Carbon in the Cloud*

Tingyu Yu†

August 2025

Abstract

This paper examines the climate impacts of corporate AI adoption using a large global sample. By analyzing annual reports and LinkedIn profiles, we identify corporate capabilities in AI infrastructure and workforce. Firms with AI infrastructure bear most of the energy burden and partially offset by using renewable electricity. In contrast, user firms with AI workers shift energy demands upstream and, more importantly, reduce their own operational (Scope 1) emissions. However, this reduction effect is concentrated in already low-emitting firms and largely absent in brown sectors, where AI is sometimes deployed to maximize the utilization of fossil assets and increase emissions. The findings suggest that AI functions as a form of directed technical change that can amplify sectoral heterogeneity in emission trajectories.

Keywords: AI adoption, carbon emissions, directed technical change, climate risks.

JEL classifications: Q54, O33, G11, D24.

*I am deeply grateful to Markus Leippold, Kai Li, and Zacharious Sautner for their continuous guidance. I thank Francisco Amaral, Olimpia Carradori, Glen Gostlow, Michel Habib, Felix von Meyerinck, Jingwei Ni, Steven Ongena, Stefano Ramelli (discussant), Chi-Yang Tsou, Alexander Wagner, Yucheng Yang, and participants at the UZH Rising Scholar Conference in Finance 2025 for detailed discussions and helpful comments. This research has been funded by the Swiss National Science Foundation (SNSF) under Grant Agreement No. 100018 207800.

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

1 Introduction

Technological change has long been central to economic growth and shapes both corporate production and its environmental externalities (Bresnahan and Trajtenberg, 1995; Acemoglu et al., 2012). Artificial intelligence (AI) represents a new wave of innovation with unusually high climate salience. As an emerging general-purpose technology, it is rapidly permeating business operations across industries.[1] This breadth of applications gives AI a wide range of potential climate impacts, depending on the direction of innovation and how production is reorganized. At the same time, AI training and inference are exceptionally electricity-intensive, with the resulting carbon footprint lying sharply on where and when computation takes place.[2] This combination of transformative scope and high energy intensity creates a distinct channel for climate externalities from corporate AI adoption. Understanding this channel is critical for guiding firms' sustainability strategies and aligning digital innovation with global decarbonization goals.

Quantifying these climate impacts presents significant challenges. First, the effects are inherently multidimensional. While running AI models consumes substantial electricity and directly increases emissions, AI can also change how firms operate, which may lower emissions through productivity gains (Brynjolfsson et al., 2025) or raise them by expanding output and enabling access to previously inaccessible brown resources (Bolton et al., 2024).[3] These opposing mechanisms complicate efforts to isolate net effects. Second, corporate disclosures on AI-driven emissions remain sparse and opaque. Important details, such as server investments and the business functions augmented by AI, are often proprietary and not available from sustainability reports, which hinders robust and comparable carbon footprint assessments.

To address the research gap, this paper applies natural language processing to annual

[1] For example, in PwC's 2024 Global CEO Survey, nearly half of the 4,701 executives identified AI integration as a top strategic focus for the next three years. See here.

[2] System-wide projection from IEA suggests data-center electricity use is set to more than double by 2030 and AI is the main driver. See here.

[3] See AI applications used to expand production capacity in carbon-intensive sectors here and here.

reports and LinkedIn profiles to distinguish firms between AI infrastructure developers and users, inspired by the regulatory categories in the EU Artificial Intelligence Act. Within this framework, we provide novel evidence on **(i)** the allocation of AI-related energy costs among adopters, and **(ii)** the significant but highly heterogeneous emission reduction effects of AI integration across user firms in diverse industries. First, firms that build or operate AI infrastructure bear the majority of the electricity consumption, with many partially offsetting related carbon footprints by using renewable electricity. Users show little change in electricity use as they outsource computation and shift the energy burden upstream in the supply chain. Second, AI user firms experience significant Scope 1 emissions reductions following adoption, with distributed lead-lag models and instrumental variable estimates supporting a causal interpretation. This effect, however, is concentrated among already low-emitting firms, whereas some brown sectors even experience increases. Reductions appear to stem from improved forecasting that optimizes resource allocation and cuts waste. This channel is largely absent in the most carbon-intensive sectors, where AI is sometimes deployed to maximize fossil fuel utilization and thus raises emissions. Lastly, these AI-related climate risks are beginning to influence portfolio decisions among climate-conscious investors.

We measure firm-level AI capabilities along two dimensions—infrastructure and talent—using regulated corporate disclosures as well as workforce profiles for a large global sample. For infrastructure, we analyze annual reports, which are legally required to disclose material capital expenditures such as data centers and computing facilities that indicate AI-related physical capacity. For talent, we draw on LinkedIn's individual-level employment data to identify AI-related workers and classify them as market-facing or internal-operations roles, depending on which part of the firm's activities they affect. Both measures are constructed via a two-step textual analysis: first, we apply keyword searches to filter potentially relevant text from large corpora; second, we assess the semantic context using a fine-tuned LLaMA model (Meta, 2024), which approximates GPT-4's classification accuracy while preserving the reproducibility and scalability of open-source models. Combining the infrastructure and

talent measures allows us to distinguish user firms from those that operate energy-intensive AI facilities.

Firms with AI infrastructure and talent differ markedly in their carbon profiles. Both have expanded over the past decade, with AI talent users expanding far more rapidly than infrastructure developers and growing fourfold from 7% in 2010 to 28% in 2024. Developers are concentrated among low-emission firms, with nearly 60% in the Computer Software and Electronic Equipment sector, which includes cloud providers such as Microsoft and hardware manufacturers like NVIDIA. In contrast, AI user firms are more evenly distributed across middle- and high-emission groups and span a wide range of industries like Retail and Medical Equipment. Carbon-intensive sectors such as Oil & Gas and Chemicals are also adopting AI, mainly for internal processes rather than customer-facing applications. Operational needs, instead of climate considerations, are the primary drivers of AI adoption: larger firms with higher capital expenditures and stronger AI-related competitive pressure are more likely to hire AI workers and internalize computing resources, while the presence of approved emission-reduction targets has little influence.

The energy burden of AI adoption falls primarily on firms with in-house infrastructure. These firms consume about 5% more electricity than industry peers in the first year after adoption, increasing to 12% in the second year. This change corresponds to an estimated aggregate increase of roughly 41 TWh over our sample in the past decade and is equivalent to over 1% of annual U.S. electricity consumption. In contrast, the share of AI workers has no significant effect, with coefficients close to zero. Infrastructure developers increasingly rely on renewable energy to *green* their reported emissions and decouple market-based Scope 2 emissions from actual electricity use. Nevertheless, these efforts only partially offset demand growth and some energy-hungry firms expand non-renewable electricity consumption by over 40% within two years. As a result, AI infrastructure development continues to place upward pressure on the reliability of electricity and the associated emissions.

A supplier-customer-level analysis further supports the disconnect between AI worker

integration and energy accountability. Firms with more AI workers are significantly more likely to source from infrastructure-intensive suppliers. Moreover, 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, with a t-statistic of 2.94. This suggests that pure AI users outsource their computational load (particularly to cloud or colocation providers) and effectively shift the energy burden upstream in the supply chain. Meanwhile, firms with more AI infrastructure disclosure tend to serve downstream customers with greater AI worker intensity and these customer firms do not exhibit higher electricity usage themselves.

More importantly, AI user firms experience subsequent reductions in Scope 1 carbon emissions after adoption. A one standard deviation increase in the share of AI-related workers is associated with a 4% decline in Scope 1 emissions in the following year and widens to 10% by the third year. Both the market-facing and internal AI roles contribute to this effect. The decomposition analysis shows that the reduction stems mainly from decreasing carbon intensity rather than changes in sales. Using a distributed lead-lag model to analyze the dynamic effect, we find no evidence of pre-trends and the observed effects are unlikely to be driven by reverse causality. The results remain robust when controlling for non-AI workers in the IT, robotics, and data analytics, as well as for green innovation engagement and emission reduction pledges. Finally, we instrument AI worker supply using both lottery-driven excess H-1B visa allocations and pre-existing connections to AI-strong universities, and reinforce the causal link between AI adoption and emission reductions.

Notably, the reduction effects are highly heterogeneous across sectors and currently concentrated among low-emitting firms. Several industries, such as Finance and Retail, show significant emission declines, but their low baseline levels make these changes less environmentally consequential in absolute terms. In contrast, the effects become more varied among carbon-intensive counterparts. While the Chemical sector records a negative effect, the Energy sector exhibits a positive coefficient, pointing to higher emissions. More strikingly, in the Utilities sector, which ranks highest in average Scope 1 emissions, AI adoption is linked

to a statistically significant 3.5% increase in emissions per one-standard-deviation rise in related worker share. Given its high baseline emissions, this represents a material increase in absolute carbon output relative to reductions in other sectors. Narrative evidence from job descriptions at estimation-influential firms supports the view that the greenness of AI applications is highly context-dependent, with industry-specific incentives and constraints shaping their direction of adoption.

AI adoption reflects directed technical change, with firms deploying AI to complement their current business models. In some brown sectors, AI adoption directs production and investments towards carbon-intensive activities with increasing emissions. By contrast, in low-emitting sectors, AI-enabled forecasting can lead to smarter resource allocation and tighter production planning, which lowers waste and Scope 1 emissions. The latter channel has limited leverage in brown sectors, where emissions primarily stem from combustion and core products are intrinsically high-carbon and difficult to substitute. In our sample, we find no evidence that AI lower emissions in these sectors via green innovation or renewable expansion; instead, it is sometimes deployed to maximize the capacity utilization of existing fossil assets. Unlike dedicated green technologies, AI is a general-purpose optimizer, making its emission impact path-dependent. Firms should therefore explicitly embed climate objectives into AI strategy, accounting not only for the direct energy use of AI computing but also for its broader impact on carbon-intensive activities.

As a final observation, climate-conscious investors, proxied by PRI signatories, show early signs of incorporating these AI-related climate risks into capital allocation decisions. They are more likely to divest from firms developing AI infrastructure powered by non-renewable electricity and from firms integrating AI workers while increasing carbon emissions. Non-PRI signatories show no consistent response.

Literature Review: This paper contributes to the growing literature on quantifying corporate AI exposure and linking it to firm operations and performance. Brynjolfsson et al.

(2025) show that generative AI enhance worker productivity. Eisfeldt et al. (2023) provide the first measure of firms' workforce exposure to it and document effects on firm value. Beyond generative AI, Babina et al. (2023, 2024a,b) provide key early measures of AI talent investment and analyze its effects on firm growth, workforce composition, and systematic risk from 2010 to 2018; Acemoglu et al. (2022) document related post-2010 changes in hiring and skill requirements; and Abis and Veldkamp (2023) provide structural estimates of how AI affects knowledge production. We extend this literature by shifting the lens from labor and financial outcomes of corporate AI adoption to its climate externalities by jointly considering the roles of infrastructure development and talent integration that reshape operations—an important yet underexplored dimension. Our analysis of carbon emissions as a byproduct of production also offers additional empirical evidence on AI's impact on corporate productivity, which is difficult to capture in broad-sector measures.

