ARTIFICIAL INTELLIGENCE AND URBAN CARBON-EMISSION INTENSITY: PRODUCTION-SIDE EVIDENCE FROM CHINA

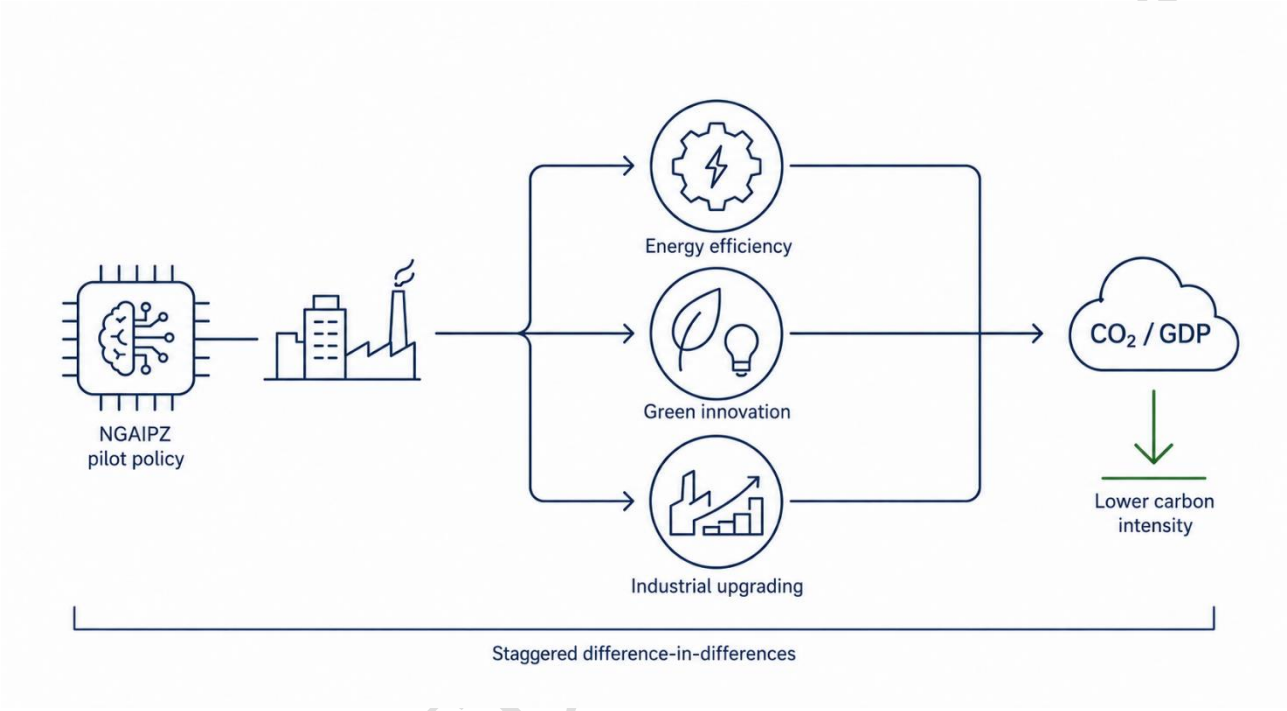
Liping Wang[1]; Paijie Wan[2*]; Hongjie Zhang[1]

[1] College of Management, Ocean University of China, Qingdao 266100, China

[2] School of Public Administration and Policy, Shanghai University of Finance and Economics, Shanghai 200433, China

* Corresponding author: 2025310090@stu.sufe.edu.cn

Graphical Abstract



ACCEPTED

Abstract. This paper examines how artificial intelligence affects urban carbon-emission intensity from a production-side perspective. We exploit the staggered rollout of China's National New Generation Artificial Intelligence Innovation and Development Pilot Zone (NGAIPZ) program and use a balanced panel of 288 prefecture-level cities from 2014 to 2023. Difference-in-differences estimates show that the NGAIPZ program significantly reduces urban carbon-emission intensity. Further results provide evidence consistent with potential production-side channels involving improved energy efficiency, greater green innovation, and industrial upgrading. The carbon-intensity effect is stronger in cities with weaker environmental regulation, indicating that intelligent production and management systems create larger reductions in carbon intensity where conventional regulatory pressure is limited. The effect is also more pronounced in cities with better internet infrastructure, which supports data circulation and the diffusion of AI applications. These findings indicate that AI-oriented policy can improve urban carbon efficiency alongside changes in energy use, green innovation, and the organization of production.

Keywords: *Artificial intelligence; Carbon-emission intensity; Energy efficiency; Green innovation; Difference-in-differences*

ACCEPTED MANUSCRIPT

Introduction

Limiting global warming requires a rapid decline in greenhouse-gas emissions, yet current mitigation efforts remain insufficient to meet established climate targets (IPCC, 2023). Global energy-related carbon dioxide emissions reached a new high in 2023, with China accounting for approximately 35% of the world total (IEA, 2024). China therefore faces the difficult task of reducing CO_2 emissions while maintaining the productive capacity of its cities, where industrial production, infrastructure use, and technological change are concentrated (Dong et al., 2024; Zhong et al., 2024). Identifying technologies that can lower urban emissions without constraining economic activity has consequently become an important empirical question (Ding et al., 2023; Guo et al., 2025).

Artificial intelligence has become a prominent candidate for reconciling economic growth with decarbonization, although its environmental consequences remain uncertain (Kaack et al., 2022; Rolnick et al., 2022). Machine learning can improve demand forecasting, system control, predictive maintenance, and the matching of energy supply with consumption, thereby reducing waste in production, buildings, transport, and power systems (Rolnick et al., 2022; Kaack et al., 2022). AI also expands the capacity to search complex technological spaces and has become increasingly embedded in climate-related inventions (Verendel, 2023). Its training and operation require substantial computing infrastructure and electricity, while productivity gains may expand output and generate rebound effects (Kaack et al., 2022; Lee and Yan, 2024). The net effect on CO_2 emissions depends on whether improvements in production, energy use, and technology adoption exceed the additional energy demand generated by intelligent production and digital infrastructure (Lee and Yan, 2024; Zhao et al., 2024).

Recent empirical studies generally associate AI development with lower emissions and cleaner production, although the estimated effects vary across technologies, sectors, and regions. Industrial robot applications have been linked to carbon reduction and improved green productivity, with the magnitude of these effects shaped by absorptive capacity, industrial characteristics, and the surrounding environmental system (Li et al., 2022; Yang and Liu, 2024). Firm-level evidence further

ACCEPTED MANUSCRIPT

shows that AI and robotics can improve energy and resource use by changing production techniques and input allocation (Li et al., 2023). Provincial and city-level studies document reductions in CO_2 emissions, gains in energy efficiency, and improvements in energy-environmental performance as AI develops (Ding et al., 2023; Wang et al., 2024a; Li et al., 2025; Guo et al., 2025). Other work identifies nonlinear effects, regional variation, and changes in the distribution of carbon emissions, indicating that the environmental benefits of AI cannot be assumed to arise uniformly (Lee and Yan, 2024; Dong et al., 2024; Zhao et al., 2024). Much of this evidence relies on robot penetration, patent counts, textual indicators, or composite AI indices, which capture different dimensions of intelligent development and may also reflect the endogenous concentration of technology-intensive activity (Ding et al., 2023; Wang et al., 2024b; Guo et al., 2025).

Three production-side channels may connect AI development with lower urban carbon-emission intensity. First, intelligent sensing, forecasting, and automated control can reduce idle energy use and improve the allocation of electricity and other inputs across production processes (Li et al., 2023; Wang et al., 2024a; Li et al., 2025). AI can also facilitate the substitution of cleaner energy for carbon-intensive inputs by improving the flexibility and coordination of energy systems (Tao et al., 2024). Second, AI lowers the cost of processing technical information and conducting experiments, which can raise the quantity and efficiency of green innovation (Lee et al., 2022; Tian et al., 2023; Wang et al., 2024b). Evidence from China further shows that AI pilot-zone policies promote firm-level green innovation through technology adoption and operational improvement (Mijit et al., 2025). Third, intelligent technologies can change the composition of local production by increasing the productivity of knowledge-intensive activities and supporting the digital upgrading of traditional industries (Yang and Liu, 2024; Lyu et al., 2024). Existing research commonly studies these channels in isolation, leaving limited evidence on whether a common AI-oriented policy is associated with adjustments in all three dimensions of urban production.

China's National New Generation Artificial Intelligence Innovation and Development Pilot Zone (NGAIPZ) program provides a useful setting for examining this question. The program supports local

ACCEPTED MANUSCRIPT

demonstrations of AI technology, policy and social experimentation, deeper integration between AI and economic activity, and the development of replicable approaches to intelligent growth (Ministry of Science and Technology of China, 2020). The program combines technological application, industrial support, infrastructure development, and institutional experimentation within a unified policy framework (Ministry of Science and Technology of China, 2020). The successive rollout across cities generates policy variation over time and across locations, offering a basis for identifying the environmental consequences of locally supported AI development.

The study closest to ours is Zhang et al. (2025), which uses the same NGAIPZ program and finds that it reduces total urban carbon emissions. Our analysis extends this work by examining changes in production-related outcomes that may accompany the carbon-intensity response to AI-oriented policy. Within the same policy framework, we examine changes in energy intensity and efficiency, green patenting, and industrial upgrading, capturing complementary adjustments in energy use, technological activity, and the sectoral composition of output (Li et al., 2023; Tao et al., 2024; Mijit et al., 2025). We also investigate whether environmental regulation and internet infrastructure shape the policy effect, since these factors influence the incentives and technical capacity associated with differences in the carbon-intensity response to policy (Meng et al., 2024; Lyu et al., 2024).

