artificial intelligence in your role," which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. We consolidate these answers into three groups: often/frequent (daily and a few times a week), sometimes (a few times a month, a few times a year), and never (once a year, less often than once a year, and never). Before asking this question, respondents are given a definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." Panel B plots the interactions for non-finance workers. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Figure 4: Composition of AI Activities Among Finance and Non-Finance



Notes.—Sources: Gallup Panel, 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 25-65, and the observations are weighted.

Figure 5: Effects of AI Utilization on Hours Worked, Finance v. Non-finance



Notes.—Sources: Gallup Panel, 2025. Panel A plots the coefficients associated with regressions of log hours worked per week on AI utilization indicators interacted with time dummies and an indicator for being in finance, conditional on controls. "How often do you use artificial intelligence in your role," which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. We consolidate these answers into three groups: often (daily and a few times a week), sometimes (a few times a month, a few times a year), and never (once a year, less often than once a year, and never). Never is the omitted group. Before asking this question, respondents are given a definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." Panel B plots the interactions for non-finance workers. Demographic factors include: age, male, White, Black, some college, college or more, has children, and married. The sample is restricted to full and part time workers ages 25-65, and the observations are weighted.

Table 1: Examining the Determinants of Frequent AI Utilization

Dep. var. =	Using AI Frequently (1/0 Indicator)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	.283***			.260***	.108***			.077***
	[.010]			[.014]	[.012]			[.027]
× Finance/Insurance	-.021			-.009	-.038			-.050
	[.028]			[.041]	[.036]			[.058]
Customer Capital		.005		-.002		-.001		.019
		[.009]		[.009]		[.017]		[.018]
× Finance/Insurance		.031		.025		-.006		-.034
		[.034]		[.039]		[.050]		[.061]
log(Restrictions/1000 Words)								
× Finance/Insurance			3.775	-12.532			-12.565	-87.205**
			[10.855]	[25.771]			[13.897]	[38.940]
Adj. R2	.17	.08	.09	.15	.59	.50	.55	.61
N	38736	22247	43872	17464	28528	10157	33232	5895
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) *permissioning/governance*, proxied by an indicator for whether the respondent reports that their organization has a clear AI strategy; (ii) *customer-relationship capital*, proxied by an indicator for high customer capital (top-two box on the customer-delivery item); 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.

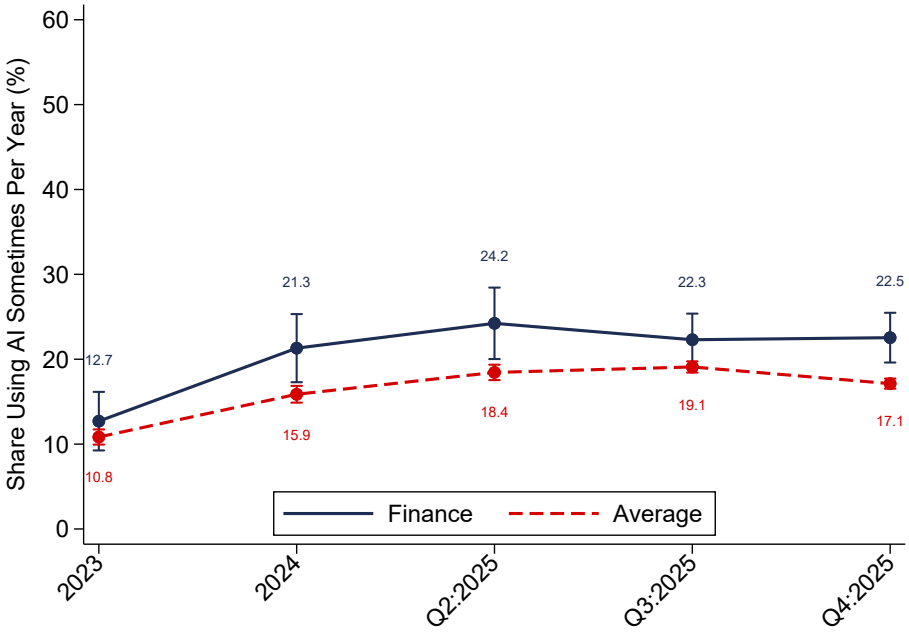
Table 2: Examining the Determinants of Frequent AI Utilization (Robustness Using Index)