We build on studies of directed technical change and its environmental impacts (Acemoglu et al., 2012; Aghion et al., 2016; Casey, 2023). Acemoglu et al. (2016) show how policy mix steers the direction of innovation from "dirty" to "clean." Aghion et al. (2025) develop a growth theory in which economic activity endogenously shifts from material production to quality improvements with declining environmental intensity. Closely related, Bolton et al. (2024) study the role of green technology in shaping future corporate carbon emissions and uncover unintended consequences such as rebound effects and cross-sector displacement, while Hege et al. (2024) document emission reduction effects by analyzing supply chain networks. Distinct from their work, we focus on artificial intelligence, which is an emerging general-purpose technology with the potential to transform business operations across diverse industries. This transformative scope with the capacity to reshape operations and associated emissions gives AI unusually high climate salience and warrants systematic study. Although several studies have examined related environmental footprint, most focus narrowly on a single language model or provide qualitative assessments based on broad assumptions.[4] We

[4]See Luccioni et al. (2022), Schneider et al. (2025), Bogmans et al. (2025), Kaack et al. (2022), and more.

address this gap by combining various datasets with natural language processing to conduct the first large-sample, firm-level analysis of AI's real-world climate impact. By capturing both direct electricity use and broader operational emission outcomes, we offer critical early insights into how AI adoption has already shaped corporate climate trajectories.

We also relate to the literature on climate risks and investor responses in financial markets (Pastor et al., 2022; Hsu et al., 2023). Bolton and Kacperczyk (2021) show that investors demand compensation for carbon risk and some institutions divest from carbon-intensive sectors. Ilhan et al. (2023) document that climate-conscious institutional investors value climate risk disclosure, and Gibson et al. (2022) examine whether such investors allocate capital responsibly. We present a new angle by examining how institutional ownership varies with firm-level AI-related climate risks, a setting that attracts heightened attention from responsible investors.

2 Detecting Corporate AI Capabilities

Given that AI infrastructure deployment and workforce engagement represent distinct channels through which AI can affect firm operations and, in turn, carbon emissions, we construct separate measures for each of the two dimensions and detail their data sources in this section.

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 quantify firm-level AI adoption, we combine information from regulated corporate disclosures and workforce profiles. On the one hand, we identify investments in AI infrastructure by analyzing corporate 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. To extend the scope beyond U.S. firms, we include non-U.S. companies listed in the MSCI World Index that are not subject to SEC reporting requirements and manually collect the English version of their annual reports from company 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 necessary as many of them play a role in the diffusion and thus offer a more complete view of cross-border patterns in infrastructure-intensive AI adoption. On the other hand, we utilize worldwide individual-level employment data from LinkedIn, obtained through the Bright Initiative in October 2024. LinkedIn hosts more than one billion member profiles 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. We rely on Scope 1 emissions from MSCI due to its broad global coverage and consistent estimation methodology.

We obtain 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' 10-K filings. These narratives often include descriptions of related infrastructure like data centers or GPU procurement, offering valuable insights that are otherwise not captured in quantitative datasets. 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). For non-U.S. firms without 10-K filings, we analyze their full annual reports (as they do not follow the same standardized structure). To identify potential AI infrastructure discussions, we conduct an initial keyword-based search using terms designed to capture computing resources. 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 (LLM) 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 to 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 in our method) that acts as the student. The resultant student model is capable of approximating GPT-4's task-specific classification performance while remaining fully reproducible. Internet Appendix B provides the technical details of our NLP method and its performance evaluation, which gives us confidence in the validity of our

approach. Based on the fine-tuned model, 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. 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. 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 three-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 (see 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 U.S. resume data. The close alignment supports the validity of the constructed measure while with broader country coverage.

Moreover, we distinguish AI worker roles into two types: *market-facing* 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 and enhancing user experience, 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. We adopt this categorization because it reflects the operational structure of firms and aligns with the top business AI applications reported in company surveys.[5] It provides a potentially useful lens to examine if 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 fine-tuned large language model (see more details in Internet Appendix B). We correspondingly calculate shares of the two types of AI workers:

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

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

This step also helps filter out workers whose role cannot fall clearly into either category, often due to insufficient textual information in the position descriptions. The firm-year level AI

[5] See McKinsey Global Survey on artificial intelligence here.

worker engagement is quantified as:

$$AIWorker_{i,t} = \frac{\#Workers_{i,t}^{AI,\,Market} + \#Workers_{i,t}^{AI,\,Internal}}{\#Workers_{i,t}}.$$

2.2.3 Illustrative Examples

Table IA1 presents examples of AI infrastructure disclosures and workforce profiles, validating that our textual methods effectively capture corporate AI adoption across its life cycle. Panel A displays example paragraphs of infrastructure involvement from 10-K filings. Microsoft and Equinix report substantial investments in data center capacity and IT infrastructure, reflecting direct operational roles in delivering large-scale AI computing power. Other firms contribute to critical physical components enabling this infrastructure; for instance, NVIDIA details its specialized GPUs for AI model training. The other two panels shift focus to the deployment of AI talent during the use phase in diverse business functions. 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.

3 Allocation of AI Energy Costs Among Adopters

AI infrastructure is the main source of energy demand in AI adoption (IEA, 2025). As the computational backbone of model development, it directly contributes to electricity consumption. Importantly, while almost all AI-adopting firms require skilled workers, not all maintain their own infrastructure. Many of them rely on cloud-based or third-party services and thus effectively outsource the associated energy costs. Building on this fact and following the EU AI Act's emphasis on differentiating roles within the AI ecosystem, we propose a framework

that separates AI infrastructure developers from pure users, which allows for a systematic assessment of energy costs across firms. This section compares the climate profiles of these two types of adopters, assesses the magnitude and distribution of related electricity usage, and examines how this consumption is reflected in corporate emission reports.

3.1 Comparing AI Adopters in the Climate Context

Time Evolution: Figure 1 shows the evolution of corporate AI adoption from 2010 to 2024. The share of AI infrastructure developers—firms reporting material infrastructure involvement—remains relatively flat for most of the period but rises from 13% to 18% in 2024, driven by post-ChatGPT adoption. We identify these firms as having an above-median infrastructure disclosure in annual reports to avoid capturing random mentions. In contrast, AI user firms—those integrating AI talent without substantial infrastructure—grow steadily and accelerate from about 7% to 29% in the past decade. This upward trajectory aligns with three technological milestones: the 2012 deep learning breakthrough with AlexNet, the 2017 introduction of the transformer architecture, and the 2022 release of ChatGPT. It also coincides with increased climate awareness after the 2016 Paris Agreement, raising the possibility that climate considerations influence AI adoption (a question that we examine later). Among user firms, internal operational applications consistently outpace market-facing ones, indicating a broader integration into the organization beyond consumer interactions.

Emission Profiles: The distinction between developer and user firms is more pronounced when comparing Scope 1 carbon emissions in Figure 2. While developers are concentrated among the lowest-emitting firms, pure user firms are disproportionately found in medium and high-emission groups, with over 20% of them hiring AI workers. Panel B analyzes the average share of AI workers across use domains, excluding developer firms from the sample. Workers in market-facing applications are more common in lower-emission firms, whereas internal AI employees are distributed evenly across emission quintiles (0.7-1.0 basis points).

Figure IA1 reinforces the pattern through industry distribution using the Fama-French 12 classification. AI developers are heavily clustered in a few low-emission tech industries, most notably the Computer, Software, and Electronic Equipment sector, which accounts for over 57% of all developers. This sector includes cloud service providers such as Microsoft and hardware manufacturers like NVIDIA. Conversely, AI users span a broader set of industries, including carbon-intensive sectors such as Oil & Gas and Chemicals, where AI is often applied to internal operations. These patterns motivate an assessment of AI adoption's impact on subsequent Scope 1 emissions, particularly in high-emission industries where operational changes can significantly alter emissions.

Determinants: Table IA2 examines the determinants of AI capabilities and suggests that operational characteristics emerge as the primary drivers, with climate considerations playing a less clear role. We use infrastructure involvement as the dependent variable in Columns 1–6 and AI worker shares in Columns 7–10. Columns 1–2 show a strong link between infrastructure involvement and the presence of AI-skilled employees, especially in internal operations. Market-facing AI roles show no significant relationship, implying they are less compute-intensive or more easily outsourced to cloud providers. Column 3 indicates that firms facing greater AI-related competitive pressure (based on 10-K filings) tend to adopt infrastructure, likely to maintain a technological edge, whereas this effect is absent for employee adoption. In addition, electricity prices in the states where these firms are headquartered tend to be higher, indicating that local energy costs have limited influence on their compute-related energy demand. These states also exhibit relatively higher shares of renewable energy generation. Lastly, setting approved emission targets has an insignificant positive association with infrastructure involvement but a marginally significant one for employee adoption. In terms of financial characteristics, both adopters tend to be larger in size, exhibit higher capital investment, and operate with lower property tangibility.

3.2 Electricity Consumption

We examine the climate impact of AI infrastructure and worker adoption separately by estimating the following regression:

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

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 firm-level electricity consumption data from CDP and thus limits the sample size. Nevertheless, it includes major technology firms and semiconductor providers in the AI ecosystem (see Table IA3). We further restrict the sample to firms reporting emissions and energy use under the operational control boundary of the GHG Protocol, which provides the most accurate measure of a firm's own energy footprint. This choice is critical in the context of AI deployment. For instance, electricity consumption from leased data centers (a common arrangement among AI adopters) is reported by the lessee if the facility falls under their operational control. By contrast, under boundaries such as financial control, this usage would often go unreported unless the lessee holds majority financial ownership of the data center. In our sample, approximately 79% of firms use the operational control reporting standard. As shown in Panel A of Table 1, the 816 sample firms, on average, devote 0.23% of annual report paragraphs to AI infrastructure and employ 5.33 AI workers per 10,000

employees.

Table 2 presents the estimates. Columns 1–3 show that firms with AI infrastructure involvement consume significantly 5% more electricity than their industry peers in the following year, rising to 12% in the second year. Given the sample mean electricity use of 1,322,968 MWh (Panel A, Table 1), this corresponds to an additional 66,148 MWh per treated firm-year and roughly 41 TWh when aggregated over our sample firm-years with positive $\mathbb{1}(AIInfras)$ (622) in the past decade. While this back-of-the-envelope calculation could understate the true aggregate consumption as it only reflects firms with available CDP electricity data, the estimation remains economically meaningful and is equivalent to over 1% of annual U.S. electricity consumption. Similar patterns hold when using electricity intensity as the dependent variable in Columns 4–6, with an increase of 8% to 16% over the three-year horizon, all significant at the 1% level. In contrast, the share of AI workers shows no significant effect and the coefficients are near zero. This divergence suggests that physical infrastructure, rather than employee acquisition, drives the observed changes in energy demand.