Using a balanced panel of 288 Chinese prefecture-level cities from 2014 to 2023, we estimate the effect of the staggered rollout of the NGAIPZ program on urban carbon-emission intensity. The results show that the NGAIPZ program significantly reduces carbon-emission intensity. Additional estimates show improvements in energy efficiency, green innovation, and industrial upgrading, consistent with potential production-side channels. The effect is stronger in cities with weaker environmental regulation and in those with better internet infrastructure.

This paper contributes to the literature in three respects. First, it provides policy-based evidence on the carbon-intensity effects of AI-oriented policy by exploiting the staggered rollout of the NGAIPZ program across Chinese cities (Ding et al., 2023; Zhong et al., 2024; Zhang et al., 2025). Second, it brings energy use, green innovation, and industrial transformation into a unified empirical framework,

ACCEPTED MANUSCRIPT

providing channel-consistent evidence on how AI policy relates to production-side adjustments and urban carbon efficiency (Li et al., 2023; Tao et al., 2024; Mijit et al., 2025). Third, it identifies the local conditions that shape the environmental returns to the NGAIPZ program. The contrasting roles of environmental regulation and internet infrastructure indicate that carbon-intensity responses differ with the scope for operational improvement and the digital capacity to support intelligent applications (Lyu et al., 2024; Meng et al., 2024).

The remainder of the paper is organized as follows. The next section develops the theoretical framework and research hypotheses. The Materials and Methods section presents the econometric model, variable definitions, and data sources. The Results section reports the baseline estimates, dynamic and treatment-timing-robust results, placebo test, and robustness checks. The Mechanism Analysis section examines energy efficiency, green innovation, and industrial upgrading, while the Heterogeneity Analysis section considers differences across environmental regulation and internet infrastructure. The final section summarizes the main findings, discusses policy implications, and outlines the limitations and directions for future research.

Theoretical Analysis and Hypothesis Development

Artificial intelligence changes urban production by expanding the capacity to collect information, generate predictions, and coordinate decisions across firms and public systems. As a general-purpose technology, AI can raise productivity and stimulate product innovation, but these gains depend on complementary investments in organizational processes, skills, and intangible capital (Brynjolfsson et al., 2021; Babina et al., 2024; Czarnitzki et al., 2023). Automation and algorithmic assistance can also alter task allocation and production scale, so their environmental consequences depend on how firms reorganize production after adoption (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Brynjolfsson et al., 2025). China's national AI strategy provides the broader policy foundation for intelligent development (State Council of China, 2017). Within this framework, the NGAIPZ program combines technological demonstration, industrial support, infrastructure provision, and

ACCEPTED MANUSCRIPT

institutional experimentation, creating a coordinated setting in which AI applications can be linked to urban emission reduction (Ministry of Science and Technology of China, 2020).

The NGAIPZ program creates a local environment for the wider application of AI. NGAIPZ cities are expected to develop application scenarios, strengthen local innovation systems, and promote the integration of intelligent technologies with existing industries (Ministry of Science and Technology of China, 2020). AI can influence innovation and industrial organization more broadly when technology developers interact closely with producers, research institutions, and users (Cockburn et al., 2019; Babina et al., 2024). The resulting changes in production and resource allocation provide three potential routes to lower urban carbon-emission intensity: improved energy efficiency, greater green innovation, and industrial upgrading.

Improving Energy Efficiency

Energy use is shaped by information about demand, equipment conditions, production schedules, and the availability of inputs. Incomplete information, uncertainty over future savings, and organizational frictions can prevent firms from adopting cost-effective energy-saving practices (Jaffe and Stavins, 1994; Allcott and Greenstone, 2012; Martin et al., 2012). AI reduces these frictions by processing high-frequency operating data and producing forecasts that can be updated as conditions change (Czarnitzki et al., 2023; Brynjolfsson et al., 2025). More accurate forecasts allow firms and utilities to align energy supply with actual demand, limiting idle consumption and unnecessary reserve capacity.

Intelligent monitoring and automated control also improve the management of equipment and production lines. Digital management practices can reveal inefficient energy use, while automation increases the precision and consistency of production processes (Martin et al., 2012; Graetz and Michaels, 2018). AI-assisted systems can identify equipment deterioration, schedule maintenance, and adjust production tasks before faults generate substantial energy losses. Evidence on automation and firm productivity further shows that the returns to new technologies arise through changes in task allocation and production organization (Acemoglu and Restrepo, 2020; Czarnitzki et al., 2023).

ACCEPTED MANUSCRIPT

The NGAIPZ program can magnify these gains by supporting digital infrastructure, application platforms, and demonstration projects that lower the cost of adopting intelligent systems. The productivity effects of a new general-purpose technology depend on complementary investments and process redesign, both of which can be facilitated by coordinated local policy support (Brynjolfsson et al., 2021; Ministry of Science and Technology of China, 2020). As AI applications spread, firms can produce a given level of output with less electricity and lower aggregate energy input. Studies of manufacturing emissions show that improvements in production technology and productivity account for a substantial share of long-run pollution reductions (Levinson, 2009; Shapiro and Walker, 2018).

Promoting Green Innovation

Green innovation requires firms to search across technical alternatives, evaluate uncertain outcomes, and adapt new solutions to existing production processes. AI can expand the range of technical information processed by researchers and assist with prediction, design, and the screening of potential inventions (Cockburn et al., 2019; Babina et al., 2024). These capabilities reduce the cost of experimentation and may increase the productivity of research activities. The directed-technical-change literature shows that shifts in the relative returns to clean and dirty technologies can redirect inventive activity toward environmentally beneficial innovation (Popp, 2002; Acemoglu et al., 2012). The NGAIPZ program strengthens this innovation process through local coordination. Policy support can connect AI firms with manufacturers, universities, research institutes, and public agencies, allowing technical knowledge to move more quickly between developers and users (Ministry of Science and Technology of China, 2020). Evidence from environmental policy shows that targeted incentives, demonstration effects, and higher returns to clean research can alter the direction of technological development (Johnstone et al., 2010; Calel and Dechezleprêtre, 2016; Brown et al., 2022). Local platforms that improve access to data, computing resources, and specialized services can therefore expand firms' capacity to develop and absorb green technologies.

An expansion of green innovation can lower urban emissions through several forms of technical change. Process innovations reduce the energy and material inputs required for production, while

ACCEPTED MANUSCRIPT

cleaner equipment and products support substitution away from carbon-intensive technologies (Aghion et al., 2016; Dechezleprêtre et al., 2011; Noailly and Smeets, 2015). Market demand and competitive pressure can reinforce this shift by increasing the returns to cleaner technological choices (Aghion et al., 2023). As the stock of applicable green technologies increases, firms gain more options for meeting production targets with lower CO_2 emissions.

Advancing Industrial Structure

AI can reshape the industrial structure of a city by changing the relative productivity and growth prospects of different activities. AI adoption is associated with higher firm productivity, product innovation, and the expansion of technology-intensive businesses (Czarnitzki et al., 2023; Babina et al., 2024). City-level evidence from China also links AI development to industrial upgrading and a more rational allocation of production across sectors (Zou and Xiong, 2023). The expansion of advanced manufacturing and knowledge-intensive services raises the share of activities that generate more value from a given amount of energy and material input.

Intelligent technologies also support the upgrading of traditional industries. Automation improves quality control, machinery utilization, and the consistency of production, while complementary intangible investments allow firms to reorganize existing production processes (Graetz and Michaels, 2018; Brynjolfsson et al., 2021). These changes can reduce dependence on low-productivity activities that rely heavily on energy and raw materials. Historical evidence shows that changes in production techniques and industrial composition are central to explaining declines in manufacturing pollution (Levinson, 2009; Shapiro and Walker, 2018).

The NGAIPZ program can accelerate these structural changes by concentrating AI investment, skilled labor, and complementary services within selected cities. The policy encourages the integration of AI with existing industries and supports the formation of new technology-based activities (Ministry of Science and Technology of China, 2020; Cockburn et al., 2019). Capital and labor may consequently move toward sectors with stronger productivity growth and lower energy requirements per unit of value added. The environmental effect of this reallocation depends on the balance among changes in

ACCEPTED MANUSCRIPT

economic scale, industrial composition, and production techniques (Copeland and Taylor, 2004; Cherniwchan, 2017; Glaeser and Kahn, 2010).

Based on the above analysis, we propose the following hypotheses:

Hypothesis 1. The NGAIPZ program reduces urban carbon-emission intensity.

Hypothesis 2. The NGAIPZ program reduces urban carbon-emission intensity by improving energy efficiency.

Hypothesis 3. The NGAIPZ program reduces urban carbon-emission intensity by promoting green innovation.

Hypothesis 4. The NGAIPZ program reduces urban carbon-emission intensity by advancing the industrial structure.

Materials and Methods

Model Specification

The NGAIPZ program was introduced in successive batches across Chinese cities. We exploit the variation in policy exposure across cities and over time and estimate the following difference-in-differences model:

$$CEI_{it} = \alpha + \beta\, AITrial_{it} + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (1)$$

where *i* and *t* denote city and year, respectively. *CEI* is urban carbon-emission intensity, and *AITrial* identifies whether a city has been approved for the NGAIPZ program. X_{it} denotes a vector of time-varying city characteristics. The model includes city fixed effects μ_i and year fixed effects λ_t.

The coefficient of interest is β. A negative estimate indicates that the NGAIPZ program reduces carbon emissions per unit of economic output. City fixed effects absorb persistent differences in resource endowments, industrial foundations, geographic conditions, and other local characteristics. Year fixed effects account for nationwide changes that affect all cities in a given year. Standard errors are clustered at the city level.

