AI and Operational Losses: Evidence from U.S. Bank Holding Companies

Ping McLemore Atanas Mihov*

September 3, 2025

Abstract

This study demonstrates that banking organizations with higher AI investments are exposed to more operational risk. Using comprehensive supervisory data on operational losses from large U.S. bank holding companies (BHCs) combined with detailed company-level data on AI-skilled human capital, we show that BHCs with more AI investments suffer higher operational losses per dollar of total assets. The impact of AI investments on operational losses significantly varies by loss type and is driven by external fraud, client-related issues, and system failures. These losses stem not only from small, frequent incidents but also from severe, tail-risk events. The risk-enhancing effect of AI is more pronounced for BHCs with weaker risk management practices. Our findings have important implications for banking performance, risk, and supervision.

Keywords: Banking organizations; Risk; Operational risk; Artificial intelligence
JEL Classification: D22, G21, J23, J24, O33

---

*Ping McLemore is a Senior Financial Economist at the Federal Reserve Bank of Richmond and can be reached at ping.mclemore@rich.frb.org. Atanas Mihov is an Associate Professor of Finance at the University of Kansas and can be reached at amihov@ku.edu. The authors thank Tania Babina (referee), Robert Marquez (editor), Allen Berger, Tony Cookson, Filippo Curti, Isil Erel, Joan Farre-Mensa, Scott Frame, Michael Gelman, Will Gornall, Alex He, Vivian Hwa, Steve Karolyi, Sophia Kazinnik, Michelle Lowry, Lira Mota, Anthony Murphy, Matt Plosser, Gustavo Schwenkler, Charles Wallace, Yabo Zhao and seminar and conference participants at the 2025 RCFS Winter Conference, the SFS Cavalcade North America 2025, the 2025 Interagency Risk Quantification Forum, the Bank of Canada, the Federal Reserve Bank of Richmond, and the University of Kansas for helpful comments and suggestions. The authors thank Bryson Alexander for research assistance. The authors also thank Tania Babina, Anastassia Fedyk, Alex He, and James Hodson for providing data. All remaining errors are the authors' alone. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond, or the Federal Reserve System.

1 Introduction

The rapid advancement and widespread adoption of artificial intelligence (AI) technologies has fundamentally transformed how firms operate across industries. In the U.S. banking sector, this transformation has been particularly profound, with financial institutions increasingly deploying AI solutions across their operations — from customer service and fraud detection to trading and risk management (Adhaen et al., 2024). While a growing body of research documents various benefits of AI adoption, including enhanced sales (Czarnitzki et al., 2023; Babina et al., 2024), increased innovation (Cockburn et al., 2018; Babina et al., 2024), and improved product quality (Fedyk et al., 2022), recent research strikingly finds no support for AI increasing operational efficiency (Babina et al., 2024). This failure to find expected efficiency gains suggested by prior literature (e.g., Agrawal et al., 2019; Acemoglu and Restrepo, 2018) points to potential hidden costs of AI adoption. To shed light on this puzzle, we examine the U.S. banking industry and study the relation between AI investments and operational losses at large banking organizations.

Operational losses can be traced to inadequate or failed internal processes, people, and systems or from external events (Basel Committee on Banking Supervision, 2006). Examples include losses from fraud, employment practices and workplace safety, unfulfilled obligations to clients, faulty product design, system failures, process management and transactions failures, and relations with counter-parties and vendors. Operational risk has emerged as a critical concern for financial institutions over recent years as large operational losses wreaked havoc on the banking industry (e.g., Afonso et al., 2019; Berger et al., 2022b). Indeed, the magnitude of these events can be staggering — for example, JPMorgan Chase's 2012 "London Whale" trading scandal resulted in losses exceeding \$6.2 billion.[1] Value-at-risk models further suggest that the largest U.S. banking organizations face potential multibillion-

[1]See *Bloomberg*: "The London Whale" (P. Hurtado, Oct. 16, 2013).

dollar operational losses in any given quarter (Curti and Mihov, 2021; Curti et al., 2022).

The theoretical relation between AI investments and operational losses at banking organizations is ex ante uncertain. On one hand, as a prediction technology (Agrawal et al., 2019; Hansen et al., 2025), AI could reduce operational losses by enhancing efficiency and accuracy in banking operations. AI technologies can automate both routine and complex tasks, thereby reducing human errors and increasing processing speed. For example, machine learning algorithms can detect fraudulent transactions more effectively than traditional methods, minimizing losses due to fraud. Additionally, AI-driven predictive analytics can identify potential operational risks before they materialize, allowing banks to implement preventive measures. The enhanced decision-making capabilities provided by AI can thus contribute to lower operational losses.

On the other hand, the integration of AI systems could introduce new sources of operational risk. Complex AI models might suffer from specification errors or biases in training data, leading to systematic mistakes in decision-making. The increasing reliance on AI systems also creates potential points of failure through cybersecurity vulnerabilities, system downtimes, or model degradation over time as market conditions change. The reduced human oversight in AI-driven processes might delay the detection of novel types of fraud or operational issues that fall outside the AI system's training parameters. Additionally, the transition period during AI implementation could temporarily increase operational losses as organizations adapt their processes and staff learn to work with new systems. In this study, we bridge the literatures on AI, operational risk, and risk management by examining whether investment in AI by financial institutions amplifies or attenuates these organizations' operational risk outcomes.

A key strength of our research is the use of detailed supervisory data on operational losses, which large U.S. bank holding companies (BHCs) report to the Federal Reserve System for regulatory purposes. De Fontnouvelle et al. (2006) and Abdymomunov et al. (2020) note that public data sources often exclude major operational loss events. In contrast to the publicly available data typically used in the operational risk literature, we utilize confidential supervisory data that are significantly richer and more comprehensive. We pair these data with Babina et al. (2024)'s novel measure of firm-level AI investments based on firms' AI-skilled human capital, which uses Cognism's extensive resume data to identify AI-skilled employees.[2] Although combining these data restricts our sample to 36 large BHCs, these institutions account for close to 82% of U.S. banking industry assets as of 2018:Q4.

Our main result is a positive and statistically significant relation between operational losses (as a share of total assets) and AI investments at banking organizations. A one standard deviation increase in our AI investments measure is associated with a 24% increase in quarterly operational losses. In dollar terms, this translates into a $68,416 increase in quarterly operational losses per $1 billion of BHC assets on average, or $12 million per quarter for the median BHC in our sample (with $174.373 billion in total assets). To address concerns regarding unobserved shocks driving both operational losses and AI investments, we follow Babina et al.'s (2024) instrumental variables (IV) strategy of instrumenting for firm-level AI investments using variation in banking organizations' ex-ante exposure to the subsequent supply of AI talent from universities that are historically strong in AI research. The core idea is that the scarcity of AI-trained labor is one of the most important constraints to firms' AI adoption and universities that are historically strong in AI research have been able to train more AI-skilled graduates in recent years, enabling firms that historically hired from those universities to more readily recruit AI talent. The instrument has a strong first

[2] The heavy reliance of AI on human expertise makes the human-capital-based approach particularly well-suited in this setting. Specifically, Babina et al. (2024) construct their AI measure by searching for highly AI-related keywords in each employment record to see if: (1) the job title or description directly includes AI terms, (2) the person obtained patents with AI terms during that year or the following two years, or (3) the person had publications or awards with AI terms during that year or the following year. If any of these conditions are met, that person is classified as an AI-related employee for that firm-year. The firm-level measure is then calculated as the percentage of a firm's U.S.-based employees who are classified as AI-related.

stage, and we show that the instrumented BHC-level increase in AI investments significantly predicts operational losses.[3]

Our additional empirical analyses provide deeper understanding into the relation between AI investments and operational losses at banking organizations. First, we show that AI investments are also positively associated with the frequency of tail operational risk events — those low-probability, high-impact occurrences that can pose significant threats to banking institutions. Tail risk poses difficulties for banking organization capital planning and management, and is particularly relevant for the risk of BHC failure.

Second, our granular analysis of loss categories reveals that three categories of loss events are the principal driver of the relation between operational risk and AI investments: External Fraud (EF) (e.g., due to new attack vectors created by AI systems and digital infrastructure); Clients, Products and Business Practices (CPBP) (e.g., due to challenges in ensuring AI-driven services meet regulatory requirements and customer expectations); and Business Disruption and Systems Failure (BDSF), indicating increased vulnerability to technical failures and system outages. In contrast, we find no significant relation between AI investments and losses from Internal Fraud (IF), Employment Practices and Workplace Safety (EPWS), Damage to Physical Assets (DPA), and Execution, Delivery and Process Management (EDPM).

Third, we document significant interaction effects between AI investments and the quality of BHCs' risk management. Specifically, BHCs with weaker risk management frameworks and internal control systems suffer disproportionately more operational losses associated with AI investments. This finding indicates that AI may introduce new complexities and risks that can exacerbate vulnerabilities in institutions lacking strong risk governance. Effective risk

[3]Beyond the IV strategy, we conduct additional robustness checks to support the validity of our findings. In Section 5, we show that our results hold under alternative data aggregation, additional controls, and fixed effects. We further explore the timing of AI investments and operational losses using a distributed lead-lag model, and find no evidence of pre-trends or reverse causality. Finally, we demonstrate that our results are robust to controlling for past operational losses and to alternative definitions of key variables.

management may thus be an important prerequisite for AI deployment, as robust internal controls enable institutions to curtail risks associated with AI use.

Recent studies document AI's impact across diverse corporate and financial settings. These include the introduction of robo-advisers in wealth management (D'Acunto et al., 2019), the value creation from AI and fintech innovation (Chen et al., 2019, 2024), the transformation of loan underwriting practices (Fuster et al., 2022), the changing nature of financial analysis (Grennan and Michaely, 2020; Abis and Veldkamp, 2023; Cao et al., 2024), the changes in firms' workforce composition and organization (Babina et al., 2023b), and spillover effects on entrepreneurship (Gofman and Jin, 2024). Most relevant to our work, Babina et al. (2024) develop a novel measurement framework using comprehensive employee data to quantify firm-level AI investments, establishing AI's role as a general purpose technology driving economic growth.

An emerging stream of research has also begun to examine the risk implications of AI adoption. Babina et al. (2025) find that AI reduces the volatility of firm fundamentals—sales, earnings, and cash flows—consistent with its role as a predictive technology that improves forecasting accuracy. Similarly, Han et al. (2025) show that AI enhances firms' resilience to natural disasters by enabling faster operational recovery. Ebrahimitorki and Kim (2025) show that AI adoption can improve loan performance by reducing non-performing loans at commercial banks, suggesting that AI may help mitigate credit risk in traditional lending activities. Babina et al. (2023a) document that AI increases firms' systematic risk by expanding growth options that amplify market comovement, particularly on the upside. Durongkadej et al. (2024) find that AI-related incidents at financial institutions trigger negative short-run stock market reactions and elevate bankruptcy risk. At the macroprudential level, Daníelsson et al. (2022) argue that AI may amplify financial instability by introducing novel tail risks, reinforcing cyclical dynamics, and undermining regulatory effectiveness.

Our paper contributes to this literature by documenting a new distinct risk channel through which AI investments affect firms. We find that increased AI adoption at banking organizations leads to higher operational losses, particularly through external fraud, client-related issues, and system failures. Our results contribute to the idea that AI's impact on firm risk is multifaceted: beyond the risk channels identified in prior research, AI adoption can introduce significant idiosyncratic operational risks that are particularly concerning for financial institutions given their potential implications for the risk of organizational failure and wider spillovers to systemic stability (e.g., Berger et al., 2022b).[4] It is important to highlight, however, that the risk-enhancing effects of AI on operational risk should not be generalized to other risk domains, as the consequences of AI adoption can vary depending on the specific functional area in which it is applied (e.g., Ebrahimitorki and Kim, 2025).

Our study also contributes to the literature on operational risk at financial institutions. Cummins et al. (2006) and Gillet et al. (2010) analyze stock market reactions to operational loss announcements at financial institutions, while Chernobai et al. (2024) study insider trading around operational loss announcements. Cope and Carrivick (2013), Abdymomunov et al. (2020), and Frame et al. (2024) analyze financial industry operational losses during the global financial crisis and explicitly link operational risk to the state of the macroeconomic environment. Additionally, research by Chernobai et al. (2011), Wang and Hsu (2013), Abdymomunov and Mihov (2019), and Curti et al. (2023) demonstrates that enhanced corporate governance, improved risk management, and employee training at financial institutions lead to a reduction in operational losses. Frame et al. (2025) document that larger and faster growing banking organizations have higher operational losses per dollar of total assets. Chernobai et al. (2021) show that BHC expansions into non-banking activities result in more operational risk. Berger et al. (2022a) show that banking organizations exposed to severe weather events incur elevated operational losses due to damage to physical assets

[4]More broadly, our research also extends the new literature on financial, economic and technological risks in the new data economy (e.g., Florackis et al., 2022; Bian et al., 2023; Curti et al., 2024; Gomes et al., 2024).

and business disruption. Lastly, Frame et al. (2023) shed light on the adverse operational risk externalities associated with financial innovation. Our study expands this literature by proposing AI as an important source of operational risk at large financial institutions. The staggering size of operational losses, as well as the challenges around measurement and monitoring of operational risk both within organizations and by outside investors, highlight the importance of understanding the organizational drivers of operational risk.

