University of Zurich
Department of Finance
Plattenstrasse 14
8032 Zürich
Switzerland
https://www.df.uzh.ch/en.html

Tingyu Yu

Postdoctoral Researcher in Finance

tingyu.yu@df.uzh.ch

To Committee Members of the NBER PhD Student Workshop on Climate Finance
Zurich, August 7, 2025

Dear Committee Members,

My name is Tingyu Yu and I am a postdoctoral researcher in finance at the University of Zurich. I am pleased to submit my job market paper, "Carbon in the Cloud," for consideration at the NBER PhD Student Workshop on Climate Finance. I believe presenting at this workshop would provide an invaluable opportunity to receive feedback from leading researchers and peers in the field.

My research lies at the intersection of climate finance, corporate innovation, and natural language processing. I develop AI-driven methods to measure corporate climate exposures and to price climate-related risks across financial assets. Moving forward, my research will continue to focus on firm-level exposure arising from the transition to a low-carbon economy as well as explore the climate externalities of corporate AI adoption.

I obtained my PhD in Financial Mathematics from the Hong Kong University of Science and Technology in 2022.

Professor Zacharias Sautner (zacharias.sautner@df.uzh.ch) at the University of Zurich has kindly agreed to serve as my reference.

Thanks a lot for your time and consideration.

Sincerely,

Tingyu Yu

TINGYU YU

Email: tingyu.yu@uzh.ch ⋄ Address: PLR-J-101 Plattenstr. 14, 8032 Zürich.

Research Interests

Climate finance, corporate innovation, empirical asset pricing, natural language processing (NLP).

Employment

University of Zurich *Switzerland*
Postdoctoral researcher, Department of Finance *2022 - Present*

Education

The Hong Kong University of Science and Technology *Hong Kong SAR*
Ph.D. in Financial Mathematics *2018 - 2022*
Southern University of Science and Technology *China*
Bachelor of Economics *2014 - 2018*
King's College London *United Kingdom*
Summer Program in International Commercial Law *2017*

Working Papers

Carbon in the Cloud (Job Market Paper)

- Abstract: This paper examines the climate impact of corporate artificial intelligence (AI) adoption in a large global sample. AI infrastructure developers bear most of AI's energy burden but partially offset it with renewable electricity. In contrast, user firms with AI talent experience plausibly causal reductions in subsequent Scope 1 emissions, as AI improves demand forecasting and enables shifts toward lower-carbon products and processes. Notably, these cuts are concentrated among low-emitting firms, while some carbon-intensive sectors experience emission increases due to capacity maximization. This underscores the need to consider not only AI's computing energy demands but also the broader operational contexts in which it is deployed. Climate-conscious investors appear to incorporate these climate risks into their capital allocation.

Corporate Climate Lobbying with M. Leippold and Z. Sautner

- *Honorable Mention for the 2024 Moskowitz Prize.*

Firm-Level Green Innovation Beyond Patents with M. Leippold

- *Accepted by the Review of Finance.*

The Anatomy of Corporate Climate Exposure with W. Hu and K. Li
Real-Time Climate Controversy Detection with D. Jaggi, N. Jamet, and M. Leippold

Publications or Conference Papers in NLP

ChatClimate: Grounding Conversational AI in Climate Science
(with S. Vaghefi et. al.) *Nature Communications Earth & Environment*, 4:480, 2023.
ChatReport: Democratizing Sustainability Disclosure Analysis through LLM-based Tools
(with J. Ni et. al.) *EMNLP Demo*, 2023
Exploring Nature: Datasets and Models for Analyzing Nature-Related Disclosures
(with T. Schimanski et. al.) *AAAI Workshop on AI in Finance for Social Impact*, 2024
Integrating Artificial Intelligence with Expert Knowledge in Global Environmental Assessments
(with V. Muccione et. al.) *Regional Environmental Change*, 24(3):121, 2024
Automated Fact-Checking of Climate Claims with Large Language Models
(with M. Leippold et. al.) *Nature npj Climate Action*, 2025

Research Activities

Presentation:

- 2025: Financial Regulation - Going Green in Halle, ZEW Conference on Ageing and Sustainable Finance, Wolfe Research 8th Annual European Quant and Macro Investment Conference, Moskowitz Prize Research Webinar at Kellogg, UZH Rising Scholar Conference in Finance.
- 2024: AFA Poster, Baruch Climate ESG Conference Poster, University of Zurich, Future of Financial Information Conference, SFI Research Days, CEPR-ESSEC-Luxembourg Conference on Sustainable Financial Intermediation, UZH Rising Scholar Conference on Sustainable Finance, EFA, Alpine Finance Summit, GRASFI, Stanford Institute for Theoretical Economics Conference: New Research in Asset Pricing, ESSEC Amundi Webinar: Harnessing Large Language Models in Asset Management, NFA, NTHU-UNSW Symposium on Sustainable Finance and Economics.
- 2023: University of Zurich, European Commission Summer School on Sustainable Finance, ETH AI+X Summit, ICGS Conference (×2), Düsseldorf CSR Conference (×2), TUM Sustainable Finance Workshop, Australasian Finance and Banking Conference.
- 2022: AFA Poster, Chinese Finance Annual Meeting.

Discussant: Düsseldorf CSR Conference 2023 (×2), TUM Sustainable Finance Workshop 2023.

Referee: Management Science, Journal of Banking and Finance, Journal of Corporate Finance.

Teaching Experience

University of Zurich *Switzerland*
Guest Lecturer *2025 Spring*

- Course: Sustainable Finance and Asset Pricing (PhD).

The Hong Kong University of Science and Technology *Hong Kong SAR*
Teaching Assistant *2019 - 2021*

- Courses: Calculus, Linear Algebra, Introduction to Actuarial Mathematics (Undergraduate).

Awards and Honors

AFA Travel Grant Award	*2024*
Postgraduate Studentship, HKUST	*2018 - 2022*
Din-Yu Hsieh Teaching Award, HKUST	*2019 - 2020*
Excellent Graduate of SUSTech	*2018*

Skills and Activities

Programming: Python, SAS, Stata, Matlab, MS Office.
Language: Chinese (native), English (fluent).

Reference

Zacharias Sautner
Professor
Department of Finance
University of Zurich
zacharias.sautner@df.uzh.ch

Markus Leippold
Professor
Department of Finance
University of Zurich
markus.leippold@df.uzh.ch

Kai Li
Professor
Department of Finance
Peking University
kaili@phbs.pku.edu.cn

Carbon in the Cloud*

Tingyu Yu†

August 2025

Abstract

This paper examines the climate impact of corporate artificial intelligence (AI) adoption in a large global sample. AI infrastructure developers bear most of AI's energy burden but partially offset it with renewable electricity. In contrast, user firms with AI talent experience plausibly causal reductions in subsequent Scope 1 emissions, as AI improves demand forecasting and enables shifts toward lower-carbon products and processes. Notably, these cuts are concentrated among low-emitting firms, while some carbon-intensive sectors experience emission increases due to capacity maximization. This underscores the need to consider not only AI's computing energy demands but also the broader operational contexts in which it is deployed. Climate-conscious investors appear to incorporate these climate risks into their capital allocation.

Keywords: AI, carbon emissions, path dependence, renewable electricity, climate risks.

*I am deeply grateful to Zacharious Sautner, Markus Leippold, and Kai Li for their continuous guidance and support. I also thank Francisco Amaral, Olimpia Carradori, Glen Gostlow, Felix von Meyerinck, Jingwei Ni, Stefano Ramelli (discussant), Chi-Yang Tsou, Alexander Wagner, and participants at the UZH Rising Scholar Conference in Finance 2025 for detailed discussions and helpful comments.

†University of Zurich, Department of Finance. Contact: tingyu.yu@df.uzh.ch.

1 Introduction

In PwC's 2024 Global CEO Survey, nearly half of the 4,701 executives identified artificial intelligence (AI) integration as a top strategic focus for the next three years. While AI rapidly becomes embedded in core business operations, its climate implications remain poorly understood. The International Energy Agency (IEA) warns of the steep carbon costs associated with widespread AI deployment, projecting that electricity demand from data centers could more than double by 2030. Despite this, only a few companies that have integrated AI report tracking these emissions, and many are beginning to reconsider their climate targets in light of AI's energy demands. This growing disconnect between digital innovation and environmental accountability underscores an urgent need for a deeper understanding of the real climate effect of AI adoption, which is critical for informing corporate sustainability strategies and also for ensuring alignment with global decarbonization goals.

Quantifying such climate impacts presents significant challenges. First, the effects are inherently ambiguous and multidimensional: while AI can drive up emissions through energy-intensive computation, data storage, and hardware manufacturing, it also has the potential to reduce emissions by improving operational efficiency, optimizing supply chains, and enabling smarter resource allocation. This dual nature complicates efforts to isolate net effects. Second, corporate disclosures related to AI-driven emissions remain sparse and opaque. Key details, such as model size, training duration, and infrastructure type, are often proprietary and not available from sustainability reporting. This lack of transparency hinders efforts to develop robust and comparable assessments of corporate AI adoption's carbon footprints.

To address the research gap, this project leverages a framework that distinguishes between *AI infrastructure developers* and *talent users*, with the help of large language models and inspired by the regulatory categories outlined in the EU Artificial Intelligence Act. AI developer firms are directly involved in building or operating AI infrastructure, such as data centers or proprietary computing systems, and thus are more likely to bear the energy costs

associated with large-scale computation. In contrast, AI user firms primarily integrate AI through workforce applications, often relying on external providers for computational needs, thereby shifting the energy burden off the balance sheet. This distinction allows for evaluating distributional climate impacts of AI adoption across the corporate value chain.

Within this framework, we observe clear asymmetries: **(i)** AI developers experience significant increases in electricity consumption due to their infrastructure involvement, though many partially offset this through greater reliance on renewable energy; **(ii**) AI users, by contrast, show no comparable rises in electricity use but often see reductions in Scope 1 carbon emissions, with a lead-lag dynamic model and instrumental variable estimates support a causal interpretation. The reductions arise as AI improves forecasting to limit resource waste and enables shifts toward lower-carbon products and processes. However, such cuts are currently concentrated among low-emitting firms, while some carbon-intensive sectors even see emission increases as AI is used to maximize capacity. This asymmetry underscores the need to account not only for the direct energy demands of AI computing but also for the broader operational contexts in which AI is deployed, which can lead to unintended and uneven emission outcomes.

Corporate AI adoption has evolved rapidly over the past decade. The proportion of firms identified as AI developers grows from 6% in 2010 to 12% after 2018, then stabilizes, before peaking at 17% in 2024. This peak is largely driven by the emergence of generative AI, which has significantly increased demand for computing power and infrastructure. The proportion of AI users rises steadily, from 7% in 2010 to approximately 28% by 2024. AI developer firms are concentrated in key few sectors. Firms in the Computer Software industry (e.g., Microsoft and other cloud computing service providers) most frequently mention AI infrastructure in their 10-K filings, followed by those in Electronic Equipment (e.g., Nvidia) and Trading (e.g., Equinix and other data center lessors). In contrast, user firms are more widely dispersed across industries. Sectors such as Business Services and Financials deploy AI to optimize both client interactions and internal operations. Notably, several carbon-intensive industries,

including Oil & Gas and Chemicals also report growing AI usage, mainly for internal process optimization rather than customer-facing applications.

AI developers and users differ markedly in their emissions profiles. Although both groups consume similar amounts of electricity, this usage does not translate into comparable Scope 2 emissions. The gap reflects the tendency of developer firms to locate infrastructure in regions with cleaner electricity grids and to employ contractual instruments (such as renewable energy certificates) to lower market-based emissions. In terms of Scope 1 carbon footprint, AI developers are primarily concentrated among low-emission firms, with user firms more frequently found in the middle- and high-emission categories. Additionally, AI workers in internal operations are evenly spread across emission quintiles, while those in market-facing roles are concentrated in lower-emission firms.

We examine potential determinants that influence firms' decision to invest in in-house AI infrastructure. Operational characteristics emerge as the primary drivers. Firms that are larger, have higher capital expenditures, possess richer data assets, and face greater AI-related competitive pressure are significantly more likely to internalize computing resources. While these firms are often headquartered in regions with cleaner electricity grids, the presence of approved emission reduction targets does not significantly affect infrastructure decisions. This suggests that climate considerations play a secondary role, with infrastructure investments more closely aligned with firms' operational needs and broader business strategies.

Our analysis obtains empirical evidence on the distributional climate impact of corporate AI adoption using a large and globally representative firm-level sample. We show that the energy burden of AI adoption is highly concentrated among firms with in-house infrastructure capabilities. Such firms consume significant 5% more electricity in the following year, increasing to 11% over a three-year horizon. In contrast, the share of AI workers shows no significant effect, with coefficients near zero. This divergence suggests that physical infrastructure, rather than talent acquisition, drives the observed changes in energy demand.

Renewable electricity plays a central role in the observed increase in energy use associated with AI infrastructure. Firms progressively use renewable electricity to "green" their reported emissions and decouple market-based Scope 2 emissions from actual electricity use. However, this increase does not fully offset the total rise in electricity usage. Many energy-hungry firms also significantly expand their non-renewable electricity consumption. The overall growth in power demand from AI infrastructure continues to exert upward pressure on emissions.

A supplier-customer-level analysis provides further support for the disconnect between AI integration and energy accountability. Firms with greater AI talent are significantly more likely to source from infrastructure-intensive suppliers. Specifically, a one standard deviation increase in a customer's AI workforce share is associated with a 5% increase in the supplier's electricity consumption in the following year. This suggests that pure AI users may outsource their computational load—particularly to cloud or colocation providers—effectively shifting the energy burden upstream in the supply chain. Conversely, firms with higher AI infrastructure disclosure are more likely to serve downstream customers with greater AI worker intensity. These customer firms do not exhibit higher electricity usage, reinforcing earlier findings that AI talent alone does not drive significant on-balance-sheet energy consumption.

On the other hand, AI user firms experience immediate reductions in their Scope 1 emissions after adoption. A one standard deviation increase in the share of AI-related workers is associated with an 8% reduction in Scope 1 emissions three years later. Decomposing this effect shows it stems mainly from lower carbon intensity rather than sales changes. We employ a distributed lead-lag model to analyze the dynamic effect of AI adoption on carbon emissions and find no evidence of pre-trends, suggesting the observed effects are unlikely to be driven by reverse causality. The results are consistent when controlling for non-AI workers in the IT, robotics, and data domains, as well as for pre-existing green innovation and emission-pledge initiatives. Finally, using lottery-driven excess H-1B visa allocations to instrument AI-worker supply reinforces the causal link between AI adoption and emission reductions.

Notably, the effects are quite different across sectors and are more concentrated on low-emitting firms. While many industries (like Finance) experience emission reduction, their low baseline emissions render these changes less environmentally consequential in absolute terms. In contrast, neither type of AI worker role has a significant association with emissions in the Oil sector; coefficients are small and not statistically distinct from zero. More strikingly, for Utility firms, internal operations-related AI adoption is associated with a statistically significant 2.4% increase in emissions per one standard deviation increase in related worker share. Given the high baseline emissions in this sector, this seemingly modest percentage change represents a material increase in absolute carbon output.

We identify two channels through which AI reduces emissions. First, improved forecasting of sales, demand, and inputs enables tighter production planning, which lower waste and Scope 1 emissions. Consistent with this, AI's effect is larger when forecast accuracy is higher. Second, AI facilitates low-carbon digitalization, shifting what and how firms produce and moving the product and process mix toward options with lower on-site carbon intensity. These channels translate poorly to carbon-intensive sectors, where waste-related emissions are a small share of the total emissions and core products and processes are intrinsically high-carbon and hard to substitute. Over our sample, we find no evidence that AI has lowered emissions in these sectors via green innovation or renewable expansion. Instead, AI is often deployed to maximize capacity utilization and raise near-term emissions. Unlike green technologies, AI is a general-purpose optimizer, so its emission impacts are path-dependent. Hence, firms should more explicitly embed climate objectives in AI strategy, accounting not only for the direct energy use of AI computing but also for its broader impact on carbon-intensive business activities.

Lastly, climate-conscious investors appear to have been cautiously attentive to the climate risks associated with certain forms of AI adoption. They appear more likely to divest from firms involved in AI infrastructure development that rely on non-renewable electricity, as well as from firms integrating AI talent while experiencing rising carbon emissions. Currently,

the overall investor response remains relatively modest. One possible explanation is that the climate implications of AI adoption are context-dependent, which makes it challenging even for responsible investors to incorporate AI-related climate risks into their investment strategies in a systematic way.

Literature Review: This project contributes to the growing literature on corporate exposure to AI and its implications for firm performance. Eisfeldt et al. (2023) develop the first firm-level measure of labor exposure to generative AI, building on the methodology of Eloundou and Rock (2023), and find firms with higher AI exposure earn higher returns following the release of ChatGPT. Using a similar measure, Bertomeu et al. (2023) show that highly exposed Italian firms experienced significant market value declines after Italy temporarily banned ChatGPT. Beyond market reactions, Jia et al. (2024) analyze earnings calls and find a sharp increase in managerial discussion of generative AI, primarily framed in positive tones and linked to action initiatives. Expanding the scope to general AI, Babina et al. (2024b) develop an AI investment measure based on employee resumes and find that AI-investing firms experience higher growth in sales and employment. Babina et al. (2024a) show these firms have also experienced increases in systematic risk over time. While these studies examine the labor and financial impact of AI, we extend the literature by shifting the lens to the dual climate impact of corporate AI adoption.