3.3 Role of Renewable Electricity

We next examine how increased electricity usage translates into firms' reported carbon emissions by replacing the dependent variable with Scope 2 emissions, calculated using both location- and market-based accounting methods. As shown in Panel B of Table 2, location-based emissions, which reflect actual grid-average emission factors, rise significantly, mirroring the increase in electricity consumption. This indicates that the energy demands of AI infrastructure have a direct and measurable carbon footprint. Market-based Scope 2 emissions, which incorporate firms' contractual purchases of renewable energy, show a positive but less significant change. Firms appear to actively use 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 split total electricity consumption into renewable and non-renewable sources. Panel C shows a sharp rise in renewable use among AI infrastructure firms, from 44% to 57% over three years. However, this growth does not fully offset the overall increase in electricity use. Columns 4–6 reveal that non-renewable consumption also climbs, particularly over longer horizons, reaching 47% above baseline after two years. The larger percentage changes for both subtypes relative to Panel A arise from their smaller baseline levels, where similar absolute changes lead to bigger percentage shifts. Overall, AI-driven electricity usage, while increasingly sourced from renewables, still significantly raises non-renewable use and thus continues to exert upward pressure on energy demand.

3.4 Supply Chain Spillovers

Do firms with AI employees but no infrastructure indirectly bear the energy costs of AI adoption? Table 3 addresses this question by examining supply chain linkages and testing if firms with AI workers tend to source from suppliers that operate AI infrastructure. Using a supplier-customer-level regression with firm controls and supplier-customer pair fixed effects, we find that firms with greater AI worker shares 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, with a t-statistic of 2.94. This provides direct evidence that pure AI users outsource their computational load to cloud or colocation providers to effectively shift the energy burden upstream in the supply chain.

Consistently, firms with greater AI infrastructure disclosure, at the same time, are more likely to serve downstream customers with higher AI worker intensity. These customer firms do not exhibit higher electricity usage themselves, aligning with earlier findings that AI workers alone do not drive on-balance-sheet energy consumption. In fact, we observe a

marginally significant decline in their Scope 1 carbon emissions, which we explore further in the next section.

4 How AI Adoption Affects Operational Emissions

As suggested in the previous section, AI adoption influences corporate emissions not only through electricity use but also by changing operational processes. On the one hand, AI can raise productivity and operational efficiency, which may reduce emissions per unit of output (Brynjolfsson et al., 2025; Fuster et al., 2019). At the same time, by lowering marginal production costs, AI may expand output (a rebound effect known as the Jevons paradox, Jevons, 1865; Bolton et al., 2024) or enable firms to access previously inaccessible carbon-intensive resources (Kellogg, 2011; Acemoglu et al., 2012), both of which increase total emissions. These opposing mechanisms make the net impact of AI adoption on operational emissions theoretically ambiguous and require empirical investigation.

This section examines the relationship between AI adoption and Scope 1 carbon emissions and observes three key findings. First, on average, AI integration leads to significant reductions in Scope 1 emissions. Distributed lead-lag regressions show no evidence of emissions trends preceding AI adoption, and the results remain robust when controlling for other IT-related technologies and corporate green engagement. We strengthen causal interpretation by instrumenting AI adoption with excess H-1B visa selections and the connection to AI-strong universities. Second, the reduction effect is highly heterogeneous across sectors and currently concentrated among already low-emitting firms. Some carbon-intensive sectors, such as Energy and Utilities, even experience emission increases. Narrative evidence from job descriptions at estimation-influential firms supports that AI applications are highly context-dependent across industries. Third, we investigate the mechanisms behind this heterogeneity and find that the greenness of AI adoption reflects directed technical change, with brown firms tending to adopt AI for carbon-intensive use. For lower-emitting firms, AI-induced

emission reductions stem from improved forecasting, which enhances resource allocation and reduces waste. This abatement channel is largely absent in brown sectors, where emissions are driven primarily by combustion rather than operational inefficiencies.

4.1 Baseline Results

We use the share of AI workers as our baseline measure as it more directly captures how firms employ AI in their operations and extends coverage to all firms with workforce data, without requiring electricity or annual report information. To better isolate AI user firms, we follow Babina et al. (2024b) to exclude companies in the Computer Software and Electronic Equipment sector (Fama-French 12 classification), which is the primary sector involved in AI infrastructure. Panel B of Table 1 presents summary statistics for this sample. On average, roughly one in 10,000 employees works in AI, with a slightly higher share in internal operational roles, indicating that AI employees remain scarce among our global user sample. This share is smaller than in the Panel A sample with electricity data, which includes larger firms and the Computer Software sector, where developer firms employ substantially more AI workers. We estimate the regression following Bolton et al. (2024):

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

The dependent variable is Scope 1 carbon emissions in years t+1, t+2, and t+3. The coefficient of interest, β, measures the relationship between AI workforce adoption and future direct emissions. Controls include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, and the current-year carbon emissions. Firm fixed effects absorb all time-invariant firm characteristics, so β is identified from within-firm time-series variation. Industry-year fixed effects capture sectoral differences. Standard errors are clustered at the firm level.

Table 4 reports the baseline results. A one standard deviation increase in the share of AI-related workers is associated with a 4% reduction in Scope 1 emissions in the following

year. This effect intensifies over time, reaching an 10% reduction by the third year. Both market-facing and internal AI roles contribute to emission reductions, with slightly larger effects from the former roles. Table IA4 explores the heterogeneous effects across time and geography. The emission reduction effect is stronger after 2018, following the introduction of transformer models. Geographically, the effect is more pronounced among U.S. firms, though results remain statistically significant for firms in other regions. Table IA5 extends the analysis to other emission scopes. The reduction persists for Scope 2 emissions, while effects on Scope 3 are less clear. Given our interest in emissions outcomes driven by AI integration into business operations, the subsequent analysis focuses on Scope 1 emissions.

As a second step, we follow Bolton et al. (2024) to decompose total Scope 1 emissions into two components, i.e., emission intensity and sales, and use them as dependent variables in Panel B of Table 4. Columns 1–3 show that the decline in emissions is driven primarily by reductions in carbon intensity, with both the magnitude and statistical significance closely matching the results in Panel A. By contrast, we find no pattern for future sales in Columns 4–6. This indicates that the observed emission decline stems mainly from efficiency gains in carbon intensity, suggesting that AI adoption enhances operational efficiency with associated climate benefits.

Robustness Checks: Table IA6 demonstrates the robustness of our findings with respect to various confounding factors and alternative empirical specifications.

First, the estimated effects could be confounded by contemporaneous changes at AI 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 analytics workers. The coefficient on AI worker share remains robust across these specifications, indicating that the observed effects are specific to AI technology. 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, Science-Based Targets initiative (SBTi) or net-zero emission target status, 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 (Bolton and Kacperczyk, 2025). Panel C further controls for external factors such as local renewable energy availability, electricity prices, and corporate financial constraints (Whited and Wu, 2006). These adjustments obtain qualitatively and quantitatively similar results.

Next, we test the sensitivity of our results to alternative specifications. Panel D considers different AI adoption measures, including a binary indicator for the presence of AI workers, the logarithm of the total number of AI workers, and AI worker numbers scaled by sales. Panel E additionally uses country-year 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.

Dynamic Effects: We estimate the dynamic effects of AI adoption on Scope 1 emissions to address potential reverse causality. Given that firms typically adopt AI incrementally rather than through one-time investments, which makes standard event-study designs unsuitable, we employ a distributed lead-lag model that accommodates continuous variation in AI adoption (Aghion et al., 2020; Babina et al., 2024b):

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

where $\Delta AIWorkers_{i,t-k-1\to t-k}$ 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. Control variables and fixed effects follow the baseline specification. Each β_k captures the cumulative response of the emission 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 reports the coefficients from the lead-lag regressions and shows that Scope 1 emissions decline following AI adoption, with the effect intensifies over the next two to three years. Importantly, there is no evidence of pre-trends: firms that adopt AI have similar emissions to others in the years prior to adoption, and significant divergence occurs only afterward. This suggests that the results are unlikely to be driven by reverse causality or omitted variables.

Instrument Variable: To further address potential bias from unobserved shocks and strengthen causal inference, we employ an instrument variable approach. The instrument is excess H-1B visa selections for firm-year changes in AI worker supply (Cen et al., 2024; Chen et al., 2021; Dimmock et al., 2022). This choice rests on two facts: first, a large share of U.S. AI employees is foreign-born and typically requires an H-1B visa to work there legally; second, 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 year after 2014, while 2010 to 2013 only saw lotteries for a subset of filings (Chen et al., 2021; Dimmock et al., 2022); hence, this instrument variable estimate focuses 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 worker 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 AI workers, 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 firm receives eight selections, with the other none. Their $ExcessVisa$ values would be $+3$ and -5, respectively, to 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 5 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 the 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 workers. 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 university u is in the top 5% by number of AI researchers during 2005–2009. Table IA7 presents the estimates for the U.S. sample from 2010 to 2023. Consistent with Babina et al. (2024b), this instrument has a reliable first stage and show a robust reduction on carbon emissions over the subsequent two years; the third-year coefficient is negative but insignificant.

4.2 Sector Heterogeneity: Brown vs. Green Firms

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 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 at 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. Sectors are ranked by their average Scope 1 carbon emissions to highlight differences in baseline carbon emission levels. Interestingly, the sectors with low baseline emissions are often AI adopters with effective emission reductions. The Finance sector stands out as the most intensive AI adopter among user industries (see Figure IA1). A one standard deviation increase in its AI worker share corresponds to a 2.5% 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 Consumer Durables sector shows a high level of AI workforce hiring, with AI workers 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 emission impacts of AI adoption become more varied among carbon-intensive industries. In Chemicals, the negative but insignificant coefficient suggests AI may help reduce emissions in certain heavy processes. In Oil, however, the coefficient is positive, pointing to higher emissions. 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, a one

standard deviation increase in AI worker share is associated with a statistically significant 3.5% increase in emissions. Given the sector's massive baseline, this implies a considerable absolute increase in CO_2 output.

Table 6 compares firms in the Energy and Utilities sectors with those in other sectors by introducing an interaction term into the baseline regression. We find a positive and significant coefficient over the subsequent two years, indicating that emission reductions are substantially smaller—and in some cases reversed—in the two most carbon-intensive sectors. This effect is primarily driven by internal operational applications of AI.

Narrative Retrieval: As the effectiveness of AI adoption in mitigating emissions is highly industry-dependent and even counterproductive in certain sectors with inherently high carbon intensity, it is essential to understand the economic mechanisms behind the observed heterogeneity. One key advantage of our AI adoption measure is that it is derived from job descriptions and contains rich narratives about how firms actually integrate AI into operations. By mining this textual information, it has the potential to provide qualitative insight into why AI mitigates emissions in some industries but not others.

To do so, we identify the firms that exert the greatest influence on the regression estimates within each sector and analyze the job descriptions of their AI-related positions. Influence is measured with DFBETA, which captures the change in a coefficient when a single firm's observations are removed; large DFBETA values indicate firms that disproportionately shape sector-level results. We select the 50 firms with the highest DFBETA value in each sector and create word clouds to visualize the most frequently occurring terms in their AI worker descriptions. Figure 5 displays the word clouds for four sectors with diverse carbon profiles and AI integration, including Wholesale and Retail, Chemical, Oil and Gas, and Utilities. Figure IA2 displays the word clouds for the remaining eight sectors. We discuss these sector-specific potential channels in detail in the next subsection.

