An Anatomy of China's Industrial Cleanup: Biased Technical Change and Emission Cost*

Ohyun Kwon† Hongsong Zhang‡ Hao Zhao§ Min Qiang Zhao¶

June 1, 2026

Abstract

This paper investigates the microeconomic drivers of China's industrial emission reduction. We develop a structural framework to jointly recover clean- and dirty-input-biased technical change and unobserved emission costs. Our estimation exploits firm-level production and emission data alongside profit-maximizing first-order conditions. We find that while rising emission costs primarily drove the decline in emission intensity between 2001 and 2008, both forms of biased technical change provided significant contributions. The induced innovation response to environmental regulation (the Porter Hypothesis) is robust and rapid: innovation offsets 40 percent of compliance costs within five years. Finally, we show that a uniform national emission cost yields substantial allocative efficiency gains but generates significant regional heterogeneity.

JEL classification: D24, O33, Q56

Keywords: Biased technical change, Production function estimation, Shadow cost of emissions, Emission intensity

Disclosure of research support and conflicts of interest: Hao Zhao acknowledges financial support from the National Natural Science Foundation of China (Grant No. 72303231). Min Qiang Zhao acknowledges financial support from the National Natural Science Foundation of China (Grant No. 71773102). The authors have no relevant or material financial interests, nor other relationships, that relate to the research described in this paper.

*We thank Nicolai Kuminoff, Devesh R. Raval, Konstantinos Serfes, Tianshi Mu, and Honghao Zheng for valuable discussion. We also thank participants at the International Industrial Organization Conference (IIOC), the Annual Conference of the Chinese Association of Environmental and Resource Economists (CAERE), the NCER-CCER China Economic Symposium, and the 2025 Workshop on Resource Security and Economic Sciences for helpful comments and suggestions. Any remaining errors are our own.

†LeBow College of Business, Drexel University, theekwonomist@gmail.com

‡Faculty of Business and Economics, The University of Hong Kong, hszhang@hku.hk

§School of Ecology and Environment, Renmin University of China, haozhao@ruc.edu.cn

¶MOE Key Laboratory of Econometrics, Wang Yanan Institute for Studies in Economics, Xiamen University, kent_zhao@xmu.edu.cn

1 Introduction

Over the past three decades, emission intensity, defined as pollution per unit of economic output, has declined in nearly all major economies.[1] A lower emission intensity allows the economy to grow with a less-than-proportional increase in emissions. Grossman and Krueger (1991) shows that this decline can be decomposed into *composition effects*, arising from shifts in the mix of industries or products, and *technique effects*, arising from cleaner production within a given activity.

Subsequent empirical studies consistently attribute most of this decline to the technique effect, which operates through four firm-level channels: (i) stringent environmental regulation; (ii) Hicks-neutral technical change; (iii) clean-input-biased technical change; and (iv) induced innovation, whereby (i) leads to (ii) and (iii). The last channel is widely known as the Porter Hypothesis, under which environmental policy can stimulate innovation that partly offsets compliance costs (Porter and van der Linde, 1995).[2]

Despite its intuitive appeal, quantifying the role of biased technical change in firm-level emission intensity decline remains difficult. The central empirical challenge is that both biased technology and the shadow cost of emissions (henceforth, emission cost) are unobserved, yet they jointly determine a firm's optimal mix of dirty and clean inputs. Existing approaches typically avoid this joint identification problem by assuming Hicks-neutral technology or by proxying the unobserved emission cost with observed factor prices (Copeland and Taylor, 2003; Doraszelski and Jaumandreu, 2018; Zhang, 2019; Hassler et al., 2021; Jo, 2024). We address this challenge by exploiting the production-function structure to jointly recover biased technical change and unobserved emission costs, allowing us to quantify the relative contributions of technology and regulation to industrial

[1] According to the Emissions Database for Global Atmospheric Research (EDGAR), global greenhouse-gas emission intensity fell by 41 percent between 1990 and 2023. For evidence across pollutants and contexts, see Grether et al. (2009); Levinson (2015); Brunel (2017).

[2] See Porter and van der Linde (1995); Acemoglu et al. (2012, 2016) for theoretical and Popp (2002); Calel and Dechezleprêtre (2016); Calel (2020) for empirical studies in this literature. A few important exceptions study directed technical change with respect to energy or fossil fuel inputs (Hassler et al., 2021; Jo, 2024). These studies typically infer technology from observed energy prices, abstracting from the shadow costs arising from emission tax or quantity-based regulatory constraints such as emission quotas.

cleanup.

This paper makes four contributions. First, we generalize the Copeland and Taylor (2003) (henceforth, CT) framework to incorporate biased technical change within a constant elasticity of substitution (CES) production structure. Our motivation stems from the observation that substantial heterogeneity in firm-level emission intensity persists even after controlling for industry, Hicks-neutral technology, and regulation. This residual heterogeneity is inconsistent with the Cobb-Douglas technology implied by the standard CT model. To reconcile this, we extend the linear abatement cost in the CT framework to a nonlinear form. This extension yields a micro-founded CES production function that accommodates two distinct biased technologies: *dirty-input-augmenting* (henceforth, dirty) and *clean-input-augmenting* (henceforth, clean) technical changes.[3] Moreover, our framework preserves the equivalence result in Copeland and Taylor (2003) that emissions can be viewed as either a by-product or a production input. We show that this generalized framework captures the firm-level environmental heterogeneity.

Our second contribution is to estimate this framework using firm-level data on sulfur dioxide (SO_2) emissions from China's manufacturing sector (2000–2008). We focus on SO_2 not only because it is a harmful industrial pollutant in its own right, but also because its strong correlation with CO_2 makes it highly relevant for climate policy. Furthermore, this period captures a critical transition: China's manufacturing sector experienced a rapid decline in SO_2 emission intensity driven by increasingly stringent environmental regulations.[4]

A central challenge in estimating a production function where emissions can be viewed as an input is that the emission cost is unobserved. Since this cost reflects regulatory stringency and effectively functions as a factor price, it is a key determinant of a firm's emission intensity. To address this, we demonstrate that the unobserved emission cost is

[3]While recent studies have employed CES production functions with emissions as an input (e.g., Duan et al. (2021); Kwon et al. (2021)), we are the first to provide a micro-foundation for this structure by extending the CT framework to accommodate two non-neutral technologies.

[4]As shown in Appendix Figure A.1, industrial SO_2 emission intensity fell by 60 percent between 2001 and 2008. This decline coincides with significant regulatory tightening, notably the elevation of SO_2 abatement to a binding target in the 11th Five-Year Plan (2006–2010).

structurally identified by variation in the labor-to-emissions ratio and the labor share of revenue. This approach allows us to jointly identify biased technology and emission costs, overcoming a persistent challenge in the industrial organization literature. Prior studies have typically relied on observed factor prices to infer biased technology (Doraszelski and Jaumandreu, 2018; Raval, 2019; Zhang, 2019) or recovered unobserved input prices by assuming a single Hicks-neutral technology (Grieco et al., 2016).

Third, we apply this framework to quantify the drivers of the marked decline in SO_2 emission intensity in China (2001–2008). We estimate an elasticity of substitution of 2.76, indicating that firms possess substantial flexibility to substitute away from dirty inputs. A structural decomposition reveals that the rising emission cost was the primary driver, accounting for 59.6 percent of the total reduction in emission intensity. Biased technical change also played a critical role, with clean and dirty technologies contributing 43.5 percent and 26.0 percent, respectively. We validate these structural changes by demonstrating that our estimated emission costs track local emission fees, while clean technology progress correlates with firm-level innovation. In contrast, rising wages offset the decline by 35.9 percent, since they induce substitution toward the relatively cheaper dirty input.

Fourth, we leverage our framework to quantify the impact of environmental regulations. First, we evaluate the magnitude and dynamics of the Porter Hypothesis. Our findings suggest a robust and rapid induced innovation response: firms offset approximately 40 percent of compliance costs through technology progress within five years, with the majority of the adjustment occurring in the first two years. The partial offset is consistent with prior estimates, whereas the speed of adjustment we document has received limited attention in the literature.

Finally, we explore the improvement of economic and environmental performances under counterfactual policy scenarios. We find that a nationwide doubling of emissions discharge fees lowers aggregate emission intensity by 22.7 percent in 2008. This decline is driven primarily by the direct cost channel rather than induced innovation. We also simulate uniform pricing schemes that equalize emission costs at either the national or

provincial level. Given the substantial heterogeneity in emission costs across prefectures, we estimate that national and provincial harmonization would increase aggregate value added by approximately 50 percent and 19 percent, respectively.

Related Literature Our analysis contributes to two strands of literature. First, we add to the environmental economics literature examining the determinants of industrial emissions. Following the seminal decomposition by Grossman and Krueger (1991), empirical research consistently identifies the technique effect as the primary driver of declining emission intensity.[5] We advance this literature by elucidating the micro-channels of the technique effect. While existing studies typically attribute industrial cleanup to either regulation (Shapiro and Walker, 2018) or Hicks-neutral technical improvement (Cherniwchan et al., 2017), we refine this dichotomy by decomposing the technique effect into clean- and dirty-input-biased components.

Second, we contribute to the literature on production function estimation with biased technical change. While standard methods focus on input endogeneity under Hicks-neutral shocks (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015), recent advances identify non-neutral change by leveraging observed factor prices or structural optimality conditions (Grieco et al., 2016; Doraszelski and Jaumandreu, 2018; Zhang, 2019; Hassler et al., 2021; Demirer, 2024). We address a distinct identification challenge: unobserved emission costs and biased technology jointly determine input choices. To resolve this, we develop a framework that simultaneously recovers both the biased technology parameters and the unobserved emission cost. This approach improves upon recent analyses that potentially conflate technology with policy costs (Lee et al., 2023; Guo et al., 2024) and eliminates the reliance on energy-price proxies to infer dirty technology (Hassler et al., 2021; Jo, 2024).

[5]The technique effect has been documented at the industry (Levinson, 2009), product (Shapiro and Walker, 2018), and establishment levels (Holladay and LaPlue III, 2021). Its primary drivers are typically identified as improved regulatory stringency (Shapiro and Walker, 2018; Najjar and Cherniwchan, 2021), Hicks-neutral productivity growth (Copeland and Taylor, 2003; Cherniwchan et al., 2017; Forslid et al., 2018), and pollution offshoring (Levinson, 2023; Choi et al., 2025). In the Chinese context, see He et al. (2020); Fan et al. (2025); Chen et al. (2025) for the effects of regulation; Wang et al. (2018) for the effects of productivity; and Kwon et al. (2023); Rodrigue et al. (2022) for trade and pollution.