ACCEPTED MANUSCRIPT

Variable Definitions

Carbon-Emission Intensity

The dependent variable is urban carbon-emission intensity (*CEI*), calculated as city-level carbon dioxide emissions divided by real gross domestic product (GDP), expressed in 2000 prices. Total carbon emissions are strongly affected by the scale of local economic activity. Carbon-emission intensity instead measures the emissions associated with producing one unit of output and therefore reflects the carbon efficiency of urban economic development. A lower value indicates that economic output is generated with fewer carbon emissions. Nominal city GDP is deflated using a GDP deflator with 2000 as the base year. The numerator is measured in tonnes of CO2, so the ratio expresses emissions per unit of real output. As a ratio, CEI can decline through lower emissions, higher real GDP, or both.

NGAIPZ Program

The core explanatory variable is *AITrial*. Policy approval dates are collected from official announcements issued by the Ministry of Science and Technology. For a city approved for the NGAIPZ program, *AITrial* equals one in the year of approval and all subsequent years. It equals zero before approval. Cities that are not approved for the program retain a value of zero throughout the sample period. The estimation sample comprises 17 approved cities, including six approved in 2019, six in 2020, and five in 2021; the remaining 271 cities were not approved during the sample period (Table A2).

Control Variables

The baseline specification controls for several city characteristics associated with carbon-emission intensity. Economic development (*PGDP*) is measured by the natural logarithm of real per capita GDP, while population density (*DEN*) is the number of permanent residents per square kilometer. Fiscal expenditure intensity (*FIS*) is calculated as general public budget expenditure divided by GDP. Financial development (*FIN*) is measured by the outstanding loan balance of financial institutions relative to GDP. Internet development (*NET*) is measured as the ratio of internet users to the

ACCEPTED MANUSCRIPT

permanent-resident population, calculated by dividing the number of internet users per 10,000 residents by 10,000. Higher values indicate greater internet penetration. These controls capture differences in local development, urban concentration, fiscal activity, financial depth, and digital infrastructure. The population denominators used for population density and per capita emissions follow the same permanent-resident population definition.

Sample and Data

The sample consists of a balanced panel of 288 Chinese prefecture-level cities from 2014 to 2023. City-level carbon dioxide emissions are constructed from the EDGAR_2026_GHG inventory. Annual CO2 emissions are provided on a 0.1° × 0.1° grid in tonnes per grid cell. The inventory covers major sources from electricity generation, industrial combustion and processes, transport, buildings, fuel production, agriculture, and waste management. Grid-level emissions are spatially aggregated to prefecture-level administrative areas using GIS overlay. The city totals measure territorial, production-based emissions. EDGAR allocates sectoral emissions to grids using spatial proxies, so the city estimates may contain allocation errors correlated with local economic activity. Its consistent coverage across cities and years permits analysis over the full policy period (European Commission, 2026). The 2023 observations are taken from this published inventory release. For combustion sources, the release uses the detailed IEA series through 2023. Policy information is collected from official documents issued by the Ministry of Science and Technology. The remaining variables are drawn mainly from the China City Statistical Yearbook, statistical bulletins published by prefecture-level governments, and the CNRDS database.

Cities without complete histories for the core variables are excluded before the estimation panel is formed. Within the retained sample, linear interpolation fills 76 missing control-variable observations, accounting for approximately 2.6% of the 2,880 city-year observations. The final balanced panel contains 288 cities and 2,880 city-year observations.

We cross-check EDGAR against the CEADs city-level emissions inventory over their overlapping observations (CEADs, 2022). The Pearson correlation between the natural logarithms of total CO2

ACCEPTED MANUSCRIPT

emissions is 0.8423 (Figure A1), indicating substantial agreement in relative emission levels across matched city-year observations. The CEADs inventory ends in 2019, when NGAIPZ approval begins, and therefore provides insufficient post-approval coverage for a comparable policy-effect estimate. Differences in absolute emission levels and spatial allocation remain relevant to interpretation.

Table 1. Descriptive Statistics

Variable	N	Mean	SD	Min	Max
CEI	2,880	1.618	1.201	0.302	4.613
AITrial	2,880	0.024	0.153	0.000	1.000
PGDP	2,880	10.908	0.491	10.101	11.822
DEN	2,880	419.371	283.116	40.000	995.542
FIS	2,880	0.208	0.089	0.102	0.431
FIN	2,880	1.171	0.538	0.542	2.516
NET	2,880	0.3026	0.1213	0.1080	0.5389

Results

Baseline Results

Table 2 reports the baseline estimates. All specifications include city and year fixed effects, while the control variables are introduced sequentially across columns. The coefficient on *AITrial* remains negative and statistically significant throughout. It ranges from −0.1321 in Column (1) to −0.0988 in the fully specified model, and its magnitude changes only moderately as additional city characteristics are included. This stability indicates that the estimated reduction in carbon-emission intensity is not driven by a particular set of observed covariates.

Column (5) presents the preferred specification. The coefficient on *AITrial* is −0.0988 and is significant at the 5% level. This corresponds to a decline of 0.0988 units of carbon-emission intensity, or approximately 6.1% of its sample mean of 1.618, providing empirical support for Hypothesis 1.

With 17 treated cities, we supplement conventional city-clustered inference with null-imposed wild-cluster bootstrap-t tests. Using 9,999 replications, the two-sided p-values for the preferred estimate are 0.0218 with Webb weights and 0.0226 with Rademacher weights (Table A3). The estimated

ACCEPTED MANUSCRIPT

decline in carbon-emission intensity remains significant at the 5% level under both weight distributions.

Table 2. Baseline Difference-in-Differences Estimates

Variables	(1)	(2)	(3)	(4)	(5)
AITrial	-0.1321***	-0.1301***	-0.0935**	-0.0988**	-0.0988**
	(0.0371)	(0.0365)	(0.0395)	(0.0396)	(0.0396)
PGDP	-1.1624***	-1.1614***	-1.0060***	-0.9848***	-0.9847***
	(0.0997)	(0.0998)	(0.1146)	(0.1157)	(0.1160)
DEN		0.0002	0.0002	0.0002	0.0002
		(0.0003)	(0.0003)	(0.0003)	(0.0003)
FIN			1.2482***	1.1781***	1.1778***
			(0.4283)	(0.4354)	(0.4354)
FIS				0.0577	0.0577
				(0.0566)	(0.0565)
NET					-0.0010
					(0.1193)
City FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	2,880	2,880	2,880	2,880	2,880
Adj. R-squared	0.9654	0.9654	0.9659	0.9660	0.9659

Notes: Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Dynamic Effects

We next estimate an event-study specification using the year before NGAIPZ approval as the reference period. Figure 1 plots the coefficients from six years before approval to four years after approval. The estimates for the pre-treatment years are small and statistically indistinguishable from zero. A joint test of the five pre-treatment coefficients yields $F(5, 287) = 0.5152$ ($p = 0.7647$), indicating no statistically significant pre-treatment differences within the event-study window.

The coefficient turns negative in the approval year and remains below zero throughout the observed post-treatment period. The decline is already visible at the time of approval and becomes larger at longer horizons. This timing is consistent with a policy whose effect operates through the gradual

ACCEPTED MANUSCRIPT

diffusion of AI technologies and their integration into urban production and management. The dynamic estimates therefore indicate that the reduction in carbon intensity emerges following approval and persists throughout the observed post-treatment period.

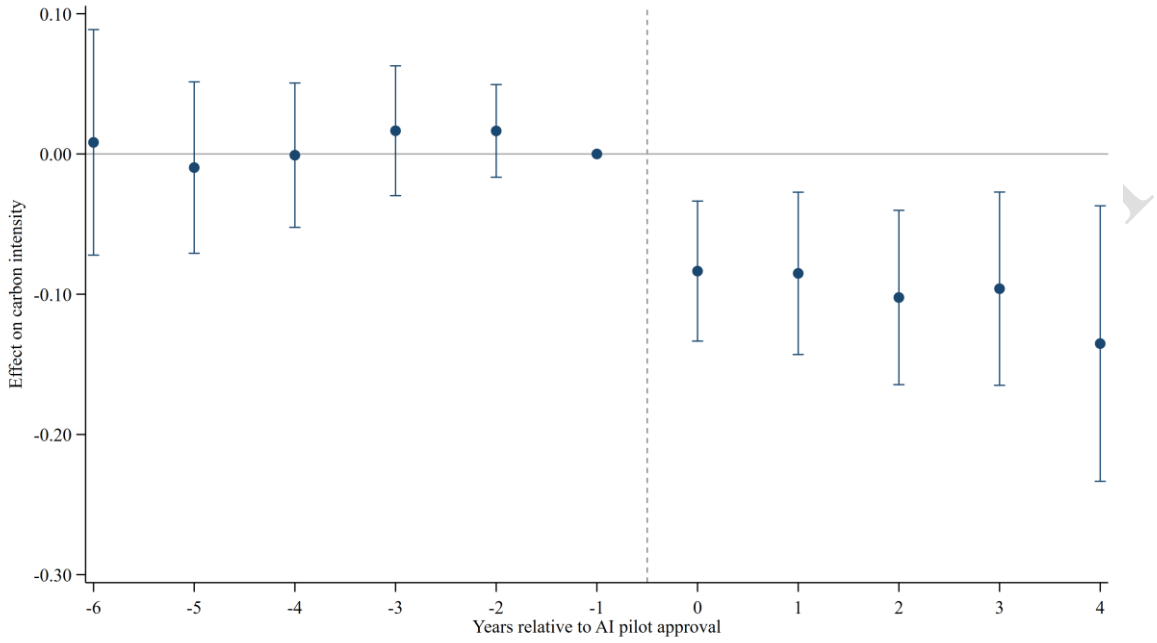


Figure 1. *Dynamic Effects of NGAIPZ Approval*

Notes: The figure reports event-study estimates relative to the year immediately before NGAIPZ approval, which is omitted. The vertical bars show 95% confidence intervals.

