Smarter Market Signals: Artificial Intelligence as a Policy Signal Accelerator for Carbon Market Decarbonization in China

Jianxian Wu[a, b] Hui Jin[c] Yaping Luo[d*]

Author Affiliations:

[a] *School of Government, Shenzhen University, Shenzhen, China*

[b] *Global Megacity Governance Institute, Shenzhen University, Shenzhen, China*

[c] *Faculty of Humanities and Social Sciences, Macao Polytechnic University, Macao*

[d] *School of Government, Beijing Normal University, Beijing, China*

Corresponding author: Yaping Luo

Email: luoyaping@mail.bnu.edu.cn

Abstract

Market-based climate policies such as carbon emission trading systems transmit regulatory price signals to guide firm-level abatement behavior, yet their effectiveness is frequently attenuated by signal lag, signal noise, and systematic mispricing arising from the structural complexity of climate policy information. Grounded in signaling theory and market microstructure scholarship, this study conceptualizes artificial intelligence (AI) not as an algorithmic trading tool but as a policy signal accelerator—an information-processing layer that compresses the time and informational loss between policy signal generation and price-level internalization. We develop a theoretical framework identifying two causal pathways through which AI improves carbon market price discovery—accelerating overnight signal instantiation into opening prices and filtering non-informational noise from intraday price movement. Exploiting the staggered rollout of China's carbon trading pilots across 288 cities, we employ a quasi-experimental design to estimate the causal effect of AI adoption on policy-driven emission reductions. Results show that AI amplifies the carbon trading emission-reduction effect by 2.1 percentage points. Mechanism analysis confirms that AI significantly raises the price efficiency ratio, indicating that carbon permit prices move more decisively in the direction of net informational content. These findings provide evidence that AI structurally enhances the signaling architecture of market-based climate policy.

Keywords: artificial intelligence; carbon emission trading; signaling theory; policy signal accelerator; price efficiency

Señales de mercado más inteligentes: la inteligencia artificial como acelerador de señales políticas para la descarbonización del mercado del carbono en China

Resumen

Las políticas climáticas basadas en el mercado, como los sistemas de comercio de emisiones de carbono, transmiten señales de precios regulatorios para guiar el comportamiento de reducción a nivel empresarial; sin embargo, su efectividad se ve frecuentemente atenuada por el retraso de la señal, el ruido de la señal y la fijación errónea sistemática de precios que surge de la complejidad estructural de la información sobre políticas climáticas. Basado en la teoría de la señalización y la investigación sobre la microestructura del mercado, este estudio conceptualiza la inteligencia artificial (IA) no como una herramienta de negociación algorítmica, sino como un acelerador de señales de política: una capa de procesamiento de información que comprime el tiempo y la pérdida de información entre la generación de la señal de política y la internalización del nivel de precios. Desarrollamos un marco teórico que identifica dos vías causales a través de las cuales la IA mejora el descubrimiento de precios en el mercado del carbono: acelerando la instanciación de la señal nocturna en precios de apertura y filtrando el ruido no informativo del movimiento de precios intradiario. Aprovechando el despliegue escalonado de los programas piloto de comercio de carbono de China en 288 ciudades, empleamos un diseño cuasiexperimental para estimar el efecto causal de la adopción de la IA en las reducciones de emisiones impulsadas por políticas. Los resultados muestran que la IA amplifica el efecto de reducción de emisiones del comercio de carbono en 2,1 puntos porcentuales. El análisis del mecanismo confirma que la IA aumenta significativamente el índice de eficiencia de precios, lo que indica que los precios de los permisos de carbono se mueven de manera más decisiva hacia el contenido informativo neto. Estos hallazgos demuestran que la IA mejora estructuralmente la arquitectura de señalización de las políticas climáticas basadas en el mercado.

Palabras clave: inteligencia artificial; comercio de emisiones de carbono; teoría de la señalización; acelerador de señales políticas; eficiencia de precios

更智慧的市场信号：人工智能作为中国碳市场脱碳的政策信号加速器

摘要

以碳排放权交易为代表的市场型气候政策，通过传递规制性价格信号引导企业层面的减排行为；然而，气候政策信息的结构复杂性所引致的信号滞后、信号噪声与系统性定价偏误，常使其政策效力受到削弱。本文立足信号传递理论与市场微观结构研究，将人工智能（AI）界定为“政策信号加速器”，而非单纯的算法交易工具——它作为一个信息处理层，能够压缩政策信号从生成到内化于价格之间的时间损耗与信息损耗。本文构建理论框架，识别出AI改善碳市场价格发现的两条因果路径：其一，加速隔夜信号向开盘价的转化；其二，过滤日内价格波动中的非信息性噪声。利用中国碳排放权交易试点在288个城市间推开所构成的准自然实验，本文估计了AI应用对政策驱动型减排的因果效应。结果表明，AI使碳交易的减排效应放大2.1个百分点。机制分析证实，AI显著提升了价格效率比率，表明碳配额价格能够更加果断地沿净信息含量方向变动。上述发现为AI从结构层面强化市场型气候政策的信号架构提供了经验证据。

关键词：人工智能；碳排放权交易；信号传递理论；政策信号加速器；价格效率

---

1. Introduction

Climate change ranks among the most urgent challenges confronting the global community, with far-reaching implications for ecosystems, economic development, and social stability (Jang and Yi, 2022; Xiao and Yi, 2025; Yi and Yuan, 2023). In response, governments worldwide have increasingly turned to market-based climate policies—including carbon emission trading systems, carbon taxes, carbon border adjustment mechanisms, renewable energy subsidies, and green financial instruments—to steer firms and consumers toward low-carbon behavior. Yet the effectiveness of these instruments hinges critically on the precision of their design, the responsiveness of market participants, and the efficiency with which they are implemented. In practice, conventional policy tools often fall short of their theoretical potential because of persistent information asymmetries, elevated transaction costs, and behavioral biases among regulated entities, all of which attenuate the price signals on which market-based mechanisms depend (Bellassen and Shishlov, 2017; Falkner et al., 2021; van der Heijden, 2020).

The rapid advancement of artificial intelligence (AI) has opened a promising frontier for overcoming these constraints (Kaack et al., 2022; Luers et al., 2024; Plésiat et al., 2024; Salekpay et al., 2024; Wu, 2025; Wu, 2026). Through big-data analytics, machine learning, and intelligent decision-support systems, AI can enhance the informational architecture underlying market-based climate policy in at least three ways: refining the design parameters of policy instruments, reshaping the decision-making processes of market participants, and improving the operational efficiency of policy implementation. In the specific context of carbon pricing, for example, AI-enabled forecasting and high-frequency information aggregation may allow carbon permit prices to incorporate new information more rapidly and accurately, thereby transmitting cleaner cost signals that strengthen firm-level abatement incentives. Despite growing enthusiasm for these possibilities, the existing literature lacks both a systematic theoretical framework and rigorous quasi-experimental evidence on whether—and through what channels—AI tangibly enhances the decarbonization potential of market-based climate policy.