The paper proceeds as follows. Section 2 documents the empirical puzzle motivating our framework. Section 3 develops the structural model and estimation strategy. Section 4 describes the data. Section 5 presents the structural estimates and analyzes the evolution of biased technologies and emission costs. Section 6 investigates the economic drivers of these components, focusing on environmental regulation and innovation. Section 7 presents counterfactual policy simulations, and Section 8 concludes.

2 Emission CES Production: Motivation and Micro-foundation

The Copeland and Taylor (2003) (CT) framework serves as the workhorse model for analyzing firm-level environmental performance. In this section, we document a key empirical puzzle that this standard model fails to rationalize, motivating its extension to a flexible CES framework that incorporates biased technical change.

A central theoretical appeal of the CT framework is its micro-foundation, which establishes an equivalence between modeling emissions as an input and as a by-product. CT assume that the production of a final good (Q) generates emissions (Z) as a by-product. Firms endogenously allocate a fraction, $\pi \in [0, 1]$, of the clean input (labor L) to abatement activities. The production technologies for output and emissions are given by:

$$\begin{aligned} &\text{(Production)} \quad Q &&= \omega L(1-\pi) \\ &\text{(Emissions)} \quad Z &&= \omega L(1-\pi)^{1/\alpha}, \end{aligned} \tag{1}$$

where ω represents Hicks-neutral technology, and $\alpha \in (0, 1)$, as it will turn out in the next formula, is the output elasticity of emissions, which governs emission intensity. The firm pays a regulatory cost proportional to emissions and therefore has an incentive to increase π. However, doing so entails a trade-off, as allocating labor to abatement reduces the production of final output. Eliminating π from the system yields the following key equation in the CT framework:

$$Q = Z^{\alpha}(\omega L)^{1-\alpha}. \tag{2}$$

Equations (1) and (2) demonstrate the equivalence of modeling emissions either as a production by-product or as a dirty input. This theoretical appeal, combined with the tractability of the Cobb-Douglas specification, has led to the widespread adoption of this framework in empirical studies (Forslid et al., 2018; Shapiro and Walker, 2018; Watabe et al., 2024).

Let W and τ denote the wage and the emission cost, respectively. This framework yields sharp predictions regarding key environmental moments in the data:[6]

$$\begin{aligned} &\text{(Emission-Input Ratio)} \quad \frac{Z}{L} &&= \frac{\alpha}{1-\alpha}\frac{W}{\tau} \\ &\text{(Emission Intensity)} \quad \frac{Z}{Q} &&= \omega^{\alpha-1}\left(\frac{\tau}{\alpha}\right)^{\alpha-1}\left(\frac{W}{1-\alpha}\right)^{1-\alpha}. \end{aligned} \tag{3}$$

Under the standard assumption that firms within the same location-industry-year cell face common factor prices (including emission cost),[7] the CT framework generates two sharp testable predictions: (1) the emission-to-labor ratio (Z/L) should be highly similar across firms within a cell; and (2) residual variation in emission intensity is determined solely by Hicks-neutral technology (ω).[8]

The empirical patterns in our data, however, reject these predictions. Following Lyubich et al. (2018), we quantify dispersion in firm performance as the 90th-10th percentile difference of log values within a location-industry-year cell and present the distribution across cells in Figure 1.[9] First, the emission-labor ratio exhibits substantial heterogeneity, contradicting the degenerate distribution implied by equation (3). Second, the dispersion in emission intensity markedly exceeds that of labor productivity or TFP—our empirical

[6] Equation (3) derives from the firm's cost minimization problem:

$$\begin{aligned} \min_{L,Z} \quad & WL + \tau Z \\ \text{s.t.} \quad & Q = Z^{\alpha}(\omega L)^{1-\alpha} \end{aligned}$$

[7] As emissions are equivalently viewed as an input in our framework, we use the term "input price" to include the emission cost.

[8] Specifically, estimates from Duan et al. (2021) indicate that $\alpha \approx 0.01$. Given this parameter value, equation (3) implies that the dispersion in Z/Q should closely mirror that of ω.

[9] We focus on SO_2 in the main text. Appendix Figure A.3 reports dispersion patterns for other pollutants.

proxies for ω. These findings, consistent with evidence from the United States (Lyubich et al., 2018), indicate that the Cobb-Douglas specification of the CT framework cannot rationalize the observed firm-level heterogeneity in environmental outcomes.

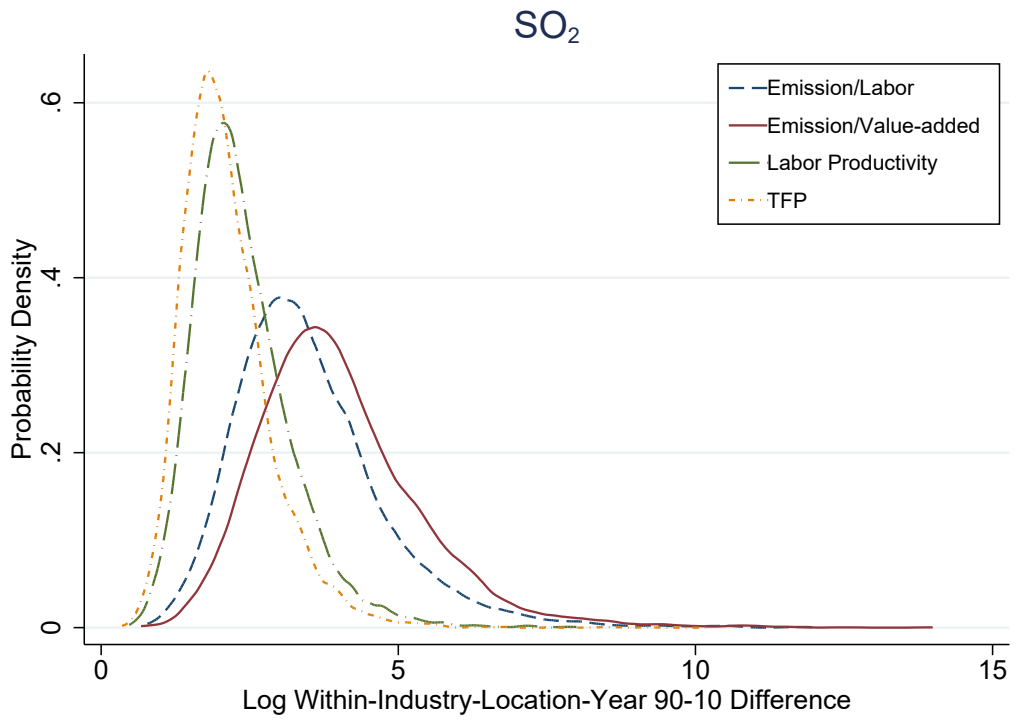


Figure 1: Dispersion of Environmental and Economic Performance Measures within Industries, Locations, and Years

Notes: This figure plots the dispersion of key firm-level performance metrics. Dispersion is defined as the 90th-10th percentile difference of the log value of each metric, calculated within each four-digit industry-province-year cell. Real value added is computed using industry-specific output deflators from Brandt et al. (2012). Total Factor Productivity (TFP) is estimated following Ackerberg et al. (2015).

To resolve this empirical puzzle, we generalize the CT framework by modifying the production system as follows:

$$\begin{aligned} \text{(Production)} \quad Q &= \tilde{\omega}_H L \cdot \underbrace{\left[(1-\beta) + \beta \left((1-\pi)^{\frac{1}{\alpha}}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}}_{C(\pi)} \\ \text{(Emissions)} \quad Z &= \tilde{\omega}_H L \cdot \underbrace{\tilde{\omega}_Z^{-1}(1-\pi)^{\frac{1}{\alpha}}}_{E(\pi)} \end{aligned} \tag{4}$$

where $C(\pi)$ and $E(\pi)$ denote the *abatement cost* and *abatement effectiveness* functions, re-

spectively.[10] The key departure from the standard CT framework lies in the specification of $C(\pi)$. While the CT model assumes a linear cost structure ($C(\pi) = 1 - \pi$), our specification allows for non-linearity. In Appendix A, we show that the production function (the first line of equation (4)) can be approximated by:

$$Q(\pi) \approx \tilde{\omega}_H L \left[1 - c\pi + \left(\frac{c^2}{\sigma}\left(1 + \beta(\sigma - 1)\right) - c\right)\pi^2\right], \tag{5}$$

where $c \equiv \beta/\alpha$. This expansion reveals two sources of flexibility: it relaxes the restriction that the linear coefficient must equal unity (as in CT, where $c = 1$), and it introduces a quadratic term. This additional flexibility allows the model to capture the complex, non-linear abatement cost structures that firms face in practice.

Solving the system in equation (4) yields a CES production function with emissions as an input. This formulation extends equation (2) of the CT framework to

$$Q = \omega_H \left[L^{\frac{\sigma-1}{\sigma}} + (\omega_Z Z)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{6}$$

where ω_H and ω_Z denote Hicks-neutral and dirty-input-augmenting technology, respectively.[11]

In summary, our framework naturally extends the CT model. It (1) generalizes the abatement cost function to allow for nonlinearity; (2) yields a CES production structure that permits flexible substitution between clean and dirty inputs; and (3) preserves the theoretical equivalence between modeling emissions as a dirty input and as a by-product.[12] When we analyze the distinct roles of clean and dirty technologies, we adopt

[10] Regarding the parameters: $\tilde{\omega}_H$ represents Hicks-neutral technology; L denotes labor; and $\tilde{\omega}_H L$ represents potential gross output. β is a normalization parameter ensuring $C(0) = 1$. $\tilde{\omega}_Z$ captures baseline abatement effectiveness when no effort is exerted ($\pi = 0$). σ governs the curvature of the abatement cost function, and α represents the output elasticity of emissions, as in the CT framework.

[11] Specifically, the parameter transformations are given by:

$$\omega_H \equiv \tilde{\omega}_H (1 - \beta)^{\frac{\sigma}{\sigma-1}}$$
$$\omega_Z \equiv \left(\frac{\beta}{1 - \beta}\right)^{\frac{\sigma}{\sigma-1}} \frac{\tilde{\omega}_Z}{\tilde{\omega}_H}.$$

[12] We use the terms "dirty input" and "emissions" interchangeably throughout the paper.

an equivalent form $Q = \left[(\tilde{\omega}_C L)^{\frac{\sigma-1}{\sigma}} + (\tilde{\omega}_Z Z)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$.