Treatment-Timing-Robust Estimates

In a staggered design, the conventional two-way fixed-effects estimator may combine comparisons across treatment cohorts in a way that is sensitive to heterogeneous policy effects. We therefore apply the interaction-weighted estimator of Sun and Abraham (2021). Figure 2 presents the resulting event-time estimates alongside the conventional estimates.

The two series display a similar pattern. For the Sun–Abraham estimator, the pre-treatment joint test yields $F(5, 287) = 0.3424$ ($p = 0.8869$). After approval, the Sun-Abraham estimates are negative in every period and become more pronounced over time. Their confidence intervals overlap those of the conventional estimates across the event window. The Sun–Abraham post-treatment average, weighted by the observed treated city-years, is −0.1056 (SE = 0.0307; 95% CI [−0.1661, −0.0451]), comparable in direction and magnitude to the baseline coefficient of −0.0988 (Table A1).

ACCEPTED MA

The wild-cluster bootstrap p-values for this weighted average are 0.0031 and 0.0047 under Webb and Rademacher weights, respectively. Restricting the average to years 0–2, when all approval cohorts contribute, yields an ATT of −0.0977, with corresponding p-values of 0.0039 and 0.0056 (Table A3). These tests condition on the observed cohort composition and aggregation weights. The interaction-weighted standard errors in Table A1 also account for uncertainty in estimated cohort shares.

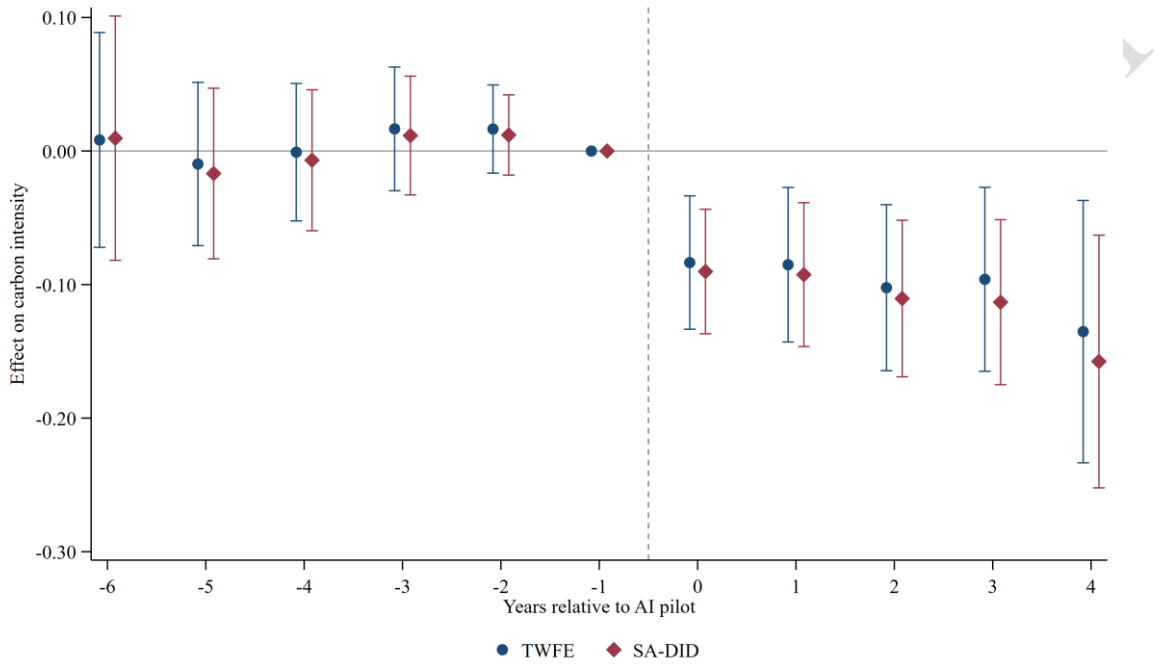


Figure 2. *Dynamic Effects Estimated by TWFE and Sun-Abraham DID*

Notes: The figure compares conventional two-way fixed-effects event-study estimates with interaction-weighted estimates following Sun and Abraham (2021). The year immediately before NGAIPZ approval is the omitted period. Vertical bars show 95% confidence intervals.

Placebo Test

We conduct 300 placebo replications by randomly reassigning pseudo-treated cities while preserving the original approval-cohort sizes: six cities in 2019, six in 2020, and five in 2021. The baseline model is re-estimated for each draw. This procedure yields a reference distribution of coefficients under randomized policy assignment.

Figure 3 shows a placebo-coefficient distribution with a mean of 0.0004 and a standard deviation of 0.0819. Of the 300 replications, 61 produce coefficients with an absolute value at least as large as the actual estimate of −0.0988. The empirical two-sided tail proportion is therefore 20.33%, and the

ACCEPTED

finite-simulation-adjusted p-value is (61 + 1)/(300 + 1) = 0.2060. This coefficient-based placebo test does not provide significance support for the baseline estimate. The placebo distribution reflects variation across reassigned cities and approval cohorts. The clustered standard error of 0.0396 and the wild-cluster bootstrap tests in Table A3 assess uncertainty conditional on the observed policy assignment.

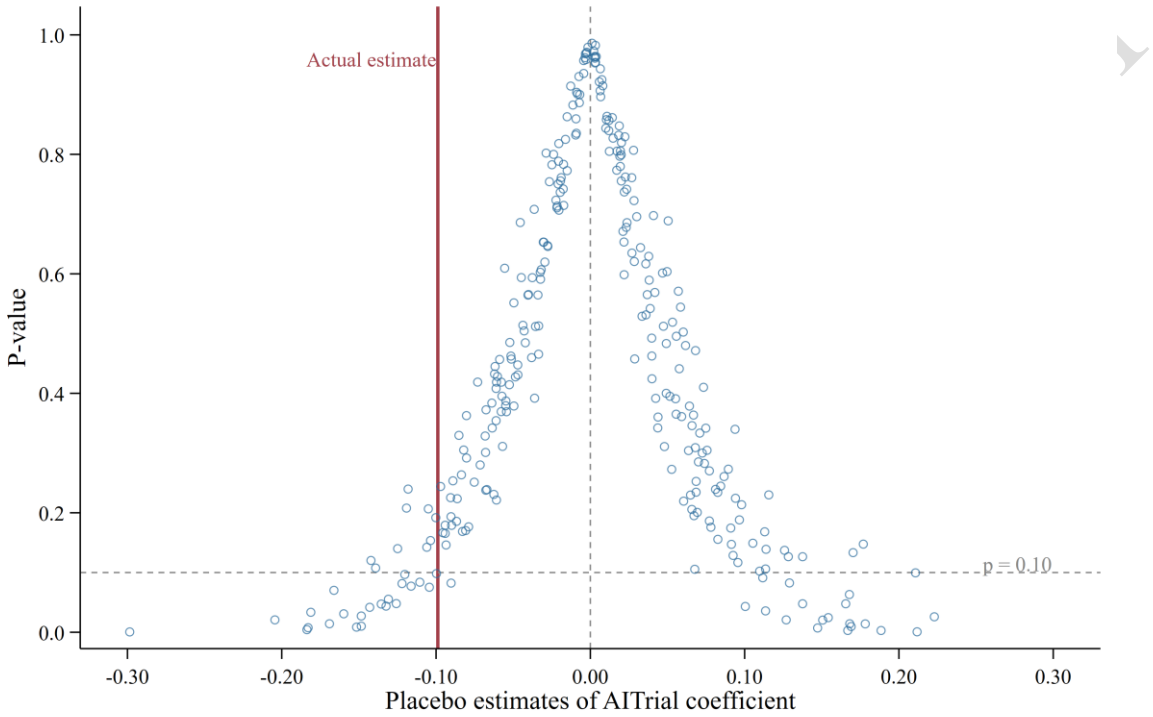


Figure 3. *Placebo Estimates*

Notes: The figure plots coefficients and conventional regression p-values from 300 random assignments of pseudo-treatment status and approval cohorts. The vertical line marks the actual baseline coefficient, and the horizontal line denotes a regression p-value of 0.10. The coefficient-based two-sided tail proportion is 61/300 = 20.33%; its finite-simulation-adjusted p-value is 0.2060.

Robustness Checks

Alternative Measures of Carbon Emissions

The baseline analysis uses carbon emissions per unit of GDP as the dependent variable. To assess whether the estimated effect is sensitive to this measurement, we employ four alternative carbon-emission outcomes. Column (2) uses the natural logarithm of carbon-emission intensity. Columns (3) and (4) consider the natural logarithms of total and per capita carbon emissions, respectively, while

ACCEPTED

Column (5) uses carbon emissions per capita without logarithmic transformation. Column (1) reproduces the baseline estimate for comparison.

The coefficient on *AITrial* is negative and statistically significant in every specification. The estimates are -0.1776 for *ln CEI*, -0.0373 for *ln CE*, -0.3976 for *ln PCE*, and -2.4861 for *PCE*. The results remain consistent when carbon emissions are measured relative to economic output, in aggregate terms, or on a per capita basis. The baseline finding is therefore not dependent on the construction of the carbon-emission outcome.

To examine changes in the GDP denominator, we also estimate models for ln GDP, ln total CO2 emissions, and ln CEI using the same sample and remaining controls, excluding ln real per capita GDP (Table A5). The ln GDP coefficient is 0.0520 (SE = 0.0332; p = 0.1179), while the coefficients for ln total emissions and ln CEI are −0.0454 (p = 0.0009) and −0.0975 (p = 0.0104), respectively. The negative emissions estimate provides evidence of a numerator contribution to the decline in carbon intensity. The GDP confidence interval, [−0.0133, 0.1174], also allows an economically relevant expansion in output, so the point estimates permit contributions from both lower emissions and higher GDP. These results distinguish the carbon-intensity effect from the separate estimate for absolute emissions.