The rest of this paper is organized as follows. Section 2 discusses potential channels through which AI may result in higher operational losses. Section 3 describes our data, the construction of variables, and descriptive statistics. Section 4 presents our main empirical results. Section 5 checks for robustness. Finally, Section 6 concludes.

2 Channels for Elevated Operational Losses

While, AI can improve product quality, expand offerings, and help banking organizations better meet customer expectations, it also creates operational risks. These risks are not necessarily unique to AI but often amplify traditional operational risks, especially when banks lack proper internal controls to support the deployment of AI tools and technological frameworks. We next discuss some specific channels that relate AI to operational risks. Our intent here is to illustrate and contextualize the link between AI and operational risk.

AI deployment at banks can increase cyber risk (e.g., U.S. Department of the Treasury, 2024; European Central Bank, 2025). By increasing banks' digital footprints, AI tools can expose new entry points for cyber threats and external fraud, especially when adversaries exploit sophisticated techniques or gain access to compromised data. AI implementation frequently depends on an extended technology "supply chain," involving external data providers, third-party cloud services, or outsourced development teams. These connections, while valuable for rapid deployment, also expand the network through which breaches, manipulated data, or other security lapses can propagate, potentially triggering widespread operational disruptions and losses (New York State Department of Financial Services, 2024).

AI-driven processes may further increase compliance and regulatory risks. AI algorithms trained on historical banking data may inadvertently learn and perpetuate existing biases, leading to unfair or discriminatory outcomes. In credit and lending decisions, this risk is acute: if the training data reflect past prejudices or structural inequalities, the AI may systematically favor or disfavor certain groups of customers (e.g., Bartlett et al., 2022; Cook and Kazinnik, 2024). Biased models expose banks to potentially significant regulatory fines and legal losses as they may trigger lawsuits and regulatory sanctions under fair lending and equal opportunity laws (e.g., The Equal Credit Opportunity Act).[5]

Technical and systemic failures represent another significant risk with poorly designed or monitored AI. Because of the speed and scale at which automated systems operate, minor errors can build into catastrophic events quickly. A 2012 Knight Capital incident—when a single glitch in an automated trading algorithm caused hundreds of millions of dollars in losses within an hour—provides an illustration.[6] This example also underscores the fact that reliance on intricate automated processes leaves financial institutions less time and fewer resources to correct errors once they begin to unfold. Technical complexity may also arise when AI tools must integrate with legacy platforms that were never intended to handle advanced analytics. Integration failures or incompatibilities can lead to system outages and downtime, halting critical services. Likewise, AI models can degrade as market conditions shift, performing unreliably outside their trained scope. Reduced human oversight in automated processes may delay the detection of operational issues.

[5]For example, the case of *Williams v. Wells Fargo Bank* consolidates six separate class-action lawsuits against Wells Fargo. The lawsuits allege that the bank's algorithm-driven approach to residential mortgage and refinancing decisions violates the Fair Housing Act and the Equal Credit Opportunity Act. Plaintiffs argue that Wells Fargo's automated underwriting system, CORE, operates with minimal human oversight and that its algorithm and machine learning processes are inherently biased against certain racial groups. See *Tech Policy Press*: "AI Lawsuits Worth Watching: A Curated Guide" (B. Barcott, Jul. 1, 2024).

[6]See *Risk.net*: "Knight Capital losses spur focus on algo risk management" (C. Davidson, Sep. 6, 2012).

3 Data Sample and Variable Definitions

3.1 Operational Loss Data

We employ supervisory data on operational losses reported by large banking organizations in accordance with the FR Y-14Q form requirements (current as of December 2022).[7] The Federal Reserve System gathers and uses these data to assess the capital adequacy of large firms, to support supervisory stress test models, and continuous monitoring efforts, as well as to inform the Federal Reserve's operational decision making in implementing the Dodd-Frank Wall Street Reform and Consumer Protection Act. Schedule E specifically captures detailed operational loss event data, requiring institutions to report comprehensive information about each operational loss above their established collection thresholds.

The data are provided by 35 financial institutions with consolidated assets of $100 billion or more. We supplement these data with data for five additional institutions (Comerica, CIT Group, Zions Bancorporation, BBVA USA Bancshares, and SunTrust Banks), which used to participate in Dodd-Frank Act Stress Tests (DFAST), but no longer do so pursuant to the Economic Growth, Regulatory Relief, and Consumer Protection Act of 2018 or due to merger activity. While the original data contains losses from 40 institutions, the availability of data on firm-level AI human capital discussed in Section 3.3 reduces the number of institutions in our sample from 40 to 36.

Per FR Y-14Q reporting instructions, BHCs must report a complete history of operational losses "starting from the point-in-time at which the institution began capturing operational loss event data in a systematic manner." These data are subject to significant data quality checks, including regular data exams conducted by Federal Reserve staff and BHC internal audit functions. The data are at the individual loss event level and provide information such

[7]More information about FR Y-14Q reporting requirements, instructions and forms can be found at: http://www.federalreserve.gov/apps/reportforms/.

as loss amounts, loss dates, and loss classifications.

Operational losses are categorized into seven event types consistent with the Basel II Accord: Internal Fraud (IF), External Fraud (EF), Employment Practices and Workplace Safety (EPWS), Clients, Products and Business Practices (CPBP), Damage to Physical Assets (DPA), Business Disruption and System Failures (BDSF), and Execution, Delivery and Process Management (EDPM). Table 1, Panel A provides definitions of each loss type. IF losses involve acts of a fraudulent nature by internal staff. EF losses stem from events such as cyberattacks, theft, or fraudulent activity conducted by third parties. EPWS includes losses related to employee discrimination, workplace safety violations, and other HR-related legal exposures. CPBP typically involves violations of regulatory standards, breaches of fiduciary duty, product misselling, or customer complaints—often resulting in substantial fines or legal settlements. DPA captures losses due to natural disasters, vandalism, or other damage to physical assets. BDSF refers to losses arising from system outages, hardware failures, or technological disruptions. Finally, EDPM covers transaction failures, processing errors, vendor management breakdowns, and documentation deficiencies.

[Insert Table 1 and Figure 1 about here]

The Basel classification system is relevant for understanding how AI adoption may generate operational risk. Several loss categories—CPBP, EF, BDSF, and EDPM—are arguably more exposed to AI-related risks because they reflect areas where AI can amplify vulnerabilities through biased decision-making, increased digital exposure, transaction automation, and system integration failures. For example, biased outcomes produced by AI-driven credit or underwriting algorithms may trigger CPBP losses by violating fair lending laws or generating discriminatory treatment that results in customer complaints, legal action, or regulatory sanctions. The expanded digital footprint created by AI tools can expose banks to new cyber threats—such as data breaches—while failures in AI-based fraud detection systems can further increase the risk of EF losses. The technical complexity of AI systems and their integration with legacy infrastructure can lead to system outages or instability, contributing to BDSF losses. And, automation errors in trading or processing systems may produce EDPM losses. In contrast, the remaining event types—IF, EPWS, and DPA—are less directly affected by AI adoption, as they primarily involve misconduct, workplace conditions, or physical hazards that do not typically arise from AI-based systems. These mappings guide our empirical analysis by linking specific AI applications to distinct operational risk channels, allowing us to test whether the effects of AI investments are concentrated in the more plausibly affected event types.

Figure 1 shows the distribution of operational losses across the seven Basel-defined event types. CPBP dominates the sample, accounting for 70.2% of total losses (approximately $128.7 billion), followed by EDPM at 18.6% (or $32.2 billion). The remaining five categories collectively account for just 12.2% of losses ($22.4 billion). This distribution highlights the central role of client-facing business practices in driving operational risk at large BHCs. It also underscores the relevance of this category for understanding how AI adoption may affect operational risk profiles—particularly in areas such as regulatory compliance and automated customer interactions, which map closely to CPBP. Notably, many of the most severe and high-dollar-loss events in our sample also fall under CPBP.

The banking organizations in our sample have different thresholds for collecting individual operational losses. To mitigate the impact of firm heterogeneity in collection thresholds on our results, we follow Abdymomunov et al. (2020) and discard operational losses below $20,000, which is the highest threshold across reporting institutions. We next aggregate loss data at the BHC-quarter level, where we use the quarter of the date when an operational loss event occurred (or began) for aggregation purposes. We finally merge loss data with financial data from FR Y-9C and AI investments data from Babina et al. (2024). Our final sample has 852 observations from 36 large BHCs over the period 2010:Q1-2018:Q4. While our combined data contain losses from only 36 BHCs, these institutions account for the majority of U.S. banking industry assets (82% as of 2018:Q4).

It is important to note that, for the covered institutions, our data is substantially more comprehensive than operational loss data offered by private vendors that rely on publicly available information. For example, Hess (2011) uses loss data from SAS OpRisk Global Data, which consists of around 7,300 loss events. Chernobai et al. (2011) analyze loss data from Algo FIRST, which consists of 2,426 events. In contrast, we have 259,408 individual loss events in our sample. As discussed in de Fontnouvelle et al. (2006), public sources of data compiled from press accounts omit substantial operational losses otherwise contained in the supervisory data used in this study. The comprehensive coverage is thus important for studying AI-related operational risks, which may not always result in public disclosure. The supervisory data captures operational risk events systematically, which is further useful for understanding how AI adoption may relate to operational risk across different loss categories and severity levels.

3.2 Operational Loss Measures

Our main measure of operational risk is the total dollar value of operational losses that occur at a BHC in a given quarter. We follow Curti et al. (2023), Frame et al. (2025), and other studies in the literature on bank risk and performance (e.g., James, 1991; Ahmed et al., 1999; Ellul and Yerramilli, 2013), and scale losses by BHC asset size. In doing so, we use lagged total assets, but our results are robust to using contemporaneous measurements of losses and assets. For presentation purposes, we multiply the loss-to-assets ratio by 10,000 and label it LtA. In some of our regression specifications, we also use log-transformed inflation-adjusted dollar losses (2020 constant dollars) that occur at an institution in a given quarter — *Log(Loss)*.

[Insert Table 2 about here]

Table 2 presents descriptive statistics. On average, the BHCs in our sample lose $181 million or the equivalent of 0.04% of their assets per quarter to operational risk. Further, the standard deviations of both dollar losses ($645 million) and asset-scaled operational losses (5.98) are high relative to the means, indicating substantial variation in operational losses.

A well-known property of operational risk is the extremely heavy tails of the empirical loss distributions (Berger et al., 2022b). Indeed, only a few "catastrophic" operational risk events account for a large proportion of the total dollar losses in our sample. Thus, while we focus on quarterly operational losses at BHCs, we also analyze tail operational risk. We use three measures of tail risk frequency, constructed as follows.

We start with the 259,408 individual loss events in our sample and scale dollar loss amounts by BHC total assets. We calculate the 90^{th}, 95^{th}, and 99^{th} quantiles of the resulting empirical distribution and categorize all loss events with severities above the respective quantiles as "tail losses." We then "collapse" the sample of losses at the BHC-quarter level by counting the number of tail events that occur at a given institution during a given quarter for each tail threshold definition, resulting in the variables *N Tail 90*, *N Tail 95*, and *N Tail 99*. Using the 90^{th} quantile definition, Table 2 shows that a BHC experiences an average of 23.312 tail operational losses per quarter, each representing 0.002% of assets on average. In contrast, tail losses under the 99^{th} quantile definition are less frequent—averaging 2.580 per quarter—but more severe, with each loss accounting for 0.012% of assets on average.

To better capture the severity of extreme losses, we also construct three additional measures—*LtA Tail 90*, *LtA Tail 95*, and *LtA Tail 99*—which sum the dollar losses from tail events at each BHC-quarter and scale them by BHC total assets (multiplied by 10,000). While the frequency measures reflect how often extreme events occur, these severity measures capture their magnitude relative to institution size.