4.3 Economic Mechanisms Underlying Emission Heterogeneity

4.3.1 Directed Technical Change in AI Adoption

The core product and technology in brown industries are intrinsically high-carbon and hard to substitute in the short run, so AI can tighten operations but cannot, by itself, change the carbon content of the underlying process. This raises a natural question given frequent claims that AI can serve as a green technology: is AI in carbon-intensive sectors being used to substitute toward lower-carbon processes and outputs, or to accelerate the utilization of existing fossil assets? The word clouds for these sectors in Figure 5 point to AI being deployed primarily to raise the utilization of existing assets rather than to substitute away from them. In Oil and Gas sector (Panel C), job descriptions emphasize asset integrity and reliability functions, such as reservoir characterization and anomaly detection, which are used to plan drilling campaigns and lift throughput. In Utilities (Panel D), the AI role centers on power-plant operations, including predictive maintenance and real-time operations, aimed at boosting availability and capacity factors of thermal units.

Both the regression analyses and narrative descriptions from AI roles point to directed technical change in AI adoption, where firms deploy AI to complement current business models and generate heterogeneous emissions effects across sectors (Hémous and Olsen, 2021). To directly test this mechanism, we ask whether the emissions effect of AI is linked to within-firm reallocation toward higher- or lower-carbon segments. If AI is complementary to brown reallocation, AI-using firms that shift production or investment toward dirtier segments should exhibit larger emissions increases; conversely, shifts toward cleaner activities should amplify emissions reductions.

We measure segment reallocation using firm-year segment-level sales and capital expenditures from Compustat and map each segment to a carbon-intensity benchmark from the EPA's Supply Chain Greenhouse Gas Emission Factors by NAICS-6 industry.[6] Using these

[6]This database is based on environmentally extended input-output modeling that links economic activity

benchmarks, we calculate a sales-weighted carbon intensity per firm-year,

$$BrownShift_{i,t}^{Sale} = \frac{\sum_s Sale_{i,t}^s \times CI^s}{Sale_{i,t}} - \frac{\sum_s Sale_{i,t-1}^s \times CI^s}{Sale_{i,t-1}},$$

which captures the realized emissions profile implied by the firm's current sales composition across segments, and a capex-weighted carbon intensity,

$$BrownShift_{i,t}^{Capex} = \frac{\sum_s Capex_{i,t}^s \times CI^s}{Capex_{i,t}} - \frac{\sum_s Capex_{i,t-1}^s \times CI^s}{Capex_{i,t-1}},$$

which reflects the expected future profile implied by its current investment allocation. CI^s is the carbon intensity of segment s.

To assess whether brown sectors direct AI adoption towards brown uses, we employ *BrownShift* as the dependent variable in Columns 1–4 of Table 7. We observe a positive and statistically significant coefficient on the interaction between AI worker share and the Energy and Utilities sector dummy, suggesting AI acts as a directed technology. We then examine whether this direction of AI adoption explains the heterogeneous emission outcomes by estimating the following regression:

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

Columns 5–8 of Table 7 reports that the interaction terms for both sales and investment-based changes are positive and statistically significant, implying that the marginal effect of AI on emissions rises with *BrownShift*. Thus, when production or investment tilts toward higher-carbon segments (*BrownShift*>0), the abatement effect of AI integration weakens and can turn emission-increasing; when activity shifts toward cleaner segments (*BrownShift*<0), AI adoption delivers larger emission reductions. Columns 6 and 8 show that this amplification

in each sector to the associated supply-chain emissions in CO_2-equivalent terms. See the database here.

is concentrated in internal operation AI roles, consistent with complementarity between AI tools and operational processes. Overall, these results provide evidence of directed technical change: firms deploy and benefit from AI along the margin where it complements their evolving segment mix, with brown (green) reallocation directing AI toward brown (green) uses and producing heterogeneous emission outcomes.

As a case study, we examine fossil fuel utilization in the Energy sector. Specifically, we test whether AI is deployed to maximize capacity—for example, by optimizing drilling schedules. Such applications can increase fuel consumption and offset efficiency gains, and thus contribute to the higher emissions observed in certain firms. To proxy for capacity usage in core carbon-intensive activities, we use year-over-year changes in oil and natural gas, with higher values indicating greater utilization. Columns 1–4 of Table 8 show positive associations between AI adoption and capacity maximization of these fossil assets. Panel B reports positive and statistically significant interaction coefficients for both oil and natural gas capacity. Consistent with our mechanism, AI integration enables firms to push capacity harder with increased emissions, and this effect is driven by adoption in internal operations.

4.3.2 Emission Reduction in Low-Emitting Sectors

The channel through which AI can reduce emissions is by improving forecast accuracy and optimizing resource allocation (Farboodi and Veldkamp, 2022, 2023; Mihet and Philippon, 2019; 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 emission output because nearly all corporate decisions, such as how much to produce and when to maintain equipment, 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 outcomes, AI adoption lower the uncertainty of decision-

making and such efficiency gains directly translate into lower operational emissions.

This mechanism is supported by the narrative evidence in the word clouds shown in Figure 5. Panel A indicates that the Wholesale and Retail sector uses AI for predictive analytics in customer service. By analyzing historical transaction data, AI identifies customer preferences and enables these firms to tailor products more effectively and reduce waste from unsold or misallocated inventory. Interestingly, predictive applications appear in some carbon-intensive industries as well, such as Chemicals (Panel B), where AI prediction is applied to demand planning and supply chain optimization, which are functions that can reduce overproduction and associated emissions.

Precisely quantifying the prediction accuracy of AI in every domain is challenging due to data limitations. Here, we focus on one measurable and economically important form of prediction—demand forecasting (Corhay et al., 2025). We measure firm-year forecast accuracy as the absolute analyst sales forecast error divided by actual sales, following the method of Feng et al. (2009):

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

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. We use the mean value of analyst sales forecasts one year before the sales announcement as a proxy for managerial forecasting capability because analysts often quickly incorporate management guidance into their predictions (Cotter et al., 2006). To strengthen this link, we restrict the sample to firms with managerial guidance events recorded in Capital IQ Key Developments, to ensure analyst forecasts are more likely to embed management guidance. Columns 1–2 of Panel A in Table 9 show a positive and statistically significant association between AI adoption and analyst forecast accuracy. This pattern is consistent with AI improving firms' internal forecasting and, in turn, the information environment faced by

analysts, supporting our measure as a reasonable proxy for managerial forecasting skill.

We then replace *BrownShift* in the previous regression with *ForecastAccuracy*. Our coefficient of interest is β_1 on the interaction term, which captures whether the emission effect of AI integration strengthens with better demand forecasts. Columns 4–5 of Panel A report statistically significant and negative β_1 and indicate that AI reduces emissions more when forecasts are more accurate. This mechanism is straightforward, as sharper forecasts of sales and input needs enable tighter production planning, lower safety stocks, and fewer energy-intensive peaks, and thus cut emissions. By role, the effect is stronger in market-facing applications that directly interact with customers and refine the demand signal; the corresponding estimate for AI workers in internal operations is statistically indistinguishable from zero.

We assess if the prediction gains from AI adoption extend to the Energy and Utilities sectors. In Column 3, the interaction coefficient is negative and insignificant, suggesting no detectable improvement in sales-forecast accuracy among firms from the two sectors. Moreover, according to Column 6, the association between AI adoption and subsequent emissions is significantly positive within brown sectors (with a t-stat of 2.61), even in subsamples with relatively better sales-forecast performance (with a t-stat of 1.31). Overall, the forecasting channel appears weak in the two most carbon-intensive sectors, possibly because their emissions are dominated by combustion and process chemistry, which are fixed by production methods, rather than by operational wastes that improved prediction could mitigate.

5 Response of Climate Conscious Investors

The critical following question is whether investors recognize and respond to the climate impacts of corporate AI adoption. The earlier analysis highlights two dimensions of such influences. First, AI infrastructure can substantially raise energy demand, with its carbon footprint being mitigated through 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 brown sectors seeing an emission rise. In this section, we investigate if investors, particularly those who are climate-conscious, incorporate the two types of AI-related climate risks into their portfolio decisions.

Table 10, Panel A, analyzes investors' response to AI infrastructure involvement using the following regression, which explicitly considers the role of renewable electricity:

$$IO_{i,t+k} = \alpha + \beta_1 \mathbb{1}(AIInfras_{i,t}) \times Electricity^{Renew} + \beta_2 \mathbb{1}(AIInfras_{i,t}) + \beta_3 Electricity^{Renew} + \mathbf{X}_{i,t} + a_i + b_{j,t} + \epsilon_{i,t}, \quad k \in \{1, 2, 3\}.$$

We use institutional ownership as the dependent variable and distinguish 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 et al., 2022). We use the same fixed effects and control variables as in the baseline regression. Columns 1–3 report the results for PRI investors. The coefficients for AI infrastructure involvement are negative but only become statistically significant after two years, with such involvement associated with a 2.12% decline in PRI institutional ownership. Notably, the interaction terms between AI infrastructure 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, explicitly considering the role of Scope 1 emissions:

$$IO_{i,t+k} = \alpha + \beta_1 \, AIWorker_{i,t} \times CE^{S1} + \beta_2 \, AIWorker_{i,t} + \beta_3 \, CE^{S1} + \mathbf{X}_{i,t} + a_i + b_{j,t} + \epsilon_{i,t}, \quad k \in \{1, 2, 3\}.$$

Among PRI investors in Columns 1–3, the coefficient on AI worker share is positive and statistically significant 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 are responsive to climate risks associated with certain forms of AI adoption. They tend to divest from firms involved in AI infrastructure powered by non-renewable electricity and from firms integrating AI workers while increasing carbon emissions. Investor responses materialize rapidly and are strongest in the first year following such adoption, in contrast to the emission effects of AI that accumulate more gradually and intensify over time.

6 Conclusion

Using a framework that distinguishes between AI infrastructure developers and talent users, this paper provides early evidence on the climate impact of corporate AI adoption for a large global sample. The energy burden of AI adoption is concentrated among firms with in-house infrastructure, which tend to offset their carbon footprint by increasing renewable

electricity use. In contrast, pure user firms without such infrastructure show no significant rise in electricity consumption and sometimes even record declines. More importantly, AI integration among user firms is associated with immediate reductions in Scope 1 carbon emissions, with the effect widening over the three years following adoption. These reductions occur mainly because AI serves as a forecasting tool that improves prediction accuracy and resource allocation. However, these channels primarily benefit low-emitting firms whose emissions stem largely from operational waste rather than chemical combustion. In high-emission sectors such as Oil and Utilities, AI is sometimes used to maximize the capacity of existing fossil assets and lead to increasing emissions. Unlike true green technologies, AI is a general-purpose optimizer whose environmental benefits are path-dependent. Therefore, it is important to explicitly integrate climate considerations into corporate AI strategy, as the impact extends beyond computing energy demands to the broader operational and sectoral contexts of deployment, which can produce unintended and uneven emission outcomes.