We add to the literature on the environmental footprint of artificial intelligence. Luccioni et al. (2022) provide the first lifecycle carbon analysis of a specific large language model, while Google's recent TPU study (Schneider et al., 2025) presents a cradle-to-grave assessment of AI hardware by incorporating embodied emissions from manufacturing. Bogmans et al. (2025) use a general equilibrium model to project how the expansion of data centers affects future electricity demand and prices. Kaack et al. (2022) propose a systematic framework for describing the effects of machine learning on greenhouse gas emissions. However, existing research largely relies on model-level estimates or qualitative assessment based on broad

assumptions, with limited empirical evidence at the firm level. This paper addresses the gap by leveraging historical electricity and emission data from AI-adopting firms to conduct the first large-sample, firm-level analysis of AI's real-world climate impact. By capturing heterogeneity in corporate responses and emission outcomes, we offer critical insights into how AI adoption has already shaped corporate climate trajectories.

2 Detecting Corporate AI Capabilities

To systematically assess the climate implications of corporate AI adoption, I develop a framework that distinguishes between two core dimensions of AI capability: *infrastructure deployment* and *workforce engagement*. These dimensions represent distinct channels through which AI can affect firm operations and, in turn, influence carbon emissions. Infrastructure deployment reflects a firm's role in building or operating the computational backbone required for AI development. Such activity is typically energy-intensive and directly contribute to electricity consumption and carbon emissions through computing and data storage. On the other hand, workforce engagement captures a firm's ability to integrate AI via skilled employees, often relying on cloud-based or third-party services. These firms adopt AI without maintaining in-house computational infrastructure and effectively outsource the energy consumption. By separating these dimensions, the framework differentiates *AI infrastructure developers* (firms with substantial infrastructure) from *AI talent users* (firms that adopt AI through labor while largely externalizing the associated energy demands), following the regulatory guidance of the EU Artificial Intelligence Act. This section details the construction of the measures to differentiate these AI firms and the underlying data sources.

2.1 Data

This project draws on data sources across three domains: (i) corporate AI adoption; (ii) climate and energy use; and (iii) financial fundamentals.

To measure firm-level AI adoption, I combine information from regulated corporate disclosures and workforce profiles. On the one hand, I identify investments in AI infrastructure by analyzing firms' annual 10-K filings submitted to the U.S. Securities and Exchange Commission (SEC), accessed via the EDGAR database. These filings are legally required to disclose material capital expenditures, including data centers and computing facilities, thereby offering information on AI-related physical infrastructure. On the other hand, I utilize individual-level employment data from LinkedIn, obtained through the Bright Initiative in October 2024. LinkedIn hosts more than one-billion-member profiles worldwide and covers an estimated 70% of the adult working population. The infrastructure and labor information together allow us to distinguish firms that merely pilot AI software from those that operate energy-intensive facilities.

The primary source of energy use data is the Carbon Disclosure Project (CDP), which collects voluntary climate disclosures from more than 23,000 companies (roughly two-thirds of global market capitalization). This includes firm-reported electricity consumption (disaggregated by renewable and non-renewable sources) and Scope 2 greenhouse gas emissions (both location-based, which reflects average emissions intensity of the grid where electricity is consumed, and market-based, which accounts for contractual instruments like renewable energy purchases). Comparing the two types of Scope 2 emissions can isolate the decarbonization efforts of firms more precisely. I rely on Scope 1 emissions from MSCI due to its broad global coverage and consistent estimation methodology.

I use financial data from Compustat and further incorporate supplier-customer-relationships from the Segment file. This dataset identifies customers that represent at least 10% of a supplier's revenue and allows the construction of primary supply chain networks. The baseline

analysis focuses on global public firms and the sample period spans 2010 to 2024 to reflect the expansion of AI over the past decade.

2.2 Measure Construction

2.2.1 AI Infrastructure Involvement

Quantifying corporate involvement in AI infrastructure is methodologically challenging, as such investments are rarely reported in a structured way. To address this, we extract information from unstructured textual disclosures in firms' annual 10-K filings, where companies are legally required to discuss material aspects of their operations and capital expenditures. These narratives often include descriptions of related infrastructure like data centers or GPU procurement, offering valuable insights that are not otherwise captured in quantitative datasets. To broaden the scope beyond U.S. firms, we extend this approach to include non-U.S. companies listed in the MSCI World Index that are not subject to SEC reporting requirements. For these firms, we manually collect the English version annual reports from their websites. Despite jurisdictional differences in reporting standards, these documents typically contain detailed disclosures on operational infrastructure that are comparable to 10-K filings. Including these global firms is essential, as many of them play a central role in the global diffusion of AI and thus offer a more complete view of cross-border patterns in infrastructure-intensive AI adoption.

Within the 10-K filings, we extract paragraph-level text from four key sections that are most likely to contain relevant discussions: Item 1 (Business), Item 2 (Properties), Item 7 (Management's Discussion and Analysis), and Item 8 (Financial Statements and Supplementary Data). To identify potential AI infrastructure discussions, we conduct an initial keyword-based search using terms designed to capture both on-premises computing resources and off-premises services. The full list of keywords is provided in the Internet Appendix B.

Considering keyword methods likely introduce false positives, we implement a second-

stage refinement using a large language model to assess the semantic context of each paragraph flagged in the first stage. This step aims to filter out references to general IT topics that lack clear relevance to AI infrastructure. We adopt a teacher-student learning approach that balance the high accuracy of state-of-the-art LLMs with the reproducibility of open-source models. In this setup, a high-capacity model (e.g., GPT-4.1 in this paper) serves as the teacher and generate good quality labels and rationales for a sample of candidate paragraphs. These annotations are then used to fine-tune a smaller open-source model (e.g., Meta-LLaMA-3.1-8B-Instruct) that acts as the student. The resultant student model is capable of approximating GPT-4's task-specific classification performance while, more importantly, remaining fully reproducible. Internet Appendix B provides a detailed description of the model fine-tuning. Based on its outputs, we classify whether each paragraph discusses AI infrastructure and construct a firm-year measure as the proportion of relevant paragraphs within the entire 10-K filing:

$$AIInfra_{i,t} = \frac{\sum_{p \in (10K_{i,t})} \#Words_p \times I(p = AI\ infrastructure)}{\sum_{p \in (10K_{i,t})} \#Words_p}.$$

2.2.2 AI Worker Engagement

The second dimension of corporate AI capability is based on human expertise. This component is particularly important in the post-generative AI era, where access to advanced models via API has lowered the need for proprietary infrastructure. We quantify this engagement based on the proportion of AI-related workers within a firm's workforce using data from LinkedIn profiles. To map LinkedIn users to public firms, we link profiles to companies from Compustat through a combination of fuzzy string matching and manual verification of company names. We restrict the sample to firms with English-language LinkedIn profiles and where at least 10% of the reported workforce (as per Compustat) can be matched to LinkedIn users (the median employee coverage is 72%). This results in a final sample of 13,617 global public firms, with 8,106 from Compustat North America and 5,511 from Compustat Global.

This dataset provides detailed individual-level information such as job title, position descriptions, start and end dates (at the month level), and, in many cases, information of project involvement, research publications, and patent contributions, which we use to identify whether a worker is AI-related or not. A worker is classified as AI-related in a given year if the LinkedIn profile satisfies at least one of the following criteria, following the method from Babina et al. (2024b): (1) their job title or description contains AI-related keywords; (2) they report involvement in AI-related projects in that year; (3) they have authored a research paper containing AI-related terms in that year; or (4) they are listed as an inventor on patents that reference AI-related terms within a one-year window (previous, current, or subsequent year). The keyword list is adapted from Babina et al. (2024b) and expanded to reflect recent AI development (e.g., generative AI) and software tools (listed in Internet Appendix B). This approach identifies 288,080 AI workers between 2010 to 2024, of whom 84% (or 241,186) are classified based on job title or description alone. To reduce measurement noise from sparsely recorded AI roles, we set the firm-level AI workforce to zero in years where fewer than five AI workers are identified. The value shows a high correlation (0.89) with the AI workforce proxy developed in Babina et al. (2024b), which is based on proprietary US resume data. The close alignment supports the validity of the constructed measure while with broader country coverage. The firm-year level AI worker engagement is quantified as:

$$AIWorker = \frac{\#Workers_{i,t}^{AI}}{\#Workers_{i,t}}.$$

Moreover, we further distinguish AI worker roles into two types: *market-facing* AI use and *internal-operation* AI use, based on which part of the firm's operations they affect. Market-facing AI refers to applications embedded in customer-facing processes, such as personalizing offerings, enhancing user experience, or sensing demand, which typically operate at the point of external interaction. In contrast, internal-operation AI refers to applications that support internal workflows, such as logistic optimization, energy management, or administrative automation, which interact more directly with Scope 1 and 2 emissions by targeting

energy-intensive processes and physical assets. This operational split provides a potentially useful lens to examine whether the climate impact of AI adoption varies with the part of the business it augments. The classification task is implemented in a similar approach used for AI infrastructure identification through a large language model. 31% of AI workers are engaged in market-facing applications and 38% in internal operations, while the remaining share does not fall clearly into either category, often due to insufficient textual information in the position descriptions. We correspondingly calculate shares of the two types of AI workers:

$$AIWorker^{Market} = \frac{\#Workers_{i,t}^{AI,\,Market}}{\#Workers_{i,t}},$$

$$AIWorker^{Internal} = \frac{\#Workers_{i,t}^{AI,\,Internal}}{\#Workers_{i,t}}.$$

2.2.3 Illustrative Examples

Table 1 presents illustrative examples of AI infrastructure disclosures and workforce profiles to validate the effectiveness of our textual methods and to demonstrate the different roles companies play in the AI ecosystem. Panel A displays example paragraphs from 10-K filings. For instance, Microsoft and Equinix report substantial investments in data center capacity and IT infrastructure, reflecting direct operational roles in delivering large-scale computing power for AI. Meanwhile, other firms contribute to critical physical components enabling this infrastructure: NVIDIA details its specialized GPUs for AI model training, whereas Optical Cable supplies high-speed connectivity products essential for data center performance.

Panels B and C shift focus to the deployment of AI talent during the use phase and showcase how firms integrate AI capabilities across diverse business functions, not solely within IT departments. Panel B features market-facing applications, such as audience behavior prediction at Time, customer engagement optimization at JPMorgan, and personalized marketing at Warner Music. Panel C emphasizes internal operational uses, where KLA leverages AI for manufacturing diagnostics, Schlumberger for drilling optimization, and UPL for automating

plant maintenance.

Overall, these cases validate our method as a structured lens for understanding corporate AI engagement across its life cycle, from infrastructure development to talent application.

3 Distribution of Energy Burden from AI Adoption

This section examines the energy burden of corporate AI adoption by comparing the roles of infrastructure and talent. Specifically, we assess the magnitude of electricity consumption, its reflection in firms' emission reports, and how the energy burden is distributed across different types of AI adopters.

3.1 Comparing AI Developers and Users in a Climate Context

3.1.1 Diffusion Over Time

Figure 1 traces the evolution of corporate AI adoption from 2010 to 2024. A key trend is the steady and accelerating rise of AI user firms—companies that integrate AI talent without owning substantial infrastructure. This upward trajectory (from approximately 7% to 29%) aligns with three major technological milestones: the 2012 deep learning breakthrough with AlexNet, the 2017 introduction of the transformer architecture, and the 2022 release of ChatGPT. Each progressively expanded the scope and accessibility of AI applications, moving from early successes in visual recognition to advanced language understanding, and ultimately to general-purpose tools. Internal operational applications consistently outpace market-facing uses, indicating a broader integration of AI into the organization beyond consumer interactions.

In stark contrast, the proportion of AI infrastructure developers, those reporting material infrastructure involvements, remains relatively flat over most of the period. This di-

vergence highlights a fundamental asymmetry in AI adoption: while a growing number of firms embrace AI through cloud services, only a smaller group commits to the capital- and energy-intensive task of building and operating the computational backbone. Nevertheless, AI developers increase from 13% to 18% in 2024, triggered by the adoption of ChatGPT. This raises immediate questions about the future trajectory of energy consumption tied to infrastructure expansion.

3.1.2 Emission Profiles

The distinction between AI developers and users extends to their emission profiles. Figure IA1 compares their average electricity consumption and related Scope 2 emissions, based on a balanced sample of firms reporting all four relevant metrics. While AI-related computing likely constitutes a large share of total electricity use for developers, these firms do not consume more electricity overall. On average, they report an annual electricity consumption of 1.68 million MWh, compared to 2.01 million MWh for AI users. This difference reflects the larger industrial footprint of many user firms, such as those in oil and gas or heavy manufacturing, where electricity demand is driven primarily by core operational processes rather than AI activity. Crucially, AI developers demonstrate a more proactive approach to decarbonization. They rely less on non-renewable electricity and report much lower Scope 2 emissions under both location-based and market-based accounting. This suggests their infrastructure is often sited in regions with cleaner electricity grids, and moreover, they actively procure renewable energy through contractual instruments to further reduce reported emissions.

The emission gap becomes even more pronounced when examining Scope 1 carbon emissions, as shown in Figure 2 by sorting firms into quintiles each year. Panel A presents the proportion of AI firms within each emission group. While AI developers are typically concentrated among the lowest-emitting firms, AI users are disproportionately found in medium and high emission categories. Specifically, more than 20% of the high- and medium-emitting firm-

year observations are classified as AI users, compared to just 7% among the lowest-emitting firms. Panel B analyzes the average share of AI workers across use domains, excluding AI developers from the sample. The share of AI talent focused on internal operations is distributed evenly across emission quintiles, ranging from 0.7 to 1.0 basis points. In contrast, talent dedicated to market-facing applications is more concentrated in lower-emission firms.

Figure IA2 reinforces these findings by examining industry-level patterns based on the Fama-French 12 industry classification. Panel A reports the number of firm-year observations across industries. A clear concentration is evident among AI developers, who are clustered in a small number of low emission tech sectors. The most prominent one is the Computer, Software, and Electronic Equipment sector, which alone accounts for more than 57% of all AI providers. This includes cloud service developers like Microsoft and firms that manufacture specialized AI hardware like NVIDIA. The second top sector is Finance, which includes some data center operators like Equinix. Conversely, AI users are distributed much more broadly across the industrial spectrum. Sectors like Manufacturing, Business Services, and Finance deploy AI to optimize both client interactions and internal workflows. Meanwhile, several carbon-intensive industries, including Oil & Gas and Chemicals, demonstrate AI usage for internal process optimization, consistent with the result based on emission-sorted groups.

3.1.3 Determinants of AI Adoption

Given the central role of AI infrastructure in shaping the climate impact of AI adoption, it is important to examine the factors that influence firms' decisions to invest in such infrastructure. Table IA2 presents results from regressions where the dependent variable is a dummy equal to one if a firm reports substantial AI infrastructure in its annual report. We define it as firms having an above-median proportion of infrastructure-related discussion to avoid capturing incidental mentions. The analysis considers both operational characteristics and climate-related factors that may drive adoption. Control variables include the logarithm of total assets, ROA, investment rate, leverage, tangibility, and sales growth, along with

industry-year and country fixed effects.

Firm-level operational characteristics appear to be the main drivers. Column 1 shows a strong association between infrastructure involvement and the presence of AI-skilled employees, particularly those in internal operations roles, which likely require more frequent or intensive model training. In contrast, market-facing AI roles show no significant association, suggesting these applications are less compute-intensive and can be more easily outsourced to cloud providers. Column 3 shows that firms with richer data assets, based on their 10-K filings, are also more inclined to invest, with a one standard deviation increase in such disclosure associated with a 3% higher likelihood of adoption. Column 4 indicates that firms under greater pressure from AI-related competition more often adopt infrastructure, likely as a response to maintain a technological edge. Interestingly, AI infrastructure developers are often located in states with higher electricity prices. This implies that compute-related energy cost represents a relatively minor factor in adoption decisions and that such firms are largely inelastic to energy prices. Lastly, adopting firms tend to be headquartered in states with higher shares of renewable energy generation. While this likely reflects broader regional energy trends rather than deliberate firm choices, operating in such environments may shape managerial perceptions of climate responsibility.

As for firm characteristics, adopters of AI infrastructures tend to be larger in size, exhibit higher capital investment, operate with lower property tangibility, and experience stronger sales growth.

3.2 Electricity Consumption

To analyze relationship with electricity usage, we estimate the following regression:

$$\Delta Electricity_{i,(t+k,t)} = \alpha + \beta_1 \mathbb{1}(AIInfras_{i,t}) + \beta_2 \, AIWorkers + \mathbf{X}_{i,t} + b_{j,t} + c_c + \epsilon_{i,t}.$$

The dependent variable captures changes in electricity consumption from year t to $t+k$, measured over one-, two-, and three-year horizons. The main explanatory variables are an indicator for AI infrastructure involvement and the share of AI-skilled workers. $\mathbb{1}(AIInfras)$ equals one if a firm has above-median AI infrastructure disclosures, and zero otherwise. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year electricity usage. The specification includes industry-year fixed effects to capture sector-specific AI trends and macroeconomic shocks, and country fixed effects to control for national energy policies and decarbonization efforts. Standard errors are clustered at the firm level.