3.3 Data on Employment Profiles and Job Postings

We adopt the measure of firm-level AI investments introduced in Babina et al. (2024), which is predicated on the critical role of AI-skilled human capital in AI implementation. In constructing the measure, Babina et al. (2024) leverage data from two primary sources: individual worker resumes from Cognism and job postings from Burning Glass Technologies.

The resume data from Cognism comprises approximately 535 million employee resumes globally. This dataset is widely used for lead generation and client relationship management services, and includes information derived from publicly available online profiles, collaborations with recruiting agencies, and third-party resume aggregators. The Cognism data are introduced and described in detail in Fedyk and Hodson (2022). Cognism's coverage spans roughly 64% of the U.S. workforce as of 2018, and provides detailed information about employment histories, including job titles, companies, job descriptions, patents, and awards. Through machine learning and natural language processing, the resumes are normalized to ensure accurate association of job roles with specific firms and functional divisions. This dataset enables the identification of AI workers within U.S. public firms from 2010 to 2018, capturing about 101 million person-firm-year observations.

Babina et al. (2024) supplement the resume data with over 180 million job postings from Burning Glass Technologies, spanning the years 2007, 2010, and 2010-2018. These postings are aggregated from over 40,000 online sources and company websites. Burning Glass processes these postings to extract labor market analytics, including job titles, locations, employers, and a detailed taxonomy of required skills. This allows for the identification of specific AI-related skills, which are crucial for tracking AI labor demand within firms. The job postings dataset represents about 60-70% of all U.S. job vacancies.

3.4 Measure of BHC-level AI Investments

The firm-level measure of AI investments of Babina et al. (2024) takes advantage of the granular classification of skills in job postings to empirically identify the terms most associated with AI roles. As the authors argue, this data-driven method circumvents issues arising from predefined keyword lists, such as false positives (non-AI jobs being classified as AI) and false negatives (actual AI jobs being omitted).

The construction of the firm-level AI investments measure involves three main steps. First, Babina et al. (2024) examine approximately 15,000 unique job skills and compute their co-occurrence with four core AI-related skills: artificial intelligence, machine learning (ML), natural language processing (NLP), and computer vision (CV). For each skill s, a co-occurrence metric w^{AI} is calculated as follows:

$$w_s^{AI} = \frac{\#\text{jobs requiring skill } s \text{ and (ML, NLP, CV, or AI in required skills/job title)}}{\#\text{jobs requiring skill } s} \quad (1)$$

This co-occurrence metric reflects how often a particular skill appears alongside the key AI competencies. For instance, as discussed in Babina et al. (2023a) “Long Short-Term Memory (LSTM)” demonstrates a high co-occurrence measure of 0.971, indicating that 97.1% of job postings requiring LSTM expertise also specify at least one core AI skill (i.e., “Artificial Intelligence,” “Machine Learning,” “Computer Vision,” or “Natural Language Processing”). Conversely, general professional skills like “Microsoft Office” exhibit minimal correlation with AI skills, with a co-occurrence measure of merely 0.003. This pattern extends to entirely unrelated skills such as “Snow Removal,” which shows zero co-occurrence with core AI competencies, thus serving as a useful baseline for skill independence.

In the second step, AI workers within firms are identified using resume data from the Cognism dataset. Employees whose job titles, job descriptions, or professional achievements (such as patents, publications, or awards) include AI-related terms are classified as AI-skilled.

For instance, an employee with the job title "Senior Machine Learning Developer" or one whose job description includes AI-related tasks is classified as an AI worker.

Finally, these AI-skilled employees are aggregated to the firm level to construct the AI investments measure. Specifically, *Share AI Workers* is defined as the proportion of employees at each firm identified as AI-skilled. Because the AI data is updated annually, we apply the same value of *Share AI Workers* to all quarters in a given year for a given BHC when we merge these data with other BHC data at the quarterly level.

Although operational losses are recorded at quarterly frequency while the measure of AI investments is updated annually, this frequency mismatch is unlikely to bias our results in any way. AI adoption is both gradual and strategic, so the composition of AI-skilled labor changes only slowly within any given year. By contrast, operational losses tend to be episodic: they can be lumpy and reflect distinct events rather than continuous processes. Thus, the annual measure of AI investments should meaningfully capture variation in the intensity of AI adoption across BHCs, and quarterly losses should remain informative for understanding how intensity of AI usage relates to risk outcomes. We lag the AI measure to ensure that it precedes the occurrence of operational losses.

Nonetheless, the persistence of the annually updated AI measure across quarters could potentially inflate the statistical significance of our regression estimates. To address this concern, we cluster standard errors at the BHC level, which should appropriately capture within-BHC correlations in AI intensity as well as operational losses. In Section 5.1, we confirm the robustness of our results to aggregating all data at the annual level.

Table 2 shows the average *Share AI Workers* for the BHCs in our sample – 0.065, which is higher than the typical Compustat firm documented by Babina et al. (2024). This difference can be contributed to two key factors. First, larger firms like BHCs, which accumulate substantial data through their operations, tend to invest more heavily in AI. Second, the finance sector has been an early and intensive adopter of AI technology, consistently maintaining higher concentrations of AI workers.

3.5 Control Variables

Our multiple regression analysis includes several control variables that capture time varying BHC characteristics. We follow Curti et al. (2022) and use the logarithm of BHC total assets (*Log(Assets)*) to control for organizational size. Larger BHCs may have higher exposure to operational risk due to factors such as complexity, moral hazard associated with "too-big-too-fail," or higher volume of transactions among other reasons. We also include year-over-year asset growth (*Asset Growth*) because faster-growing banking organizations tend to experience higher operational losses (Frame et al., 2025).

To account for differences in business models, we include the non-interest to interest income ratio (*II-to-NII*). As documented by Chernobai et al. (2021), banks focused on traditional activities (deposit-taking and lending) exhibit different risk profiles compared to those deriving more income from non-core activities like trading and investment banking. We control for profitability using return on equity (*ROE*), calculated as the ratio of net income to book value of equity. Higher profitability may enable greater allocation of resources to risk management; alternatively, as suggested by Jin and Myers (2006), senior management might be more likely to overlook internal control failures when firms are less financially constrained.

Given that more AI-intensive firms tend to innovate more, and more innovative banking organizations experience higher operational losses (Frame et al., 2023; Babina et al., 2024), we control for financial patent innovation using (*Log(N Patents)*). We also include several risk-related controls: the proportion of non-performing loans (*Non-Performing Loans*) to account for credit risk exposure, which can be related to operational risk (Chernobai et al., 2011); the ratio of total assets to book value of equity (*Leverage*); and the log absolute difference between assets and liabilities that reprice or mature within a year (*Maturity Gap*). In Section 5.4, we extend the analysis by adding additional controls and assessing the robustness of our results to their inclusion.

4 Regression Results

4.1 Operational Losses

Figure 2 provides an initial look at the relation between AI investments and operational risk. Each quarter, we sort BHCs into terciles based on *Share AI Workers* ("Low," "Medium," and "High" AI intensity) and plot the average *LtA* for each tercile. The figure reveals a clear positive association: BHCs with greater AI intensity consistently incur higher operational losses per dollar of assets than their less AI-intensive counterparts.

To more formally examine whether more AI-intensive BHCs have more operational risk, we next estimate the following regression model:

$$Operational\ Loss_{i,t} = \beta_t + \beta_1 Share\ AI\ Workers_{i,t-1} + \beta_2 Controls_{i,t-1} + \epsilon_{i,t} \qquad (2)$$

where i indexes BHCs and t indexes time periods (quarters). *Operational Loss* is one of four operational loss measures: (i) operational losses as a proportion of total assets that occur at BHC i over a given calendar quarter; (ii) log-transformed operational dollar losses that occur at the BHC over a given calendar quarter; (iii) frequency of operational losses that occur at the BHC over a given calendar quarter; or (iv) log-transformed average severity of operational losses that occur at the BHC over a given calendar quarter. *Share AI Workers* measures AI investment intensity at banking organizations, measured in the year prior to the year of quarter t. *Controls* represents our previously discussed vector of control variables. All explanatory variables are lagged.

We include quarter fixed effects (β_t) to absorb period-specific shocks common across all BHCs (e.g., industry-level operational losses). We do not include BHC fixed effects because the average AI investment intensity of BHCs is informative about the average level of operational losses incurred by the BHCs. We deal with potential endogeneity issues introduced by omitted variables including time-invariant factors relevant to operational losses that would be absorbed by BHC fixed effects in Section 4.2. Section 5.4 shows, however, that our results are robust to the inclusion of BHC fixed effects. We cluster standard errors at the BHC level to account for within-BHC correlation of the regression error terms. Table 3 presents the results.

[Insert Figure 2 and Table 3 about here]

The regression specification in Column (1) uses our main measure of operational losses, *LtA*, as the dependent variable. The estimated coefficient on *Share AI Workers* is positive and significant at the 1% level, suggesting that banking organizations with more intensive AI investments experience more operational losses per dollar of assets. The positive relation continues to hold when we use *Log(Loss)* as the dependent variable in Column (2), suggesting the robustness of our results to redefining the operational loss measure. Columns (3) and (4) decompose operational losses into loss frequency and loss severity components. The results indicate that AI investments are positively related to both the frequency and severity of operational losses. Based on Column (1), a one standard deviation increase in *Share AI Workers* is associated with a 24% increase in operational losses. In dollar terms, this translates into a $68,416 increase in quarterly operational losses per $1 billion of BHC assets on average, or $12 million per quarter for the median BHC in our sample (with $174.373 billion in total assets).

4.2 Instrumental Variables

The relation between AI investments and operational losses at banking organizations poses significant identification challenges. Banking organizations experiencing higher operational losses may be more likely to invest in AI technologies to prevent future losses, creating reverse causality concerns. Additionally, omitted variables like risk management culture or regulatory pressure could simultaneously drive both AI adoption and operational loss patterns, which could bias our ordinary least squares estimates.

To address these identification challenges, we employ an instrumental variables (IV) approach, utilizing the instrument developed by Babina et al. (2024). Their instrument exploits variation in firms' ex-ante exposure to the supply of AI talent from universities that were historically strong in AI research. Specifically, for each company i, the instrument is constructed as follows:

$$IV_i = \sum_u s_{i,u}^{2010} \times AI\ Strong_u \tag{3}$$

where $s_{i,u}^{2010}$ is the share of STEM workers in company i in 2010 who graduated from university u, and $AI\ Strong_u$ is an indicator equal to one if university u is identified as an AI-strong university based on pre-2010 publications. Babina et al. (2024) classify a university as AI-strong if it meets either of two criteria between 2005-2009: (1) the number of AI researchers is in the top 5% across all universities, or (2) the number of AI researchers is in the top 10% and the share of AI researchers (relative to all researchers) is in the top 5%. They identify AI researchers based on publications in AI-specific journals and conferences, using data from the Open Academic Graph.

The instrument's relevance stems from the fact that universities with historically strong AI research programs were better positioned to train AI-skilled graduates when commercial interest in AI surged in the 2010s. Since firms tend to maintain persistent hiring relationships with specific universities through alumni networks, firms with stronger pre-existing connections to AI-strong universities had better access to AI talent. This access to AI-skilled labor is particularly important given that the scarcity of AI-trained workers has been identified as one of the key constraints to firms' AI adoption (Correlation One, 2019; Babina et al., 2024).

The identifying assumption (exclusion restriction) requires that a bank's historical connections to AI-strong universities only affect its operational losses through the channel of AI investments. Several institutional features support this assumption. First, commercial interest in AI only became widespread around 2012, while AI research in universities had flourished for decades prior. This temporal separation suggests that banks' pre-2010 hiring networks with AI-strong universities were unlikely to be driven by their anticipated need for AI talent. Second, Babina et al. (2024) focus on connections through STEM graduates broadly, rather than specifically AI-related hires, making the networks more likely to reflect general university relationships rather than targeted AI recruitment.

[Insert Table 4 about here]

Table 4, Columns (1) reports the first-stage estimation results. The estimated coefficient of the instrumental variable is positive and highly significant. The ex-ante exposure to the supply of AI talent from universities is relevant for the proportion of AI-skilled labor at banking organizations. The adjusted R^2 is relatively high, and the F-statistic is above the threshold of 10 prescribed by Stock et al. (2002). This result suggests our estimations do not suffer from a weak-instrument problem. Columns (2) presents second stage results and shows that the estimated coefficient on *Share AI Workers* retains its positive sign and is significant at the 1% level. These findings suggest that the results in the previous section are robust to accounting for endogeneity and reverse causality concerns, confirming the positive relation between BHC AI investments and operational losses.