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, Firm websites
$\mathbb{1}(AIInfras)$	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
$ElecIntensity$	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
$ExcessVisa$	Number of H1B visas a firm receives in the lottery in excess of the amount expected given its total registrations and the aggregate lottery selection rate for that fiscal year.	USCIS, OFLC
$AIUniSupply$	Share of the 2010 STEM workforce employed by the firm who graduated from universities in the top 5% for AI research during 2005-2009.	Babina et al. (2024b)
$ForecastAccuracy$	Negative of the absolute difference between actual sales and the mean or median analyst sales forecast made one year earlier, scaled by actual sales per share.	I/B/E/S
$LowCarbonIntangibility$	Negative change in a firm's sales- or capex-weighted carbon intensity across business segments.	Compustat Segments
$GreenOutput$:		
$GreenPatent$	Year-over-year change in the logarithm of green patent counts.	USPTO
$MethaneAbatement$	Logarithm of methane leakage in the previous year minus that in the current year.	EPA
$RenewableGeneration$	Change in the logarithm of renewable electricity generation.	EIA
$CapacityMaximization$:		
Oil	Change in the ratio of annual oil production to proved oil reserves.	Compustat Industry
$NaturalGas$	Change in the ratio of annual natural gas production to proved natural gas reserves.	Compustat Industry
$ElectricityGeneration$	Change in the ratio of total electricity generation to installed generation capacity.	EIA
IOR^{PRI}	Institutional ownership by PRI signatories.	S34, PRI
IOR^{NonPRI}	Institutional ownership by non-PRI signatories.	S34, PRI
$log(Asset)$	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

Abis, S. and Veldkamp, L. (2023). The changing economics of knowledge production. *Review of Financial Studies*, 37(1):89–118.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1):131–66.

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to clean technology. *Journal of Political Economy*, 124(1):52–104.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2020). What are the labor and product market effects of automation? new evidence from france. *Working Paper, SSRN 3547376*.

Aghion, P., Boppart, T., Peters, M., Schwartzman, M., and Zilibotti, F. (2025). A theory of endogenous degrowth and environmental sustainability. *Working Paper NBER33634*.

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.

Babina, T., Fedyk, A., He, A. X., and Hodson, J. (2023). Firm investments in artificial intelligence technologies and changes in workforce composition. *Working Paper, SSRN 4060233*.

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. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549.

Bolton, P. and Kacperczyk, M. (2025). Firm commitments. *Management Science*.

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

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *Quarterly Journal of Economics*, 140(2):889–942.

Casey, G. (2023). Energy efficiency and directed technical change: Implications for climate change mitigation. *Review of Economic Studies*, 91(1):192–228.

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*.

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.

Fuster, A., Plosser, M., Schnabl, P., and Vickery, J. (2019). The role of technology in mortgage lending. *Review of Financial Studies*, 32(5):1854–1899.

Gibson, 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*.

Hege, U., Li, K., and Zhang, Y. (2024). Climate innovation and carbon emissions: Evidence from supply chain networks. *Working Paper, SSRN 4557447*.

Hsu, P.-H., Kai, L., and Chi-Yang, T. (2023). The pollution premium. *Journal of Finance*, 78(3):1343–1392.

Hémous, D. and Olsen, M. (2021). Directed technical change in labor and environmental economics. *Annual Review of Economics*, 13(Volume 13, 2021):571–597.

IEA (2025). Energy and ai. *Report*.

Ilhan, E., Krueger, P., Sautner, Z., and Starks, L. T. (2023). Climate risk disclosure and institutional investors. *Review of Financial Studies*, 36(7):2617–2650.

Jevons, W. S. (1865). The coal question; an inquiry concerning the progress of the nation, and probable exhaustion of our coal mines. *London: Macmillan and Co.*

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*.

Kellogg, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch. *Quarterly Journal of Economics*, 126(4):1961–2004.

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

Meta (2024). The llama 3 herd of models. *arXiv:2407.21783*.

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*.

Pastor, , Stambaugh, R. F., and Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, 146(2):403–424.

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.

Whited, T. M. and Wu, G. (2006). Financial constraints risk. *Review of Financial Studies*, 19(2):531–559.

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.

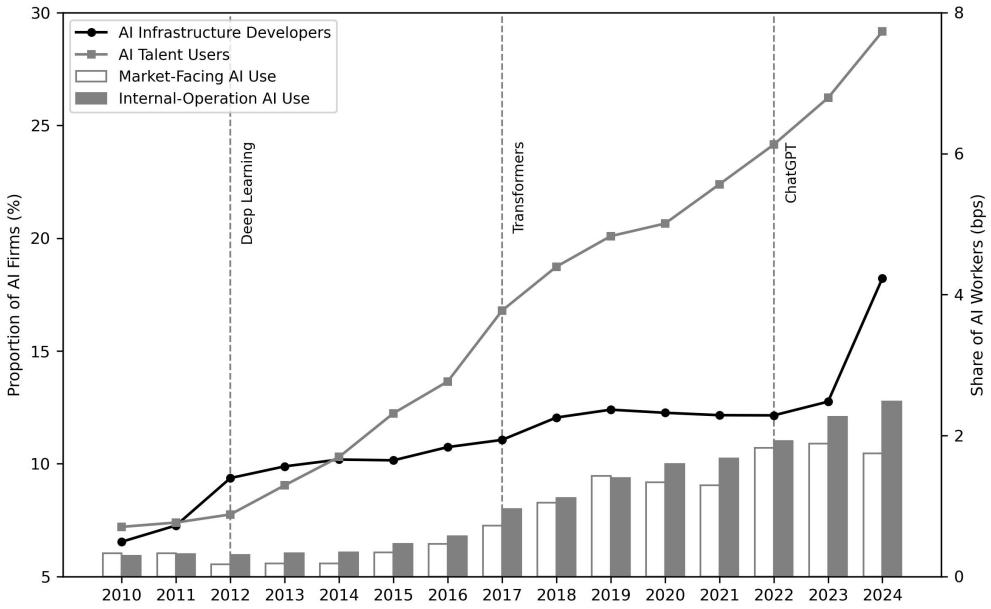
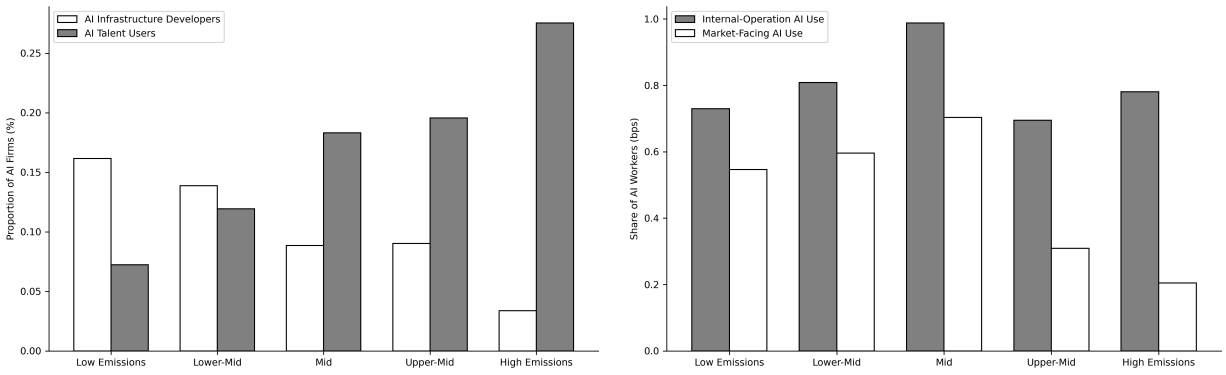


Figure 2: **Emission Profiles: Scope 1 Carbon Emissions**

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.



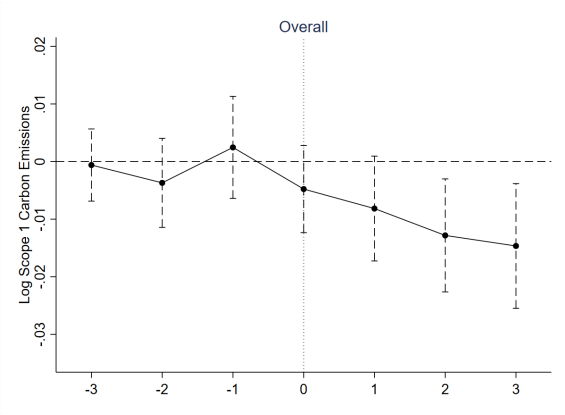
(a) (b)

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

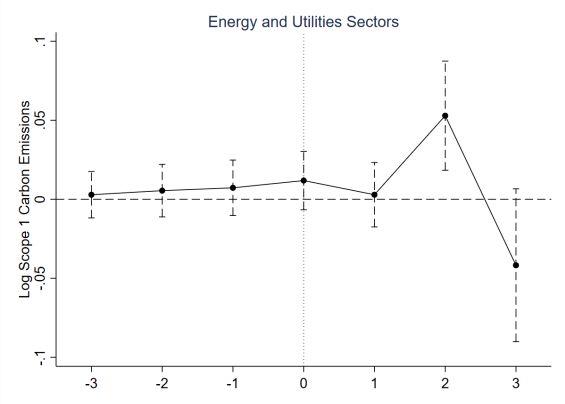
The figure plots the coefficients from the distributed lead-lag model, with vertical lines indicating 90% confidence intervals. The dependent variable is the logarithm of Scope 1 carbon emissions.

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

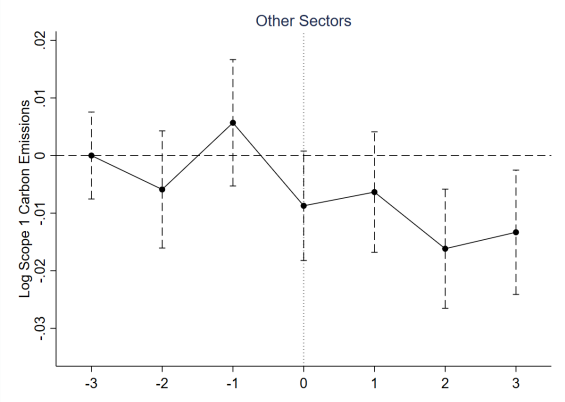
The sample period is from 2010 to 2023. Standard errors are clustered at the firm level.



(a)



(b)



(c)

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

This figure shows the effect of AI adoption on direct carbon emissions at the industry level, based on the Fama-French 12 classification. Black dots represent regression coefficients of future Scope 1 emissions on AI worker shares, with vertical lines indicating 90% confidence intervals. Gray bars indicate the average Scope 1 emissions for each industry.

