Carbon in the Cloud

Tingyu Yu*

January 2026

Abstract

This paper studies how AI investment affects corporate emissions in a large global sample. Firms with higher AI worker shares subsequently reduce Scope 1 emissions on average, but effects are heterogeneous across sectors. Emissions fall broadly through improved forecasting and operational efficiency, whereas Energy and Utilities see emissions increase as AI is disproportionately used to expand carbon-intensive output rather than accelerate the green transition. Instrumental variable estimates using H-1B visa lotteries support a causal interpretation. In the aggregate, the emissions-weighted effect is close to zero, suggesting that brown AI use in a small set of high-emitting sectors offsets emission reductions elsewhere. Estimated emissions from electricity used for AI computing remain modest relative to these operational responses.

Keywords: AI workers, carbon emissions, directed technical change, climate risk
JEL classifications: Q54, O33, G11, D24

*University of Zurich, Department of Finance. Email: tingyu.yu@df.uzh.ch.
I am deeply grateful to Markus Leippold, Kai Li, and Zacharias Sautner for their continued guidance. I thank Francisco Amaral, Olimpia Carradori, Keith Chan (discussant), Glen Gostlow, Michel Habib, David Hémous, Weiwei Hu, Felix Kübler, Felix von Meyerinck, Jingwei Ni, Steven Ongena, Per Östberg, Stefano Ramelli (discussant), Tobias Schimanski, Chi-Yang Tsou, Alexander Wagner, Yucheng Yang, Jun Yu, Xiao Zhao, and participants at the UZH Rising Scholar Conference in Finance 2025, University of Zurich, University of Manchester, Imperial CEP, the HEC-HKUST Sustainable Finance Webinar, and the 6th Annual RCF-ECGI Corporate Finance and Governance Conference for detailed discussions and helpful comments. This research has been funded by the Swiss National Science Foundation under Grant Agreement No. 100018 207800.

1 Introduction

Technological change has long been central to economic growth, reshaping both firms' production and their environmental externalities (Bresnahan and Trajtenberg, 1995; Acemoglu et al., 2012). Artificial intelligence (AI), as an emerging general-purpose technology, is diffusing rapidly across industries. According to PwC's 2024 Global CEO Survey, nearly half of the 4,701 executives identified AI integration as a top strategic focus over the next three years. This widespread adoption has raised growing concerns about its climate consequences. Understanding how AI affects firms' climate externalities is critical for aligning technological progress with global decarbonization goals.

Quantifying the climate impact of corporate AI investment is challenging for two reasons.[1] First, the effects are inherently multidimensional. Running large-scale AI models increases electricity use and can raise emissions.[2] Beyond this direct effect, AI integration can also shift emissions through operational changes. It may reduce emissions by improving productivity and resource efficiency (Brynjolfsson et al., 2025), and it may also increase emissions by expanding output from carbon-intensive operations (e.g., by making previously uneconomical resources profitable; Bolton et al., 2024).[3] The net effect is therefore theoretically ambiguous at both the firm and aggregate levels. Second, data on corporate AI investment is often proprietary or unstructured, further impeding systematic carbon footprint assessments.

To address the research gap, this paper identifies AI investment by applying natural language processing to LinkedIn profiles and performs three stages of investigation. First, I estimate how firm-level AI investment predicts subsequent Scope 1 emissions (i.e., emissions directly tied to firms' production activities) to capture operational responses. Second, I

[1]In this paper, AI refers to machine learning based technologies that learn from high-dimensional data to generate predictions and support optimization in firms' operations. This definition is broader than generative AI, a subset that uses large models to produce new content. It also differs from generic data analysis, which typically relies on rule-based or descriptive methods and is less likely to materially affect decision-making.

[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.

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

develop a theoretical framework grounded in directed technical change to generate testable implications and guide the interpretation of mechanisms. Third, I assess aggregate climate relevance by mapping firm-level effects to industries and also by comparing emissions from operational changes with those from AI-related electricity consumption.

I measure firm-level AI investment as the share of AI workers among 13,617 public firms in 70 countries from 2010 to 2024. The measure is constructed by first identifying candidate workers through a keyword-based search for AI-related terms in LinkedIn job descriptions, and then applying a fine-tuned LLaMA model (Meta, 2024) to assess the semantic context of each profile and determine whether the worker applies AI in internal operations or market-facing functions. This approach mitigates superficial AI mentions and achieves performance comparable to GPT-4 (with F1 scores around 0.94). Based on this measure, AI has expanded markedly over the past decade and extends well beyond the technology sector that develops AI tools. Among non-technology sectors, the proportion of firms employing AI workers rises from 4% in 2010 to 24% in 2024. Importantly, carbon-intensive industries are active adopters. In the Energy sector, for example, about 20% of firms employ AI workers, primarily to improve internal operations. AI workers are more prevalent among larger firms with higher capital expenditures, richer data assets, and stronger AI-related competitive pressure.

I present three main results. First, firms with more AI workers, on average, experience subsequent reductions in Scope 1 emissions. A one standard deviation increase in AI worker shares is associated with a 2% decline in Scope 1 emissions in the following year, growing to 6% by the third year. Both internal and market-facing AI roles contribute to this effect. The results remain robust when controlling for non-AI employment in IT, robotics, and data analytics, as well as firms' green innovation engagement and emission reduction pledges. Using a distributed lead-lag specification to analyze the dynamic effect (Aghion et al., 2020; Babina et al., 2024b), I find no evidence of pre-trends, suggesting the estimates are unlikely to be driven by reverse causality.

Notably, the emission response is heterogeneous across industries and, so far, the average decline is driven mainly by relatively low-emitting sectors. Several industries, such as Consumer Durables, experience significant emission reductions, but their small baseline Scope 1 emissions limit the environmental consequences in absolute terms. In contrast, effects become more varied among carbon-intensive counterparts. The Energy sector exhibits a positive coefficient, pointing to higher emissions. Moreover, in the Utilities sector, which ranks highest in Scope 1 emissions, a one standard deviation increase in AI worker share is linked to a marginally significant 3.5% increase in emissions. Given its substantial baseline emissions, this implies a material increase in absolute carbon output. In short, AI worker share is associated with lower emissions across many industries but higher emissions in the most carbon-intensive ones.

One identification concern is that some unobserved shocks could jointly affect AI worker share and emissions; to support a causal interpretation, I instrument AI worker supply using lottery-induced variation in H-1B visa allocations. The key idea is that many U.S. firms rely on foreign-born AI workers, so fluctuations in H-1B visa availability directly affect their capacity to hire and retain these workers. Because visa demand consistently exceeds the annual cap, the U.S. government allocates H-1B visas through a randomized lottery before reviewing applications, which generates quasi-random variation in firms' AI hiring capacity. I construct the instrument as the difference between a firm's realized number of selected visas and the expected number implied by the aggregate selection rate, scaled by the firm's share of STEM H-1B applications (and, in robustness, by the share of applications with AI keywords in job titles). The instrument has a strong first stage, and the instrumented AI worker share predicts subsequent reductions in Scope 1 emissions, with effects consistently concentrated in lower-emitting sectors.[4]

Second, I explore the economic mechanisms behind the average emission reduction as well as the sector heterogeneity. I develop a theoretical framework built on two economic proper-

[4]Results remain robust when using firms' pre-existing connections to AI-strong universities from Babina et al. (2024b) as an alternative instrument.

ties of AI to interpret the patterns. On one hand, AI is fundamentally a prediction technology that enables more accurate forecasting and decision-making, which reduces resource misallocation, improves operational efficiency, and thus cuts emissions per unit of output (*efficiency* channel). This mechanism is stronger in activities where data-driven coordination is central. On the other hand, AI represents an intangible capital typically deployed to complement existing physical assets and operational constraints to maximize returns. This feature implies that AI functions as a form of directed technical change, with its climate impact depending on whether it augments brown or green activities (*directed scale* channel). This mechanism is especially relevant for high-carbon producers, where AI that complements incumbent fossil assets can expand carbon-intensive output, while AI directed toward renewables can instead support the green transition, affecting both carbon intensity and total emissions.

Empirical evidence supports that differences in the relative strength of the two economic forces explain the variation in emission outcomes. In most sectors, the efficiency channel dominates. Their carbon intensity falls significantly after AI hiring, with magnitudes closely mirroring reductions in emissions levels. To sharpen the mechanism, I use demand forecast accuracy as a proxy for one dimension of operational efficiency and show that firms with a higher AI worker share forecast demand more accurately; this improved forecasting, in turn, is associated with reduced Scope 1 emission intensity in the next year, likely due to tighter production planning and lower waste. As their output is largely unchanged, these efficiency gains translate into lower emission levels.

By contrast, the directed scale channel dominates in Energy and Utilities. Utilizing firm-year data on sales and capital expenditures across operating segments, I find that these firms tend to deploy AI disproportionately toward more carbon-intensive production and investment, resulting in much weaker carbon-intensity abatement or even increases. Focusing specifically on their real activities shows that increases in AI worker share are associated with greater utilization in oil and gas operations and correspondingly higher emissions. The job descriptions further reveal extensive AI use in subsurface modeling and production optimization, which enables faster and more precise extraction from remaining carbon-intensive reserves. Over my sample period, I find no detectable relationship between AI worker share and either green innovation or renewable generation, indicating that AI has not delivered measurable emission abatement along those dimensions. This directed expansion of carbon-intensive activities accounts for the sharply different emission responses in these sectors.

Third, I evaluate the aggregate climate relevance along two dimensions. On one dimension, I relate industry-level variation in AI worker share to subsequent emissions and observe significant declines, indicating that, for many sectors, firm-level reductions are not limited to a small set of adopters but translate into industry-level changes (potentially reflecting within-industry spillovers). However, when I re-estimate the specification using WLS weighted by industries' Scope 1 emissions, the effects become statistically insignificant. This suggests that although AI enhances efficiency in many industries, aggregate carbon impacts are dominated by a small number of high-emitting sectors, where carbon-intensive AI applications currently offset emission reductions achieved elsewhere.

On the other dimension, I examine whether AI-driven changes in Scope 1 emissions are accompanied by higher electricity use and the associated carbon footprint (due to AI computing). Interestingly, hiring AI workers is not linked to a rise in electricity use; instead, increases are concentrated among firms operating in-house AI infrastructure (e.g., data centers, which I detect through annual report disclosures). Back-of-the-envelope estimates imply that such infrastructure adds roughly 41 TWh of electricity use over the past decade, corresponding to about 25.3 $MtCO_2$ of Scope 2 emissions. This magnitude remains modest and is equivalent to approximately 0.9% of total Scope 1 emissions among firms with AI workers in 2023.

Overall, the results suggest that the climate impact of AI extends beyond the electricity required to run models, with AI-induced changes in production and operations also meaningfully affecting corporate emissions. More importantly, while AI improves efficiency, the sign and magnitude of these operational effects depend critically on where and how AI is deployed, or more specifically, whether it primarily reinforces carbon-intensive activities or supports the green transition, which generates pronounced heterogeneity in emission responses and the resulting transition risk across firms.

Contribution and Related Literature. This paper contributes to the growing literature on firm-level AI investment. Recent studies have quantified corporate AI exposure and linked 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 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. Abis and Veldkamp (2023) provide structural estimates of how AI affects knowledge production. In comparison, my paper shifts the lens from labor and financial outcomes to climate externalities—an important yet underexplored dimension of corporate AI investment. The analysis of carbon emissions as a byproduct of production also offers additional empirical evidence on AI's impact on firm-level productivity, which has yielded mixed results in the existing literature (Acemoglu et al., 2022; Babina et al., 2024b). As a physical and process-based outcome, emissions serve as a complement to traditional productivity metrics and are particularly well suited to capture productivity gains from AI integration, which often manifest as process improvement beyond quantity increases.

I also 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 steers the direction of innovation from dirty to clean. Aghion et al. (2025) develop a growth theory in which economic activity 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, I focus on artificial intelligence, which is an emerging general-purpose technology with the potential to transform business activities across diverse industries. This transformative scope with the capacity to reshape operations gives AI high climate salience and warrants systematic study. Although several studies have examined AI-related environmental footprint, most of them focus narrowly on a single language model or provide qualitative assessments based on broad assumptions.[5] I address this gap by combining various datasets with natural language processing to conduct the first large-sample, firm-level analysis of AI's climate impact. By capturing both direct electricity use and broader operational emission outcomes, this paper offers critical early insights into how AI investment has already shaped corporate climate trajectories and the associated transition risk, which will be an important input for assessing the financial implications of this technological change.

2 Detecting Firm-Level AI Worker Shares

This section describes the data sources used in the paper and presents the construction, validation, and descriptive patterns for the firm-level AI worker share measure.

2.1 Data

I use data from three domains: (i) AI investment; (ii) carbon emissions and energy use; and (iii) production-related proxies and financial fundamentals.

To measure firm-level AI investment in human capital, I utilize individual-level employment data from LinkedIn, obtained through the Bright Initiative in October 2024. This dataset provides detailed employee information such as job title, position descriptions, start and end dates (at the month level), and, in many cases, information on project involvement, research publications, and patent contributions, which allows me to identify AI-related workers. I map LinkedIn users to public firms in Compustat through a combination of fuzzy

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

string matching and manual verification. I restrict the sample to firms with English-language LinkedIn profiles and with at least 10% of their reported workforce (per Compustat) matched to LinkedIn users (the median employee coverage is 72%). The data comprises 13,617 global public firms, with 8,106 from Compustat North America and 5,511 from Compustat Global. For identification, I also instrument AI workers with excess H-1B lottery selections using data from the United States Citizenship and Immigration Services (USCIS) and the Office of Foreign Labor Certification (OFLC).

I rely on Scope 1 emissions from the MSCI, which offers broad global coverage and a consistent estimation methodology. Firm-level electricity use comes from the Carbon Disclosure Project (CDP), which collects voluntary disclosures from more than 23,000 companies (roughly two-thirds of global market capitalization). These disclosures include firm-reported electricity consumption (disaggregated by renewable and non-renewable sources) and Scope 2 greenhouse gas emissions (both location-based and market-based).

I use production-related proxies that include manager sales forecasts from I/B/E/S to capture demand prediction ability; within-firm production reallocation across business segments from Compustat Segments; measures of brown activity, such as oil and natural gas utilization, from Compustat Industries; and green activity proxies including green patents from the United States Patent and Trademark Office (USPTO) and renewable generation from the Energy Information Administration (EIA). I obtain financial data from Compustat. The baseline analysis covers global public firms from 2010 to 2024 to capture AI expansion over the past decade.

2.2 Measure Construction

I measure firm-level AI investment using the share of employees with AI-specific roles. The construction is implemented by a two-step natural language processing (NLP) procedure, described in detail in the Internet Appendix B.

In the first step, I apply a keyword-based search for AI-related terms, which reduces the large set of LinkedIn profiles to a manageable pool of candidate workers. Specifically, I detect potential AI workers in a given year by assessing if its profile satisfies at least one of the following criteria, following the method of 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 in the previous, current, or subsequent years. The keyword list is adapted from Babina et al. (2024b) and expanded to reflect recent AI development (e.g., generative AI) and software tools (to capture potential reliance on external AI consultancies). This approach identifies 288,080 AI workers from Compustat firms between 2010 to 2024, of whom 84% (or 241,186) are detected based on job title or description alone.

One potential limitation of the LinkedIn-based AI worker measure is misreporting in either direction. Some employees may mention AI superficially to attract attention, even when engagement is limited, thereby overestimating AI investment; others may omit detailed job descriptions, creating selective reporting that would understate AI investment. To address superficial mentions, I use a second-step procedure in which a large language model (LLM) evaluates the full semantic context of each profile (not just keywords) and retains only profiles that describe concrete machine learning tasks. To address selective non-disclosure, I test robustness by excluding industries in which employees are less likely to disclose role details and by controlling for firms' share of employees with LinkedIn profiles as a proxy for disclosure propensity. Details are provided in Section 3.1.

Using the contextual analysis in the second step, I classify each AI worker as either *internal operations* or *market-facing* functions. These categories align with leading AI applications reported in company surveys and provide a lens to examine whether the emission impact of AI depends on the business function it augments.[6] AI in internal operations supports internal

[6]See McKinsey Global Survey on artificial intelligence.

workflows, such as logistic optimization, energy management, or administrative automation, and is more directly linked to Scope 1 emissions by targeting energy-intensive processes and physical assets. By contrast, market-facing AI refers to applications embedded in customer-facing processes, such as personalization and user experience enhancement, and operates primarily at the point of external interaction.