To address potential violations of the exclusion restriction, we conduct several robustness checks. First, we control for bank organizations' exposure to computer science (CS)-strong universities using an analogous measure of firms' ex-ante exposure to CS-strong universities

($\sum_u s_{iu}^{2010} \times CS\ Strong_u$, where $CS\ strong_u$ measures the university's strength in non-AI computer science research). This addresses concerns that AI-strong universities might excel in broader computer science education, and that computer science capabilities at banking organizations may affect their operational risk through general technology-driven channels distinct from AI—such as increased technological infrastructure complexity, greater reliance on automated processing systems, and expanded digital banking services. Second-stage results remain robust as reported in Table 4, Column (3).

Next, we account for the possibility that AI-strong universities might be concentrated in regions with particular labor market characteristics by controlling for a set of commuting zone attributes. These controls include the share of workers in IT-related occupations, the proportion of college-educated workers, log average wages, the percentage of foreign-born workers, and the share of workers in finance and manufacturing industries. As shown in Column (4), our results again remain robust after incorporating these regional labor market controls. To address concerns about persistent unobservable firm characteristics that might correlate with both operational risk profiles and historical university connections, we introduce controls for pre-period operational losses spanning 2000-2009. This control helps account for the possibility that banks with historically higher operational losses might have systematically different relationships with universities. Results remain robust as reported in Column (5). We further extend our analysis by controlling for U.S. Census Bureau geographic divisions to account for banking organization's broader regional labor market conditions.[8] The results, shown in Column (6), demonstrate that our findings are not driven by regional heterogeneity.

Finally, Column (7) presents our most stringent specification, which simultaneously includes all previously mentioned controls: CS-strong university exposure, commuting zone

[8]The nine distinct divisions are New England, Middle Atlantic, East North Central, West North Central, South Atlantic, East South Central, West South Central, Mountain, and Pacific.

characteristics, pre-period operational losses, and geographic divisions. The persistence of our main findings across this battery of robustness checks strengthens our confidence in a causal interpretation of the relation between AI investments and operational losses. The stability of our results across these various specifications suggests that our instrumental variable approach successfully isolates the effect of AI investments on operational risk from potentially confounding factors.[9]

4.3 Operational Loss Types

Operational risk is an amalgamation of various types of subcomponent risks. The Basel Committee on Banking Supervision categorizes operational risk into seven distinct event types: Internal Fraud (IF), External Fraud (EF), Employment Practices and Workplace Safety (EPWS), Clients, Products and Business Practices (CPBP), Damage to Physical Assets (DPA), Business Disruption and System Failures (BDSF), and Execution, Delivery and Process Management (EDPM). While our earlier analysis established a strong relation between total operational losses and AI investments at the BHC level, examining each category separately may reveal important variations in how AI relates to different types of operational risk. To investigate these potential differences, we re-estimate Equation 2 for each of the seven event types and present our findings in Table 5.

[Insert Table 5 about here]

The positive coefficient of *Share AI Workers* in Column (2), significant at the 1% level, indicates that the relation between AI investments and operational losses is particularly pro-

[9] A less plausible concern regarding our instrument is that BHCs that anticipated the surge in demand for AI may have started building their connections to AI-strong universities before 2010, making BHC-university hiring networks in 2010 endogenous to BHCs' demand for AI-trained students. However, this idea runs counter to the lack of both commercial interest in AI by firms and AI-skilled graduates by universities prior to 2010. Furthermore, Babina et al. (2024) present evidence that firms associated with strong AI universities in 2010 did not increase their share of new graduate hires from those institutions between 2005 and 2010 (see Babina et al.'s Appendix Table A.2).

nounced in EF, supporting the interpretation that AI systems may create new attack vectors for external actors through, for example, increased digital infrastructure complexity and new technological vulnerabilities. The strong relation may reflect heightened exposure to sophisticated cyber attacks targeting AI systems, innovative fraud schemes exploiting automated processes, and opportunities for malicious actors to manipulate AI-driven operations.

We similarly find that CPBP losses increase significantly with AI investments, as evidenced by a positive coefficient of *Share AI Workers* in Column (4), significant at the 5% level. This relation potentially reflects challenges in ensuring AI-driven services meet regulatory requirements and customer expectations. For example, the complexity of AI systems may create difficulties in maintaining transparency, managing algorithmic biases, and adequately disclosing AI use to customers. Finally, BDSF losses also show a positive and significant, at the 10% level, relation with *Share AI Workers* in Column (6). This finding indicates that greater AI adoption may increase vulnerability to technical failures and system outages. For example, as AI systems get deployed in banking operations, disruptions from model failures, integration issues with legacy systems, or AI system degradation may create more severe operational impacts.

In contrast, we find no significant relation between AI investments and losses from IF, EPWS, and DPA. The coefficient of *Share AI Workers* is indistinguishable from zero in Columns (1), (3), and (5). These losses originate from employee misconduct, inadequate workplace practices, or physical asset damage due to natural disasters, and are less likely to be related to AI technologies. Contrary to our discussion in Section 3.1, we find no significant relation between AI investments and EDPM losses. While the coefficient of *Share AI Workers* is positive, it is statistically insignificant at conventional levels (p-value = 0.176). Overall, these findings have important implications for risk management and regulatory oversight. The heterogeneous effects across loss categories suggest that banks and regulators may focus their AI-related risk management efforts on specific operational vulnerabilities rather than applying uniform approaches across all risk types.

4.4 Tail Operational Losses

Our earlier analysis focused on the relation between AI investments and the conditional average of asset-scaled operational losses. However, the distribution of operational losses is highly right-skewed, with a small number of extreme events representing potentially greater threats to institutional stability than routine losses. Distinguishing between elevated but stable operational losses and infrequent tail events is important. While persistently higher operational losses can erode profitability, they are more predictable and therefore easier to anticipate and provision for. In contrast, tail operational losses—more rare but severe events—are harder to forecast and reserve against, posing challenges for capital management and contributing more directly to failure risk.

As described in Section 3.2, we construct three measures of quarterly tail loss frequency based on different severity thresholds: *N Tail 90*, *N Tail 95*, and *N Tail 99*. We then test whether AI-intensive banking organizations experience a higher incidence of such events using a multivariate framework analogous to Equation 2. Because the dependent variables are event counts, we estimate Negative Binomial (NB) regressions.

[Insert Table 6 about here]

Table 6, Columns (1)–(3) show that BHCs with greater AI intensity experience significantly more tail operational loss events. The coefficients on *Share AI Workers* are positive and statistically significant at the 1% level for all three tail definitions. The economic magnitudes are sizable: a one–standard deviation increase in *Share AI Workers* from its mean is associated with a 22–30% increase in the number of quarterly tail events, relative to the sample mean. Columns (4)–(6) confirm these results using alternative measures—*LtA Tail*—that emphasize the size rather than frequency of tail losses. In every case, *Share AI Workers* remains positive and significant. Overall, these findings indicate that AI investments are associated not only with higher average operational losses, but also with a greater incidence of extreme loss events that are more likely to materially affect institutional stability.

4.5 Risk Management

Risk management functions play a crucial role in assessing, managing, and monitoring risks at banking organizations to ensure they remain within the limits set by management and boards of directors. Prior research has demonstrated that weak risk controls and lack of independence in risk management functions are associated with increased risk exposures at large BHCs (e.g., Ellul and Yerramilli, 2013; Abdymomunov and Mihov, 2019; Frame et al., 2020).[10] Given this established relation between risk management quality and risk outcomes, we examine whether strong risk management practices help banking organizations mitigate the risk-enhancing effects of AI.

To investigate this question, we utilize the risk management index (RMI) developed by Ellul and Yerramilli (2013). The *RMI* provides a continuous measure of the organizational strength and independence of risk management functions at large banking organizations. It is constructed as the first principal component of seven measures capturing BHC risk management quality, including whether an institution has a designated risk officer to manage enterprise-wide risk and how effectively quantitative and qualitative risk information flows

[10] The Wells Fargo cross-selling scandal, in which inadequate internal controls coupled with aggressive sales targets and incentive structures allowed employees to create millions of unauthorized accounts, further illustrates the importance of risk management quality for managing operational risk. The misconduct resulted in significant operational losses for Wells Fargo, including $3 billion in penalties, substantial litigation costs, and severe reputational damage. The banking organization's supervisors also imposed a restriction preventing it from expanding its assets beyond 2017 levels (ultimately lifted in June 2025) until it addressed the underlying risk governance and controls issues. See *Federal Reserve System*: "Responding to widespread consumer abuses and compliance breakdowns by Wells Fargo, Federal Reserve restricts Wells' growth until firm improves governance and controls. Concurrent with Fed action, Wells to replace three directors by April, one by year end" (Feb. 02, 2018); *U.S. Department of Justice*: "Wells Fargo Agrees to Pay $3 Billion to Resolve Criminal and Civil Investigations into Sales Practices Involving the Opening of Millions of Accounts without Customer Authorization" (February 21, 2020).

between business segments and senior management. Higher values of the index indicate more robust risk management practices.

The *RMI* data is available at an annual frequency from the beginning of our sample period through 2013 and covers 17 of the 36 BHCs in our baseline sample. To preserve sample size, and motivated by the view that risk management practices are relatively persistent over time (e.g., Fahlenbrach et al., 2012), we carry forward each BHC's 2013 *RMI* value through the end of the sample period in 2018. To reconcile the annual frequency of the *RMI* data with our quarterly analysis, we assign the annual *RMI* value to all quarters within the corresponding year. We then re-estimate our baseline regression model, augmented to include both the *RMI* term and its interaction with *Share AI Workers*. Table 7 presents the results.

[Insert Table 7 about here]

Column (1) shows that the coefficient estimate on this interaction term, *Share AI Workers* × *RMI*, is negative and significant at the 5% level, indicating that more AI-intensive BHCs with strong risk management functions tend to incur fewer operational losses. The economic magnitude is substantial — increasing *Share AI Workers* by one standard deviation while simultaneously increasing RMI by one standard deviation (i.e., improving BHC risk management quality) decreases *LtA* by 9% relative to its unconditional mean.

To ensure robustness, we also examine these patterns using a simplified binary version of the *RMI* measure. We create an indicator variable equal to 1 for *RMI* values above the sample median and 0 otherwise. The interaction between this binary measure, *RMI (0/1)*, and *Share AI Workers* maintains its negative sign and statistical significance at the 5% level. Collectively, these findings provide strong evidence that robust risk management serves as an important moderating factor, helping banking organizations contain the operational risks that emerge from increased AI adoption. This suggests that strengthening risk management frameworks may be an important prerequisite for banks seeking to expand their use of AI technologies.

5 Robustness Checks

This section explores the robustness of our main empirical findings through multiple angles. We confirm that the positive relation between AI investments and operational losses holds under annual data aggregation, distributed lead-lag models, controls for past losses, BHC fixed effects, additional control variables, and alternative variable definitions.

5.1 Data Aggregation at Annual Level

A potential concern in our analysis is the discrepancy in measurement frequency across variables. Our primary AI investments variable, *Share AI Workers*, is measured annually, whereas operational losses and control variables are measured quarterly. Given the slow-moving nature of *Share AI Workers* over time, this difference is unlikely to introduce substantial mismeasurement.

Nonetheless, to address this potential issue, we conduct a robustness test in which all variables are aggregated to the annual frequency and our baseline regressions are re-estimated. Quarterly operational losses are summed within each year for each BHC. Control variables are aggregated by summing "flow" variables and taking year-end values for "stock" variables, while *Share AI Workers* remains unchanged. This procedure ensures that operational losses and controls are measured on the same temporal scale as *Share AI Workers*, mitigating the risk that frequency differences drive our results.

[Insert Table 8 about here]

Table 8 presents the results of this analysis and shows that our main findings remain robust in this annual aggregation framework. Specifically, the coefficient on *Share AI Workers* remains positive and statistically significant, confirming that higher AI investment intensity is associated with greater operational losses at the annual frequency.

5.2 Distributed Lead-lag Model

In this section, we examine the temporal dynamics of operational losses around AI investments to strengthen our causal interpretation and explore whether the observed effect is transitional or persistent. Specifically, we investigate three possibilities: (i) reverse causality, where banks experiencing higher operational losses adopt AI in response; (ii) pre-existing differences in loss trajectories between high- and low-AI BHCs; and (iii) whether the observed increase in losses represents a short-lived adjustment or a longer-term shift in the bank's operational risk profile.

