Liquidity for Transformation: Bank Credit and Firms' AI Adoption*

Sharjil Haque† Yusik Andrew Kim‡ Simon Mayer § Edson Wu¶

August 7, 2026

Abstract

We study the role of bank credit in firms' adoption of artificial intelligence (AI), combining U.S. supervisory loan-level data with firm-level AI-related job postings. We find a two-way association between bank credit and AI adoption: greater bank credit predicts subsequent adoption, while adoption is followed by more credit at lower spreads, primarily through credit lines. Adopters also exhibit lower credit-line utilization. To address endogeneity in the effect of AI adoption on credit, we instrument for adoption using firms' predetermined workforce exposure to generative AI interacted with the release of ChatGPT. We show that this credit expansion crowds out lending to non-adopters, occurs without evidence of superior information production about AI adopters, and is concentrated among banks specializing in this segment. Overall, our findings suggest that banks facilitate AI adoption by providing precautionary liquidity through credit lines during firms' transformation.

*Views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the Federal Reserve System. Ryan Jin, Jacob Mitovich, and Arturo Baldares provided stellar research assistance.

†Federal Reserve Board of Governors. Email: sharjil.m.haque@frb.gov

‡Federal Reserve Board of Governors. Email: andrew.kim@frb.gov

§Carnegie Mellon University. Email: simonmay@andrew.cmu.edu

¶Federal Reserve Board of Governors. Email: edson.ananwu@frb.gov

AI is transforming the economy, creating new ventures and disrupting existing business models. Although such innovation is often associated with startups, incumbents must transform their businesses and workforces to adopt AI or risk disruption. This transformation requires financing, which for many incumbents comes from banks. Yet banks may seem ill-suited to finance AI adoption: a long literature argues that they fare poorly in financing innovative projects, which are difficult to evaluate and often lack pledgeable collateral (Ueda, 2004; Dell'Ariccia, Kadyrzhanova, Minoiu and Ratnovski, 2021)). Against this backdrop, several questions arise. Can and do banks finance the adoption of this relatively new technology? If so, do they provide funding through term loans or liquidity insurance through credit lines? Do banks produce superior information about AI-adopting borrowers?

Banks generally finance firms through term loans and credit lines, which serve distinct but complementary functions. Term loans fund large investment outlays, whereas credit lines provide liquidity for ongoing financing needs and cash-flow shocks. Our findings suggest that banks facilitate AI adoption by providing liquidity insurance during firms' transformation, consistent with adoption being an ongoing process characterized by uncertain liquidity needs rather than a one-time investment. Crucially, banks' risk assessments appear less accurate for AI borrowers. Taken together, the concentration of credit expansion in credit lines and the weaker accuracy of banks' risk assessments indicate that banks' role in supporting AI adoption lies primarily in providing liquidity rather than producing superior information.

We combine supervisory loan-level data from the Federal Reserve's Y-14 with Lightcast job postings to measure AI adoption and its relationship with bank lending across more than 50,000 U.S. firms. The Y-14 data cover 2013–2025 and about 70% of corporate (C&I) lending but are limited to larger banks and loans with commitments of at least $1 million. Borrowers are therefore generally established firms—mostly private, middle-market companies—rather than startups; the data exclude venture debt and similar financing. We observe loan terms, bank-reported credit-risk measures, loan performance, and borrower financials.

Following Babina, Fedyk, He and Hodson (2024), we construct a human-capital-based

measure of AI adoption. Using Lightcast job postings, we create a time-varying firm-level measure equal to the share of a firm's postings related to AI, machine learning, or automation. We identify these postings using a classification that builds on Acemoglu, Autor, Hazell and Restrepo (2022). Our findings are robust to using only jobs related to generative AI. We show that AI adoption is concentrated among larger firms with fewer tangible assets and among firms in IT, services, technical industries, and finance. Our measure is also associated with broader workforce changes, suggesting that AI adoption in our setting reflects an ongoing transformation rather than a one-time change.

Consistent with this view, lagged bank credit is positively associated with subsequent adoption on both the extensive margin (new adoption) and the intensive margin (adoption intensity). Decomposing bank credit into term loans and credit lines, we find that credit lines primarily drive this relationship, providing initial evidence of their role in facilitating adoption. While lagged bank credit predicts subsequent AI adoption, an event-study design shows that firms rely more heavily on bank credit following their initial adoption. Taken together, these findings suggest a strong link between AI adoption and bank credit.

Our main empirical specifications regress bank-loan-level outcomes reported in the Y-14—such as spreads and committed loan amounts, loan type, and loan performance measures—on this measure of AI adoption. We include firm and bank×quarter fixed effects, as well as industry (two-digit NAICS)×quarter fixed effects. That is, we aim to estimate the effect of AI adoption on bank loan terms. However, AI adoption is endogenous, potentially biasing the OLS estimates. For example, unobserved factors—such as growth opportunities—may both affect AI adoption and access to bank credit.

To address these concerns, we propose a shift-share instrument for firms' AI adoption. The shift is the release of ChatGPT in November 2022; the share is firms' predetermined workforce exposure to this shock, capturing the extent to which LLMs can reduce task completion time (Eloundou, Manning, Mishkin and Rock, 2024). The instrument thus reflects both the extent to which a firm can be disrupted (requiring AI adoption to stay relevant)

and the extent to which adoption can benefit it by improving efficiency. The instrument is relevant, that is, positively associated with AI adoption intensity.

The exclusion restriction requires that, conditional on these fixed effects, exposure affects bank lending only through AI adoption. Note that ChatGPT's release was plausibly exogenous to individual firms and that firms' workforce exposure was predetermined. Industry×quarter fixed effects absorb industry-wide responses to the release, while firm fixed effects absorb persistent differences across firms with different exposure levels. Identification therefore comes from differences in AI adoption after ChatGPT's release between firms with greater and lesser preexisting exposure. Our results are robust to using this instrument.

We show that AI adoption is associated with an expansion in bank credit on more favorable terms. Adopting firms obtain larger loan commitments at lower spreads. At the firm level, adoption is also associated with higher leverage, driven primarily by greater bank debt. The expansion is concentrated in credit lines: adopters are more likely to obtain a credit line and have a larger credit-line share of total bank commitments. At the same time, they exhibit lower utilization rates than non-adopters.

This pattern is consistent with a precautionary motive, i.e., adopting firms obtain credit lines to absorb potential liquidity shocks during the transformation. Unlike a conventional investment requiring a large one-time outlay, AI adoption entails a broader workforce transformation involving hiring, ongoing reorganization, workflow redesign, and uncertain implementation costs. By providing funds that firms can draw on as needed, banks supply liquidity insurance throughout the adoption process.

Since banks expand credit to AI adopters at lower spreads, a natural question is how their information production and risk assessment evolve with firms' AI adoption. We leverage banks' own ex ante estimates of the probability of default (PD), reported in the Y-14 data. We find that banks' reported PDs predict subsequent defaults and other measures of loan performance, but are less predictive for AI adopters. Thus, banks expand credit to AI adopters despite having less accurate risk assessments. We therefore find no evidence that

banks produce superior information about AI borrowers relative to other borrowers.

We show that banks with greater prior exposure to AI-intensive borrowers act as specialists, extending larger commitments at lower spreads than less-exposed banks within the same firm and quarter. We measure specialization using the average commitment-weighted AI-adoption intensity of a bank's borrowers. These results indicate that credit provision to AI-adopting firms is concentrated among specialist banks. However, we find no robust evidence that specialist lenders possess an informational advantage: their risk assessments are no more predictive for AI adopters than those of other banks.

We find that lending to AI adopters crowds out lending to non-adopters. Exploiting variation across multiple lenders to the same non-adopting firm, we show that, in a given quarter, the firm receives smaller commitments from banks directing a larger share of new credit to AI adopters. These specifications include firm×time fixed effects, so the comparison is within borrower and period. This negative association is strongest among banks specializing in AI lending.

Taken together, our findings suggest that banks facilitate AI adoption through their role as liquidity providers rather than through superior information production. Both demand- and supply-side forces likely contribute to the expansion of credit to AI adopters, and we do not seek to fully disentangle them. On the demand side, AI adoption represents a broader workforce transformation that requires ongoing working capital rather than a single upfront investment. Firms therefore obtain credit lines as precautionary liquidity insurance.

On the supply side, AI adoption is associated with more bank credit at lower spreads, consistent with increased credit supply. One potential mechanism is that banks view adopters as a strategically valuable segment and compete for their business, extending larger commitments at lower spreads despite the difficulty of assessing their credit risk. This interpretation is consistent with the concentration of credit expansion among specialist banks and the crowding out of non-adopters given limited balance-sheet capacity. A complementary mechanism is that AI adoption may erode lenders' private informational advantages, weak-

ening the winner's curse that normally hampers loan competition (see, e.g., Marquez, 2002). This channel may further intensify competition and reduce spreads. Both mechanisms are consistent with the lower accuracy of banks' risk assessments in this market segment.

Related Literature. Whether banks finance innovation is a long-standing question; see, for instance, Ueda (2004), Winton and Yerramilli (2008), Minetti (2011), and Amore, Schneider and Žaldokas (2013), and, more recently, Bircan and De Haas (2020), Dell'Ariccia et al. (2021), Degryse, Roukny and Tielens (2026b), and Degryse, De Jonghe, Gambacorta and Huylebroek (2026a). This literature generally argues that banks are ill-suited to finance innovation because innovative projects are difficult to evaluate and often lack collateralizable assets. Limited collateralizability also constrains debt financing for intangible assets (Falato, Kadyrzhanova, Sim and Steri, 2022; He, 2026; Li, 2025). Certain assets, particularly patents, can alleviate this constraint (Mann, 2018; Hochberg, Serrano and Ziedonis, 2018), including by supporting non-bank venture debt (Mann and Gonzalez-Uribe, 2024). Relatedly, Berton, Mocetti, Presbitero and Richiardi (2018) study how credit-supply shocks affect employment and workforce composition, Fonseca and Van Doornik (2022) show that greater bank credit increases firms' skill intensity in Brazil, and Sforza and Acabbi (2025) links credit availability to within-firm labor reorganization.

Our contribution is to connect these literatures by showing that AI adoption is associated with greater reliance on bank credit, particularly credit lines. As a broad transformation process, AI adoption has financing needs distinct from those of other forms of innovation: it requires ongoing liquidity rather than a term loan funding a one-time investment. Banks can therefore play a key role in financing AI adoption, even if their role in financing innovation and intangible investment more generally is limited. Our findings suggest that banks facilitate AI adoption mainly through provision of liquidity, rather than through superior information production via screening and monitoring.

We further relate to work on banks' provision of liquidity insurance through credit lines (Holmström and Tirole, 1998; Acharya, Almeida, Ippolito and Perez, 2014; Brown, Gustafson

and Ivanov, 2021; Chodorow-Reich, Darmouni, Luck and Plosser, 2022; Haque, Mayer and Stefanescu, 2024; Greenwald, Krainer and Paul, 2025). This literature emphasizes that credit lines allow firms to absorb cash-flow shocks, unlike term loans, which primarily finance investment outlays. We extend this insight by highlighting how bank-provided liquidity facilitates new technology adoption, a broader transformation characterized by ongoing and uncertain financing needs.

Our findings are consistent with banks expanding credit to AI adopters to build relationships in a potentially lucrative segment. This interpretation echoes the relationship-banking literature, which shows that banks offer favorable terms or invest in early-stage firms to establish borrower relationships (Boot and Thakor, 2000; Petersen and Rajan, 1994; Rajan, 1992; Hellmann, Lindsey and Puri, 2008). Our findings also relate to the literature on bank specialization (Paravisini, Rappoport and Schnabl, 2023; Blickle, He, Huang and Parlatore, 2025; Blickle, Parlatore and Saunders, 2026), although we find no evidence that specialist lenders possess an informational advantage.

We connect to the literature on AI and finance. Babina et al. (2024) measure firms' AI investments using job postings; we build on their approach while focusing on how adoption is financed. While many studies focus primarily on large publicly traded firms, which account for only about a quarter of aggregate U.S. employment (Babina et al., 2024; Eisfeldt, Schubert and Zhang, 2023), our data allow us to study AI adoption across more than 50,000 firms, including many private, middle-market firms. Gandhi, Lu, Pan, Plazzi and Wei (2026) study how firms' AI adoption is priced in equity and debt markets, while Buchetti, Eufinger, Lin and Miquel-Flores (2026) examine banks' risk-assessment response to adoption. We instead examine how banks facilitate the AI adoption process through credit lines.

1 Data and AI Adoption Measures

For our analysis, we combine data on bank loans and borrower financials from the Y-14 credit register, with job postings data from Lightcast, where the latter data source allows us to construct a firm-level measure of AI adoption. All variables are defined in Appendix A.1. Our sample period is 2013Q1-2025Q4.

1.1 Federal Reserve's Y-14 Data

Our primary data source is the Federal Reserve's FR Y-14Q H.1 schedule on commercial loans (commonly referred to as the Y-14 data).[1] The Y-14 data provides detailed information on all bilateral and syndicated loan facilities with committed amounts over $1 million held by Bank Holding Companies (BHCs) subject to the Federal Reserve's stress tests.[2] The Y-14 reporting banks hold over 85% of total U.S. banking assets (Caglio, Darst and Kalemli-Özcan, 2021) and account for roughly 70% of all commercial & industrial (C&I) lending (Minoiu, Zarutskie and Zlate, 2021).