This analysis relies on electricity consumption data from CDP and thus limit the sample size. Nevertheless, it includes key players in the AI ecosystem, such as major technology firms and semiconductor providers (see Table IA1), making it particularly informative for assessing the energy demands associated with AI investment. We further restrict the sample to firms reporting emissions and energy use under the operational control boundary. This boundary choice is critical for interpreting the data, especially in the context of AI deployment. For instance, electricity consumption from leased data centers, a common setup among AI adopters, is reported by the lessee only if the facility falls under their operational control. In contrast, under boundaries like financial control, such usage may go unreported unless the firm holds majority ownership, which is uncommon. Focusing on operational control thus provides a more accurate reflection of a firm's own energy footprint. In our sample, approximately 79% of firms adopt this reporting standard.

Table 3 presents the estimates. Columns 1–3 shows that firms with AI infrastructure involvement consume significant 5% more electricity in the following year, increasing to 11% over a three-year horizon. Similar patterns hold when using electricity intensity as the dependent variable in Columns 4–6. In contrast, the share of AI workers shows no significant effect, with coefficients near zero. This divergence suggests that physical infrastructure, rather than talent acquisition, drives the observed changes in energy demand.

3.3 Role of Renewable Electricity

We next examine how increased electricity usage associated with AI infrastructure adoption translates into firms' reported carbon emissions by replacing the dependent variable with Scope 2 emissions, using both location- and market-based accounting methods. As shown in Panel B of Table 3, location-based emissions, which reflect actual grid-average emission factors, rise significantly, mirroring the increase in electricity consumption. This suggests that the energy demands of AI infrastructure have a direct and measurable climate footprint. Market-based Scope 2 emissions, which incorporate firms' contractual purchases of renewable energy, show a positive but less significant change. Firms are actively using renewable energy certificates or power purchase agreements to "green" their reported emissions, partially decoupling their market-based carbon footprint from actual energy use.

To understand the composition of this shift, we decompose total electricity consumption into renewable and non-renewable sources. We find a substantial rise in renewable electricity use, from 44% to 57% over three years. This reflects strategic investments in low-carbon energy sourcing among firms most exposed to AI-related energy demands. However, the increase in renewables does not fully offset the total rise in electricity usage. Non-renewable consumption also rises, particularly over longer horizons, indicating that AI-driven power demand continues to exert upward pressure on emissions.

3.4 Supply Chain Spillovers

To assess whether firms with AI talent but no infrastructure indirectly bear the energy costs associated with AI use, we analyze their supply chain relationships in Table 4. Specifically, we test whether firms with a higher share of AI workers are more likely to source from suppliers that operate AI infrastructure. Using a supplier–customer-level regression with firm-level controls and customer–supplier pair fixed effects, we find that firms with greater AI talent are significantly more likely to be connected to infrastructure-intensive suppliers.

Furthermore, a one standard deviation increase in a customer's AI worker share is associated with a 5% increase in the supplier's electricity consumption in the following year. This suggests that pure AI users may outsource their computational load to cloud or colocation providers, to effectively shift the energy burden upstream in the supply chain. Conversely, firms with greater AI infrastructure disclosure are more likely to serve downstream customers with higher AI worker intensity. These customer firms do not exhibit higher electricity usage, consistent with earlier findings that AI talent alone does not drive on-balance-sheet energy consumption. In fact, we observe a marginally significant decrease in their Scope 1 carbon emissions, which we explore further in the next section.

4 AI Integration and Scope 1 Carbon Emissions

What are the direct climate outcomes of incorporating AI into business operations? On the one hand, AI has the potential to improve operational efficiency, which could lead to a reduction in emissions. On the other hand, unlike green technologies, AI is not inherently designed for climate mitigation. Thus, any efficiency gains may also lower production costs, potentially encouraging expanded operations and ultimately increasing overall emissions, a dynamic known as the Jevons paradox. Additionally, AI adoption may enable firms to exploit previously inaccessible resources or optimize the use of carbon-intensive inputs, further elevating carbon intensity. As a result, the unintended climate impact of AI integration remains ambiguous and needs to be addressed empirically.

4.1 Baseline Results

This subsection explores the relationship between AI integration and Scope 1 carbon emissions and expands the sample to include all firms with AI worker information, without the constraint of electricity or 10-Ks data. This enables us to analyze a more comprehensive set

of firms worldwide. To better isolate AI user firms, we exclude firms from the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code), which ranks as the top sector for AI infrastructure involvement. We estimate the following regression:

$$ClimateEmission_{i,t+k}^{Scope1} = \alpha + \beta\, AIWorkers_{i,t} + \mathbf{X}_{i,t} + a_i + b_{j,t} + c_{c,t} + \epsilon_{i,t}, \quad k \in \{1, 2, 3\},$$

where the dependent variable is Scope 1 carbon emissions in the future three years and the share of AI workers serves as the key independent variable. The coefficient β captures the relationship between AI workforce adoption and direct future emissions. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The firm fixed effects absorb all time-invariant characteristics at the firm level, ensuring that the coefficient is estimated solely from within-firm time-series variation. Industry-year and country-year fixed effects account for sectoral and regional differences in regulation and energy markets. Standard errors are clustered at the firm level.

Table 5 reports the results. A one standard deviation increase in the share of AI-related workers is associated with a 3% reduction in Scope 1 emissions in the following year. This effect intensifies over time, reaching an 8% reduction by the third year. Both market-facing and internal AI roles contribute to emission reductions, with slightly stronger effects from the former roles. This may reflect the fact that firms adopting market-facing AI often have lower baseline emissions, so even modest absolute reductions yield larger percentage declines. Table IA3 explores the heterogeneous effect by interacting AI worker shares with some indicators. The emission reduction effect is stronger after 2018, following the introduction of transformer models that is a major advancement in machine learning. Geographically, the effect is more pronounced among U.S. firms, though results remain statistically significant for firms in other regions. Table IA4 shows that the reduction extends to Scope 2 emissions, while the effect on

Scope 3 emissions is less clear. Since our primary interest lies in emissions reductions driven by AI integration into core business operations, the remaining analysis focuses on Scope 1 emissions.

As a second step, we decompose total Scope 1 emissions into two components, i.e., emission intensity and sales, and use them as dependent variables in Table 6. Panel A shows that the reduction in emissions is primarily driven by improvements in carbon intensity, with both the magnitude and statistical significance closely mirroring the results in Table 5. In contrast, there is no consistent pattern observed for future sales. This suggests that the observed decline in emissions is largely due to efficiency gains in carbon intensity, indicating that AI adoption enhances operational efficiency with climate benefits.

Dynamic Effects: We extend our analysis by estimating the dynamic effects of AI adoption on Scope 1 emissions. This approach helps address concerns about reverse causality and reveals the temporal pattern of AI's impact. Specifically, we employ a distributed lead-lag model that accommodates continuous variation in AI adoption as firms typically adopt AI incrementally rather than through one-time investments, which makes standard event-study designs unsuitable. The model is specified as:

$$ClimateEmission_{i,t}^{Scope1} = \alpha + \sum_{k=-3}^{3} \beta_k \, \Delta AIWorkers_{i,t-k-1,t-k} + \mathbf{X}_{i,t} + a_i + b_{j,t} + c_{c,t} + \epsilon_{i,t},$$

where $\Delta AIWorkers_{i,t-k-1,t-k}$ is the annual change in the share of AI workers from year t-k-1 to year t-k. The dependent variables are the logarithm of carbon emissions, intensity, or sales. Control variables and fixed effects follow the baseline specification. Each β_k captures the cumulative response of the outcome in year t to changes in AI adoption k years earlier, holding constant AI adoption in all other years. The lead coefficients serve as a pre-trend test. If firms adopting AI follow emission trends similar to those of others before adoption, these coefficients should not be statistically different from zero.

Figure 3 presents the coefficients from the lead-lag regressions. Panel A shows that Scope 1 emissions decline following AI adoption, with a one standard deviation increase in AI workforce growth associated with a 3% reduction in emissions, intensifying over the next two to three years. Panel B confirms that these reductions are primarily driven by improvements in carbon intensity. Importantly, there is no evidence of pre-trends. Firms that later adopt AI exhibit similar emission levels to others in the years prior to adoption, with significant divergence occurring only afterward. This supports that the observed effects are unlikely to be driven by reverse causality or omitted variables, and strengthens the causal interpretation of our findings.

Robustness Checks: Table IA5 demonstrates the robustness of our findings with respect to various confounding factors and alternative empirical specifications.

First, the estimated effects of AI adoption could be confounded by contemporaneous changes at adopting firms that contribute to lower emissions. To address this, we examine several potential confounding factors and provide evidence that our results are not driven by them. Panel A tests whether AI adoption is correlated with the adoption of other technologies by controlling for the share of IT workers, robotic workers, and non-AI data workers. The coefficient on AI worker share remains robust across these specifications, indicating that the observed effects are specific to AI technology. Notably, IT and data workers are only associated with lower emissions in the following one year but not further, while robotic workers are even positively associated with emissions. Panel B addresses the possibility that firms adopting AI are also more engaged in green initiatives, which could independently drive emissions reductions. We thus control for green patent share, a dummy for firms with emissions targets aligned with the Science-Based Targets initiative (SBTi) or net-zero goals, and a text-based measure of firm-level climate exposure from Sautner et al. (2023). The results remain consistent, suggesting that the effects are not simply due to pre-existing climate commitments; as expected, firms with ambitious climate targets also show declining emissions. Panel C

further controls for external factors such as local renewable energy share, electricity prices, and financial constraints. These adjustments obtain qualitatively and quantitatively similar results.

Second, we explore the sensitivity of the results to the choice of empirical specifications. Panel D tests alternative measures of AI adoption, including a binary indicator for the presence of AI workers, the logarithm of the total number of AI workers, and AI worker share scaled by sales. Panel E alternatively uses industry-year and country fixed effects, clusters standard errors by both firm and year, and clusters by industry. Across all specifications, the results remain consistent in both magnitude and significance.

Instrument Variable: To further address potential bias from unobserved shocks and strengthen causal inference, we employ an instrument variable approach. Specifically, we use excess H-1B visa selections as an instrument for firm-year changes in the supply of AI workers (Cen et al., 2024; Chen et al., 2021; Dimmock et al., 2022). This choice rests on two facts: (i) a large share of U.S. AI talent is foreign-born and typically requires an H-1B visa to work there legally; and (ii) H-1B demand routinely exceeds the annual cap, prompting the U.S. government to allocate visas by randomized lottery before reviewing applications. The lottery generates quasi-random variation in firms' ability to hire foreign AI workers, producing plausibly exogenous shocks to AI labor supply.

Internet Appendix C provides detailed information on the institutional background of the visa lottery and the measure construction. Full visa lotteries occurred in each fiscal year after 2014, while 2010 to 2013 saw lotteries only for a subset of filings (Chen et al., 2021; Dimmock et al., 2022); hence, the instrument variable estimates focus on U.S. firms from 2014 to 2023. Moreover, given the excess visas are more directly impact the flow of newly hired AI workers each year, rather than the total AI stock, we estimate a first-difference specification accordingly. We use changes in carbon emissions as the dependent variable and changes in AI worker share as the independent variable, while controlling for industry-year

fixed effects. This specification is conceptually aligned with our baseline model.

For firm i in year t, our instrument is the excess number of visa selected relative to expected selections given the aggregate lottery selection ratio, following Cen et al. (2024):

$$ExcessVisa_{i,t} = VisaSelection_{i,t} - VisaDemand_{i,t} \times SelectRate_t,$$

where $VisaSelection_{i,t}$ is the number of firm i's H1B visa registrations drawn in the lottery, $VisaDemand_{i,t}$ is the total registrations submitted, and $SelectRate_t$ is the fiscal-year lottery selection rate, computed as total selections divided by total registrations. To further target foreign AI talent, we weight the instrument by the firm-year share of STEM-related visa applications based on their SOC codes. This refinement concentrates variation on positions most relevant to AI adoption:

$$AIExcessVisa_{i,t} = ExcessVisa_{i,t} \times STEMShareVisa_{i,t}.$$

To illustrate, consider two firms that each submit ten H-1B registrations in a year with a 50% average selection rate. One receives eight selections, with the other none. Their $ExcessVisa$ values would be +3 and -5, respectively, and reflect deviations from the expected five selections. The difference is driven solely by the lottery and represents exogenous shocks to each firm's labor supply. A firm with -5 may face delays in AI initiatives due to a shortage of skilled workers, while a firm with +3 can proceed as planned.

Panel A of Table 7 presents the results. Column 1 confirms that the instrument strongly predicts changes in AI worker share, with a statistically significant positive coefficient and a first-stage F-statistic above the threshold of 10. Columns 2–4 report the second-stage (2SLS) estimates, showing that increased AI adoption significantly reduces carbon emissions. A one standard deviation increase in AI worker share leads to a 6% reduction in Scope 1 emissions in the following year, expanding to an 8% reduction after three years. Panel B uses the AI-

related excess visa selection as an instrument. Results are stronger under this more targeted specification, with significant emission reduction of 7% to 10% over the subsequent three years, reinforcing the causal link between AI adoption and emissions reduction.

For robustness, we also instrument firm-level changes in AI worker shares based on firms' ex-ante exposure to universities historically strong in AI research, using the measures provided by Babina et al. (2024b). The idea is that AI-strong universities produced more AI-skilled graduates in recent years, making it easier for firms that typically recruit from them to attract AI talent. This instrument is calculated as $\sum_u s_{iu}^{2010} AIstrong_u$, where s_{iu}^{2010} is is the share of a firm i's STEM workers in 2010 who graduated from university u, and $AIstrong_u$ is a dummy equal to one if u is in the top 5% by number of AI researchers during 2005–2009. Table IAIA6 presents the estimates for the U.S. sample from 2010 to 2023. Consistent with Babina et al. (2024b), this instrument has a strong first stage and show a robust reduction on carbon emissions over the subsequent two years; the third-year coefficient is negative but insignificant. A one standard deviation increase in AI worker shares leads to a 9% decline in future Scope 1 emissions.

4.2 Sector Heterogeneity

While the firm-level analysis establishes a robust average relationship between AI adoption and subsequent reductions in Scope 1 emissions, this aggregate effect may obscure substantial variation across sectors. Two core economic considerations make it essential to extend the analysis to the industry level. First, technological path dependence (Aghion et al., 2016; Bolton et al., 2024) implies that the climate impact of AI, as a general-purpose tool, depends on the capital stock and production processes it complements. In service-oriented industries, AI mainly augments digital workflows, and efficiency gains are likely to dominate. This is different for carbon-intensive sectors, where the algorithms are embedded in heat- and power-hungry physical operations, so their net effect depend critically on whether they substitute

for or accelerate fossil-fuel inputs. Second, a given percentage reduction in emissions carries vastly different implications depending on a sector's baseline emissions (Hartzmark and Shue, 2023). For instance, a two-percent drop in Scope 1 emissions in a high-emitting utility firm translates into a far larger absolute reduction in CO_2 than the same percentage decrease in a low-emitting service firm. It is these absolute reductions that matter most for decarbonization goals. Therefore, we must disaggregate the analysis to the industry level to identify where AI contributes materially to emissions mitigation, and where its impact may be negligible or even counterproductive.

Figure 4 explores sectoral heterogeneity by estimating regressions at the industry level, using both market-facing and internal operations-related AI worker shares as independent variables. Sectors are ranked by their average Scope 1 carbon emissions to highlight differences in baseline carbon intensity. The figure plots the estimated coefficients for each AI role (market-facing in purple, internal operations in blue), along with 90% confidence intervals.

Interestingly, the sectors with the lowest baseline emissions are often active adopters of AI, with nearly all of them showing at least one statistically significant negative coefficient for future emissions. The Finance sector stands out as the most intensive AI adopter among user industries (see Figure IA2). This sector shows marginally significant reductions in emissions for both market-facing and internal AI roles. A one standard deviation increase in market-facing AI worker share corresponds to a 1.9% reduction in Scope 1 emissions in the following year. However, due to the sector's low baseline emissions, this only translates into a tiny absolute reduction of CO_2. Similarly, the Shops and Retail sector shows a high level of AI talent hiring, with internal-operational AI roles leading to a significant emission reduction. Yet again, the sector's low baseline emissions mean that these reductions, while statistically significant, are modest in absolute environmental terms.

The emissions impact of AI adoption becomes notably more varied among carbon-intensive industries, where climate consequences are more material. In the Chemical sector, internal operations-related AI roles exhibit a statistically significant negative coefficient. This sug-

gests that in some heavy industries, AI are effectively used to lower emissions. In contrast, the Oil sector shows no meaningful relationship between AI adoption and emissions; coefficients are small and not statistically distinct from zero, indicating a lack of measurable climate benefit from AI integration in this industry. More strikingly, in the Utility sector, which ranks highest in average Scope 1 emissions and exceeds those of the Finance sector by over a thousand-fold, AI adoption related to internal operations is associated with a statistically significant 2.4% increase in emissions per one standard deviation increase in AI worker share. While this percentage may appear modest, the extremely high baseline emissions in this sector make the absolute increase in CO_2 output substantial. Overall, the effectiveness of AI adoption in mitigating emissions is highly context-dependent and may be even counterproductive in sectors with inherently high carbon intensity.