To do so, we estimate the following distributed lead-lag model (adapted from Aghion et al., 2020; Babina et al., 2024):

$$\begin{aligned} LtA_{i,t} = \beta_t + \beta_i + \beta_1 \Delta Share\ AI\ Workers_{i,[t-12,t-8]} + \\ \beta_2 \Delta Share\ AI\ Workers_{i,[t-8,t-4]} + \beta_3 \Delta Share\ AI\ Workers_{i,[t-4,t]} + \\ \beta_4 \Delta Share\ AI\ Workers_{i,[t,t+4]} + \beta_5 \Delta Share\ AI\ Workers_{i,[t+4,t+8]} + \epsilon_{i,t} \end{aligned} \tag{4}$$

where $\Delta Share\ AI\ Workers_{i,[t-k-4,t-k]}$ represents the annual change in the share of AI workers. We include quarter fixed effects to control for period-specific shocks common across all BHCs. We show specifications with and without BHC fixed effects that absorb BHC-specific time-invariant factors. This model enables us to detect both anticipatory trends and post-adoption effects. The lag coefficients $(\beta_1, \beta_2, \beta_3)$ measure the post-investment impact, and their evolution provides insight into whether the effect is short-lived (e.g., due to AI learning curves) or persistent. The lead coefficients (β_4, β_5) test for reverse causality and pre-trends—if BHCs that eventually invest in AI already exhibit increasing operational losses beforehand, we would expect significant positive coefficients. Following Babina et al. (2024), we adopt the distributed lead–lag approach because AI investments are typically phased in over time rather than made as large, discrete commitments in a single period, making standard event-study designs with sharply defined pre- and post-treatment windows less suitable.

Table 9 presents the results. Among the lagged terms, only β_2 is positive and statistically significant. This indicates that operational losses tend to rise about one year after a bank increases its AI investments. In contrast, the other lag terms (β_1 and β_3) are statistically insignificant, suggesting that the effect does not persist beyond this one-year window. Importantly, both lead coefficients (β_4 and β_5)—which measure whether future increases in AI investments predict current operational losses—are indistinguishable from zero. This pattern rules out reverse causality, where higher losses would trigger AI adoption, and also suggests that AI-intensive banking organizations were not already on an upward trajectory of losses before investing in AI. Together, the insignificance of the lead terms and the delayed timing of the loss increase support for a causal interpretation.

[Insert Table 9 about here]

These dynamics further clarify the nature of the observed losses. The one-year lag between AI investments and higher losses is consistent with transitional risks—ranging from implementation frictions, integration with legacy systems, and human–AI coordination challenges to data pipeline instability or inaccurate inputs during early deployment, workflow disruptions, regulatory issues and compliance breaches, and vulnerabilities introduced in initial rollouts. The lack of persistence in the post-investment period—i.e., the insignificant β_1 coefficient estimate—suggests that the increase in operational losses does not reflect a permanent elevation in risk exposure. Rather, AI adoption appears to trigger a temporary uptick in operational losses before institutions adapt and stabilize. This finding complements our earlier results, indicating that while AI adoption is associated with increased operational risk, its concentration in the first year after investment suggests it can be mitigated if banks anticipate the adjustment period and invest in adequate risk controls.

5.3 Additional Evidence Against Reverse Causality

Our IV regression results in Section 4.2 and distributed lead–lag model estimates in Section 5.2 indicate that reverse causality is unlikely to explain the positive relation between BHC operational losses and AI investments. In this section, we reinforce this conclusion with additional tests that assess whether banks' historical operational loss experiences influence subsequent AI adoption and whether controlling for past losses alters our main findings.

We begin with visual evidence examining whether past operational losses drive subsequent AI investments. Figure 3 presents bar charts of the change in the share of AI-skilled employees from 2010 to 2018 for BHCs sorted into terciles based on past operational losses (LtA), averaged over three different time windows: [2010:Q1-2010:Q4], [2008:Q1-2010:Q4], and [2006:Q1-2010:Q4]. If BHCs systematically respond to operational losses by increasing their AI investments as a risk mitigation strategy, we should observe institutions with higher historical losses subsequently exhibiting greater AI investments. Contrary to this prediction, BHCs in the tercile with the highest historical operational losses do not exhibit the highest subsequent AI investments over the sample period and are overall similar to those of BHCs in the lowest past-loss tercile. These patterns are inconsistent with the notion that past operational losses are a primary driver of AI investment decisions.

We complement this visual analysis with more formal regression tests that augment our baseline model from Equation 2 with controls for trailing operational losses measured over one-year [t-4, t-1], three-year [t-12, t-1], and five-year [t-20, t-1] periods. This approach directly tests whether including past operational loss measures affects the core documented relation between AI investments and (current) operational losses.

[Insert Table 10 about here]

Table 10 presents these results. Controlling for past operational losses has little effect on the coefficient of *Share AI Workers*, which remains positive and statistically significant across all three specifications. The coefficients on the past loss measures themselves are consistently positive, although not robustly significant across all specifications, indicating that past operational losses have predictive power for current risk outcomes during our sample period. Collectively, these analyses provide consistent additional evidence against reverse causality as an explanation for our main findings.

5.4 Fixed Effects, Additional Controls, and Alternative Measures

Our baseline analysis controls for a set of observable BHC characteristics and includes quarter fixed effects to absorb time-specific shocks common to all BHCs. However, it omits BHC fixed effects to allow for cross-sectional variation in AI investment intensity to explain differences in operational loss outcomes. In this section, we demonstrate that our results are robust to including BHC fixed effects, as well as to controlling for specific additional time-varying factors that could influence both AI investments and operational risk.

To assess the role of time-invariant unobserved heterogeneity, we re-estimate the four baseline regressions from Table 3, now including BHC fixed effects. As shown in Table 11, Columns (1)-(4), the coefficient on *Share AI Workers* remains positive and statistically significant in all specifications. This indicates that the documented relation between AI intensity and operational losses holds even after absorbing persistent BHC-specific characteristics.

[Insert Table 11 about here]

Next, we incorporate additional time-varying control variables. First, we account for risk management quality using the risk management index from Ellul and Yerramilli (2013), introduced in Section 4.5. Although this significantly reduces our sample due to data availability, the coefficient on *Share AI Workers* remains positive and significant, suggesting that AI-related risks are not simply driven by differences in risk management quality (Column (5)).

We then examine business complexity, which may simultaneously drive AI adoption and increase operational risk due to coordination challenges or internal control breakdowns. While our baseline specifications already include the ratio of interest to non-interest income—a proxy for business complexity (e.g., Chernobai et al., 2021)—we introduce two additional measures: (i) a Herfindahl-Hirschman Index (HHI) based on the distribution of subsidiaries across four-digit NAICS codes, and (ii) the number of distinct business segments, defined by four-digit NAICS codes, in which a BHC operates subsidiaries. Results presented in Table 11, Columns (6) and (7), show that the coefficient on *Share AI Workers* remains statistically significant and is unaffected by the inclusion of these complexity proxies.

We also evaluate the robustness of our results to controlling for corporate governance proxies. Weak governance may lead managers to pursue complex or high-risk (AI) technology investments that are not aligned with shareholder interests. We include three standard governance indicators: the share of institutional ownership, the proportion of independent directors, and an indicator for CEO–Chair non-duality. Re-estimating our baseline model with these controls (Table 11, Columns (8)-(10)), we find that the positive relation between AI intensity and operational losses persists, suggesting that corporate governance measures do not explain our main findings.

Lastly, we examine whether banks' revenue models captured by the proportion (relative to total revenue) of interest income on loans and leases, interest income on investment securities, income from fiduciary activities, trading revenue, income from investment banking, advisory, brokerage and underwriting fees, venture capital revenue, securitization income and servicing fees revenue explain the positive relation between AI investments and operational losses. We include these detailed revenue controls in Table 11, Column (11), and again find that the coefficient on *Share AI Workers* remains positive and statistically significant.

In addition to incorporating fixed effects and richer controls, we conclude this section by evaluating the robustness of our results to alternative operational risk measures. One concern is that our main operational risk measure—operational losses scaled by total assets (*LtA*)—may conflate the effects of AI investments on losses with its potential effects on the size of banking organizations. Specifically, if AI investments influence both the numerator (losses) and the denominator (assets), the observed association between AI intensity and *LtA* could be partially driven by denominator dynamics rather than a genuine link between AI and operational risk.[11]

We mitigate this concern by redefining *LtA* and holding the denominator—BHC total assets—fixed at its value in 2010:Q1, the beginning of our sample period. This alternative measure eliminates the possibility that time-varying firm size mechanically influences the loss ratio, helping to assess whether the association between AI investments and operational losses holds when firm scale is held constant. We also construct two additional normalized loss measures by scaling operational losses by total expenses and non-interest expenses, also fixed as of 2010:Q1. These definitions compare losses to a bank's initial cost structure rather than its evolving asset base.

As shown in Table 11, Columns (12)-(14), the positive relation between *Share AI Workers* and operational losses persists under all three alternative definitions. This observation is also consistent with earlier results in Table 3, Columns (2)-(4), which indicate the robustness of our results to using non-scaled operational loss metrics such as log-transformed dollar losses and loss event frequency, as used in prior studies such as Chernobai et al. (2011), Chernobai

[11]In untabulated results, we find that (log) total assets are not significantly related to *Share AI Workers* in regression specifications similar to Equation 2. Moreover, although the association is statistically insignificant, the estimated coefficients are positive. This direction of association—if anything—should attenuate the positive relation between *LtA* and *Share AI Workers*, since larger total assets in the denominator would reduce the *LtA* ratio.

et al. (2021), and Abdymomunov et al. (2020). Overall, these results suggest that the risk-increasing effect of AI adoption is not driven by BHC total asset dynamics.

6 Conclusion

This paper examines the relation between AI investments and operational risk at large U.S. banking organizations. Using comprehensive supervisory data on operational losses and a measure of AI investments based on AI-skilled human capital, we find that banking organizations with higher AI investments tend to experience larger operational losses per dollar of assets. This finding is robust to accounting for endogeneity concerns through an instrumental variables approach based on banks' historical connections to AI-strong universities, as well as numerous additional robustness checks.

Our analysis reveals several important nuances in how AI investments relate to operational risk. First, AI investments affect not only the average level of operational losses but also increase the frequency of tail risk events that can affect bank capital and stability. Second, the impact varies across different types of operational losses, with significant effects observed in external fraud, client-related issues, and system failures. This heterogeneity suggests AI adoption may create new vulnerabilities through, for example, increased digital infrastructure complexity, challenges in regulatory compliance, and greater exposure to technical failures. Third, strong risk management practices attenuate these risks, highlighting the important role of risk governance in mitigating AI-related operational vulnerabilities.

These findings have important implications for banking organization management and supervision. While prior research has documented various benefits of AI adoption, including improved productivity and innovation, our results suggest that these benefits come with meaningful operational risks that need to be carefully managed. Banking organizations may benefit from evaluating and enhancing their risk management frameworks as they increase their AI investments, while supervisors might consider incorporating AI-related operational risk into existing monitoring frameworks, particularly for institutions with less robust risk management practices or historically elevated operational losses.

References

Abdymomunov, A., Curti, F., and Mihov, A. (2020). U.S. banking sector operational losses and the macroeconomic environment. *Journal of Money, Credit and Banking*, 52(1):115–144. https://doi.org/10.1111/jmcb.12661.

Abdymomunov, A. and Mihov, A. (2019). Operational risk and risk management quality: Evidence from U.S. bank holding companies. *Journal of Financial Services Research*, 56(1):73–93. https://doi.org/10.1007/s10693-017-0284-3.

Abis, S. and Veldkamp, L. (2023). The changing economics of knowledge production. *Review of Financial Studies*, 37(1):89–118. https://doi.org/10.1093/rfs/hhad059.

Acemoglu, D. and Restrepo, P. (2018). Artificial intelligence, automation, and work. *NBER Working Paper 24196*. http://www.nber.org/papers/w24196.

Adhaen, M., Chen, W., Wadi, R. A., and Aldhaen, E. (2024). Exploring artificial intelligence adoption in the banking sector: Multiple case studies. In Hamdan, A. and Braendle, U., editors, *Harnessing AI, Machine Learning, and IoT for Intelligent Business*, page 301–314. Springer.

Afonso, G., Curti, F., and Mihov, A. (2019). Coming to terms with operational risk. *Federal Reserve Bank of New York Liberty Street Economics*.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2020). What are the labor and product market effects of automation? New evidence from France. CEPR Discussion Paper No. DP14443.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50. https://doi.org/10.1257/jep.33.2.31.