ACCEPTED MANUSCRIPT

Table 3. Alternative Measures of Carbon Emissions

Variables	(1)	(2)	(3)	(4)	(5)
	CEI	***ln CEI***	***ln CE***	***ln PCE***	***PCE***
AITrial	-0.0988** (0.0396)	-0.1776*** (0.0375)	-0.0373*** (0.0136)	-0.3976*** (0.0659)	-2.4861*** (0.5355)
Controls	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	2,880	2,880	2,880	2,880	2,880
Adj. R-squared	0.9659	0.9780	0.9891	0.9683	0.9597

Notes: *CEI* denotes carbon-emission intensity, measured as total carbon dioxide emissions divided by GDP. *CE* denotes total carbon dioxide emissions, and *PCE* denotes carbon dioxide emissions per capita. Columns (2)-(4) use the natural logarithms of the corresponding measures. All regressions include the control variables and City FE and Year FE. Standard errors clustered at the city level are reported in parentheses. ** and *** denote significance at the 5% and 1% levels, respectively.

Controlling for Province-by-Year Fixed Effects

Unobserved shocks at the provincial level may affect both NGAIPZ approval and local carbon-emission intensity. To address this concern, Column (1) of Table 4 replaces year fixed effects with province-by-year fixed effects while retaining city fixed effects. This specification absorbs annual changes common to cities within the same province, including shifts in provincial industrial policy, environmental regulation, and energy conditions. The coefficient on *AITrial* is −0.1280 and remains significant at the 1% level. The estimated reduction in carbon-emission intensity therefore persists after accounting for province-specific time-varying shocks.

Excluding Special-Status Cities

Cities with special administrative status generally possess stronger fiscal capacity, greater access to policy resources, and more developed technological infrastructure. Their inclusion may affect the estimated policy effect. Column (2) excludes the four centrally administered municipalities and the five cities with independent planning status under the national plan. Specifically, the excluded cities are Beijing, Tianjin, Shanghai, Chongqing, Shenzhen, Xiamen, Ningbo, Qingdao, and Dalian. The

ACCEPTED MANUSCRIPT

restricted sample retains 279 cities, including 12 of the original 17 treated cities. The coefficient on *AITrial* is −0.1591 and is statistically significant at the 1% level. The negative estimate also appears among cities without these administrative designations, although the restriction changes the composition of treated cities.

Excluding the COVID-19 Period

The COVID-19 pandemic caused substantial disruptions to urban production, mobility, and energy consumption. These changes may introduce temporary fluctuations in carbon-emission intensity that are unrelated to the NGAIPZ program. We therefore remove observations from 2020 to 2022 and re-estimate the baseline model. As reported in Column (3), the coefficient on *AITrial* is −0.0964 and remains significant at the 5% level. This restriction leaves 23 of the original 69 treated city-year observations: six in 2019 and 17 in 2023. The 2020 and 2021 cohorts contribute treated observations only in 2023. The estimate is a supplementary comparison based on a narrower set of post-approval years and provides less evidence on persistence across the treatment period than the full-sample results.

Table 4. Alternative Fixed Effects and Sample Restrictions

Variable	(1) Province × Year FE	(2) Drop special cities	(3) Drop COVID years
AITrial	-0.1280*** (0.0381)	-0.1591*** (0.0361)	-0.0964** (0.0441)
Controls	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Year FE	Absorbed	Yes	Yes
Province × Year FE	Yes	No	No
N	2,830	2,790	2,016
Adj. R-squared	0.9695	0.9653	0.9671

Notes: All regressions include the indicated fixed effects and control variables. Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Column (2) retains 12 treated cities. Column (3) retains 23 treated city-year observations after excluding 2020–2022.

ACCEPTED MANUSCRIPT

Propensity Score Matching

Cities selected for the NGAIPZ program may have stronger economic foundations or better digital infrastructure than other cities. We use propensity score matching to improve comparability on observed pre-policy characteristics. Propensity scores are estimated from city-level averages of the baseline controls over 2014-2018. The DID model is then re-estimated on matched city panels.

Table 5 reports results from two radius specifications and a kernel-matching specification. The coefficient of *AITrial* ranges from -0.0717 to -0.1299 and remains statistically significant in every column. The estimates show that the reduction in carbon intensity persists when greater weight is placed on cities with similar pre-policy economic and digital conditions. After matching, the standardized mean differences for the matching covariates are all below 10%. Radius matching with calipers of 0.10 and 0.20 retains 13 treated and 270 control cities, and 16 treated and 271 control cities, respectively; kernel matching with bandwidth 0.10 retains 17 treated and 271 control cities.

Table 5. PSM-DID Estimates

Variable	(1) Radius 0.10	(2) Radius 0.20	(3) Kernel bw = 0.10
AITrial	-0.1299*** (0.0401)	-0.1162*** (0.0372)	-0.0717** (0.0346)
Controls	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	2,830	2,870	2,880
Adj. R-squared	0.9655	0.9659	0.9719

Notes: Propensity scores are estimated using the 2014-2018 city-level means of *PGDP*, *DEN*, *FIS*, *FIN*, and *NET*. The matched city sample is returned to the panel before estimating the DID model. Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Controlling for Concurrent Policies

The sample period overlaps with several national programs that may affect digital development, energy use, or low-carbon production. We add indicators for the national information pilot, the

ACCEPTED MANUSCRIPT

national big-data comprehensive pilot zone, new-energy demonstration cities, and the low-carbon city pilot. The policies are introduced separately in Columns (1)-(4) and jointly in Column (5) of Table 6.

The coefficient of *AITrial* remains negative and significant at the 5% level in all specifications. Its magnitude ranges from -0.0936 to -0.0993, close to the baseline estimate of -0.0988. The estimated carbon-intensity effect of the NGAIPZ program is therefore distinct from the influence of these overlapping policy initiatives.

Table 6. Controlling for Concurrent Policy Pilots

Variable	(1) Information	(2) Big data	(3) New energy	(4) Low carbon	(5) All policies
AITrial	-0.0988** (0.0396)	-0.0993** (0.0387)	-0.0988** (0.0396)	-0.0936** (0.0384)	-0.0944** (0.0377)
Policy control	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	2,880	2,880	2,880	2,880	2,880
Adj. R-squared	0.9659	0.9660	0.9659	0.9660	0.9660

Notes: Columns (1)-(4) control separately for the national information pilot, the national big-data comprehensive pilot zone, new-energy demonstration cities, and the low-carbon city pilot. Column (5) includes all four policy indicators. Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Mechanism Analysis

The baseline results show that the NGAIPZ program reduces urban carbon-emission intensity. We next examine production-side outcomes consistent with the channels developed above. Equation (1) is re-estimated with each channel-related outcome as the dependent variable.

ACCEPTED MANUSCRIPT

Improving Energy Efficiency

Artificial intelligence can improve the way cities monitor and allocate energy. More accurate forecasting helps firms and public utilities adjust production and energy demand in a timely manner. Automated control systems also reduce idle operation and identify abnormal consumption that is difficult to detect through conventional management. As AI applications spread across local industries, the same level of output can be produced with less energy input.

We use three measures to examine this channel. Electricity intensity (*EINT*) is the logarithm of electricity consumption per unit of real GDP, and energy intensity (*ENINT*) is the logarithm of total energy consumption per unit of real GDP. Lower values indicate lower energy use relative to economic output. Energy efficiency (*EE*) is the logarithm of super-efficiency SBM energy utilization efficiency, for which a higher value denotes more efficient energy use. The super-efficiency slack-based measure (SBM) incorporates labor, capital, and energy as inputs, real GDP as the desirable output, and CO2 emissions as the undesirable output.

Columns (1) to (3) of Table 7 report the results. The coefficient of *AITrial* is -0.4225 for electricity intensity and -0.3798 for energy intensity; both estimates are significant at the 1% level. The coefficient is 0.2621 when the SBM measure is used and is also significant at the 1% level. The NGAIPZ program thus lowers the energy required for urban production while raising measured energy utilization efficiency. The agreement across these indicators is consistent with the energy-efficiency channel proposed in Hypothesis 2.

Promoting Green Innovation

The development and diffusion of AI can change the local innovation process. AI tools expand the capacity to process technical information, shorten experimentation cycles, and improve the matching of research resources with industrial needs. They can also help firms identify energy-saving technologies and environmental solutions that are suitable for existing production conditions. A stronger AI ecosystem may therefore increase the supply of green technologies available to local producers.

ACCEPTED MANUSCRIPT

Green innovation is measured in two ways. *GI* is the logarithm of one plus the number of green patent applications, while *GIPC* is the number of green patent applications per 10,000 residents. The first measure captures the overall scale of green inventive activity. The per capita measure reduces the influence of differences in city size.

The coefficient of *AITrial* in Column (4) is 0.6909 and is significant at the 1% level. Column (5) reports a coefficient of 1.3710 for green patents per 10,000 residents, significant at the 5% level. Green patenting increases after a city is approved for the NGAIPZ program, and the result remains when innovation is expressed relative to population. These results are consistent with the green-innovation channel proposed in Hypothesis 3.

Advancing Industrial Structure

AI development may also alter the composition of urban economic activity. The adoption of intelligent technologies raises demand for knowledge-intensive services and facilitates the digital upgrading of traditional industries. Resources can gradually move away from production activities that rely heavily on energy and material inputs. Such structural change lowers the carbon intensity of local output even when the scale of economic activity continues to expand.

Industrial upgrading (ISU) is measured as $ISU_{it} = \Sigma(k \times s_{kit})$, for k = 1, 2, 3, where s_{kit} is the share of sector k in city i's GDP in year t. Higher values indicate a shift toward sectors assigned higher weights, particularly the tertiary sector. As shown in Column (6), the coefficient of *AITrial* is 0.1085 and is significant at the 1% level. The NGAIPZ program therefore leads to a discernible improvement in the local industrial structure. This finding is consistent with the industrial-upgrading channel proposed in Hypothesis 4.