The Y-14 data provides a granular view of loan contracts across borrowing firms on a quarterly basis. The dataset includes time-varying information on committed exposures, utilized exposures, interest rates, spreads, maturity, collateral, bankruptcy priority, ex-ante estimates of probability of default, loan-type (e.g., credit line or term loan), and industry information through NAICS codes. We identify new loans from existing ones using the loan's origination date. Crucially we also observe ex-post loan performance measures such as flags for loans that banks report as "non-accruing," number of days a loan payment is past due, as well as the dollar amount of a facility that banks charge off. Banks also report annual financial information for borrowers. Our analysis relies on both the quarterly loan-level data

[1] For details on variables contained in schedule H.1 and reporting requirements, see the table beginning on page 170 in the publicly available reporting form.

[2] A loan facility is a lending arrangement between a bank and a borrower that may encompass multiple loans of different types, such as credit lines or term loans. Banks categorize the facility type based on the loan type that represents the majority of total commitment amount.

and annual borrower-level financials.[3] Although reporting began in 2011Q3, we start our sample in 2013Q1 when bank coverage stabilized and data structures and variables became consistent. Appendix A.2 details our data cleaning and filtering procedures.

1.2 AI Adoption Measure from Lightcast

Our second key data source comes from Lightcast (formerly Burning Glass Technologies), a commercial provider that collects and parses the near-universe of online job postings in the United States, and has been widely used to study skill demand across firms and labor markets (Brown and Matsa, 2016; Hershbein and Kahn, 2018; Deming and Kahn, 2018; Acemoglu et al., 2022; Babina et al., 2024). Lightcast collects data from roughly 40,000 company websites and online job boards, with no more than 5% of vacancies from any one source. As also described in Babina et al. (2024), Lightcast uses a sophisticated deduplication algorithm to avoid double counting vacancies that post on multiple job boards. Crucially for our study, the data contain detailed information for each job posting, including job title, job location, occupation, and employer name. Furthermore, each job posting is tagged with thousands of specific skills standardized from the open text in each job opening, providing extensive detail of the individual job postings that allow us to identify AI-related jobs.

Similar to Babina et al. (2024), we construct a human-capital based measure of AI adoption at the firm level. Following the job-postings approach of Acemoglu et al. (2022), we construct a firm-quarter measure of AI adoption from firms' vacancy postings. We classify a posting as AI-related if its text contains at least one term from a fixed dictionary adapted from the (narrow) AI skill list in Acemoglu et al. (2022). The full dictionary appears in Appendix A.3. We aggregate these posting-level indicators to the firm-quarter level and measure adoption intensity as the share of a firm's postings classified as AI-related. Related vacancy-based measures of demand for AI-skilled labor are used by Alekseeva, Azar, Giné, Samila and Taska (2021) and, as a complementary measure, by Babina et al. (2024).

[3]Borrower financials are available for about 60% of firms, with higher coverage among larger borrowers.

For each firm j in quarter t, we compute the 3-year rolling AI posting share:

$$\text{AI}_{j,t} = \frac{\sum_{s=t-12}^{t-1} \#\text{AI}_{j,s}}{\sum_{s=t-12}^{t-1} \text{Postings}_{j,s}} \times 100 \tag{1}$$

where $\#\text{AI}_{j,s}$ is the count of firm j's job postings requiring AI skills in quarter s and $\text{Postings}_{j,s}$ is its total posting count. The classification on whether a job requires AI skills follows Acemoglu et al. (2022). The measure is strictly lagged (excludes the current quarter) to avoid simultaneity. Firms with no postings in the window have $\text{AI}_{j,t} = 0$. The 3-year rolling window balances two concerns: shorter windows (1 year) are noisy because AI hiring is lumpy; longer windows (cumulative) are slow to respond to changes in AI strategy. We verify robustness to 1-year rolling and cumulative alternatives. In Appendix Section A.3, we discuss the taxonomy in greater detail and the definition of AI skills.

One concern with our approach is that job postings capture demand for AI skills but not whether these positions are ultimately filled, potentially introducing measurement error. Evidence from prior studies mitigates this concern. First, Acemoglu et al. (2022) show that the occupational and industry composition of Lightcast postings closely aligns with actual employment shares from the Occupational Employment Statistics (OES). Second, Babina et al. (2024) show that résumé-based measures of realized AI hiring are highly correlated with Lightcast posting-based measures and yield similar results for firm performance and employment. Accordingly, Babina et al. (2024) and Babina (2026) conclude that job postings provide valid data for studying the effects of AI adoption.

1.3 Exposure to GPT Release in 2022

In much of our empirical analysis that looks at the association between AI adoption and bank lending, AI adoption will serve as an explanatory variable. Since AI adoption is endogenous, we aim to construct an instrument for it based on the release of ChatGPT in November 2022, as well as its potential to transform industries.

To preview, we propose a shift-share IV, explained in greater detail later in Section 3.1.2, where we exploit the unexpected release of ChatGPT in 2022 as the shift and firms' pre-determined workforce exposure to this disruption as the share. To construct a measure of firm-level workforce exposure to GenAI disruption, we rely on occupation-level AI exposure scores. We use the β-threshold from Eloundou et al. (2024), which scores each of 923 O*NET occupations on the degree to which LLMs (with complementary software) can reduce task completion time. Scores range from 0 to 1.

For each firm j, we calculate

$$E_j = \frac{\sum_k \text{Postings}_{j,k} \cdot \beta_k}{\sum_{k'} \text{Postings}_{j,k'}} \tag{2}$$

where $\text{Postings}_{j,k}$ is firm j's stock of Lightcast postings in O*NET occupation k over the fixed window 2018Q1–2022Q3, and β_k is the Eloundou et al. (2024) β score for occupation k. The window ends before the release of ChatGPT, so the occupational shares are predetermined with respect to the shock and cannot respond to it.[4]

Consequently, E_j captures firm j's exposure to potential disruption from the release of ChatGPT in 2022 and the subsequent GenAI boom. As we also discuss later in detail (see subsection 3.1.2), it is ex-ante unclear whether such exposure is positive or negative. On the negative side, high exposure could reflect disruption risk from GenAI. On the positive side, it could also reflect the potential for efficiency gains. We will show that higher exposure is positively associated with AI adoption following the ChatGPT release.

1.4 Merged Sample and Summary Statistics

Y-14 and Lightcast lack common firm identifiers. We use the FedMatch algorithm from Cohen, Dice, Friedrichs, Gupta, Hayes, Kitschelt, Lee, Marsh, Mislang, Shaton et al. (2021) to match borrowers by name and industry (NAICS code), followed by manual verification. This

[4]Alternative constructions, such as an accumulating-from-first-posting version and an earlier 2018–2020 window that ends two years before ChatGPT, deliver similar results in our analysis.

process yields 53,648 unique firms across a broad range of sectors, as shown in Appendix Table A.1. These firms represent approximately 47 percent of all unique corporate borrowers that ever appear in the Y-14 dataset between 2013Q1-2025Q4. Aggregate bank-loan commitments to matched firms rose steadily from $50 billion at the beginning of the sample to around $650 billion at the end, underscoring the sample's economic importance. To our knowledge, this is the largest firm-level dataset linking AI-related job postings, bank loans, and borrower financial information.

Table 1 provides detailed summary statistics for our baseline loan- and firm-level samples. Consistent with prior studies, the mean loan observed in the Y-14 data has a commitment amount of $24 million and an interest-rate spread of approximately 2 percentage points over the base rate. The median loan has a commitment amount of around $10 million and a maturity of five years. Both panels show that a large share of loans are structured as (revolving) credit lines, which provide liquidity insurance to firms seeking to finance working capital or cash-flow shortfalls. The share of loan observations that are credit lines is about 70%. On the firm level, the share of credit lines to committed credit is close to 80%.

Crucially, the median firm has assets of around $32 million, while the mean is much larger, confirming that the typical Y-14 firm is considerably larger than a traditional startup. For example, Robb and Robinson (2014) show that typical venture capital targets have book assets below $100,000, although their study covers a different sample period. Thus, the borrowers in our sample are generally neither startups nor VC-backed firms. They are better understood as mature companies adopting AI to optimize their existing workforces, and the loans we examine should not be considered venture debt. About 5% of the firms in the Y-14 are publicly listed firms. As such, many firms in our data are private, middle-market firms. However, in private capital markets, they would be considered LBO or growth-capital targets rather than firms financed by VCs.

2 The Characteristics of AI Adoption

2.1 AI Adoption as Workforce Transformation

We measure AI adoption using hiring data. We show that AI adoption in our setting captures a workforce transformation process. Unlike a large one-time investment, such a transformation requires ongoing financing to cover working-capital needs. This suggests that AI adoption may create demand for continued liquidity support through credit lines, a point we elaborate upon throughout the remainder of the paper.

The details of this analysis are provided in Appendix A.6, while we provide a concise summary here. To capture this transformation, we track changes in a firm's occupational hiring mix, defined as the distribution of its job postings across detailed occupations (six-digit O*NET codes, the standard U.S. occupational taxonomy). We use three measures of how much this mix changes over a three-year window.

The first, a dissimilarity index, is one-half the sum of the absolute changes in occupation shares; it measures how much a firm reshuffles hiring across the occupations it already employs. The second, the cosine distance between the firm's occupation-share vectors at the start and end of the window, measures how much the composition of hiring shifts, independent of its scale. The third, the share of postings in occupations the firm had never posted before, measures hiring into entirely new roles. The first two capture how much a firm rearranges its existing workforce; the third, how much it moves into new activities. Thus, all of these measures capture workforce transformation.

We then regress all three measures on a firm's AI adoption, where the results are presented in Appendix Table A.2. Notably, all three measures are positively associated with AI adoption. This indicates that AI-adopting firms reorganize their workforces rather than simply expand them. Intuitively, AI adoption likely captures automation or organizational change aimed at reducing costs or improving the efficiency of existing operations. At the same time, this association is not surprising, as we measure AI adoption through job postings:

hiring for new technologies naturally changes the composition of a firm's workforce.

2.2 Firm Characteristics and AI Adoption

We examine the link between firm characteristics and AI adoption. To this end, we regress firms' AI adoption—measured by the change in AI posting intensity and an indicator for whether a firm newly adopts AI—on lagged firm financials and the lagged level of total bank credit. The outcomes capture both the extensive margin (whether a firm newly adopts) and the intensive margin (the change in adoption intensity). The specifications include industry×quarter fixed effects; our industry classification is NAICS two-digit sectors. Because we do not include firm fixed effects, the results should be interpreted as reflecting cross-sectional differences among firms with different levels of AI adoption.

We emphasize that all regressors are lagged by only one year. On the intensive margin, the dependent variable is the year-over-year change in AI adoption, which captures incremental adoption during the observation year. The new-adoption indicator is based solely on whether a firm begins adopting AI in the current quarter. In both cases, the lagged firm-level regressors are determined before the outcome. The results are presented in Table 2.

First, firm size, measured by book assets, is positively associated with AI adoption, while asset tangibility is negatively associated with it. Thus, larger, less tangible-intensive firms are more likely to adopt AI. This is intuitive, and in line with the fact that many adopters operate in service or software industries. Indeed, Appendix Table A.1 shows that adoption is most prevalent in two-digit NAICS sectors related to information technology and professional and technical services, both characterized by low tangibility, as well as in finance and insurance. It is least prevalent in construction, accommodation, and agriculture, that is, in more tangible-heavy industries.

Second, Table 2 shows that lagged total book debt (in logs) is negatively associated with subsequent AI adoption, while the debt-to-assets ratio is not statistically significant. Liquidity, measured by cash holdings, is positively associated with subsequent adoption. In

contrast, bank debt—whether scaled by assets or expressed in logs—consistently predicts future AI adoption. Although this contrast may initially appear surprising, our following analysis suggests that firms' reliance on bank debt reflects their use of credit lines as liquidity insurance rather than term loans to fund large investments.

Specifically, we show both that credit-line commitments predict subsequent AI adoption and that adoption is followed by greater reliance on credit lines. Together, these findings are consistent with banks providing liquidity that firms use to meet ongoing and precautionary needs during the transformation.

2.3 Bank Credit (Lines) and AI Adoption

We now examine the role of bank credit in financing AI adoption in greater detail. We provide both graphical and regression-based evidence on the importance of bank debt—particularly credit lines—for adoption. Banks typically finance firms through term loans and credit lines, which serve distinct but complementary functions. Term loans fund large investment outlays, whereas credit lines provide liquidity to meet ongoing financing needs and absorb cash-flow shocks. Our findings suggest that banks support AI adoption primarily through credit lines and the liquidity they provide.

Regression Evidence. We start with firm-level regressions examining whether lagged bank credit predicts AI adoption. As in the previous sections, we regress two measures of AI adoption—the year-over-year change in adoption intensity and an indicator for new adoption—on lagged bank credit. We first consider total bank credit and then decompose it into term loans and credit-line commitments. Some specifications additionally control for total debt. All debt amounts are expressed in logs. The regressions include firm fixed effects, firm-level controls (total assets in logs, tangibility, EBITDA/assets, and liquidity), and sector-quarter fixed effects. The firm fixed effects control for firms' average reliance on bank credit. As such, our estimates effectively capture the association between changes in bank credit and AI adoption.

Table 3 presents the results. Columns (1)–(3) use the change in AI intensity as the outcome, while columns (4)–(6) use the new-adoption indicator. Columns (1) and (4) show that lagged bank credit is positively associated with subsequent AI adoption, consistent with firms using bank credit to help finance adoption.

In columns (2) and (5), we decompose bank credit into credit-line commitments and term loans. Credit-line commitments are positively and significantly associated with subsequent AI adoption in both specifications. By contrast, the coefficient on term loans is insignificant in column (2) and, although significant in column (5), smaller in magnitude than that on credit lines. Thus, the association between bank credit and subsequent AI adoption is concentrated in credit lines.