For this step, I 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 particular, a high-capacity model (e.g., GPT-4.1 in this paper) serves as the teacher and generates good-quality labels and rationales for a sample of candidate paragraphs. These annotations are then used to fine-tune a smaller open-source model (i.e., Meta-LLaMA-3.1-8B-Instruct) that acts as the student. The resultant student model is capable of approximating GPT-4's task-specific classification performance while remaining fully reproducible. I calculate the share of AI workers as:

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

where a higher value indicates deeper engagement. Similarly, I construct separate worker share measures for the two categories.[7]

2.3 Validation Tests

I validate the measures in two complementary ways. For internal validation, I randomly sample 500 job descriptions and construct a human-labeled benchmark. The fine-tuned LLaMA classifier delivers strong out-of-sample performance, with an F1 score of 94% (as shown in Table IA B1). For external validation, I relate the AI worker share to established AI proxies, including AI patent shares from Pairolero et al. (2025), AI infrastructure disclosures in annual reports (see Internet Appendix F for details), and AI strategy discussions in earn-

[7]To reduce noise from sparsely recorded AI roles, I set a firm's AI worker share to zero in years with fewer than five identified AI workers. This measure is highly correlated (>0.8) with the AI worker measure in Babina et al. (2024b), which is based on proprietary resume data.

ings calls (using the AI keywords described above). Each proxy is strongly and positively associated with AI worker share (Columns 1–3 of Table IA2), indicating that the measure captures economically meaningful variation in AI investment rather than noise from LinkedIn disclosures or the NLP procedure. Meanwhile, relative to alternative AI proxies, AI worker share is particularly informative for my research question because it directly reflects firms' operational use of AI and thus represents a plausible channel through which AI can affect production emissions (see Section 3.1 for supporting evidence comparing these proxies).

Table IA B2 provides examples of detected AI worker profiles as a further validation. Panel A illustrates internal operational uses, where KLA leverages AI for manufacturing diagnostics, Schlumberger for drilling optimization, and UPL for automating plant maintenance. Panel B features market-facing applications, such as audience behavior prediction at Time, customer engagement optimization at JPMorgan, and personalized marketing at Warner Music.

2.4 Industry, Time, and Geographic Variation and Determinants

Figure 1 illustrates the industry distribution using the Fama-French 12 classification. AI investment is most prevalent in the Computers, Software, and Electronic Equipment sector, where nearly 37% of firm-years employ AI workers, often in market-facing roles as AI tools are their core commercial products. This sector includes cloud service providers such as Microsoft and hardware manufacturers like NVIDIA. Importantly, AI integration extends well beyond this technology industry. Energy firms rank third across sectors, with roughly 20% employing AI workers. In this sector, AI is more widely integrated into internal operations, such as optimizing energy production and predictive maintenance, where efficiency enhancement directly affects costs and emissions. Across all other industries, Consumer Nondurables and Utilities show the lowest intensity, yet still about 10% of their firm-years report AI employment. To focus on how AI affects the operations of user firms, I exclude firms in the Computers, Software, and Electronic Equipment sector from subsequent descriptive and empirical analyses (except for the part of electricity use in Section 5.2). Table 1, Panel A, presents summary statistics for this sample. Roughly one in 10,000 employees works in AI, indicating that AI employees remain scarce among global firms.[8]

Figure 2 shows the evolution of AI workers from 2010 to 2024. The proportion of firms that hire AI workers has grown steadily from 4% to 24% 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; each plausibly lowered the fixed cost of deployment and broadened the set of tasks amenable to AI technology. Moreover, AI applications in internal operations are more prevalent and earlier to scale, indicating diffusion into production and logistics rather than solely into customer interfaces.

Table IA1 reports the top countries by the number of firm-year observations. The dataset spans 70 countries and the U.S. accounts for 40.6% of observations, followed by India and the United Kingdom, with additional broad coverage across Europe and Asia. This distribution reflects the reliance on English-language LinkedIn profiles, which are more prevalent in certain countries. Several digitally advanced economies, including the U.S., Switzerland, the Netherlands, and Denmark, display a relatively high level of AI workforce engagement.

Table IA2 analyzes the firm-level determinants of AI worker share and suggests that operational needs emerge as the primary drivers, with climate considerations also playing a role. First, firms with richer data assets, as proxied by disclosures in annual reports, tend to have more AI workers; this is consistent with the idea that AI applications rely heavily on access to proprietary historical data for training and deployment. Second, firms are more likely to expand their AI workforce when their product-market peers, based on text-based industry classification from Hoberg and Phillips (2016), had a higher AI worker share in the previous year. This likely reflects peer effects, including investor pressure to adopt emerging

[8]This share is smaller than in the Panel B sample with electricity data, which includes a higher concentration of large firms (since electricity disclosure to CDP is skewed toward large firms) and retains the Computers, Software, and Electronic Equipment sector, where AI employment is substantially higher.

technologies and strategic imitation among operationally similar firms. In addition, firms with AI workers are more likely to have set approved emission reduction targets. In terms of financial characteristics, AI workforce integration is more common among larger firms with higher capital investment and lower property tangibility.

3 AI Worker Shares and Scope 1 Emissions

AI influences corporate emissions through changes in operations. On the one hand, AI can raise productivity and operational efficiency, which reduces 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), which increases total emissions. These opposing mechanisms make the net emission impact theoretically ambiguous and require empirical investigation. This section examines the relationship between AI worker share and Scope 1 emissions and supports the analysis with an instrumental variable (IV) strategy based on H-1B visa lottery selections among U.S. firms.

3.1 Average Effects

I estimate the regression following the method of Bolton et al. (2024):

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

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

Table 2 reports the results. A one standard deviation increase in AI worker share is associated with a 2% reduction in Scope 1 emissions in the subsequent year. This effect intensifies over time, reaching a 6% reduction by the third year. Both internal and market-facing AI roles contribute to emission reductions. The high R^2 (around 0.97) is largely because emissions are persistent and the fixed effects absorb most cross-firm differences. Given that the identification comes from within-firm changes over time, a large R^2 is not problematic per se. The dynamic patterns presented below further indicate that the emission results are hard to attribute to pure persistence.

Table IA3 examines heterogeneity in this relationship across firm characteristics. The effect strengthens after 2018 (the year after the introduction of transformer-based models) and is more pronounced among U.S. firms. Larger firms exhibit stronger reductions, while the estimates remain statistically significant for smaller firms. I also classify AI workers into different types using a similar two-step NLP procedure. The reductions are concentrated among leader (senior) AI workers, consistent with a managerial decision-making effect. Splitting AI workers into invention versus adoption shows that the reductions are primarily driven by adoption. Further, considering the measure emphasizes in-house AI workers and may overlook outsourced use, I identify AI worker profiles indicating reliance on external or third-party AI solutions and find around 10% do so, suggesting in-house AI workers and external services can be complements. The emission effects arise from the larger group that does not mention external consulting, pointing to the operational importance of in-house investment.

Table IA4 extends the analysis to other emission scopes. The reduction persists for Scope 2 emissions, suggesting AI also reduces operational electricity use. Panel B presents the results for the fifteen subcategories of Scope 3 emissions. Electricity use from AI computing provided by third-party services would most likely appear in Categories 1 and 2 (Purchased Goods and Services, and Capital Goods), yet the coefficients in these categories show no significant effects. This may reflect the fact that AI-related computing expenses constitute a small fraction of total purchased goods and services, generating limited observable variation in these categories. In contrast, several other subcategories exhibit negative and statistically significant coefficients. For example, AI worker share is associated with lower upstream transportation emissions (Category 4), reduced waste generated in operations (Category 5), fewer business travel emissions (Category 6), and decreases in downstream production (Category 10) and franchise-related emissions (Category 14). These patterns imply that AI contributes to decarbonization across supply-chain activities through coordination improvements.

Robustness Checks: Table IA5 demonstrates that the results are robust to a wide range of confounders and alternative specifications. First, Panel A controls for the use of other technologies by including the shares of IT, robotics, and non-AI data analytics workers. The coefficient on AI worker share remains robust, 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. Second, Panel B addresses the possibility that firms using AI are more engaged in green initiatives, which could independently drive emission reductions. I 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 consistent results suggest 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). Third, Panel C adds local renewable energy availability, electricity prices, and corporate financial constraints (Whited and Wu, 2006). Fourth, the results are unchanged when controlling for alternative AI proxies, such as AI patent shares, AI-related infrastructure disclosures in annual reports, and AI strategy discussion in earnings calls, consistent with emission changes being driven by operational AI use that the LinkedIn-based worker measure captures more directly. Fifth, Panel E considers alternative AI worker constructions, 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. Sixth, to address selective disclosure on LinkedIn, I exclude security-sensitive industries (Defense, Aircraft, and Shipbuilding) and trade-secret-intensive industries (Medical Equipment, Pharmaceuticals, and Measuring Equipment), and control for firm-level LinkedIn coverage (an indicator for above-median share of workers with a profile). Panel G additionally includes country-year fixed effects, clusters standard errors by both firm and year, and clusters by industry. The results remain consistent across all specifications.

Dynamic Patterns: I estimate the dynamic effects of AI workers on Scope 1 emissions to address potential reverse causality. As firms adjust AI worker share gradually rather than through discrete shocks, conventional event-study designs that rely on a single treatment time are not well-suited to this setting. I therefore employ a distributed lead-lag model that accommodates continuous variation in AI workers (Aghion et al., 2020; Babina et al., 2024b):

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

where $\Delta AIWorker_{i,t-k-1\to t-k}$ is the change in the share of AI workers from year $t-k-1$ to year $t-k$. The dependent variables are the logarithm of Scope 1 emissions in year t. Control variables and fixed effects follow the baseline specification. β_k captures the cumulative response of the emission outcome in year t to changes in AI worker share k years earlier, holding constant AI worker share in all other years. The lead coefficients serve as a pre-trend test; if firms investing in the AI workforce follow emission trends similar to those of others before investments, these coefficients should not be statistically different from zero. Figure 3, Panel A, reports the coefficients from the lead-lag regressions and shows that Scope 1 emissions decline following AI workforce investment, with the effects intensifying over the next two to three years. Importantly, there is no evidence of pre-trends: firms that use AI have similar emission changes to others in the years prior to investment, and significant divergence occurs only afterward. This suggests that the results are unlikely to be driven by reverse causality.

3.2 Sector Heterogeneity

While the firm-level estimates establish a reduction in Scope 1 emissions following AI investment, this average effect may obscure variation across sectors. It is important to note that the impact of a given percentage reduction depends heavily on a firm's baseline emissions (Hartzmark and Shue, 2023). For instance, a two-percent drop in a high-emitting energy firm translates into a much larger absolute reduction in CO_2 than the same percentage decrease in a low-emitting service firm. As absolute emissions are what matter for decarbonization goals, I explore where AI contributes materially to emission mitigation.

Figure 4 explores this heterogeneity by estimating regressions at the sector level, based on the Fama-French 12 industry classification, and considers firm and country-year fixed effects. Sectors are ranked by their average Scope 1 emissions to highlight differences in baseline emission levels. Notably, industries with emission reductions are often those with already low emissions. The Finance sector stands out as an intensive AI-using industry (see Figure 1). 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 these employees leading to a significant 4.5% emission reduction. Yet again, the sector's low emissions mean that these reductions, while statistically significant, are modest in absolute environmental terms.

The emission impacts become more varied among carbon-intensive industries. In Chemicals, the negative but insignificant coefficient suggests AI helps reduce emissions in some processes. The coefficient becomes positive for the Energy sector, pointing to higher emissions. Moreover, in the Utilities sector, which ranks highest in Scope 1 emissions, a one standard deviation increase in AI worker share is associated with a 3.5% increase in emissions. Given its massive baseline, this implies a considerable absolute increase in CO_2 output. This heterogeneity raises questions about the effectiveness of AI-related emission reductions observed before, particularly considering the critical role of Energy and Utilities in achieving global decarbonization goals, which together account for over 50% of total emissions across industries (Pastor et al., 2024).

Dynamic Patterns: Figure 3 illustrates the dynamic effects separately for firms in the Energy and Utilities (Panel B) and for all other sectors (Panel C). In both panels, the pre-period estimates are small and statistically indistinguishable from zero. Following AI investment, their trajectories diverge markedly. Among Energy and Utilities firms, the coefficients remain flat through $t+1$, spike sharply to a positive value at $t+2$, and then turn negative at $t+3$ with a wide confidence interval. This suggests no systematic abatement and, if anything, a short-run increase in Scope 1 emissions. Firms outside Energy and Utilities exhibit an accumulating decline post-investment, with the point estimates becoming more negative over time. These results confirm that the relationship between AI workforce expansion and emissions differs across sectors, with the emission reduction mainly arising from lower-emitting parts of the economy.

3.3 Instrumental Variable Estimates

To strengthen causal inference and address potential bias from unobserved shocks that jointly drive both firms' AI worker share and emission outcomes, I employ an IV approach by using excess H-1B visa selections to instrument 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 requires an H-1B visa to work there legally; second, H-1B demand routinely exceeds the annual cap and prompts 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. As full visa lotteries occurred in each year after 2014 (2010 to 2013 only saw lotteries for a subset of filings), this IV estimate focuses on U.S. firms from 2014 to 2023. Internet Appendix C provides detailed information on the institutional background and measure construction.

I construct the instrument as the realized number of visa selection minus the expected number implied by the aggregate selection rate, building on the method of Cen et al. (2024):

$$AIExcessVisa_{i,t} = (VisaSelection_{i,t} - VisaDemand_{i,t} \times SelectRate_t) \times STEMShare_{i,t},$$

where $VisaSelection_{i,t}$ is the number of firm i's H-1B 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 better target AI workers, I scale excess visa selections by the firm-year share of STEM H-1B registrations (based on SOC codes), which are more likely to correspond to AI roles. As a robustness check, I also weight excess selections by the share of registrations with AI keywords in job titles; this weight may understate true AI intensity as some AI work is not explicitly labeled in job titles, but it is the most granular text-based information available in the visa data.

The instrument satisfies the relevance condition as firms that receive more excess H-1B selections are mechanically able to hire more foreign AI capable employees, thereby increasing their subsequent AI worker share.[9] To illustrate, consider two firms that each submit ten STEM H-1B registrations in a year with a 50% selection rate, with one firm receiving eight selections and the other receiving none. Their $AIExcessVisa$ 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. The firm with -5 may face delays in AI initiatives due to a shortage of skilled workers, while the firm

[9] Although firms can use other visa categories, H-1B is a major channel for hiring high-skill foreign workers. In my sample, restricting to observations with both AI workers and H-1B registrations, the median numbers of registrations and selections are 14 and 5, respectively. Weighting selections by the firm's STEM application share (or by the share with AI keywords in job titles), the implied number of AI-related H-1B selections corresponds to roughly 80% (or 23%) of newly hired AI workers in the same firm-year. This suggests that the lottery shock can generate economically meaningful variation in firms' ability to staff AI roles.

with +3 is more likely to proceed as planned. Table 3, Column 1, confirms that the instrument $AIExcessVisa$ strongly predicts changes in AI worker share, with a statistically significant positive coefficient and a first-stage Kleibergen-Paap F-statistic above the threshold of 10. Here, as excess visas directly impact the flow of newly hired AI workers each year rather than the AI worker stock, I accordingly estimate a first-difference specification. I use changes in carbon emissions as the dependent variable and changes in AI worker share as the independent variable of interest, controlling for industry-year fixed effects. This approach is conceptually aligned with the baseline specification.

Columns 2–4 of Table 3 report the second-stage (2SLS) estimates. A one standard deviation increase in instrumented AI worker share leads to a 6% reduction in Scope 1 emissions in the following year, expanding to a 9% reduction after three years. This supports the causal link between AI integration and emission reduction. These coefficients are larger in magnitude than the corresponding OLS coefficients, which may be explained by the local average treatment effects. The IV estimates capture the effect for firms whose AI workforce investment is constrained by access to skilled foreign labor. These firms are likely to face binding shortages of technical talent, so additional AI hires are valuable and often deployed in the most productivity-enhancing applications. Consequently, the resulting emission reductions can be larger than the average effect observed in the OLS estimates.

The instrument satisfies the exclusion restriction as the random component of H-1B visa selection is highly unlikely to effect firms' carbon emissions through channels other than AI labor supply. Table IA C1 provides two sets of supporting exercises. First, Panel A assesses AI specificity using alternative instrument constructions and placebo tests. While the baseline instrument scales the excess visas by firms' STEM application share, it could still capture broader tech hiring beyond AI. To sharpen the AI channel, I reweight the excess visas by the share of visa applications with AI-related keywords in job titles and obtain consistent results for both stages of the IV test. As placebo instruments, I reweight the same excess visas by the share of Non-STEM applications and by the share of applications without AI keywords. These placebos yield weak first stages and no effect on emissions, consistent with the interpretation that the baseline results operate through AI-related labor. I further show that the lottery shock does not predict changes in other worker shares, including IT, robotics, non-AI data analytics, and the second-stage estimates are unchanged when controlling for these other shares.[10] Excluding firms that provide IT services leaves the results the same, alleviating concerns that high-filing consultancies drive the findings.

Second, Panel B addresses concerns related to endogenous filing behavior and within-firm workforce substitution. Firms choose how many registrations to submit, so the level of H-1B selections is not fully random. My identification instead comes from the lottery-driven deviation of realized selections from expected selections, conditional on a firm's registrations and the aggregate selection rate. Holding registrations fixed, this residual variation is quasi-random. I also control for filing intensity using visa registrations scaled by total employment, which mitigates the concern that strategic variation in filing intensity drives the results. To assess within-firm reallocation, I obtain similar estimates using an AI headcount proxy rather than worker share, suggesting the results are not driven by denominator effects or worker reshuffling. Results are also robust to excluding firms with substantial foreign exposure (more than 25% non-U.S. sales), which could substitute toward offshore AI capacity following unsuccessful visa selection. Placebo timing tests using leads of the instrumented change in AI worker share show no predictive power for past emission changes, reducing concerns about anticipatory trends. I also construct the instrument using pre-period STEM exposure (2010–2013, so that only pre-existing propensity to hire technical talent loads on lottery variation) and restrict the sample to the pre-2021 period (before the electronic registration changes enabled multi-employer registrations for the same beneficiary). Across these specifications, the first stage remains strong and the second-stage estimates are stable, supporting both robustness and the exclusion-based interpretation of the instrument.