ACCEPTED MANUSCRIPT

Table 7. Mechanism Analysis

Variables	Energy Efficiency			Green Innovation		Industrial Upgrading
	(1) *EINT*	(2) *ENINT*	(3) *EE*	(4) *GI*	(5) *GIPC*	(6) *ISU*
AITrial	-0.4225***	-0.3798***	0.2621***	0.6909***	1.3710**	0.1085***
	(0.0565)	(0.0515)	(0.0642)	(0.0669)	(0.6371)	(0.0306)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	2,820	2,820	2,820	2,880	2,880	2,880
Adj. R-squared	0.7474	0.7499	0.7974	0.9397	0.7632	0.9045

Notes: Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Heterogeneity Analysis

The baseline estimates describe the average effect of the AI pilot-zone program, but the scope for carbon reduction may depend on the local institutional and digital environment. We examine two dimensions that are closely connected to the implementation of artificial intelligence: environmental regulation and internet infrastructure. Cities are classified using the sample medians of their 2014–2018 pre-policy averages for the corresponding indicators, and Equation (1) is estimated separately for the two groups. Table 8 also reports formal tests of the differences between subgroup coefficients. Environmental regulation is proxied by the proportion of environmental-protection words in city government work reports, collected from the reports through web scraping; a higher proportion indicates greater emphasis on environmental protection. Group membership is fixed throughout 2014–2023. Each subgroup contains 144 cities and 1,440 observations, with no city-year observations lost after classification. The weak- and strong-regulation groups contain eight and nine treated cities, respectively; the low- and high-internet groups contain 12 and five. We therefore report wild-cluster bootstrap tests for the subgroup estimates and coefficient differences (Table A4).

ACCEPTED MANUSCRIPT

Environmental Regulation

Environmental regulation may alter the incremental contribution of AI development to low-carbon transition. Where regulation is already stringent, local governments and firms face stronger incentives to control emissions, and part of the available reduction potential may have been realized before the AI pilot is introduced. The additional effect of the policy can therefore be limited. In cities with weaker regulation, by contrast, AI-enabled monitoring, production optimization, and energy management may compensate for deficiencies in conventional environmental governance. The pilot can create a new route to emission reduction even when administrative pressure is relatively modest. Columns (1) and (2) of Table 8 report the results. The coefficient on *AITrial* is −0.1649 in the weak-regulation group and −0.0283 in the strong-regulation group. The corresponding Webb bootstrap p-values are 0.0046 and 0.3661. The strong-minus-weak coefficient difference is 0.1366 (SE = 0.0567), with Webb and Rademacher p-values of 0.0172 and 0.0181. The formal comparison indicates a larger decline in carbon-emission intensity in cities with weaker environmental regulation. This pattern is consistent with greater scope for improvements in production and energy management where initial regulatory pressure is limited.

Internet Infrastructure

The use of artificial intelligence relies on the collection, transmission, and processing of large volumes of data. A well-developed internet network lowers the cost of connecting firms, research organizations, and public agencies, allowing AI applications to move beyond isolated demonstration projects. When digital connectivity is poor, data cannot circulate efficiently and the adoption of intelligent technologies remains fragmented. The carbon-reduction effect of the pilot should consequently be stronger where internet infrastructure provides the necessary technological base.

The estimates in Columns (3) and (4) support this argument. In cities with low internet development, the coefficient on *AITrial* is −0.0127 and the Webb bootstrap p-value is 0.4207. The estimate is −0.1198 in the high-internet group, with Webb and Rademacher p-values of 0.0003 and 0.0005. The high-minus-low difference is −0.1071 (SE = 0.0328), with corresponding bootstrap p-values of

ACCEPTED MANUSCRIPT

0.0014 and 0.0017. The larger reduction in carbon-emission intensity in cities with stronger internet infrastructure is consistent with digital connectivity supporting the use of production data and the diffusion of AI applications.

Table 8. Heterogeneity Analysis

Variable	Environmental Regulation		Internet Infrastructure	
	(1) Weak	(2) Strong	(3) Low	(4) High
AITrial	-0.1649*** (0.0475)	-0.0283 (0.0309)	-0.0127 (0.0155)	-0.1198*** (0.0289)
Between-group difference	Strong - Weak: 0.1366**		High - Low: -0.1071***	
Controls	Yes	Yes	Yes	Yes
City FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	1,440	1,440	1,440	1,440
Treated cities	8	9	12	5
Adj. R-squared	0.9658	0.9670	0.9657	0.9661

Notes: Standard errors clustered at the city level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Groups are fixed using 2014–2018 city averages. Each group contains 144 cities observed for ten years. Differences are defined as strong minus weak regulation and high minus low internet development. Wild-cluster bootstrap p-values are reported in Table A4.

Conclusions

This study evaluates the environmental effects of China's NGAIPZ program using a staggered difference-in-differences design. From a production-side perspective, the analysis focuses on whether AI changes the energy requirements, technological content, and industrial composition of urban economic activity. The results show that the NGAIPZ program significantly lowers urban carbon-emission intensity, with a preferred estimate of −0.0988, equivalent to approximately 6.1% of the sample mean CEI. Wild-cluster bootstrap inference also supports this carbon-intensity effect. Separate estimates for total emissions are negative, including when real per capita GDP is removed from the controls. The coefficient-based placebo p-value of 0.2060 provides weaker evidence under randomized policy reassignment.

ACCEPTED MANUSCRIPT

The channel-related results identify three production-side outcomes that respond to the program. First, the NGAIPZ program lowers electricity and energy use relative to economic output and improves energy-utilization efficiency, consistent with more economical energy use in production. Second, the program increases green patenting, indicating an expansion of green inventive activity. Third, it advances the industrial structure as reflected in a higher industrial-upgrading index. Together, these changes are consistent with the proposed energy-efficiency, green-innovation, and industrial-upgrading channels.

The heterogeneity analysis further identifies the local conditions under which these production-side adjustments are more effective. The policy effect is larger in cities with weaker environmental regulation, where AI-enabled monitoring and production management have greater scope to improve existing practices. Cities with stronger internet infrastructure also experience larger reductions in carbon-emission intensity because digital connectivity facilitates the transmission of production data and the deployment of intelligent systems. The capacity of AI policy to lower urban carbon-emission intensity therefore depends on both the initial room for production improvement and the infrastructure needed to support technological diffusion.

Policy Implications

First, AI policy should be tied to measurable improvements in energy use and production processes. Pilot cities should prioritize intelligent energy-management systems, equipment monitoring, production scheduling, electricity-demand forecasting, and predictive maintenance in energy-intensive industries. Policy assessment should include changes in energy consumption, equipment utilization, and verified CO_2 reductions. Shared AI service platforms can lower adoption costs for small and medium-sized firms that lack independent access to computing resources, technical teams, and intelligent equipment.

Second, local governments should strengthen the link between AI applications and green technological development. Pilot cities can establish joint platforms connecting AI firms, manufacturers, universities, research institutes, energy suppliers, and environmental authorities.

ACCEPTED MANUSCRIPT

These platforms should provide industrial data, computing capacity, testing facilities, and real production scenarios for the development of energy-saving technologies. Fiscal support should place greater weight on technologies that achieve verified improvements in production efficiency and emissions, including intelligent energy control, low-carbon process optimization, and clean production equipment.

Third, implementation of the NGAIPZ program should follow differentiated local strategies. Cities with weaker environmental regulation should use intelligent monitoring, automated reporting, and anomaly detection to improve the coverage and precision of environmental governance, while also strengthening enforcement and disclosure. Cities with limited digital foundations should invest in broadband networks, industrial internet platforms, interoperable data standards, and computing infrastructure before expanding large-scale AI applications. Future pilot evaluation and financial support should therefore account for both governance capacity and digital readiness.

Limitations and Future Research

This study is subject to several limitations. The city-level data do not fully capture how individual firms adopt AI technologies and reorganize production, while the mechanism analysis cannot trace the complete sequence from policy support to AI application, production adjustment, and emission reduction. Because the outcomes are measured within city boundaries, the estimates do not capture cross-city adjustments through technology diffusion, industrial relocation, or energy and supply-chain linkages. Future research could combine firm-level production and energy-use data with information on AI investment and application scenarios, and extend the analysis to intercity and supply-chain spillovers to provide a more complete assessment of the environmental effects of AI development.

REFERENCES

[1] Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012): The environment and directed technical change. – American Economic Review 102(1): 131-166. https://doi.org/10.1257/aer.102.1.131.

ACCEPTED MANUSCRIPT

[2] Acemoglu, D., and Restrepo, P. (2020): Robots and jobs: Evidence from US labor markets. – Journal of Political Economy 128(6): 2188-2244. https://doi.org/10.1086/705716.

[3] Aghion, P., Bénabou, R., Martin, R., and Roulet, A. (2023): Environmental preferences and technological choices: Is market competition clean or dirty? American Economic Review: Insights, 5(1), 1-20. https://doi.org/10.1257/aeri.20210014.

[4] Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016): Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. – Journal of Political Economy 124(1): 1-51. https://doi.org/10.1086/684581.

[5] Allcott, H., and Greenstone, M. (2012): Is there an energy efficiency gap? Journal of Economic Perspectives, 26(1), 3-28. https://doi.org/10.1257/jep.26.1.3.

[6] Babina, T., Fedyk, A., He, A., and Hodson, J. (2024): Artificial intelligence, firm growth, and product innovation. – Journal of Financial Economics 151: 103745. https://doi.org/10.1016/j.jfineco.2023.103745.

[7] Brown, J. R., Martinsson, G., and Thomann, C. (2022): Can environmental policy encourage technical change? Emissions taxes and R&D investment in polluting firms. – Review of Financial Studies 35(10): 4518-4560. https://doi.org/10.1093/rfs/hhac003.