Based on our CES framework, the counterpart to equation (3) is given by:

$$\begin{aligned} &\text{(Emission-Input Ratio)} \quad \frac{Z}{L} &&= \omega_Z^{\sigma-1}\left(\frac{W}{\tau}\right)^{\sigma} \\ &\text{(Emission Intensity)} \quad \frac{Z}{Q} &&= \omega_H^{-1}\omega_Z^{-1}\left(\omega_Z^{1-\sigma}\left(\frac{W}{\tau}\right)^{1-\sigma} + 1\right)^{-\frac{\sigma}{\sigma-1}}. \end{aligned} \tag{7}$$

These equations illustrate how biased technology can rationalize the heterogeneous environmental outcomes observed in the data. Crucially, variation in dirty technology (ω_Z) generates heterogeneity in emission-input ratios (Z/L) even among firms facing identical factor prices. Furthermore, the inclusion of distinct technology components allows emission intensity (Z/Q) to vary independently of Hicks-neutral technology.

In Appendix A.2, we demonstrate that the impact of biased technology on emission intensity can be decomposed into *efficiency* and *substitution* effects. The efficiency effect captures the direct reduction in emission intensity resulting from higher output for a given input mix. Conversely, the substitution effect captures the increase in emission intensity as firms substitute toward relatively more efficient dirty input (increasing Z/L). While improvements in neutral and clean technologies unambiguously lower emission intensity, the net effect of dirty technology is ambiguous. Thus, we employ the structural model to quantify the magnitude and direction of these potentially competing channels.

3 Structural Model and Estimation

3.1 Model Setup

Production Technology We model the production technology of firm j at time t as a CES aggregate of clean inputs and emissions:

$$Q_{jt} = \left[(\exp(\tilde{\omega}_{Cjt})X_{jt})^{\frac{\sigma-1}{\sigma}} + (\exp(\tilde{\omega}_{Zjt})Z_{jt})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\nu\sigma}{\sigma-1}}, \tag{8}$$

where Q_{jt} denotes real value added. Production combines a composite clean input, X_{jt}, and emissions (dirty input), Z_{jt}. The clean input is a Cobb-Douglas composite of capital (K_{jt}) and labor (L_{jt}), such that $X_{jt} = K_{jt}^{\theta} L_{jt}^{1-\theta}$.[13]

The model characterizes technical change through two firm-specific, time-varying parameters: $\tilde{\omega}_{Cjt}$ represents *clean technology*, which augments the capital-labor bundle, and $\tilde{\omega}_{Zjt}$ represents *dirty technology*, which augments emissions. The parameter $\sigma > 0$ governs the elasticity of substitution between clean and dirty inputs, while ν determines returns to scale. Observed output, Y_{jt}, differs from planned output, Q_{jt}, due to an i.i.d. shock e_{jt}, such that $Y_{jt} = Q_{jt} \exp(e_{jt})$. This term captures either classical measurement error or a productivity shock that is unanticipated by the firm at the time of input choice.

Market Structure and Firm Decisions We assume firms operate under monopolistic competition in the product market and perfect competition in input markets. Firms face a constant elasticity demand curve for their differentiated product:[14]

$$Q_{jt} = \Phi_{jt} P_{jt}^{\eta},$$

where P_{jt} denotes price, $\eta < -1$ is the price elasticity of demand, and Φ_{jt} is a firm-specific demand shifter.

The firm's decision problem is dynamic regarding capital accumulation but static regarding variable inputs. At the beginning of period t, the firm observes its state variables: the capital stock K_{jt} (determined in $t-1$) and the vector of current shocks $(\tilde{\omega}_{Cjt}, \tilde{\omega}_{Zjt}, \xi_{Djt})$. Given this state, the firm chooses variable inputs (L_{jt}, Z_{jt}) to maximize current profits and chooses investment to determine the next period's capital stock (K_{jt+1}).

Our estimation relies on the firm's static optimization problem. The firm chooses labor

[13] While we adopt this functional form for tractability, we demonstrate in Appendix C that our identification strategy is robust to more general production structures.

[14] Following standard practice in the structural production function literature (De Loecker, 2011; Ackerberg et al., 2015; Gandhi et al., 2020), we specify this demand relationship using real value added and its implicit price. As demonstrated in Appendix B, this specification naturally derives from a gross output demand equation under the assumption of a Leontief technology for intermediate inputs. This theoretical equivalence ensures that our firm-specific demand shifter can be consistently estimated using observed real value added and the industry-level value-added price index.

and emissions to solve:

$$\begin{aligned} \max_{L_{jt}, Z_{jt}} \quad & P_{jt}Q_{jt} - W_{jt}L_{jt} - \tau_{pt}Z_{jt} \\ \text{s.t.} \quad & Q_{jt} = \left[\left(\exp(\tilde{\omega}_{Cjt}) K_{jt}^{\theta} L_{jt}^{1-\theta} \right)^{\frac{\sigma-1}{\sigma}} + \left(\exp(\tilde{\omega}_{Zjt}) Z_{jt} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\nu\sigma}{\sigma-1}} \\ & Q_{jt} = \Phi_{jt} P_{jt}^{\eta} \end{aligned}$$

where W_{jt} denotes the wage rate and τ_{pt} is the location-year-specific emission cost for a firm in location p. We interpret τ_{pt} as the *shadow cost of emissions* implied by environmental regulations. Specifically, under price-based policies such as discharge fees or cap-and-trade systems, τ_{pt} corresponds to the fee or the market price of emission rights. Under quantity-based policies such as command-and-control quotas, τ_{pt} corresponds to the Lagrange multiplier associated with the emission quantity constraint. In either case, a higher emission cost τ_{pt} captures the stringency of environmental regulation.

3.2 Identification Strategy

3.2.1 Identifying Dirty Technology and Emission Cost

To facilitate estimation, we express the production function in terms of Hicks-neutral (ω_{Hjt}) and dirty (ω_{Zjt}) technologies:

$$Q_{jt} = \exp(\omega_{Hjt}) \left[\left(K_{jt}^{\theta} L_{jt}^{1-\theta} \right)^{\frac{\sigma-1}{\sigma}} + \left(\exp(\omega_{Zjt}) Z_{jt} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\nu\sigma}{\sigma-1}}. \tag{9}$$

The original structural parameters are recovered post-estimation via the mapping $\tilde{\omega}_{Cjt} = \nu^{-1}\omega_{Hjt}$ and $\tilde{\omega}_{Zjt} = \omega_{Zjt} + \nu^{-1}\omega_{Hjt}$.

The first-order conditions for labor and emissions, derived from the firm's static profit

maximization problem, are:

$$W_{jt} = \nu(1-\theta)\frac{1+\eta}{\eta}\Phi_{jt}^{-\frac{1}{\eta}}Q_{jt}^{1+\frac{1}{\eta}-\frac{\sigma-1}{\sigma\nu}}\exp\left(\frac{\sigma-1}{\sigma\nu}\omega_{Hjt}\right)\left(K_{jt}^{\theta}L_{jt}^{1-\theta}\right)^{\frac{\sigma-1}{\sigma}}L_{jt}^{-1}, \tag{10}$$

$$\tau_{pt} = \nu\frac{1+\eta}{\eta}\Phi_{jt}^{-\frac{1}{\eta}}Q_{jt}^{1+\frac{1}{\eta}-\frac{\sigma-1}{\sigma\nu}}\exp\left(\frac{\sigma-1}{\sigma\nu}\omega_{Hjt}\right)\exp\left(\frac{\sigma-1}{\sigma}\omega_{Zjt}\right)Z_{jt}^{-\frac{1}{\sigma}}. \tag{11}$$

Combining these conditions yields an explicit expression for dirty technology ω_{Zjt}:

$$\omega_{Zjt} = \frac{\sigma}{\sigma-1}\ln(1-\theta) + \frac{\sigma}{\sigma-1}\ln\left(\frac{\tau_{pt}}{W_{jt}}\right) + \theta\ln\left(\frac{K_{jt}}{L_{jt}}\right) + \frac{1}{\sigma-1}\ln\left(\frac{Z_{jt}}{L_{jt}}\right). \tag{12}$$

Equation (12) highlights the fundamental identification challenge: data on input ratios (K_{jt}/L_{jt} and Z_{jt}/L_{jt}) are insufficient to separately identify dirty technology (ω_{Zjt}) from emission cost (τ_{pt}). Intuitively, a high emission-to-labor ratio can be rationalized either by lenient regulation (low τ_{pt}/W_{jt}) or by advanced dirty technology, which induces substitution toward the relatively more efficient dirty input. Because the emission cost is unobserved, we cannot employ standard methods that rely on observed factor prices to identify non-neutral technologies (Doraszelski and Jaumandreu, 2018; Raval, 2019).

To resolve this indeterminacy, we leverage the information contained in factor revenue shares.[15] Because ω_{Zjt} and τ_{pt} affect the labor share of revenue differentially, they can be identified with an additional equation governing the "true" labor share of revenue, $S_{Ljt}^{*} \equiv \frac{W_{jt}L_{jt}}{P_{jt}Q_{jt}}$.[16] Plugging the production function and first order condition for labor yields the following share equation:

$$S_{Ljt}^{*} = \frac{\nu(1-\theta)\frac{1+\eta}{\eta}\left(K_{jt}^{\theta}L_{jt}^{1-\theta}\right)^{\frac{\sigma-1}{\sigma}}}{\left(K_{jt}^{\theta}L_{jt}^{1-\theta}\right)^{\frac{\sigma-1}{\sigma}} + \left(\exp(\omega_{Zjt})Z_{jt}\right)^{\frac{\sigma-1}{\sigma}}}. \tag{13}$$

Substituting equation (12) into this expression eliminates ω_{Zjt} and yields a closed-form

[15] We define the revenue share as the ratio of factor expenditure to value added.

[16] We use the term "true" to emphasize that the revenue in the denominator is based on Q_{jt}, which is not directly observed due to the noise term e_{jt}.

solution for the unobserved emission cost:

$$\tau_{pt} = \frac{W_{jt}}{Z_{jt}/L_{jt}} \left(\nu \frac{1+\eta}{\eta} \frac{1}{S^*_{Ljt}} - \frac{1}{1-\theta} \right). \tag{14}$$

Equations (13) and (14) demonstrate that the observed input ratio Z_{jt}/L_{jt} and labor share S^*_{Ljt} jointly identify the unobserved technology ω_{Zjt} and emission cost τ_{pt}, conditional on $\{K_{jt}, L_{jt}, Z_{jt}, W_{jt}\}$.