[10] This suggests that non-AI technical roles (e.g., IT, robotics, non-AI data analytics) can be filled more readily through domestic hiring or other channels, so their shares respond little to H-1B lottery shocks. In contrast, AI workers are harder to replace through domestic hiring given the large share of foreign AI workers.

Importantly, Panel A of Table IA6 demonstrates that the estimated effect is primarily driven by firms outside the Energy and Utilities industries, consistent with the presence of meaningful sectoral heterogeneity in how AI workers affect emissions.[11]

4 Economic Mechanisms

Why do emissions fall in many sectors after AI integration, yet abatement is much weaker or even reversed in Energy and Utilities? This section develops a simple theoretical framework adapted from Acemoglu et al. (2012) and formalizes two economic mechanisms. I empirically test these mechanisms in Energy and Utilities versus other sectors and supplement the analysis with evidence from firms' real activities and AI-related job descriptions.

In particular, AI affects emissions through two channels. First, AI is fundamentally a prediction technology that extracts complex patterns from high-dimensional data (Farboodi and Veldkamp, 2022, 2023; Mihet and Philippon, 2019). Since most operating and investment choices (such as production planning and supply chain management) rely on forecasts, higher predictive accuracy reduces misallocation of labor and capital. This lowers waste, raises efficiency, and ultimately reduces emission intensity in many settings. Second, AI is an intangible capital that is typically deployed to complement installed physical assets (Mihet and Philippon, 2019; Brynjolfsson et al., 2021), and thus it functions as a form of directed technical change that tends to reinforce the emission profile of the processes it augments (Hemous and Olsen, 2021; Aghion et al., 2016). In high-emission sectors, this complementarity can work in two directions. When AI is deployed to improve intrinsically carbon-intensive

[11] For robustness, Panel B of Table IA6 also instruments 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), and obtains consistent results. 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 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 IA6 presents the estimates for the U.S. sample from 2010 to 2023. This instrument has a reliable first stage and shows a robust reduction in Scope 1 emissions over the subsequent two years; the third-year coefficient is negative but insignificant.

processes, it cannot by itself alter the carbon content and the productivity gains may scale up carbon output and emissions. When AI instead complements green capital, such as renewable generation, it can facilitate reallocation from brown to green activities and accelerate both the transition and emission reductions in these sectors. The net emission effect depends on where AI is deployed and on the balance between efficiency gains and directed reallocation across activities.

4.1 Theoretical Framework

Let a denote the AI intensity of a firm, with a convex adoption cost $C(a)$. The firm produces a single final good by combining brown and green composite inputs X_j in a constant elasticity of substitution (CES) production function with elasticity $\sigma > 1$:

$$Y = \left(Y_b^{\frac{\sigma-1}{\sigma}} + Y_g^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \qquad Y_j = A_j(a)X_j = a^{\eta}\Theta_j(a)X_j, \qquad j \in \{b, g\}.$$

AI improves productivity in two ways: a neutral efficiency term a^{η}, which improves overall efficiency and reduces total inputs per unit of output, and a directed term $\Theta_j(a)$, which captures complementarity with brown or green inputs with elasticity $\mu_j(a) = \frac{\partial \ln \Theta_j}{\partial \ln a}$. Cross-firm differences in AI complementarities $(\mu_b - \mu_g)$ naturally arise from differences in their production structure and in the availability of relevant historical data for training and deploying AI tools. Input prices are normalized to $p_b = p_g = 1$ for simplicity. On the demand side, I assume monopolistic competition with CES preferences and elasticity $\xi > 1$. This implies a constant markup so that AI-driven cost reductions pass through proportionally to prices and hence affect output scale.

The cost of producing one unit of output and the associated expenditure shares on brown and green inputs are (see the proofs in Internet Appendix D.1 and D.2):

$$c(a) = \frac{1}{a^{\eta}}\left[\Theta_b^{\sigma-1} + \Theta_g^{\sigma-1}\right]^{\frac{1}{1-\sigma}}, \qquad s_b = \frac{\Theta_b^{\sigma-1}}{\Theta_b^{\sigma-1} + \Theta_g^{\sigma-1}}, \qquad s_g = 1 - s_b.$$

AI reduces unit cost through the neutral efficiency term a^{η} and through directed improvements in $\Theta_j(a)$. The firm chooses input shares to minimize cost, which are endogenously adjusted with AI intensity according to:

$$\frac{ds_b}{d\ln a} = (\sigma - 1)s_b s_g(\mu_b - \mu_g).$$

Thus, the brown share tends to rise when $\mu_b > \mu_g$, and vice versa.

Corporate emissions are a by-product of production and equal the sum of inputs weighted by their respective emissions factors (e_j):

$$E = e_b X_b + e_g X_g, \qquad 0 < e_g < e_b.$$

The firm's emission intensity is emissions per unit of output (see the proofs in Internet Appendix D.3):

$$\bar{e} = \frac{E}{Y} = c\left(e_b s_b + e_g s_g\right),$$

that is, unit cost (inputs per unit of output) times emissions per dollar of inputs.

The elasticity of emission intensity with respect to AI intensity is:

$$\frac{d\ln \bar{e}}{d\ln a} = \underbrace{-\tilde{\eta}}_{\text{efficiency term (-)}} + \underbrace{(\sigma - 1)s_b s_g(\mu_b - \mu_g)\Delta_e}_{\text{directed term (+/-)}}, \quad \Delta_e \equiv \frac{e_b - e_g}{e_b s_b + e_g s_g}. \tag{1}$$

The first component, $-\tilde{\eta} = -\eta - (s_b\mu_b + s_g\mu_g)$, is the efficiency term. AI-enabled prediction reduces the inputs (either brown or green) required per unit of output, so emission intensity falls correspondingly at a rate η. Meanwhile, $-(s_b\mu_b + s_g\mu_g)$ captures that when AI raises productivity more for inputs with larger expenditure shares, unit input use and associated emissions fall further. The second component, $(\sigma - 1)s_b s_g(\mu_b - \mu_g)\Delta_e$, is the directed term. With substitutable inputs, if AI is more complementary to brown activities ($\mu_b > \mu_g$), the cost-minimizing composition shifts toward brown inputs. This reallocation raises emissions per unit of output in proportion to the emission gap Δ_e. Conversely, if AI complements green inputs more ($\mu_g > \mu_b$), the mechanism works in the opposite direction and intensifies the decline in emission intensity. Whether intensity ultimately falls or rises depends on the relative strength of the efficiency and directed terms.

The output elasticity with respect to AI intensity is

$$\frac{d \ln Y}{d \ln a} = \xi \left[\eta + s_b \mu_b + s_g \mu_g \right], \tag{2}$$

where the demand elasticity $\xi > 1$ captures the extent to which cost savings translate into higher quantities rather than lower prices.

The elasticity of the emission level follows from $E = \bar{e} \cdot Y$:

$$\frac{d \ln E}{d \ln a} = \frac{d \ln \bar{e}}{d \ln a} + \frac{d \ln Y}{d \ln a} = \underbrace{-\tilde{\eta}}_{\text{efficiency channel (-)}} + \underbrace{(\sigma - 1) s_b s_g (\mu_b - \mu_g) \Delta_e + \xi \left[\eta + s_b \mu_b + s_g \mu_g \right]}_{\text{directed scale channel (+/-)}}.$$

An increase in AI intensity a affects emissions through (i) an *efficiency* channel that lowers emissions per unit of output through improved prediction and optimization, and (ii) a *directed scale* channel that captures input reallocation across activities and the associated output expansion. If the market expansion component (governed by ξ) is similar across sectors, differences in emission responses are driven mainly by where AI is deployed (summarized by the directed term $\propto (\mu_b - \mu_g)\Delta_e$). In most sectors, the emission gap (Δ_e) across inputs is small, so the directed term is muted and the efficiency channel dominates. In high-carbon sectors, Δ_e from different energy sources can be large and the directed term is more relevant. The propositions below formalize these implications.

Proposition 1 (Efficiency-dominated channel):

If the emission gap across inputs, Δ_e, is small, the directed term in Eq. (1) is weak and $\left.\frac{\partial \ln \bar{e}}{\partial \ln a}\right|_s < 0$. Total emissions fall if the induced scale response is sufficiently muted, so that $\frac{d \ln \mathcal{E}}{d \ln a} \simeq \frac{d \ln \bar{e}}{d \ln a} < 0$.

Proposition 2 (Directed scale-dominated channel):

If Δ_e is large and AI is more complementary to brown (green) inputs, the directed term in Eq. (1) is positive (negative) and attenuates (amplifies) the decline in emission intensity. In the brown case, total emissions rise when the induced reallocation and scale-up of brown input use outweigh the efficiency-driven reduction in intensity.

4.2 Energy and Utilities vs. Other Sectors

Guided by the framework, I test whether the two channels operate differently in Energy and Utilities relative to other industries. I focus on these sectors because they display distinct emission responses to AI investment and play a central role in decarbonization. They also feature large emissions gaps across production inputs (e.g., energy sources), making the directed scale channel more likely to be quantitatively important. To isolate the mechanisms, I decompose emissions into emission intensity, which reflects both the efficiency and directed terms, and sales growth, which captures short-term output expansion, as dependent variables:

$$CarbonIntensity^{Scope1}_{i,t+k} \text{ or } SalesGrowth_{i,t+k} = \alpha + \beta_1\, AIWorker_{i,t} \times EnergyUtil_{i,t} + \beta_2\, AIWorker_{i,t} + \gamma \mathbf{X}_{i,t} + a_i + b_{j,t} + \epsilon_{i,t},$$

where $EnergyUtil_{i,t}$ is an indicator equal to one for firms in the Energy or Utilities sectors. The coefficient β_2 identifies the association between AI worker share and the outcome outside Energy and Utilities, and β_1 measures how this association differs within these industries.

Table 4 reports the results. Columns 1–3 show that $\beta_2 < 0$ and is statistically significant at all horizons, with magnitudes comparable to those in Table 2. This pattern indicates that, outside Energy and Utilities, a higher share of AI workers is associated with reduced carbon intensity, aligned with improved operational efficiency. Columns 4–6 do not show a significant effect on sales growth, suggesting limited scale expansion. The combined evidence is consistent with the efficiency channel dominating in most sectors.

In contrast, the interaction term β_1 in Columns 1–3 is positive and statistically significant over the next two years, indicating that, relative to the rest of the economy, the carbon intensity reduction associated with AI is substantially weaker for Energy and Utilities. Meanwhile, Columns 4–6 show a positive (marginally significant) β_1 at $t+2$, which is robust to using deflated value (based on OECD country-industry output deflators) in Panel A of Table IA7.[12] These patterns suggest a positive directed term, where efficiency gains are partly offset by reallocation toward browner activities. I evaluate this directed scale channel with two tests: first, I examine whether a higher AI worker share predict within-firm reallocation toward more carbon-intensive segments (i.e., *BrownShift*); second, I assess whether such reallocation translates into higher emissions.

I measure *BrownShift* using firm-year sales and capital expenditures across different business segments at the NAICS-6 level and map each segment to a carbon-intensity benchmark from the EPA's Supply Chain Greenhouse Gas Emission Factors.[13] I calculate the year-over-year change in sales-weighted carbon intensity as:

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

where $Sale_{i,t}^s$ is the sales of firm i in year t in segment s and $CarbonIntensity^s$ is the carbon intensity of segment s according to the EPA. This measure captures the shift in realized

[12]This implies the sales response reflects changes in real quantities rather than purely price movements. Why does the scale response concentrate in Energy and Utilities? The evidence points to a combination of supply- and demand-side forces. On the supply side, AI is an intangible capital that tends to generate the highest returns when paired with installed physical assets. Energy and Utilities are among the most asset-intensive sectors, where AI applications can relax operational constraints and directly raise utilization of existing capacity. On the demand side, energy (especially electricity) is an essential upstream input into nearly all production activities; when capacity is relaxed, increases in supply are more likely to translate into higher quantities because demand is broadly derived from downstream activity rather than limited to a narrow end market. Panel B of Table IA7 shows that the sales growth effect is stronger for firms with higher tangibility and a higher intermediate use share (i.e., the fraction of sector output sold as intermediate inputs rather than final demand, based on OECD Inter-Country Input-Output tables). Importantly, as shown in Panel C, Energy and Utilities have substantially higher tangibility and intermediate input shares than other sectors, consistent with these characteristics helping account for their AI-driven scale expansion.

[13]This emission factor dataset is based on environmentally extended input-output modeling that links economic activity in each sector to the associated supply-chain emissions in CO_2-equivalent terms.

emission profile implied by the firm's current sales composition across segments. I similarly construct the capital expenditure-weighted carbon intensity, which reflects the change in expected future emission profile implied by investment allocation. The two *BrownShift* measures can serve as empirical counterparts to the directed term, capturing whether a larger share of production and investment comes from browner inputs.

Table 5, Panel A, implements the first test by estimating:

$$BrownShift_{i,t} = \alpha + \beta_1 \, AIWorker_{i,t} \times EnergyUtil_{i,t} + \beta_2 \, AIWorker_{i,t} + \gamma \mathbf{X}_{i,t} + a_i + b_{j,t} + \epsilon_{i,t}.$$

The strongly positive β_1 indicates that AI integration in Energy and Utilities is disproportionately associated with shifts toward browner segments, whereas other sectors show no comparable pattern, as reflected by the near-zero β_2.

Panel B tests if this direction of AI deployment explains the emission outcomes using:

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

β_1 is positive and statistically significant for both sales- and investment-based *BrownShift*. This implies that the marginal emission effect of AI depends on the direction of reallocation; when AI is directed toward more carbon-intensive activities (*BrownShift* > 0), the abatement weakens or can reverse. Columns 2 and 4 show that this pattern is driven by AI roles in internal operations. The combined evidence suggests that the directed scale channel operating through brown activities dominates in Energy and Utilities.

After establishing how the two channels differ across sectors, I link each to firms' real activities in the subsections below to illustrate how they translate into emission changes.

4.3 Efficiency Channel

While AI can enhance efficiency across many domains, these effects are hard to quantify systematically. I therefore focus on a measurable and economically important dimension—demand forecasting (Corhay et al., 2025). The intuition is that sharper sales forecasts enable tighter production planning and lower emissions per unit of output. I evaluate whether firms with more AI workers generate more accurate demand forecasts and whether these forecast improvements (if any) contribute to lower emission intensity.

I measure forecasting ability as follows, with a higher value indicating greater accuracy:[14]

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

Panel A of Table 6 relates demand forecast accuracy to AI worker share. A one standard deviation increase in AI worker share is associated with a 0.02 standard deviation improvement in forecasting accuracy. Panel B uses the logarithm of carbon intensity as the dependent variable. The negative interaction term indicates that improved demand forecasting further reduces Scope 1 emission intensity. Interestingly, while forecast improvements are driven by market-facing AI roles, the additional emission reductions stem from AI workers in internal operations, suggesting complementarities between external demand prediction and internal process optimization.

Evidence for efficiency gains beyond demand forecasting comes from job descriptions at firms that experience AI-enabled emission reductions. To connect the narratives to my estimates, I identify the 50 firms that contribute most to each sector's AI coefficient (in Figure 4) and construct word clouds from their AI worker job descriptions. The Internet Appendix E describes the text retrieval procedure and presents illustrative excerpts from highly influential firms in four representative sectors that differ in AI functions and carbon profiles. Figure 5

[14] This calculation follows Feng et al. (2009). The absolute forecast error captures the magnitude rather than the direction of the error. I use the first management sales forecasts issued within one year before the sales announcement as a proxy for firms' forecasting capability.

shows that emission reductions in Consumer Durables stem from AI prediction embedded in product design and manufacturing. For example, General Motors applies machine learning to enhance advanced driving systems, and Ford develops predictive quality models. In Wholesale and Retail, AI is integrated into customer-facing systems and logistics. Nordstrom uses AI to optimize inventory consolidation and supply-chain management, and Zalando runs daily machine-learning pipelines to improve warehouse order processing, both of which reduce operational inefficiencies and associated emissions.

4.4 Directed Scale Channel

I examine the directed scale channel that dominates in Energy and Utilities by testing whether AI enhances fossil fuel utilization or instead supports the green transition.

To proxy for fossil asset utilization, I construct year-over-year changes in oil and natural gas production scaled by proved developed reserves; higher values reflect a greater extraction intensity. Panel A of Table 7 presents that a higher AI worker share is associated with increased fuel utilization. A one standard deviation increase in AI workers in internal operations correlates to a 0.09 standard deviation increase in natural gas utilization. Panel B interacts AI worker share with fossil utilization to predict next-year Scope 1 emissions and shows that internal AI capacity amplifies the emission increase from greater fossil utilization.

To proxy for the green transition, I leverage the shares of green patents and renewable generation, which the IEA identifies as critical for long-run emission mitigation (IEA, 2025). If AI supports these transition efforts, we would expect a higher AI worker share to be positively associated with these outcomes and their interaction to correspond with lower future Scope 1 emissions. However, Panels C and D find no statistically significant effects during my sample period, suggesting that AI-driven improvements in these dimensions have not yet produced measurable emission reductions. This may reflect the early stage of green AI applications or the longer time horizon over which their environmental benefits materialize.