[8] Brynjolfsson, E., Li, D., and Raymond, L. (2025): Generative AI at work. – Quarterly Journal of Economics 140(2): 889-942. https://doi.org/10.1093/qje/qjae044.

[9] Brynjolfsson, E., Rock, D., and Syverson, C. (2021): The productivity J-curve: How intangibles complement general purpose technologies. – American Economic Journal: Macroeconomics 13(1): 333-372. https://doi.org/10.1257/mac.20180386.

[10] Calel, R., and Dechezleprêtre, A. (2016): Environmental policy and directed technological change: Evidence from the European carbon market. – Review of Economics and Statistics 98(1): 173-191. https://doi.org/10.1162/REST_a_00470.

[11] CEADs. (2022): The emission inventories for 290 Chinese cities from 1997 to 2019. https://www.ceads.net.cn/data/carbon-inventory/city-level/.

ACCEPTED MANUSCRIPT

[12]Cherniwchan, J. (2017): Trade liberalization and the environment: Evidence from NAFTA and U.S. – manufacturing. Journal of International Economics 105: 130-149. https://doi.org/10.1016/j.jinteco.2017.01.005.

[13]Cockburn, I. M., Henderson, R., and Stern, S. (2019): The impact of artificial intelligence on innovation: An exploratory analysis. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), The Economics of Artificial Intelligence: An Agenda (pp. 115-146). University of Chicago Press.

[14]Copeland, B. R., and Taylor, M. S. (2004): Trade, growth, and the environment. – Journal of Economic Literature 42(1): 7-71. https://doi.org/10.1257/002205104773558047.

[15]Czarnitzki, D., Fernández, G. P., and Rammer, C. (2023): Artificial intelligence and firm-level productivity. – Journal of Economic Behavior & Organization 211: 188-205. https://doi.org/10.1016/j.jebo.2023.05.008.

[16]Dechezleprêtre, A., Glachant, M., Haščič, I., Johnstone, N., and Ménière, Y. (2011): Invention and transfer of climate change-mitigation technologies: A global analysis. – Review of Environmental Economics and Policy 5(1): 109-130. https://doi.org/10.1093/reep/req023.

[17]Ding, T., Li, J., Shi, X., Li, X., and Chen, Y. (2023): Is artificial intelligence associated with carbon emissions reduction? Case of China. – Resources Policy 85: 103892. https://doi.org/10.1016/j.resourpol.2023.103892.

[18]Dong, M., Wang, G., and Han, X. (2024): Impacts of artificial intelligence on carbon emissions in China: In terms of artificial intelligence type and regional differences. – Sustainable Cities and Society 113: 105682. https://doi.org/10.1016/j.scs.2024.105682.

[19]European Commission, Joint Research Centre. (2026): EDGAR Community GHG database, version EDGAR_2026_GHG. https://edgar.jrc.ec.europa.eu/dataset_ghg2026.

[20]Glaeser, E. L., and Kahn, M. E. (2010): The greenness of cities: Carbon dioxide emissions and urban development. – Journal of Urban Economics 67(3): 404-418. https://doi.org/10.1016/j.jue.2009.11.006.

[21]Graetz, G., and Michaels, G. (2018): Robots at work. – Review of Economics and Statistics 100(5): 753-768. https://doi.org/10.1162/rest_a_00754.

ACCEPTED MANUSCRIPT

[22]Guo, Q., Peng, Y., and Luo, K. (2025): The impact of artificial intelligence on energy environmental performance: Empirical evidence from cities in China. – Energy Economics 141: 108136. https://doi.org/10.1016/j.eneco.2024.108136.

[23]IEA. (2024): CO_2 emissions in 2023. International Energy Agency. https://www.iea.org/reports/co2-emissions-in-2023.

[24]IPCC. (2023): Climate change 2023: Synthesis report. Intergovernmental Panel on Climate Change. https://doi.org/10.59327/IPCC/AR6-9789291691647.

[25]Jaffe, A. B., and Stavins, R. N. (1994): The energy-efficiency gap: What does it mean? Energy Policy, 22(10), 804-810. https://doi.org/10.1016/0301-4215(94)90138-4.

[26]Johnstone, N., Haščič, I., and Popp, D. (2010): Renewable energy policies and technological innovation: Evidence based on patent counts. – Environmental and Resource Economics 45: 133-155. https://doi.org/10.1007/s10640-009-9309-1.

[27]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 12: 518-527. https://doi.org/10.1038/s41558-022-01377-7.

[28]Lee, C.-C., and Yan, J. (2024): Will artificial intelligence make energy cleaner? Evidence of nonlinearity. – Applied Energy 363: 123081. https://doi.org/10.1016/j.apenergy.2024.123081.

[29]Lee, C.-C., Qin, S., and Li, Y. (2022): Does industrial robot application promote green technology innovation in the manufacturing industry? Technological Forecasting and Social Change, 183, 121893. https://doi.org/10.1016/j.techfore.2022.121893.

[30]Levinson, A. (2009): Technology, international trade, and pollution from U.S. – manufacturing. Review of Economic Dynamics 12(1): 1-22. https://doi.org/10.1016/j.red.2008.02.003.

[31]Li, J., Ma, S., Qu, Y., and Wang, J. (2023): The impact of artificial intelligence on firms' energy and resource efficiency: Empirical evidence from China. – Resources Policy 82: 103507. https://doi.org/10.1016/j.resourpol.2023.103507.

ACCEPTED MANUSCRIPT

[32]Li, X., Li, S., Cao, J., and Spulbar, A. C. (2025): Does artificial intelligence improve energy efficiency? Evidence from provincial data in China. – Energy Economics 142: 108149. https://doi.org/10.1016/j.eneco.2024.108149.

[33]Li, Y., Zhang, Y., Pan, A., Han, M., and Veglianti, E. (2022): Carbon emission reduction effects of industrial robot applications: Heterogeneity characteristics and influencing mechanisms. – Technology in Society 70: 102034. https://doi.org/10.1016/j.techsoc.2022.102034.

[34]Lyu, Y., Xiao, X., and Zhang, J. (2024): Does the digital economy enhance green total factor productivity in China? The evidence from a national big data comprehensive pilot zone. – Structural Change and Economic Dynamics 69: 183-196. https://doi.org/10.1016/j.strueco.2023.12.009.

[35]Martin, R., Muûls, M., de Preux, L. B., and Wagner, U. J. (2012): Anatomy of a paradox: Management practices, organizational structure and energy efficiency. – Journal of Environmental Economics and Management 63(2): 208-223. https://doi.org/10.1016/j.jeem.2011.08.003.

[36]Meng, C., Wang, L., and Lin, Y. (2024): Digital governance and carbon emission reduction: Evidence from the National Pilot Policy of Information Benefiting the People in China. – Journal of Environmental Management 368: 122179. https://doi.org/10.1016/j.jenvman.2024.122179.

[37]Mijit, R., Hu, Q., Xu, J., and Ma, G. (2025): Greening through AI? The impact of Artificial Intelligence Innovation and Development Pilot Zones on green innovation in China. – Energy Economics 146: 108507. https://doi.org/10.1016/j.eneco.2025.108507.

[38]Ministry of Science and Technology of China. (2020, September 29): Guidelines for the construction of national new generation artificial intelligence innovation and development pilot zones (revised edition). https://www.most.gov.cn/xxgk/xinxifenlei/fdzdgknr/fgzc/gfxwj/gfxwj2020/202012/t20201224_171987.html.

ACCEPTED MANUSCRIPT

[39]Noailly, J., and Smeets, R. (2015): Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. – Journal of Environmental Economics and Management 72: 15-37. https://doi.org/10.1016/j.jeem.2015.03.004.

[40]Popp, D. (2002): Induced innovation and energy prices. – American Economic Review 92(1): 160-180. https://doi.org/10.1257/000282802760015658.

[41]Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C., Ng, A. Y., Hassabis, D., Platt, J. C., Creutzig, F., Chayes, J., and Bengio, Y. (2022): Tackling climate change with machine learning. – ACM Computing Surveys 55(2): Article 42. https://doi.org/10.1145/3485128.

[42]Shapiro, J. S., and Walker, R. (2018): Why is pollution from U.S. – manufacturing declining? The roles of environmental regulation, productivity, and trade. American Economic Review 108(12): 3814-3854. https://doi.org/10.1257/aer.20151272.

[43]State Council of China. (2017, July 20): Notice on issuing the new generation artificial intelligence development plan. https://www.gov.cn/zhengce/content/2017-07/20/content_5211996.htm.

[44]Sun, L., and Abraham, S. (2021): Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. – Journal of Econometrics 225(2): 175-199. https://doi.org/10.1016/j.jeconom.2020.09.006.

[45]Tao, W., Weng, S., Chen, X., ALHussan, F. B., and Song, M. (2024): Artificial intelligence-driven transformations in low-carbon energy structure: Evidence from China. – Energy Economics 136: 107719. https://doi.org/10.1016/j.eneco.2024.107719.

[46]Tian, H., Zhao, L., Li, Y., and Wang, W. (2023): Can enterprise green technology innovation performance achieve corner overtaking by using artificial intelligence? Evidence from Chinese manufacturing enterprises. – Technological Forecasting and Social Change 194: 122732. https://doi.org/10.1016/j.techfore.2023.122732.

ACCEPTED MANUSCRIPT

[47]Verendel, V. (2023): Tracking artificial intelligence in climate inventions with patent data. – Nature Climate Change 13: 40-47. https://doi.org/10.1038/s41558-022-01536-w.

[48]Wang, Y., Zhao, W., and Ma, X. (2024a). The spatial spillover impact of artificial intelligence on energy efficiency: Empirical evidence from 278 Chinese cities. Energy, 312, 133497. https://doi.org/10.1016/j.energy.2024.133497.