Finally, columns (3) and (6) include lagged total debt. Its coefficient is positive but smaller than that on credit-line commitments and insignificant in column (3), reinforcing the particularly strong association between credit lines and subsequent AI adoption.

Graphical Evidence. Next, we conduct an event study of firms' reliance on bank debt and credit lines around their first recorded AI adoption, including firm and industry×quarter fixed effects. Figure 1 presents the results, with the borrower's total committed bank credit in the left panel and total committed credit-line amount in the right panel.

Figure 1 shows that both total bank debt and credit-line commitments increase around the initial adoption date and remain elevated thereafter. This pattern suggests a persistent rise in firms' reliance on bank debt, especially credit lines, once they start adopting AI. Thus, while the regression evidence shows that higher bank credit predicts subsequent AI adoption, the event study highlights firms' increased reliance on bank credit following adoption.

Taken together, the findings suggest that banks support AI adoption by providing liquidity through credit lines. Hiring-based AI adoption does not necessarily require a large lump-sum investment; instead, it creates ongoing working-capital needs during the transformation process, for which credit lines may be particularly well suited.[5]

[5]Intuitively, investment in AI differs from technologies that require capital investments, which may require

3 AI Adoption and Bank Credit Supply

We next examine in greater detail how AI adoption is associated with firms' access to and reliance on bank credit, as well as the terms on which that credit is provided. To do so, we regress bank-loan characteristics on AI adoption. Thus, the empirical design in Section 3 differs from the regressions in Tables 2 and 3: whereas those regressions treat AI adoption as the outcome and bank credit as the key regressor, we now treat AI adoption as the key regressor and bank-credit measures as the outcomes.

3.1 Main Empirical Specification

We employ both an OLS specification and an instrumental-variables design.

3.1.1 OLS

We employ the following OLS specification, where the dependent variable is a bank-loan-level outcome at observation date t (the unit of time is a quarter, consistent with the reporting frequency in the Y-14 data):

$$Y_{i,b,t} = \beta \cdot \mathrm{AI}_{j,t} + \gamma' X_{j,t} + \alpha_j + \alpha_{\mathrm{NAICS2}\times t} + \alpha_{b\times t} + \varepsilon_{i,b,t}, \tag{3}$$

where i indexes loans, b banks, j firms, and t quarters. We include firm fixed effects (α_j), which account for time-invariant firm-level factors. We also include industry (NAICS 2-digit)×quarter fixed effects ($\alpha_{\mathrm{NAICS2}\times t}$), accounting for time-varying industry level factors and shocks. Finally, we also include bank×quarter fixed effects ($\alpha_{b\times t}$), which account for time-varying bank level characteristics (including their capital constraints and deposit growth). The coefficient of interest is β, which captures the differential effect of AI adoption on bank loan outcomes, $Y_{i,b,t}$.

Given our fixed effects, our loan-level regressions are effectively comparing, within a given

term loans or bonds, such as industrial robots (Benmelech and Zator, 2025).

quarter, loans made by the same bank to borrowers in the same industry that primarily differ in their level of AI adoption. Other specifications employ Y-14-reported firm-level outcomes, which are available at an annual frequency.

3.1.2 Instrumental Variable

Firms' AI adoption is endogenous, which may bias the OLS estimate of the effect of AI adoption. First, unobserved factors related to firm quality or growth opportunities could affect both a firm's AI adoption decision and its access to (and demand for) bank credit; this would overstate the association between AI adoption and bank credit. Second, firms might adopt AI when facing operational or efficiency problems, or competitive pressure, that also limit their access to bank credit; this would bias the association downward.

To mitigate these endogeneity concerns, we employ throughout a shift-share instrument for AI adoption, using firms' predetermined exposure to generative AI as of 2022 as the "share" and the release of ChatGPT as the "shift." The predetermined exposure of firm j, denoted by E_j and constructed in Section 1.3, captures the extent to which generative AI and LLMs can reduce task-completion times for the firm's workforce and thus its exposure to this technology. We then construct the instrument as:

$$Z_{j,t} = E_j \times \mathbf{1}[t \geq 2022\text{Q4}] \tag{4}$$

The post period begins in 2022Q4 (ChatGPT was released on November 30, 2022). Before 2022Q4, $Z_{j,t} = 0$ for all firms: the technology shock has not yet arrived. After 2022Q4, cross-sectional variation in E_j generates variation in treatment intensity. We then estimate a two-stage least squares (2SLS) specification in which we instrument for AI adoption.

The instrument can be viewed as a shock to firms' potential gains and losses from AI. On the positive side, it captures the potential gains from adoption, for instance, through cost reductions from automation or other efficiency gains. On the negative side, it may also

capture disruption risk, as a higher value indicates that the firm's services can be more easily automated, allowing competitors to enter and offer the same services more cheaply. Both forces may induce adoption: firms may adopt AI to realize its potential gains or to remain competitive in the face of disruption. Crucially, the instrument is relevant: in the first stage, a higher instrument value is associated with greater AI adoption, with F-statistics as high as 100 and consistently above 40 across all specifications.

Regarding the exclusion restriction, the instrument is valid as long as there is no other channel, conditional on controls and fixed effects, through which it affects bank-loan-level outcomes. In our context, the shift (the ChatGPT release) is plausibly exogenous, while the share is predetermined. Moreover, our fixed effects absorb the most natural competing channels. Since LLM exposure is highly correlated across firms within an industry, the NAICS2×quarter fixed effects absorb industry-specific shocks through which the ChatGPT release could affect bank credit supply, for instance, banks repricing credit to exposed industries as a whole. The firm fixed effects absorb any persistent differences between more and less exposed firms. Identification thus comes from within-industry, within-firm variation in exposure, in that our controls and fixed effects account for common industry-level responses to the arrival of generative AI.

3.2 Amount, Pricing, and Loan Terms

We examine the association between bank lending and AI adoption using both our OLS specification, in which we regress loan-level outcomes on the firm-level measure of AI adoption, and the IV specification; our results are robust across the two. Specifically, we run regression (3), alongside the IV 2SLS regression, with committed loan amount and loan spread—both expressed in logarithms—as outcome variables. In other words, we examine the association between AI adoption and the quantity and pricing of bank credit.

Table 4 presents the results. Across all specifications, AI adoption is associated with lower loan spreads and larger committed amounts. These findings hold in both the OLS and

IV specifications, although the IV estimates are larger in magnitude. This difference could indicate that the OLS estimates are biased toward zero. For example, unobserved factors—such as competitive pressure or operational inefficiencies—may simultaneously encourage AI adoption and impair firms' access to bank credit. A firm may adopt AI precisely because it faces competitive pressure, which also makes banks less willing to lend. Alternatively, the larger IV estimates may reflect a local average treatment effect (LATE) for firms whose adoption decisions respond to the instrument.

Our results suggest that AI adoption is associated with increased bank credit at lower spreads. The simultaneous increase in loan commitments and decrease in spreads is consistent with an outward shift in credit supply. However, we acknowledge that we cannot cleanly disentangle supply- and demand-side explanations. Indeed, as we argue below, the increase in bank credit is also consistent with firms' greater need for liquidity insurance.

Additional Results. We present two sets of additional results that complement our findings on AI-adopting firms' increased reliance on bank debt. First, we confirm this pattern in firm-level regressions examining the debt-to-assets ratio and the total amount of bank credit (in logs). Table 5 presents estimates from OLS and IV specifications in which we regress these outcomes on our measure of AI adoption. Columns (1) and (2) show that AI adoption is associated with higher leverage. Columns (3) and (4) show that it is associated with greater levels of bank debt, indicating increased reliance on bank credit.

Second, we have shown that AI adoption is associated with lower loan spreads. A similar pattern appears regarding loans' non-price terms, such as seniority. Appendix Table A.3 shows that AI adoption is associated with a greater likelihood that a bank loan is unsecured or junior, suggesting that borrowers obtain less stringent non-price terms in addition to lower spreads. Within firms, however, these non-price margins are at most marginally significant.

3.3 The Role of Bank Credit Lines

A natural question is whether the increased credit provision to AI-adopting borrowers occurs through term debt or credit lines, which serve different purposes. Term loans are typically used to finance lump-sum investments, while credit lines are used as liquidity to cover cash flow shocks. These two instruments are commonly bundled (Haque et al., 2024). We now examine the association between firms' AI adoption and the type of bank credit. We first look at credit line commitment amounts, before analyzing their utilization.

Credit Line Commitments. To this end, we estimate our baseline regression (3) with two new outcome variables that capture firms' reliance on credit lines. The first is an indicator for whether a given loan is a credit line (as opposed to a term loan or another loan type). The second is the share of credit line commitments in a given borrower's total bank commitments. In columns (3) and (4), we additionally weight loan observations by their commitment size when employing the indicator. As before, we present OLS and IV estimates alongside each other. Note that columns (1)-(4) run regressions on the bank-loan-time level, while columns (5)-(6) aggregate data to the firm-time level. Table 6 presents the results of these regressions in columns (1)-(6).

The coefficient on AI adoption is positive and significant in columns (1)–(5) and retains the same sign, though it is no longer significant, in column (6). Notably, this result holds across all outcome variables capturing firms' reliance on credit lines. As before, the IV estimates have the same sign as the OLS estimates, but tend to be larger in magnitude than the OLS estimates, consistent with either a LATE interpretation or OLS estimates that are biased toward zero. Overall, we find that AI adoption is associated with greater credit-line commitments.

Credit Line Utilization. We next examine the link between AI adoption and credit-line utilization, using the utilized share of a given credit line as the outcome variable. Columns (7) and (8) of Table 6 present the results. The coefficient on AI adoption is negative and

significant in both specifications. Thus, AI adoption is associated with both larger credit-line commitments and lower utilization rates.

Taking stock, our findings suggest that AI-adopting firms use credit lines primarily as precautionary liquidity insurance rather than to fund a large lump-sum investment. Adoption as we measure it involves hiring specialists and sustaining ongoing expenditures and working-capital needs throughout the transformation process. Banks therefore appear to facilitate AI adoption by providing committed liquidity that firms can draw on as needed. Crucially, the combination of larger commitments and lower utilization is consistent with firms securing precautionary liquidity during the transformation.

3.4 Banks' Information in AI Lending

Our results indicate that banks provide credit lines to AI-adopting firms on favorable terms. A natural question is what drives this credit provision. One possibility is that banks possess superior information about AI adopters.

We examine this possibility using bank-reported probabilities of default (PDs), which capture banks' risk assessments. We conduct two analyses. First, using a linear regression design similar to (3), we show that reported PDs predict realized default—that is, ex-post loan performance—but that their predictive power weakens for AI-adopting firms.

Second, we quantify this weaker predictive power using a rescaling-invariant metric standard in the bank information-production literature (Weitzner and Howes, 2023): the area under the ROC curve (AUC) of bank-assigned PDs in predicting realized default, computed separately for AI and non-AI borrowers.

Overall, our findings suggest that banks' information in AI lending is no better than in other types of lending. If anything, their risk assessments are less accurate for AI adopters, consistent with banks either possessing less information or producing less of it. Thus, banks' expansion of credit to AI adopters does not appear to be driven by an informational advantage or particularly intensive information production. Instead, the findings are consistent

with banks supporting AI adoption through liquidity provision rather than superior information production.

Bank Estimates of Default. Banks report estimated probabilities of default (PDs) in the Y-14. We assess the extent to which these estimates predict actual loan performance. To this end, we run OLS regression (3) and include the PD as well as an interaction between the PD and our AI adoption measure. As outcome variables, we employ indicators of loan performance: whether the loan defaults at some point in the future, is placed on non-accrual status, or is past due (by 30 or 90 days). As before, we employ both the OLS regression and the 2SLS regression, in which we instrument for AI adoption.

Table 7 presents the results, which are similar across specifications. The coefficient on the PD is positive and significant. Accordingly, the bank-estimated default probability—an ex-ante measure—predicts actual ex-post default. However, the interaction term between the AI adoption measure and the PD is negative and significant: for AI-adopting borrowers, a given increase in the bank's assigned PD translates into a smaller increase in realized default. In other words, higher AI adoption weakens the predictive power, and hence the informativeness, of bank-estimated PDs.

ROC Curves and Area under the Curve (AUC). Next, we assess this weaker predictiveness using the rescaling-invariant metric standard in the bank information-production literature (Weitzner and Howes, 2023): the area under the ROC curve (AUC). Specifically, we compute the AUC of the bank-assigned PD in predicting realized default separately for loans to AI-adopting and non-adopting borrowers, and test whether the two areas differ. The AUC measures the bank's ability to rank-order borrowers by default risk: an AUC of 0.5 corresponds to an uninformative signal, and higher values indicate better discrimination. Because the AUC is invariant to any monotone transformation of the PD, a lower AUC among AI borrowers cannot be an artifact of banks scaling or calibrating PDs differently for this segment—it therefore reflects a genuine loss of discriminatory power.

Table 8 and Figure 2 present the findings. Table 8 reports the AUC separately for AI and

non-AI loans, where an AI loan is defined as a loan extended to a borrower with a positive level of AI adoption, and a non-AI loan as one made to a borrower without observable adoption. The AUC is higher for non-AI borrowers than for AI borrowers. This echoes the previous finding: banks' risk assessments are more accurate for non-AI borrowers, and loans associated with AI adoption are assessed less accurately.