This brown-side directed scale channel is also supported by narrative evidence. As reflected in Panel C of Figure 5, in the Energy sector, AI is often embedded in subsurface modeling and production optimization. At Schlumberger, engineers deploy the industry's first AI-driven autonomous wireline logging in Kuwait and apply neural networks for reservoir characterization in mature oil fields, which enable faster and more precise identification of remaining reserves; at Halliburton, reservoir managers implement signal analysis algorithms during the production phase to improve geophysical analysis and cut turnaround times by 30%. While such tools improve efficiency, they also increase capacity utilization by allowing firms to extract more output from existing carbon-intensive assets. In the Utilities sector, AI is mainly applied to predictive maintenance and load forecasting. Employees at Tata Power implement AI-based solutions for reliability-centered maintenance. These cases show AI's potential to reduce emissions when applied to renewable sources but also the possibility of sustaining fossil-based infrastructure by optimizing thermal plants and pipelines.

5 Aggregate Implications

This section assesses the aggregate climate relevance of the firm-level emission results along two dimensions. First, I relate industry-level variation in AI worker share to subsequent emissions using WLS regressions weighted by Scope 1 emission levels to test whether emission reductions survive after accounting for heterogeneity between high- and low-emitting sectors. This design also captures potential within-industry spillovers that a firm-level panel may miss. Second, given that AI computing is electricity-intensive, I evaluate whether the observed Scope 1 emission changes are accompanied by increased electricity use and its associated carbon footprint. These tests combined shed light on whether firm-level average emission declines scale into meaningful aggregate decarbonization.

5.1 Industry-Level Emission Estimates

To examine the link between AI worker share and aggregate Scope 1 emissions, I estimate the following specification at the industry-country-year level:

$$CarbonEmission^{Scope1}_{jc,t+k} = \alpha + \beta\, AIWorker_{jc,t} + \gamma \mathbf{X}_{jc,t} + a_{jc} + b_t + \epsilon_{jc,t}, \quad k \in \{1, 2, 3\}.$$

$CarbonEmission^{Scope1}_{jc,t+k}$ denotes total Scope 1 emissions of industry j from country c in year $t + k$, based on Fama-French 49 industry classification. The key independent variable, $AIWorker_{jc,t}$, is the share of AI workers among all employees in the same industry-country-year group. The control variables include industry-country-level aggregates or averages of firm characteristics used in the baseline regressions. I include industry-country and year fixed effects, and standard errors are clustered at the industry-country level.

Table 8 reports the results. Columns 1–3 present the WLS estimates weighted by the number of employees as a proxy for firm size. A one standard deviation increase in an industry's AI worker share relates to a 3% reduction in industry-level Scope 1 emissions in the following year, which deepens to about 7% by the third year. This suggests that emission reductions are not confined to the small subset of firms with AI workers, but extend to lower emissions at the industry level. These magnitudes are larger than those found in the firm-level analysis. One potential reason is that larger firms (with relatively larger carbon emissions within the same industry) are more likely to integrate AI into their operations. Meanwhile, within-industry spillovers may further contribute to the stronger industry-level response. For instance, AI-driven logistics has the potential to lower costs for peer firms through shared contractors, even when those peers do not directly employ AI workers.

Columns 4–6 present the results when observations are weighted by Scope 1 emission levels. This approach effectively evaluates the effect per ton of CO_2 emitted and emphasizes the outcomes in more carbon-intensive sectors. Notably, under this weighting, the coefficients become statistically insignificant and close to zero. This suggests that although AI improves efficiency and reduces emissions in many sectors, its overall carbon impact remains heavily shaped by results in a few high-emitting industries, where brown AI applications currently outweigh the efficiency gains achieved elsewhere.

5.2 Electricity Consumption

Given that the energy demand from AI computing is a widely discussed concern, I complement the main analysis by testing whether firms' AI investment predicts subsequent growth in electricity use (with additional details in the Internet Appendix F):

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

The dependent variable captures changes in electricity consumption from year t to $t+k$, measured over one-, two-, and three-year horizons. $\Delta AIWorker$ represents the change in the share of AI workers relative to the previous year. $\mathbb{1}(AIInfras)$ equals one if the firms' disclosures of AI infrastructure (e.g., data centers and server rooms) in annual reports are above the median among samples with positive disclosures. I add $\mathbb{1}(AIInfras)$ because AI infrastructure is the input that pins down the energy profile of AI computing workloads (IEA, 2025), and controlling for it helps isolate the workforce channel. Meanwhile, I include the Computers, Software, and Electronic Equipment sector in this analysis.

Table 9 shows no statistically significant relationship between AI worker share and electricity use. This does not mean AI integration is energy-free. Instead, incremental electricity demand is concentrated among firms that build AI infrastructure. Relative to their industry peers, these firms consume 5% more electricity in the following year, rising to 12% in the second year.[15] This pattern implies that increased energy demand is driven by physical in-

[15] Table IA F3 decomposes electricity changes into renewable and non-renewable components. Firms with AI workers experience a modest composition shift in their electricity use toward renewable sources. On the other hand, electricity demand among infrastructure-heavy firms is increasingly met by non-renewable sources. Columns 1–3 of Panel A show their renewable electricity increases by 18% (with a t-stat of 2.25) in $t+1$, but is not statistically different from zero thereafter, whereas non-renewable electricity use rises

frastructure investment rather than hiring AI employees. Consistent with this interpretation, many AI user firms appear to outsource computation and the associated electricity footprint to external infrastructure providers, as reflected in Table IA F4.[16]

I estimate back-of-the-envelope electricity consumption attributable to AI infrastructure using the coefficient of Column 1 in Table 9. With the sample mean electricity use of 1,322,968 MWh (Panel B, Table 1), a 5% annual increase corresponds to approximately 66,148 MWh per treated firm-year and roughly 41 TWh when aggregated across firm-years with positive $\mathbb{1}(AIInfras)$ (622) over the past decade. Renewables account for 20% of total electricity use in the sample (see Table 1); applying this share implies an additional 25.3 $MtCO_2$ of Scope 2 emissions.[17] This magnitude remains modest in aggregate terms and is equivalent to approximately 0.9% of Scope 1 emissions from firms with AI workers in a single year of 2023. The comparison suggests that AI-driven changes in production and operating decisions are likely to be the first-order channel to affect corporate emissions. That said, the additional electricity load is concentrated in a relatively small set of infrastructure firms (or providers) and can create meaningful local pressure on generation and grid capacity, which is a separate issue from this paper's focus on carbon emissions.

6 Conclusion

This paper provides early evidence on the climate impact of AI investment using a global firm-level sample. On average, a one standard deviation increase in a firm's AI worker shares

persistently, reaching 20% above baseline in $t+2$ and 22% by $t+3$. Panel B shows that Scope 2 emissions increase significantly with AI infrastructure investment, especially for location-based measures.

[16]Table IA F4 assesses whether firms with AI employees indirectly bear the energy costs by analyzing supply chain linkages. Firms with a higher AI worker share 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 (t = 2.94). Similarly, firms disclosing greater AI infrastructure involvement are more likely to serve downstream customers with a higher AI worker share; these customers do not use more electricity themselves but experience reduced Scope 1 emissions.

[17]I use a fossil average emission factor of 0.77 $kgCO_2/kWh$ for non-renewable electricity and a zero emission factor for renewables.

predicts a 2% decline in emissions in the following year, with the magnitude increasing over time. The reductions arise from AI-enabled prediction and optimization that enhance operational efficiency and reduce carbon intensity. However, this average effect masks substantial heterogeneity. In high-emitting sectors, most notably Energy and Utilities, abatement is much weaker and emissions sometimes increase, consistent with AI being directed toward optimizing existing fossil-fuel operations rather than facilitating a transition toward greener technologies. A plausible reason for this directed use is that these firms remain heavily invested in fossil assets and have decades of accumulated operational data that make legacy AI applications more immediately profitable; by contrast, greener applications often require new capital investment and data, limiting the short-run response. Overall, the emissions-weighted aggregate effect is close to zero, indicating that brown AI use in a small set of large emitters offsets emission reductions across the rest of the economy.

These results have two implications. First, the climate footprint of AI extends beyond the electricity required to train and run the models; another key channel operates through how AI is deployed within firms' production and operations. This channel is especially important for large emitters, where even modest AI-induced changes can generate large absolute emission effects given high baseline levels. Second, unlike climate technologies explicitly designed to reduce emissions, AI is a general-purpose optimizer. It can improve efficiency, but these gains do not automatically translate into lower emission levels because it depends on what that efficiency is used to scale. So far, AI deployment within high-carbon sectors appears to have tilted toward scaling brown rather than green activities, outweighing efficiency gains elsewhere. Looking ahead, aligning AI investment with decarbonization will require embedding climate objectives into corporate AI strategies and strengthening incentives that direct AI toward the green transition. This becomes even more important when considering the growing energy demand associated with the widespread adoption of AI in the future.

Data Appendix

Variables	Definitions	Sources
$AIWorker$	Share of employees in AI-related roles.	LinkedIn
$AIWorker^{Internal}$	Share of AI employees in internal operations.	LinkedIn
$AIWorker^{Market}$	Share of AI employees in market-facing functions.	LinkedIn
$AIWorker^{Leader}$	Share of AI workers in leadership (senior) roles.	LinkedIn
$AIWorker^{NonLeader}$	Share of AI workers in non-leadership roles.	LinkedIn
$AIWorker^{Invention}$	Share of AI workers with invention-related activities.	LinkedIn
$AIWorker^{Adoption}$	Share of AI workers with adoption (implementation) activities.	LinkedIn
$AIWorker^{Consulting}$	Share of AI workers indicating reliance on external or third-party AI solutions.	LinkedIn
$AIWorker^{NonConsulting}$	Share of AI workers not indicating reliance on external or third-party AI solutions.	LinkedIn
$AIWorker^{Peer}_{i,t-1}$	Average share of AI workers in peer firms based on Hoberg and Phillips (2016) from the previous year.	LinkedIn
$IT/Robot/DataWorker$	Share of employees in IT/robot/non-AI data-related roles.	LinkedIn
$LinkedInCoverage$	Indicator equal to 1 if the firm's share of workers with a LinkedIn profile is above the sample median, and 0 otherwise.	LinkedIn
$\Delta AIWorker$	Year-over-year change in the AI worker share.	LinkedIn
$AIExcessVisa$	Excess STEM-related H-1B lottery selections, defined as the number of selected STEM H-1B registrations minus the expected number given the firm's total registrations and the fiscal-year aggregate selection rate.	USCIS, OFLC
$AIExcessVisa^{Text}$	Excess AI-keyword H-1B lottery selections, defined as the number of selected registrations with AI-keyword in job titles minus the expected number given the firm's total registrations and the fiscal-year aggregate selection rate.	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)
$AIInfras$	Proportion of AI infrastructure-related disclosure in 10-K filings (U.S. firms) or annual reports (non-U.S. firms).	SEC EDGAR, Firm websites
$\mathbb{1}(AIInfras)$	Indicator for firm-years with above-median $AIInfras$ disclosures.	SEC EDGAR, Firm websites
$AIPatent$	Share of a firm's patents classified as AI-related, based on the Artificial Intelligence Patent Dataset (Pairolero et al., 2025).	USPTO
$CarbonEmission^{Scope1}$	Logarithm of Scope 1 emissions.	MSCI
$CarbonIntensity^{Scope1}$	Logarithm of Scope 1 emissions scaled by sales.	MSCI
$\Delta CarbonEmission^{Scope1}$	Year-over-year change in the logarithm of Scope 1 emissions.	MSCI
$CarbonEmission^{Scope2}$	Logarithm of Scope 2 emissions.	MSCI
$CarbonEmission^{Scope3}$	Logarithm of Scope 3 emissions.	MSCI
$Electricity$	Logarithm of total electricity consumption.	CDP
$\Delta Electricity$	Change in the logarithm of total electricity consumption.	CDP
$Electricity^{Renew}$	Logarithm of renewable electricity consumption.	CDP
$Electricity^{NonRenew}$	Logarithm of non-renewable electricity consumption.	CDP
$EnergyUtil$	Indicator equal to one for firms in the Energy and Utilities sectors.	Compustat
$DemandForecast$	Negative of the difference between actual sales and the first management forecast issued within one year before the sales announcement, scaled by actual sales.	I/B/E/S
$BrownShift$	Within-firm changes in sales-weighted or capex-weighted carbon intensity across business segments.	Compustat Segments
$FossilUtilization$	Year-over-year change in oil or natural gas production scaled by proved developed reserves.	Compustat Industry
$GreenTransition$	Share of green patents or renewable electricity generation.	USPTO, EIA

(Continued)

(Continued)

Variables	**Definitions**	**Sources**
$EmissionTarget$	Indicator for firms with approved science-based or net-zero targets.	MSCI
$ElectricityPrice$	Annual average commercial electricity price at the state level.	EIA
$RenewableGrids$	Share of renewable electricity in local grids at the state level (inside the U.S.) or country level (outside the U.S.).	EIA
$CCExposure$	Firm-level climate change exposure derived from earnings calls.	Sautner et al. (2023)
$DataAssets$	Proportion of sentences related to data assets in annual reports.	SEC EDGAR, Firm websites
$log(Asset)$	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
$SalesGrowth^{Deflated}$	Percentages change in deflated sales using OECD country-industry output deflators.	Compustat, OECD
$DownstreamDemand$	Fraction of sector output sold as intermediate inputs rather than final demand, based on OECD Inter-Country Input-Output tables.	Compustat, OECD
WW	Firm-level measure of financial constraints from Whited and Wu (2006).	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 NBER 33634*.

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. (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.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–72.

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*.

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.

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.

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*.

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

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

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

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*.

Pairolero, N. A., Giczy, A. V., Torres, G., Erana, T. I., Finlayson, M. A., and Toole, A. A. (2025). The artificial intelligence patent dataset (aipd) 2023 update. *The Journal of Technology Transfer*.

Pastor, L., Stambaugh, R. F., and Taylor, L. A. (2024). Carbon burden. *Working Paper, SSRN 4998860*.

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*.

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

Figure 1: **Industry Distribution**

This figure illustrates the distribution of AI workers at the industry level, based on the Fama-French 12 classification. The line showcases the proportion of firms with AI workers (left axis). The two bars represent the average share of AI workers engaged in internal operations and market-facing functions (right axis).

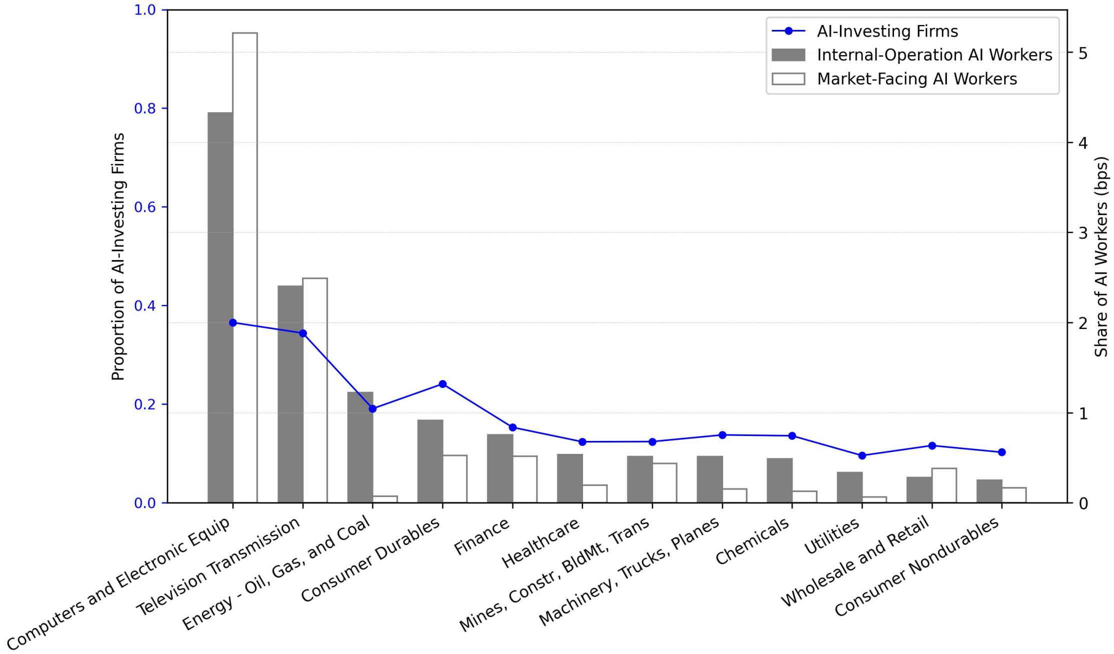

Figure 2: **Time-Series Variation**

This figure illustrates the time-series variation in AI workers. The line showcases the proportion of firms with AI workers (left axis). The two bars represent the average share of AI workers engaged in internal operations and market-facing functions (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. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) to focus on AI user firms.