Dep. var. =	AI Utilization Index (0-10 Index)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	2.988***			2.916***	1.169***			1.198***
	[.081]			[.110]	[.104]			[.217]
× Finance/Insurance	-.291			-.391	-.241			-.084
	[.213]			[.319]	[.276]			[.444]
Customer Capital		.050		-.047		.081		.245*
		[.080]		[.083]		[.137]		[.145]
× Finance/Insurance		.291		.246		-.204		-.332
		[.299]		[.342]		[.404]		[.499]
log(Restrictions/1000 Words)			-1.888	-10.400*			.093	-.353**
			[2.093]	[6.062]			[.084]	[.164]
× Finance/Insurance			13.728	-168.416			-109.498	-524.707**
			[94.868]	[202.714]			[115.914]	[234.384]
Adj. R2	.26	.13	.15	.25	.71	.60	.65	.72
N	38736	22247	43872	17464	28528	10157	33232	5895
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 AI utilization index on measures capturing three hypothesized channels for the rise of AI in finance: (i) *permissioning/governance*, proxied by an indicator for whether the respondent reports that their organization has a clear AI strategy; (ii) *customer-relationship capital*, proxied by an indicator for high customer capital (top-two box on the customer-delivery item); and (iii) *regulatory constraints*, proxied by the log of regulatory restrictions per 1,000 words matched at the NAICS2-by-year level. The AI frequency count is defined as follows: Daily=10, Few Times Week=8, Few Times Month=6, Few times Year=4, Once a Year=2, Less than Once a Year=1, Never/Other=0. 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.

A Online Appendix

Figure A.1: Sometimes AI Users in Finance and the Aggregate



Notes.—Sources: Gallup Workforce Panel, 2023-2025. 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," which include: (a) daily, (b) a few times a week, (c) a few times a month, (d) few times a year, (e) once a year, (f) less often than once a year, and (g) never. Before asking this question, respondents are given a definition of AI: "Artificial intelligence is when a computer does or enhances a task that is normally performed by a person." The series is aggregated using sample weights.

Table A.1: Placebo Determinants of Frequent AI Utilization Using Occupation

Dep. var. =	Using AI Frequently (1/0 Indicator)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Clear AI Strategy	.268***			.247***	.104***			.082***
	[.010]			[.014]	[.014]			[.030]
× Finance Role	.044			.045	-.009			-.076
	[.031]			[.045]	[.043]			[.072]
Customer Capital		-.004		-.011		-.008		.004
		[.009]		[.010]		[.018]		[.020]
× Finance Role		.015		.022		-.059		.006
		[.040]		[.042]		[.060]		[.083]
log(Restrictions/1000 words)			.161***	.049			-.123	-.547***
			[.043]	[.058]			[.121]	[.206]
× Finance Role			-.083	-.179			-.405	.566
			[.168]	[.215]			[.445]	[.362]
Adj. R2	.20	.11	.12	.19	.60	.51	.55	.60
N	32886	18881	37233	14821	23252	8264	27186	4762
Controls	Yes	Yes	Yes	Yes	No	No	No	No
SOC3 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) *permissioning/governance*, proxied by an indicator for whether the respondent reports that their organization has a clear AI strategy; (ii) *customer-relationship capital*, proxied by an indicator for high customer capital (top-two box on the customer-delivery item); 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 being in a finance-related occupation based on SOC3 classifications. Columns (1)–(4) include SOC3 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.

References

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.

Bertrand, M., Goldin, C., and Katz, L. F. (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, 2(3):228–255.

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.

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.

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

Célérier, C. and Vallée, B. (2019). Returns to Talent and the Finance Wage Premium. *The Review of Financial Studies*, 32(10):4005–4040.

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.

Dillon, E. W., Jaffe, S., Immorlica, N., and Stanton, C. T. (2025). Shifting work patterns with generative AI. Working Paper 33795, National Bureau of Economic Research.

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: An early look at the labor market impact potential of large language models. *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.

Gallup, Inc. (2025). Gallup Workplace Panel AI Use Module. Accessed January 2026.

Greenwood, R. and Scharfstein, D. (2013). The growth of finance. *Journal of Economic Perspectives*, 27(2):3–28.

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.

King, R. G. and Levine, R. (1993). Finance and growth: Schumpeter might be right. *The Quarterly Journal of Economics*, 108(3):717–737.

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

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

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. (2025). AI and Operational Losses: Evidence from US Bank Holding Companies. Technical report, Working paper.

Philippon, T. and Reshef, A. (2012). Wages and Human Capital in the U.S. Finance Industry: 1909–2006. *The Quarterly Journal of Economics*, 127(4):1551–1609.

Rajan, R. G. and Zingales, L. (1998). Financial dependence and growth. *The American Economic Review*, 88(3):559–586.

Stanford HAI (2025). The AI Index Report 2025. https://aiindex.stanford.edu/report/. Stanford Institute for Human-Centered Artificial Intelligence.

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