Ahmed, A. S., Takeda, C., and Thomas, S. (1999). Bank loan loss provisions: A re-examination of capital management, earnings management and signaling effects. *Journal of Accounting and Economics*, 28(1):1–25. https://doi.org/10.1016/s0165-4101(99)00017-8.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2023a). Artificial intelligence and firms' systematic risk. *Working Paper*. http://dx.doi.org/10.2139/ssrn.4868770.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2023b). Firm investments in artificial intelligence technologies and changes in workforce composition. *Technology, Productivity, and Economic Growth, forthcoming*, 83. NBER Studies in Income and Wealth editted by Susanto Basu and Lucy Eldridge and John Haltiwanger and Erich Strassner.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2025). Artificial intelligence makes firm operating performance less volatile. *AEA Papers and Proceedings*, 115:35–39. https://doi.org/10.1257/pandp.20251002.

Bartlett, R., Morse, A., Stanton, R., and Wallace, N. (2022). Consumer-lending discrimination in the fintech era. *Journal of Financial Economics*, 143(1):30–56.

Basel Committee on Banking Supervision (2006). International convergence of capital measurement and capital standards. Bank of International Settlements.

Berger, A. N., Curti, F., Lazaryan, N., Mihov, A., and Roman, R. A. (2022a). Climate risks in the U.S. banking sector: Evidence from operational losses and extreme storms. *Working Paper*. http://dx.doi.org/10.2139/ssrn.4294026.

Berger, A. N., Curti, F., Mihov, A., and Sedunov, J. (2022b). Operational risk is more systemic than you think: Evidence from U.S. bank holding companies. *Journal of Banking & Finance*, 143:106619. https://doi.org/10.1016/j.jbankfin.2022.106619.

Bian, B., Pagel, M., Tang, H., and Raval, D. (2023). Consumer surveillance and financial fraud. Working Paper 31692, National Bureau of Economic Research. https://doi.org/10.3386/w31692.

Cao, S., Jiang, W., Wang, J., and Yang, B. (2024). From man vs. machine to man + machine: The art and AI of stock analyses. *Journal of Financial Economics*, 160:103910. https://doi.org/10.1016/j.jfineco.2024.103910.

Chen, M. A., Wu, Q., and Yang, B. (2019). How valuable is fintech innovation? *Review of Financial Studies*, 32(5):2062–2106. https://doi.org/10.1093/rfs/hhy130.

Chen, W. X., Shi, T. T., and Srinivasan, S. (2024). The value of AI innovations. *Harvard Business School Working Paper, No. 24-069.*

Chernobai, A., Curti, F., Mihov, A., and Xiong, X. (2024). Operational losses and insider trading: Evidence from U.S. financial institutions. *Working Paper*. http://dx.doi.org/10.2139/ssrn.4834650.

Chernobai, A., Jorion, P., and Yu, F. (2011). The determinants of operational risk in U.S. financial institutions. *Journal of Financial and Quantitative Analysis*, 46(6):1683–1725. https://doi.org/10.1017/s0022109011000500.

Chernobai, A., Ozdagli, A., and Wang, J. (2021). Business complexity and risk management: Evidence from operational risk events in U.S. bank holding companies. *Journal of Monetary Economics*, 117:418–440. https://doi.org/10.1016/j.jmoneco.2020.02.004.

Cockburn, I. M., Henderson, R., and Stern, S. (2018). The impact of artificial intelligence on innovation. Working Paper w24449, National Bureau of Economic Research. http://dx.doi.org/10.3386/w24449.

Cook, T. R. and Kazinnik, S. (2024). Social bias in financial applications of large language models. *Working paper.*

Cope, E. W. and Carrivick, L. (2013). Effects of the financial crisis on banking operational losses. *Journal of Operational Risk*, 8(3):3–29. https://doi.org/10.21314/jop.2013.125.

Correlation One (2019). Future of data talent. *2019 Annual Report.*

Cummins, J. D., Lewis, C. M., and Wei, R. (2006). The market value impact of operational loss events for US banks and insurers. *Journal of Banking & Finance*, 30(10):2605–2634. https://doi.org/10.1016/j.jbankfin.2005.09.015.

Curti, F., Fauver, L., and Mihov, A. (2023). Workforce policies and operational risk: Evidence from U.S. bank holding companies. *Journal of Financial and Quantitative Analysis*, 58(7):3085–3120. https://doi.org/10.1017/S0022109022000989.

Curti, F., Frame, W. S., and Mihov, A. (2022). Are the largest banking organizations operationally more risky? *Journal of Money, Credit and Banking*, 54(5):1223–1259. https://doi.org/10.1111/jmcb.12933.

Curti, F., Macchiavelli, M., Mihov, A., and Pisciotta, K. (2024). Corporate espionage and innovation: Evidence from the theft of trade secrets. *Working paper.* http://dx.doi.org/10.2139/ssrn.4613975.

Curti, F. and Mihov, A. (2021). Catch the thief! fraud in the U.S. banking industry. *Working Paper.* http://dx.doi.org/10.2139/ssrn.3535934.

Czarnitzki, D., Fernández, G. P., and Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211:188–205.

Daníelsson, J., Macrae, R., and Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140:106290. https://doi.org/10.1016/j.jbankfin.2021.106290.

de Fontnouvelle, P., Dejesus-Rueff, V., Jordan, J. S., and Rosengren, E. S. (2006). Capital and risk: New evidence on implications of large operational losses. *Journal of Money, Credit and Banking*, 38(7):1819–1846. https://doi.org/10.1353/mcb.2006.0088.

Durongkadej, I., Hu, W., and Wang, H. E. (2024). How artificial intelligence incidents affect banks and financial services firms? a study of five firms. *Finance Research Letters*, 70:106279. https://doi.org/10.1016/j.frl.2024.106279.

D'Acunto, F., Prabhala, N., and Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *Review of Financial Studies*, 32(5):1983–2020. https://doi.org/10.1093/rfs/hhz014.

Ebrahimitorki, M. and Kim, H. H. (2025). Harnessing artificial intelligence: Impact of AI adoption on bank loan performance. *Working Paper*. http://dx.doi.org/10.2139/ssrn.5226964.

Ellul, A. and Yerramilli, V. (2013). Stronger risk controls, lower risk: Evidence from U.S. bank holding companies. *Journal of Finance*, 68(5):1757–1803. https://doi.org/10.1111/jofi.12057.

European Central Bank (2025). The rise of artificial intelligence: Benefits and risks for financial stability. *Financial Stability Review, May 2024*.

Fahlenbrach, R., Prilmeier, R., and Stulz, R. M. (2012). This time is the same: Using bank performance in 1998 to explain bank performance during the recent financial crisis. *Journal of Finance*, 67(6):2139–2185.

Fedyk, A. and Hodson, J. (2022). Trading on talent: Human capital and firm performance. *Review of Finance*, 27(5):1659–1698. https://doi.org/10.1093/rof/rfac068.

Fedyk, A., Hodson, J., Khimich, N., and Fedyk, T. (2022). Is artificial intelligence improving the audit process? *Review of Accounting Studies*, 27:938–985. https://doi.org/10.1007/s11142-022-09697-x.

Florackis, C., Louca, C., Michaely, R., and Weber, M. (2022). Cybersecurity risk. *Review of Financial Studies*, 36(1):351–407. https://doi.org/10.1093/rfs/hhac024.

Frame, W. S., Lazaryan, N., McLemore, P., and Mihov, A. (2024). Operational loss recoveries and the macroeconomic environment: Evidence from the U.S. banking sector. *Journal of Banking & Finance*, 165:107220. https://doi.org/10.1016/j.jbankfin.2024.107220.

Frame, W. S., McLemore, P., and Mihov, A. (2023). Financial innovation and risk: Evidence from operational losses at U.S. banking organizations. *Working Paper*. http://dx.doi.org/10.2139/ssrn.4621295.

Frame, W. S., McLemore, P., and Mihov, A. (2025). Haste makes waste: Banking organization growth and operational risk. *Review of Corporate Finance Studies, forthcoming*.

Frame, W. S., Mihov, A., and Sanz, L. (2020). Foreign investment, regulatory arbitrage, and the risk of U.S. banking organizations. *Journal of Financial and Quantitative Analysis*, 55(3):955–988. https://doi.org/10.1017/S0022109019000267.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance*, 77(1):5–47. https://doi.org/10.1111/jofi.13090.

Gillet, R., Hubner, G., and Plunus, S. (2010). Operational risk and reputation in the financial industry. *Journal of Banking & Finance*, 34(1):224–235. https://doi.org/10.1016/j.jbankfin.2009.07.020.

Gofman, M. and Jin, Z. (2024). Artificial intelligence, education, and entrepreneurship. *Journal of Finance*, 79(1):631–667. https://doi.org/10.1111/jofi.13302.

Gomes, O., Mihet, R., and Rishabh, K. (2024). Data risk, firm growth, and innovation. *Swiss Finance Institute Research Paper No. 23-86*. http://dx.doi.org/10.2139/ssrn.4559921.

Grennan, J. and Michaely, R. (2020). Artificial intelligence and high-skilled work: Evidence from analysts. *Swiss Finance Institute Research Paper No. 20-84*. http://dx.doi.org/10.2139/ssrn.3681574.

Han, M., Shen, H., Wu, J., and Zhang, X. M. (2025). Artificial intelligence and firm resilience: Empirical evidence from natural disaster shocks. *Information Systems Research, forthcoming*. 10.1287/isre.2022.0440.

Hansen, A. L., Horton, J. J., Kazinnik, S., Puzzello, D., and Zarifhonarvar, A. (2025). Simulating the survey of professional forecasters. *Working Paper*.

Hess, C. (2011). The impact of the financial crisis on operational risk in the financial services industry: Empirical evidence. *Journal of Operational Risk*, 6(1):23–35. https://doi.org/10.21314/jop.2011.087.

James, C. (1991). The losses realized in bank failures. *Journal of Finance*, 46(4):1223–1242. https://doi.org/10.1111/j.1540-6261.1991.tb04616.x.

Jin, L. and Myers, S. C. (2006). R^2 around the world: New theory and new tests. *Journal of Financial Economics*, 79(2):257–292. https://doi.org/10.1016/j.jfineco.2004.11.003.

New York State Department of Financial Services (2024). Industry letter: Cybersecurity risks arising from artificial intelligence and strategies to combat related risks.

Stock, J. H., Wright, J. H., and Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics*, 20(4):pp. 518–529. https://doi.org/10.1198/073500102288618658.

U.S. Department of the Treasury (2024). Managing artificial intelligence-specific cybersecurity risks in the financial services sector.

Wang, T. and Hsu, C. (2013). Board composition and operational risk events of financial institutions. *Journal of Banking & Finance*, 37(6):2042–2051. https://doi.org/10.1016/j.jbankfin.2013.01.027.