3.2.2 Identifying Hicks-Neutral Technology

The true labor share, S^*_{Ljt}, is unobserved due to the output shock e_{jt}. It relates to the observed share, S_{Ljt}, via the identity $S^*_{Ljt} = S_{Ljt} \exp(e_{jt})$. As derived in Appendix D.1, we combine the logarithmic transformation of the production function with the first-order conditions and the labor share identity, yielding a structural equation for output that depends solely on factor inputs, the neutral technology ω_{Hjt}, the true labor share S^*_{Ljt}, and the error term e_{jt}:

$$\ln Y_{jt} = \omega_{Hjt} + \nu\theta \ln\left(\frac{K_{jt}}{L_{jt}}\right) + \nu \ln L_{jt} + \frac{\nu\sigma}{\sigma-1} \ln\left(\nu \frac{(1+\eta)(1-\theta)}{\eta}\right) - \frac{\nu\sigma}{\sigma-1} \ln S^*_{Ljt} + e_{jt}. \tag{15}$$

Equation (15) simplifies the identification problem considerably. By conditioning on the true labor share, we reduce the system from two unknown technologies to a single unobserved scalar, ω_{Hjt}. This allows us to apply standard control function methods (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015).[17] Specifically, we follow Ackerberg et al. (2015) and use intermediate input expenditure (m_{jt}) as a proxy for ω_{Hjt}.[18] The validity of this approach requires that intermediate input demand be strictly monotonic in ω_{Hjt}, a condition generally satisfied in this class of models.[19]

[17] Alternatively, Demirer (2024) propose a non-parametric approach that replaces the proxy variable with exogenous controls for Hicks-neutral technology.

[18] As demonstrated by Ackerberg et al. (2015) and Gandhi et al. (2020), under a Leontief structure, intermediate input demand is structurally proportional to intended value-added output. Consequently, conditioning non-parametrically on m_{jt} in the first stage perfectly absorbs the predictable structural components of output. This equivalence allows us to consistently isolate the unanticipated shock e_{jt}.

[19] As shown by Levinsohn and Petrin (2003), this condition holds if ω_{Hjt} increases the marginal product of inputs and the marginal product of intermediate inputs is non-decreasing in labor.

The firm's static optimization yields an intermediate input demand function of the form $m_{jt} = m(\omega_{Hjt}, \omega_{Zjt}, k_{jt}, W_{jt}, \tau_{pt})$, where lowercase letters denote logarithmic values.

Under the monotonicity assumption, this intermediate input demand function can be inverted. Substituting equations (12) and (14) for ω_{Zjt} and τ_{pt} in the inverted function expresses ω_{Hjt} as a function of observables and the true labor share of revenue, S^*_{Ljt}:

$$\omega_{Hjt} = \omega_H(m_{jt}, k_{jt}, z_{jt}, l_{jt}, W_{jt}, S^*_{Ljt}).$$

Substituting this inverted proxy function into equation (15) yields a general expression for observed output:

$$\ln Y_{jt} = \phi(m_{jt}, k_{jt}, z_{jt}, l_{jt}, W_{jt}, S^*_{Ljt}) + e_{jt} \equiv \tilde{\phi}(\cdot, e_{jt}). \tag{16}$$

Because the error term e_{jt} enters the non-parametric component $\phi(\cdot)$ non-separably through S^*_{Ljt} in addition to its additive role, standard identification results do not immediately apply. Following Imbens and Newey (2009), identification in such non-separable models is achieved if the structural function $\tilde{\phi}(\cdot, e_{jt})$ is strictly monotonic in the unobserved shock, e_{jt}. In Appendix D.2, we establish that a sufficient condition for $\tilde{\phi}(\cdot, e_{jt})$ to be strictly decreasing in e_{jt} is $\nu > \frac{\sigma-1}{\sigma} > 0$, a theoretical restriction that we verify empirically with our model estimates.

In summary, our structural model achieves identification through a synthesis of multiple sources of variation: the static first-order conditions for profit maximization, the observed labor share of revenue, and the control function linking Hicks-neutral technology to intermediate input demand. This approach addresses the unique challenges of our empirical setting by integrating insights from three distinct strands of the production function literature, as detailed below.

Remark 1. Standard methods identify biased technology with observed input prices (Doraszelski and Jaumandreu, 2018; Raval, 2019; Wang, 2025). We extend this literature

by exploiting the labor share of revenue to identify biased technology when a factor price is unobserved. In our context, the unobserved factor price is emission cost.

Remark 2. Leveraging revenue shares to infer unobserved heterogeneity is an established identification strategy. Share-based first-order conditions are frequently employed to recover unobserved input prices in settings *without* biased technology (Grieco et al., 2016; Wang et al., 2025) or to identify structural parameters in other contexts (Zhang, 2019; Hassler et al., 2021; Demirer, 2024). We apply this strategy to the environmental context, where accounting for biased technical change is crucial.

Remark 3. Finally, our strategy integrates the canonical proxy variable framework (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015). We use the control function approach to proxy for ω_{Hjt}, constructing a semi-linear equation that isolates the unanticipated output shock e_{jt}.

3.3 Estimation Procedure

Our estimation proceeds in two stages. First, we estimate the unanticipated output shock, e_{jt}. Second, we use this estimate as an input for the iterative GMM procedure to recover the full set of structural parameters.

First Stage: Estimation of the Output Shock The objective of the first stage is to recover the unanticipated output shock, e_{jt}, which is required to construct the true labor share S^*_{Ljt}. Following Zhang (2019), we approximate the unknown function $\phi(\cdot)$ in equation (16) using a semi-linear specification. We model the non-parametric component using a third-order polynomial $\varphi(\cdot;\chi)$, where χ denotes the vector of coefficients to be estimated. The specification also includes industry and county-year fixed effects.[20] This yields the

[20]Constructing real value added ($\ln Y_{jt}$) in equation (17) requires deflating nominal value added by a firm-level value-added price index (P^{VA}_{jt}). Because this firm-level index is unavailable, we proxy it using the industry-level gross output price index ($P_{Output,jt}$). As shown in Appendix B, these two deflators are related by $P_{Output,jt} = P^{VA}_{jt}/\lambda_{\text{ind}(j)}$. Deflating by the gross output price index introduces an additive industry-specific term, $\ln(\lambda_{\text{ind}(j)})$. We absorb this discrepancy in the first stage by including industry fixed

estimation equation:

$$\ln Y_{jt} = \varphi(m_{jt}, k_{jt}, z_{jt}, l_{jt}, W_{jt}; \chi) + \chi_{S_L} \ln S_{Ljt} + \varepsilon_{jt}, \tag{17}$$

where the composite error term is proportional to the true output shock: $\varepsilon_{jt} = (1+\chi_{S_L})e_{jt}$.

The observed labor share, $\ln S_{Ljt}$, is endogenous due to its mechanical correlation with the output shock e_{jt}. We address this by estimating equation (17) via Two-Stage Least Squares (2SLS), using the lagged wage $\ln(W_{jt-1})$ as an instrument. The lagged wage is a valid instrument as it is correlated with current factor prices (and thus the labor share) but is orthogonal to the unanticipated period-t output shock. The 2SLS estimation yields consistent estimates of the polynomial coefficients $\hat{\chi}$ and the labor share parameter $\hat{\chi}_{S_L}$. We then recover the output shock for each observation:

$$\hat{e}_{jt} = \frac{\ln Y_{jt} - \varphi(m_{jt}, k_{jt}, z_{jt}, l_{jt}, W_{jt}; \hat{\chi}) - \hat{\chi}_{S_L} \ln S_{Ljt}}{1 + \hat{\chi}_{S_L}}. \tag{18}$$

Using this estimate, we construct the "true" labor share, $\hat{S}^*_{Ljt} = S_{Ljt} \exp(\hat{e}_{jt})$, and "true" output, $\hat{\phi}_{jt} = \ln Y_{jt} - \hat{e}_{jt}$, which serve as inputs for the second stage.

Second Stage: Iterative GMM Estimation The second stage estimates the structural parameters using iterative GMM. We parameterize the demand shifter as a function of observable characteristics–firm size (proxied by capital stock K_{jt}), age, and state ownership status (SOE)–along with year (γ^D_t) and industry ($\gamma^D_{\text{ind}(j)}$) fixed effects and an i.i.d. demand shock (ξ_{Djt}):

$$\ln \Phi_{jt} = \alpha_s \text{Size}_{jt} + \alpha_a \ln(\text{Age})_{jt} + \alpha_o \text{SOE}_{jt} + \gamma^D_t + \gamma^D_{\text{ind}(j)} + \xi_{Djt}. \tag{19}$$

effects within the function $\varphi(\cdot)$. Furthermore, this term is consistently absorbed in the second stage by the industry fixed effects ($\gamma^D_{\text{ind}(j)}$, $\gamma^H_{\text{ind}(j)}$, $\gamma^Z_{\text{ind}(j)}$) included in the demand equation (19) and the VAR processes (20)–(21), leaving the structural parameters consistently identified.

We assume that technologies evolve according to an interdependent VAR(1) process:

$$\omega_{Hjt} = \gamma_t^H + \gamma_{\text{ind}(j)}^H + \gamma_{HH}\omega_{Hjt-1} + \gamma_{HZ}\omega_{Zjt-1} + \gamma_{H\tau}\tau_{pt-1} + \xi_{Hjt}, \tag{20}$$

$$\omega_{Zjt} = \gamma_t^Z + \gamma_{\text{ind}(j)}^Z + \gamma_{ZZ}\omega_{Zjt-1} + \gamma_{ZH}\omega_{Hjt-1} + \gamma_{Z\tau}\tau_{pt-1} + \xi_{Zjt}, \tag{21}$$

where γ_t and $\gamma_{\text{ind}(j)}$ denote year and industry fixed effects, respectively, and ξ_{Hjt} and ξ_{Zjt} represent i.i.d. technology shocks. We include the lagged emission cost τ_{pt-1} to capture induced technical change.

Identification relies on the assumption that the productivity innovations (ξ_{Hjt}, ξ_{Zjt}), along with the demand shock ξ_{Djt}, are orthogonal to the firm's information set at time $t-1$. This yields the moment condition:

$$E\left[\mathbf{M}'_{jt-1}\xi_{jt}\right] = 0,$$

where $\xi_{jt} = (\xi_{Hjt}, \xi_{Zjt}, \xi_{Djt})'$ is the vector of shocks and $\mathbf{M}_{jt-1}$ is a matrix of instruments belonging to the information set at $t-1$.