This study aims to fill that gap. We first develop a theoretical framework articulating how AI can empower market-based climate policy through the lens of signaling theory. We then exploit China's staggered rollout of its carbon emission trading system pilot programs across 288 cities over nearly two decades to construct a quasi-experimental research design, enabling us to estimate the causal effect of AI adoption on policy-driven decarbonization.

Our findings yield three principal contributions. First, the baseline results demonstrate that AI adoption significantly amplifies the emission-reduction effect of the carbon emission trading systems. Relative to the control group, pilot cities in the treatment group experienced a statistically significant reduction in carbon emissions of approximately 2.1 percentage points —an effect that remains robust across a battery of specifications with progressively richer sets of city-level controls. Second, a mechanism analysis reveals that AI enhances the price efficiency of carbon trading markets: within pilot cities, greater AI penetration is associated with a significantly higher price efficiency ratio, indicating that carbon permit prices move more decisively in the direction of net informational content and are less distorted by noise trading. This improved price discovery transmits more credible abatement cost signals to regulated firms, thereby strengthening their incentive to invest in low-carbon technologies and production processes. Third, by situating these empirical results within an integrated theoretical framework, we provide one of the first pieces of quasi-experimental evidence that AI serves not merely as a technological convenience but as a structural enhancer of the informational and incentive architecture on which market-based climate policy relies.

Taken together, these results carry important implications for policymakers seeking to harness digital technologies in the pursuit of climate goals. As jurisdictions around the world expand or refine their carbon pricing instruments, understanding how AI can sharpen the market signals embedded in these mechanisms is essential for unlocking their full decarbonization potential.

2. Theoretical Framework

2.1 Theoretical Foundation: Signaling Theory in a Public Policy Context

Signaling theory, as originally formulated by Spence (1973) in the context of labor markets, addresses a fundamental problem of information asymmetry—when one party holds private information that another party cannot directly observe, observable and differentially costly actions can serve as credible signals capable of closing the informational gap. Since its introduction, the framework has been extended across organization theory, financial economics, and, more recently, public administration scholarship, where it is deployed to analyze how governments use policy instruments as signals to shape governance behavior (Zhu and Wang, 2025).

The application of signaling theory to emissions trading systems is institutionally well-grounded. A carbon market operates, at its core, as a mechanism through which government authorities transmit signals about the regulatory price of carbon over a forward-looking policy horizon (Wu et al., 2024). The cap-and-trade architecture—comprising total quantity limits, allocation rules, and compliance schedules—encodes a sequence of policy commitments that market participants must continuously interpret and price. Carbon prices, therefore, are not merely the outcome of supply-and-demand equilibria; they are the market's running valuation of policy credibility, regulatory intent, and abatement cost expectations (Hintermann, 2010). The accuracy and timeliness with which prices incorporate new policy information constitutes the primary measure of how effectively the signaling mechanism is functioning.

Yet a fundamental tension arises when signaling theory is transported into the carbon market context. Unlike private market signals—which are often standardized, numerically expressed, and rapidly disseminated—climate policy signals are structurally complex, textually dense, and institutionally embedded. Legislative amendments, regulatory guidance documents, international negotiation outcomes, and enforcement decisions are overwhelmingly non-structured in form, requiring interpretation against deep technical and jurisdictional background knowledge. Market participants face systematic cognitive constraints in processing this informational complexity, producing characteristic inefficiencies: delayed signal incorporation (signal lag), distortion through heterogeneous interpretation (signal noise), and systematic mispricing relative to policy fundamentals. It is against this background that the present study poses the core explanatory framework illustrated in Figure 1.

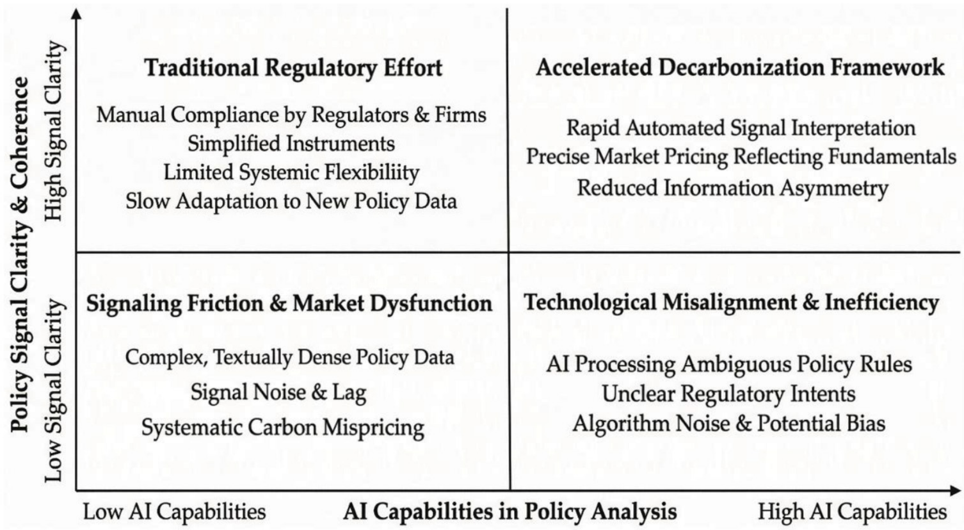

Figure 1. Theoretical framework

2.2 Core Conceptual Contribution: AI as a Policy Signal Accelerator

This study conceptualizes the function of artificial intelligence in carbon markets not as an algorithmic trading tool but as a policy signal accelerator—an information processing layer that compresses the time and informational loss between policy signal generation and complete price-level internalization.

This conceptualization draws on two complementary bodies of scholarship. From the public policy tradition, Majone (1989) argues that the effectiveness of any policy instrument depends critically on the legibility of the signals it encodes—the degree to which the target audience can decode the policy message with accuracy and low ambiguity. When legibility is low, even a well-designed instrument fails to produce the behavioral responses it intends. From the market microstructure tradition, Grossman and Stiglitz (1980) demonstrate that the allocative efficiency of any price-based mechanism depends on the capacity of informed traders to arbitrage informational gaps, a process that is itself bounded by the cost and difficulty of information acquisition and processing.