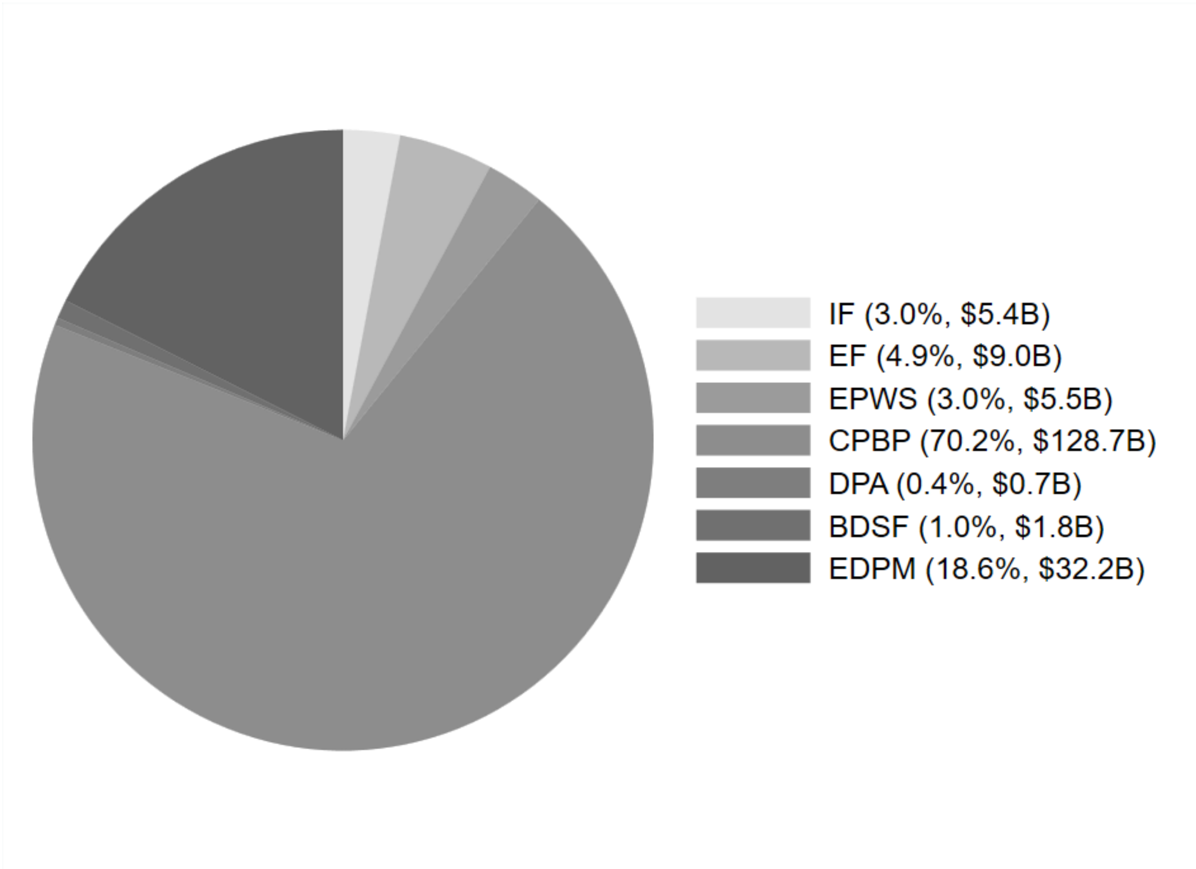

Figure 1: **Operational Losses by Event Type**
This figure presents the percentage allocation of losses among the seven operational risk event type categories. Losses in each category are first averaged within bank holding companies (BHCs) and then averaged across BHCs. The nomenclature for event types is as follows: Internal Fraud (IF), External Fraud (EF), Employment Practices and Workplace Safety (EPWS), Clients, Products and Business Practices (CPBP), Damage to Physical Assets (DPA), Business Disruption and System Failures (BDSF), and Execution, Delivery and Process Management (EDPM). The sample includes 852 operational losses incurred by 36 large U.S. BHCs over the period [2010:Q1-2018:Q4].

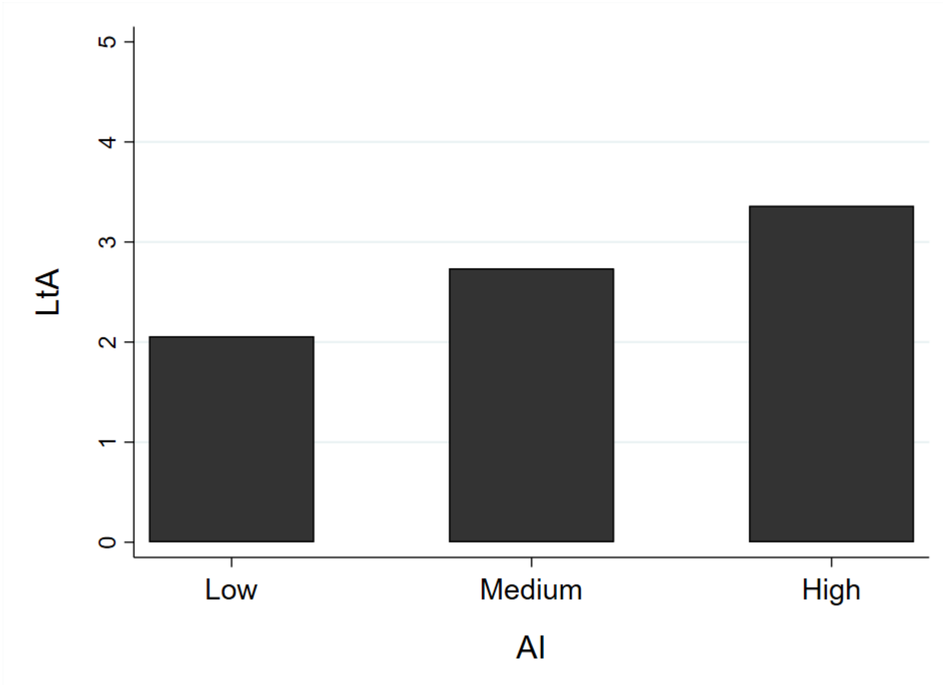

Figure 2: **Operational Losses by AI Investment Groups**
This figure presents a bar chart of the average ratio of operational losses to (lagged) total assets (multiplied by 10,000), *LtA*, for bank holding companies (BHCs) sorted in terciles based on share of AI-skilled employees: “Low”, “Medium,” and “High”. The chart presents the average *LtA* for each tercile. The sample comprises an unbalanced panel of 852 quarterly observations of 36 large U.S. BHCs over the period [2010:Q1-2018:Q4].

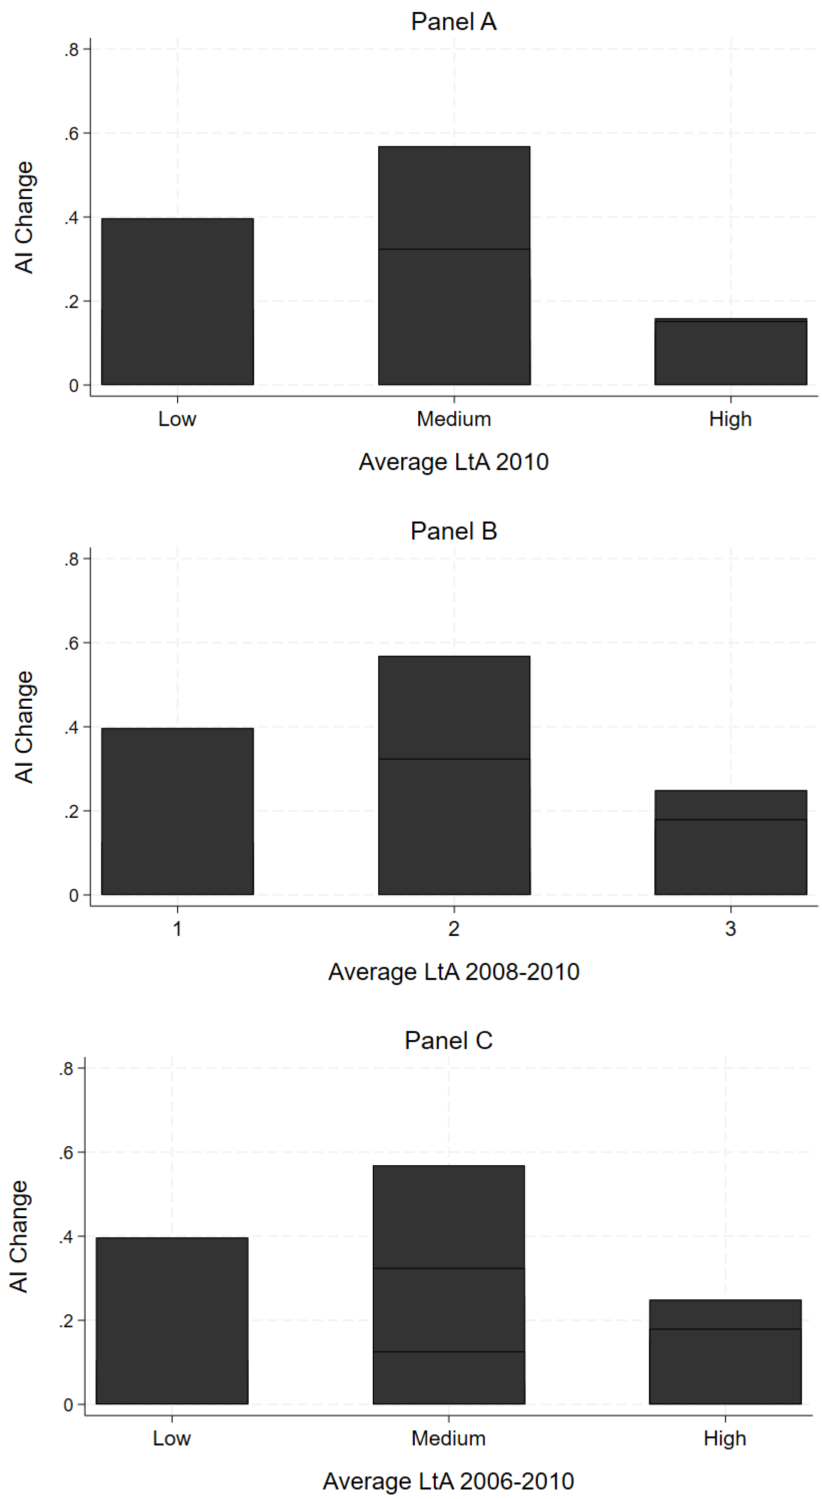

Figure 3: **AI Investments and Past Operational Losses**
This figure presents bar charts of the change in the share of AI-skilled employees from 2010 to 2018 for bank holding companies (BHCs) sorted in terciles ("Low", "Medium," and "High") based on past operational losses, *LtA*, measured over different time windows. *LtA* is defined as the ratio of operational losses to (lagged) total assets (multiplied by 10,000). In Panel A, *LtA* is averaged over [2010:Q1, 2010:Q4]. In Panel B, *LtA* is averaged over [2008:Q1, 2010:Q4]. In Panel C, *LtA* is averaged over [2006:Q1, 2010:Q4]. The sample comprises observations from 36 large U.S. BHCs.

Table 1: **Operational Loss Event Type and Variable Definitions**
This table presents operational loss event type definitions according to Basel Committee on Banking Supervision (2006) in Panel A and variable definitions in Panel B.

Panel A: Operational Loss Event Types

Event Type Category	Short	Description
Internal Fraud	IF	Acts of a type intended to defraud, misappropriate property or circumvent regulations, which involves at least one internal party
External Fraud	EF	Acts of a type intended to defraud, misappropriate property or circumvent the law, by a third party
Employment Practices and Workplace Safety	EPWS	Acts inconsistent with employment, health or safety laws or agreements, from payment of personal injury claims, or from diversity / discrimination events
Clients, Products and Business Practices	CPBP	An unintentional or negligent failure to meet a professional obligation to specific clients, or from the nature or design of a product
Damage to Physical Assets	DPA	Damage to physical assets from natural disasters or other events
Business Disruption and System Failures	BDSF	Disruption of business or system failures
Execution, Delivery and Process Management	EDPM	Failed transaction processing or process management, from relations with trade counterparties and vendors

Panel B: Variable Definitions	
Variable	**Definition**
Operational Risk Variables	
Avg Sev	The average severity of operational losses that occur at a BHC over a given calendar quarter in millions of U.S. Dollars
Loss	Operational losses that occur at a BHC over a given calendar quarter in millions of U.S. Dollars
LtA	Operational losses that occur at a BHC over a given calendar quarter as a proportion of the BHC's lagged total assets, multiplied by 10,000
LtA_{2010}	Operational losses that occur at a BHC over a given calendar quarter as a proportion of the BHC's total assets measured as of 2010:Q1, multiplied by 10,000
LtE_{2010}, $LtNIE_{2010}$	Operational losses that occur at a BHC over a given calendar quarter as a proportion of either the BHC's total expenses or the BHC's non-interest expenses measured as of 2010:Q1, multiplied by 10,000
LtA Tail (90, 95, 99)	Tail operational losses at the 90th, 95th, or 99th percentile, respectively, that occur at a BHC over a given calendar quarter as a proportion of the BHC's lagged total assets, multiplied by 10,000
N Evts	The number of operational losses that occur at a BHC over a given calendar quarter
N Evts Tail (90, 95, 99)	The frequency of total assets-scaled tail operational losses at the 90th, 95th, or 99th percentile, respectively, that occur at a BHC over a given calendar quarter
AI and Baseline Control Variables	
Share AI Workers	The share of a BHC's AI-skilled employees relative to total number of BHC employees
Assets	BHC total assets in billions of U.S. Dollars
Asset Growth	Year-over-year growth in BHC consolidated total assets
II-to-NII	The ratio of BHC interest income to non-interest income
Leverage	The ratio of BHC total assets to book value of equity
Maturity Gap	A natural log transformation of the absolute difference between all assets that either reprice or mature within a year and all liabilities that reprice or mature within a year
Non-Performing Loans	Loans 90 days or more past due as a proportion of total loans

Variable	Definition
N Patents	The number of (successful) financial patent applications by a BHC
ROE	BHC return on equity, defined as the ratio of net income to book value of equity
RMI (0/1)	RMI is the risk-management index developed by Ellul and Yerramilli (2013). RMI (0/1) is an indicator variable equal to 1 if RMI is greater than the sample median, and 0 otherwise
Additional Control Variables	
CEO Non-Duality	An indicator variable equal to 1 if the CEO of the BHC is different than the Board Chair
Fiduciary Activities	Income from fiduciary activities (as a proportion of total income)
HHI Business Segments	A Herfindahl-Hirshman index defined as the sum of squared shares of subsidiaries in distinct business segments (4-digit NAICS industry codes), where the shares are expressed as a percentage of total number of subsidiaries
Independent Directors	The number of independent directors divided by total Board members defined in percent
Institutional Ownership	Total institutional ownership as a percent of the total shares outstanding
Investment Banking	Income from investment banking, advisory, brokerage and underwriting fees (as a proportion of total income)
Investment Securities	Interest income on investment securities (as a proportion of total income)
Loans and Leases	Interest income on loans and leases (as a proportion of total income)
N Business Segments	Number of distinct business segments (4-digit NAICS industry codes) of subsidiaries owned by the BHC
Securitization	Securitization income (as a proportion of total income)
Servicing Fees	Revenue from servicing fees (as a proportion of total income)
Trading	Trading revenue (as a proportion of total income)
Venture Capital	Venture capital revenue (as a proportion of total income)

Table 2: **Descriptive Statistics**
This table presents variable descriptive statistics. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B.