Figure 2 illustrates this finding graphically by plotting the ROC curves for different outcome variables (ever-default in the left panel and 90 days past due in the right panel). The ROC curve for loans to non-AI borrowers lies consistently above that for AI borrowers, so the AUC is larger for the former, in line with Table 8. Again, these results indicate that banks' risk assessments are less predictive of actual loan performance for AI borrowers.

Two interpretations are consistent with our finding that banks lend to AI adopters, despite their weaker information. First, AI adoption projects may be intrinsically harder for banks to evaluate: the technology is new, its returns are uncertain, and banks have not yet accumulated experience in assessing such investments. Second, banks may choose to invest less in screening and monitoring these borrowers, in that they produce less information. Or, at least, they do not display superior information production about AI borrowers. A possible reason is that banks compete for this segment on strategic grounds, so that winning the borrower, rather than pricing its risk precisely, is the primary concern. Both interpretations are consistent with banks' lending to AI adopters not being driven by superior information about this segment. We return to these explanations when discussing potential mechanisms behind our findings in Section 5.

4 Other Results and Robustness

4.1 Bank Specialization in AI Lending

Our previous results documented that firms' AI adoption is associated with increased reliance on bank debt, specifically credit lines. This suggests that banks do provide liquidity during

the AI transformation process. We now examine whether such credit provision is driven by banks, specializing in this domain. Indeed, bank specialization in industries or specific lending practices is relevant for loan outcomes (Paravisini et al., 2023; Blickle et al., 2026). To this end, we construct a measure of bank specialization, which is based on our job postings data (it thus differs from other measures used in the literature constructed on different data).

We measure specialization using the average commitment-weighted AI-adoption intensity of a bank's borrowers. For each bank b and $t-1$, we calculate the loan commitment-weighted AI hiring exposure of its borrowers. A bank b is considered a specialist (in lending to AI borrowers) at time t if this measure is above the median across other banks. Note that specialization at time t is calculated on lagged loan commitments and variables. The measure is also time-varying and binary. The formal construction of the measure is described in Appendix A.5.

Specialist Bank Credit Predicts AI Adoption. We now use this measure to test whether credit expansion to AI adopters is concentrated among specialist banks. To this end, we examine whether credit from specialist and non-specialist banks predicts AI adoption on both the extensive margin (whether a firm adopts in a given year) and the intensive margin (the year-over-year change in adoption intensity). We follow a specification similar to that in Section 2.3, particularly Table 3, regressing both measures of AI adoption on lagged bank credit from specialist and non-specialist banks.

Table 9 presents the results. Across all specifications, greater credit from specialist banks is positively and significantly associated with subsequent AI adoption. By contrast, credit from non-specialist banks is significant only in column (2) and is negative in column (3); in columns (1), (2), and (4), its coefficient is positive but smaller than that on specialist bank credit. Thus, the relationship between bank-credit expansion and subsequent AI adoption is concentrated among banks specializing in lending to AI adopters. Overall, these findings suggest that specialist banks are the primary providers of credit for firms' AI adoption.

AI Adoption and Reliance on Specialist Bank Credit. We next examine whether

the relationship between AI adoption and loan terms is stronger for specialist banks. We augment both the OLS regression in (3) and the corresponding IV specification with an interaction between AI adoption and bank specialization. Thus, we follow the regression design in Table 4 but introduce this interaction term.

Table 10 presents the results. As in Table 4, the coefficient on AI adoption is negative and significant for spreads and positive, though not significant, for loan commitments. These estimates are consistent with AI adoption being associated with larger bank loans at lower spreads. Crucially, the interaction coefficient is negative for spreads and positive for loan commitments and is significant across specifications.

Thus, while AI adoption is generally associated with more bank credit at lower spreads, this relationship is particularly strong for specialist banks. In other words, specialist banks drive credit provision to AI-adopting firms.

Specialist Banks and Information. Finally, one might wonder whether the less accurate information in AI lending also extends to specialized banks. Appendix Table A.4 repeats the analysis of Table 7, splitting the sample into loans by specialist and non-specialist banks. The results are qualitatively similar in both subsamples: the PD predicts actual loan performance, while the interaction term is negative (and generally significant, except in some specifications). Specialists' activity in this segment thus does not reflect an intrinsic information advantage—relative to their other lending or to other banks—but likely other motives. In other words, we do not find evidence that specialists have superior information in lending to AI adopters.

4.2 Does Lending to AI Adopters Crowd Out Other Lending?

Our findings thus far suggest that banks increase lending after firms adopt AI, primarily by providing credit lines that support the adoption process. Banks, however, face limits on their lending capacity. Those that lend more actively to AI adopters may therefore reduce credit to firms that do not adopt AI. We next test whether lending to AI adopters crowds

out lending to non-adopters.

To this end, we construct measures of banks' AI lending activity, denoted $\text{NewAIFlow}_{b,t}$, which can take two forms. For each bank b and quarter t, we employ two measures of the bank's new-AI lending flow. The first, the *new-AI share*, is the fraction of the bank's total committed credit that flows into newly started lending relationships with AI-adopting firms in that quarter. The second, the *new-AI intensity*, replaces the adopter indicator with the borrower's AI posting share: it is the committed-exposure-weighted mean AI intensity of the bank's newly started relationships, standardized to unit variance. The share treats a marginal and an intensive adopter alike, while the intensity weighs by adoption intensity. Appendix A.4 provides the formal definitions for these outcome variables.

Armed with these measures of AI lending, we now focus on non-AI firms—that is, firms that do not adopt AI according to our measure. We compare these firms' loan commitments at a given point in time across banks with different levels of AI-lending activity. Specifically, for the subsample of loans to non-AI firms, we estimate loan-level regressions of the form:

$$Y_{f,b,t} \;=\; \beta\,\text{NewAIFlow}_{b,t} \;+\; \alpha_{f\times t} \;+\; \varepsilon_{f,b,t}. \tag{5}$$

We estimate the regression using either the full sample of loans or the subsample of newly originated loans. The outcome variable is the log of total commitments.

Crucially, we include firm×quarter fixed effects, following, for example, Khwaja and Mian (2008), Paravisini et al. (2023), and Greenwald et al. (2025). These fixed effects allow us to compare credit commitments to the same non-AI borrower in the same quarter across banks with different levels of AI-lending activity. They absorb firms' time-varying characteristics and overall credit demand, though not necessarily bank-specific credit demand, as emphasized by Paravisini et al. (2023).

Table 11 presents the results. Over the full sample period, the coefficient on a bank's new-AI share is negative and significant across all specifications. Thus, in the same quarter,

a given non-AI borrower obtains less credit from a bank with a higher new-AI share than from one with a lower share. This pattern is consistent with lower credit supply at banks more actively lending to AI adopters and therefore with a crowding-out effect.

Table 11 also reports estimates separately for the pre- and post-ChatGPT periods (i.e., before and after November 2022). The coefficients on the new-AI share are consistently larger in magnitude after ChatGPT's release, suggesting that the crowding out of lending to non-AI borrowers became more pronounced as banks expanded credit to AI adopters.

Overall, our findings suggest that firms rely more heavily on bank credit as they adopt AI and that this expansion partly crowds out lending to non-adopters.

Specialists and Crowding Out. A natural question is whether this crowding-out effect is more pronounced among banks specializing in lending to AI-adopting firms, as discussed in Section 4.1. To examine this, we reestimate the regressions separately for specialist and non-specialist banks. Appendix Table A.6 shows that crowding out of lending to non-AI borrowers is stronger among AI-specialized banks, consistent with their deeper engagement in this segment. Indeed, the coefficient on AI lending activity is consistently negative and significant for specialized banks, while being smaller or even positive for non-specialists.

Taken together, specialist banks appear to drive both the expansion of credit to AI adopters and the corresponding reduction in lending to non-adopters.

4.3 Generative AI Adoption and Bank Lending

Throughout the paper, we measure AI adoption from AI-related job postings, following Acemoglu et al. (2022). As detailed in Appendix A.3, these "narrow" AI postings are identified using keywords such as machine learning, deep learning, and neural networks. For the robustness test in this section, we instead construct a separate, narrower keyword list that captures only *generative*-AI postings—terms such as language models, prompt engineering, and ChatGPT. Naturally, generative-AI postings were uncommon before the release of ChatGPT.

As a robustness test, we reconstruct our measure using only generative-AI postings, while leaving the remainder of the methodology described in Section 1 unchanged. This measure takes lower values because generative AI accounts for a much smaller share of job postings than our baseline classification for AI jobs.

Using this measure, we re-estimate our main regressions and show that the findings in Tables 3 and 4 are robust. First, consistent with Table 3, Appendix Table A.7 shows that greater bank credit—particularly in the form of credit lines—predicts subsequent generative-AI adoption on both the extensive margin (whether a firm newly adopts) and the intensive margin (adoption intensity).

Second, consistent with Table 4, Appendix Table A.8 shows that generative-AI adoption is associated with larger bank-loan commitments at lower spreads. The estimated coefficients are larger when we use the generative-AI measure, reflecting the substantially smaller scale of this measure. Overall, our key findings continue to hold, when defining AI adoption based on generative AI postings, that is, when explicitly measuring generative AI adoption.

5 Mechanism: Supply- vs. Demand-Side Explanations

Our central question is whether and how banks finance firms' AI adoption. We document two complementary patterns. First, firms' reliance on bank credit—particularly credit lines—predicts subsequent AI adoption. Second, AI adoption is associated with larger bank-credit commitments, primarily through credit lines, as well as lower spreads and lower credit-line utilization. Term loans typically fund large investment outlays, whereas credit lines meet ongoing financing needs and absorb cash-flow shocks.

Our findings suggest that banks facilitate AI adoption by providing liquidity insurance during an ongoing and uncertain transformation process. This credit expansion occurs despite banks' less accurate risk assessments and partly crowds out lending to non-adopters. Taken together, the concentration of credit expansion in credit lines and the weaker accuracy

of banks' risk assessments suggest that banks' role lies primarily in providing liquidity rather than producing superior information.

A natural next question is whether the provision of bank credit to AI adopters is driven by supply- or demand-side channels. Although our analysis does not fully disentangle these forces, the evidence suggests that both contribute to the expansion of credit. We discuss demand- and supply-side mechanisms consistent with our findings.

Demand-Side Mechanism. On the demand side, firms may require working capital and liquidity during AI adoption. Indeed, we show in Section 2.1—and in greater detail in Appendix A.6—that AI adoption corresponds to a broader transformation of the workforce and business model. Such a transformation entails ongoing financing needs whose timing and magnitude are uncertain, rather than a single upfront investment. Credit lines are well suited to meeting these needs. The combination of larger credit-line commitments and lower utilization rates is therefore consistent with firms securing precautionary liquidity that they can draw on as needed during the adoption process.

Supply-Side Mechanisms. On the supply side, we propose two complementary mechanisms. The first is a strategic reallocation of bank credit toward firms undergoing AI transformation. Banks may view AI adopters as a strategically valuable borrower segment and therefore extend larger commitments on more favorable terms, predominantly through credit lines. This expansion occurs despite the difficulty of assessing adopters' credit risk and is concentrated among banks with prior experience financing AI-intensive borrowers and those with existing borrower relationships. Moreover, the associated reduction in lending to non-adopters is consistent with banks reallocating limited balance-sheet capacity toward AI adopters. This interpretation relates to the relationship-lending literature (Boot and Thakor, 2000): banks offer favorable loan terms to AI adopters to build or deepen their presence in this segment.

A second mechanism may amplify this reallocation. AI adoption may reduce the value of lenders' private information by making borrowers' future business models and cash flows

harder to infer from historical performance. If this transformation reduces informational asymmetries across competing banks, rather than merely increasing uncertainty for all lenders, it may weaken the winner's curse faced by less-informed banks. The resulting increase in competition could contribute to lower spreads and greater credit supply even as the accuracy of banks' credit-risk assessments deteriorates. The winner's curse in bank competition is studied theoretically by, among others, Broecker (1990), Marquez (2002), Hauswald and Marquez (2006), and Von Thadden (2004); Beyhaghi, Fracassi and Weitzner (2026) provide evidence for this mechanism in bilateral lending to smaller borrowers. The informational channel also parallels models of lending booms in which changes in the information environment intensify bank competition (Dell'Ariccia and Marquez, 2006). Our evidence is consistent with this channel, although it does not directly isolate it.

The two supply-side mechanisms are complementary: the erosion of informational advantages may lower barriers to competing for AI borrowers, while banks' strategic incentives may explain the intensity and allocation of that competition. Banks do not appear to expand credit because adopters become easier to evaluate; rather, they compete to finance a potentially valuable segment despite greater uncertainty, with specialist banks playing a leading role.

Demand- and supply-side mechanisms also interact in shaping banks' credit expansion to AI adopters. Overall, firms' demand for liquidity during transformation helps explain the form and quantity of the credit expansion, while banks' willingness to supply that liquidity helps explain its pricing and allocation.

6 Conclusion

Combining loan-level data from the Federal Reserve's Y-14 with job-posting data from Lightcast, we study the role of bank credit in firms' adoption of artificial intelligence (AI). We construct a measure of AI adoption from firms' job postings. We show that AI adoption cap-

tures a broader workforce transformation with ongoing liquidity needs rather than a one-time change, which shapes the role of banks in the transformation process.

We find a two-way association between bank credit and AI adoption: greater bank credit predicts subsequent adoption, while adoption is followed by more credit at lower spreads, primarily through credit lines. Adopters also exhibit lower credit-line utilization. At the same time, banks' information—measured by the ability of their risk assessments to predict loan performance—is weaker for loans to AI-adopting firms than for other loans. That is, we find no evidence of superior information production about AI adopters (e.g., through banks' screening and monitoring). Lending to AI adopters also crowds out lending to other firms. We further document that specialist banks provide more credit to AI adopters but do not display superior information. These results are robust to instrumenting for AI adoption using firms' predetermined workforce exposure to large language models interacted with the release of ChatGPT.