The theoretical bridge between these two traditions is provided by AI's differential advantage in processing non-structured, high-dimensional policy information (Schiff and Schiff, 2023; Wellstead et al., 2024). Natural language processing capabilities allow AI systems to parse regulatory announcements, legislative records, and multilateral agreement texts at a speed and scale that is categorically inaccessible to human analysts or conventional quantitative tools. Real-time sentiment and policy stance monitoring systems can continuously update market participants' policy priors without the latency inherent in human review cycles. And multi-source information fusion algorithms can synthesize signals across jurisdictional boundaries—a capacity of particular relevance in carbon markets that are increasingly integrated across regulatory regimes.

2.3 Causal Mechanisms: Two Pathways of Price Discovery Improvement

The theoretical framework generates two distinct but complementary causal pathways through which AI-augmented information processing improves carbon market price discovery. These pathways differ in their temporal locus—one operates primarily at the boundary between the overnight non-trading period and the market open, the other operates across the full intraday price formation process—and in the aspect of signal quality they most directly address.

2.3.1 Signal Instantiation and Intraday Price Correction

The first pathway concerns the speed and completeness with which policy signals are incorporated into opening prices. In carbon markets, the overnight non-trading period is a particularly information-rich interval: regulatory announcements, court rulings on compliance decisions, legislative committee proceedings, and international negotiating updates routinely accumulate during hours when no trading is possible. The degree to which these signals are reflected in the opening price, rather than generating protracted intraday drift as traders process and arbitrage the informational gap, is a direct measure of the market's signal instantiation capacity.

AI improves this capacity through two sub-mechanisms. First, by enabling near-real-time automated parsing of overnight policy documents, AI-assisted platforms allow market participants to arrive at the open with substantially more complete policy priors, reducing the structural informational asymmetry between informed and uninformed traders that drives mispricing at the open (Admati and Pfleiderer, 1988). Second, by generating continuous monitoring outputs that update directional policy stance estimates throughout the overnight period, AI narrows the confidence intervals around policy expectations (Silva, 2025), reducing the uncertainty premium that otherwise inflates intraday correction activity. The net result is a compression of the gap between opening and closing prices that reflects genuine informational updating rather than liquidity-driven noise—a compression that is theoretically measurable through the intraday price correction magnitude indicator.

2.3.2 Noise Filtering and the Price Efficiency Ratio

The second pathway operates through the composition of intraday price movement rather than its overall magnitude. Even in markets with robust overall liquidity, a significant share of intraday price volatility reflects non-informational forces: liquidity shocks, short-term speculative positioning, herding behavior among retail participants, and the mechanical responses of algorithmic strategies to price-level triggers rather than fundamental signals. When these noise components are large relative to information-driven price movement, the price path becomes directionality-poor—prices traverse a wide range intraday but close near their opening level, suggesting that most of the recorded movement was ultimately reversed and carried no policy-informational content.

AI addresses this problem by improving the signal-to-noise ratio in the information environment available to market participants—a concept directly continuous with Grossman and Stiglitz's (1980) foundational formulation. By enabling sophisticated participants to more precisely discriminate between policy-relevant signals and market microstructure noise, AI-assisted information processing raises the proportion of price movement that is directionally informed (Wang et al., 2024). Critically, this effect is amplified in carbon markets by the structural opacity of climate policy signals—the same regulatory development that is immediately legible to an AI-assisted analyst may remain ambiguous to a conventional market participant for several hours, creating temporary informational wedges that noise traders can exploit. As AI adoption increases market-wide, these wedges narrow, and the ratio of directional to total intraday price movement should rise correspondingly—a dynamic captured by the price efficiency ratio.

2.4 Research Hypotheses

Building on these theoretical foundations, this study advances two research hypotheses:

Hypothesis 1: AI adoption facilitates accelerated decarbonization under climate policy, particularly within carbon emission trading markets.

Hypothesis 2: AI accelerates policy-driven decarbonization through a signaling effect mechanism.

3. Methods

Section 3 outlines the research methods and data. Section 3.1 presents the modeling of the benchmark estimation model, including the key variables and data. Section 3.2 provides the modeling of the event study model and the descriptive statistics of the key variables.

3.1 Baseline Estimation Model

Drawing on annual panel data from 288 Chinese cities, this study constructs a triple-difference model to quantify the impact of AI on the decarbonization potential of carbon emission trading systems.

$$Y_{it} = \beta_1 AI_{it} \times CETS_{it} + \beta_2 AI_{it} + \beta_3 CETS_{it} + \gamma X_{it} + \delta_i + \eta_t + \epsilon_{it} \tag{1}$$

where the core explanatory variable is the interaction term $AI_{it} \times CETS_{it}$, which captures the policy effect of AI in pilot cities i under the carbon emission trading system in year t. Specifically, AI_{it} is a continuous variable measured by the logarithm of the number of AI-related invention patent applications. Patent data are sourced from the National Intellectual Property Administration. $CETS_{it}$ is a dummy variable indicating the implementation of the carbon emission trading system; it takes the value of 1 for cities designated as pilot regions for the market-based climate policy from 2013 onward, and 0 otherwise. Policy implementation data are derived from the National Energy Administration, with detailed information on the treatment group presented in Appendix Table A1[1].

The dependent variable Y_{it} denotes the logarithm of carbon emissions for city i in year t. Raw carbon emission data are obtained from the Carbon Emission Accounts and Datasets (CEADs) and matched to city-year observations spanning 2000 to 2016 (Chen et al., 2020). For robustness testing, we employ an alternative dataset from the Emissions Database for Global Atmospheric Research, which covers the period from 2000 to 2022, to re-examine our findings (Crippa et al., 2023). To mitigate omitted variable bias, the analysis controls for key city characteristics such as gross domestic product (GDP), population size, industrial output, and infrastructure (e.g., airports and high-speed rail).

The reasons for controlling these variables are as follows: GDP reflects economic development levels and resource accessibility, which are closely related to energy consumption and carbon emissions; population size represents the fundamental scale effect of energy demand; industrial output directly affects energy structure and emission intensity; airports and high-speed rail represent transportation infrastructure, which is related to urban energy consumption patterns, technological diffusion, and AI application environments. Previous research indicates that these variables are important determinants of carbon emissions, and neglecting them may lead to omitted variable bias (Hickel and Kallis, 2020; Lin et al., 2021; Pu and Xia, 2025; Sibanda and Ndlela, 2020; Whitmarsh et al., 2020), thereby affecting the accuracy of the causal relationship between AI and market-based climate policy.