To evaluate the moment conditions, we must recover the unobserved shocks. The demand shock ξ_{Djt} is recovered directly from the demand equation:

$$\xi_{Djt} = \ln Q_{jt} - \eta \ln P_{jt} - \alpha_s \text{Size}_{jt} - \alpha_a \ln(\text{Age})_{jt} - \alpha_o \text{SOE}_{jt} - \gamma_t^D - \gamma_{\text{ind}(j)}^D. \tag{22}$$

Recovering the technology shocks requires deriving the unobserved technology levels. Using the first-stage estimate $\hat{S}^*_{Ljt}$, we compute τ_{pt} via equation (14). We then invert the production system to express ω_{Zjt} and ω_{Hjt} in terms of observables and parameters:

$$\omega_{Zjt} = \frac{\sigma}{\sigma-1}\ln(1-\theta) + \frac{\sigma}{\sigma-1}\ln\left(\frac{\tau_{pt}}{W_{jt}}\right) + \theta\ln\left(\frac{K_{jt}}{L_{jt}}\right) + \frac{1}{\sigma-1}\ln\left(\frac{Z_{jt}}{L_{jt}}\right), \tag{23}$$

$$\omega_{Hjt} = \hat{\phi}_{jt} - \nu\theta\ln\left(\frac{K_{jt}}{L_{jt}}\right) - \nu\ln L_{jt} - \frac{\nu\sigma}{\sigma-1}\ln\left(1 + (1-\theta)\frac{\tau_{pt}}{W_{jt}}\frac{Z_{jt}}{L_{jt}}\right). \tag{24}$$

We calculate the current and lagged values of $(\omega_{Hjt}, \omega_{Zjt})$ using these expressions and

substitute them into the VAR(1) system to recover the technology shocks (ξ_{Hjt}, ξ_{Zjt}).

We define the matrix of instrumental variables $\mathbf{M}_{jt-1}$ using three sets of lagged observables corresponding to the three shock equations. The instruments for the Hicks-neutral technology process are $\mathbf{M}_{H,jt-1} = (\hat{\phi}_{jt-1}, \ln K_{jt-1}, \ln L_{jt-1}, \hat{S}^*_{Ljt-1})$. The instruments for the dirty technology process are $\mathbf{M}_{Z,jt-1} = (\ln Z_{jt-1}, \ln K_{jt-1}, \ln L_{jt-1}, \hat{S}^*_{Ljt-1})$. The instruments for the demand equation are $\mathbf{M}_{D,jt-1} = (\ln P_{jt-1}, \text{Size}_{jt-1}, \ln(\text{Age})_{jt-1}, \text{SOE}_{jt-1})$.

Let Θ denote the full vector of parameters to be estimated. The sample analogue of the moment condition is given by $g(\Theta) = \frac{1}{NT} \sum_{j,t} \mathbf{M}'_{jt-1} \xi_{jt}(\Theta)$. The GMM estimator is defined as the parameter vector that minimizes the criterion function $J(\Theta) = g(\Theta)' \mathbf{W} g(\Theta)$, where $\mathbf{W}$ is a positive definite weighting matrix.

Directly minimizing this criterion function to estimate the full system simultaneously is computationally intractable due to the high-dimensional fixed effects present in both the VAR(1) technology process and the demand equation. For tractability, we employ an iterative procedure that separates the estimation of these fixed effects from the estimation of the core structural parameters.[21] Each iteration takes the production parameters (ν, σ, θ) and demand $(\eta, \alpha_s, \alpha_a, \alpha_o)$ parameters as given to recover the implied technology series $(\omega_{Hjt}, \omega_{Zjt})$ and estimate the linear parameters of the VAR(1) and demand equations via OLS. This separation allows the computationally intensive GMM optimization to focus exclusively on those core structural parameters.

Specifically, our estimation proceeds via the following iterative loop:

(i) Initialization: We initialize the algorithm with a guess for the structural parameters and VAR coefficients.

(ii) GMM Estimation: In iteration k, taking the emission costs $\tau^{(k)}$ and linear coefficients $\gamma^{(k)}$ as given, we minimize the GMM objective function to obtain updated production and demand parameters $(\nu, \sigma, \theta, \eta, \alpha)$.

[21] The iterative separation of high-dimensional fixed effects from structural parameters is a standard technique in empirical industrial organization (e.g., Berry et al., 1995). For instance, Bayer et al. (2007) use a contraction mapping to separate the estimation of location-specific fixed effects from household preference parameters.

(iii) Parameter Update: Using the updated structural estimates from Step (ii), we perform two updates:

(a) *Update Emission Costs:* We compute the firm-level emission cost, $\hat{\tau}_{jpt}$, via equation (14). To ensure numerical stability, negative values are replaced with a small positive constant (10^{-5}). We then update the aggregate emission cost, $\tau_{pt}^{(k+1)}$, by taking the trimmed mean of these firm-level estimates within each prefecture-year cell.[22]

(b) *Update Linear Parameters:* Using $\tau_{pt}^{(k+1)}$, we reconstruct the technology series ω_{Hjt} and ω_{Zjt}. We then update the VAR coefficients ($\gamma^{(k+1)}$) by estimating equations (20) and (21) via OLS. Similarly, we update the demand fixed effects ($\gamma_t^{D(k+1)}, \gamma_{\text{ind}(j)}^{D(k+1)}$) by estimating equation (19), given the α parameters.

(iv) Convergence: We repeat steps (ii) and (iii) until the parameter vector converges.

(v) Inference: Standard errors are computed via bootstrapping.

4 Data Description

This section outlines the data sources, sample construction, and variable definitions used in our empirical analysis.

4.1 Data Sources

We rely on two primary firm-level datasets. The first is China's Environmental Statistical Database (CESD), compiled by the Ministry of Ecology and Environment. The CESD provides comprehensive plant-level pollution data, covering firms that cumulatively account for at least 85 percent of emissions for major pollutants—including sulfur dioxide (SO_2), chemical oxygen demand (COD), and nitrogen oxides (NO_x)—within each county-year cell.

[22] We trim the top and bottom 10% of firm-level τ estimates within each cell to mitigate the influence of outliers.

The second source is the Annual Survey of Industrial Enterprises (ASIE), maintained by China's National Bureau of Statistics. This dataset covers all state-owned enterprises (SOEs) and non-SOEs with annual sales exceeding 5 million RMB, providing detailed production and financial information. We merge these datasets for the 2000–2008 period using the firm identifiers and matching algorithms developed in Kwon et al. (2023).[23] The resulting matched panel contains detailed information on firm-level emissions, output, value added, employment, wage, capital stock, and intermediate inputs.

4.2 Sample Construction and Key Variables

Our analysis focuses on SO_2 emissions. To ensure data quality and consistency with our theoretical framework, we impose three sample restrictions. First, we trim the top and bottom 1 percent of the annual distribution for employment, wages, real capital, intermediate inputs, and value added. Second, we restrict the sample to firms with a labor share of revenue between 0.5 and 99 percent. Finally, we exclude the bottom 5 percent of the annual SO_2 emissions distribution, as these firms likely face negligible abatement incentives.

We construct real value added and real intermediate inputs using industry-specific output and input price deflators from Brandt et al. (2012) and provincial price indices from the *China Statistical Yearbook*. Real capital stock is constructed using the perpetual inventory method. Labor is measured as the total number of employees, and the wage rate is calculated as the total wage bill divided by employment. The labor share is defined as the ratio of the total wage bill to nominal value added. Appendix Table H.1 reports summary statistics for these variables.

[23] We restrict our analysis to the pre-2008 period because of significant changes in survey methodology, namely the omission of value-added and intermediate input data, together with growing concerns regarding the reliability of reported output in later years (Brandt et al., 2023). Since our identification strategy requires lagged instruments, the estimation sample begins in 2001.

5 Empirical Results

5.1 Structural Parameter Estimates

Table 1 presents the structural parameter estimates from the second-stage iterative GMM procedure. Panel A reports the production and demand parameters. The returns to scale parameter ($\nu = 0.992$) indicates that production exhibits effectively constant returns to scale. The elasticity of substitution is estimated at $\sigma = 2.763$, implying a high degree of substitutability between the composite clean input and emissions. The capital share parameter ($\theta = 0.291$) is consistent with standard estimates in the literature.

The price elasticity of demand is estimated at $\eta = -1.411$, which is consistent with our assumption of imperfect competition and indicates that firms possess considerable market power. The demand shifters behave as expected: conditional on price, demand is increasing in firm size (α_s) and age (α_a), but is lower for state-owned enterprises (α_o).

Panels B and C detail the technology evolution processes. Panel B reports the VAR(1) estimates for the auxiliary technology variables used for estimation (ω_H, ω_Z), while Panel C presents the implied dynamics for the original technology variables ($\tilde{\omega}_C, \tilde{\omega}_Z$). In both specifications, technology exhibits a high degree of persistence, evidenced by the large and significant autoregressive coefficients ($\gamma_{HH}, \gamma_{ZZ}, \gamma_{\tilde{C}\tilde{C}}$ and $\gamma_{\tilde{Z}\tilde{Z}}$). This strong persistence aligns with prior findings using Chinese firm-level data (Zhang, 2019; Wang, 2025).

Crucially, the lagged emission cost has a large, positive, and statistically significant effect on the subsequent evolution of dirty technology ($\gamma_{Z\tau} = 0.977$ and $\gamma_{\tilde{Z}\tau} = 1.051$). This finding provides empirical support for the Porter Hypothesis (Porter and van der Linde, 1995) and the theory of directed technical change (Acemoglu et al., 2012; Calel and Dechezleprêtre, 2016), suggesting that regulation incentivizes firms to direct innovation toward improving emission efficiency.

In contrast, the effect of the emission cost on neutral or clean technologies ($\gamma_{H\tau}$ and $\gamma_{\tilde{C}\tau}$, respectively) is economically small. This asymmetry suggests that the emission cost directs technical change toward dirty inputs rather than broad productivity gains. We

Table 1: GMM Estimation Results

	Estimate	Std. Error
Panel A: Production and Demand Parameters		
Returns to Scale (ν)	0.992***	(0.012)
Elasticity of Substitution (σ)	2.763***	(0.210)
Capital Share in Clean Input (θ)	0.291***	(0.030)
Price Elasticity of Demand (η)	-1.411***	(0.014)
Demand Shifter: Size (α_s)	0.691***	(0.001)
Demand Shifter: Age (α_a)	0.017***	(0.003)
Demand Shifter: SOE (α_o)	-0.168***	(0.006)
Panel B: Technology Evolution (Estimation Stage VAR(1))		
Dependent Variable: ω_H		
Lagged ω_H (γ_{HH})	0.618***	(0.005)
Lagged ω_Z (γ_{HZ})	-0.016***	(0.001)
Lagged τ ($\gamma_{H\tau}$)	0.074**	(0.030)
Dependent Variable: ω_Z		
Lagged ω_H (γ_{ZH})	-0.076***	(0.016)
Lagged ω_Z (γ_{ZZ})	0.581***	(0.008)
Lagged τ ($\gamma_{Z\tau}$)	0.977***	(0.160)
Panel C: Technology Evolution (Applying $\tilde{\omega}_C = \nu^{-1}\omega_H$ and $\tilde{\omega}_Z = \omega_Z + \nu^{-1}\omega_H$)		
Dependent Variable: $\tilde{\omega}_C$		
Lagged $\tilde{\omega}_C$ ($\gamma_{\tilde{C}\tilde{C}}$)	0.634***	(0.006)
Lagged $\tilde{\omega}_Z$ ($\gamma_{\tilde{C}\tilde{Z}}$)	-0.016***	(0.001)
Lagged τ ($\gamma_{\tilde{C}\tau}$)	0.074**	(0.030)
Dependent Variable: $\tilde{\omega}_Z$		
Lagged $\tilde{\omega}_C$ ($\gamma_{\tilde{Z}\tilde{C}}$)	-0.023	(0.015)
Lagged $\tilde{\omega}_Z$ ($\gamma_{\tilde{Z}\tilde{Z}}$)	0.565***	(0.008)
Lagged τ ($\gamma_{\tilde{Z}\tau}$)	1.051***	(0.155)

Notes: This table reports the main parameter estimates from the iterative GMM procedure. Bootstrap standard errors are reported in parentheses. Panel B reports the VAR(1) coefficients for the auxiliary technologies used in estimation. Panel C reports the implied VAR(1) process for the structural technologies of interest. All VAR regressions include year and industry fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

explore the drivers of these technology responses further in Section 6.