	Mean	SD	P25	P50	P75	N
Op Risk Variables:						
LtA	2.811	5.980	0.454	0.964	2.242	852
Loss ($M)	181.385	644.672	6.496	20.755	115.371	852
N Evts	252.288	370.721	36.500	86.000	271.000	852
Avg Sev	0.774	2.252	0.128	0.228	0.493	852
N Tail 90	23.312	21.450	12.000	17.000	26.000	852
N Tail 95	11.851	11.566	6.000	9.000	13.000	852
N Tail 99	2.580	3.030	1.000	2.000	3.000	852
LtA Tail 90	2.619	5.949	0.340	0.762	1.998	852
LtA Tail 95	2.534	5.938	0.282	0.687	1.892	852
LtA Tail 99	2.299	5.898	0.109	0.460	1.525	852
Other Variables:						
Share AI Workers	0.065	0.107	0.000	0.028	0.089	852
Assets ($B)	504.212	687.214	118.137	174.373	385.799	852
Asset Growth	0.037	0.104	-0.013	0.027	0.069	852
II-to-NII	1.791	2.443	0.625	1.547	2.255	852
ROE	0.018	0.024	0.010	0.019	0.027	852
N Patents	0.352	0.732	0.000	0.000	0.000	852
Leverage	0.884	0.025	0.870	0.886	0.899	852
Non-Performing Loans	0.297	0.409	0.074	0.173	0.351	852
Maturity Gap	18.119	1.239	17.336	17.845	18.723	852

Table 3: **Operational Losses and Share of AI-Skilled Employees**
This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees and control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Columns (1), (2), and (4) are estimated with Ordinary Least Squares regressions. Column (3) is estimated with a Negative Binomial regression. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) Log(Loss)	(3) N Evts	(4) Log(Avg Sev)
Share AI Workers	6.394***	2.445***	1.440***	1.166**
	(0.000)	(0.001)	(0.000)	(0.030)
Log(Assets)	0.660	1.139***	0.793***	0.335**
	(0.100)	(0.000)	(0.000)	(0.018)
Asset Growth	3.065	0.227	−0.077	0.233
	(0.314)	(0.745)	(0.711)	(0.729)
II-to-NII	−0.095	−0.018	0.002	−0.029
	(0.115)	(0.360)	(0.780)	(0.272)
ROE	9.523	5.692*	5.929***	−0.391
	(0.332)	(0.052)	(0.000)	(0.808)
Log(N Patents)	0.041	0.155	0.273***	−0.157*
	(0.872)	(0.116)	(0.000)	(0.096)
Leverage	−8.954	−5.342	−7.604***	1.303
	(0.376)	(0.121)	(0.000)	(0.663)
Non-Performing Loans	−0.110	−0.157	−0.468***	0.306***
	(0.912)	(0.405)	(0.000)	(0.003)
Maturity Gap	0.056	0.150	0.108**	0.055
	(0.834)	(0.325)	(0.012)	(0.635)
N Obs	852	852	852	852
Adj R^2	0.07	0.67		0.17

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: **Instrumental Variables**

This table reports results of instrumental variables regressions of operational losses on the share of AI-skilled employees and control variables. We instrument the share of AI-skilled employees with BHCs' ex-ante exposure to the subsequent supply of AI talent from universities that are historically strong in AI research. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Control variables (*Log(Assets)*, *Asset Growth*, *II-to-NII*, *ROE*, *Log(N Patents)*, *Leverage*, *Non-Performing Loans*, and *Maturity Gap*) are included, but their coefficient estimates are omitted for brevity. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) Share AI Workers	(2) LtA	(3) LtA	(4) LtA	(5) LtA	(6) LtA	(7) LtA
IV	0.281***						
	(0.000)						
Share AI Workers		12.841***	11.993***	9.619*	11.178**	17.159**	11.464***
		(0.004)	(0.009)	(0.072)	(0.013)	(0.013)	(0.009)
Baseline Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
CS-Strong University	*No*	*No*	*Yes*	*No*	*No*	*No*	*Yes*
Commuting Zone Attributes	*No*	*No*	*No*	*Yes*	*No*	*No*	*Yes*
Loss 2000-2009	*No*	*No*	*No*	*No*	*Yes*	*No*	*Yes*
Division Fixed Effects	*No*	*No*	*No*	*No*	*No*	*Yes*	*Yes*
N Obs	852	852	852	675	802	852	633
Adj R^2	0.41	0.02	0.02	0.02	0.02	0.02	0.01

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: **Operational Loss Types**

This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees and control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Operational losses are categorized into seven categories: Internal Fraud (IF), External Fraud (EF), Employment Practices and Workplace Safety (EPWS), Clients, Products and Business Practices (CPBP), Damage to Physical Assets (DPA), Business Disruption and System Failures (BDSF), and Execution, Delivery and Process Management (EDPM). The definitions of operational loss event types are presented in Table 1, Panel A. Control variables (*Log(Assets)*, *Asset Growth*, *II-to-NII*, *ROE*, *Log(N Patents)*, *Leverage*, *Non-Performing Loans*, and *Maturity Gap*) are included, but their coefficient estimates are omitted for brevity. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA IF	(2) LtA EF	(3) LtA EPWS	(4) LtA CPBP	(5) LtA DPA	(6) LtA BDSF	(7) LtA EDPM
Share AI Workers	1.178	1.612***	−0.064	2.236**	−0.027	0.065*	1.395
	(0.313)	(0.005)	(0.161)	(0.047)	(0.294)	(0.086)	(0.176)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N Obs	852	852	852	852	852	852	852
Adj R^2	0.00	0.03	0.06	0.06	0.11	0.05	0.04

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: **Tail Operational Losses**

This table reports coefficients from panel regressions of tail operational losses on the share of AI-skilled employees and control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Columns (1)-(3) are estimated with Negative Binomial regressions and Columns (4)-(6) are estimated with Ordinary Least Square regressions. Control variables (*Log(Assets)*, *Asset Growth*, *II-to-NII*, *ROE*, *Log(N Patents)*, *Leverage*, *Non-Performing Loans*, and *Maturity Gap*) are included, but their coefficient estimates are omitted for brevity. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) N Evts Tail 90	(2) N Evts Tail 95	(3) N Evts Tail 99	(4) LtA Tail 90	(5) LtA Tail 95	(6) LtA Tail 99
Share AI Workers	1.823***	2.204***	2.450***	6.112***	5.979***	5.358***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
N Obs	852	852	852	852	852	852
Adj R^2				0.06	0.06	0.05

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: **Risk Management Quality**

This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees, measures of risk management quality, their interactions and control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Control variables (*Log(Assets)*, *Asset Growth*, *II-to-NII*, *ROE*, *Log(N Patents)*, *Leverage*, *Non-Performing Loans*, and *Maturity Gap*) are included, but their coefficient estimates are omitted for brevity. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) LtA
Share AI Workers	38.806**	4.320***
	(0.014)	(0.006)
Share AI Workers × RM	−45.000**	
	(0.030)	
RMI	−0.551	
	(0.697)	
Share AI Workers × RMI (0/1)		−13.649**
		(0.048)
RMI (0/1)		0.368
		(0.476)
Controls	*Yes*	*Yes*
N Obs	486	486
Adj R^2	0.07	0.07

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: **Data Aggregation at the Annual Level**

This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees and control variables. The sample includes 216 annual observations of 36 large U.S. bank holding companies (BHCs) over the period [2010-2018] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Columns (1), (2), and (4) are estimated with Ordinary Least Squares regressions. Column (3) is estimated with a Negative Binomial regression. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) Log(Loss)	(3) N Evts	(4) Log(Avg Sev)
Share AI Workers	19.497***	2.213***	0.878**	1.286**
	(0.000)	(0.001)	(0.017)	(0.035)
Controls	*Yes*	*Yes*	*Yes*	*Yes*
N Obs	216	216	216	216
Adj R^2	0.17	0.74		0.23

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: **Distributed Lead-Lag Model**
This table reports coefficients from distributed lead-lag panel regressions of operational losses on (lead and lag) changes in the share of AI-skilled employees. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Columns (1) and (4) use 3 lag terms. Columns (2) and (5) use 2 lead terms. Columns (3) and (6) use 5 lead-lag terms. Columns (1)–(3) include quarter fixed effects, while Columns (4)–(6) include both quarter and BHC fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) LtA	(3) LtA	(4) LtA	(5) LtA	(6) LtA
ΔShare AI Workers $[t-12, t-8]$	-5.467		-6.168	-12.566		-11.988
	(0.425)		(0.571)	(0.128)		(0.180)
ΔShare AI Workers $[t-8, t-4]$	38.525***		33.978***	33.809**		30.585***
	(0.009)		(0.006)	(0.021)		(0.009)
ΔShare AI Workers $[t-4, t]$	-5.778		-10.825	-12.679		-13.653
	(0.635)		(0.293)	(0.367)		(0.270)
ΔShare AI Workers $[t, t+4]$		21.847	15.960		13.622	13.101
		(0.114)	(0.107)		(0.307)	(0.185)
ΔShare AI Workers $[t+4, t+8]$		-3.976	-3.013		-8.277	-6.078
		(0.690)	(0.741)		(0.515)	(0.605)
N Obs	770	733	729	770	732	728
Adj R^2	0.06	0.06	0.06	0.06	0.06	0.06

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: **Controlling for Past Operational Losses**
This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees, controls for past operational losses and other control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. All specifications include quarter fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) LtA	(3) LtA
Share AI Workers	5.463***	5.198***	5.549***
	(0.000)	(0.000)	(0.000)
LtA $[t-4, t-1]$	0.172		
	(0.234)		
LtA $[t-12, t-1]$		0.297***	
		(0.000)	
LtA $[t-20, t-1]$			0.240***
			(0.002)
Controls	*Yes*	*Yes*	*Yes*
N Obs	852	852	852
Adj R^2	0.07	0.08	0.07

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: **Robustness Checks**

This table reports coefficients from panel regressions of operational losses on the share of AI-skilled employees and control variables. The sample includes 852 quarterly observations of 36 large U.S. bank holding companies (BHCs) over the period [2010:Q1-2018:Q4] for which requisite data is available. The definitions of all variables are reported in Table 1, Panel B. Control variables (*Log(Assets)*, *Asset Growth*, *II-to-NII*, *ROE*, *Log(N Patents)*, *Leverage*, *Non-Performing Loans*, and *Maturity Gap*) are included, but their coefficient estimates are omitted for brevity. All specifications include quarter and BHC fixed effects. The error terms are clustered at the BHC level. p-values are presented in parentheses.

	(1) LtA	(2) Log(Loss)	(3) N Evts	(4) Log(Avg Sev)	(5) LtA	(6) LtA	(7) LtA	(8) LtA	(9) LtA	(10) LtA	(11) LtA	(12) LtA_{2010}	(13) LtE_{2010}	(14) $LtNIE_{2010}$
Share AI Workers	4.306**	1.078***	0.570***	0.546*	4.695***	4.284**	4.650**	4.872**	4.252**	4.409**	4.308**	6.299***	576.138***	781.406***
	(0.022)	(0.001)	(0.000)	(0.075)	(0.000)	(0.036)	(0.017)	(0.013)	(0.023)	(0.017)	(0.032)	(0.010)	(0.004)	(0.003)
Controls	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
N Obs	852	852	852	852	486	852	852	766	850	852	824	840	835	835
Adj R^2	0.06	0.71		0.26	0.07	0.06	0.06	0.07	0.06	0.06	0.06	0.06	0.07	0.07

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$