To address the issue of unobserved factors and potential omitted variable bias, the study further controls for city fixed effects δ_i and year fixed effects η_t. In terms of estimation coefficients, β_1 captures the contribution of AI to the decarbonization potential of the carbon emission trading system. If the estimated coefficient β_1 is negative and statistically significant, it indicates that AI contributes to enhancing the decarbonization potential of the carbon emission trading system. The estimation uses ordinary least squares with robust standard errors to conduct valid statistical inference.

1 All supporting materials can be obtained through the following channels: https://doi.org/10.13140/RG.2.2.27808.34563

3.2 Event Study Model

Additionally, the validity of the baseline model estimation relies on the assumption that the carbon emission trends of the treatment and control groups are similar prior to the quasi-experimental intervention. If the carbon emission differences between the treatment and control groups are not statistically significant before the policy implementation, the baseline estimate can be interpreted as the causal effect of the intervention. This assumption is known as the parallel trends assumption (Liu et al., 2024). Given that the AI variable is originally continuous, and to facilitate the use of an event study framework, this study transforms it into a binary indicator. Specifically, cities with AI levels above the median are assigned a value of 1 (high AI group), and cities with AI levels below the median are assigned a value of 0 (low AI group). This transformation ensures that the interaction term in the model becomes binary, thus meeting the identification requirements of the event study design. The corresponding econometric specification is constructed as follows:

$$Y_{it} = \sum_{k\geq -5,\neq -1}^{4} \vartheta_k AI\_dummy_{ik} \times CETS_{ik} + \rho_1 AI\_dummy_{it} + \rho_2 CETS_{it} + \gamma X_{it} + \delta_i + \eta_t + \epsilon_{it} \quad (2)$$

where the interaction term captures the variation in carbon emission trends between the treatment and control groups before and after the quasi-experimental intervention. When k is less than 0 and the estimated coefficient is not statistically significant, it indicates that there are no systematic differences in carbon emissions between the treatment and control groups prior to the quasi-experimental intervention. When k is greater than or equal to 0 and the estimated coefficient is negative and statistically significant, the reduction in carbon emissions in the treatment group relative to the control group can be interpreted as the net effect of the quasi-experimental intervention. The other variables remain consistent with the baseline estimation model. Table A2 reports the descriptive statistics of the main variables.

4. Empirical Results

Section 4 presents the empirical results in detail. Section 4.1 reports the baseline estimation results, Section 4.2 provides the mechanism analysis results, and Section 4.3 presents the robustness check results.

4.1 Baseline Results

Table 1 reports the baseline impact of incorporating AI into the carbon emission trading system on carbon reduction. Column (1) controls for the main effects, city fixed effects, and year fixed effects, showing that compared to the control group, AI reduces carbon emissions in the treatment group by 2.53%, which is statistically significant at the 1% level. To further mitigate the impact of omitted variables, columns (2) to (6) progressively add control variables related to city characteristics. The baseline estimated coefficient of the carbon emission trading system's impact on carbon reduction gradually decreases to 2.1%, still statistically significant at the 1% level. These baseline estimation results provide preliminary support for *Hypothesis 1*, indicating that AI contributes to enhancing the decarbonization potential of the carbon emission trading system.

Regarding the control variables, GDP, industrial output, and airports all significantly increase carbon emissions at the statistical level, while high-speed rail significantly reduces carbon emissions. The relationship between population and carbon emissions is not statistically significant. The estimated coefficients for most control variables align with expectations. However, it should be noted that since the quasi-experimental intervention focuses on the core explanatory variable, and the control variables serve to enhance the accuracy of the quasi-experimental intervention estimates, the coefficients for these control variables do not require extensive interpretation.

Table 1. Baseline estimation of AI's role in promoting decarbonization in carbon emission trading system

	Dependent variable: Carbon emissions (logged)					
	(1)	(2)	(3)	(4)	(5)	(6)
AI×CETS	-0.0253***	-0.0271***	-0.0267***	-0.0214***	-0.0215***	-0.0210***
	(0.0038)	(0.0039)	(0.0039)	(0.0038)	(0.0037)	(0.0037)
GDP		0.2316***	0.2363***	0.1586***	0.1512***	0.1503***
		(0.0161)	(0.0166)	(0.0208)	(0.0203)	(0.0202)
Population			-0.0333	-0.0101	-0.0026	-0.0009
			(0.0282)	(0.0263)	(0.0258)	(0.0259)
Industry output				0.0663***	0.0685***	0.0680***
				(0.0077)	(0.0077)	(0.0077)
Airport					0.0537***	0.0517***
					(0.0102)	(0.0102)
High-speed rail						-0.0149***
						(0.0051)
Main effects	Yes	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4447	4445	4445	4438	4438	4438
Adjusted R^2	0.9856	0.9873	0.9873	0.9877	0.9878	0.9878

Note: This table reports the baseline estimation of AI's role in promoting decarbonization in carbon emission trading system. The dependent variable is carbon emissions (logged), and the core explanatory variable is AI×CETS, which represents the degree to which AI is integrated into the carbon emission trading system. The control strategy includes main effects, city control variables, city fixed effects, and year fixed effects. Numbers in parentheses represent robust standard errors. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

4.2 Mechanism Investigation

To illuminate the pathway through which AI adoption translates into reduced carbon emissions within pilot regions, this study posits that AI enhances the information efficiency of carbon trading markets, thereby strengthening the price signal that disciplines firm-level abatement behavior. We follow the structure of the baseline specification in Equation (1) and substitute the primary outcome variable with the price efficiency ratio (Barclay et al., 2008), defined as . This ratio, derived from daily carbon market transaction records sourced from the China Stock Market & Accounting Research (CSMAR) database and aggregated from the daily city level to the annual city level, captures the degree to which intraday price movement reflects net directional information rather than random noise. A value approaching unity indicates that prices move decisively from open to close relative to the full daily range—signaling that informed trading dominates and that market participants are efficiently incorporating new information into price discovery. Conversely, lower values reflect greater mid-session price reversals, consistent with elevated noise trading and diminished informational content.