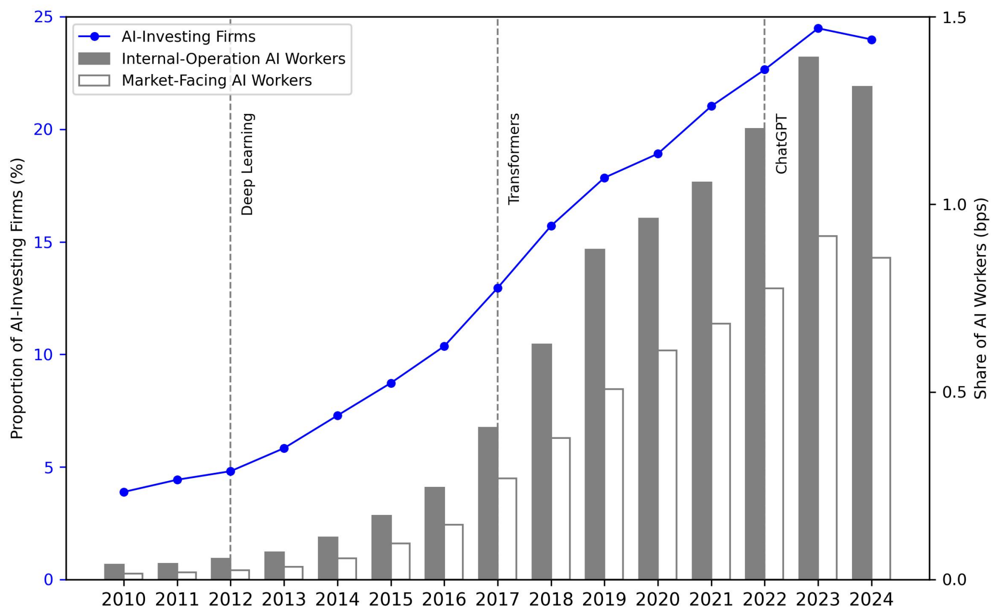

Figure 3: **AI Worker Shares and Scope 1 Emissions: Dynamic Effects**

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.

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

Panel A includes firms from all industries (excluding the Computers, Software, and Electronic Equipment firms), whereas Panels B and C focus on the Energy & Utilities sectors versus other sectors, respectively. The sample period is from 2010 to 2023. Standard errors are clustered at the firm and year level.

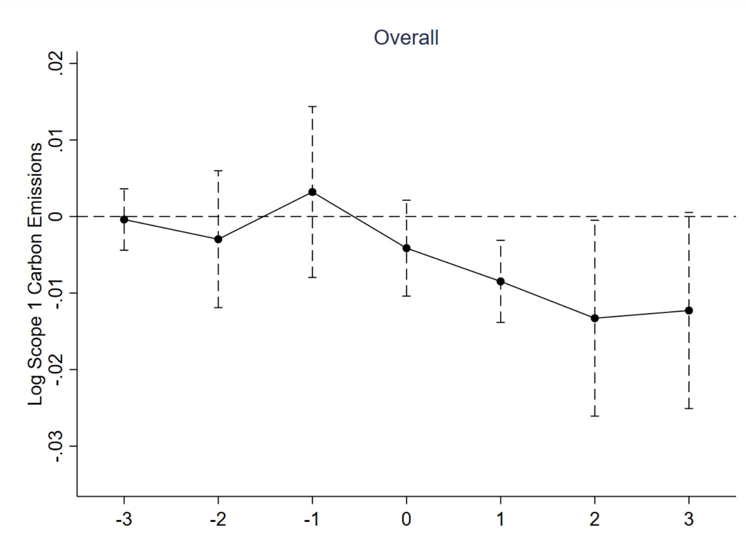

(a)

(continued)

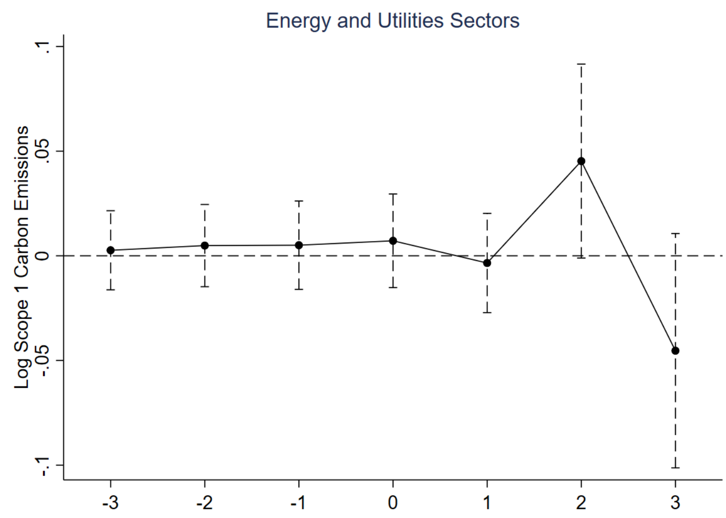

(b)

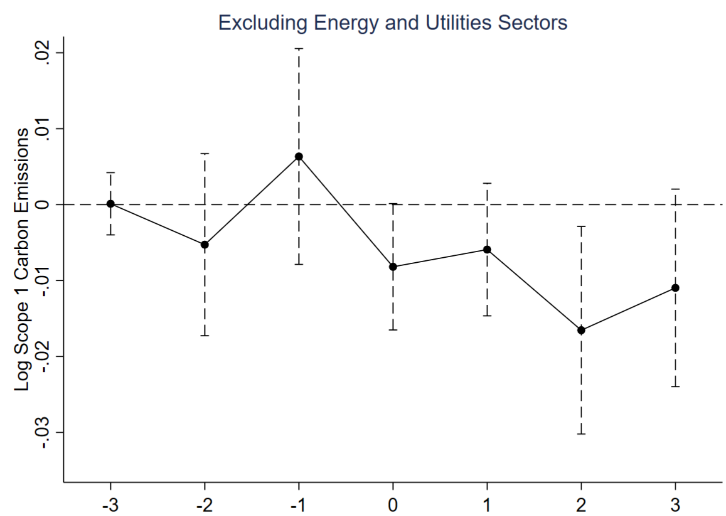

(c)

Figure 4: **AI Worker Shares and Scope 1 Emissions: Sector Heterogeneity**

This figure shows the impact of AI worker share on Scope 1 carbon emissions for each industry, based on the Fama-French 12 classification. Black dots represent regression coefficients for each sector, with vertical lines indicating 90% confidence intervals. Gray bars report the average Scope 1 emission levels.

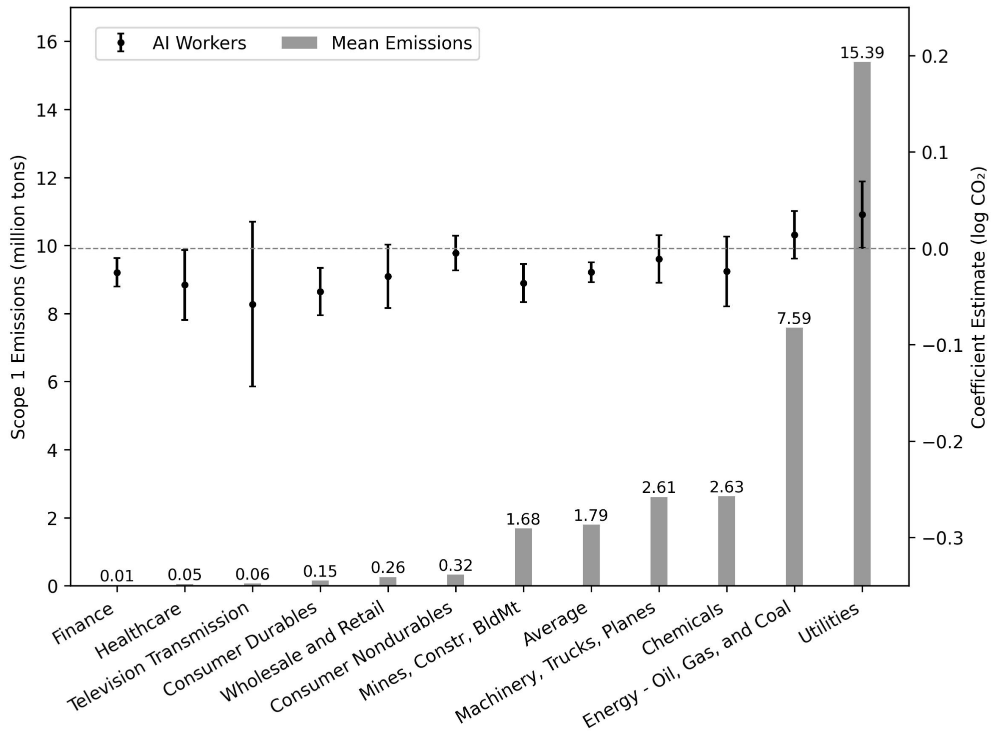

Figure 5: **AI Worker Shares and Scope 1 Emissions: Narrative Retrieval**

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 Consumer Durables, Wholesale and Retail, Energy, and Utilities, respectively. The estimated coefficients are negative for Panel A–B and positive for Panel C–D (circled in red).

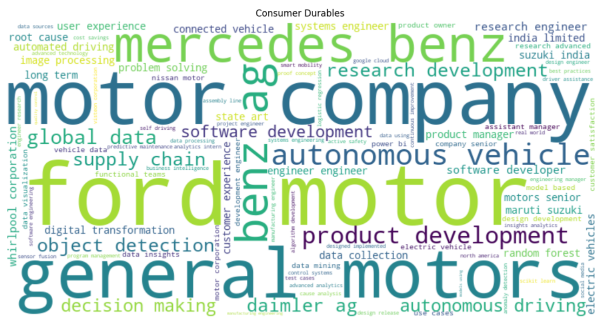

(a)

(b)

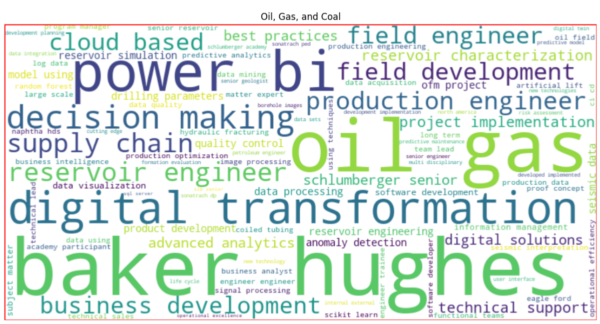

(c)

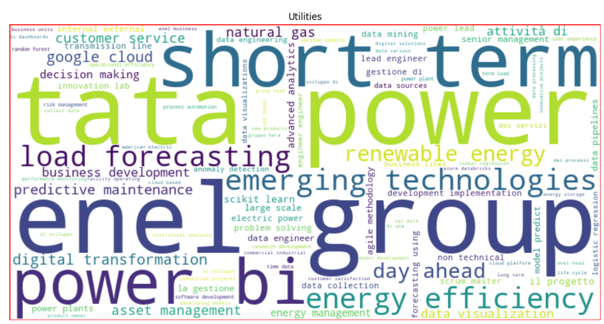

(d)

Table 1: **Summary Statistics**

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

Variable	Mean	SD	p10	p50	p90	N
Panel A: Scope 1 Emission Sample						
$AIWorker\,(bp)$	0.98	4.32	0	0	1.69	38,969
$AIWorker^{Internal}\,(bp)$	0.60	2.39	0	0	1.11	38,969
$AIWorker^{Market}\,(bp)$	0.36	2.18	0	0	0.26	38,969
$CarbonEmission^{Scope1}\,(level)$	1,357,820	5,373,256	284	16,063	1,782,281	38,969
$CarbonIntensity^{Scope1}\,(level)$	195.41	641.31	0.37	8.53	405.35	38,936
Panel B: Electricity Sample						
$AIWorker\,(bp)$	5.33	10.59	0	0.55	15.47	4,272
$\mathbb{1}(AIInfras)$	0.28	0.45	0	0	1	4,272
$Electricity\,(level)$	1,322,968	2,810,351	35,430	369,738	3,489,033	4,272
$Electricity^{Renew}\,(level)$	256,831	721,716	0	27,284	549,720	4,034
$Electricity^{NonRenew}\,(level)$	1,058,110	2,475,267	4,236	234,720	2,789,262	4,140

Table 2: **AI Worker Shares and Scope 1 Emissions**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. I use the logarithm of Scope 1 emissions in years t+1, t+2, and t+3 as dependent variables. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year logarithm 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.

	$CarbonEmission^{Scope1}$					
	$t+1$		$t+2$		$t+3$	
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	-0.02***		-0.04***		-0.06***	
	(-3.98)		(-4.06)		(-4.22)	
$AIWorker^{Internal}$		-0.01*		-0.02*		-0.03**
		(-1.72)		(-1.91)		(-1.99)
$AIWorker^{Market}$		-0.02***		-0.04***		-0.05***
		(-2.79)		(-2.59)		(-2.73)
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 3: **AI Worker Shares and Scope 1 Emissions: Instrumental Variable**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are changes in the logarithm of Scope 1 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 instrumental variable is the excess H-1B visa selections a firm receives due to the lottery in year t-1, which are weighted by the share of visa applications for STEM positions. Control variables include changes in $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, and $SalesGrowth$. The sample excludes firms in the Computers, Software, and Electronic Equipment sector (based on Fama-French 12 classification) 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.

	$\Delta AIWorker$	$\Delta CarbonEmission^{Scope1}$		
		$t \to t+1$	$t \to t+2$	$t \to t+3$
	(1)	(2)	(3)	(4)
$AIExcessVisa$	0.14***			
	(4.59)			
$\Delta AIWorker$		-0.06**	-0.09**	-0.09*
		(-2.50)	(-2.18)	(-1.69)
N	9,972	8,578	7,244	5,961
F-Stat	21.11	22.01	20.34	16.80
Control	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Table 4: **AI Worker Shares and Scope 1 Emissions: Carbon Intensity vs. Sales**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are the logarithm of Scope 1 carbon intensity or the sales growth in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$. 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 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.

	$CarbonIntensity^{Scope1}$			$SalesGrowth$		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times EnergyUtil$	0.03**	0.05*	0.03	0.02	0.03*	0.01
	(2.31)	(1.87)	(1.15)	(1.31)	(1.95)	(0.71)
$AIWorker$	-0.02***	-0.04***	-0.05***	-0.00	0.00	-0.00
	(-3.16)	(-3.49)	(-3.60)	(-0.71)	(0.02)	(-0.33)
N	29,351	26,230	23,189	29,134	26,230	23,014
R^2	0.972	0.962	0.960	0.383	0.372	0.364
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 Worker Shares and Scope 1 Emissions: Brown Shift**

This table analyzes the relationship between future Scope 1 emissions and AI worker share, with a focus on AI as a form of directed technical change. I measure *BrownShift* as the year-over-year changes in sales-weighted or capital expenditure-weighted carbon intensity across business segments to capture within-firm shift toward browner production or investment. Panel A uses *BrownShift* as the dependent variable, and Panel B uses the logarithm of Scope 1 emissions in year $t+1$ as the dependent variable. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year logarithm 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: Production and Investment Reallocation

	$BrownShift$			
	SalesShift		CapexShift	
	(1)	(2)	(3)	(4)
$AIWorker \times EnergyUtil$	0.27*		0.37**	
	(1.74)		(2.09)	
$AIWorker$	0.00		0.00	
	(0.20)		(0.03)	
$AIWorker^{Internal} \times EnergyUtil$		0.09		0.33**
		(0.44)		(2.15)
$AIWorker^{Internal}$		-0.00		-0.02
		(-0.21)		(-0.60)
$AIWorker^{Market} \times EnergyUtil$		0.74		-1.09
		(0.49)		(-1.09)
$AIWorker^{Market}$		0.00		0.02
		(0.39)		(1.08)
N	12,421	12,421	9,434	9,434
R^2	0.153	0.153	0.136	0.136
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Panel B: Subsequent Carbon Emissions

	$CarbonEmission^{Scope1}_{t+1}$			
	SalesShift		CapexShift	
	(1)	(2)	(3)	(4)
$AIWorker \times BrownShift$	0.02***		0.02**	
	(3.43)		(2.24)	
$AIWorker$	-0.04***		-0.04**	
	(-4.41)		(-2.53)	
$AIWorker^{Internal} \times BrownShift$		0.02***		0.01**
		(3.02)		(2.34)
$AIWorker^{Internal}$		-0.02		-0.02
		(-1.58)		(-1.12)
$AIWorker^{Market} \times BrownShift$		0.00		0.04
		(0.10)		(0.96)
$AIWorker^{Market}$		-0.03***		-0.02**
		(-3.15)		(-2.02)
$BrownShift$	-0.00	-0.00	0.00	0.01
	(-0.95)	(-1.17)	(1.02)	(1.04)
N	11,184	11,184	8,502	8,502
R^2	0.985	0.985	0.985	0.985
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Table 6: **AI Worker Shares and Scope 1 Emissions: Efficiency Channel**

This table analyzes the relationship between future Scope 1 emissions and AI worker share, with a focus on AI as a prediction technology to improve demand forecasting. I measure forecast accuracy as the negative of the absolute sales forecast error divided by actual sales. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$. 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: Forecast of Sales

	$DemandForecast$	
	(1)	(2)
$AIWorker$	0.02**	
	(2.12)	
$AIWorker^{Internal}$		-0.01
		(-0.52)
$AIWorker^{Market}$		0.03**
		(2.09)
N	9,954	9,954
R^2	0.118	0.118
Control	Yes	Yes
Firm FE	Yes	Yes
Industry-Year FE	Yes	Yes