Finally, the coefficients governing the interaction between technologies are generally negative and significant. This pattern implies a substitution effect in innovation: im-

provements in one technology tend to crowd out advancements in the other in subsequent periods.

5.2 Aggregate Trends in Biased Technology and Emission Costs

We use the structural estimates from Table 1 to recover firm-specific measures of clean and dirty technology ($\tilde{\omega}_{Cjt}, \tilde{\omega}_{Zjt}$) and the prefecture-level emission cost (τ_{pt}). To characterize the aggregate evolution of these variables, we estimate the following specification:

$$y_{jt} = \Psi_t + \Omega_j + \epsilon_{jt}, \tag{25}$$

where y_{jt} denotes the logarithm of clean technology, dirty technology, emission cost, or emission intensity (Z_{jt}/Q_{jt}). The term Ψ_t captures year fixed effects, and Ω_j absorbs firm (or prefecture) heterogeneity.

Figure 2 plots the estimated time trends ($\hat{\Psi}_t$). All three structural components—clean technology, dirty technology, and emission cost—exhibit substantial growth over the sample period, coinciding with a marked decline in aggregate emission intensity. However, the growth rates diverge sharply: while clean technology grew by less than 40 percent, both dirty technology and emission cost rose by more than 200 percent. As discussed in Appendix A.2, an increase in clean technology or emission cost unambiguously lowers emission intensity, whereas the net effect of dirty technical change is theoretically ambiguous. We quantify the relative contribution of each channel via structural decomposition in Section 7.

In Appendix E, we document the substantial heterogeneity underlying these aggregate trends. To ensure our comparisons reflect fundamental structural divides rather than compositional differences, we partial out the appropriate temporal, industrial, and spatial fixed effects before comparing clean technology ($\tilde{\omega}_C$), dirty technology ($\tilde{\omega}_Z$), emission intensity, and emission costs. We find that non-state-owned and larger enterprises possess superior clean technology and operate with significantly lower emission intensity than

their state-owned and smaller counterparts. Geographically, the economically developed Eastern region leads in clean technical efficiency, while the Central region relies most heavily on dirty technology and faces the highest regulatory emission costs. These distributional patterns underscore the highly uneven allocation of environmental capabilities and regulatory burdens across China's industrial landscape.

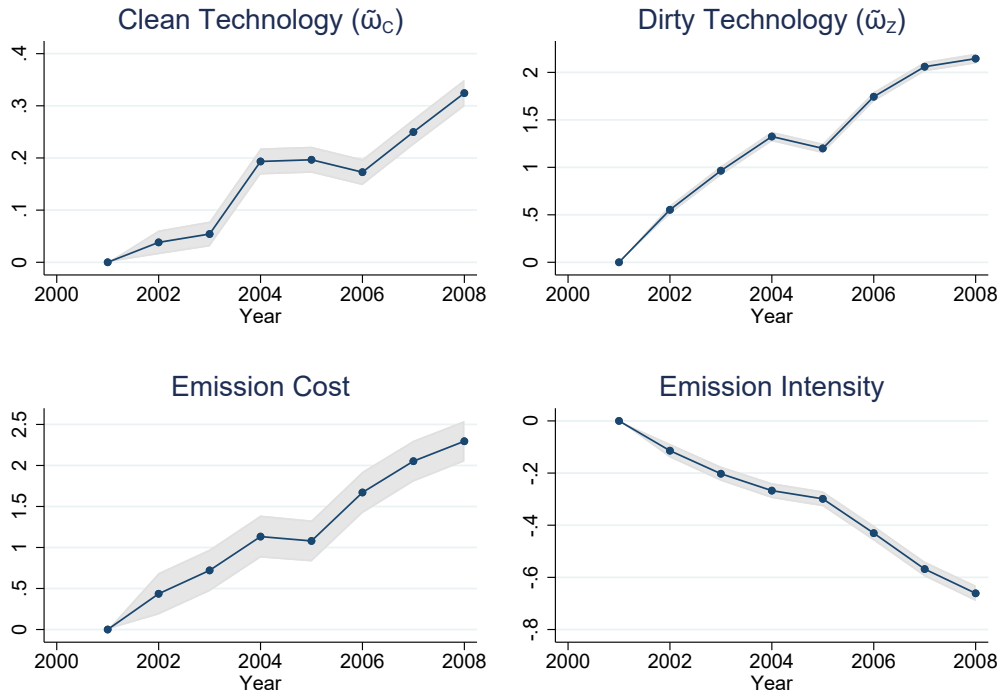


Figure 2: Aggregate Trends in Technology, Emission Cost, and Emission Intensity

Notes: The figure plots the estimated time trends for clean and dirty technologies ($\tilde{\omega}_C$ and $\tilde{\omega}_Z$), log emission cost ($\ln \tau$), and log emission intensity ($\ln(Z/Q)$). Each series corresponds to the year fixed effects (Ψ_t) estimated from equation (25). All series are normalized to zero in 2001.

6 The Drivers of Emission Costs and Technology

This section investigates the determinants of the emission cost and the two technologies, focusing on the roles of environmental regulation and firm innovation.

6.1 Environmental Regulation and the Emission Cost

We first examine whether our estimated emission cost, τ_{pt}, reflects observable variation in regulatory stringency. During our sample period, China's primary price-based instrument for SO_2 control included emission discharge fees collected by local environmental protection agencies (EPAs) (Wang and Wheeler, 2005).

To validate our estimates, we compare τ_{pt} against the actual discharge fees paid by firms. Following Wang and Wheeler (2005), we proxy for *de facto* enforcement stringency using firm-level emission fee payments reported in the 2004 ASIE to compute the average discharge fee per unit of SO_2 at the prefecture level.[24]

Figure 3 plots our estimated emission cost against this measure. We find a positive and statistically significant correlation: a 10 percent increase in the local discharge fee is associated with a 2.8 percent increase in the estimated emission cost. This relationship provides external validation for our structural estimates, confirming that τ_{pt} captures a key component of the firm's marginal cost of emissions.[25]

6.2 Innovation and Technical Changes

To characterize the drivers of technical progress, we augment the technology evolution process to directly incorporate observable measures of firm-level innovation. Following the structural literature on endogenous productivity (Aw et al., 2011; Doraszelski and Jaumandreu, 2013), we include measures of innovation output (patents) and input (R&D) directly in the autoregressive equations for both clean and dirty technologies. Table 2 presents the results. We find a robust, positive, and statistically significant correlation be-

[24]The "emissions fee" reported in the ASIE aggregates payments for multiple pollutants, notably COD. To address potential measurement error, we verify robustness by recalculating the prefecture-level fee after excluding water-pollution-intensive firms (i.e., those in the top quartile of the COD emissions distribution). Our findings are insensitive to this exclusion.

[25]We also investigate the impact of major command-and-control policies during this period, specifically the "Two Control Zones" (TCZ) designation (Chen et al., 2018; Lu and Pless, 2026) and the SO_2 reduction targets in the 11th Five-Year Plan. As detailed in Appendix F.1, we find mixed and largely insignificant effects on the shadow cost. This likely reflects the endogenous and dynamic nature of policymaking, where areas previously under loose regulation (and thus facing lower initial shadow costs) are subsequently targeted for stricter enforcement, complicating the cross-sectional identification of these policies without structural controls.

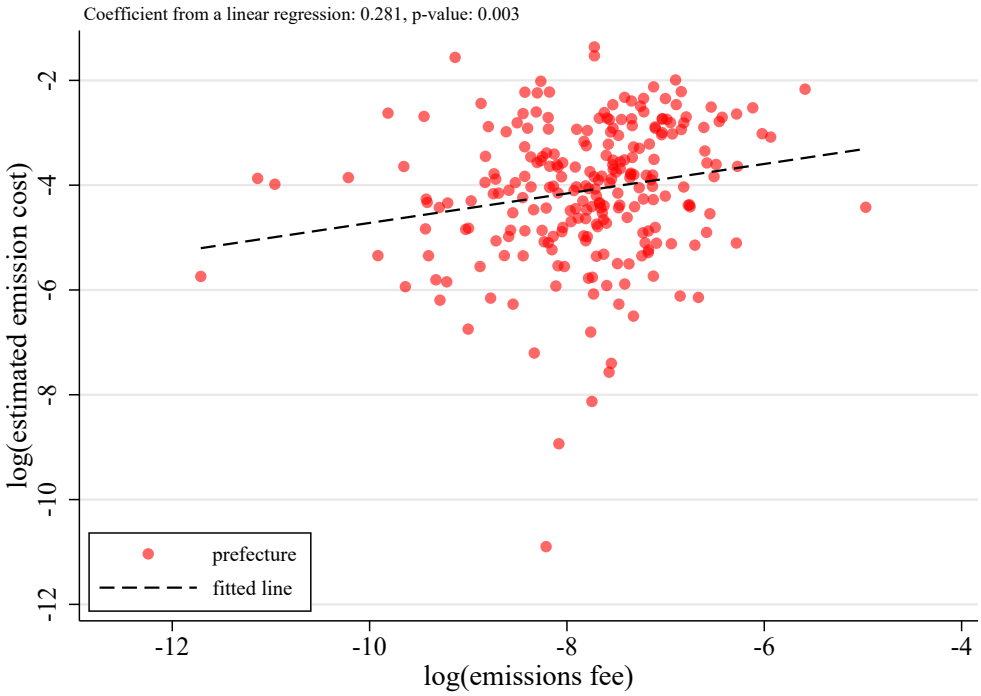


Figure 3: Estimated Emission Cost versus Average Discharge Fees (2004)

Notes: The figure plots the logarithm of the estimated emission cost against the logarithm of the prefecture-level average discharge fee. Each observation represents a prefecture in 2004. The dashed line represents the linear fit. Discharge fees are calculated using firm-level payment data from the 2004 ASIE.

tween clean technology and all innovation metrics, including current patent applications, patenting history, and R&D expenditure. In contrast, we find no statistically significant association for dirty technology.