The estimation sample is restricted to cities covered by the carbon emission trading system pilot programs, and the key explanatory variable remains the city-level AI patent number. Results are presented in Figure 2, which reports the estimated coefficient on AI under both the baseline specification without additional controls and the fully controlled specification. The estimated coefficient on AI is positive and statistically significant in both specifications, with point estimates of approximately 0.116 in each case and 95% confidence intervals lying largely above zero. This finding indicates that greater AI adoption is associated with a measurable improvement in carbon market price efficiency, specifically a higher price efficiency ratio, within carbon emission trading system pilot cities. These results lend empirical support to *Hypothesis 2*. AI technologies deployed by market participants, including machine learning-based forecasting models, high-frequency information aggregation, and automated trading systems, appear to reduce the share of noise-driven transactions and enhance the extent to which carbon permit prices reflect underlying supply and demand fundamentals. As market prices become more informationally efficient, they transmit cleaner and more credible cost signals to regulated entities, strengthening the incentive to invest in abatement technologies and low-carbon production processes.

4.3 Robustness Checks

4.3.1 Parallel Trends Test

Figure 3 reports the results of the parallel trends test based on the event study specification in Equation (2). Prior to the quasi-experimental intervention, the differences in carbon emissions between the treatment and control groups are not statistically significant. However, following the intervention, a noticeable decline in carbon emissions is ob-

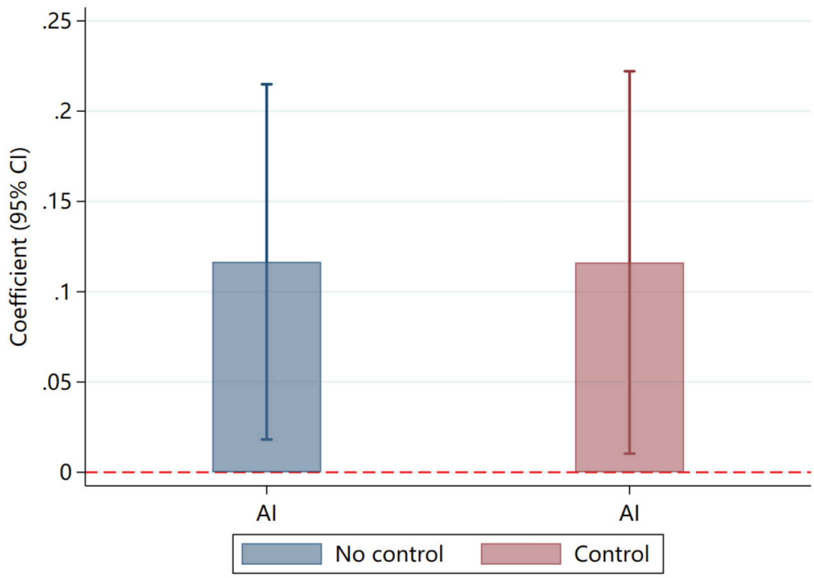

Figure 2. Mechanism analysis results

Note: The two sets of estimates differ in that the former excludes control variables and fixed effects, whereas the latter incorporates the full set of controls and fixed effects.

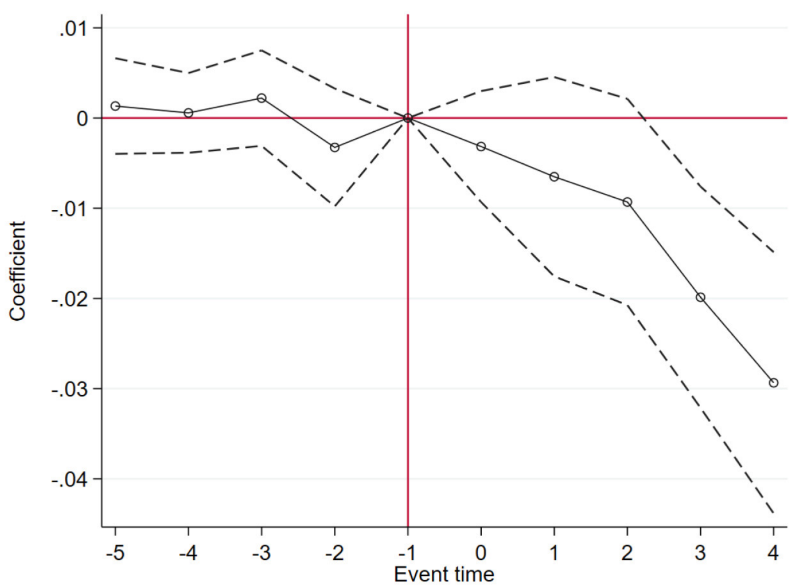

Figure 3. Dynamic effect estimation

Note: This figure reports the dynamic effect estimation of AI on decarbonization in carbon emission trading system. The horizontal axis represents the event time, relative to the year of the quasi-experimental shock. The vertical axis represents the estimated coefficients, showing the difference in carbon emissions (logged) between the treatment and control groups. The control strategy includes main effects, city control variables, city fixed effects, and year fixed effects, along with the 95% confidence intervals.

served in the treatment group relative to the control group. These post-intervention differences, compared with the pre-intervention differences, provide evidence consistent with the parallel trends assumption, suggesting that the observed reduction in carbon emissions in the treatment group is attributable to AI's integration into the carbon emission trading system rather than to pre-existing differences between the two groups.

4.3.2 Propensity Score Matching Estimation

Figure A1 reports the reduction in bias of city-level characteristic variables before and after 1:1 nearest-neighbor matching, showing that the average bias decreases to below 25%, indicating a satisfactory matching quality. After removing samples with substantial differences between the treatment and control groups, the baseline model is re-estimated. The specification remains consistent with the baseline estimation, incorporating main effects, city fixed effects, year fixed effects, and city-level characteristic variables. As shown in Table A3, the propensity score matching estimate is robust, with AI contributing to a 2.7% reduction in carbon emissions in the treatment group relative to the control group, statistically significant at the 1% level.

4.3.3 Placebo Test and Controlling for Other Climate Policies

Figure A2 presents the results of 100 placebo test re-estimations, indicating that the vast majority of placebo effects are smaller than the estimated coefficient for the treatment group. This provides preliminary evidence that omitted variables pose limited threats to the validity of the results. We further control for three climate policies that may confound the baseline estimation: the low-carbon city pilot program, pollution information disclosure, and central environmental inspections. These policies are coded as binary variables, assigned a value of one for the specific cities and years in which they were implemented and zero otherwise, and are included as control variables in the baseline model. As shown in Table A4, after controlling for the effects of these additional policies, AI contributes to a 2.09% reduction in carbon emissions in the treatment group relative to the control group, statistically significant at the 1% level.