Our findings suggest that banks facilitate firms' AI adoption primarily by providing liquidity through credit lines. Both demand- and supply-side explanations are consistent with this credit expansion. On the demand side, firms require precautionary liquidity throughout an ongoing transformation. On the supply side, banks are willing to accommodate and compete for this segment despite less accurate risk assessments, whether because AI transformation is inherently difficult to assess or because banks produce less information about these borrowers.

Thus, contrary to the conventional view that banks are ill-suited to finance innovation, banks play an important role in financing incumbent firms' AI transformation. Our finding suggest that they support AI adoption as liquidity providers rather than through superior information production.

Our study contributes to the growing literature on the interaction between AI and finance by providing early evidence on the role of banks in facilitating firms' AI adoption. Several important questions remain open. First, given that banks lend extensively to AI adopters

despite having a weaker informational advantage in this segment, does such lending create additional risks for the banking sector or increase banks' broader exposure to AI-related technological risk? We find no evidence thus far that loans to AI adopters perform worse, but assessing the systemic implications of this lending requires a longer time horizon. Second, how does the reallocation of credit away from non-adopters ultimately affect their investment and competitiveness relative to adopting firms? Third, although we document a weaker informational advantage, further research could shed light on how banks screen and monitor AI-adoption projects. We view these questions as fruitful avenues for future research.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

Acharya, Viral, Heitor Almeida, Filippo Ippolito, and Ander Perez, "Credit lines as monitored liquidity insurance: Theory and evidence," *Journal of financial economics*, 2014, *112* (3), 287–319.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska, "The Demand for AI Skills in the Labor Market," *Labour Economics*, 2021, *71*, 102002.

Amore, Mario Daniele, Cedric Schneider, and Alminas Žaldokas, "Credit supply and corporate innovation," *Journal of financial economics*, 2013, *109* (3), 835–855.

Babina, Tania, "Understanding Firms' AI Efforts and Their Economic Impact," 2026.

_ , Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Benmelech, Efraim and Michał Zator, "Robots and firm investment," *Journal of Financial Economics*, 2025, *174*, 104183.

Berton, Fabio, Sauro Mocetti, Andrea F Presbitero, and Matteo Richiardi, "Banks, firms, and jobs," *The Review of Financial Studies*, 2018, *31* (6), 2113–2156.

Beyhaghi, Mehdi, Cesare Fracassi, and Gregory Weitzner, "Adverse selection in corporate loan markets," *The Journal of Finance*, 2026, *81* (1), 239–284.

Bircan, Çağatay and Ralph De Haas, "The limits of lending? Banks and technology adoption across Russia," *The Review of Financial Studies*, 2020, *33* (2), 536–609.

Blickle, Kristian, Cecilia Parlatore, and Anthony Saunders, "Specialization in banking," *The Journal of Finance*, 2026, *81* (3), 1531–1572.

_ , Zhiguo He, Jing Huang, and Cecilia Parlatore, "Information-based pricing in specialized lending," *Journal of Financial Economics*, 2025, *172*, 104135.

Boot, Arnoud WA and Anjan V Thakor, "Can relationship banking survive competition?," *The journal of Finance*, 2000, *55* (2), 679–713.

Broecker, Thorsten, "Credit-worthiness tests and interbank competition," *Econometrica: Journal of the Econometric Society*, 1990, pp. 429–452.

Brown, James R, Matthew T Gustafson, and Ivan T Ivanov, "Weathering cash flow shocks," *The Journal of Finance*, 2021, *76* (4), 1731–1772.

Brown, Jennifer and David A Matsa, "Boarding a sinking ship? An investigation of job applications to distressed firms," *The Journal of Finance*, 2016, *71* (2), 507–550.

Buchetti, Bruno, Christian Eufinger, Luca X Lin, and Ixart Miquel-Flores, "Headwind or Tailwind? How Banks Interpret Borrowers' AI Exposure," *How Banks Interpret Borrowers' AI Exposure (July 01, 2026)*, 2026.

Caglio, Cecilia R, R Matthew Darst, and ebnem Kalemli-Özcan, "Risk-taking and monetary policy transmission: Evidence from loans to smes and large firms," Technical Report, National Bureau of Economic Research 2021.

Chodorow-Reich, Gabriel, Olivier Darmouni, Stephan Luck, and Matthew Plosser, "Bank liquidity provision across the firm size distribution," *Journal of Financial Economics*, 2022, *144* (3), 908–932.

Cohen, Gregory J, Jacob Dice, Melanie Friedrichs, Kamran Gupta, William Hayes, Isabel Kitschelt, Seung Jung Lee, W Blake Marsh, Nathan Mislang, Maya Shaton et al., "The US syndicated loan market: Matching data," *Journal of Financial Research*, 2021, *44* (4), 695–723.

Degryse, Hans, Olivier De Jonghe, Leonardo Gambacorta, and Cédric Huylebroek, "Bank specialization and corporate innovation," *Review of Finance*, 2026, *30* (4), 1365–1402.

_ , Tarik Roukny, and Joris Tielens, "Asset overhang and technological change," *The Review of Financial Studies*, 2026, *39* (7), 2115–2178.

Dell'Ariccia, Giovanni and Robert Marquez, "Lending booms and lending standards," *The journal of finance*, 2006, *61* (5), 2511–2546.

Dell'Ariccia, Giovanni, Dalida Kadyrzhanova, Camelia Minoiu, and Lev Ratnovski, "Bank lending in the knowledge economy," *The Review of Financial Studies*, 2021, *34* (10), 5036–5076.

Deming, David and Lisa B. Kahn, "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals," *Journal of Labor Economics*, 2018, *36* (S1), S337–S369.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, "Generative AI and firm values," Technical Report, National Bureau of Economic Research 2023.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 2024, *384* (6702), 1306–1308.

Falato, Antonio, Dalida Kadyrzhanova, Jae Sim, and Roberto Steri, "Rising intangible capital, shrinking debt capacity, and the US corporate savings glut," *The Journal of Finance*, 2022, *77* (5), 2799–2852.

Fonseca, Julia and Bernardus Van Doornik, "Financial development and labor market outcomes: Evidence from Brazil," *Journal of Financial Economics*, 2022, *143* (1), 550–568.

Gandhi, Priyank, Juntai Lu, Jasper Pan, Alberto Plazzi, and Jia Wei, "Equity Prices the Opportunity, Debt Prices the Risk: Generative AI and Corporate Credit Spreads," *Debt Prices the Risk: Generative AI and Corporate Credit Spreads (March 31, 2026)*, 2026.

Greenwald, Daniel L, John Krainer, and Pascal Paul, "The credit line channel," *the Journal of Finance*, 2025, *80* (6), 3137–3183.

Haque, Sharjil, Simon Mayer, and Irina Stefanescu, "Private debt versus bank debt in corporate borrowing," in "Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting" 2024.

Hauswald, Robert and Robert Marquez, "Competition and strategic information acquisition in credit markets," *The Review of Financial Studies*, 2006, *19* (3), 967–1000.

He, Bianca, "Financing intangibles," *Available at SSRN 4692149*, 2026.

Hellmann, Thomas, Laura Lindsey, and Manju Puri, "Building relationships early: Banks in venture capital," *The Review of Financial Studies*, 2008, *21* (2), 513–541.

Hershbein, Brad and Lisa B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, July 2018, *108* (7), 1737–72.

Hochberg, Yael V, Carlos J Serrano, and Rosemarie H Ziedonis, "Patent collateral, investor commitment, and the market for venture lending," *Journal of Financial Economics*, 2018, *130* (1), 74–94.

Holmström, Bengt and Jean Tirole, "Private and public supply of liquidity," *Journal of political Economy*, 1998, *106* (1), 1–40.

Khwaja, Asim Ijaz and Atif Mian, "Tracing the impact of bank liquidity shocks: Evidence from an emerging market," *American Economic Review*, 2008, *98* (4), 1413–1442.

Li, Ye, "Fragile new economy: intangible capital, corporate savings glut, and financial instability," *American Economic Review*, 2025, *115* (4), 1100–1141.

Mann, William, "Creditor rights and innovation: Evidence from patent collateral," *Journal of Financial Economics*, 2018, *130* (1), 25–47.

_ **and Juanita Gonzalez-Uribe**, "Venture debt as bridge financing," *Available at SSRN 5002114*, 2024.

Marquez, Robert, "Competition, adverse selection, and information dispersion in the banking industry," *The Review of Financial Studies*, 2002, *15* (3), 901–926.

Minetti, Raoul, "Informed finance and technological conservatism," *Review of Finance*, 2011, *15* (3), 633–692.

Minoiu, Camelia, Rebecca Zarutskie, and Andrei Zlate, "Motivating banks to lend? Credit spillover effects of the Main Street lending program," 2021.

Paravisini, Daniel, Veronica Rappoport, and Philipp Schnabl, "Specialization in bank lending: Evidence from exporting firms," *The Journal of Finance*, 2023, *78* (4), 2049–2085.

Petersen, Mitchell A and Raghuram G Rajan, "The benefits of lending relationships: Evidence from small business data," *The journal of finance*, 1994, *49* (1), 3–37.

Rajan, Raghuram G, "Insiders and outsiders: The choice between informed and arm's-length debt," *The Journal of finance*, 1992, *47* (4), 1367–1400.

Robb, Alicia M and David T Robinson, "The capital structure decisions of new firms," *The Review of Financial Studies*, 2014, *27* (1), 153–179.

Sforza, Alessandro and Edoardo Maria Acabbi, "Credit and firms' organization," *Journal of International Economics*, 2025, p. 104172.

Thadden, Ernst-Ludwig Von, "Asymmetric information, bank lending and implicit contracts: the winner's curse," *Finance Research Letters*, 2004, *1* (1), 11–23.

Ueda, Masako, "Banks versus venture capital: Project evaluation, screening, and expropriation," *The Journal of Finance*, 2004, *59* (2), 601–621.

Weitzner, Gregory and Cooper Howes, "Bank information production over the business cycle," *Available at SSRN 3934049*, 2023.

Winton, Andrew and Vijay Yerramilli, "Entrepreneurial finance: Banks versus venture capital," *Journal of financial economics*, 2008, *88* (1), 51–79.

Figure 1: Committed bank credit around firms' AI adoption.
This figure plots event-study coefficients of firms' committed bank credit around the first quarter in which the firm's AI posting share turns positive ($AI_{3yr} > 0$), for log total committed credit (left) and log credit-line committed credit (right). Estimates are from a single regression of the outcome on event-time indicators with firm and NAICS2×quarter fixed effects; the dashed vertical line marks the adoption quarter. Shaded areas report 90% (dark) and 95% (light) confidence intervals based on standard errors clustered at the firm level.

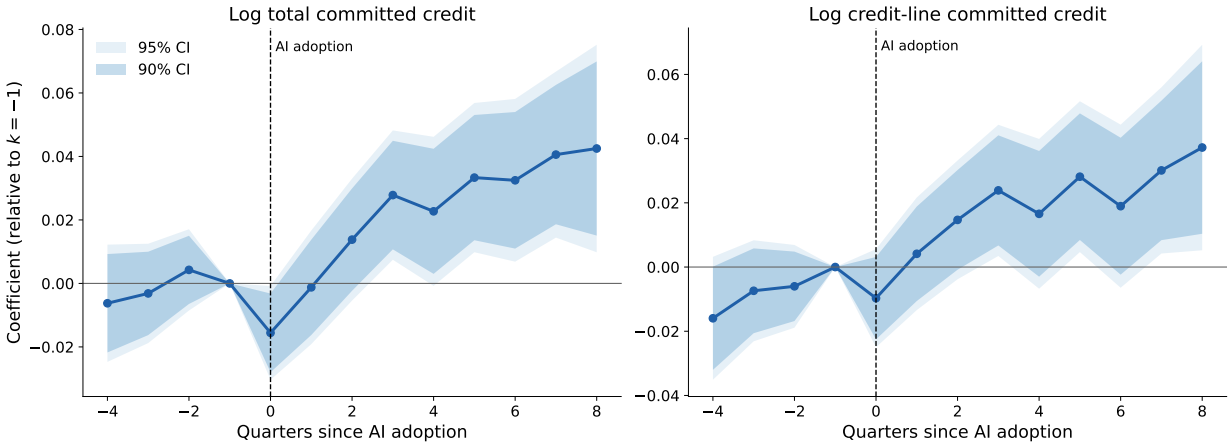


Figure 2: ROC curves of ex-ante probability of default predicting ex-post default. Each panel plots the ROC curve of the bank's assigned PD predicting the realized default outcome (left: ever-default; right: 90+ days past due), separately for AI-adopting and non-AI borrowers, with the 45° random-ranking reference line. Curves are drawn on the loan-quarter sample.