Panel B: Subsequent Carbon Intensity

	$CarbonIntensity_{t+1}^{Scope1}$	
	(1)	(2)
$AIWorker \times DemandForecast$	-0.05**	
	(-2.35)	
$AIWorker$	-0.04***	
	(-2.83)	
$AIWorker^{Internal} \times DemandForecast$		-0.06***
		(-3.22)
$AIWorker^{Internal}$		-0.03**
		(-2.31)
$AIWorker^{Market} \times DemandForecast$		0.01
		(0.80)
$AIWorker^{Market}$		-0.01
		(-0.82)
$DemandForecast$	0.02	0.02
	(0.52)	(0.53)
N	6,306	6,306
R^2	0.940	0.940
Control	Yes	Yes
Firm FE	Yes	Yes
Industry-Year FE	Yes	Yes

Table 7: **AI Worker Shares and Scope 1 Emissions: Directed Scale Channel**

This table analyzes the relationship between future Scope 1 emissions and AI worker share, with a focus on the utilization of fossil fuels versus green transition. Panels A and B measure fossil fuel utilization as the year-over-year change in oil production scaled by proved developed reserves in Columns 1–2 and natural gas production scaled by proved developed reserves in Columns 3–4. Panels C and D measure corporate green transition as the green patent share in Columns 1–2 and the renewable generation share in Columns 3–4. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year logarithm 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: Fossil Fuel Utilization

	$FossilUtilization$			
	Oil		NaturalGas	
	(1)	(2)	(3)	(4)
$AIWorker$	0.04		0.06*	
	(1.10)		(1.81)	
$AIWorker^{Internal}$		0.04		0.09**
		(1.17)		(2.32)
$AIWorker^{Market}$		0.02		0.01
		(1.02)		(0.42)
N	186	186	174	174
R^2	0.065	0.065	0.261	0.264
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Panel B: Fossil Fuel Utilization and Subsequent Carbon Emissions

	$CarbonEmission_{t+1}^{Scope1}$			
	Oil		NaturalGas	
	(1)	(2)	(3)	(4)
$AIWorker \times FossilUtilization$	0.14*		0.06***	
	(1.97)		(3.92)	
$AIWorker$	-0.02		-0.06**	
	(-0.70)		(-2.63)	
$AIWorker^{Internal} \times FossilUtilization$		0.22**		0.05***
		(2.44)		(5.19)
$AIWorker^{Internal}$		-0.04		-0.06
		(-1.52)		(-1.68)
$AIWorker^{Market} \times FossilUtilization$		-0.17		-0.02
		(-0.64)		(-0.93)
$AIWorker^{Market}$		-0.03		-0.02
		(-0.93)		(-1.72)
$FossilUtilization$	0.04**	0.04	0.01	0.01
	(2.21)	(0.97)	(1.41)	(0.87)
N	160	160	149	149
R^2	0.946	0.947	0.962	0.962
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

(Continued)

(Continued)

Panel C: Green Transition

	$GreenTransition$			
	GreenPatent		RenewableGeneration	
	(1)	(2)	(3)	(4)
$AIWorker$	0.00		0.06	
	(0.07)		(0.31)	
$AIWorker^{Internal}$		-0.02		0.10
		(-0.83)		(0.58)
$AIWorker^{Market}$		0.01		-0.06
		(1.21)		(-0.37)
N	5,167	5,167	370	370
R^2	0.704	0.704	0.456	0.456
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Panel D: Green Transition and Subsequent Carbon Emissions

	$CarbonEmission_{t+1}^{Scope1}$			
	GreenPatent		RenewableGeneration	
	(1)	(2)	(3)	(4)
$AIWorker \times GreenTransition$	-0.02		0.01	
	(-0.53)		(1.18)	
$AIWorker$	-0.08***		-0.05	
	(-3.14)		(-1.13)	
$AIWorker^{Internal} \times GreenTransition$		-0.04		0.01
		(-1.36)		(0.68)
$AIWorker^{Internal}$		-0.07***		-0.02
		(-3.10)		(-0.36)
$AIWorker^{Market} \times GreenTransition$		0.03		-0.00
		(0.66)		(-0.05)
$AIWorker^{Market}$		-0.02		-0.05
		(-0.78)		(-1.26)
$GreenTransition$	-0.03*	-0.03*	-0.01	-0.01
	(-1.68)	(-1.72)	(-0.66)	(-0.76)
N	4,139	4,139	321	321
R^2	0.977	0.977	0.994	0.994
Control	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes

Table 8: **AI Worker Shares and Scope 1 Emissions: Industry-Level Estimates**

This table analyzes the relationship between industry-country-level Scope 1 emissions and AI worker share. The dependent variable is the logarithm of total Scope 1 emissions among firms within each industry-country group in years $t+1$, $t+2$, and $t+3$. Columns 1–3 report WLS estimates weighted by the number of employees, and Columns 4–6 report WLS estimates weighted by the level of Scope 1 emissions. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$. 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 period is from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the industry-country group. *** p<0.01, ** p<0.05, * p<0.1.

	$CarbonEmission^{Scope1}$					
	Number of Employees			Level of Scope 1 Emissions		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker$	-0.03**	-0.05***	-0.07***	0.01	0.00	-0.01
	(-2.47)	(-2.77)	(-2.71)	(0.65)	(0.03)	(-0.34)
N	7,020	6,384	5,750	7,020	6,384	5,750
R^2	0.986	0.981	0.980	0.967	0.955	0.956
Control	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 9: **AI Worker Shares and Electricity Consumption**

This table analyzes the relationship between changes in electricity consumption and AI worker share. The dependent variables are changes in the logarithm of electricity consumption from year t to t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$, as well as the current-year logarithm 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.

	$\Delta Electricity$		
	$t \rightarrow t$+1	$t \rightarrow t$+2	$t \rightarrow t$+3
	(1)	(2)	(3)
$\Delta AIWorker$	-0.00	0.00	0.02
	(-0.46)	(0.13)	(1.19)
$\mathbb{1}(AIInfras)$	0.05**	0.12**	0.12*
	(2.50)	(2.35)	(1.91)
N	2,762	2,205	1,783
R^2	0.073	0.135	0.151
Control	Yes	Yes	Yes
Country FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes

Internet Appendix

for

Carbon in the Cloud

This Internet Appendix provides supplementary material for the study. Section A includes additional figures and tables for the empirical analysis. Section B details the NLP approach used to measure corporate AI investment. Section C outlines the institutional background of the H-1B visa lottery and instrument construction. Section D provides the theoretical framework linking AI with emissions. Section E presents narrative evidence from job descriptions related to sectoral heterogeneity in emissions effects. Section F reports additional details on emissions from AI-related electricity consumption.

A Additional Figures and Tables

Table IA1: **Geographical Distribution**

This table reports the distribution of the sample across the top 30 countries, ranked by number of observations.

Country	$AIWorker$(bp)	$AIWorker^{Internal}$(bp)	$AIWorker^{Market}$(bp)	$\#Obs$
United States	0.99	0.59	0.41	18,905
India	0.70	0.45	0.25	4,705
United Kingdom	1.08	0.59	0.49	3,365
Canada	0.63	0.47	0.15	2,171
Australia	0.63	0.41	0.22	1,762
Sweden	1.14	0.60	0.54	1,261
Germany	0.92	0.59	0.33	1,252
Switzerland	1.27	0.93	0.34	1,166
France	1.05	0.77	0.28	1,143
Italy	0.67	0.50	0.17	784
Japan	1.01	0.42	0.59	773
Netherlands	1.88	1.35	0.53	602
South Africa	1.20	0.74	0.46	568
Spain	1.14	0.63	0.51	533
Denmark	2.06	1.31	0.74	493
Brazil	0.53	0.45	0.07	468
Norway	1.26	0.65	0.61	465
Ireland	1.18	0.77	0.41	429
Finland	1.09	0.84	0.26	418
Hong Kong	0.47	0.36	0.11	399
New Zealand	0.50	0.25	0.25	363
Turkey	4.13	2.55	1.58	360
Belgium	1.57	0.98	0.60	328
China	0.45	0.15	0.30	302
Austria	0.23	0.19	0.03	294
Mexico	0.05	0.04	0.01	255
Bermuda	0.49	0.44	0.05	246
Singapore	2.94	2.07	0.87	241
Luxembourg	0.29	0.24	0.05	206
Israel	0.22	0.21	0.02	197

Table IA2: **Determinants of AI Worker Shares**

This table examines the determinants of firms' decisions to hire AI workers. The dependent variable is the share of AI workers. $AIPatent$ is the share of a firm's patents classified as AI-related, based on the Artificial Intelligence Patent Dataset (Pairolero et al., 2025). $AIInfras$ is the proportion of AI infrastructure-related disclosures in firms' annual reports. $AIStrategies^{ECC}$ is the share of discussion devoted to AI strategies in earnings call transcripts. $DataAssets$ is the discussion proportion of data assets in annual reports. $AIWorker^{Peer}_{i,t-1}$ represents the average share of AI workers in peer firms (based on Hoberg and Phillips (2016)) from the previous year. $EmissionTarget$ is an indicator for firm-years with approved science-based or net-zero emission targets from the MSCI. 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.

	$AIWorker$					
	(1)	(2)	(3)	(4)	(5)	(6)
$AIPatent$	0.30***					
	(6.35)					
$AIInfras$		0.09***				
		(3.20)				
$AIStrategies^{ECC}$			0.30***			
			(6.87)			
$DataAssets$				0.12***		
				(3.65)		
$AIWorker^{Peer}_{i,t-1}$					0.62***	
					(5.97)	
$EmissionTarget$						0.08*
						(1.89)
$log(Asset)$	0.09***	0.13***	0.14***	0.13***	0.14***	0.12***
	(6.83)	(9.34)	(11.71)	(7.77)	(8.50)	(8.33)
ROA	0.16	0.08	-0.01	0.01	0.08	0.06
	(1.21)	(0.61)	(-0.13)	(0.08)	(0.58)	(0.26)
$Capex/Assets$	1.53***	2.08***	1.16***	2.22***	1.96***	1.45***
	(4.17)	(5.42)	(3.48)	(4.36)	(4.48)	(3.18)
$Leverage$	-0.15*	-0.20**	-0.22***	-0.20*	-0.13	-0.29**
	(-1.78)	(-2.27)	(-2.78)	(-1.90)	(-1.46)	(-2.25)
$Tangibility$	-0.79***	-0.84***	-0.67***	-0.94***	-0.75***	-0.82***
	(-5.09)	(-5.43)	(-5.76)	(-4.50)	(-4.22)	(-4.68)
$SalesGrowth$	-0.08***	-0.07***	-0.04*	-0.07***	-0.05**	0.01
	(-4.03)	(-3.62)	(-1.81)	(-3.12)	(-2.48)	(0.26)
N	15,944	17,720	24,527	13,068	13,927	16,637
R^2	0.220	0.169	0.227	0.160	0.258	0.208
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table IA3: **AI Worker Shares and Scope 1 Emissions: Variation by Time, Geography, Firm Size, and AI Worker Type**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are the logarithm of Scope 1 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 logarithm 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.

	$CarbonEmission^{Scope1}$								
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Time, Countries, and Firm Size Variation									
$AIWorker \times \mathbb{1}(Before2018)$	-0.01	-0.03**	-0.06***						
	(-1.02)	(-2.03)	(-3.24)						
$AIWorker \times \mathbb{1}(After2018)$	-0.03***	-0.05***	-0.06***						
	(-4.20)	(-4.10)	(-4.04)						
$AIWorker \times \mathbb{1}(US)$				-0.03***	-0.06***	-0.07***			
				(-3.51)	(-3.80)	(-3.90)			
$AIWorker \times \mathbb{1}(NonUS)$				-0.02**	-0.03**	-0.05**			
				(-2.47)	(-2.31)	(-2.49)			
$AIWorker \times \mathbb{1}(Small)$							-0.02	-0.03*	-0.04*
							(-1.61)	(-1.76)	(-1.84)
$AIWorker \times \mathbb{1}(Medium)$							-0.02**	-0.04**	-0.05**
							(-2.44)	(-2.39)	(-2.01)
$AIWorker \times \mathbb{1}(Large)$							-0.03***	-0.05***	-0.08***
							(-3.85)	(-3.85)	(-4.97)
N	29,363	26,240	23,200	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	0.980	0.973	0.971
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: AI Worker Types									
$AIWorker^{Leader}$	-0.02**	-0.04***	-0.03*						
	(-2.20)	(-2.67)	(-1.88)						
$AIWorker^{NonLeader}$	-0.01	-0.01	-0.03						
	(-0.72)	(-0.39)	(-1.59)						
$AIWorker^{Invention}$				-0.01	-0.01	-0.00			
				(-1.11)	(-1.10)	(-0.05)			
$AIWorker^{Adoption}$				-0.02***	-0.04***	-0.06***			
				(-3.48)	(-3.63)	(-4.01)			
$AIWorker^{Consulting}$							-0.00	-0.01	0.00
							(-0.23)	(-0.43)	(0.31)
$AIWorker^{NonConsulting}$							-0.02***	-0.04***	-0.07***
							(-2.71)	(-2.81)	(-4.22)
N	29,363	26,240	23,200	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	0.977	0.968	0.966
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA4: **AI Worker Shares and Indirect Carbon Emissions**

This table analyzes the relationship between future indirect emissions and AI worker share. The dependent variables are the logarithm of Scope 2, combined Scope 1+2, and Scope 3 emissions in Panel A, and fifteen Scope 3 subcategories in Panel B. 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: Indirect Carbon Emissions

	$CarbonEmission^{Scope2}$			$CarbonEmission^{Scope1+2}$			$CarbonEmission^{Scope3}$		
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$AIWorker$	-0.02***	-0.04***	-0.06***	-0.02***	-0.04***	-0.06***	-0.06	-0.10*	-0.06
	(-2.80)	(-3.13)	(-3.19)	(-4.34)	(-4.56)	(-4.76)	(-1.57)	(-1.84)	(-0.95)
N	29,456	26,330	23,280	29,456	26,330	23,280	9,398	7,816	6,479
R^2	0.958	0.940	0.934	0.983	0.977	0.975	0.843	0.812	0.806
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Subcategories of Scope 3 Emissions

	$CarbonEmission_{i,t+1}^{Scope3}$						
	$C1+2$	$C3$	$C4$	$C5$	$C6$	$C7$	$C8$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$AIWorker$	-0.00	-0.01	-0.10**	-0.05*	-0.03**	-0.03	-0.14*
	(-0.04)	(-0.43)	(-2.42)	(-1.67)	(-2.07)	(-1.46)	(-1.68)
N	8,645	10,004	9,717	9,904	11,276	9,586	7,642
R^2	0.782	0.764	0.742	0.763	0.793	0.716	0.612
	$C9$	$C10$	$C11$	$C12$	$C13$	$C14$	$C15$
	(8)	(9)	(10)	(11)	(12)	(13)	(14)
$AIWorker$	-0.04	-0.27***	-0.04	0.03	-0.10	-0.33***	-0.07
	(-0.67)	(-3.12)	(-0.63)	(0.46)	(-0.98)	(-3.09)	(-0.77)
N	7,642	7,374	8,473	8,238	7,416	7,180	7,434
R^2	0.612	0.701	0.773	0.766	0.633	0.731	0.630
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table IA5: **AI Worker Shares and Scope 1 Emissions: Robustness Checks**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are the logarithm of Scope 1 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 logarithm of carbon emissions. *ITWorker, RobotWorker, DataWorker* are the share of employees in related roles based on the LinkedIn profiles. *GreenPatents* is the share of green patents. *EmissionTarget* is an indicator for firm-years with approved science-based or net-zero emission targets from the MSCI. *CCExposure* is the firm-level climate change exposure derived from earnings calls provided by Sautner et al. (2023). *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. 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.

	$CarbonEmission^{Scope1}$								
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Confounding Factors - Non-AI Skills									
$AIWorker$	-0.02***	-0.04***	-0.06***	-0.03***	-0.05***	-0.06***	-0.02***	-0.04***	-0.06***
	(-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									
$AIWorker$	-0.03***	-0.05***	-0.06***	-0.02**	-0.04***	-0.05**	-0.03***	-0.05***	-0.07***
	(-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									
$AIWorker$	-0.03***	-0.04***	-0.06***	-0.03***	-0.05***	-0.06***	-0.02***	-0.04***	-0.06***
	(-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 Proxies									
$AIWorker$	-0.03***	-0.05***	-0.07***	-0.03***	-0.05***	-0.07***	-0.03***	-0.05***	-0.07***
	(-3.40)	(-3.43)	(-3.85)	(-3.81)	(-3.96)	(-4.44)	(-3.64)	(-4.10)	(-4.61)
$AIPatent$	0.00	-0.01	-0.01						
	(0.33)	(-0.78)	(-1.03)						
$AIInfras$				-1.10	0.70	5.05			
				(-0.37)	(0.16)	(0.99)			
$AIStrategies^{ECC}$							-0.01	-0.01	-0.01
							(-0.85)	(-1.26)	(-1.09)

(Continued)

(Continued)

Panel E: 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 F: Selective Disclosure and LinkedIn Coverage									
	Exclude Security Sectors			Exclude Trade Secret Sectors			LinkedIn Coverage		
$AIWorker$	-0.02***	-0.04***	-0.06***	-0.02***	-0.04***	-0.05***	-0.02***	-0.04***	-0.06***
	(-3.99)	(-4.08)	(-4.27)	(-3.52)	(-3.64)	(-3.58)	(-3.86)	(-4.01)	(-4.20)
$LinkedInCoverage$							-0.05***	-0.04	-0.02
							(-3.05)	(-1.56)	(-0.77)
Panel G: Alternative Specifications									
	Plus Country-Year FE			SE Cluster by Firm and Year			SE Cluster by Industry		
$AIWorker$	-0.03***	-0.04***	-0.06***	-0.02***	-0.04***	-0.06***	-0.02***	-0.04***	-0.06***
	(-3.89)	(-3.98)	(-4.12)	(-3.84)	(-3.78)	(-3.78)	(-4.98)	(-5.66)	(-5.07)

Table IA6: **AI Worker Shares and Scope 1 Emissions: IV Robustness**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are changes in the logarithm of Scope 1 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 instrumental variable is the excess H-1B visa selections a firm receives due to the lottery in year t-1, which are weighted by the share of visa applications for STEM positions. Panel A excludes firms in the Energy and Utilities sectors. Panel B uses an alternative instrumental variable, which 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) to examine the effect on AI user firms. The sample consists of U.S. firms from 2010 to 2023. t-statistics reported in parentheses are based on standard errors clustered by the firm. *** p<0.01, ** p<0.05, * p<0.1.