4.3.4 Alternative Indicators and Extended Time Period Tests

We examine the sensitivity of the results using alternative indicators, including the logarithm of AI-related utility model patent applications, the number of authorized invention patents, and the number of authorized utility model patents. Tables A5 through A7 show that, even when employing different measurement indicators, AI contributes to reducing carbon emissions in the treatment group by 2.3%, 1.99%, and 2.42%, respectively, all statistically significant at the 1% level. Furthermore, the baseline model is re-estimated using the most recent carbon emission data from the Emissions Database for Global Atmospheric Research, extending the analysis through 2022 (covering the full period from 2000 to 2022). The results remain consistent with the main findings, as demonstrated in Figure A3, thereby providing additional support for the robustness of the study's conclusions.

4.3.5 Additional Robustness Checks

Entropy Balancing. A potential threat to the baseline estimates arises from pre-treatment differences between AI-adopting and non-adopting cities that may confound the estimated effect of $AI \times CETS$. To address this, we implement entropy balancing, a covariate-balancing preprocessing technique that reweights control-group observations so that the weighted moments of key covariates—GDP, population, industry output, airport infrastructure, and high-speed rail access—match those of the treatment group. Table A8 reports the covariate distributions before and after balancing. After reweighting, the means, variances, and skewness statistics of all five covariates are closely aligned across the treatment and control groups, confirming that systematic pre-treatment heterogeneity has been effectively eliminated. As shown in Column (1) of Table A9, re-estimating the baseline specification on the entropy-balanced sample yields a coefficient on $AI \times CETS$ of −0.0175 ($p < 0.01$), consistent in sign, magnitude, and statistical significance with the baseline result. This finding confirms that the main effect is not an artifact of pre-existing compositional differences between cities.

Instrumental Variable (IV) Estimation. Two interrelated concerns motivate an IV approach. First, using AI-related invention patent applications as a proxy for actual AI adoption introduces measurement error: patents capture innovative activity rather than deployment, and a substantial share of filed patents may never translate into operational applications in climate governance contexts. Second, a reverse causality concern exists whereby cities confronting more severe carbon emission pressure may proactively invest in AI technology, conflating the proposed causal pathway (AI adoption → emissions reduction) with its reverse. To jointly address these issues, we instrument for $AI \times CETS$ using the interaction between the number of post offices per million inhabitants in 1984 and the one-year lagged value of the core explanatory variable. Historical postal infrastructure serves as a plausible instrument because it predicts the early diffusion of information and communication technologies, and by extension contemporary AI capacity, while being otherwise unrelated to current carbon emission levels conditional on the included controls. The Cragg–Donald Wald F-statistic of 1,093.532 substantially exceeds conventional weak-instrument thresholds, validating instrument relevance. Column (2) of Table A9 reports the IV estimates; the coefficient on $AI \times CETS$ is −0.0141 ($p<0.1$), confirming that the negative effect of AI-enabled policy enforcement on carbon emissions is not driven by measurement error in the proxy variable or by reverse causality.

Additional Control Variables. Omitting determinants of carbon emissions that are correlated with both AI adoption and carbon market participation risks biasing the estimated interaction effect. Two such confounders are of particular concern: energy structure and environmental regulatory intensity. We therefore augment the baseline specification with province-year level coal consumption (sourced from the Chinese Research Data Services Platform and merged to the city-year panel) as a measure of energy mix, and with the annual count of environmental administrative penalty cases at the city level (sourced from the Peking University Law database) as a proxy for regulatory enforcement stringency. Column (3) of Table A9 reports the results. The coefficient on $AI \times CETS$ remains negative and statistically significant at −0.0126 ($p < 0.01$), while GDP, industry output, and airport infrastructure continue to exhibit positive and significant associations with carbon emissions, consistent with prior expectations. The stability of the core estimate after the inclusion of these additional controls suggests that omitted variable bias is unlikely to account for the main findings.

5. Discussion

5.1 AI as a Complement to, not a Substitute for, Policy Design

The findings of this study speak to a central question in the public policy literature: what determines whether a policy instrument achieves its intended behavioral effects? As Capano and Howlett (2020) argue, the success of any instrument depends not only on its technical architecture—the precision of its targets, the stringency of enforcement, the clarity of compliance schedules—but also on the design attributes that shape how target populations perceive and respond to the policy. Among these attributes, policy legibility (Chen and Greitens, 2022) and credibility (Fan et al., 2021) are especially consequential—a well-designed regulatory instrument that is opaque in communication or inconsistent in enforcement will fail to elicit the behavioral responses it intends.

Our results suggest that AI functions as an amplifier of these design attributes within the carbon emission trading system. By enabling near-real-time parsing of regulatory texts and policy announcements, AI-assisted information systems reduce the cognitive cost that market participants must incur to decode carbon market signals. This does not simplify the underlying regulatory content; rather, it lowers the barrier to understanding it. At the same time, the improvement in the price efficiency ratio documented in our mechanism analysis indicates that policy commitments are reflected in carbon prices more rapidly and completely. When market participants observe that prices respond coherently to new policy information, confidence in the regulatory regime's forward-looking commitments is reinforced, strengthening incentives for long-horizon abatement investment.

This interpretation carries an important implication: AI does not render good policy design irrelevant; it raises the returns to it. A poorly designed carbon market—one with opaque allocation rules, inconsistent enforcement, or politically driven cap adjustments—would generate noisy and contradictory signals that even sophisticated AI systems could not coherently synthesize. The emission reductions documented here are therefore best understood as a joint product of institutional design quality and technological information-processing capacity, not as evidence that technology can substitute for institutional quality.

5.2 Market Participation and Price Discovery

Our mechanism evidence also sheds light on how AI affects the quality of market participation. The improvement in the price efficiency ratio indicates that a larger share of intraday price movement in AI-augmented pilot cities reflects genuine informational updating rather than noise or liquidity-driven reversals. This pattern is consistent with a market in which informational barriers between informed and uninformed participants have been reduced, enabling a broader set of actors to engage with policy-relevant information when making trading decisions. The finding connects to Grossman and Stiglitz's (1980) foundational insight that the allocative efficiency of price-based mechanisms is bounded by the cost of information acquisition—by reducing this cost, AI tilts the composition of market activity toward fundamentals-driven participation. In practical terms, AI appears to shift the margin of market engagement from minimal compliance toward more informed trading behavior, thereby strengthening the price discovery function that gives emissions trading its dynamic efficiency advantage over command-and-control alternatives.

5.3 Contextualizing the Magnitude of Effects

The estimated reduction of approximately 2.1 percent in carbon emissions, while modest in absolute terms, should be interpreted against two benchmarks. First, it represents an incremental gain layered on top of an already functioning cap-and-trade system—a marginal efficiency improvement that compounds over time as AI adoption deepens. Second, the estimate captures only the informational channel of AI's impact; the full decarbonization potential of AI, encompassing direct applications in energy management, industrial process optimization, and grid coordination, is likely substantially larger. The modesty of the estimated effect is a feature of the identification strategy, which isolates one theoretically precise mechanism, rather than a limitation of the analysis.