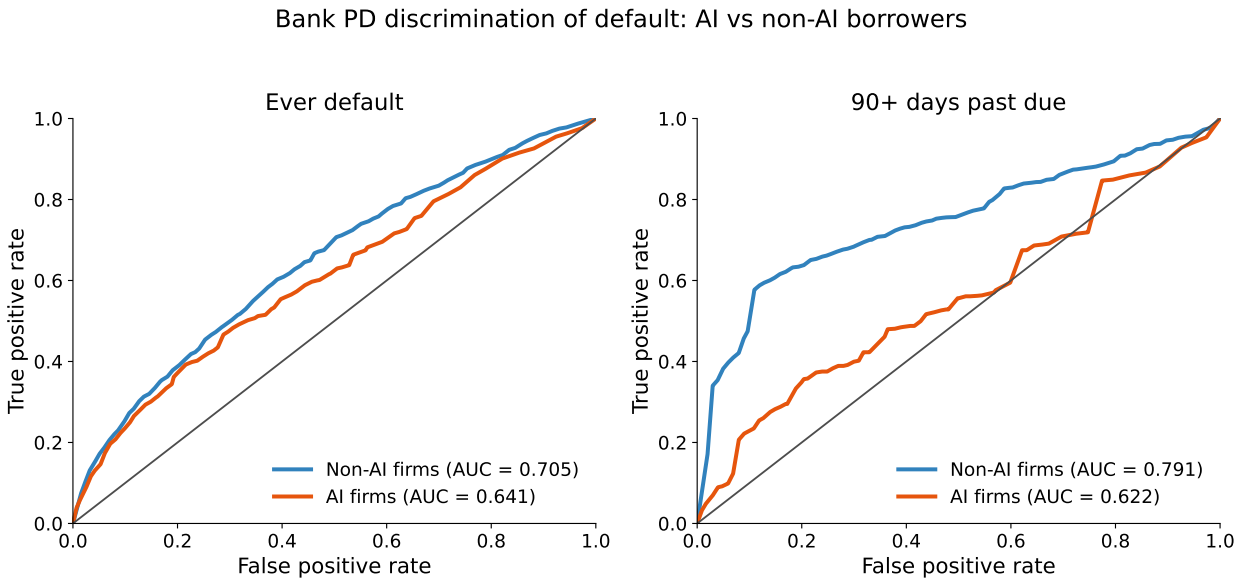


Table 1: Summary Statistics

(a) This table reports summary statistics for the estimation samples. Panel A describes the loan-level regression sample (loan×quarter observations); Panel B describes the firm×year sample used in the firm-level regressions, and the credit-line share of committed credit measured on the firm×quarter panel. Sample period is 2013Q1-2025Q4. All variables are defined in Appendix A.1.

	N	Mean	SD	p25	p50	p75
Panel A: Loan level						
Committed exposure ($M)	1,723,027	24.95	35.94	3.00	10.00	30.57
Utilized exposure ($M)	1,723,027	12.44	28.32	0.82	3.67	14.27
Interest rate spread (pp)	1,723,027	2.01	1.13	1.25	1.80	2.50
1(Credit line)	1,723,027	0.72	0.45	0	1	1
Maturity (years)	1,577,719	6.03	4.28	3.76	5.00	7.00
Panel B: Firm level						
Total assets ($M)	279,549	793.0	2,961.9	10.6	32.1	151.1
Log assets	279,549	17.65	2.23	16.18	17.29	18.83
Tangibility	279,549	0.90	0.18	0.90	0.99	1.00
EBITDA/Assets	278,951	0.17	0.28	0.06	0.12	0.21
Liquidity	279,475	0.14	0.17	0.02	0.07	0.20
Debt/Assets	279,549	0.28	0.26	0.07	0.23	0.42
Debt/EBITDA	231,298	5.18	12.53	0.68	2.00	4.56
Credit-line share of committed credit	865,723	0.79	0.37	0.73	1.00	1.00

Table 2: Characteristics of AI-Adopting Firms: Cross-Sectional Variation

(a) This table reports regressions of AI adoption on the lagged level of a firm's total bank credit (total committed exposure across all the firm's banks, at $t-4$). Data is at the firm-quarter level. Bank credit is divided by total assets in columns (1)–(4) and expressed in logs in columns (5)–(8). In columns (1)–(2) and (5)–(6) the dependent variable is the year-over-year change in firm's AI posting share. In columns (3)–(4) and (7)–(8), the dependent variable is an indicator equal to one in the first quarter the firm's AI posting share turns positive, estimated on firms that have never posted an AI job; firms exit the sample after adopting. Columns (2) and (4) add Debt/Assets as a control; columns (6) and (8) add the log level of the firm's total debt to the same regression. Sample period is 2013Q1-2025Q4. Sector x Qtr FE is defined using 2-digit NAICS code. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	ΔAI	ΔAI	Newly adopt	Newly adopt	ΔAI	ΔAI	Newly adopt	Newly adopt
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\left(\frac{\text{Bank credit}}{\text{Assets}}\right)_{t-4}$	0.0012***	0.0012***	0.0004***	0.0004***				
	(0.0001)	(0.0001)	(0.0000)	(0.0000)				
$\left(\frac{\text{Debt}}{\text{Assets}}\right)_{t-4}$		−0.0002		−0.0001				
		(0.0007)		(0.0002)				
log(Bank credit$_{t-4}$)					0.0031***	0.0032***	0.0011***	0.0012***
					(0.0002)	(0.0002)	(0.0001)	(0.0001)
log(Total Debt$_{t-4}$)						−0.0005***		−0.0003***
						(0.0001)		(0.0000)
log(Assets$_{t-4}$)	0.0051***	0.0051***	0.0017***	0.0017***	0.0028***	0.0032***	0.0010***	0.0012***
	(0.0002)	(0.0002)	(0.0000)	(0.0000)	(0.0002)	(0.0002)	(0.0000)	(0.0001)
Tangibility$_{t-4}$	−0.0104***	−0.0104***	−0.0031***	−0.0031***	−0.0113***	−0.0116***	−0.0033***	−0.0035***
	(0.0015)	(0.0015)	(0.0004)	(0.0004)	(0.0015)	(0.0015)	(0.0004)	(0.0004)
EBITDA/Assets$_{t-4}$	−0.0017**	−0.0017**	0.0003	0.0003	−0.0015**	−0.0016**	0.0003	0.0002
	(0.0008)	(0.0008)	(0.0002)	(0.0002)	(0.0008)	(0.0008)	(0.0002)	(0.0002)
Liquidity$_{t-4}$	0.0185***	0.0184***	0.0031***	0.0030***	0.0204***	0.0194***	0.0037***	0.0032***
	(0.0016)	(0.0016)	(0.0004)	(0.0004)	(0.0016)	(0.0017)	(0.0004)	(0.0004)
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	1,268,508	1,268,508	1,164,616	1,164,616	1,268,508	1,268,508	1,164,616	1,164,616

Table 3: Bank Credit (Lines) Predict AI Adoption

(a) This table reports within-firm regressions of AI adoption on the firm's lagged bank credit, split by type of loan. Data is at the firm-quarter level. In columns (1)–(3) the dependent variable is the year-over-year change in the firm's AI posting share, $\Delta \text{AI}_t = \text{AI}_{3yr,t} - \text{AI}_{3yr,t-4}$; in columns (4)–(6) it equals one in the first quarter the firm's AI posting share turns positive, estimated on firms that have never posted an AI job. Columns (1) and (4) use the log of the firm's total bank credit at $t-4$; columns (2) and (5) split it into credit-line commitment and term-loan commitment. Firms with no credit of a given type at $t-4$ enter with that log set to zero and are absorbed by unreported indicators for positive credit of each class. Controls include lagged firm characteristics (log assets, tangibility, EBITDA/Assets, liquidity). Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	ΔAI			Newly adopt		
	(1)	(2)	(3)	(4)	(5)	(6)
Bank credit $(\log)_{t-4}$	0.0018***			0.0010***		
	(0.0004)			(0.0001)		
Credit-line credit $(\log)_{t-4}$		0.0017***	0.0017***		0.0010***	0.0010***
		(0.0004)	(0.0004)		(0.0001)	(0.0001)
Term-loan credit $(\log)_{t-4}$		−0.0001	−0.0001		0.0004**	0.0011***
		(0.0005)	(0.0005)		(0.0002)	(0.0003)
Total debt $(\log)_{t-4}$			0.0003			0.0003***
			(0.0002)			(0.0001)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y
Observations	1,264,218	1,264,218	1,264,218	1,160,369	1,160,369	1,160,369

Table 4: Bank Loan Terms and AI Adoption

(a) This table reports loan-level regressions of loan quantities and pricing on AI adoption. Data is at the loan-quarter level. AI_{3yr} is the three-year rolling share of narrow-AI job postings. The outcome variables are loan spreads and loan commitment size (both expressed in logs). Controls are debt ratio, tangibility, and log assets; IV columns instrument AI_{3yr} with predetermined occupational AI exposure (2018–2022Q3 shares) interacted with a post-ChatGPT (2022Q4) indicator. Sample period is 2013Q1-2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Loan Spread (log)		Loan Commitment (log)	
	OLS	IV	OLS	IV
AI_{3yr}	-0.025^{***}	-0.148^{***}	0.017^{***}	0.108^{**}
	(0.004)	(0.047)	(0.005)	(0.052)
KP F		112		112
Controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y
Observations	1,723,027		1,723,027	

Table 5: Firm-Level Outcomes: Leverage and Bank Credit

(a) This table reports firm-year level regressions to test how firm-level financial characteristics respond to AI adoption. Data is at the firm-year level. AI_{3yr} is measured in the quarter of the firm's financial statement. The outcome variables are Debt/Asset ratio and the amount of committed bank credit. Log Bank Credit is the log of the firm's total committed bank credit summed over the statement year. Sample period is 2013-2025. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Debt/Assets		Log Bank Credit	
	OLS	IV	OLS	IV
AI_{3yr}	0.001***	0.056***	0.003**	0.179***
	(0.000)	(0.011)	(0.001)	(0.047)
KP F		50		50
Controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
NAICS2×Year FE	Y	Y	Y	Y
Observations	279,549		279,549	

Table 6: AI Adoption and Credit Line Share of Bank Debt

(a) This table reports regressions of credit-line use on AI adoption. Data are at the loan-quarter level in columns (1)–(4) and (7)–(8) and at the firm-quarter level in columns (5)–(6). Columns (1)–(4) are loan-level linear probability models for whether a loan is a credit line, unweighted and weighted by committed exposure. The dependent variable in columns (5)–(6) is the credit-line share of the firm's total committed bank credit on the firm×quarter panel. The dependent variable in columns (7)–(8) is the utilization ratio—the ratio of utilized to committed exposure, restricted to credit lines. AI_{3yr} is the firm's three-year rolling narrow-AI posting share. Controls are debt ratio, tangibility, and log assets. IV columns instrument AI_{3yr} with the firm's pre-ChatGPT occupational AI exposure (posting-weighted over 2018Q1–2022Q3) interacted with a post-2022Q4 indicator. The sample period is 2013Q1–2025Q4. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	Loan level: Pr(Credit Line)				Firm×quarter: Credit Line Share		Credit Line utilization	
	Unweighted		Commit-weighted					
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI_{3yr}	0.0071***	0.0440**	0.0120***	0.0422**	0.0059***	0.0192	−0.0076***	−0.0849***
	(0.0019)	(0.0204)	(0.0028)	(0.0193)	(0.0015)	(0.0266)	(0.0019)	(0.0225)
KP F	—	112	—	101	—	170	—	124
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y	—	—	Y	Y
Observations	1,723,027	1,723,027	1,723,027	1,723,027	865,695	865,695	1,243,434	1,243,434

Table 7: Bank-Estimated Probabilities of Default and Loan Performance

(a) This table reports loan-level regressions testing the sensitivity of bank's own ex-ante estimate of default probability on ex-post loan outcomes. Data is at the loan-quarter level. Controls are debt ratio, tangibility, and log assets. The IV uses the optimal instrument (PD×AI ⇐ PD×$\hat{Z}$). Sample period is from 2013Q1-2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Ever Default		Non-Accrual		Past Due 30+		Past Due 90+	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
PD	0.403***	0.411***	0.220***	0.223***	0.158***	0.165***	0.074***	0.076***
	(0.025)	(0.026)	(0.019)	(0.019)	(0.016)	(0.017)	(0.010)	(0.010)
PD×AI$_{3yr}$	−0.043**	−0.068***	−0.044***	−0.055***	−0.026**	−0.049***	−0.015**	−0.023***
	(0.017)	(0.023)	(0.009)	(0.011)	(0.011)	(0.014)	(0.007)	(0.008)
AI$_{3yr}$	−0.002*	0.004	0.000	0.000	0.001*	0.009	0.000	0.002
	(0.001)	(0.014)	(0.000)	(0.006)	(0.001)	(0.010)	(0.000)	(0.005)
KP F		41		41		41		41
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	1,324,321	1,324,321	1,324,321	1,324,321	1,324,321	1,324,321	1,324,321	1,324,321

Table 8: Bank Probability of Default (PD) Discrimination of Default (AUC): AI vs. Non-AI Borrowers

(a) This table reports the area under the ROC curve of the bank's assigned PD predicting each realized default outcome, computed separately for AI and non-AI borrowers, with a two-sample DeLong z.