	Panel A: Excluding Energy and Utilities			
	$\Delta AIWorker$	$\Delta CarbonEmission^{Scope1}$		
		$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$
	(1)	(2)	(3)	(4)
$AIExcessVisa$	0.14***			
	(4.33)			
$\Delta AIWorker$		-0.07***	-0.09**	-0.10*
		(-2.62)	(-2.25)	(-1.74)
N	9,077	7,798	6,579	5,407
F-Stat	18.75	19.63	18.17	15.26
Control	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
	Panel B: Alternative Instrument			
	$\Delta AIWorker$	$\Delta CarbonEmission^{Scope1}$		
		$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$
	(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

Table IA7: **AI Worker Shares and Sales Growth**

This table analyzes the relationship between sales growth and AI worker share. The dependent variables are the sales growth in years t+1, t+2, and t+3. Control variables include $log(Asset)$, ROA, $Capex/Assets$, $Leverage$, $Tangibility$, $SalesGrowth$. 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 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: Deflated Sales Growth

	$SalesGrowth^{Deflated}$		
	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)
$AIWorker \times EnergyUtil$	0.00	0.03*	-0.04
	(0.09)	(1.76)	(-1.43)
$AIWorker$	-0.01	0.01	-0.01
	(-1.50)	(0.98)	(-1.36)
N	17,733	15,714	13,793
R^2	0.618	0.599	0.662
Control	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes

Panel B: Scale Response by Tangibility and Downstream Demand

	$SalesGrowth$					
	$t+1$	$t+2$	$t+3$	$t+1$	$t+2$	$t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$AIWorker \times Tangibility$	0.01**	0.01*	0.00			
	(1.97)	(1.65)	(0.45)			
$AIWorker \times DownstreamDemand$				0.00	0.01**	0.00
				(0.39)	(2.18)	(0.92)
$AIWorker$	0.00	0.00	0.00	-0.00	0.00	-0.00
	(0.42)	(1.20)	(0.20)	(-0.06)	(0.12)	(-0.54)
N	29,134	26,230	23,014	28,793	25,926	22,761
R^2	0.383	0.372	0.363	0.385	0.373	0.365
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 C: Comparing Sector Characteristics

	$EnergyUtil$	$Others$	$Difference$	t-$stat$
	(1)	(2)	(3)	(4)
$Tangibility$	0.60	0.22	0.38	96.01
$DownstreamDemand$	0.74	0.56	0.17	59.28

B NLP Details for AI Worker Detection

B.1 NLP Method Description

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

I fine-tune the LLaMA model within a teacher-student learning framework (Ni et al., 2024; Dettmers et al., 2023; Wadhwa et al., 2024). Specifically, I 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 710 and 577 for the two tasks, respectively). For each sample, GPT-4.1 provides both a classification label and a one-sentence rationale for its decision. I 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.

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

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

in the same manner as the training data (with 240 positive samples for each of the two tasks). GPT-4.1 delivers consistently high performance for both tasks, with F1-scores of 0.97 and 0.96, which suggests its suitability to serve as a teacher; meanwhile, the fine-tuned LLaMA3 model closely approximates GPT-4.1's performance and achieves a F1-score of 0.94. The results demonstrate that LLaMA3 effectively learns from GPT-4.1 and transfers its classification ability to the two tasks.

Importantly, the fine-tuned LLaMA3 model offers two advantages over GPT-based models. 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 inference. Second, the fine-tuned model is substantially more computationally efficient, enabling large-scale annotation in less than 2.5% of GPT-4.1's runtime for my tasks (the latter exceeds 100 hours per task).

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

	Accuracy	Precision	Recall	F1
Task 1: 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
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

B.2 Keywords for Preliminary AI Worker Detection

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:  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, 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.