6. Conclusion

6.1 Summary of Findings

This study investigates whether and how AI accelerates decarbonization within China's carbon emission trading system. Drawing on signaling theory transposed into a public policy context, we conceptualize AI as a policy signal accelerator—an information-processing layer that compresses the temporal and cognitive distance between the generation of climate policy signals and their internalization into carbon market prices. The empirical analysis exploits a quasi-experimental design in which the staggered rollout of carbon emission trading system pilots across Chinese cities provides plausibly exogenous variation in exposure to AI-augmented carbon trading infrastructure.

Three principal findings emerge. First, triple-difference estimates indicate that AI integration into the carbon emission trading system reduces carbon emissions in treatment-group cities by approximately 2.1 percent relative to control-group cities, a result that is statistically significant at the one percent level and robust across a comprehensive battery of sensitivity checks. Second, mechanism analysis reveals that AI adoption is associated with a statistically significant increase in the price efficiency ratio within carbon trading pilot cities, indicating that AI compresses the share of noise-driven transactions in intraday price movements. Third, these findings validate the study's two core hypotheses: AI contributes to accelerating decarbonization through the carbon emission trading system (Hypothesis 1), and it does so via a signaling-efficiency mechanism that improves market price discovery (Hypothesis 2).

6.2 Policy Implications

The findings yield several actionable implications. First, governments should invest in the digital information infrastructure surrounding carbon markets, not only in market architecture itself. The evidence that AI improves the price efficiency ratio—and that this improvement translates into emission reductions—implies that policies promoting AI-assisted market monitoring and real-time regulatory text analysis can complement conventional cap-and-trade design parameters. Public investment in open-access, AI-powered policy information platforms could help democratize informational advantages that currently accrue to technologically advanced participants.

Second, the finding that AI amplifies the returns to good policy design underscores the importance of regulatory clarity. In an environment where AI systems continuously parse and price regulatory signals, ambiguity or politically motivated deviations from announced policy trajectories will be identified and priced more rapidly than in a pre-AI regime. Well-communicated commitments will be rewarded with faster market adjustment and stronger abatement incentives; policy reversals will be penalized more swiftly.

Third, as carbon markets expand—including through the ongoing national rollout of China's emission trading system—policymakers should proactively address the distributional implications of AI-augmented information processing. This may involve establishing minimum digital infrastructure standards for market participation, providing technical assistance to smaller emitters, and designing surveillance mechanisms capable of detecting whether AI-driven informational advantages produce systematic patterns of market power.

6.3 Limitations and Future Research

This study is subject to several limitations. First, the measurement of AI adoption relies on patent-based indicators at the city level—a proxy that captures the intensity of technological activity but does not directly observe the deployment of specific AI applications within carbon trading operations. Future research employing firm-level data on AI tool adoption, or exploiting natural experiments associated with the rollout of specific AI-powered trading platforms, would permit more granular identification.

Second, the analysis is confined to the Chinese carbon emission trading system, which operates within a distinctive institutional context characterized by strong state capacity and centralized regulatory authority. While the theoretical framework is generalizable in principle, the empirical magnitudes may differ in systems with different institutional configurations, such as the EU Emissions Trading System or nascent carbon markets in developing economies. Cross-country comparative research is needed to assess external validity.

Third, the study documents an average treatment effect without disaggregating across firm types, industrial sectors, or regional development levels. As noted in the discussion, the distributional dimension of AI-augmented market efficiency is both theoretically important and policy-relevant, and heterogeneous treatment effect analysis should be a priority for future work.

Fourth, the temporal scope of the analysis does not capture the most recent advances in generative AI and large language models, which have dramatically expanded the capacity of AI systems to process complex, unstructured policy information. As these technologies diffuse into carbon market operations, their effects on price discovery may be qualitatively different from those estimated here. Longitudinal research tracking the co-evolution of AI capabilities and carbon market performance will be important.

Finally, this study focuses on the informational channel but does not examine potential feedback effects between AI-augmented price discovery and the political economy of cap-setting. If improved price signals make it easier for regulators to tighten caps by demonstrating that abatement is cost-effective, the long-run impact of AI integration may exceed what the current estimates suggest. Exploring these recursive dynamics between technology, market efficiency, and policy ambition represents a productive frontier for research at the intersection of public policy, environmental economics, and computational social science.

References

Admati, A.R., & Pfleiderer, P. 1988. A Theory of Intraday Patterns: Volume and Price Variability. *The Review of Financial Studies*. 1, 3-40.

Barclay, M.J., Hendershott, T., & Jones, C.M. 2008. Order Consolidation, Price Efficiency, and Extreme Liquidity Shocks. *Journal of Financial and Quantitative Analysis*. 43, 93-121.

Bellassen, V., & Shishlov, I. 2017. Pricing Monitoring Uncertainty in Climate Policy. *Environmental & Resource Economics*. 68, 949-974.

Capano, G., & Howlett, M. 2020. The Knowns and Unknowns of Policy Instrument Analysis: Policy Tools and the Current Research Agenda on Policy Mixes. *Sage Open*. 10.

Chen, H.R., & Greitens, S.C. 2022. Information capacity and social order: The local politics of information integration in China. *Governance: An International Journal of Policy Administration and Institutions*. 35, 497-523.

Chen, J., Gao, M., Cheng, S., Hou, W., Song, M., Liu, X., Liu, Y., & Shan, Y. 2020. County-level CO2 emissions and sequestration in China during 1997–2017. *Scientific Data*. 7, 391.

Crippa, M., Guizzardi, D., Pagani, F., Banja, M., Muntean, M., Schaaf, E., Becker, W., Monforti-Ferrario, F., Quadrelli, R., & Risquez, M. 2023. GHG emissions of all world countries, Luxembourg.

Falkner, K., Schmid, E., & Mitter, H. 2021. Integrated modelling of cost-effective policies to regulate Western Corn Rootworm under climate scenarios in Austria. *Ecological Economics*. 188.

Fan, S.Y., Zhang, T.Y., & Li, M.Y. 2021. The credibility and bargaining during the process of policy implementation-a case study of China's prohibition of open burning of crop straw policy. *Journal of Chinese Governance*. 6, 283-306.

Grossman, S.J., & Stiglitz, J.E. 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review*. 70, 393-408.

Hickel, J., & Kallis, G. 2020. Is Green Growth Possible? *New Political Economy*. 25, 469-486.