Narrative Retrieval: One key advantage of our AI adoption measure is that it is built from job descriptions, which carry rich narratives about how firms actually integrate AI into their operations. Mining such textual data allows us to uncover the economic mechanisms behind the observed heterogeneity in emissions effects across industries.

For each sector, we first locate the firms that most influence the regression estimates. Influence is measured with DFBETA, which quantifies how much observations from a single firm shift a coefficient when they are removed; large DFBETA values flag firms that drive the sector-level result. We select the 50 firms with the highest DFBETA in each sector and extract the full text of their AI job descriptions. We then create word clouds that visualize the most frequent terms in those descriptions. Figure 5 displays word clouds for the six sectors with statistically significant coefficients; most of them are negative, except for Utilities (circled in red). Figure IA3 shows the remaining six sectors with non-significant effects. Such narrative retrieval reveals the types of AI tasks most closely linked to emission outcomes and offers qualitative insight into why AI mitigates emissions in some industries but not others. The next subsection discusses these sector-specific mechanisms in detail.

4.3 Economic Mechanisms

We document that AI-adopting firms can significantly reduce emissions through efficiency gains from (i) improved forecasting and (ii) low-carbon digitalization of products and processes, while we find no evidence of improvement through climate-related output like green innovation. By contrast, the most carbon-intensive sectors sometimes deploy AI for capacity maximization that leads to future emission increases.

4.3.1 AI to Improve Forecasting

The first channel through which AI can reduce emissions is by improving forecast accuracy and optimizing resource allocation (Farboodi and Veldkamp, 2022; Mihet and Philippon, 2019; Farboodi and Veldkamp, 2023; Babina et al., 2024b). As a prediction technology, AI excels at processing large volumes of high-dimensional data to detect complex patterns, which outperform traditional statistical methods in predictive performance. This capability is effective in reducing emissions because nearly all operational decisions, such as how much to produce, when to maintain equipment, and how to manage supply chains, rely on forecasts. When forecasts are inaccurate, either overly optimistic or pessimistic, they can lead to misallocation of labor, capital, and inputs, ultimately increasing operational waste (Farboodi and Veldkamp, 2023; Tanaka et al., 2020). Thus, by shrinking the gap between expected and actual outcome, AI adoption lower the uncertainty of decision-making and such efficiency gains can directly translate into lower operational emissions.

This mechanism is supported by the narrative evidence in the word clouds shown in Figure 5. For example, Panel A indicates that the Shops and Retail sector employs AI workers for predictive analytics in areas such as customer service and fraud prevention. By analyzing historical transaction data, AI can identify customer preferences and enable firms to tailor products more effectively and reduce waste from unsold or misallocated inventory. Panel B shows that the Finance sector makes extensive use of AI for asset management, data-

driven decision-making, credit risk prediction, and broader risk management. Interestingly, similar predictive applications appear in more carbon-intensive industries, such as Chemicals in Panel C, where AI is applied to demand planning, supply chain optimization, inventory management, and even talent acquisition, functions that can reduce overproduction, idle capacity, and associated emissions.

Precisely quantifying the prediction accuracy of AI in different domains is challenging due to data limitations. Here, we focus on one measurable and economically important form of prediction—demand forecasting (Corhay et al., 2025). More accurate forecasts of sales, demand, and input needs enable tighter production schedules and more efficient inventory management, reducing overproduction and other waste that directly leads to on-site fuel use. We measure firm-year forecast accuracy as the absolute analyst sales forecast error divided by actual sales per share, following the method of Feng et al. (2009):

$$ForecastAccuracy_{i,t} = -\frac{|ForecastSales_{i,t-1} - ActualSales_{i,t}|}{SalesPerShare_{i,t-1}}.$$

We use the absolute forecast error to capture the magnitude rather than the direction of the error. By multiplying this value by -1, higher values indicate greater forecast accuracy (i.e., lower errors). While we do not have access to the management guidance data, the analyst forecast can serve as a reasonable proxy, as analysts often quickly incorporate management guidance into their predictions (Cotter et al., 2006). We focus on sales rather than earnings because sales more directly reflect consumer demand. We then estimate the following regression:

$$\begin{aligned} ClimateEmission_{i,t+1}^{Scope1} = \alpha + \beta_1\, AIWorkers_{i,t} \times ForecastAccuracy_{i,t} + \beta_2\, AIWorkers_{i,t} \\ + \beta_3\, ForecastAccuracy_{i,t} + \mathbf{X}_{i,t} + a_i + b_{j,t} + c_{c,t} + \epsilon_{i,t}. \end{aligned}$$

Our coefficient of interest is β_1, which captures how AI adoption affects emissions condi-

tional on improved demand forecasting. In Table 8, Panel A, the interaction term is negative and statistically significant, indicating that AI's emissions-reducing effect is larger when forecasts are more accurate. The results are robust when forecast accuracy is constructed using either the mean or the median analyst forecast. By role, the effect is stronger for market-facing AI workers, while the estimate for internal-operations AI workers is also negative but marginally insignificant, suggesting both groups contribute to better predictive analytics.

4.3.2 AI to Enable Low-Carbon Digitalization

Beyond tightening plans, AI can reconfigure what and how firms produce. Machine learning-based design and process control (e.g., digital twins, set-point optimization, computer vision) enable products that use less material and processes that substitute lower-emission inputs. These shifts move the firm's product/process mix toward options with lower on-site carbon intensity, delivering emissions reductions beyond forecasting-driven efficiency. To test this mechanism, we assemble firm-year sales and capex by segment from Compustat and map each segment to an EPA-based carbon intensity. We then aggregate to a firm-level weighted carbon intensity using (i) sales weights (realized mix) and (ii) capex weights (forward-looking shift). Our low-carbon digitalization measure is

$$LowCarbonDigit_{i,t} = -\left(\frac{\sum_{s\in Segment} Sale^{s}_{i,t} \times CE^{s}}{Sale_{i,t}} - \frac{\sum_{s\in Segment} Sale^{s}_{i,t-1} \times CE^{s}}{Sale_{i,t-1}}\right).$$

So higher values indicates a move to cleaner segments. In Panel B, we add this measure and its interaction with AI worker share to the baseline regression. We observe a negative and significant coefficient on the interaction term, suggesting that AI's emission impact is stronger when the firm's product/process mix becomes cleaner. Columns 3–4 show that this effect mainly comes from the internal operation roles.

4.3.3 Other Mechanisms

To understand why carbon-intensive firms are less effective at AI-driven emission reductions, we examine mechanisms more salient in these sectors. First, we test whether AI adoption is associated with greater climate-related outputs, including more green patenting, lower methane leakage, and higher renewable generation. Panel A of Table 9 augments the baseline with interactions between AI worker share and the log of each outcome. Despite frequent claims that AI should boost green innovation, curb methane releases, or accelerate renewables, we do not observe such effects in the data over our sample period. Carbon-intensive firms may not realize AI-related emission reductions via climate-output pathways currently.

Second, unlike purpose-built "green" technologies, general-purpose AI can be deployed to maximize throughput, such as optimizing drilling schedules, routing, and plant utilization. In carbon-intensive industries, such usage can raise fuel burn and process emissions, offsetting efficiency gains and producing the emission increases we observe. We proxy capacity usage with year-over-year changes in three ratios: (i) oil capacity use = oil production / proved reserves; (ii) gas capacity use = natural-gas production / proved reserves; and (iii) electricity capacity factor = total generation / installed capacity. We add each change and its interaction with AI worker share to the baseline (Panel B). The AI and capacity-use interactions for oil and gas are positive and significant, indicating that AI adoption helps these firms push capacity harder, leading to higher emissions. The effect is concentrated among internal-operations AI roles.

Overall, AI reduces emissions mainly through forecasting-driven efficiency and AI-enabled low-carbon digitalization, but these gains are concentrated among lower-emitting firms. In carbon-intensive sectors, the channels translate poorly because waste-related emissions are a small share of their total emissions and their core products/processes are intrinsically high-carbon and hard to substitute. We also find no evidence that AI lowers emissions in these sectors via green innovation or renewable expansion over our sample period. Instead, AI is

often deployed to maximize capacity utilization, which can raise emissions.

5 Perception of Climate Conscious Investors

A critical question is whether investors recognize and respond to the climate implications of corporate AI adoption. Based on the above analysis, AI adoption has at least two key dimensions of climate impacts. First, the infrastructure supporting AI can significantly increase energy consumption, although the associated carbon footprint can be mitigated through the use of renewable electricity. Second, user firms that integrate AI by hiring related talent experience heterogeneous emission outcomes, which tend to be negative overall but with certain carbon-intensive sectors seeing a notable emission rise. In this section, we investigate whether investors, particularly those who are climate-conscious, account for the two types of AI-related climate risks in their portfolio decisions.

Table 10 uses institutional ownership as the dependent variable and distinguishes between investors that are signatories to the Principles for Responsible Investment (PRI) and those that are not. PRI signatories serve as a clear proxy for climate-conscious investors, as they commit to integrating environmental, social, and governance factors, particularly the assessment and mitigation of climate-related risks, into their investment decision-making (Dyck et al., 2019; Gibson Brandon et al., 2022). Panel A examines their response to AI infrastructure involvement. The independent variables include a dummy indicating whether a firm is involved in AI infrastructure, the firm's renewable electricity consumption, and the interaction between the two. We use the same fixed effects and control variables as in the baseline regression. Columns 1–3 report the results for PRI institutional investors. The coefficients for AI infrastructure involvement are negative but only become statistically significant after three years. Specifically, such involvement is associated with a 2.45% decline in PRI institutional ownership. Notably, the interaction terms between AI infrastructure involvement and renewable electricity consumption are positive and statistically significant

throughout the subsequent three years. This suggests that climate-conscious investors are less likely to divest from firms with higher renewable electricity usage, even if those firms are heavily engaged in AI infrastructure. Non-PRI investors do not show clear patterns in Columns 4–6.

Panel B explores how investors respond to AI adoption through talent integration among user firms. The key independent variables include the share of AI workers, the logarithm of Scope 1 carbon emissions, and their interaction term. Among PRI investors in Columns 1–3, the coefficient on AI worker share is positive and becomes statistically significant only in the following one year, suggesting an initial increase in investment. More importantly, the interaction terms between AI worker share and Scope 1 emissions are negative and significant. This indicates that climate-conscious investors reduce their holdings in AI-adopting firms when those firms exhibit rising carbon emissions. Additionally, Scope 1 emissions alone have negative and significant coefficients across the three-year period, further confirming that PRI investors penalize firms with worsening emission profiles. By contrast, non-PRI investors show no consistent response and their investment decisions are not meaningfully influenced by carbon emissions and AI integration.

These results suggest that climate-conscious institutional investors exhibit cautious attention to the climate risks associated with certain forms of AI adoption. Specifically, they appear more likely to divest from firms involved in AI infrastructure development that rely on non-renewable electricity, as well as from firms integrating AI talent while experiencing rising carbon emissions. Currently, the overall investor response remains relatively modest. One possible explanation is that the climate implications of AI adoption are context-dependent, which makes it challenging even for responsible investors to incorporate AI-related climate risks into their investment strategies in a systematic way.

6 Conclusion

This project provides empirical evidence on the dual climate impact of corporate AI adoption, leveraging a large and globally representative firm-level dataset. We find that the energy burden of AI adoption is highly concentrated among firms with in-house infrastructure capabilities, which tend to mitigate their carbon footprint through greater use of renewable electricity. In contrast, firms without such infrastructure show no significant increase in energy use, and in some cases, even exhibit declines. At the same time, AI adoption is associated with immediate reductions in Scope 1 emissions, though the magnitude and persistence of this effect vary across sectors and are most pronounced among low-emitting firms. These findings highlight the importance of integrating climate considerations into corporate AI strategy, recognizing that the climate impact of AI extends beyond computing energy demands to the broader operational and sectoral contexts in which AI is deployed. Given AI adoption accelerates, understanding and managing such climate footprints will be critical for aligning innovation with global decarbonization goals.

Data Appendix

Variables	Definitions	Sources
$AIInfra$	Proportion of AI infrastructure-related content in 10-K filings (U.S. firms) or annual reports (non-U.S. firms).	SEC EDGAR
$\mathbb{1}(AIInfras)$	Indicator for above-median $AIInfra$ disclosures.	SEC EDGAR
$AIWorker$	Share of employees in AI-related roles.	LinkedIn
$AIWorker^{Market}$	Share of AI employees in market-facing functions.	LinkedIn
$AIWorker^{Internal}$	Share of AI employees in internal operations.	LinkedIn
$Electricity$	Logarithm of total electricity consumption.	CDP
$ElecInt$	Logarithm of electricity consumption scaled by sales.	CDP
$Electricity^{Renew}$	Logarithm of renewable electricity consumption.	CDP
$Electricity^{NonRenew}$	Logarithm of non-renewable electricity consumption.	CDP
$CE^{S2(Loc)}$	Logarithm of location-based Scope 2 emissions.	CDP, MSCI
$CE^{S2(Mkt)}$	Logarithm of market-based Scope 2 emissions.	CDP, MSCI
CE^{S1}	Logarithm of Scope 1 emissions.	MSCI
$DataAssets$	Proportion of 10-K paragraphs containing data asset-related keywords. These include: big data, data analytic, data driven, customer insight, consumer insight, data asset, data resource, data platform, data warehouse, data lake, and data science.	SEC EDGAR
$CompetePressure$	Proportion of 10-K paragraphs mentioning both competition- and technology-related terms. Competition-related keywords include: intense competition, intense competitive, competitive pressure, competitive threat. Technology-related keywords include: technology, machine learning, artificial intelligence, digital disruption, and technological disruption.	SEC EDGAR
$ElectricityPrice$	Annual average commercial electricity price at the state level.	EIA
$RenewableGrids$	Share of renewable electricity in local grids at the state level (inside the U.S.) or country level (outside the U.S.).	EIA
$EmissionTarget$	Indicator for firms with approved science-based or net-zero emission targets.	MSCI
$log(Asset)$	Logarithm of total assets.	Compustat
ROA	Operating income before depreciation divided by total assets.	Compustat
$Capex/Assets$	Capital expenditures divided by total assets.	Compustat
$Leverage$	Total debt divided by total assets.	Compustat
$Tangibility$	Net property, plant, and equipment divided by total assets.	Compustat
$SalesGrowth$	Percentages change in sales.	Compustat

References

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Reenen, J. V. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024a). Artificial intelligence and firms' systematic risk. *Working Paper, SSRN 4868770*.

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

Bertomeu, J., Lin, Y., Liu, Y., and Ni, Z. (2023). Capital market consequences of generative AI: Early evidence from the ban of chatgpt in italy. *Working Paper, SSRN 4452670*.

Bogmans, C., Gomez-Gonzalez, P., Ganpurev, G., Melina, G., Pescatori, A., and Thube, S. (2025). Power hungry: How AI will drive energy demand. *IMF Working Papers*, 2025(081):1.

Bolton, P., Kacperczyk, M. T., and Wiedemann, M. (2024). The co2 question: Technical progress and the climate crisis. *Working Paper, SSRN 4212567*.

Cen, L., Han, Y., Qiu, J., and Wu, J. (2024). Artificial intelligence along the supply chain. *Working Paper, SSRN 4929988*.

Chen, J., Hshieh, S., and Zhang, F. (2021). The role of high-skilled foreign labor in startup performance: Evidence from two natural experiments. *Journal of Financial Economics*, 142(1):430–452.

Corhay, A., Hu, K., Li, J., Tong, J., and Tsou, C.-Y. (2025). Data, markups, and asset prices. *Working Paper, SSRN 5093509*.

Cotter, J., Tuna, I., and Wysocki, P. D. (2006). Expectations management and beatable targets: How do analysts react to explicit earnings guidance? *Contemporary Accounting Research*, 23(3):593–624.

Dimmock, S. G., Huang, J., and Weisbenner, S. J. (2022). Give me your tired, your poor, your high-skilled labor: H-1b lottery outcomes and entrepreneurial success. *Management Science*, 68(9):6950–6970.

Dyck, A., Lins, K. V., Roth, L., and Wagner, H. F. (2019). Do institutional investors drive corporate social responsibility? international evidence. *Journal of Financial Economics*, 131(3):693–714.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., and Taska, B. (2023). Generative AI and firm values. *Working Paper, SSRN 4436627*.

Eloundou, Tyna, S. M. P. M. and Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv:2303.10130*.

Farboodi, M. and Veldkamp, L. (2022). A model of the data economy. *Working Paper, NBER 28427*.

Farboodi, M. and Veldkamp, L. (2023). Data and markets. *Annual Review of Economics*, 15:23–40.

Feng, M., Li, C., and McVay, S. (2009). Internal control and management guidance. *Journal of Accounting and Economics*, 48(2):190–209.

Gibson Brandon, R., Glossner, S., Krueger, P., Matos, P., and Steffen, T. (2022). Do responsible investors invest responsibly? *Review of Finance*, 26(6):1389–1432.

Hartzmark, S. M. and Shue, K. (2023). Counterproductive sustainable investing: The impact elasticity of brown and green firms. *Working Paper, SSRN 4359282*.

Jia, N., Li, N., Ma, G., and Xu, D. (2024). Corporate responses to generative AI: Early evidence from conference calls. *Working Paper, SSRN 4736295*.

Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., and Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*.

Luccioni, A. S., Viguier, S., and Ligozat, A. L. (2022). Estimating the carbon footprint of BLOOM, a 176B parameter language model. *arXiv:2211.02001*.

Mihet, R. and Philippon, T. (2019). The economics of big data and artificial intelligence. In *Disruptive Innovation in Business and Finance in the Digital World*.

Sautner, Z., van Lent, L., Vilkov, G., and Zhang, R. (2023). Firm-level climate change exposure. *Journal of Finance*, 78(3):1449–1498.

Schneider, I., Xu, H., Benecke, S., Patterson, D., Huang, K., Ranganathan, P., and Elsworth, C. (2025). Life-cycle emissions of AI hardware: A cradle-to-grave approach and generational trends. *arXiv:2502.01671*.

Tanaka, M., Bloom, N., David, J. M., and Koga, M. (2020). Firm performance and macro forecast accuracy. *Journal of Monetary Economics*, 114:26–41.

Figure 1: **Time-Series Variation**

This figure illustrates the time-series variation in corporate AI adoption. The two lines showcase the proportion of firms classified as AI infrastructure developers (those with substantial AI infrastructure disclosures in annual reports) and AI talent users (those with AI workers but without substantial AI infrastructure disclosures), using the left axis. The two bars represent the average share of AI workers engaged in market-facing and internal operations among firms, based on the right axis. The vertical dotted lines mark three key technological milestones: the 2012 deep learning breakthrough with AlexNet, the 2017 introduction of the transformer architecture, and the 2022 release of ChatGPT.

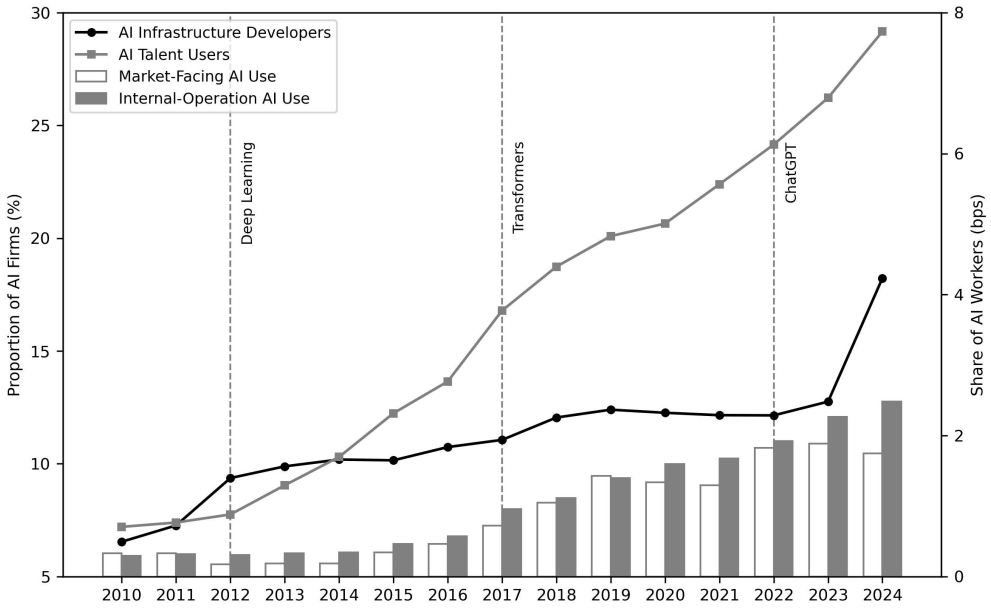
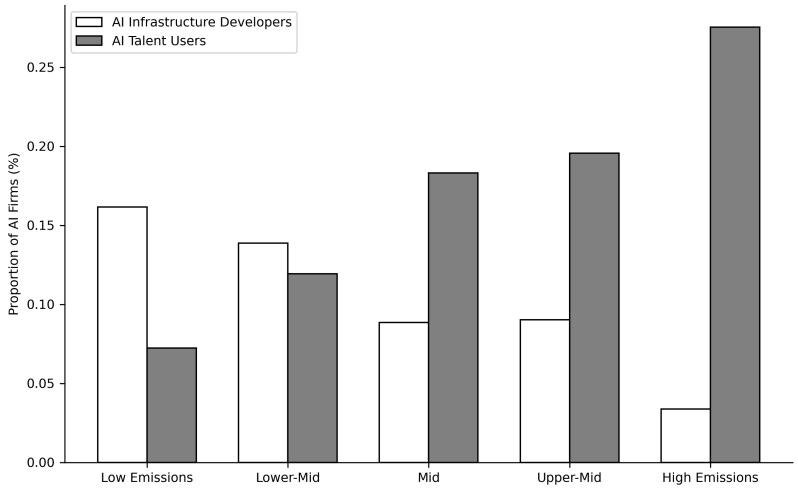
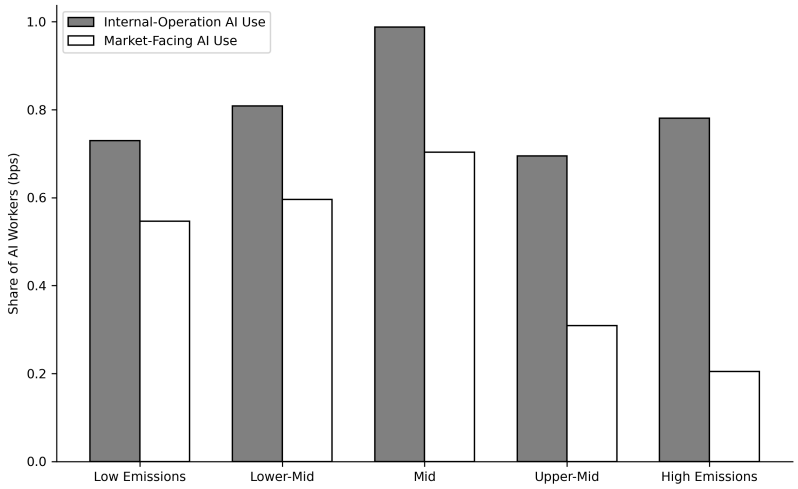


Figure 2: **Emission Profiles: Scope 1 Carbon Emissions**

The figure sorts companies into quintiles by Scope 1 emission levels and displays the distribution of AI firms across these groups. Panel A presents the proportion of AI infrastructure developers and talent users, respectively. Panel B displays the average share of AI workers across use domains.



(a)



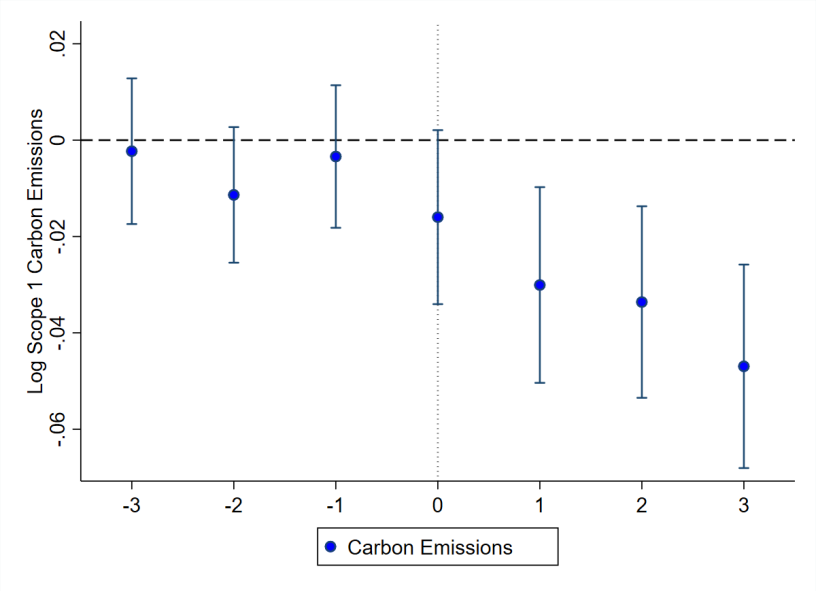
(b)

Figure 3: **AI Integration and Scope 1 Emissions: Dynamic Effects**

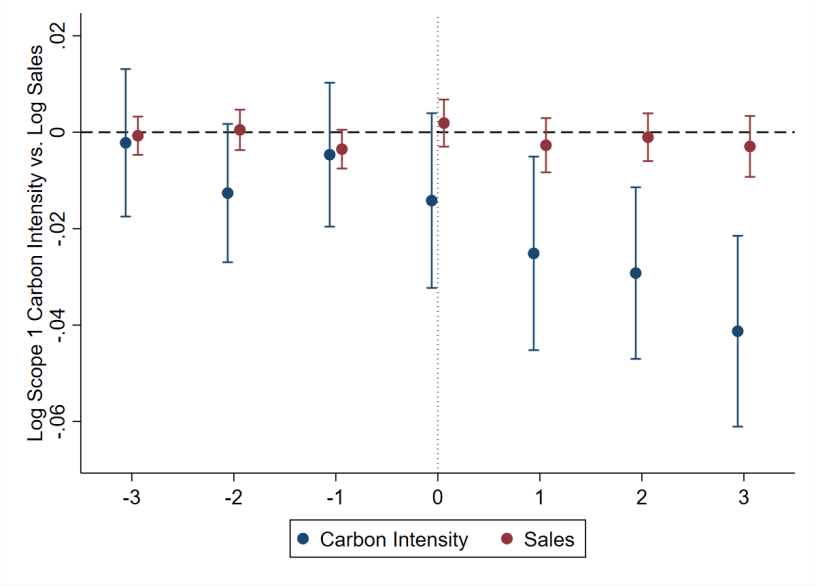
The figure plots the coefficients from the distributed lead-lag model, with vertical bars indicating 95% confidence intervals:

$$CE_{i,t}^{S1} = \alpha + \sum_{k=-3}^{3} \beta_k \, \Delta AIWorker_{i,t+k-1,t+k} + \mathbf{X}_{i,t} + a_i + b_{j,t} + c_{c,t} + \epsilon_{i,t}.$$

In Panel A, the dependent variable is the logarithm of Scope 1 carbon emissions; in Panel B, the dependent variables are the logarithm of Scope 1 carbon intensity and the logarithm of sales. The sample period is from 2010 to 2023. Standard errors are clustered at the firm level.



(a)



(b)

Figure 4: **AI Integration and Scope 1 Emissions: Industry Heterogeneity**

This figure shows the effect of AI adoption on direct carbon emissions at the industry level, based on the Fama-French 12 classification code. Purple and blue dots represent regression coefficients of future Scope 1 emissions on AI worker shares in market-facing and internal roles, respectively. Gray bars indicate the average Scope 1 emissions for each industry.

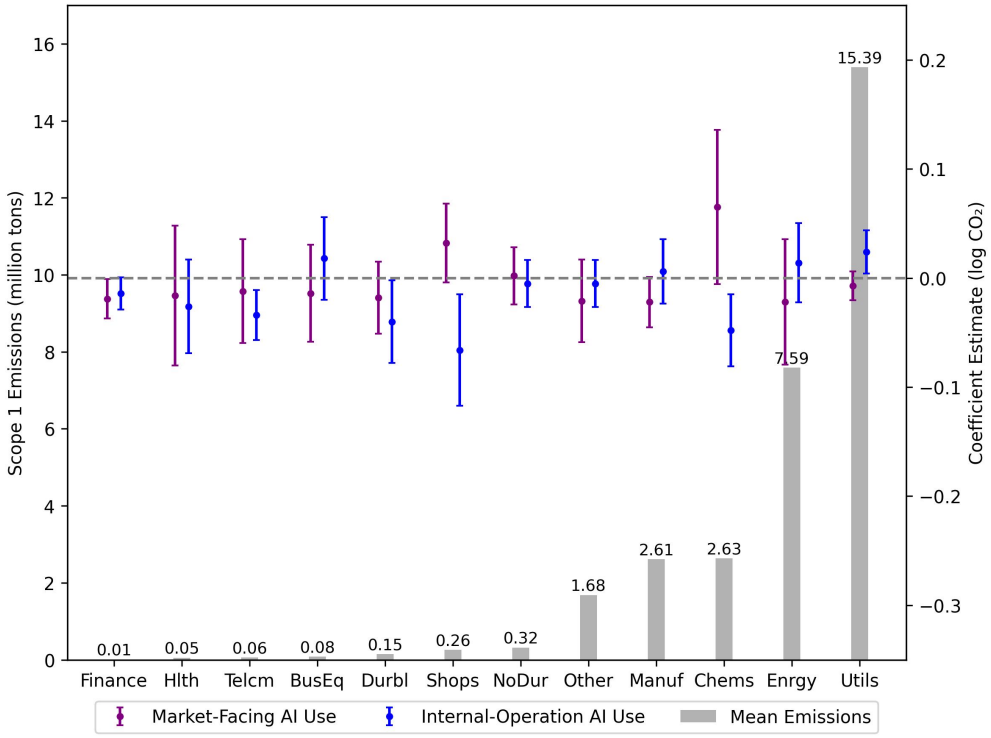


Figure 5: **AI Integration and Scope 1 Emissions: Narrative Retrieval**

This figure displays word clouds of AI worker job descriptions from the 50 firms with the highest DFBETA influence in the sector-level regression of future carbon emissions on AI worker share, based on the Fama-French 12 classification code. DFBETA measures how much the estimated regression coefficient would change if a specific firm were removed and thus identifies firms with the greatest influence on the estimated relationship. Panel A–F corresponds to Wholesale and Retail, Finance, Chemicals, Consumer Durables, Oil and Gas, and Utilities, respectively. The estimated coefficients are negative for Panel A–E and positive for Panel F (circled in red).



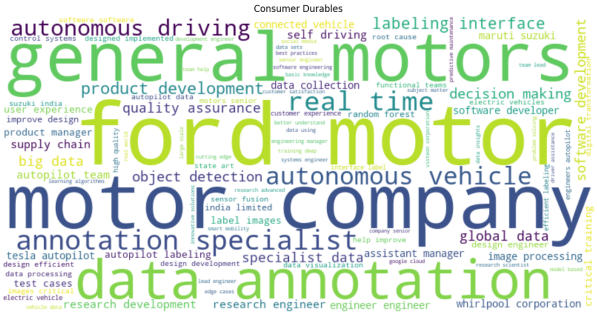
(a)



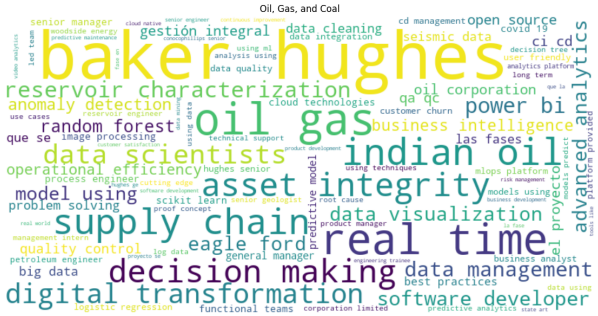
(b)



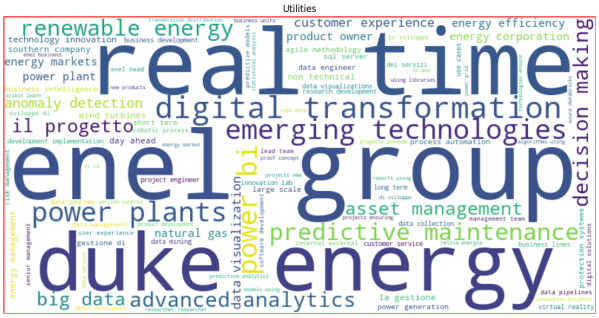
(c)



(d)



(e)



(f)

Table 1: **Illustrative Excerpts of AI Infrastructure and Worker Roles**

Panels A present excerpts from 10-K filings on AI infrastructures. Panels B and C show LinkedIn job descriptions for AI roles in market-facing and internal operations, respectively.