[49]Wang, Z., Zhang, T., Ren, X., and Shi, Y. (2024b). AI adoption rate and corporate green innovation efficiency: Evidence from Chinese energy companies. Energy Economics, 132, 107499. https://doi.org/10.1016/j.eneco.2024.107499.

[50]Yang, S., and Liu, F. (2024): Impact of industrial intelligence on green total factor productivity: The indispensability of the environmental system. – Ecological Economics 216: 108021. https://doi.org/10.1016/j.ecolecon.2023.108021.

[51]Zhang, K., Kou, Z.-X., Zhu, P.-H., Qian, X.-Y., and Yang, Y.-Z. (2025): How does AI affect urban carbon emissions? Quasi-experimental evidence from China's AI innovation and development pilot zones. – Economic Analysis and Policy 85: 426-447. https://doi.org/10.1016/j.eap.2024.12.013.

[52]Zhao, C., Li, Y., Liu, Z., and Ma, X. (2024): Artificial intelligence and carbon emissions inequality: Evidence from industrial robot application. – Journal of Cleaner Production 438: 140817. https://doi.org/10.1016/j.jclepro.2024.140817.

[53]Zhong, W., Liu, Y., Dong, K., and Ni, G. (2024): Assessing the synergistic effects of artificial intelligence on pollutant and carbon emission mitigation in China. – Energy Economics 138: 107829. https://doi.org/10.1016/j.eneco.2024.107829.

[54]Zou, W., and Xiong, Y. (2023): Does artificial intelligence promote industrial upgrading? Evidence from China. – Economic Research-Ekonomska Istraživanja 36(1): 1666-1687. https://doi.org/10.1080/1331677X.2022.2092168.

ACCEPTED MANUSCRIPT

Appendix. Additional Estimation Details

Table A1. Comparison of Average Treatment-Effect Estimates

Estimator	ATT	SE	95% CI	p-value
Baseline TWFE	−0.0988	0.0396	[−0.1768, −0.0208]	0.0132
Sun–Abraham (weighted)	−0.1056	0.0307	[−0.1661, −0.0451]	0.0007
Sun–Abraham (years 0–2)	−0.0977	0.0304	[−0.1576, −0.0379]	0.0015

Notes: All estimates use 288 cities and 2,880 city-year observations, with city and year fixed effects and city-clustered standard errors. The principal Sun–Abraham average weights event-time estimates by their observed post-treatment city-year counts. The common-window estimate covers years 0–2 after approval, when all three cohorts contribute. The weighted summary covers 69 treated city-year observations, while the common-window summary covers 51.

Table A2. NGAIPZ Approval Cohorts and Sample Composition

Approval cohort	Cities	City-year observations	Observed post-approval years	Post-approval city-years
2019	6	60	0–4	30
2020	6	60	0–3	24
2021	5	50	0–2	15
Not approved	271	2,710	—	0
Total	288	2,880	—	69

Notes: The panel covers 2014–2023. The post-approval observation counts are the actual treated city-years available for the event study and the treatment-timing-robust average effect.

Table A3. Wild-Cluster Bootstrap Inference for Average Effects

Estimator	Estimate	Clustered SE	Webb p	Rademacher p
Baseline TWFE	−0.0988	0.0396	0.0218	0.0226
Sun–Abraham, years 0–4	−0.1056	0.0306	0.0031	0.0047
Sun–Abraham, years 0–2	−0.0977	0.0301	0.0039	0.0056

Notes: Two-sided, null-imposed wild-cluster bootstrap-t tests use 9,999 replications for each weight distribution, with clustering at the city level. All models use 2,880 observations from 288 cities, including 17 treated cities, and retain city and year fixed effects and the baseline controls. The Sun–Abraham tests condition on observed cohort composition and aggregation weights. Their clustered standard errors are 0.0306 and 0.0301 under fixed weights; the original interaction-weighted standard errors in Table A1 are 0.0307 and 0.0304. Years 0–4 and 0–2 cover 69 and 51 treated city-year observations, respectively.

ACCEPTED MANUSCRIPT

Table A4. Wild-Cluster Bootstrap Inference for Heterogeneity

Group or difference	Estimate	Clustered SE	Webb p	Rademacher p
Weak regulation	−0.1649	0.0475	0.0046	0.0058
Strong regulation	−0.0283	0.0309	0.3661	0.3694
Strong minus weak	0.1366	0.0567	0.0172	0.0181
Low internet	−0.0127	0.0155	0.4207	0.4244
High internet	−0.1198	0.0289	0.0003	0.0005
High minus low	−0.1071	0.0328	0.0014	0.0017

Notes: Groups are fixed using 2014–2018 city means and the cross-city median. Each subgroup contains 144 cities and 1,440 observations. The weak- and strong-regulation groups contain eight and nine treated cities; the low- and high-internet groups contain 12 and five. Difference tests use all 288 cities and 2,880 observations. All specifications retain controls and city and year fixed effects. Bootstrap tests are two-sided, impose the null, cluster at the city level, and use 9,999 replications for each weight distribution.

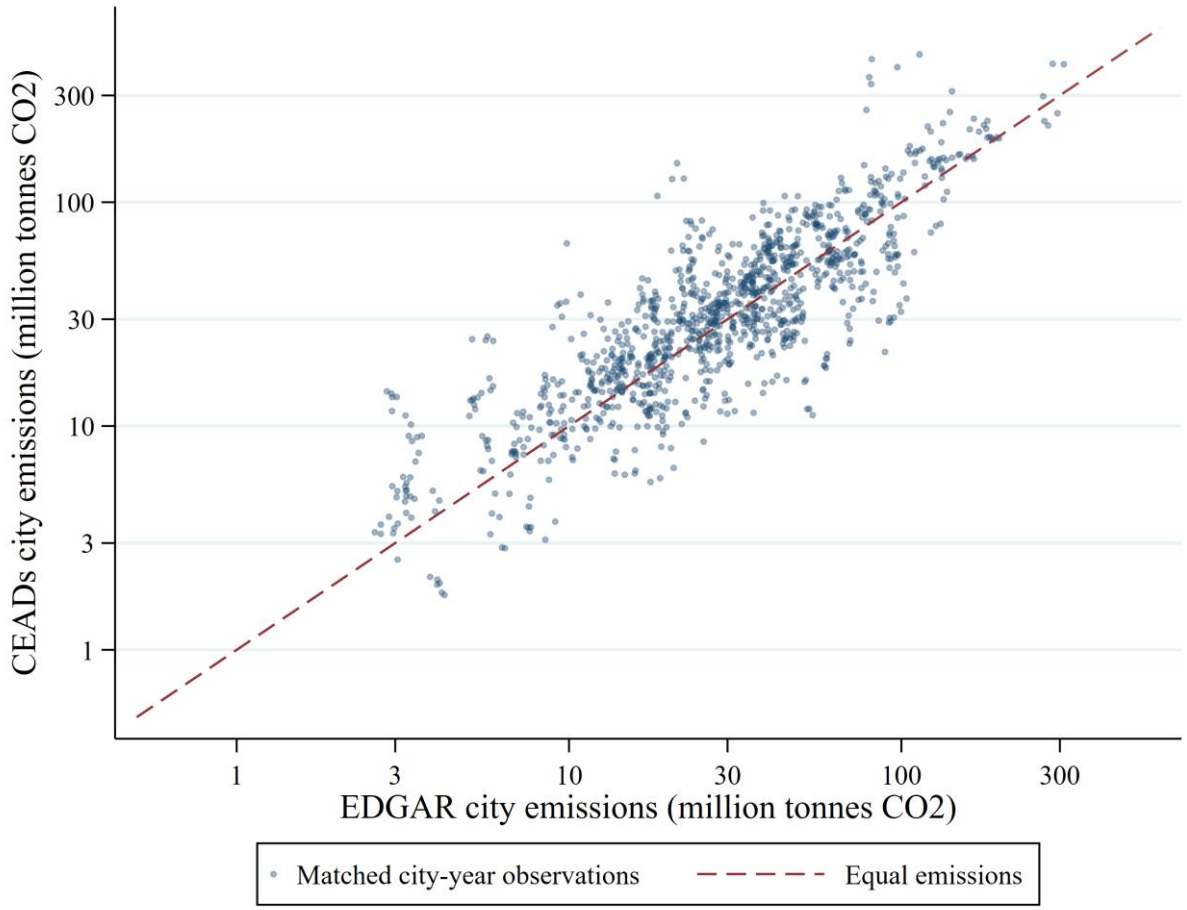
Table A5. GDP and Emissions Outcomes without the Per Capita GDP Control

	ln GDP	ln total CO2	ln CEI
AITrial	0.0520	−0.0454***	−0.0975**
Clustered SE	(0.0332)	(0.0135)	(0.0378)
Two-sided p-value	0.1179	0.0009	0.0104
95% CI	[−0.0133, 0.1174]	[−0.0721, −0.0188]	[−0.1718, −0.0231]
Remaining controls	Yes	Yes	Yes
City FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	2,880	2,880	2,880
Adjusted R-squared	0.9911	0.9891	0.9743
Cluster level	City	City	City

Notes: The three models use the same sample, controls, and fixed effects. The natural logarithm of real per capita GDP is excluded from the controls in every column. City-clustered standard errors appear in parentheses. ** and *** denote significance at the 5% and 1% levels.

ACCEPTED MANUSCRIPT

Figure A1. Comparison of CEADs and EDGAR City Emissions



Notes: Each point represents a matched city-year observation. Both axes use logarithmic scales, with emissions shown in million tonnes of CO2. The dashed line denotes equal emissions in the two inventories. The Pearson correlation between the natural logarithms of total emissions is 0.8423. The CEADs inventory ends in 2019 and does not cover the subsequent NGAIPZ treatment period.

ACCEPTED