Hintermann, B. 2010. Allowance price drivers in the first phase of the EU ETS.

Journal of Environmental Economics and Management. 59, 43-56.

Jang, S.J., & Yi, H.T. 2022. Organized elite power and clean energy: A study of negative policy experimentations with renewable portfolio standards. *Review of Policy Research*. 39, 8-31.

Kaack, L.H., Donti, P.L., Strubell, E., Kamiya, G., Creutzig, F., & Rolnick, D. 2022. Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*. 12, 518-527.

Lin, Y., Qin, Y., Wu, J., & Xu, M. 2021. Impact of high-speed rail on road traffic and greenhouse gas emissions. *Nature Climate Change*. 11, 952-957.

Liu, L., Wang, Y., & Xu, Y. 2024. A Practical Guide to Counterfactual Estimators for Causal Inference with Time-Series Cross-Sectional Data. *American Journal of Political Science*. 68, 160-176.

Luers, A., Koomey, J., Masanet, E., Gaffney, O., Creutzig, F., Ferres, J.L., & Horvitz, E. 2024. Will AI accelerate or delay the race to net-zero emissions? *Nature*. 628, 718-720.

Majone, G. 1989. Evidence, Argument, and Persuasion in the Policy Process. Yale University Press.

Plésiat, É., Dunn, R.J.H., Donat, M.G., & Kadow, C. 2024. Artificial intelligence reveals past climate extremes by reconstructing historical records. *Nature Communications*. 15.

Pu, J.F., & Xia, F.Z. 2025. Exploring the relationship between city size and carbon emissions: An integrated population-land framework. *Applied Geography*. 177.

Salekpay, F., van den Bergh, J., & Savin, I. 2024. Comparing advice on climate policy between academic experts and ChatGPT. *Ecological Economics*. 226.

Schiff, D.S., & Schiff, K.J. 2023. Narratives and expert information in agenda- setting: Experimental evidence on state legislator engagement with artificial intelligence policy. *Policy Studies Journal*. 51, 817-842.

Sibanda, M., & Ndlela, H. 2020. The Link Between Carbon Emissions, Agricultural Output And Industrial Output: Evidence From South Africa. *Journal of Business Economics and Management*. 21, 301-316.

Silva, G.C. 2025. Articulating AI futures for Brazil: On different regimes of technological solutionism. *Critical Policy Studies*. 19, 637-656.

Spence, M. 1973. Job Market Signaling. *The Quarterly Journal of Economics*. 87, 355-374.

van der Heijden, J. 2020. Urban climate governance informed by behavioural insights: A commentary and research agenda. *Urban Studies*. 57, 1994-2007.

Wang, C.H., Zeng, Y.P., & Yuan, J.C. 2024. Two-stage stock portfolio optimization based on AI-powered price prediction and mean-CVaR models. *Expert Systems with Applications*. 255.

Wellstead, A.M., Mechling, S.M., Carter, A., & Gofen, A. 2024. Artificial intelligence possibilities to improve analytical policy capacity: The case of environmental policy innovation labs and sustainable development goals. *Policy Design and Practice*. 7, 456-467.

Whitmarsh, L., Capstick, S., Moore, I., Köhler, J., & Le Quéré, C. 2020. Use of aviation by climate change researchers: Structural influences, personal attitudes, and information provision. *Global Environmental Change-Human and Policy Dimensions*. 65.

Wu, J. 2025. Digital Jevons paradox in urban data center energy systems. *Nature Cities*. 2, 677-677.

Wu, J. 2026. Using nighttime light data to assess the impact of AI pilot cities on government overtime in China. *Urban Studies*. 1-27 (in press).

Wu, J., Nie, X., & Wang, H. 2024. Are policy mixes in energy regulation effective in curbing carbon emissions? Insights from China's energy regulation policies. *Journal of Policy Analysis and Management*. 43, 1152-1184.

Xiao, Y., & Yi, H. 2025. The Evolution of Chinese Policy Studies: A Bibliometric Analysis from 2000 to 2024. *China Policy Journal*. doi: 10.18278/cpj.3.1.5.

Yi, H.T., & Yuan, M. 2023. Policy advocacy, network formation and network governance. *Chinese Public Administration Review*. 14, 83-93.

Zhu, X., & Wang, Y. 2025. Credible signaling to promote local compliance: Evidence from China's multiwave inspection of environmental protection. *Public Administration*. 103, 47-65.

Author Contributions: Conceptualization, Methodology, Formal analysis, Writing—original draft, and Writing—review & editing: J.W. Conceptualization, Formal analysis, Writing—original draft, and Writing—review & editing: H.J. Conceptualization, Methodology, Writing—original draft, and Writing—review & editing: Y.L.

Competing Interest Statement: The authors declare no conflict of interest.

Acknowledgments: ***We extend our sincere appreciation to Prof. Ping Xu, Dr. Yan Yu, Dr. Yue Guo, and Dr. Lili Liu for their professional handling of our manuscript and for providing valuable constructive feedback.*** This research was supported by the National Social Science Fund of China (Grant No. 24ZDA051), the National Social Science Fund of China (Grant No. 22AZD141), and the China Postdoctoral Science Foundation (Grant No. 2025M780694).

Data Availability Statement: Data will be made available on request.

Bios:

Dr. Wu is a Research Fellow at the School of Government, Shenzhen University, and the Global Megacity Governance Institute. His research primarily focuses on public policy, government management, environmental governance, and digitalization. His scholarly works have been published in the *Nature Cities, Journal of Institutional Economics, Journal of Institutional and Theoretical Economics- Zeitschrift Fur Die Gesamte Staatswissenschaft (forthcoming), Journal of Policy Analysis and Management, Land Economics, Urban Studies, Review of Policy Research, Forest Policy and Economics, World Development, Social Indicators Research, Journal of Rural Studies, Telecommunications Policy, International Journal of Production Economics, Ecological Economics, Journal of Industrial Ecology, Environmental Science & Policy, Utilities Policy, Energy Policy, Technological Forecasting and Social Change, Journal of Environmental Management, Sustainable Development, International Journal of Sustainable Development & World Ecology, Marine Policy, and Habitat International.*

Dr. Jin is a PhD candidate in the Faculty of Humanities and Social Sciences at Macao Polytechnic University. Her primary research interests focus on human resource management and urban public affairs.

Dr. Luo is a PhD candidate at the School of Government, Beijing Normal University, with a primary focus on public policy and government management. Her works have been published in journals such as the *Social Science & Medicine, Journal of Asian Public Policy, Ecological Economics, Marine Policy, and Energy Policy.*