Panel A: AI Infrastructure Disclosure
1. Microsoft Corp., Computer Software, 2014 We own approximately three million additional square feet of office and data center space domestically and lease many sites domestically, totaling approximately **five million square feet**. Our strategy requires continuing **investment in data centers and other infrastructure** to support our devices and services.
2. Equinix, Inc., Trading, 2010 Equinix provides **49 data centers in 18 markets** around the world that protect and connect the most valued information assets. Equinix offers the following services: premium data center colocation, interconnection and exchange services, and **outsourced IT infrastructure** services.
3. NVIDIA Corp., Electronic Equipment, 2023 We introduced the Hopper architecture of data center **GPUs**, designed to **accelerate the training of AI transformer models** by an order of magnitude over the prior generation. It is ideal for accelerating applications such as large language models, deep recommender systems, and complex digital twins.
Panel B: AI Worker Profile (Market-Facing)
1. Time Inc., Printing and Publishing, Vice President Technology, Sep 2015 - Sep 2017 Executive-level role to build a product engineering team in Bangalore. My team is responsible for the publisher analytics and **advertising revenue**: i) 360-degree anonymous consumer navigation history for better audience insights; ii) Scalable big data platform using Hadoop stack on AWS; iii) NLP and Machine learning algorithms for **user behavior prediction**.
2. JPMorgan Chase & Co., Banking, Data Scientist, Sep 2018 - Feb 2020 Worked on several innovation projects and utilized cutting-edge technology like Machine Learning, Recommendation Engine, and NLP to enhance specialist performance to **improve customers' experience**.
3. Warner Music Group Corp., Recreation, Research & Analysis, Director, Jun 2021 - Feb 2024 Led a global team (data engineers, data analysts, data scientists) to develop a marketing mix model that boosted revenue by 27%. Built the team from scratch and delivered a machine learning-based **Sales Forecasting** web application. **Optimized customer acquisition**, reducing Cost Per Action (CPA) by 30% through campaign enhancements, uplift models, and A/B testing.
Panel C: AI Worker Profile (Internal-Operation)
1. KLA Corp., Measuring and Control Equipment, Metrology Systems Engineer, May 2012 - Oct 2024 Make systems perform beyond requirements by **advanced diagnostics design** based on supervised learning from MFG Data. Develop methods of alignment of complex optical systems utilizing image processing and deep learning techniques.
2. Schlumberger Ltd., Petroleum and Natural Gas, Division Manager, Jan 2022 - Aug 2024 Managing a one billion dollar Drilling Program. Responsible for deploying Drilling Digital to **automate drilling operations**, remote monitoring analysis, and developing Digital programs, as well as for deploying visual analytics, AI, and ML for **safety operations** at the rig floor.
3. UPL Ltd., Chemicals, Data Scientist, Oct 2020 - Apr 2024 In the Manufacturing domain, I have worked with Plant Maintenance data to extract valuable insights using NLP approaches like n-grams model and procurement Dashboard. I have managed and delivered an end-to-end **Automation Pipeline** to fetch Data from SAP BW to QLiksense.

Table 2: **Summary Statistics**

This table reports summary statistics for key variables used in the analysis. Panel A covers firms with electricity data (2013–2023) and Panel B includes those used for Scope 1 emission analysis (2010–2023). Variable definitions are provided in the Data Appendix.

Variable	Mean	SD	p10	p50	p90	N
Panel A: Electricity Sample						
$AIInfra\,(\%)$	0.23	0.68	0	0	0.66	4,272
$AIWorker\,(bp)$	5.33	10.59	0	0.55	15.47	4,272
$AIWorker^{Market}\,(bp)$	2.05	5.49	0	0.03	5.09	4,272
$AIWorker^{Internal}\,(bp)$	3.21	5.66	0	0.36	10.01	4,272
$Electricity$	1,322,968	2,810,351	35,430	369,738	3,489,033	4,272
$Electricity^{Renew}$	256,831	721,716	0	27,284	549,720	4,034
$Electricity^{NonRenew}$	1,058,110	2,475,267	4,236	234,720	2,789,262	4,140
$CE^{S2(Loc)}$	590,204	1,296,271	12,292	150,265	1,551,129	4,160
$CE^{S2(Mkt)}$	469,603	1,090,026	1,286	103,475	1,232,313	3,702
Panel B: Direct Emission Sample						
$AIWorker\,(bp)$	0.98	4.32	0	0	1.69	38,969
$AIWorker^{Market}\,(bp)$	0.36	2.18	0	0	0.26	38,969
$AIWorker^{Internal}\,(bp)$	0.60	2.39	0	0	1.11	38,969
CE^{S1}	1,357,820	5,373,256	284	16,063	1,782,281	38,969
CEI^{S1}	195.41	641.31	0.37	8.53	405.35	38,936

Table 3: **AI Infrastructure and Electricity Consumption**

This table analyzes the relationship between changes in electricity consumption and AI infrastructure disclosures. The dependent variables are changes from year t to t+1, t+2, and t+3 in total electricity consumption level/intensity (Panel A), location- and market-based Scope 2 emissions (Panel B), and renewable and non-renewable electricity consumption (Panel C). Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of electricity consumption. The sample period is from 2013 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Electricity Consumption						
	$\Delta Electricity$			$\Delta ElecInt$		
	$t+1$ - t	$t+2$ - t	$t+3$ - t	$t+1$ - t	$t+2$ - t	$t+3$ - t
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}(AIInfras)$	0.05**	0.12**	0.11*	0.08***	0.12***	0.16***
	(2.41)	(2.29)	(1.86)	(3.28)	(2.67)	(2.90)
$AIWorker$	0.00	0.00	0.01	-0.01	-0.00	-0.01
	(0.18)	(0.37)	(0.90)	(-1.19)	(-0.04)	(-0.53)
N	2,762	2,205	1,783	2,725	2,205	1,755
R^2	0.073	0.135	0.151	0.124	0.176	0.211
Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Scope 2 Carbon Emissions						
	$\Delta CE^{S2(Loc)}$			$\Delta CE^{S2(Mkt)}$		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}(AIInfras)$	0.09***	0.16***	0.19**	0.10	0.18*	0.31*
	(2.77)	(2.73)	(2.36)	(1.43)	(1.67)	(1.76)
$AIWorker$	-0.00	-0.00	0.01	-0.04***	-0.08***	-0.16***
	(-0.69)	(-0.03)	(0.32)	(-2.59)	(-2.97)	(-3.43)
N	2,699	2,149	1,741	2,252	1,772	1,424
R^2	0.081	0.119	0.157	0.056	0.132	0.176
Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel C: Renewable vs. Non-renewable Electricity						
	$\Delta Electricity^{Renew}$			$\Delta Electricity^{NonRenew}$		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}(AIInfras)$	0.44**	0.65**	0.57	0.19	0.47*	0.76**
	(2.24)	(2.08)	(1.49)	(1.18)	(1.85)	(2.12)
$AIWorker$	0.05	0.05	0.01	-0.08**	-0.14*	-0.21*
	(1.03)	(0.69)	(0.12)	(-2.03)	(-1.95)	(-1.94)
N	2,519	1,993	1,603	2,639	2,090	1,686
R^2	0.251	0.397	0.486	0.061	0.087	0.147
Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 4: **AI Infrastructure and Electricity Consumptions: Supply Chain Spillovers**

This table examines the relationship between energy consumption and AI adoption across the supply chain. Panel A examines how customers' AI adoption affects their suppliers. The dependent variables are suppliers' AI infrastructure disclosures, electricity consumption, and Scope 1 emissions; the key independent variable is the AI worker share of customer firms. Panel B examines the reverse direction. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	AI Users' Suppliers			*AI Developers' Customers*		
	$AIInfra^S$	$Electricity^S$	$CE^{S1,S}$	$AIWorker^C$	$Electricity^C$	$CE^{S1,C}$
	$t+1$	$t+1$	$t+1$	$t+1$	$t+1$	$t+1$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker^C$	0.11***	0.05***	-0.01			
	(3.64)	(2.94)	(-0.33)			
$AIInfra^C$	-0.01	-0.01	-0.01			
	(-1.12)	(-1.54)	(-1.51)			
$AIInfra^S$				0.08**	-0.01	-0.02*
				(2.19)	(-0.49)	(-1.86)
$AIWorker^S$				0.04	0.02*	-0.01*
				(1.33)	(1.81)	(-1.88)
N	10,759	1,763	7,720	13,824	7,531	12,738
R^2	0.754	0.961	0.963	0.866	0.951	0.986
Control	Yes	Yes	Yes	Yes	Yes	Yes
Supplier-Customer FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: **AI Integration and Scope 1 Emissions**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are Scope 1 carbon emissions in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	CE^{S1}					
	$t+1$		$t+2$		$t+3$	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	-0.03***		-0.06***		-0.08***	
	(-3.83)		(-4.14)		(-4.47)	
$AIWorker^{Market}$		-0.03**		-0.05***		-0.06***
		(-2.36)		(-2.60)		(-2.82)
$AIWorker^{Internal}$		-0.01*		-0.03**		-0.04**
		(-1.73)		(-2.11)		(-2.17)
$log(Asset)$	0.07***	0.07***	0.07***	0.07***	0.06**	0.06**
	(5.36)	(5.36)	(3.71)	(3.70)	(2.54)	(2.54)
ROA	-0.09	-0.09	0.03	0.03	0.18	0.18
	(-0.82)	(-0.82)	(0.30)	(0.29)	(1.23)	(1.22)
$Capex/Assets$	-0.04	-0.04	0.28	0.29	0.15	0.15
	(-0.28)	(-0.27)	(1.25)	(1.27)	(0.53)	(0.55)
$Leverage$	0.07	0.07	0.16**	0.16**	0.16**	0.17**
	(1.57)	(1.59)	(2.36)	(2.37)	(2.00)	(2.01)
$Tangibility$	0.08	0.08	0.01	0.01	-0.12	-0.12
	(0.93)	(0.93)	(0.05)	(0.06)	(-0.61)	(-0.61)
$SalesGrowth$	-0.03	-0.03	0.00	0.00	-0.03	-0.03
	(-1.02)	(-1.02)	(0.12)	(0.12)	(-0.86)	(-0.86)
$CE^{S1}_{i,t}$	0.66***	0.66***	0.43***	0.43***	0.24***	0.24***
	(48.15)	(48.11)	(27.07)	(27.04)	(14.43)	(14.41)
N	29,217	29,217	26,107	26,107	23,078	23,078
R^2	0.977	0.977	0.969	0.969	0.966	0.966
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: **AI Integration and Scope 1 Emissions: Carbon Intensity vs. Sales**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are the logarithm of Scope 1 carbon intensity (Panel A) and the logarithm of sale (Panel B) in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Carbon Intensity			
	CEI^{S1}		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker$	-0.03***	-0.05***	-0.07***
	(-2.89)	(-3.55)	(-3.95)
N	29,205	26,097	23,067
R^2	0.969	0.956	0.953
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes
Panel B: Sales			
	$logSales$		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker$	-0.02	-0.01	-0.02*
	(-1.59)	(-1.03)	(-1.84)
N	30,737	28,183	24,979
R^2	0.962	0.976	0.965
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes

Table 7: **AI Integration and Scope 1 Emissions: Instrument Variable**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are changes in Scope 1 carbon emissions from year t to years t+1, t+2, and t+3. The independent variable is the change in AI worker share from year t-1 to year t. The instrument variable is the "excess" H1B visa selections in year t-1 in Panel A and the AI-related selections in Panel B:

$$ExcessVisa_{i,t} = VisaSelections_{i,t} - VisaDemand_{i,t} \times SelectRate_t.$$

$VisaSelections_{i,t}$ is the number of firm i's H1B visa registrations drawn in the lottery, $VisaDemand_{i,t}$ is the total registrations submitted, and $SelectRate_t$ is the average annual selection rate. AI-related selections are weighted by the share of visa applications for STEM positions. Control variables include changes in $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, and $SalesGrowth$. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample consists of U.S. firms from 2014 to 2023 (as only partial lottery selection occurred between 2010 to 2013). t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Excess H1B Visa Selection				
	$\Delta AIWorker$	$\Delta CE^{S1}_{t,t+1}$	$\Delta CE^{S1}_{t,t+2}$	$\Delta CE^{S1}_{t,t+3}$
	(1)	(2)	(3)	(4)
$ExcessVisa$	0.14***			
	(4.88)			
$\Delta AIWorker$		-0.06**	-0.08**	-0.08*
		(-2.40)	(-2.22)	(-1.72)
N	9,972	8,578	7,244	5,961
F-Stat	23.8	23.5	20.2	18.7
Control	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
Panel B: AI-related Excess H1B Visa Selection				
	$\Delta AIWorker$	$\Delta CE^{S1}_{t,t+1}$	$\Delta CE^{S1}_{t,t+2}$	$\Delta CE^{S1}_{t,t+3}$
	(1)	(2)	(3)	(4)
$AIExcessVisa$	0.14***			
	(4.70)			
$\Delta AIWorker$		-0.07***	-0.09**	-0.10*
		(-2.68)	(-2.32)	(-1.83)
N	9,972	8,578	7,244	5,961
F-Stat	22.1	22.9	20.5	17.5
Control	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Table 8: **AI Integration and Scope 1 Emissions: Economic Mechanisms**

This table examines how AI integration reduces emissions via two channels. In Panel A, we measure predictive improvements using $ForecastAccuracy_{i,t} = -\frac{|SalesForecast_{i,t-1} - ActualSales_{i,t}|}{SalesPerShare_{i,t-1}}$. Columns 1–2 report results using the mean analyst forecast one year before the actual sales announcement, while Columns 3–4 use the median forecast. In Panel B, we measure low-carbon digitalization as the year-over-year changes in sales-weighted carbon intensity across business segments to capture firms' shift toward low-carbon products or processes. Columns 1–2 report the sales-weighted intensity, and Columns 3–4 use capital expenditure-weighted intensity. The dependent variables are Scope 1 carbon emissions in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Predictive Improvements

	CE^{S1}			
	Mean		Median	
	(1)	(2)	(3)	(4)
$AIWorker \times ForecastAccuracy$	-0.10**		-0.09**	
	(-2.24)		(-2.34)	
$AIWorker$	-0.04***		-0.04***	
	(-3.95)		(-3.95)	
$AIWorker^{Market} \times ForecastAccuracy$		-0.10*		-0.09**
		(-1.90)		(-2.04)
$AIWorker^{Market}$		-0.03**		-0.03**
		(-2.38)		(-2.38)
$AIWorker^{Internal} \times ForecastAccuracy$		-0.03		-0.02
		(-1.50)		(-1.44)
$AIWorker^{Internal}$		-0.02*		-0.02
		(-1.68)		(-1.64)
$ForecastAccuracy$	0.03	0.03	0.03	0.03
	(1.41)	(1.40)	(1.31)	(1.31)
N	22,752	22,752	22,752	22,752
R^2	0.979	0.979	0.979	0.979
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes

Panel B: Low-Carbon Digitalization

	CE^{S1}			
	Sales		Capex	
	(1)	(2)	(3)	(4)
$AIWorker \times LowCarbonDigit$	-0.04**		-0.03*	
	(-2.57)		(-1.79)	
$AIWorker$	-0.06***		-0.06***	
	(-4.66)		(-3.09)	
$AIWorkerr^{Market} \times LowCarbonDigit$		0.00		-0.12
		(0.11)		(-1.33)
$AIWorker^{Market}$		-0.05***		-0.05*
		(-2.85)		(-1.96)
$AIWorker^{Internal} \times LowCarbonDigit$		-0.03*		-0.01*
		(-1.74)		(-1.90)
$AIWorker^{Internal}$		-0.02		-0.03
		(-1.62)		(-1.50)
$LowCarbonDigit$	-0.01	-0.00	-0.01	-0.03
	(-1.35)	(-0.98)	(-1.50)	(-1.45)
N	11,077	11,077	8,358	8,358
R^2	0.982	0.982	0.982	0.982
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes

Table 9: **AI Integration and Scope 1 Emissions: Alternative Mechanisms**

This table explores alternative channels through which AI integration may reduce emissions, with a focus on those related to carbon-intensive sectors. Panel A measures carbon-related improvements using the year-over-year change in the logarithm of green patent numbers in Columns 1–2, methane leakage in Columns 3–4, and renewable electricity generation in Columns 5–6. Panel B measures capacity maximization using the year-over-year change in oil production scaled by proved developed reservior in Columns 1–2, natural gas production scaled by proved developed reservior in Columns 3–4, and electricity generation scaled by capacity in Columns 5–6. The dependent variables are Scope 1 carbon emissions in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample consists of U.S. firms from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Carbon Output Improvement

	CE^{S1}					
	GreenPatent		MethaneLekage		RenewableGeneration	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times CarbonOutput$	0.00		0.08		0.00	
	(0.12)		(1.36)		(0.76)	
$AIWorker$	-0.03***		-0.03*		-0.10	
	(-3.14)		(-1.86)		(-1.55)	
$AIWorker^{Market} \times CarbonOutput$		-0.01		-0.05		-0.00
		(-1.06)		(-0.10)		(-0.43)
$AIWorker^{Market}$		-0.02**		0.17		-0.15*
		(-2.54)		(1.02)		(-1.85)
$AIWorker^{Internal} \times CarbonOutput$		0.01		0.06		0.00
		(1.50)		(1.19)		(0.95)
$AIWorker^{Internal}$		-0.01		-0.02*		-0.03
		(-1.53)		(-1.90)		(-0.67)
$CarbonOutput$	-0.02*	-0.02*	0.03	0.01	0.00	-0.00
	(-1.68)	(-1.76)	(1.29)	(0.20)	(0.14)	(-0.24)
N	12,225	12,225	583	583	310	310
R^2	0.980	0.980	0.966	0.966	0.992	0.992
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Capacity Maximization

	CE^{S1}					
	OilReservior		NGReservior		GenerationCapacity	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times CapacityMaximization$	0.01*		0.04***		0.11	
	(2.02)		(3.27)		(0.49)	
$AIWorker$	-0.02		-0.04**		-0.09	
	(-0.95)		(-2.20)		(-1.52)	
$AIWorker^{Market} \times CapacityMaximization$		-2.56		-2.65		0.23
		(-1.04)		(-1.16)		(0.16)
$AIWorker^{Market}$		-2.98		-1.81		-0.12*
		(-0.91)		(-1.60)		(-1.99)
$AIWorker^{Internal} \times CapacityMaximization$		0.01**		0.02***		0.04
		(2.45)		(4.50)		(0.15)
$AIWorker^{Internal}$		-0.03		-0.03		-0.03
		(-1.50)		(-1.43)		(-0.61)
$CapacityMaximization$	0.00**	-0.46	0.01	-0.48	0.02	0.05
	(2.32)	(-1.03)	(0.75)	(-1.15)	(0.23)	(0.25)
N	163	163	152	152	310	310
R^2	0.928	0.927	0.948	0.947	0.992	0.992
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 10: **AI Adoption and Climate Conscious Investors**

This table examines the relationship between institutional ownership (IO) and corporate AI adoption. The dependent variables are IO from climate-conscious investors (i.e., PRI signatories, Columns 1–3) and non-climate-conscious investors (non-PRI signatories, Columns 4–5) in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$. Panel B excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2013 to 2023 in Panel A and from 2010 to 2023 in Panel B. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: AI Infrastructure Developers						
	IO^{PRI}			IO^{NonPRI}		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}(AIInfras) \times Electricity^{Renew}$	0.18*	0.18*	0.15*	-0.20	0.01	0.13
	(1.80)	(1.78)	(1.85)	(-1.10)	(0.04)	(0.91)
$\mathbb{1}(AIInfras)$	-1.71	-1.91	-2.45**	-0.89	-0.59	-1.91
	(-1.31)	(-1.46)	(-2.24)	(-0.42)	(-0.26)	(-1.00)
$Electricity^{Renew}$	-0.04	-0.05	-0.05	0.00	-0.00	-0.01
	(-1.04)	(-1.40)	(-1.28)	(0.01)	(-0.06)	(-0.24)
N	1,769	1,383	1,095	1,769	1,383	1,095
R^2	0.862	0.863	0.855	0.758	0.804	0.796
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: AI Talent Users						
	IO^{PRI}			IO^{NonPRI}		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times CE^{S1}$	-0.28***	-0.17**	-0.11	0.12	0.13	0.09
	(-4.24)	(-2.16)	(-1.27)	(1.02)	(1.22)	(0.83)
$AIWorker$	1.56***	0.85	0.71	-1.57	-1.59*	-1.16
	(3.00)	(1.30)	(0.94)	(-1.56)	(-1.80)	(-1.33)
CE^{S1}	-0.37**	-0.37**	-0.37**	0.05	0.10	0.06
	(-2.03)	(-2.10)	(-2.31)	(0.13)	(0.34)	(0.22)
N	13,350	12,101	10,815	13,350	12,101	10,815
R^2	0.842	0.835	0.832	0.729	0.755	0.780
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Internet Appendix

for

Carbon in the Cloud

This Internet Appendix provides supplementary material for this study.