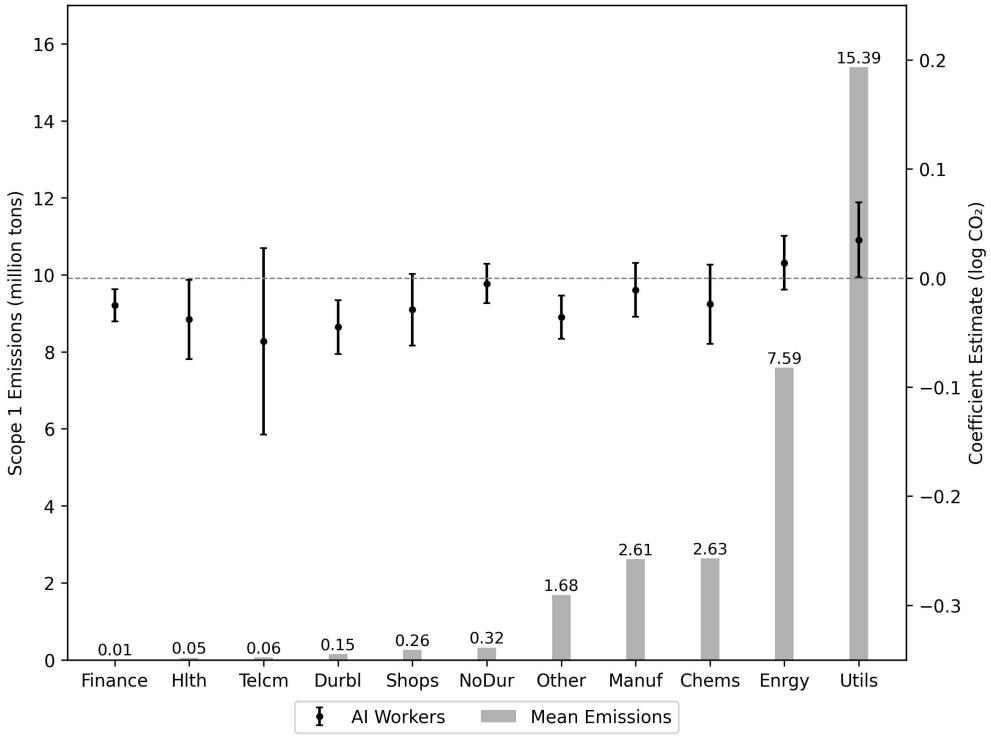


Figure 5: **AI Integration and Scope 1 Emissions: Narrative Retrieval**

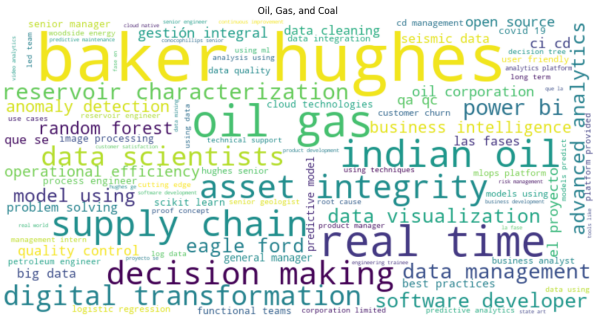
This figure displays word clouds of AI worker job descriptions from the 50 firms with the highest DFBETA influence in the sector-level regression of future carbon emissions on AI worker share, based on the Fama-French 12 classification. DFBETA measures how much the estimated regression coefficient would change if a specific firm were removed and thus identifies firms with the greatest influence on the estimated relationship. Panel A–D corresponds to Wholesale and Retail, Chemicals, Oil and Gas, and Utilities, respectively.



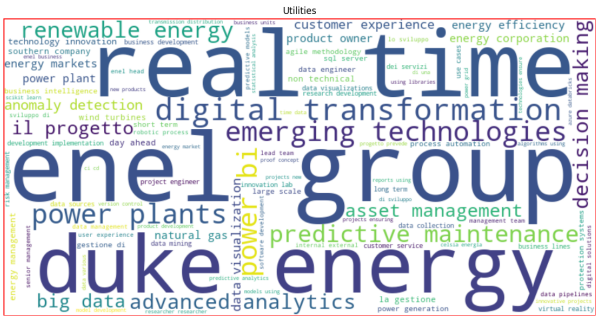
(a)



(b)



(c)



(d)

Table 1: **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 2: **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						
	$\Delta Electricity$			$\Delta ElecIntensity$		
	$t \to t$+1	$t \to t$+2	$t \to t$+3	$t \to t$+1	$t \to t$+2	$t \to t$+3
	(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
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Scope 2 Carbon Emissions						
	$\Delta CE^{S2(Loc)}$			$\Delta CE^{S2(Mkt)}$		
	$t \to t$+1	$t \to t$+2	$t \to t$+3	$t \to t$+1	$t \to t$+2	$t \to t$+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
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Panel C: Renewable vs. Non-renewable Electricity						
	$\Delta Electricity^{Renew}$			$\Delta Electricity^{NonRenew}$		
	$t \to t$+1	$t \to t$+2	$t \to t$+3	$t \to t$+1	$t \to t$+2	$t \to t$+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
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 3: **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
Supplier Control	Yes	Yes	Yes	Yes	Yes	Yes
Customer 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 4: **AI Integration and Scope 1 Emissions**

This table analyzes the relationship between future direct emissions and AI worker shares. Panel A uses the logarithm of Scope 1 carbon emissions in years t+1, t+2, and t+3 as the dependent variable; Panel B uses either log Scope 1 carbon intensity or log sales. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of carbon emissions. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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 Emissions						
	CE^{S1}					
	$t+1$		$t+2$		$t+3$	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	-0.04***		-0.07***		-0.10***	
	(-3.98)		(-4.06)		(-4.22)	
$AIWorker^{Market}$		-0.03***		-0.06***		-0.08***
		(-2.79)		(-2.59)		(-2.73)
$AIWorker^{Internal}$		-0.01*		-0.03*		-0.04**
		(-1.72)		(-1.91)		(-1.99)
N	29,363	29,363	26,240	26,240	23,200	23,200
R^2	0.977	0.977	0.968	0.968	0.966	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

Panel B: Carbon Intensity vs. Sales						
	CEI^{S1}			$Sales$		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	-0.03**	-0.05***	-0.08***	-0.01	-0.00	-0.01
	(-2.04)	(-3.03)	(-3.56)	(-0.70)	(-0.51)	(-0.89)
N	29,352	26,231	23,190	30,885	28,328	25,107
R^2	0.964	0.954	0.952	0.957	0.974	0.960
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: **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 a firm receives due to the lottery in year t-1 in Panel A and the AI-related selections, which are weighted by the share of visa applications for STEM positions, in Panel B. 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) 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				
	$\Delta AIWorker$	$\Delta CE^{S1}_{t\to t+1}$	$\Delta CE^{S1}_{t\to t+2}$	$\Delta CE^{S1}_{t\to t+3}$
	(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				
	$\Delta AIWorker$	$\Delta CE^{S1}_{t\to t+1}$	$\Delta CE^{S1}_{t\to t+2}$	$\Delta CE^{S1}_{t\to t+3}$
	(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 6: **AI Integration and Scope 1 Emissions: Brown vs. Other Sectors**

This table analyzes the relationship between future direct emissions and AI worker shares. The dependent variables are the logarithm of 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 carbon emissions. $EnergyUtil$ is a dummy variable that equals one for firms in the Energy or Utilities sectors. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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)
$AIWorker \times EnergyUtil$	0.05**		0.07**		0.06	
	(2.41)		(2.15)		(1.24)	
$AIWorker$	-0.04***		-0.08***		-0.10***	
	(-4.09)		(-4.08)		(-4.10)	
$AIWorker^{Market} \times EnergyUtil$		-0.14		-0.28		-0.31
		(-0.71)		(-0.97)		(-0.85)
$AIWorker^{Market}$		-0.03**		-0.05**		-0.07**
		(-2.24)		(-2.12)		(-2.38)
$AIWorker^{Internal} \times EnergyUtil$		0.04**		0.05**		0.04
		(2.30)		(2.02)		(0.92)
$AIWorker^{Internal}$		-0.02**		-0.04**		-0.04*
		(-2.18)		(-2.12)		(-1.91)
N	29,363	29,363	26,240	26,240	23,200	23,200
R^2	0.977	0.977	0.968	0.968	0.966	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

Table 7: **AI Integration and Scope 1 Emissions: Economic Mechanism**

This table examines how AI integration affects emissions as a directed technology. We measure *BrownShift* as the year-over-year changes in sales-weighted and capital expenditure-weighted carbon intensity across business segments to capture firms' shift toward browner production or investments. Columns 1–4 use *BrownShift* as the dependent variable and Columns 5–8 use the logarithm of Scope 1 carbon emissions in years $t+1$ as the dependent variable. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of carbon emissions. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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.

	$BrownShift$				CE^{S1}_{t+1}			
	SalesShift		CapexShift		SalesShift		CapexShift	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$AIWorker \times EnergyUtil$	0.43*		0.58**					
	(1.74)		(2.09)					
$AIWorker \times BrownShift$					0.04***		0.04**	
					(3.43)		(2.24)	
$AIWorker$	0.00		0.00		-0.07***		-0.06**	
	(0.20)		(0.03)		(-4.41)		(-2.53)	
$AIWorker^{Market} \times EnergyUtil$		1.25		-1.85				
		(0.49)		(-1.09)				
$AIWorker^{Market} \times BrownShift$						0.00		0.07
						(0.10)		(0.96)
$AIWorker^{Market}$		0.01		0.03		-0.05***		-0.04**
		(0.39)		(1.08)		(-3.15)		(-2.02)
$AIWorker^{Internal} \times EnergyUtil$		0.12		0.47**				
		(0.44)		(2.15)				
$AIWorker^{Internal} \times BrownShift$						0.02***		0.01**
						(3.02)		(2.34)
$AIWorker^{Internal}$		-0.00		-0.03		-0.03		-0.03
		(-0.21)		(-0.60)		(-1.58)		(-1.12)
$BrownShift$					0.00	-0.00	0.01*	0.01
					(0.19)	(-0.44)	(1.77)	(1.15)
N	12,421	12,421	9,434	9,434	11,184	11,184	8,502	8,502
R^2	0.153	0.153	0.136	0.136	0.982	0.982	0.982	0.982
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 8: **AI Integration and Scope 1 Emissions: Fossil Fuel Utilization**

This table analyzes the relationship between future direct emissions and AI worker shares, with a focus on the utilization of fossil fuels. We measure capacity maximization as the year-over-year change in oil production scaled by proved developed reserves in Columns 1–2/5–6 and natural gas production scaled by proved developed reserves in Columns 3–4/7–8. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year value of carbon emissions. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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.

	$CapacityMaximization$				CE^{S1}_{t+1}			
	Oil		NaturalGas		Oil		NaturalGas	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$AIWorker \times CapacityMaximization$					0.01*		0.06***	
					(1.97)		(3.92)	
$AIWorker$	3.85		0.24*		-0.02		-0.06**	
	(1.10)		(1.81)		(-0.70)		(-2.63)	
$AIWorker^{Market} \times CapacityMaximization$						-0.02		-0.02
						(-0.64)		(-0.93)
$AIWorker^{Market}$		1.60		0.04		-0.03		-0.02
		(1.02)		(0.42)		(-0.93)		(-1.72)
$AIWorker^{Internal} \times CapacityMaximization$						0.02**		0.05***
						(2.44)		(5.19)
$AIWorker^{Internal}$		3.65		0.38**		-0.04		-0.06
		(1.17)		(2.32)		(-1.52)		(-1.68)
$CapacityMaximization$					0.00**	0.00	0.01	0.01
					(2.21)	(0.97)	(1.41)	(0.87)
N	186	186	174	174	160	160	149	149
R^2	0.065	0.065	0.261	0.264	0.946	0.947	0.962	0.962
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 9: **AI Integration and Scope 1 Emissions: Other Sectors**

This table analyzes the relationship between future direct emissions and AI worker shares, with a focus on AI as a prediction tool to improve demand forecasting. We measure predictive accuracy as the negative of the absolute sales forecast error divided by actual sales. 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) 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.