Outcome	AUC, Non-AI	AUC, AI	Difference	DeLong z
Panel A: AI (any) vs. non-AI				
Ever default	0.705	0.641	−0.065	−13.70
Charge-off	0.828	0.714	−0.114	−7.53
90+ days past due	0.791	0.622	−0.169	−12.82
Loans	273,352	81,364		
Panel B: High-AI (top quartile) vs. non-AI				
Ever default	0.705	0.643	−0.063	−7.15
Charge-off	0.828	0.682	−0.146	−5.71
90+ days past due	0.791	0.601	−0.190	−8.19
Loans	273,352	20,343		

Table 9: Bank Specialization and AI Adoption

(a) This table reports within-firm regressions of AI adoption on the firm's lagged bank credit, split by lender type. Data is at the firm-quarter level. Bank credit is the firm's total committed exposure across all its banks, at $t-4$ and in logs; a bank is a specialist if the committed-exposure-weighted AI intensity of its loan book is above the within-quarter median (lagged one quarter). In columns (1) and (3) the dependent variable is the year-over-year change in the firm's AI posting share; in columns (2) and (4) it is an indicator equal to one in the first quarter the firm's AI posting share turns positive, estimated on firms that have never posted an AI job; firms exit the sample after adopting. Columns (1)–(2) retain the full sample: each type's log is set to zero where the firm holds no credit of that type, with has-type indicators included. Columns (3)–(4) restrict to firm-quarters with credit from both types. All columns include log total debt and lagged firm controls (log assets, tangibility, EBITDA/Assets, and liquidity). The bottom row reports the one-sided test H_0: Specialist − Non-specialist > 0 (coefficient and significance). Sample period is 2013Q1–2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Full sample		Both lender types > 0	
	ΔAI	Newly adopt	ΔAI	Newly adopt
	(1)	(2)	(3)	(4)
Specialist bank credit $(\log)_{t-4}$	0.0019***	0.0012***	0.0019*	0.0014*
	(0.0004)	(0.0002)	(0.0011)	(0.0008)
Non-specialist bank credit $(\log)_{t-4}$	0.0002	0.0007***	−0.0007	0.0004
	(0.0004)	(0.0002)	(0.0012)	(0.0005)
Total debt (log) control	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y
Observations	1,264,218	1,160,369	147,947	109,276
H_0: Specialist − Non-spec. > 0	0.0016***	0.0005***	0.0026**	0.0010

Table 10: Bank Specialization: Loan Terms

(a) This table adds the interaction of AI adoption with the lender's specialist indicator to the baseline loan-level specification. Data is at the loan-quarter level. $Spec_{b,t}$ is a measure of bank specialization in AI lending. Its main effect is absorbed by the bank×quarter fixed effects. Controls are debt ratio, tangibility, and log assets. Sample period is from 2013Q1-2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Loan Spread (log)		Loan Commitment (log)	
	OLS	IV	OLS	IV
AI_{3yr}	-0.016^{***}	-0.141^{***}	0.003	0.088
	(0.005)	(0.047)	(0.006)	(0.054)
$AI_{3yr} \times Spec_{b,t}$	-0.014^{***}	-0.012^{***}	0.021^{***}	0.032^{***}
	(0.004)	(0.005)	(0.006)	(0.008)
KP F		56		56
Firm FE	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y
Observations		1,723,027		

Table 11: Crowding-Out: Bank New-AI Lending and Credit to Non-AI Borrowers

(a) This table reports regressions of log committed credit to non-AI borrowers on the bank's new-AI lending flow (the share of its book flowing into just-started AI relationships, or its standardized intensity version). Standard errors are clustered at the bank level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Credit to the non-AI firm			
	Log Total Commit		Log New-Orig. Commit	
Bank new-AI exposure:	Share	Intensity (per SD)	Share	Intensity (per SD)
	(1)	(2)	(3)	(4)
Full sample	−1.857***	−0.027**	−0.927**	−0.027***
	(0.563)	(0.010)	(0.420)	(0.009)
Observations	921,608	921,608	49,188	49,188
Pre-ChatGPT	−1.781***	−0.040*	−0.837*	−0.024**
	(0.596)	(0.020)	(0.423)	(0.010)
Observations	748,035	748,035	42,522	42,522
Post-ChatGPT	−2.556***	−0.018***	−4.755***	−0.037**
	(0.630)	(0.004)	(1.602)	(0.016)
Observations	173,573	173,573	6,666	6,666
Firm×Qtr FE	Y	Y	Y	Y

A Appendix

A.1 Variable Definition

Variable definitions are categorized into Y-14 loan-level, Y-14 firm-level, and Lightcast variables.

Y-14 Loan-Level Variables.

- *Committed Loan Amount:* Reported total loan commitment in a given credit facility.
- *Utilized Loan Amount:* Reported drawn amount in a given credit facility.
- *Utilization Rate:* Ratio of utilized to committed loan amount.
- *Credit Line* takes value of 1 for facility types values reported by banks from 1-6.
- *Term loan* takes value of 1 for facility types values reported by banks from 7-13.
- *Maturity:* Years between the facility's origination and maturity dates.
- *Interest Rate Spread:* Reported interest rate spread in a given credit facility, expressed in logs in the regressions.
- *Probability of Default (PD):* The bank's expected annual probability of default of the borrower over the life of the loan.
- *Ever Default* takes value of 1 if loan ever shows disposition flag $\in \{2, 3, 4\}$, PD $= 1$, past due > 0, or chargeoff > 0.
- *Non-Accrual* takes value of 1 if loan is ever placed on non-accrual status as assigned by the bank.
- *Charge-off* takes value of 1 if cumulative chargeoffs are ever positive.
- *Past Due 30+ (90+)* takes value of 1 if loan is ever 30 (90) or more days past due.
- *Junior* takes value of 1 if loan is senior-unsecured or contractually subordinated.
- *Cash-Flow loan* takes value of 1 if no collateral is pledged.
- *New Bank* takes value of 1 if first-ever loan between firm and bank.

Y-14 Firm-Level Variables.

- *Debt/Asset* is the ratio of total short and long term debt to total assets in the financial statement year.
- *Debt/EBITDA* is the ratio of total debt to EBITDA, for firm-years with positive EBITDA.
- *EBITDA/Assets* is the ratio of EBITDA to total assets.
- *Total Debt* is Debt/Assets multiplied by total assets, expressed in logs where indicated.
- *Bank Credit* is the firm's total committed loan amount summed across all its banks in the quarter; expressed in logs, or divided by total assets, where indicated.
- *Credit-Line (Term-Loan) Credit* is the component of the firm's bank credit from facilities of types 1–6 (7–13); in logs, firms with no credit of a class enter with that log set to zero, with indicators for positive credit of each class.
- *Credit-Line Share* is the ratio of the firm's credit-line committed amount to its total committed amount, at the firm-quarter level.
- *Sales* is the net sales reported by the firm in the financial statement year, expressed in logs.
- *Total Assets* is the total book assets reported by the firm in the financial statement year, expressed in logs.
- *Tangibility* is the ratio of tangible to total assets.
- *Liquidity* is the ratio of cash and marketable securities to total assets.

Lightcast Variables.

- AI_{3yr} is the 3-year rolling narrow AI posting share (%), quarters $t-1$ to $t-12$ for a given firm observed in Y14.
- ΔAI is the year-over-year change in the firm's AI posting share, $\Delta \text{AI}_t = \text{AI}_{3yr,t} - \text{AI}_{3yr,t-4}$.
- *Ever/Newly Adopt* AI share is positive at any point; Newly=1 in the quarter the firm's AI share first turns positive for a given firm observed in Y14. Estimated on the risk set of firms that have never posted an AI job; a firm exits the sample after adopting.

- *Dissimilarity index* One half the sum of absolute differences between the firm's current-quarter and three-year-prior occupation posting shares (see Appendix A.6).
- *Cosine distance* One minus the cosine similarity between the firm's current-quarter and three-year-prior occupation posting shares (see Appendix A.6).
- $Spec_{b,t}$ takes value of 1 if the committed-exposure-weighted AI posting share of bank b's borrowers is above the within-quarter median across banks, lagged one quarter; the classification is updated each quarter, so a bank's specialist status varies over time (see Appendix A.5).
- *New AI Flow* Share of bank b's newly started lending relationships (committed \$) going to AI firms in quarter t.
- *New AI Intensity* Same as New AI Flow with the adopter indicator replaced by the borrower's AI posting share, standardized to unit variance.
- E_j is the posting-share-weighted Eloundou β exposure over 2018Q1–2022Q3 (predetermined occupational shares)
- $Z_{j,t}$=$E_j \times \mathbf{1}[t \geq 2022Q4]$: Instrument for AI adoption

A.2 Y-14 Data Cleaning and Sample Construction

Our sample consists of Federal Reserve's FR Y-14Q H.1 schedule borrowers matched to Lightcast. We merge the data using a name-based crosswalk keyed to the bank's internal obligor identifier, dropping loan records without a Lightcast match. We restrict the Y-14 data to traditional loan facilities — revolving credit facilities, lines of credit, and term loans — and exclude non-loan exposures such as capitalized leases, standby letters of credit, and fronting arrangements.

Finally, our regression samples are further restricted to firms with sufficient Lightcast posting history to compute the occupational AI exposure measure. That is, firms with job postings over 2018Q1–2022Q3. Financial ratios are winsorized at the 1st and 99th percentiles.

A.3 AI Skill Taxonomy

Following Acemoglu et al. (2022), our primary measure of AI-related skills in job postings relies on a narrow category, which we complement with a distinct generative AI category. We construct these measures as binary indicators: a job posting is coded as AI (narrow or generative) if its raw job description contains any term from the corresponding category's list, and zero otherwise. The set of skills associated with each one of these measures is given below:

- **Narrow AI**: machine learning, deep learning, neural networks, computer/machine vision, natural language processing, speech/image/pattern/object recognition, supervised and unsupervised learning, chatbots and virtual agents, machine translation, sentiment analysis and text mining, recommender systems, predictive models, and specific methods and libraries (support vector machines, random forests, gradient boosting, XGBoost, Keras, OpenCV, word2vec, latent Dirichlet allocation, kernel methods).
- **Generative AI**: generative artificial intelligence, language model, large language modeling, ChatGPT, prompt engineering, chatbot.

For our baseline AI adoption measure, whose construction is discussed in Section 1, we only consider the so-called Narrow AI job postings, following the analysis in Acemoglu et al. (2022). Indeed, Acemoglu et al. (2022) use this definition for their main analysis, but show robustness to broadening the scope. Similarly, we show in Section 4.3 that our key findings remain robust when measuring AI adoption through Generative AI postings.

A.4 Bank New-AI Lending Measures

For each bank b and quarter t, the *new-AI share* is the flow of credit into just-started relationships with AI-adopting firms as a fraction of the bank's book:

$$\text{NewAIShare}_{b,t} \;=\; \frac{\sum_{j\in\mathcal{A}^{\text{new}}_{b,t}} \text{Committed}_{j,b,t}}{\sum_{j} \text{Committed}_{j,b,t}}, \tag{6}$$

where $\mathcal{A}^{\text{new}}_{b,t}$ is the set of new bank-borrower relationships with AI-adopting firms in quarter t. Here, f is a firm, and the total commitment between b and firm f at observation date t is calculated based on the loan observations in Y-14. The *new-AI intensity* replaces the adopter indicator in the numerator with the borrower's AI posting share:

$$\text{NewAIIntensity}_{b,t} \;=\; \frac{\sum_{j\in\mathcal{N}_{b,t}} \text{Committed}_{j,b,t} \times \text{AI}_{3yr,j,t}}{\sum_{j} \text{Committed}_{j,b,t}}, \tag{7}$$

where $\mathcal{N}_{b,t}$ is the set of all firms whose relationship with bank b begins in t; non-AI borrowers carry $\text{AI}_{3yr} = 0$ and contribute zero to the numerator. The intensity measure is standardized to unit variance in the regressions, so its coefficients are per standard deviation.

A.5 Bank Specialization Measure

We calculate the AI exposure as the committed-exposure-weighted mean AI hiring intensity of the bank's borrowers:

$$x_{b,t} \;=\; \frac{\sum_{f\in\mathcal{B}_{b,t}} \text{Committed}_{f,b,t} \;\times\; AI^{3yr}_{f,t}}{\sum_{f\in\mathcal{B}_{b,t}} \text{Committed}_{f,b,t}}, \tag{8}$$

where $\mathcal{B}_{b,t}$ is the bank's borrower set (facility types 1–13, matched dollars only). Here, f denotes a firm, and the total commitment of bank b to firm f at time t is calculated based on the Y-14 loan observations.

The specialist indicator is a within-quarter cross-bank median split, lagged one quarter: $Spec_{b,t} = \mathbf{1}\{x_{b,t-1} > \text{median}_{b'}\,(x_{b',t-1})\}$.

A.6 Workforce Transformation

This Appendix presents details to Section 2.1. We show that AI adoption transforms a firm's workforce composition, i.e., it is associated with a change in workforce composition. To do that, we leverage firm-level O*NET occupation job posting composition mix in the Lightcast

data. Broadly, we compare a firm's current O*NET occupation job posting composition to its composition three years earlier using similarity and dissimilarity measures.