A Additional Figures and Tables

Figure IA1: **Emission Profiles: Electricity Consumption**

This figure presents the average electricity consumption and Scope 2 emissions for AI infrastructure developers and AI talent users, based on a balanced sample of firms reporting all four metrics.

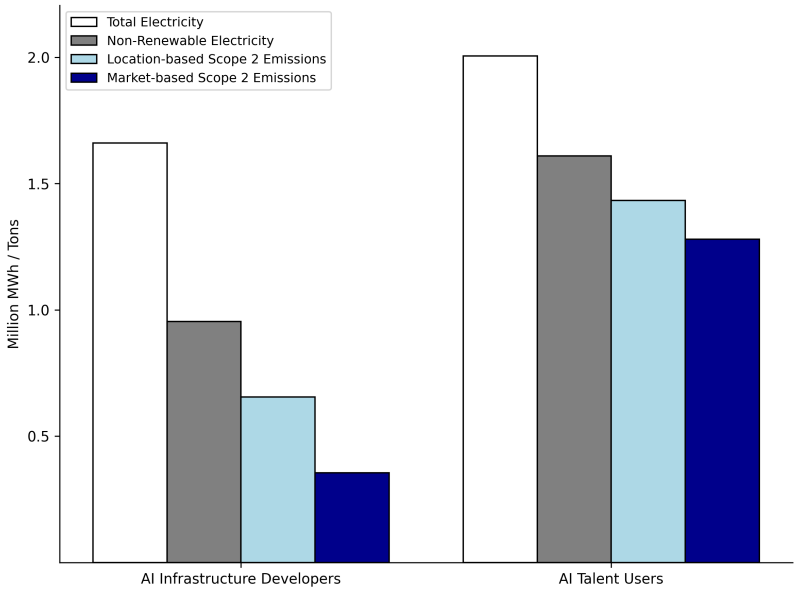
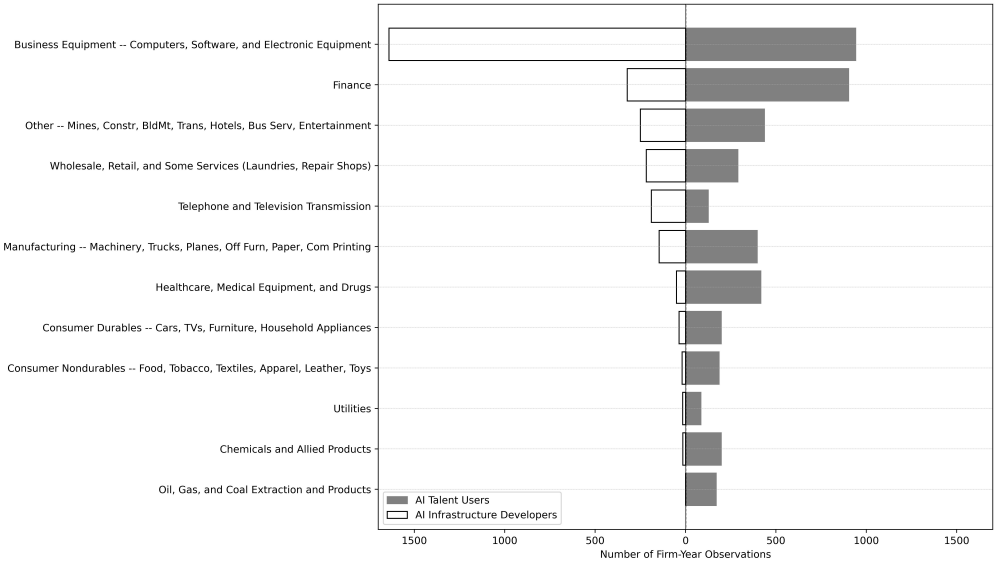
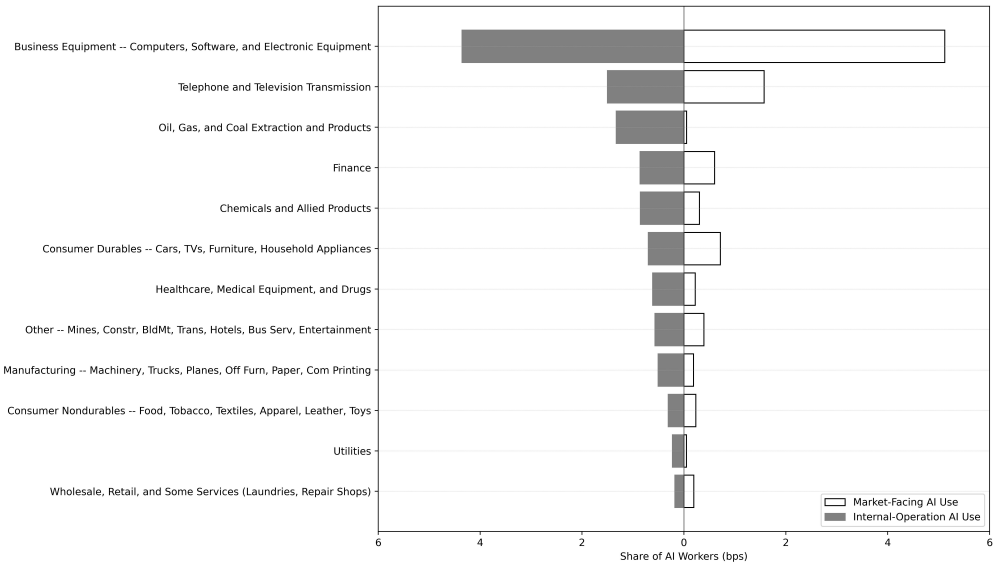


Figure IA2: **Industry Distribution**

This figure illustrates the distribution of AI firms at the industry level, based on the Fama-French 12 classification code. Panel A presents the industry distribution of firm-year observations for AI infrastructure developers (left bars) and talent users (right bars). Panel B reports, within each industry, the average share of AI workers engaged in market-facing and internal operations, in the sample excluding AI infrastructure developers.



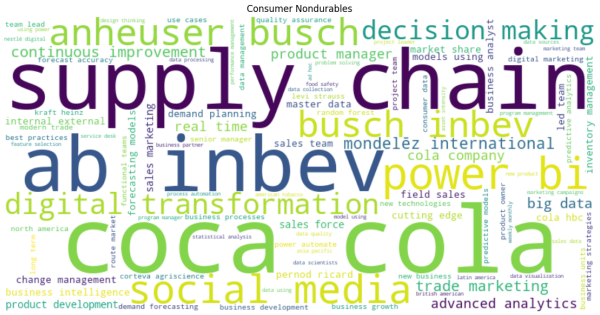
(a)



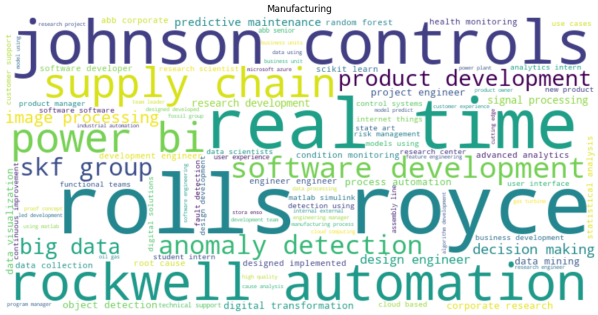
(b)

Figure IA3: **AI Integration and Scope 1 Emissions: Narrative Retrieval**

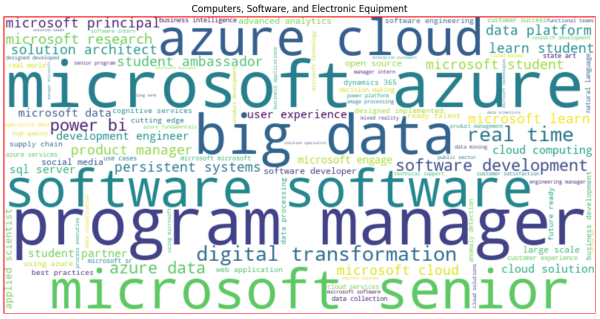
This figure displays word clouds of AI worker job descriptions from the 50 firms with the highest DFBETA influence in the sector-level regression of future carbon emissions on AI worker share, based on the Fama-French 12 classification code. DFBETA measures how much the estimated regression coefficient would change if a specific firm were removed and thus identifies firms with the greatest influence on the estimated relationship. Panel A–F corresponds to Consumer Nondurables, Manufacturing, Computers, Telephone Transmission, Healthcare, and Others, respectively. The estimated coefficients are positive for Panel C (circled in red) and negative for the other Panels.



(a)



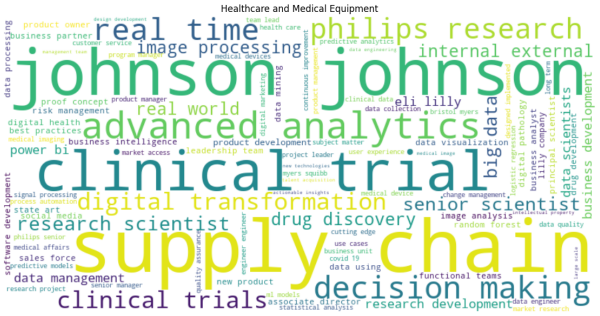
(b)



(c)



(d)



(e)



(f)

Table IA1: **Top Firms**

This table presents the top 30 AI-adopting firms, ranked by their average number of AI workers.

Company	$\#AIWorker$	$AIWorker$	$AIInfras$	$Electricity$	$Electricity^{NonRenew}$	CE^{S1}
Amazon.com Inc	3,454	0.23%	2.36%	42,436,459	1,873,573	13,795,000
International Business Machines	2,782	0.84%	0.31%	3,765,497	2,590,332	233,377
Microsoft Corporation	1,944	1.30%	1.68%	7,461,774	777,808	88,574
Accenture	1,679	0.32%	0.20%	400,348	270,672	21,086
Cognizant Technology Solutions	693	0.24%	0.48%	202,192	151,995	22,160
Intel Corporation	684	0.65%	1.91%	5,335,574	1,323,857	1,081,237
Apple Inc.	515	0.35%	0.31%	1,725,373	18,858	52,558
Meta	466	0.65%	3.39%	9,421,839	0	55,173
Oracle Corporation	456	0.32%	2.43%	1,420,625	853,492	13,071
SAP SE	451	0.46%	1.11%	327,278	22,965	134,816
Hewlett Packard Enterprise	349	0.54%	1.45%	945,725	601,911	58,676
Genpact International	334	0.32%	0.00%	99,848	95,611	5,235
Cisco Systems, Inc.	324	0.41%	1.71%	1,561,584	451,935	45,099
Siemens AG	303	0.09%	0.23%	1,739,845	720,563	670,897
Salesforce, Inc.	303	0.38%	3.26%	819,349	160,272	4,000
QUALCOMM Inc.	290	0.70%	0.93%	385,838	319,576	83,564
Nokia Group	266	0.31%	0.33%	775,488	462,392	91,833
NVIDIA Corporation	264	2.08%	4.46%	162,041	132,797	2,843
HP Inc	220	0.41%	0.06%	567,469	313,237	57,738
Adobe	217	0.86%	1.98%	175,201	100,817	9,148
DXC Technology	216	0.16%	0.72%	1,284,522	835,909	27,316
Citigroup Inc.	214	0.10%	0.01%	1,340,411	794,572	28,419
CGI Inc.	208	0.23%	0.05%	151,622	58,348	15,790
UBS Group AG	207	0.29%	0.00%	362,232	158	8,570
EXL Service Holdings	206	0.45%	0.71%	26,607	25,772	1,449
EPAM	201	0.37%	0.34%	15,874	14,958	219
Ford Motor Company	196	0.11%	0.00%	7,366,817	6,360,269	1,336,865
HSBC Holdings plc	175	0.08%	0.06%	707,287	484,572	35,919
AstraZeneca	167	0.21%	0.00%	692,742	2,539	238,590
CGI Group Inc.	163	0.21%	0.02%	188,310	99,162	17,649

Table IA2: **Determinants of AI Infrastructure Involvement**

This table examines the determinants of firms' decisions to develop AI infrastructure. The dependent variable, $\mathbb{1}(AIInfras)$, is a dummy variable equal to one if the firm has substantial AI infrastructure discussions (above-median) in annual reports, and zero otherwise. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	$\mathbb{1}(AIInfras)$					
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	0.02**					
	(2.51)					
$AIWorker^{Market}$		-0.00				
		(-0.13)				
$AIWorker^{Internal}$		0.02***				
		(2.60)				
$DataAssets$			0.03***			
			(3.93)			
$CompetePressure$				0.01**		
				(2.30)		
$ElectricityPrice$					0.02***	
					(3.35)	
$RenewableGrids$						0.01**
						(2.56)
$log(Asset)$	0.01**	0.01*	0.01**	0.01***	0.01**	0.01***
	(2.19)	(1.80)	(2.56)	(2.69)	(2.45)	(2.92)
ROA	0.01	0.01	0.01	0.01	0.03	0.02
	(0.34)	(0.21)	(0.32)	(0.32)	(0.69)	(0.60)
$Capex/Assets$	0.55***	0.54***	0.53***	0.56***	0.55***	0.57***
	(4.10)	(4.04)	(3.10)	(3.28)	(3.63)	(3.98)
$Leverage$	-0.00	0.00	-0.01	0.00	0.01	0.00
	(-0.05)	(0.02)	(-0.17)	(0.03)	(0.30)	(0.02)
$Tangibility$	-0.10***	-0.10***	-0.09**	-0.10**	-0.11**	-0.12***
	(-2.78)	(-2.72)	(-1.99)	(-2.24)	(-2.53)	(-3.18)
$SalesGrowth$	0.01**	0.02**	0.02***	0.02**	0.02**	0.01**
	(2.46)	(2.52)	(2.61)	(2.45)	(2.29)	(2.14)
N	20,974	20,974	15,699	15,699	17,325	20,313
R^2	0.223	0.224	0.225	0.218	0.226	0.218
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA3: **AI Integration and Scope 1 Emissions: Heterogeneous Effects**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are Scope 1 carbon emissions in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	CE^{S1}					
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times \mathbb{1}(After2018)$	-0.04***	-0.07***	-0.09***			
	(-4.14)	(-4.25)	(-4.30)			
$AIWorker \times \mathbb{1}(Before2018)$	-0.02	-0.04*	-0.08***			
	(-1.00)	(-1.80)	(-3.20)			
$AIWorker \times \mathbb{1}(US)$				-0.05***	-0.08***	-0.10***
				(-3.59)	(-4.07)	(-4.38)
$AIWorker \times \mathbb{1}(NonUS)$				-0.02*	-0.04**	-0.07**
				(-1.92)	(-2.04)	(-2.30)
N	29,217	26,107	23,078	29,217	26,107	23,078
R^2	0.977	0.969	0.966	0.977	0.969	0.966
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA4: **AI Integration and Indirect Carbon Emissions**