	$ForecastAccuracy$			CE^{S1}_{t+1}		
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times ForecastAccuracy$				-0.04**		-0.05**
				(-2.04)		(-2.19)
$AIWorker$	0.02**		0.02**	-0.05***		-0.05***
	(1.98)		(1.98)	(-4.69)		(-4.81)
$AIWorker^{Market} \times ForecastAccuracy$					-0.04*	
					(-1.89)	
$AIWorker^{Market}$		0.02			-0.04***	
		(1.62)			(-3.07)	
$AIWorker^{Internal} \times ForecastAccuracy$					0.00	
					(0.09)	
$AIWorker^{Internal}$		0.01			-0.01	
		(0.92)			(-1.43)	
$AIWorker \times ForecastAccuracy \times EnergyUtil$						0.09
						(1.31)
$AIWorker \times EnergyUtil$			-0.02			0.06***
			(-0.85)			(2.61)
$ForecastAccuracy \times EnergyUtil$						0.02
						(0.59)
$ForecastAccuracy$				-0.03***	-0.02***	-0.03***
				(-5.31)	(-4.59)	(-5.35)
N	25,934	25,934	25,934	23,446	23,446	23,446
R^2	0.095	0.095	0.095	0.978	0.978	0.978
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-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 (PRI signatories in Columns 1–3) and non-climate-conscious investors (non-PRI signatories in 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) 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)
$\mathbb{1}(AIInfras) \times Electricity^{Renew}$	0.17*	0.19**	0.15*	-0.19	0.02	0.15
	(1.74)	(1.98)	(1.89)	(-1.11)	(0.13)	(1.07)
$\mathbb{1}(AIInfras)$	-1.69	-2.12*	-2.38**	-0.70	-0.64	-2.00
	(-1.35)	(-1.71)	(-2.26)	(-0.36)	(-0.28)	(-1.10)
$Electricity^{Renew}$	-0.03	-0.05	-0.04	-0.03	-0.01	-0.01
	(-0.97)	(-1.35)	(-1.11)	(-0.45)	(-0.25)	(-0.11)
N	1,918	1,482	1,156	1,918	1,482	1,156
R^2	0.861	0.859	0.850	0.763	0.810	0.800
Control	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-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)
$AIWorker \times CE^{S1}$	-0.27***	-0.17**	-0.12	0.11	0.11	0.06
	(-4.33)	(-2.23)	(-1.36)	(0.91)	(1.00)	(0.63)
$AIWorker$	1.53***	0.87	0.75	-1.42	-1.36	-0.94
	(3.07)	(1.38)	(1.05)	(-1.46)	(-1.56)	(-1.10)
CE^{S1}	-0.37**	-0.37**	-0.38**	0.04	0.10	0.05
	(-2.07)	(-2.17)	(-2.41)	(0.10)	(0.32)	(0.18)
N	13,406	12,153	10,863	13,406	12,153	10,863
R^2	0.842	0.835	0.833	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

Internet Appendix

for

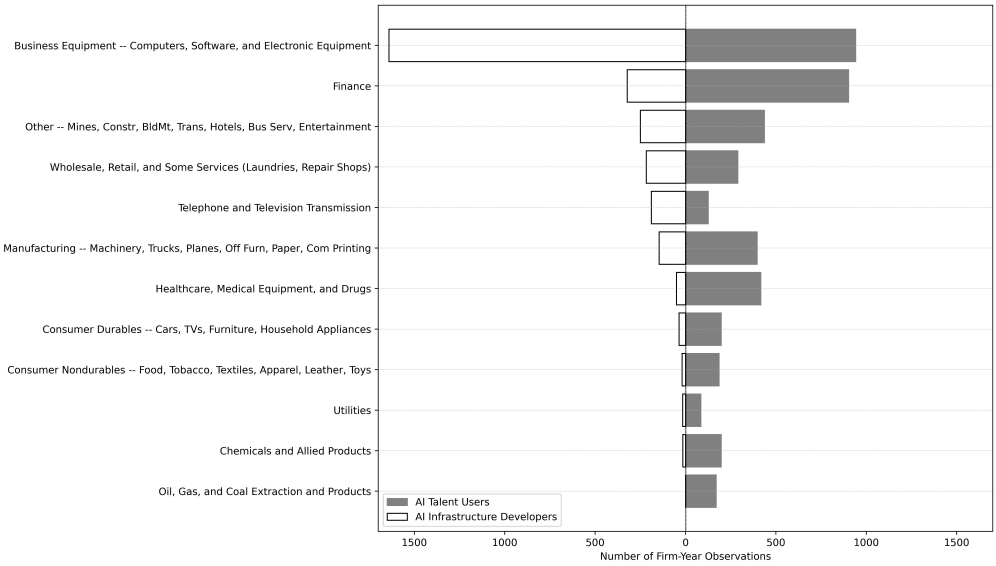
Carbon in the Cloud

This Internet Appendix provides supplementary material for this study.

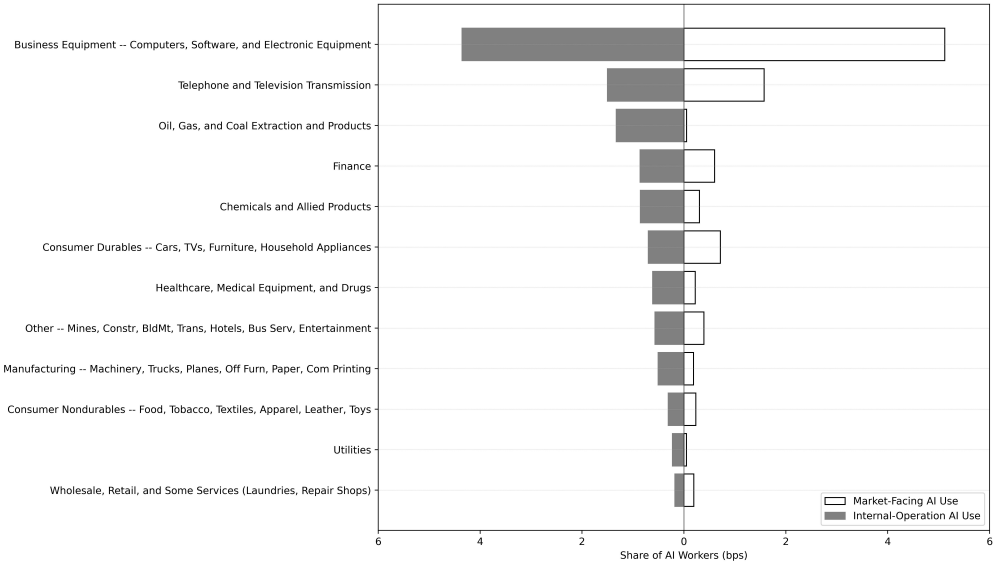
A Additional Figures and Tables

Figure IA1: **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 number 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.



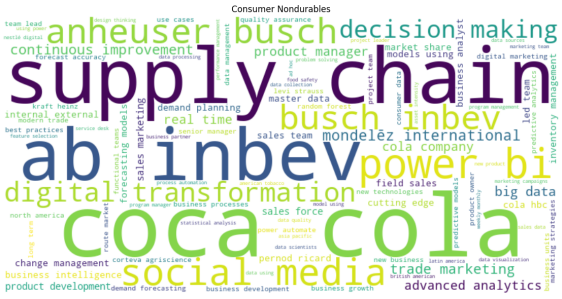
(a)



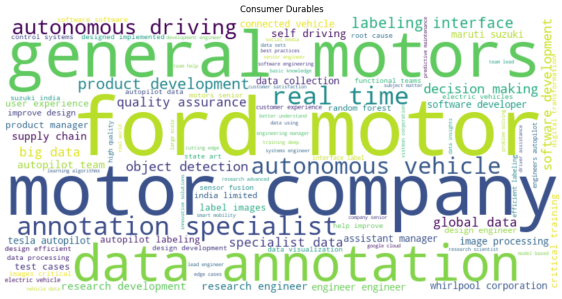
(b)

Figure IA2: **AI Integration and Scope 1 Emissions: Narrative Retrieval**

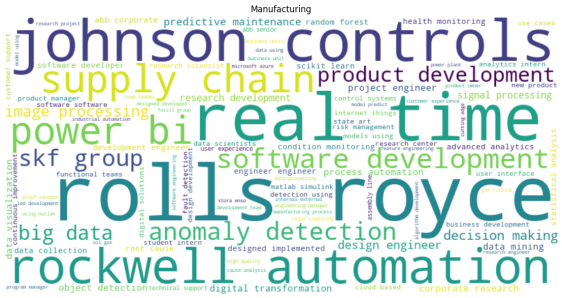
This figure displays word clouds of AI worker job descriptions from the 50 firms with the highest DFBETA influence in the sector-level regression of future carbon emissions on AI worker share, based on the Fama-French 12 classification code. DFBETA measures how much the estimated regression coefficient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



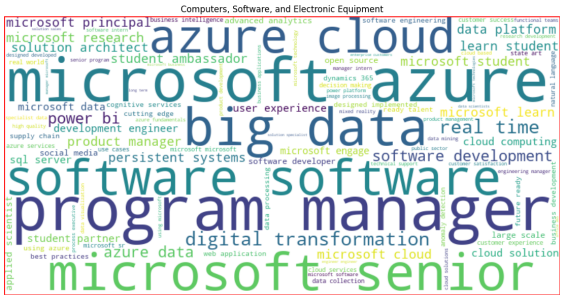
(a)



(b)



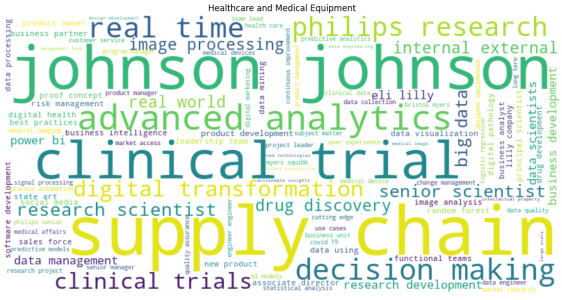
(c)



(d)



(e)



(f)



(g)



(h)

Table IA1: **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 IA2: **Determinants of AI Adoption**

This table examines the determinants of firms' decisions to adopt AI infrastructure and workers. The dependent variables are $\mathbb{1}(AIInfras)$, a dummy variable equal to one if the firm has substantial AI infrastructure discussions (above-median) in annual reports, in Columns 1–4 and the share of AI workers in Columns 5–8. *CompetePressure* represents the proportion of 10-K paragraphs mentioning both competition- and technology-related terms. *ElectricityPrice* is the annual average commercial electricity price at the state level from the EIA. *RenewableGrids* is the share of renewable electricity in local grids at the state level (inside the U.S.) or country level (outside the U.S.) from the EIA. *EmissionTarget* is an indicator for firm-years with approved science-based or net-zero emission targets from the MSCI. 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.