*** 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.
```

Table IA B2: **Excerpts of AI Worker Roles**

Panels A and B show LinkedIn job descriptions for AI roles in internal operations and market-facing functions.

Panel A: Internal-Operation AI Workers
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.
Panel B: Market-Facing AI Workers
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.

C Instrumental Variable Construction

I 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 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 AI-related PhDs have gone to non-citizens since 2003, compared to about 20% in other fields. The 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 supply of advanced AI expertise.

[2] See the White House report.

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. 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 my 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 USCIS for each 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 2014 to 2024, 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

I construct firm-level measures of excess H-1B visa selections using multiple data sources. For fiscal years after 2021, I 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.

For years before 2021, I estimate demand and supply using LCA and I-129 petition data, following the methodology of Chen et al. (2021) and Dimmock et al. (2022). I 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, I infer this from the timing of the filing. Specifically, I 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 within 180 days of certification). Certifying too early may trigger an earlier start date and reduce the effective

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

[4]See LCA data from the Office of Foreign Labor Certification.

employment period.

I 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] I include 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. I manually match them to Compustat identifiers and then verify matches using the EIN suffix. A limitation is that the data do not include lottery selections that never resulted in a petition (e.g., selected registrations where the employer chose not to file Form I-129). Hence, this measure serves as a proxy for realized supply.

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

$$AIExcessVisa_{i,t} = (VisaSelections_{i,t} - VisaDemand_{i,t} \times SelectRate_t) \times STEMShareLCA_{i,t}.$$

$VisaSelections_{i,t}$ 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 and proxied using submitted new cap-subject I-129 petitions for pre-2021 years. $VisaDemand_{i,t}$ is the total registrations submitted by the firm. This is also directly available after 2021 and 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, I scale by the firm-year share of STEM workers in related LCAs based on their SOC codes. This refinement concentrates variation on positions most relevant to AI adoption. I thus use $AIExcessVisa_{i,t}$ to serve as a supply-side instrument to identify the effect of foreign AI talent on subsequent emission outcomes.

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

Table IA C1: **AI Worker Shares and Scope 1 Emissions: IV Exclusion Restriction**

This table analyzes the relationship between future Scope 1 emissions and AI worker share. The dependent variables are changes in Scope 1 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 instrumental variable is the excess H1B visa selections a firm receives due to the lottery in year t-1, weighted by the share of STEM visa applications. Panel A evaluates AI specificity using alternative instruments and placebo tests (AI-keyword-weighted H-1B shock, Non-STEM-weighted placebo, and non-AI-keyword-weighted placebo), verifies that the lottery shock does not predict other worker share measures, and excludes IT-services firms (SIC-based). Panel B addresses filing endogeneity and within-firm substitution through (i) adding controls for filing intensity (visa demand) and the STEM-share weight, (ii) constructing the instrument using pre-period STEM exposure (2010–2013), (iii) placebo timing tests using leads of the (instrumented) change in AI worker share, (iv) alternative AI measures based on headcount instead of share, (v) excluding firms with substantial foreign exposure (>25% non-U.S. sales), and (vi) restricting the sample to the pre-2021 period (prior to the electronic registration change). 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.

	$\Delta CarbonEmission^{Scope1}$								
	$t \to t+1$	$t \to t+2$	$t \to t+3$	$t \to t+1$	$t \to t+2$	$t \to t+3$	$t \to t+1$	$t \to t+2$	$t \to t+3$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Alternative Instruments and Placebo Tests (AI Specificity)									
	AI-Keyword Share (Text)			Non-STEM Share			Non-AI-Keyword Share		
$\Delta AIWorker$	-0.05**	-0.07*	-0.05	0.23	0.41	0.09	-0.02	-0.06	-0.10
	(-2.09)	(-1.77)	(-1.07)	(0.67)	(0.51)	(0.44)	(-0.27)	(-0.51)	(-0.59)
N	8,578	7,244	5,961	8,495	7,164	5,893	8,495	7,164	5,893
F-Stat	17.01	16.29	15.35	0.90	0.43	3.36	7.64	5.95	4.61
	Predicting Other Worker Shares			Adding Other Worker Share			Excluding IT services		
	IT	Robot	Data						
$AIExcessVisa$	0.02*	0.01	0.02						
	(1.77)	(0.61)	(1.22)						
$\Delta AIWorker$				-0.06**	-0.09**	-0.10*	-0.06**	-0.09**	-0.10*
				(-2.50)	(-2.31)	(-1.85)	(-2.43)	(-2.14)	(-1.76)
$\Delta ITWorker$				-0.00	0.01	0.02			
				(-0.59)	(1.40)	(1.33)			
$\Delta RobotWorker$				0.00	0.01	0.02			
				(0.75)	(1.08)	(1.39)			
$\Delta DataWorker$				0.01	0.02	0.01			
				(1.48)	(1.55)	(1.11)			
N	9,972	9,972	9,972	8,578	7,244	5,961	8,439	7,124	5,860
F-Stat	3.13	0.37	1.49	22.50	20.10	16.31	20.98	20.16	16.44
Panel B: Addressing Endogenous Filing and Within-Firm Substitution									
	Adding Visa Demand			Pre-Period STEM Exposure			Lead AI Worker Share		
$\Delta AIWorker$	-0.11**	-0.14**	-0.16**	-0.05**	-0.10***	-0.12**			
	(-2.39)	(-2.23)	(-1.99)	(-2.15)	(-2.64)	(-2.43)			
$VisaDemand/Emp$	0.01*	0.01	0.01						
	(1.65)	(1.21)	(0.95)						
$STEMShare$	0.02	0.03	0.05						
	(0.84)	(0.81)	(1.04)						
$\Delta AIWorker_{t+1 \to t+2}$							-0.01		
							(-0.45)		
$\Delta AIWorker_{t+2 \to t+3}$								-0.06	
								(-1.01)	
$\Delta AIWorker_{t+3 \to t+4}$									-0.04
									(-0.43)
N	8,504	7,173	5,900	8,578	7,244	5,961	7,240	5,949	4,724
F-Stat	15.61	16.40	14.63	19.53	17.29	15.53	15.47	15.48	12.70
	AI Headcount Proxy			Excluding Big Multinationals			Pre-2021		
$\Delta AIWorker$	-0.04**	-0.06*	-0.06	-0.06**	-0.08**	-0.08	-0.05**	-0.08**	-0.08*
	(-1.99)	(-1.75)	(-1.40)	(-2.43)	(-2.05)	(-1.57)	(-2.11)	(-2.18)	(-1.83)
$\Delta ForeignSaleShare$				0.08	0.02	-0.00			
				(1.64)	(0.39)	(-0.00)			
N	8,578	7,244	5,961	8,510	7,191	5,921	6,752	6,687	5,961
F-Stat	23.17	20.78	22.06	23.31	20.98	17.00	20.46	19.90	18.77
Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

D Theoretical Framework of AI and Emissions

Setup. Let a denote a firm's AI intensity, with a convex adoption cost $C(a)$. The firm produces a single final good by combining brown and green composite inputs X_j in a constant elasticity of substitution (CES) production function with elasticity $\sigma > 1$.

D.1 Cost Minimization

Firm problem. The firm minimizes input cost to produce one unit of output Y:

$$\min_{X_b, X_g} \; p_b X_b + p_g X_g \quad \text{s.t.} \quad \left((A_b X_b)^{\frac{\sigma-1}{\sigma}} + (A_g X_g)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \geq 1.$$

Lagrangian:

$$\mathcal{L} = p_b X_b + p_g X_g + \lambda \left[1 - (A_b X_b)^{\frac{\sigma-1}{\sigma}} - (A_g X_g)^{\frac{\sigma-1}{\sigma}} \right].$$

FOCs w.r.t. X_j, $j \in \{b, g\}$:

$$\frac{\partial \mathcal{L}}{\partial X_j} = p_j - \lambda \frac{\sigma - 1}{\sigma} A_j^{\frac{\sigma-1}{\sigma}} X_j^{-\frac{1}{\sigma}} = 0 \quad \Rightarrow \quad X_j = \left[\frac{\lambda(\sigma-1)}{\sigma} \frac{A_j^{\frac{\sigma-1}{\sigma}}}{p_j} \right]^{\sigma}.$$

Let $K = \frac{\lambda(\sigma-1)}{\sigma}$, which is a constant determined by the constraint.

Plug X_j into the constraint $(A_b X_b)^{\frac{\sigma-1}{\sigma}} + (A_g X_g)^{\frac{\sigma-1}{\sigma}} = 1$,

$$\Rightarrow \quad K = \left(\frac{A_b^{\sigma-1}}{p_b^{\sigma-1}} + \frac{A_g^{\sigma-1}}{p_g^{\sigma-1}} \right)^{-\frac{1}{\sigma-1}}.$$

The unit cost is:

$$c = p_b X_b + p_g X_g = \left(\frac{A_b^{\sigma-1}}{p_b^{\sigma-1}} + \frac{A_g^{\sigma-1}}{p_g^{\sigma-1}} \right) \left(\frac{A_b^{\sigma-1}}{p_b^{\sigma-1}} + \frac{A_g^{\sigma-1}}{p_g^{\sigma-1}} \right)^{-\frac{\sigma}{\sigma-1}} = \left[\left(\frac{p_b}{A_b} \right)^{1-\sigma} + \left(\frac{p_g}{A_g} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

With $A_j(a) = a^\eta \Theta_j(a)$ and $p_b = p_g = 1$,

$$c(a) = \frac{1}{a^\eta}\left[\Theta_b^{\sigma-1}(a) + \Theta_g^{\sigma-1}(a)\right]^{\frac{1}{1-\sigma}}.$$

Let $Z(a) = \Theta_b^{\sigma-1}(a) + \Theta_g^{\sigma-1}(a)$, then $c(a) = \frac{1}{a^\eta} Z(a)^{\frac{1}{1-\sigma}}$. Then

$$\ln c(a) = -\eta \ln a + \frac{1}{1-\sigma} \ln Z(a).$$

The elasticity of unit cost with respect to AI intensity is:

$$\frac{d\ln c}{d\ln a} = -\eta + \frac{1}{1-\sigma}\frac{1}{Z}\frac{dZ}{d\ln a}, \quad \frac{dZ}{d\ln a} = (\sigma-1)\Theta_b^{\sigma-1}\mu_b + (\sigma-1)\Theta_g^{\sigma-1}\mu_g, \quad \mu_j = \frac{d\ln\Theta_j}{d\ln a}.$$

The expenditure share is

$$s_b = \frac{p_b X_b}{p_b X_b + p_g X_g} = \frac{1}{1 + \left(\frac{p_g}{p_b}\right)\cdot\left(\frac{X_g}{X_b}\right)} = \frac{\Theta_b^{\sigma-1}}{\Theta_b^{\sigma-1} + \Theta_g^{\sigma-1}}, \quad s_g = 1 - s_b.$$

$$\Rightarrow \frac{d\ln c}{d\ln a} = -\eta - \frac{\Theta_b^{\sigma-1}\mu_b + \Theta_g^{\sigma-1}\mu_g}{Z(a)} = -\eta - (s_b\mu_b + s_g\mu_g).$$

The intuition is that AI improves overall efficiency by reducing total input needs per unit of output. In addition, AI raises the productivity of specific technologies, and the resulting cost reduction reflects a share-weighted average of these technology-specific improvements. AI lowers costs more when it complements the inputs that account for larger spending shares.

D.2 Expenditure Shares Across Brown and Green Inputs

The relative ratio of expenditure shares across brown and green inputs is:

$$\frac{s_b}{s_g} = \left(\frac{\Theta_b}{\Theta_g}\right)^{\sigma-1} \quad \Rightarrow \quad \ln\frac{s_b}{s_g} = (\sigma-1)\big(\ln\Theta_b - \ln\Theta_g\big).$$

Differentiate w.r.t. $\ln a$:

$$\frac{d}{d\ln a}\big(\ln s_b - \ln s_g\big) = (\sigma-1)(\mu_b-\mu_g).$$

With $d\ln s_b = \dfrac{ds_b}{s_b}$ and $d\ln s_g = \dfrac{d(1-s_b)}{1-s_b} = -\dfrac{ds_b}{s_g}$,

$$\frac{ds_b}{s_b}+\frac{ds_b}{s_g} = (\sigma-1)(\mu_b-\mu_g)\,d\ln a \quad \Rightarrow \quad \frac{ds_b}{d\ln a} = (\sigma-1)\,s_b s_g\,(\mu_b-\mu_g).$$

The intuition is that with $\sigma > 1$ (substitutes), when AI is more complementary to brown inputs ($\mu_b > \mu_g$), the brown share increases.

D.3 Carbon Emissions

Carbon intensity. Given the emission factor e_j for respective inputs, a firm's carbon intensity is defined as the emissions per unit output:

$$\bar{e} = \frac{E}{Y} = \frac{e_b X_b + e_g X_g}{Y} = c\,(e_b s_b + e_g s_g).$$

This is because $s_j = \frac{p_j X_j}{p_b X_b + p_g X_g} = \frac{p_j X_j}{cY}$, thus $\frac{X_j}{Y} = s_j c$ for $j \in \{b, g\}$ when $p_j = 1$.

The elasticity of carbon intensity w.r.t. *ln* a:

$$\begin{aligned}\frac{d\ln\bar{e}}{d\ln a} &= \frac{d\ln c}{d\ln a} + \frac{1}{e_b s_b + e_g s_g}\left(e_b\frac{ds_b}{d\ln a} + e_g\frac{ds_g}{d\ln a}\right) = \frac{d\ln c}{d\ln a} + \frac{e_b - e_g}{e_b s_b + e_g s_g}\cdot\frac{ds_b}{d\ln a}\\ &= -\eta - (s_b\mu_b + s_g\mu_g) \;+\; (\sigma-1)\,s_b s_g\,(\mu_b-\mu_g)\,\Delta_e, \qquad \Delta_e = \frac{e_b - e_g}{e_b s_b + e_g s_g}.\end{aligned}$$

The first part, $-\eta - (s_b\mu_b + s_g\mu_g)$, is the efficiency term. Better prediction enabled by AI reduces the total inputs required per unit of output, so emissions per unit of output fall correspondingly with a scale of η. Meanwhile, $-(s_b\mu_b + s_g\mu_g)$ implies that AI can raise the productivity of particular technologies; when AI helps the technologies that account for a larger share of spending, unit inputs decline more, further lowering emission intensity. The second part, $(\sigma-1)s_b s_g(\mu_b-\mu_g)\Delta_e$, is the directed term. With substitutable inputs ($\sigma>1$), AI that is more complementary to brown activities ($\mu_b>\mu_g$) shifts the cost-minimizing composition toward the brown inputs. The resultant reallocation raises emissions per unit of output in proportion to the emission gap Δ_e. If AI is instead more complementary to green inputs ($\mu_g>\mu_b$), the same mechanism works in the opposite direction and intensifies the decline in emission intensity. Whether intensity ultimately falls or rises depends on the balance between efficiency and directed terms.

Output response. For the product with price p, the CES demand implies

$$Y = Dp^{-\xi}, \qquad \xi > 1,$$

where a 1% increase in price makes quantity fall by ξ%. D can be viewed as everything that determines the quantity, holding the fixed price. With monopolistic competition, the firm sets a constant markup over marginal cost (unit cost) c: $p=\varphi c$, with $\varphi=\frac{\xi}{\xi-1}$. This is because the firm maximizes profits $\pi=(p-c)Y$ subject to $Y=Dp^{-\xi}$. The FOC w.r.t. p is

$$\frac{d\pi}{dp} = Y + (p-c)\frac{dY}{dp} = 0 \quad \Rightarrow \quad 1 + \frac{p-c}{p}\frac{d\ln Y}{d\ln p} = 0 \Rightarrow \quad \frac{p}{c} = \varphi = \frac{\xi}{\xi-1}.$$

Hence φ is constant and $d\ln p = d\ln c$. Using the chain rule through price:

$$\frac{d\ln Y}{d\ln a} = \frac{d\ln Y}{d\ln p}\cdot\frac{d\ln p}{d\ln a} = -\xi\left(\frac{d\ln c}{d\ln a}\right) = \xi\left[\eta + s_b\mu_b + s_g\mu_g\right].$$

Thus, AI that makes production cheaper induces a proportional price cut and, through elastic demand, an increase in scale. Markets with more elastic demand (larger ξ) exhibit a stronger output response to the same technological improvement.

Emission levels. The emission level elasticity follows from $E = \bar{e} \cdot Y$:

$$\frac{d \ln E}{d \ln a} = \frac{d \ln \bar{e}}{d \ln a} + \frac{d \ln Y}{d \ln a}.$$

Hence, the total emission response to AI adoption depends on both the change in carbon intensity (efficiency and directed effects) and the resulting expansion in output (scale effect).

D.4 Optimal AI Intensity

While the firm minimizes cost to determine the input composition for a given output, it chooses AI intensity a that maximizes profit as a affects both production costs and the scale of sales through its impact on demand. The cost of AI adoption is $C(a)$. The profit is:

$$\Pi(a) = pY - cY - C(a) = (\frac{\xi}{\xi - 1} - 1)c \cdot Y - C(a).$$

FOC w.r.t. a is:

$$\frac{d\Pi}{da} = (\frac{\xi}{\xi - 1} - 1)\big[c\frac{dY}{da} + Y\frac{dc}{da}\big] - C'(a) = -Y\frac{dc}{da} - C'(a) = 0.$$

Using $-\frac{dc}{da} = \frac{c}{a}\big[\eta + s_b\mu_b + s_g\mu_g\big]$, the FOC is

$$C'(a) = \frac{cY}{a}\big[\eta + s_b\mu_b + s_g\mu_g\big].$$

The economic intuition is that the firm chooses a so that marginal adoption cost equals the marginal increase in operating profit generated by the AI-induced cost reduction.

E Narrative Evidence for Sectoral Emission Effects

The AI investment measure is grounded in job descriptions and contains rich narratives about how firms actually integrate AI into operations. To retrieve relevant text and shed light on the underlying economic mechanisms for emission outcomes, I identify the 50 firms with the largest influence on each sector's AI coefficient in emission regressions (in Figure 4) and create word clouds to visualize the most frequently occurring terms in their AI worker descriptions. 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. Figure 5 displays word clouds for four sectors with diverse carbon profiles and AI integration, including Consumer Durables, Wholesale and Retail, Energy, and Utilities. Figure IA E1 displays the word clouds for the remaining eight sectors. Here, I provide job description excerpts from highly influential firms in the four selected sectors.

E.1 Consumer Durables

<u>General Motors, Advanced Radar Group, May 2020 - Mar 2022.</u>

Led the development of a radar software stack, incorporating advanced algorithms and simulations to optimize data processing for Advanced Driver Assistance Systems features and machine learning-based perception software. Conduct in depth business analysis with respect to cost, time, and quality.

<u>Ford Motor Company, Software Analysts, Dec 2011 - Feb 2023.</u>

Build a model to predict vehicle and sub-system performance as functions of input vectors and parameters to guide future product design. Build a decision tree model for the prediction of car quality given other attributes about the car. Shifting from modeling physics simply moves the burden of extracting the physics onto the data.

E.2 Wholesale and Retails

Nordstrom Inc., Lead Product Manager, Jun 2015 - Aug 2017.

Built a machine learning algorithm to optimize inventory consolidation and operational readiness planning. Led supply chain integration between Nordstrom and Hautelook to realize an average annual benefit of $15M in transfer costs and off-price demand growth.

Zalando SE, Cloud-Data-Engineer, Feb 2018 - Sep 2018.

Creation of a daily running machine learning pipeline to predict (1) the return of shipped goods (return rate) and (2) the cancellation of an order before shipping (cancellation rate). Building the architecture for Zalando's logistics algorithms and optimizing how customer orders are processed together to improve the productivity of the warehouse.

E.3 Energy

Schlumberger Ltd., General Field Engineer, May 2021 - Sep 2021.

The industry's first AI-driven autonomous wireline logging deployment to Kuwait. Successfully building artificial neural network models, which were used in the estimation of a synthetic photoelectric factor log, permeability and reservoir facies. Worked on an innovative project to perform reservoir characterization with integrated static and dynamic uncertainties using machine learning for a giant mature oil field.

Halliburton, Formation and Reservoir Solutions Manager, Oct 2011 - Nov 2017.

Utilization of the most advanced technologies, such as machine learning, analytics, and cloud, around the completion and production phases of the well life cycle. Developed signal analysis algorithms and implemented machine learning with Tensorflow to detect acoustic wavelets. This improved the analysis process done by the geophysicist team and reduced turnaround times by up to 30%.

E.4 Utilities

Tata Power, Engineering Group, Jan 2017 - Dec 2019.

Working as a group head for the implementation of reliability-centered maintenance across generation, transmission & distribution. Implementing AI and ML-based OT & IoT solutions for performance improvement of thermal & renewable power assets. Worked on forecasting of electrical load needed (day ahead) for the west and central Delhi areas using data from previous years with the help of artificial neural networks.

Enbridge Inc., Senior Advisor, Aug 2019 - Mar 2020.

Strategic advisor between the lab and core business functions on initiatives focused on the development of advanced technological solutions to enhance process efficiency, operations reliability, and safety, utilizing modern techniques (AI, ML, etc.) in an Agile fashion. Analyzed drag-reducing agent usage for day-to-day pipeline operations that maximize revenues.

Figure IA E1: **AI Worker Shares 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 Scope 1 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–H corresponds to Consumer Nondurables, Manufacturing, Chemicals, Computers, Telephone Transmission, Healthcare, Finance, and Others.

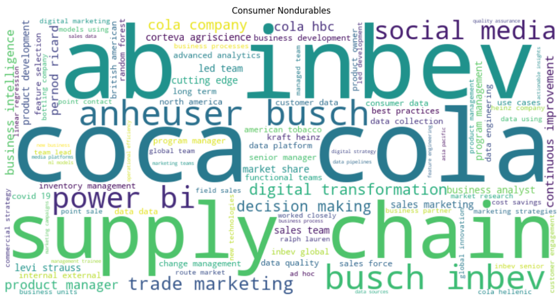

(a)

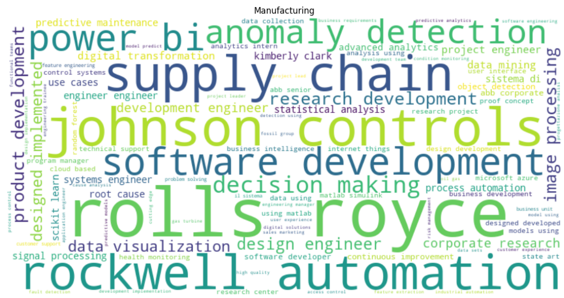

(b)

(c)

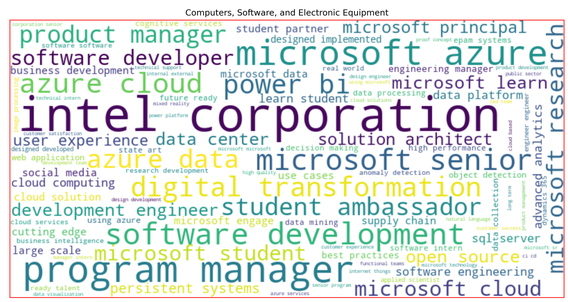

(d)

(e)

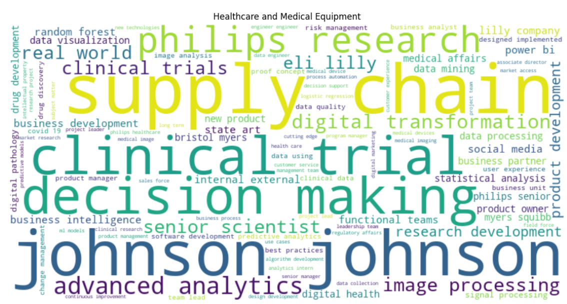

(f)

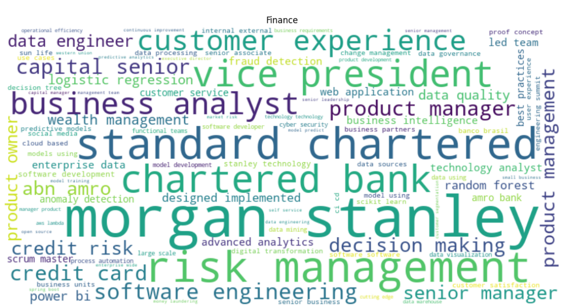

(g)

(h)

F AI Worker Shares and Electricity Consumption

Although the paper focuses on AI's emissions impact through operational changes at user firms, a separate concern is the electricity required to train and run AI models. As a complementary analysis, I also examine emissions through this electricity channel.

This analysis relies on firm-level electricity consumption data from CDP and thus limits the sample size. Nevertheless, it covers major technology firms and semiconductor providers in the AI ecosystem (see Table IA F1). I further restrict the sample to firms reporting energy consumption under the operational control boundary (guided by GHG Protocol). This boundary assigns electricity use to the entity that operates the facility rather than to the one that merely owns it. For example, electricity from a leased data center is reported by the lessee if it directs operations, but would be excluded under the financial control approach unless the lessee holds majority ownership. Using operational control is thus critical to ensure that reported emissions reflect the firm's own operations. In the sample, approximately 79% of firms use the operational control reporting standard. Panel B of Table 1 shows that, on average, the 816 sample firms disclose AI infrastructure in 28% of firm-years and employ 5.33 AI workers per 10,000 employees.

Given that AI infrastructure, like data centers and server rooms, is the input that largely determines the energy profile of computing workloads (IEA, 2025), I construct a firm-level measure of AI infrastructure involvement and include it as a control beyond AI worker share to isolate the workforce effect. I build this measure using the same two-step NLP procedure applied to annual reports. The classifier performs well out of sample, with GPT-4.1 achieving an F1 score of 0.98 and the fine-tuned LLaMA3 achieving 0.97 on a test set of 500 manually annotated samples (keywords, prompts, and sample excerpts are provided below). Specifically, I leverage 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, I 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. Within 10-K filings, I analyze 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, I analyze their full annual reports (as they do not follow the same standardized structure). I identify paragraphs that discuss AI infrastructure and then construct a firm-year measure defined as the proportion of such paragraphs. This measure captures the degree to which a firm emphasizes AI-related computing facilities and serves as a proxy for its infrastructure capacity.

Keywords for Preliminary AI Infrastructure Detection:

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.

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

```
*** Task:  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.
```

Table IA F1: **Key AI-investing Firms with Electricity Consumption Data**

This table lists the 30 largest AI-investing firms in the sample with electricity consumption data from CDP, ranked by their average number of AI workers.

Company	$\#AIWorker$	$AIWorker$	$AIInfras$	$Electricity$	$Electricity^{NonRenew}$	$CarbonEmission^{Scope1}$
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 IA F2: **Excerpts of AI Infrastructure Disclosure**

This table presents excerpts from annual reports on AI infrastructure.

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.

Table IA F3: **AI Worker Shares and Electricity Consumption: Energy Sources**

This table analyzes the relationship between changes in electricity consumption and AI worker share. The dependent variables are changes from year t to $t+1$, $t+2$, and $t+3$ in renewable and non-renewable electricity consumption (Panel A), and location- and market-based Scope 2 emissions (Panel B). 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: Renewable vs. Non-Renewable Electricity

	$\Delta Electricity^{Renew}$			$\Delta Electricity^{NonRenew}$		
	$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$	$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta AIWorker$	0.02	0.05*	0.06*	-0.04	-0.05	-0.05
	(1.32)	(1.78)	(1.91)	(-1.51)	(-1.28)	(-1.08)
$\mathbb{1}(AIInfras)$	0.18**	0.15	0.13	0.07	0.20**	0.22*
	(2.25)	(1.20)	(0.83)	(1.09)	(2.18)	(1.96)
N	1,769	1,344	1,023	2,470	1,946	1,564
R^2	0.257	0.416	0.466	0.055	0.113	0.161
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

	$\Delta CarbonEmission^{S2(Loc)}$			$\Delta CarbonEmission^{S2(Mkt)}$		
	$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$	$t \rightarrow t+1$	$t \rightarrow t+2$	$t \rightarrow t+3$
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta AIWorker$	-0.01	0.00	0.01	0.01	0.01	-0.01
	(-0.90)	(0.10)	(0.51)	(0.50)	(0.43)	(-0.29)
$\mathbb{1}(AIInfras)$	0.09***	0.16***	0.20**	0.06	0.11	0.19
	(2.80)	(2.77)	(2.43)	(0.87)	(0.99)	(1.11)
N	2,699	2,149	1,741	2,252	1,772	1,424
R^2	0.081	0.119	0.157	0.053	0.125	0.161
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 IA F4: **AI Worker Shares and Electricity Consumption: Supply Chain Spillover**

This table examines the relationship between energy consumption and AI investment across the supply chain. Panel A examines how customers' AI investment 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.

	$AI\,Users'\,Suppliers$			$AI\,Providers'\,Customers$		
	$AIInfra^{S}$	$Electricity^{S}$	$CarbonEmission^{S1,S}$	$AIWorker^{C}$	$Electricity^{C}$	$CarbonEmission^{S1,C}$
	$t+1$	$t+1$	$t+1$	$t+1$	$t+1$	$t+1$
	(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

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.

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

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*.

Pairolero, N. A., Giczy, A. V., Torres, G., Erana, T. I., Finlayson, M. A., and Toole, A. A. (2025). The artificial intelligence patent dataset (aipd) 2023 update. *The Journal of Technology Transfer*.

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.