This table analyzes the relationship between future indirect emissions and AI worker shares. The dependent variables are Scope 2 carbon emissions in Panel A and Scope 3 emissions in Panel B in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

Panel A: Scope 2 Carbon Emissions			
	CE^{S2}		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker$	-0.03***	-0.06***	-0.10***
	(-3.66)	(-3.71)	(-3.78)
N	29,310	26,197	23,158
R^2	0.959	0.942	0.937
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes
Panel B: Scope 3 Carbon Emissions			
	CE^{S3}		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker$	-0.08	-0.14**	-0.09
	(-1.58)	(-2.00)	(-1.20)
N	9,265	7,693	6,374
R^2	0.843	0.813	0.808
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Country-Year FE	Yes	Yes	Yes

Table IA5: **AI Integration and Scope 1 Emissions: Robustness Check**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are Scope 1 carbon emissions in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of the dependent variable. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	CE^{S1}								
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Confounding Factors: Non-AI Skills									
$AIWorker$	-0.03***	-0.06***	-0.08***	-0.04***	-0.06***	-0.09***	-0.03***	-0.06***	-0.08***
	(-3.05)	(-3.87)	(-4.12)	(-3.85)	(-4.26)	(-4.69)	(-3.63)	(-4.05)	(-4.39)
$ITWorker$	-0.05**	-0.03	-0.05						
	(-2.43)	(-0.99)	(-1.40)						
$RobotWorker$				0.01	0.02	0.04**			
				(0.60)	(1.58)	(1.97)			
$DataWorker$							-0.03**	-0.03	-0.03
							(-2.50)	(-1.30)	(-1.08)
Panel B: Confounding Factors: Green Engagement									
$AIWorker$	-0.04***	-0.06***	-0.08***	-0.03**	-0.06***	-0.07***	-0.04***	-0.06***	-0.09***
	(-2.93)	(-3.23)	(-3.63)	(-2.15)	(-2.80)	(-2.58)	(-4.51)	(-4.27)	(-4.73)
$GreenPatent$	-0.01**	-0.01	-0.01						
	(-2.31)	(-0.78)	(-1.64)						
$EmissionTarget$				-0.01	-0.04*	-0.07***			
				(-0.54)	(-1.84)	(-2.72)			
$CCExposure$							0.00	0.01	0.02
							(0.33)	(1.14)	(1.49)
Panel C: Confounding Factors: Local Energy Source, Energy Price and Financial Constraints									
$AIWorker$	-0.04***	-0.06***	-0.08***	-0.04***	-0.06***	-0.08***	-0.03***	-0.06***	-0.08***
	(-3.86)	(-4.18)	(-4.39)	(-2.91)	(-3.21)	(-3.62)	(-3.77)	(-4.03)	(-4.37)
$ElectricityRenewable$	0.00	0.03	0.05						
	(0.09)	(0.75)	(1.12)						
$ElectricityPrice$				0.01	0.02	-0.01			
				(0.46)	(0.43)	(-0.16)			
WW							-0.00	-0.00	-0.00
							(-0.07)	(-0.27)	(-0.80)
Panel D: Alternative AI Worker Measures									
$I\{AIWorker\}$	-0.02*	-0.07***	-0.08***						
	(-1.68)	(-2.68)	(-2.75)						
$log(\#AIWorker)$				-0.02***	-0.05***	-0.06***			
				(-3.45)	(-4.56)	(-4.65)			
$\#AIWorker/Sales$							-0.02***	-0.04***	-0.05***
							(-3.61)	(-3.78)	(-3.73)
Panel E: Alternative Specifications									
	Industry-Year and Country FE			SE Cluster by Firm and Year			SE Cluster by Industry		
$AIWorker$	-0.01***	-0.02**	-0.04***	-0.03***	-0.06***	-0.08***	-0.03***	-0.06***	-0.08***
	(-2.65)	(-2.31)	(-2.74)	(-3.33)	(-3.57)	(-3.85)	(-4.24)	(-4.92)	(-4.71)

Table IA6: **AI Integration and Scope 1 Emissions: Alternative Instrument Variable**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are changes in Scope 1 carbon emissions from year t to years t+1, t+2, and t+3. The independent variable is the change in AI worker share from year t-1 to year t. The instrument variable is ... Control variables include changes in $log(Asset)$, ROA, $Capex/Assets$, and $SalesGrowth$. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification code) to examine the effect on AI user firms. The sample consists of U.S. firms from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	$\Delta AIWorker$	$\Delta CE^{S1}_{t,t+1}$	$\Delta CE^{S1}_{t,t+2}$	$\Delta CE^{S1}_{t,t+3}$
	(1)	(2)	(3)	(4)
$AIUniSupply$	0.09***			
	(4.09)			
$\Delta AIWorker$		-0.09*	-0.19*	-0.15
		(-1.84)	(-1.92)	(-0.95)
N	12,112	10,914	9,763	8,640
F-Stat	16.7	18.5	18.8	16.7
Control	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

B Additional Details for NLP Methods

B.1 Keywords for Preliminary Detection

B.1.1 AI-Infrastructure Involvement

data center, datacenter, data centre, datacentre, server room, server farm, server infrastructure, data facility, colocation, co-location, public cloud, cloud service, cloud infrastructure, amazon web service, microsoft azure, google cloud, ibm cloud, edge computing, edge infrastructure, edge server, edge node, computing infrastructure, edge device.

B.1.2 AI-Worker Engagement

artificial intelligence, ai, computer vision, machine learning, natural language processing, nd4j, kernel methods, xgboost, sentiment classification, long short-term memory, libsvm, semi-supervised learning, neural networks, word2vec, mxnet, autoencoder, autoencoders, mlpack, keras, theano, vowpal wabbit, jung framework, opennlp, neural network, natural language toolkit, unsupervised learning, dlib, scikit-learn, latent semantic analysis, latent dirichlet allocation, stochastic gradient descent, gradient boosting, dimensionality reduction, deep learning, dbscan, recommender systems, random forests, deeplearning4j, support vector machines, svm, unstructured information, apache uima, maximum entropy classifier, naive bayes, expectation-maximization, em algorithm, weka, speech recognition, clustering algorithms, object recognition, image recognition, bert, roberta, distilbert, bayesian networks, supervised learning, k-means, sentiment analysis, opinion mining, machine translation, openai, gpt-2, gpt-3, gpt-3.5, gpt-4, chatgpt, gpt2, gpt3, gpt3.5, gpt4, claude, google bard, bing chat, machine vision, lstm, rnn, nltk, reinforcement learning, genai, llama, llama2, llama3, llama-2, llama-3, hyperparameter tuning, transfer learning, hugging face, huggingface, cognitive computing, language models, prompt learning, generative adversarial networks, generative adversarial network, generative pre-trained transformer, llms, prompt engineer, prompt engineering, few-shot learning, zero-shot learning, meta-learning, representation learning, instruction tuning, diffusion models, large language models, prompt tuning, fine-tuning, dall-e, decision-making algorithms, autonomous vehicles, decision-making algorithm, facial recognition, microsoft cognitive toolkit, microsoft 365 copilot, language model, diffusion model, tensorflow, pytorch, generative models, transformer models, text generation, image generation, generative model, transformer model.

B.2 Prompts Used in GPT-4 and LLaMA-3 Fine-Tuning

B.2.1 Task 1: AI Infrastructure or Not

You are a sustainability analyst examining a paragraph from a company's 10-K. Your task is to determine whether the paragraph refers to anything plausibly related to data centers or AI infrastructures, even if without explicit detail.
Answer "Yes" if the paragraph mentions or implies the physical assets, construction, manufacturing, or operation of:
- hyperscale or edge data centres, server farms, GPU/TPU clusters,
- semiconductor fabrication for AI chips,
- networking/power/cooling systems supporting model training or inference,
- capital expenditure, energy purchases or efficiency measures specific to those assets.
Answer "No" if the paragraph:
- Only describes the company using AI infrastructures or services provided by others,
- The paragraph is irrelevant or does not mention any infrastructure, manufacturing, or provision of AI capabilities.
Please return your output strictly in the format of: CATEGORY (EXPLANATION).
Where CATEGORY is either "Yes" or "No", and EXPLANATION is one sentence explaining your reasoning.

B.2.2 Task 2: Market-Facing AI Use or Not

You are a financial analyst classifying LinkedIn job descriptions for an academic study on corporate AI adoption. Your goal is to determine whether the following job description belongs to Market-AI (label = 1) or Not-Market-AI (label = 0).
Market-AI roles build, deploy or maintain AI systems whose primary purpose is to enhance a firm's products, services, or customer relationships.
Typical use-cases (non-exhaustive):
- Recommender systems, personalization, search ranking, ad-tech targeting, pricing & demand prediction for customers.
- Generative-AI product features (chatbots, content creation, marketing copy, image or video synthesis offered to users).
- Customer-facing NLP/NLU (voice assistants, multilingual support bots).
- Fraud-detection or credit-scoring models sold to end customers (e.g., banking apps).
- Data-science or ML roles in Product, Growth, Marketing, UX, Ad-tech, CRM, Sales Ops.
Not-Market-AI covers everything else: operational optimization, predictive maintenance, cor-

porate back-office analytics, pure data engineering, non-AI IT, etc.
Return exactly a valid JSON object in this format:
"Market AI": 1 or 0, "Explanation": one-sentence explanation highlighting the key elements that influenced your decision.

B.2.3 Task 3: Internal-Operation AI Use or Not

You are a financial analyst classifying LinkedIn job descriptions for an academic study on corporate AI adoption. Your goal is to determine whether the following job description belongs to Operational-AI (label = 1) or Not-Operational-AI (label = 0).
Operational-AI roles focus on the development, deployment, or application of AI systems for internal use and improve the efficiency of operations, workflows, or teams.
Typical use-cases (non-exhaustive):
- AI for internal business processes (e.g., finance, HR, compliance).
- AI for internal IT automation or document generation.
- AI used in logistics, supply chain, scheduling, or production planning.
- Internal fraud detection, forecasting, or analytics dashboards.
- Internal LLMs or NLP tools used by employees.
- Workplace or facility optimization tools.
- HVAC or building energy optimization.
- Smart manufacturing, digital twins.
- AI in EHS (environment, health, safety) or facilities operations.
Not-Operational-AI covers everything else: customer-facing AI, generic AI tooling, or vague/ undocumented tasks.
Return exactly a valid JSON object in this format:
"Operational AI": 1 or 0, "Explanation": one-sentence explanation highlighting the key elements that influenced your decision.

C Construction of the Instrument Variable

We use excess H-1B visa selections as an instrument to capture firm-year variation in the supply of AI workers. This approach is valid for two main reasons. First, a substantial share of AI talent in the United States is foreign-born and typically require H-1B visas to work there legally. Second, the demand for H-1B visas has consistently exceed the annual cap and lead the U.S. government to allocate visas through a randomized lottery before evaluating applications. This lottery introduces quasi-random variation in firms' ability to hire foreign AI workers and creates plausibly exogenous shifts in the supply of AI talent at the firm level. This section first explains the two reasons, then describes how the measure is constructed and details the data sources used.

C.1 International AI Talent in the U.S.

Non-U.S. citizens represent a significant source of potential AI talent in the United States. According to the White House's AI talent report, 40%–60% of AI-related master's degrees earned at U.S. universities each year are awarded to non-citizens.[1] At the doctoral level, the share is even higher, with over half of all AI-related PhDs have gone to non-citizens since 2003, compared to about 20% in other fields. Our LinkedIn data further underscores this trend, as 15% of AI professionals working in the U.S. earned their first degree abroad and 44% hold at least one degree from a non-U.S. institution. These statistics highlight the U.S. AI sector's strong reliance on international talent. Whether these individuals are able to remain in the U.S. after graduation, which often relies on visa policies, plays a key role in shaping the effective domestic supply of advanced AI expertise.

[1] See the White House report here.

C.2 Institutional Background of the H-1B Visa Lottery

The H-1B visa program is the primary legal channel for U.S. firms to hire skilled foreign workers in specialized fields like AI and other STEM disciplines. Since the H-1B Visa Reform Act of 2004, the program has been subject to an annual cap of 65,000 visas, with an additional 20,000 visas reserved for applicants holding advanced degrees from U.S. institutions. This cap has consistently been binding; between 2010 and 2024 in our sample, USCIS often received more than twice as many applications as available slots. In such cases, the United State Citizenship and Immigration Service (USCIS) allocates visas through a randomized lottery, which introduces quasi-random variation in firms' access to foreign skilled labor.

To hire a foreign worker under the H-1B program, a U.S. employer must submit two filings. First, the employer files a Labor Condition Application (LCA) with the U.S. Department of Labor at the position level. It details the number of hires, job classification (SOC code), and employment dates. The LCA, which ensures compliance with wage and labor standards, can be submitted at any time and is typically processed within a few business days. Second, the employer files an H-1B petition (Form I-129) with U.S. Citizenship and Immigration Services (USCIS) for each individual worker. If approved, both filings authorize employment for up to three years, with a one-time extension of up to three additional years.

USCIS begins accepting H-1B petitions along with certified LCAs on the first business day of April each year for the upcoming fiscal year, which runs from October 1 to September 30. By regulation, the application window must remain open for at least five business days, regardless of how quickly the cap is reached. If the annual quota is not met within the five-day window, USCIS continues accepting petitions until it is filled and then conducts a lottery only among petitions received on the final day. This occurred in fiscal years 2010 to 2013, when the cap was reached only after several weeks or months. In contrast, from FY2014 to FY2024, the cap was reached within the initial five-day window. In these years, USCIS conducted a full lottery among all petitions submitted during that period using a

computer-based randomized algorithm. In that case, the exact cutoff date for inclusion in the lottery is determined internally by USCIS and is not disclosed in advance. As a result, firms typically submit petitions as early as possible after April 1 to maximize their chances.

C.3 The Measure of Excess H-1B Visa Selection

We construct firm-level measures of excess H-1B visa selections using multiple data sources. For fiscal years after 2021, we rely on a dataset Bloomberg News obtained via the Freedom of Information Act from USCIS.[2] This data includes the number of electronic lottery registrations, selections, subsequent I-129 petitions, and the employer identification number (EIN). It is available only from 2021 onward, when USCIS introduced an electronic registration system. This system provides a complete record of both demand (registrations) and lottery outcomes (selections), which was not available in earlier years.

There was no centralized registration system before 2021 and unselected petitions were returned without being recorded. For those years, we estimate demand and supply using LCA and I-129 petition data, following the methodology of Chen et al. (2021) and Dimmock et al. (2022). We proxy a company's demand for cap-subject high-skilled foreign workers using the number of intended hires reported in its LCAs.[3] Since the LCA form does not specify whether a position is cap-subject, we infer this from the timing of the filing. Specifically, we classify an LCA as likely cap-subject if it was certified between January and April and has an employment start date four to six months later. This approach is based on two considerations. First, since demand consistently exceeds the cap and only petitions submitted during the early April window are entered into the lottery, LCAs certified after this period are unlikely to support cap-subject filings for that fiscal year. Second, employers are incentivized to certify LCAs as late as possible and align start dates close to October 1 to maximize the overlap between the H-1B's three-year validity and the LCA's three-year term (which must begin

[2]See the data that Bloomberg requested from USCIS here.
[3]See LCA data from the Office of Foreign Labor Certification here.

within 180 days of certification). Certifying too early may trigger an earlier start date and reduce the effective employment period.

We proxy realized H-1B supply using the USCIS H-1B Employer Data Hub, which reports all processed H-1B cases and their action types.[4] We include only petitions coded as new employment (approved or denied) and exclude amendments, transfers, extensions, and filings from cap-exempt employers (e.g., universities or nonprofits). The dataset includes each employer's legal name and the last four digits of the EIN. We manually match them to Compustat identifiers and then verify matches using the EIN suffix. A limitation is that we do not observe lottery selections that never resulted in a petition (e.g., selected registrations for which the employer chose not to file Form I-129). Hence, this measure serves as a proxy for realized supply rather than for selections per se.

For firm i in year t, our instrument is the excess number of visa selected relative to expected selections given the aggregate lottery selection ratio, following Cen et al. (2024):

$$ExcessVisa_{i,t} = VisaSelections_{i,t} - VisaDemand_{i,t} \times SelectRate_t.$$

$VisaSelections_{i,t}$ is the number of firm i's H-1B visa registrations drawn in the lottery. This is observed directly in the Bloomberg/USCIS registration data from 2021 to 2024; for pre-2021 years, we proxy selections with submitted new cap-subject I-129 petitions. $VisaDemand_{i,t}$ is the total registrations submitted by the firm. This is also observed directly after 2021 and was proxied by LCAs for earlier years. $SelectRate_t$ is the fiscal-year lottery selection rate, computed as total selections divided by total registrations.

To better target foreign AI talent, we scale the instrument by the firm-year share of STEM workers in related LCAs based on their SOC codes. This refinement concentrates variation

[4]See USCIS H-1B Employer Data Hub here.

on positions most relevant to AI adoption:

$$AIExcessVisa_{i,t} = ExcessVisa_{i,t} \times STEMShareLCA_{i,t}.$$

We use both $ExcessVisa_{i,t}$ and $AIExcessVisa_{i,t}$ to serve as supply-side instruments and exploit quasi-random lottery shocks to identify the effect of foreign AI talent on subsequent emission outcomes.