	$\mathbb{1}(AIInfras)$				$AIWorker$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CompetePressure	0.01**				0.04			
	(2.30)				(1.57)			
ElectricityPrice		0.02***				0.06**		
		(3.35)				(2.44)		
RenewableGrids			0.01**				0.05**	
			(2.56)				(2.47)	
EmissionTarget				0.00				0.05*
				(0.14)				(1.89)
log(Asset)	0.01***	0.01**	0.01***	0.01*	0.09***	0.08***	0.07***	0.07***
	(2.69)	(2.45)	(2.92)	(1.94)	(8.05)	(7.88)	(11.06)	(8.32)
ROA	0.01	0.03	0.02	0.10	0.03	0.00	0.02	0.04
	(0.32)	(0.69)	(0.60)	(0.94)	(0.25)	(0.02)	(0.35)	(0.26)
Capex/Assets	0.56***	0.55***	0.57***	0.74***	1.49***	1.42***	0.61***	0.91***
	(3.28)	(3.63)	(3.98)	(2.79)	(4.54)	(4.94)	(3.81)	(3.18)
Leverage	0.00	0.01	0.00	-0.00	-0.11*	-0.09	-0.13***	-0.18**
	(0.03)	(0.30)	(0.02)	(-0.07)	(-1.68)	(-1.64)	(-2.99)	(-2.25)
Tangibility	-0.10**	-0.11**	-0.12***	-0.14**	-0.62***	-0.62***	-0.43***	-0.52***
	(-2.24)	(-2.53)	(-3.18)	(-2.17)	(-4.72)	(-5.14)	(-6.79)	(-4.68)
SalesGrowth	0.02**	0.02**	0.01**	0.02	-0.05***	-0.05***	-0.02	0.01
	(2.45)	(2.29)	(2.14)	(1.37)	(-3.10)	(-3.68)	(-1.62)	(0.26)
N	15,699	17,325	20,313	9,166	13,068	15,015	31,018	16,637
R^2	0.218	0.226	0.218	0.273	0.153	0.136	0.161	0.208
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA3: **Key AI Adopters with Electricity Consumption Data**

This table lists the 30 largest AI-adopting firms in our sample with electricity consumption data from CDP, 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 IA4: **AI Integration and Scope 1 Emissions: Time and Geographic Variation**

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) 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***	-0.10***			
	(-4.20)	(-4.10)	(-4.04)			
$AIWorker \times \mathbb{1}(Before2018)$	-0.02	-0.05**	-0.10***			
	(-1.02)	(-2.03)	(-3.24)			
$AIWorker \times \mathbb{1}(US)$				-0.05***	-0.09***	-0.11***
				(-3.51)	(-3.80)	(-3.90)
$AIWorker \times \mathbb{1}(NonUS)$				-0.03**	-0.05**	-0.08**
				(-2.47)	(-2.31)	(-2.49)
N	29,363	26,240	23,200	29,363	26,240	23,200
R^2	0.977	0.968	0.966	0.977	0.968	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

Table IA5: **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) 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***
	(-2.80)	(-3.13)	(-3.19)
N	29,456	26,330	23,280
R^2	0.958	0.940	0.934
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Panel B: Scope 3 Carbon Emissions			
	CE^{S3}		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker$	-0.09	-0.15*	-0.09
	(-1.57)	(-1.84)	(-0.95)
N	9,398	7,816	6,479
R^2	0.843	0.812	0.806
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes

Table IA6: **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. *ITWorker, RobotWorker, DataWorker* are the share of employees in related roles based on the LinkedIn profiles. *GreenPatents* is the share of green patents. *CCExposure* is the firm-level climate change exposure derived from earnings calls provided by Sautner et al. (2023). WW is the firm-level measure of financial constraints based on Whited and Wu (2006). The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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***	-0.09***	-0.04***	-0.07***	-0.10***	-0.04***	-0.07***	-0.09***
	(-3.20)	(-3.75)	(-3.89)	(-3.99)	(-4.17)	(-4.43)	(-3.79)	(-3.96)	(-4.15)
ITWorker	-0.05**	-0.03	-0.05						
	(-2.41)	(-1.16)	(-1.40)						
RobotWorker				0.01	0.02	0.03*			
				(0.54)	(1.46)	(1.82)			
DataWorker							-0.03**	-0.03	-0.03
							(-2.45)	(-1.41)	(-1.02)
Panel B: Confounding Factors - Green Engagement									
AIWorker	-0.04***	-0.08***	-0.10***	-0.03**	-0.06***	-0.08**	-0.05***	-0.08***	-0.11***
	(-3.01)	(-3.24)	(-3.46)	(-2.36)	(-2.83)	(-2.56)	(-4.50)	(-4.32)	(-4.96)
GreenPatent	-0.01**	-0.01	-0.01						
	(-2.31)	(-0.78)	(-1.64)						
EmissionTarget				-0.00	-0.04*	-0.08***			
				(-0.22)	(-1.88)	(-3.19)			
CCExposure							-0.00	0.01	0.01
							(-0.15)	(0.59)	(0.88)
Panel C: Confounding Factors - Local Energy Source, Energy Price and Financial Constraints									
AIWorker	-0.04***	-0.07***	-0.09***	-0.04***	-0.07***	-0.10***	-0.04***	-0.07***	-0.09***
	(-4.00)	(-4.07)	(-4.12)	(-2.86)	(-2.91)	(-3.18)	(-3.91)	(-3.94)	(-4.10)
ElectricityRenewable	-0.02	-0.00	0.00						
	(-1.19)	(-0.13)	(0.03)						
ElectricityPrice				0.01	0.02	-0.01			
				(0.49)	(0.44)	(-0.12)			
WW							0.00	-0.00	-0.00
							(0.37)	(-0.19)	(-0.97)
Panel D: Alternative AI Worker Measures									
$I\{AIWorker\}$	-0.03*	-0.07***	-0.09***						
	(-1.80)	(-2.82)	(-2.93)						
$log(\#AIWorker)$				-0.02***	-0.05***	-0.06***			
				(-3.72)	(-4.79)	(-4.83)			
$\#AIWorker/Sales$							-0.02***	-0.04***	-0.05***
							(-3.69)	(-3.79)	(-3.73)
Panel E: Alternative Specifications									
	Plus Country-Year FE			SE Cluster by Firm and Year			SE Cluster by Industry		
AIWorker	-0.04***	-0.07***	-0.10***	-0.04***	-0.07***	-0.10***	-0.04***	-0.07***	-0.10***
	(-3.89)	(-3.98)	(-4.12)	(-3.84)	(-3.78)	(-3.78)	(-4.98)	(-5.66)	(-5.07)

Table IA7: **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 the share of the 2010 STEM workforce employed by the firm who graduated from universities in the top 5% for AI research during 2005-2009, provided by Babina et al. (2024b). 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 AI Capability Detection

B.1 NLP Method Description

For both AI infrastructure disclosure and worker detection, we implement a two-step procedure. In the first step, we conduct a keyword search to locate potentially relevant passages, which reduces the large text corpus to a targeted subsample. In the second step, we assess the semantic context of each candidate passage using a fine-tuned LLaMA model (Meta, 2024), which perform three classification tasks independently: (i) identifying whether a 10-K paragraph contains AI infrastructure disclosure; (ii) determining whether a LinkedIn profile indicates AI use in market-facing functions; and (iii) determining whether a LinkedIn profile indicates AI use in internal operations. The semantic evaluation minimizes false positives by retaining only passages that substantively address related topics based on their full contextual meaning rather than isolated keywords.

We fine-tune the LLaMA model within a teacher-student learning framework (Ni et al., 2024; Dettmers et al., 2023; Wadhwa et al., 2024). Specifically, we employ GPT-4.1 as the high-capacity *teacher* model to generate labels for 1,500 randomly selected training samples for each task. These samples are drawn from the pool of candidate texts obtained through the initial keyword-based filtering and are purposefully balanced across positive and negative cases (with positive samples numbering 759, 577, and 710 for the three tasks, respectively). For each sample, GPT-4.1 provides both a classification label and a one-sentence rationale for its decision. We then present the same prompts and training samples, together with GPT-4.1's labels and rationales, to fine-tune a smaller open-source model (i.e., Meta-LLaMA-3.1-8B-Instruct) as the *student*.[1] This approach enables the student model to learn not only the teacher's classifications but also the underlying reasoning, and thus transfers GPT-4.1's domain-specific decision logic in a structured and interpretable way.

[1]We conduct a full fine-tune without adapters on an A100 GPU with three training epochs.

Table IA B1 reports the performance of GPT-4.1 and the fine-tuned LLaMA3 model on a test set of 500 manually annotated samples per task. The test samples are constructed in the same manner as the training data (with 319, 240, and 240 positive samples across the three tasks). GPT-4.1 delivers consistently high performance for all tasks, which suggests its suitability to serve as a teacher; meanwhile, the fine-tuned LLaMA3 model closely approximates GPT-4.1's performance. For AI infrastructure classification, GPT-4.1 achieves an F1-score of 0.98, while the fine-tuned LLaMA3 reaches a comparably high score of 0.97. Similarly, GPT-4.1 records F1-scores of 0.96 and 0.97 for the two types of worker roles, whereas the fine-tuned LLaMA3 achieves 0.94. These results demonstrate that LLaMA3 effectively learns from GPT-4.1 and transfers its classification ability to the three tasks.

Importantly, the fine-tuned LLaMA3 model provides two key advantages. First, it is fully reproducible. Because LLaMA can be run locally, it always produces identical outputs for identical inputs when setting the parameter temperature to zero. In contrast, GPT-4.1's outputs may vary slightly across runs (even with temperature fixed at zero) due to stochasticity in cloud-based inference. Second, the fine-tuned model is substantially more computationally efficient, enabling large-scale annotation at less than 2.5 percent of GPT-4.1's runtime for our tasks.

Table IA B1: **Performance of Fine-Tuned LLaMA3 Models**

	Accuracy	Precision	Recall	F1
Task 1: AI Infrastructure Involvement				
GPT4.1 (teacher)	0.9780	0.9843	0.9812	0.9828
Fine-tuned LLaMA3 (student)	0.9640	0.9809	0.9625	0.9716
Task 2: Market-Facing AI Use				
GPT4.1 (teacher)	0.9620	0.9542	0.9662	0.9602
Fine-tuned LLaMA3 (student)	0.9400	0.9224	0.9536	0.9378
Task 3: Internal-Operation AI Use				
GPT4.1 (teacher)	0.9680	0.9750	0.9590	0.9669
Fine-tuned LLaMA3 (student)	0.9420	0.9654	0.9139	0.9389

B.2 Keywords for Preliminary Detection

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.

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.3 Prompts for GPT-4 and LLaMA-3 Fine-Tuning

```
*** Task 1: AI Infrastructure Involvement ***
You are a financial 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 indirectly or without explicit detail.
Answer 1 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 0 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.
Return exactly a valid JSON object in this format:
"AI Infrastructure": 1 or 0, "Explanation": one-sentence explanation highlighting
the key elements that influenced your decision.

*** Task 2: Market-Facing AI Use ***
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, corporate 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.

*** Task 3:  Internal-Operation AI Use ***
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.[2] 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.

[2]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.[3] 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.[4] 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

[3]See the data that Bloomberg requested from USCIS here.
[4]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.[5] 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

[5] 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.

IA References

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

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.

Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. (2023). Qlora: Efficient finetuning of quantized llms. *arXiv:2305.14314*.

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.

Meta (2024). The llama 3 herd of models. *arXiv:2407.21783*.

Ni, J., Schimanski, T., Lin, M., Sachan, M., Ash, E., and Leippold, M. (2024). Diras: Efficient llm annotation of document relevance in retrieval augmented generation. *arXiv:2406.14162*.

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

Wadhwa, S., Amir, S., and Wallace, B. C. (2024). Investigating mysteries of cot-augmented distillation. *arXiv:2406.14511*.

Whited, T. M. and Wu, G. (2006). Financial constraints risk. *Review of Financial Studies*, 19(2):531–559.