This dichotomy suggests that general innovation activities, as captured by standard R&D and patent data, primarily improve the efficiency of clean inputs (capital and labor) rather than emission efficiency. Appendix Figure A.2 corroborates this pattern, displaying a distinct rightward shift in the distribution of ξ_C for innovative firms, whereas the distribution of ξ_Z remains close between innovative and non-innovative firms.

Table 2: The Impact of Innovation Activities on Non-Neutral Technologies

VARIABLES	(1) $\tilde{\omega_C}$	(2) $\tilde{\omega_Z}$	(3) $\tilde{\omega_C}$	(4) $\tilde{\omega_Z}$	(5) $\tilde{\omega_C}$	(6) $\tilde{\omega_Z}$	(7) $\tilde{\omega_C}$	(8) $\tilde{\omega_Z}$
Patent Application	0.145***	0.0507						
	(0.0169)	(0.0396)						
Patent Application Ever			0.114***	0.0330				
			(0.0109)	(0.0263)				
Patent Application All					0.00577***	-0.000167		
					(0.000935)	(0.00179)		
ln(R&D)							0.0220***	-0.00471
							(0.00467)	(0.0109)
Adjusted R^2	0.53	0.49	0.53	0.49	0.53	0.49	0.59	0.54
No. of Obs.	62228	62228	62228	62228	62228	62228	4829	4829
No. of Firms	23214	23214	23214	23214	23214	23214	2962	2962

Notes: This table presents the estimated coefficients on firm-level innovation activities from the augmented VAR(1) technology evolution process. The dependent variables are the structural technologies, clean technology ($\tilde{\omega}_{Cjt}$, odd columns) and dirty technology ($\tilde{\omega}_{Zjt}$, even columns). The regressions control for lagged technologies and the lagged shadow cost of emissions as specified in the baseline model. "Patent Application" is an indicator variable equal to one if the firm applied for a patent in year t. "Patent Application Ever" is an indicator equal to one if the firm has ever applied for a patent up to year t. "Patent Application All" is the cumulative stock of patent applications by the firm up to year t. "ln(R&D)" is the natural log of the firm's R&D expenditure. Robust standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

7 Decomposition of Emission Intensity Decline and Counterfactual Policy Analysis

7.1 Decomposition of Emission Intensity Decline (2001–2008)

We quantify the drivers of the aggregate decline in emission intensity by decomposing the total change into contributions from five channels: (1) clean technical change ($\tilde{\omega}_C$), (2) dirty technical change ($\tilde{\omega}_Z$), (3) capital accumulation (K), (4) the emission cost (τ), and (5) the wage (W).

We employ a Shapley decomposition to account for the path dependency of different channels (Shorrocks, 2012; Brandt and Lim, 2024). For any subset of factors κ of the $\mathcal{K} = 5$ channels, we construct a counterfactual emission intensity $EI_t(\kappa)$ by setting the factors in κ to their period-t values while holding the remaining factors at their period-$(t-1)$ levels. Firms are assumed to adjust labor and emissions optimally to produce the observed period-t output.[26]

The contribution of channel k, denoted C_k, is quantified by its Shapley value (Shorrocks, 2012; Brandt and Lim, 2024):[27]

$$C_k = \sum_{\kappa \subseteq \mathcal{K} \setminus \{k\}} \frac{|\kappa|!(|\mathcal{K}| - |\kappa| - 1)!}{|\mathcal{K}|!} \left[EI(\kappa \cup \{k\}) - EI(\kappa) \right]. \tag{26}$$

The term in the bracket captures the marginal contribution of channel k conditional on the adjustment of the subset κ. The Shapley decomposition offers two distinct advantages in this context. First, it yields an exact additive decomposition, ensuring that the sum of individual contributions is equal to the total change. Second, it systematically addresses path dependence. Because the marginal effect of a variable depends on the concurrent

[26] To prevent outliers from skewing the decomposition, we exclude observations where the emission cost is negligible ($\tau < 0.001$), which constitutes approximately 6 percent of the sample.

[27] We calculate the contribution of channel k to the change in emission intensity between adjacent years and report the average contribution over the 2001–2008 period. An alternative approach using long differences (2001 to 2008) would require a balanced panel, significantly reducing sample size and representativeness.

levels of other determinants, the Shapley value averages these marginal effects over all possible permutation sequences, thereby generating an order-independent measure of contribution.

Table 3 reports the average decomposition results. Consistent with theory, improvements in clean technology, capital accumulation, and rising emission costs all reduced emission intensity. Dirty technical change also contributed to the decline, implying that the efficiency effect dominated the substitution effect. In contrast, the wage channel increased emission intensity, a result driven by the rapid growth of real wages during this period (Yang et al., 2010). Overall, the rising emission cost was the dominant factor, accounting for 59.6 percent of the gross reduction, while clean and dirty technical change contributed 43.5 percent and 26.0 percent, respectively.

Table 3: Decomposition of Emission Intensity Decline (2001–2008)

Decomposition of the Contribution (total: 100%)	
Clean-input-augmenting technology ($\tilde{\omega}_C$)	43.5%
Dirty-input-augmenting technology ($\tilde{\omega}_Z$)	26.0%
Capital (K)	6.8%
Shadow costs of emissions (τ)	59.6%
Wages (W)	-35.9%

Notes: This table reports the Shapley decomposition of the total change in aggregate SO_2 emission intensity, averaged over the 2001–2008 period. Each row reports the percentage contribution of the respective channel. Positive values indicate a contribution to the reduction in emission intensity; negative values indicate an increase. Contributions sum to 100% by construction.

Our finding that rising emission costs are the primary driver of industrial cleanup aligns with Shapiro and Walker (2018) for the U.S. manufacturing sector. However, our framework explicitly models biased technical change. Neglecting this dimension risks conflating technical progress with regulatory costs, thereby inflating the estimated impact of environmental policy. Consequently, while still dominant, our estimate of the contribution of τ is more attenuated relative to the literature.[28]

[28]While Shapiro and Walker (2018) do not report a specific decomposition for emission *intensity*, their results (Figure 5, Panel E) imply that regulation explains nearly the entire decline in *total* SO_2 emissions, as other shocks roughly cancel out.

7.2 Short- and Long-Run Implications of the Porter Hypothesis

Porter and van der Linde (1995) hypothesize that firms can offset the costs of stringent environmental regulation through induced innovation. We quantify the magnitude and timing of this response by simulating the evolution of marginal costs following a permanent 30 percent increase in the emission cost, using our estimated VAR(1) framework. We focus on marginal cost because it captures the net effect of two opposing forces: the direct increase in input costs from regulation and the countervailing efficiency gains from directed technical change. As detailed in Appendix G.3, we hold other primitives constant to trace the transition to the new steady state.

Figure 4 illustrates the transition dynamics. Upon impact, the direct regulatory burden raises marginal costs by 13.3 percent on average. Over time, however, induced innovation attenuates this shock, reducing the long-run cost increase to 8.1 percent relative to the baseline. This trajectory implies that directed technical change offsets approximately 40 percent of the initial compliance cost. These findings are consistent with the partial-offset interpretation of the Porter Hypothesis, a result well established in the literature (Ambec et al., 2013).

While the partial nature of this offset is well understood, our analysis reveals a novel dynamic dimension. We characterize the innovation response as robust and rapid. Marginal costs nearly converge to the new steady-state level within five years, and the majority of the adjustment occurs in the first two years. This rapid response is critical for policy-making, as the timing of the offset governs the discounted present value of compliance costs.

The subsample analysis in Figure 4 (dashed and dotted lines) reveals that the induced innovation effect is most pronounced among firms with high initial levels of dirty technology. These firms use dirty inputs more intensively and are therefore more exposed to rising emission costs.