Specifically, we use the following measures to capture the change in workforce composition: (1) index of dissimilarity; (2) cosine distance; and (3) share of job postings in never-before-posted occupations. All measures are built using the firm × occupation × quarter panel of Lightcast postings, at the 6-digit O*NET-SOC occupation level. Let $n_{j,k,t}$ be firm j's postings in occupation k and quarter t, and let

$$s_{j,k,t} = \frac{n_{j,k,t}}{\sum_{k'} n_{j,k',t}}$$

be the share of occupation k in the firm's quarter-t postings. Our measures compare the firm's occupation mix in quarter t, $\mathbf{s}_{j,t}$, with its mix three years earlier, $\mathbf{s}_{j,t-12}$; the dissimilarity and cosine measures are defined for firm-quarters with postings in both quarters, while the new-occupation measure requires postings in quarter t only. The index of dissimilarity is given by

$$\text{Dissim}_{j,t} \;=\; \tfrac{1}{2} \sum_{k} \left| \, s_{j,k,t} - s_{j,k,t-12} \, \right| \;\in\; [0,1]. \tag{9}$$

While traditionally used to measure spatial segregation across cross-sectional units[6], we adapt the index to capture within-firm intertemporal reallocation. Intuitively, it measures how a firm's current workforce composition compares to its composition three years earlier. A value of zero implies that the two occupational mixes are identical, while a value of one implies complete occupational turnover. The cosine distance is defined as

$$\text{CosDist}_{j,t} = 1 - \frac{\sum_k s_{j,k,t}\, s_{j,k,t-12}}{||\mathbf{s}_{j,t}||\;||\mathbf{s}_{j,t-12}||} \in [0,1] \tag{10}$$

where $||.||$ is the euclidean norm, and $\mathbf{s}_{j,t}$ is the vector of shares for firm j at time t. The measure captures the geometric alignment between a firm's current occupation mix and its mix three years earlier. When the two vectors are perfectly aligned—identical in composition—CosDist takes the value of zero. When they are orthogonal to each other—i.e., the firm is posting for a completely different set of occupations than it did three years earlier—it takes the value of one. Finally, the share of job postings in never-before-posted occupations is

[6]For a comprehensive methodological overview of this index in the context of spatial segregation, see the U.S. Census Bureau's *Measures of Residential Segregation* guidance: https://www.census.gov/topics/housing/housing-patterns/guidance/appendix-b.html

given by

$$\text{NewOcc}_{j,t} = \frac{\sum_k n_{j,k,t} \times \mathbf{1}\{n_{j,k,\tau} = 0, \tau < t\}}{\sum_k n_{j,k,t}}. \tag{11}$$

$\text{NewOcc}_{j,t}$ is the fraction of firm j's period-t job postings in novel occupations—specifically, occupations the firm has never advertised prior to quarter t. Therefore, if this measure is zero, all postings in period t belong to occupations that firm j had previously advertised. To show that AI adoption transforms a firm's workforce composition, we run the following regression:

$$Y_{j,t} \;=\; \beta \cdot \text{AI}_{j,t} \;+\; \gamma \ln(\text{Assets}_{j,t}) \;+\; \alpha_j \;+\; \alpha_{\text{NAICS2}\times t} \;+\; \varepsilon_{j,t}, \tag{12}$$

where $Y_{j,t} \in \{\text{Dissim}_{j,t}, \text{CosDist}_{j,t}, \text{NewOcc}_{j,t}\}$, α_j is a firm fixed effect, and $\alpha_{\text{NAICS2}\times t}$ is 2-digit NAICS and time fixed effect. The IV instruments $\text{AI}_{j,t}$ with $Z_{j,t} = E_j \times \mathbf{1}[t \geq 2022\text{Q4}]$, the same predetermined instrument as the loan-level tables. Standard errors are clustered by firm.

Table A.2 presents the results. The coefficient on the measure of AI adoption is consistently positive and significant. Crucially, this result holds both under the OLS specification and under the IV estimation, where we employ the instrument from Section 3.1.2 (explained there in greater detail). Our findings suggest that AI adoption is associated with a change in workforce composition or, put differently, with a broader transformation of the workforce.

Table A.1: AI Exposure and Adoption by Industry Sector

(a) This table reports AI exposure and adoption by 2-digit NAICS sectors, at the firm level, for the firms of the loan-level regression sample (53,648 firms). E is the firm's posting-weighted occupational AI exposure computed over 2018Q1–2022Q3; % High-Exp is the share of the sector's firms with E above the median across all firms; Mean AI is the average three-year narrow-AI firm-level posting share (in percent) from 2013Q1-2025Q4; % AI is the share of firms with a positive share from 2013Q1-2025Q4. Sectors are sorted according to Mean E column.

NAICS	Sector	Firms	Mean E	% High-Exp	Mean AI (%)	% AI
51	Information	2,628	0.580	91.6	0.88	35.3
54	Prof/Technical	6,197	0.556	88.7	0.51	27.5
52	Finance/Insurance	610	0.553	93.9	0.29	22.6
55	Management	195	0.472	67.7	0.28	19.5
22	Utilities	520	0.449	64.6	0.48	32.3
45	Retail Trade	778	0.441	55.7	0.22	18.1
92	Public Admin.	84	0.429	50.0	0.10	13.1
61	Education	716	0.425	48.2	0.17	17.6
53	Real Estate	1,881	0.420	54.9	0.05	6.7
56	Admin/Support	1,967	0.415	48.9	0.19	19.8
42	Wholesale Trade	7,785	0.412	48.7	0.05	5.0
33	Manufacturing	6,458	0.400	45.3	0.20	13.2
81	Other Services	1,010	0.395	45.4	0.05	6.6
21	Mining/Oil & Gas	612	0.392	45.9	0.11	13.7
62	Health Care	3,079	0.392	37.5	0.05	9.4
48	Transportation	1,736	0.383	34.7	0.05	7.7
23	Construction	3,850	0.373	44.6	0.02	3.4
32	Manufacturing	3,163	0.369	35.7	0.09	9.9
49	Transportation	181	0.365	28.2	0.02	6.6
31	Manufacturing	1,975	0.362	31.4	0.05	8.8
44	Retail Trade	5,410	0.353	31.0	0.02	4.1
71	Arts/Entertainment	859	0.338	32.9	0.07	7.9
11	Agriculture	546	0.309	23.4	0.02	3.1
72	Accommodation/Food	1,344	0.295	16.8	0.05	6.1
	Total	**53,648**	**0.418**	**50.0**	**0.18**	**12.3**

Table A.2: Workforce Transformation

(a) This table reports workforce-transformation regressions of the index of dissimilarity, cosine-distance, and new-occupation measures on AI adoption. The control is log assets; the IV instruments AI_{3yr} with predetermined occupational AI exposure × Post(2022Q4). Standard errors are clustered at the firm level.

	Dissimilarity		Cosine Dist.		New Occ.	
	OLS	IV	OLS	IV	OLS	IV
AI_{3yr}	0.005***	0.161***	0.008***	0.213***	0.005***	0.025*
	(0.001)	(0.015)	(0.001)	(0.019)	(0.001)	(0.013)
KP F		201		201		231
Firm FE	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y
Observations	282,104		282,104		452,332	

Table A.3: Non-Price Loan Terms: Loan Seniority and Collateral

(a) This table reports loan-level regressions of non-price loan terms on AI adoption, with firm, NAICS2×quarter, and bank×quarter fixed effects, on the loan-level regression sample. Junior takes value of 1 if the loan is senior-unsecured or contractually subordinated; Unsecured takes value of 1 if no collateral is pledged. AI_{3yr} is the three-year rolling share of narrow-AI job postings. Controls are debt ratio, tangibility, and log assets; IV columns instrument AI_{3yr} with the firm's pre-ChatGPT occupational AI exposure (posting-weighted over 2018Q1–2022Q3) interacted with a post-2022Q4 indicator. Sample period is 2013Q1-2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Junior		Cash-Flow (Unsecured)	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
AI_{3yr}	0.004*	0.027	0.004*	0.026
	(0.002)	(0.023)	(0.002)	(0.022)
KP F		111		112
Firm FE	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y
Observations	1,718,470	1,718,470	1,723,027	1,723,027

Table A.4: Probability of Default and Lender Specialization

(a) This table estimates the PD-screening regression separately on loans held by specialist and non-specialist banks. $Spec_{b,t}$ as in equation (8); banks in their first panel quarter, for which the lagged classification does not yet exist, are coded as non-specialists. The estimation sample is the PD-screening table's, split by lender type. The IV uses the optimal instrument (PD×AI ⇐ PD×$\hat{Z}$), refit within each subsample. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Ever Default		Non-Accrual		Past Due 30+		Past Due 90+	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
Panel A: Specialist banks' loans								
PD	0.330***	0.334***	0.197***	0.203***	0.146***	0.155***	0.072***	0.077***
	(0.031)	(0.032)	(0.021)	(0.021)	(0.021)	(0.021)	(0.013)	(0.013)
PD×AI$_{3yr}$	−0.023	−0.040	−0.037***	−0.051***	−0.029**	−0.050***	−0.016*	−0.025**
	(0.021)	(0.029)	(0.010)	(0.013)	(0.014)	(0.018)	(0.009)	(0.011)
Observations				801,349				
Panel B: Non-specialist banks' loans								
PD	0.415***	0.424***	0.230***	0.229***	0.123***	0.130***	0.058***	0.060***
	(0.035)	(0.035)	(0.028)	(0.028)	(0.020)	(0.020)	(0.014)	(0.014)
PD×AI$_{3yr}$	−0.072***	−0.118***	−0.054***	−0.048***	−0.007	−0.043***	−0.008	−0.019**
	(0.020)	(0.024)	(0.012)	(0.015)	(0.012)	(0.013)	(0.008)	(0.009)
Observations	519,228	519,228	519,228	519,228	519,228	519,228	519,228	519,228
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y	Y	Y	Y	Y

Table A.5: Bank PD Discrimination of Default (AUC) by Lender Type

(a) This table reports the area under the ROC curve of the bank's assigned PD predicting each realized default outcome, computed separately for AI and non-AI borrowers, on the AUC table's loan-collapsed sample split by lender type. $Spec_{b,t}$ as in equation (8), evaluated at the loan's first observed quarter; banks in their first panel quarter are coded as non-specialists. The DeLong z tests the AI–non-AI difference within each panel; the bottom rows test the difference in those gaps across panels.

Outcome	AUC, Non-AI	AUC, AI	Difference	DeLong z
Panel A: Specialist banks' loans				
Ever default	0.679	0.645	−0.034	−6.14
90+ days past due	0.763	0.617	−0.147	−9.48
Loans	150,474	55,456		
Panel B: Non-specialist banks' loans				
Ever default	0.746	0.647	−0.099	−10.56
90+ days past due	0.828	0.658	−0.170	−6.64
Loans	122,878	25,908		
Difference in AI–non-AI gaps (Panel B − Panel A)				
Ever default			−0.065	−5.98
90+ days past due			−0.023	−0.78

Table A.6: Crowding-Out by Lender Type: Specialist vs. Non-Specialist Banks

(a) This table splits the crowding-out regression of Table 11 by lender type. The regressor is the bank's new-AI lending flow (the share of its book flowing into just-started AI relationships); lender type is the above-median within-quarter book AI-intensity indicator, lagged one quarter. Panel A estimates the regression separately by lender type; Panel B reports the pooled regression with the specialist interaction (specialist main effect included). Standard errors are clustered at the bank level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Log Total Commit		Log New-Orig. Commit	
	Specialists	Non-spec.	Specialists	Non-spec.
	(1)	(2)	(3)	(4)
Panel A: split by lender type				
Full sample	-5.358^{**}	1.473^{**}	-3.565^{***}	1.630^{*}
	(2.103)	(0.627)	(1.122)	(0.956)
Observations	471,039	208,128	22,550	13,979
Post-ChatGPT	-28.591^{***}	-2.010	-38.802^{***}	-3.540
	(7.491)	(1.734)	(7.357)	(4.666)
Observations	85,296	40,287	3,092	1,988
Panel B: pooled, new-AI flow × Specialist				
Full sample	-5.777^{***}		-5.132^{***}	
	(2.028)		(1.669)	
Observations	894,439		47,231	
Post-ChatGPT	-21.497^{***}		-31.500^{***}	
	(7.544)		(8.790)	
Observations	172,639		6,639	
Firm×Qtr FE	Y	Y	Y	Y

Table A.7: Bank Credit Lines and Generative-AI Adoption

(a) Generative-AI robustness check. Within-firm regressions of generative-AI adoption on lagged bank credit, split by facility class. Columns (1)–(3): DV is the year-over-year change in the firm's *generative*-AI posting share (ΔAI). Columns (4)–(6): indicator equal to one in the first quarter it turns positive (at-risk sample). Columns (1),(4) use log total bank credit; (2),(5) split into credit-line (types 1–6) and term-loan (7–13); (3),(6) add log total debt. Firm and NAICS2×Qtr FE; controls as in the baseline. SE clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	ΔAI			Newly adopt		
	(1)	(2)	(3)	(4)	(5)	(6)
Bank credit $(\log)_{t-4}$	0.0164***			0.6500***		
	(0.0035)			(0.1171)		
Credit-line credit $(\log)_{t-4}$		0.0202***	0.0194***		0.8048***	0.7858***
		(0.0035)	(0.0035)		(0.1169)	(0.1171)
Term-loan credit $(\log)_{t-4}$		-0.0116**	-0.0119**		-0.2742*	-0.2834*
		(0.0047)	(0.0047)		(0.1635)	(0.1636)
Total debt $(\log)_{t-4}$			0.0094***			0.2204***
			(0.0015)			(0.0564)
Controls	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y	Y	Y
Observations	1,264,151	1,264,151	1,264,151	1,242,315	1,242,315	1,242,315

Table A.8: Bank Loan Terms and Generative-AI Adoption

(a) Generative-AI robustness check for the main loan-level table. AI^{gen}_{3yr} is the three-year rolling share of *generative*-AI job postings. Controls are debt ratio, tangibility, and log assets; IV columns instrument AI^{gen}_{3yr} with predetermined occupational AI exposure (2018–2022Q3 shares) × post-ChatGPT (2022Q4). Sample period is 2013Q1–2025Q4. Standard errors are clustered at the firm level. *, **, *** denote significance at the 10%, 5%, and 1% levels.

	Loan Spread (log)		Loan Commitment (log)	
	OLS	IV	OLS	IV
AI^{gen}_{3yr}	-0.012***	-0.665***	0.009**	0.483**
	(0.005)	(0.216)	(0.004)	(0.238)
KP F		75		75
Firm FE	Y	Y	Y	Y
NAICS2×Qtr FE	Y	Y	Y	Y
Bank×Qtr FE	Y	Y	Y	Y
Observations	1,723,027		1,723,027	