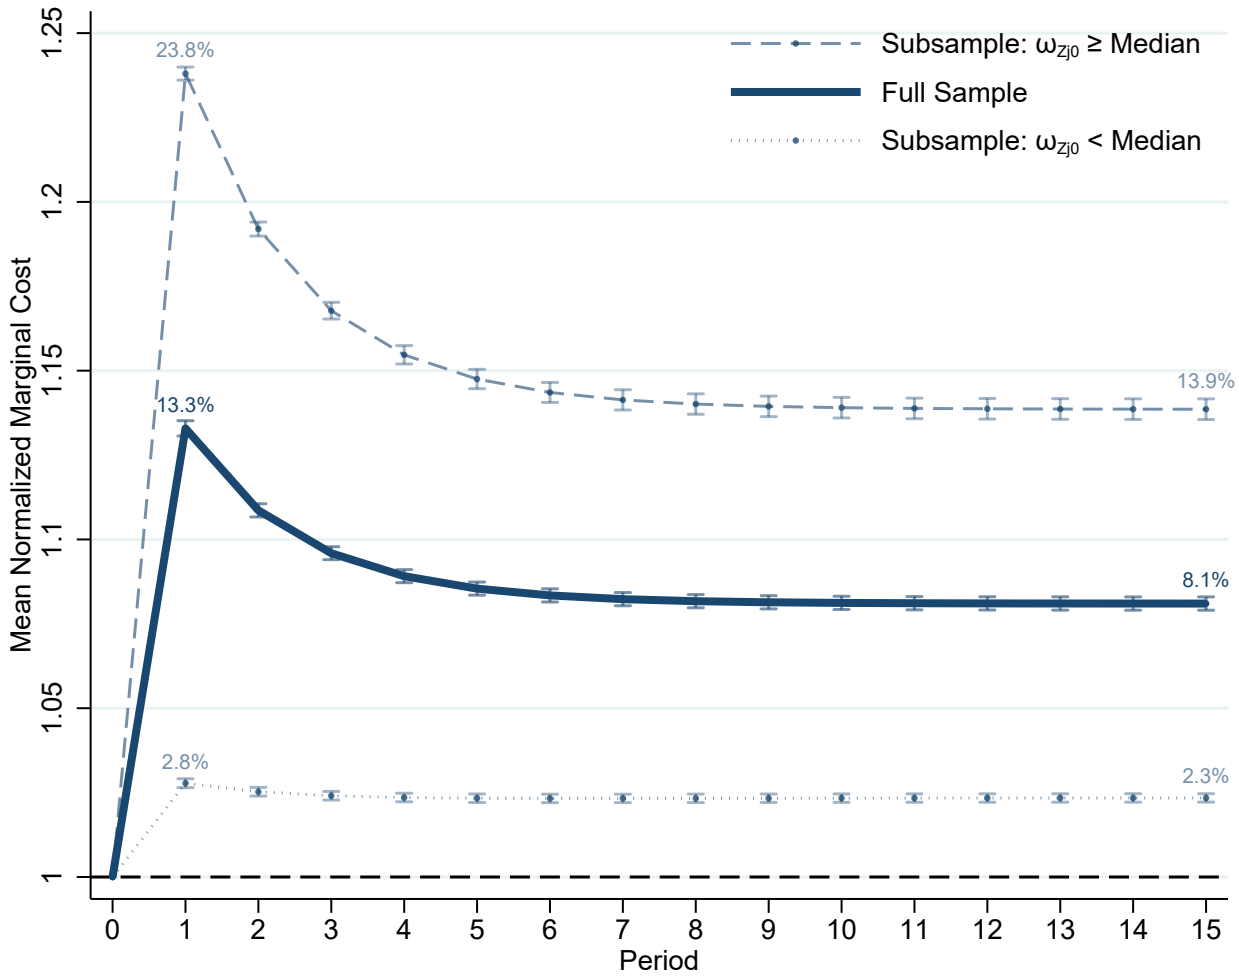


Figure 4: Transitional Dynamics of Marginal Costs under Higher Discharge Fees

Notes: The figure plots the evolution of the mean normalized marginal cost of production ($NMC_{jt} = MC_{jt}/MC_{j0}$) following a permanent 30 percent increase in the emission cost starting from $t = 1$. Period $t = 0$ represents the baseline steady state ($NMC = 1$). The solid line represents the full sample average. For heterogeneity analysis, we divide the sample based on the base-period level of ω_{Zj0}: the dashed line denotes high-ω_Z firms ($\omega_{Zj0} \geq$ Median), while the dotted line denotes low-ω_Z firms ($\omega_{Zj0} <$ Median). We also plot 95 percent confidence intervals for each estimate. Because our VAR(1) specification assumes that the emission cost affects technology with a one-year lag, the offsetting effects of induced innovation begin in $t = 2$.

7.3 Higher Discharge Fee Regime

We employ our structural framework to quantify the impact of doubling the nationwide discharge fee, accounting for both the direct cost effect and the indirect effect of induced technical change. Based on the elasticity implied by Figure 3, this policy corresponds to a 30 percent increase in the firm-level emission cost. We simulate the counterfactual evolution of aggregate outcomes, allowing biased technologies to evolve accord-

ing to the estimated VAR(1) process.

Figure 5 presents the results. The black line plots the observed aggregate emission intensity (Status Quo). The red dashed line isolates the direct cost effect, holding technology fixed at baseline levels. The blue line captures the total effect, incorporating the endogenous response of technology. By 2008, the policy reduces aggregate emission intensity by 22.7 percent relative to the status quo.

Notably, the two counterfactual series are nearly indistinguishable. This suggests that induced innovation generates negligible additional emission reductions beyond the direct cost effect. The intuition stems from the asymmetry of induced technical change. While a higher emission cost stimulates dirty technical change ($\tilde{\omega}_Z$), it has a negligible effect on clean technical change ($\tilde{\omega}_C$) (see Table 1). As discussed, the impact of dirty technical change on emission intensity is theoretically ambiguous. The efficiency effect, which lowers intensity, is countered by the substitution effect. Our results indicate that these forces largely cancel out and neutralize the contribution of induced innovation to aggregate emission intensity decline.

7.4 Uniform Emission Cost Regimes

Our structural estimates reveal substantial heterogeneity in emission costs across prefectures. As suggested by Hsieh and Klenow (2009), such dispersion implies significant allocative inefficiency. Economic theory dictates that policies equalizing marginal emission costs, such as cap-and-trade systems, minimize the aggregate cost of emission reductions (Fowlie et al., 2012). To quantify the potential efficiency gains from removing these distortions, we simulate two counterfactual pricing regimes: a *national uniform price*, which establishes a single nationwide emission cost that equates total emissions to baseline levels; and a *provincial uniform price*, which equalizes costs within each province to match provincial baselines.

Appendix Table G.1 reports the results. Both regimes generate large gains in allocative efficiency, which are reflected in lower aggregate emission intensity. Under a national

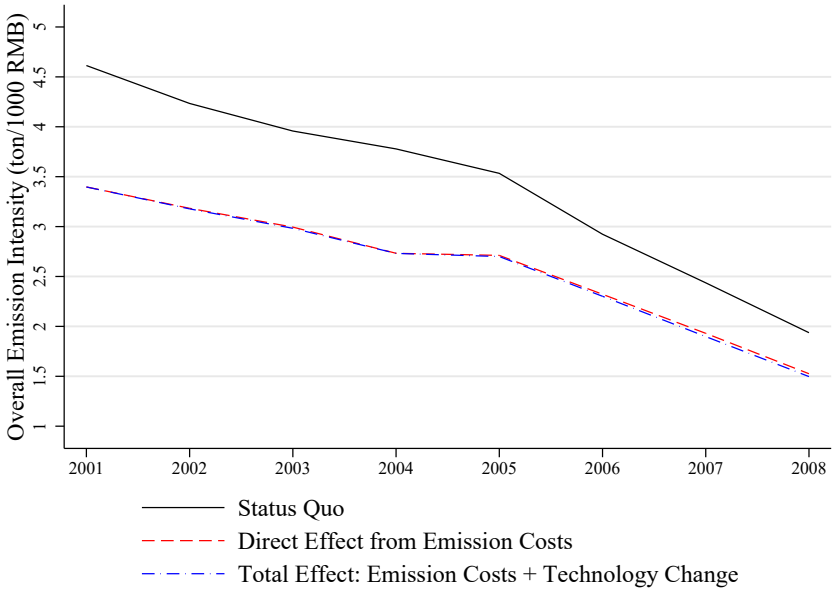


Figure 5: Impact of Doubling Emission Discharge Fee on Aggregate Emission Intensity

Notes: This figure illustrates the simulated impact of a nationwide doubling of the emission discharge fee on aggregate SO_2 emission intensity. "Status Quo" plots the observed emission intensity series. "Direct Effect" shows the counterfactual path where the emission cost (τ) is increased by 30 percent in all years, holding technologies ($\tilde{\omega}_C, \tilde{\omega}_Z$) at their baseline levels. "Total Effect" shows the full impact of the policy, allowing technology to evolve endogenously according to the estimated VAR(1) process.

uniform price, aggregate emission intensity falls by 30.8 percent relative to the baseline in 2008. This reduction stems from the reallocation of production: the uniform price relaxes constraints on firms that previously faced excessively high emission costs, allowing them to expand, while contracting firms with historically low emission costs. The provincial uniform price yields a smaller, yet still substantial, 16.0 percent reduction in emission intensity. This implies that while correcting intra-provincial misallocation delivers significant gains, a non-trivial portion of the inefficiency arises from the misalignment of emission costs across provinces.

These aggregate gains involve distinct regional distributional consequences. As illustrated in Appendix Figure G.1, although a national uniform price lowers emission intensity in most provinces, it raises it in Shandong and Jilin by 20 and 34 percent, respectively. This divergence occurs because the national emission price falls below their pre-policy local costs, inducing firms in these regions to substitute toward the now relatively cheaper

dirty inputs. In contrast, the provincial uniform price lowers emission intensity in every province. By holding provincial emissions fixed and equalizing the marginal product of emissions (MPE) within each province, this regime mechanically raises output without imposing a cross-regional redistribution of pollution.

8 Conclusion

This paper investigates the microeconomic drivers of China's industrial emission reduction by developing a structural framework that jointly recovers biased technical change and unobserved emission costs. We find that improvements in clean and dirty technology, alongside rising emission costs, drove the significant decline in emission intensity between 2001 and 2008. Validating our estimates against external data, we show that recovered emission costs correlate strongly with local discharge fees, while clean technical change tracks firm-level R&D.

We use this framework to assess the Porter Hypothesis, finding support for a robust and rapid response: induced innovation offsets approximately 40 percent of compliance costs within five years, with the majority of the adjustment occurring in the first two years. Finally, our counterfactual analysis demonstrates that while higher emission fees effectively reduce intensity through direct cost channels, the incremental contribution of induced technical change is negligible, and a uniform emission cost regime generates substantial efficiency gains at the expense of regional equity.

References

ACEMOGLU, D., P. AGHION, L. BURSZTYN, AND D. HEMOUS (2012): "The Environment and Directed Technical Change," *American Economic Review*, 102, 131–166.

ACEMOGLU, D., U. AKCIGIT, D. HANLEY, AND W. KERR (2016): "Transition to Clean Technology," *Journal of Political Economy*, 124, 52–104.

ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83, 2411–2451.

AMBEC, S., M. A. COHEN, S. ELGIE, AND P. LANOIE (2013): "The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness?" *Review of Environmental Economics and Policy*, 7, 2–22.

ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98, 1998–2031.

AW, B. Y., M. J. ROBERTS, AND D. Y. XU (2011): "R&D investment, exporting, and productivity dynamics," *American Economic Review*, 101, 1312–1344.

BAYER, P., F. FERREIRA, AND R. MCMILLAN (2007): "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 115, 588–638.

BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841.

BRANDT, L. AND K. LIM (2024): "Opening up in the 21st Century: A Quantitative Accounting of Chinese Export Growth," *Journal of International Economics*, 150, 103895.

BRANDT, L., J. VAN BIESEBROECK, AND Y. ZHANG (2012): "Creative Accounting or Creative Destruction? Firm-level Productivity Growth in Chinese Manufacturing," *Journal of Development Economics*, 97, 339–351.

BRANDT, L., J. VAN BIESEBROECK, Y. ZHANG, AND L. WANG (2023): "Where Has All the Dynamism Gone? Productivity Growth in China's Manufacturing Sector, 1998-2013," *SSRN Electronic Journal*.

BRUNEL, C. (2017): "Pollution Offshoring and Emission Reductions in EU and US Manufacturing," *Environmental and Resource Economics*, 68, 621–641.

CAI, X., Y. LU, M. WU, AND L. YU (2016): "Does Environmental Regulation Drive Away Inbound Foreign Direct Investment? Evidence from a Quasi-Natural Experiment in China," *Journal of Development Economics*, 123, 73–85.

CALEL, R. (2020): "Adopt or Innovate: Understanding Technological Responses to Cap-and-Trade," *American Economic Journal: Economic Policy*, 12, 170–201.

CALEL, R. AND A. DECHEZLEPRÊTRE (2016): "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market," *Review of Economics and Statistics*, 98, 173–191.

CHEN, Q., Z. CHEN, Z. LIU, J. C. SUÁREZ SERRATO, AND D. Y. XU (2025): "Regulating Conglomerates: Evidence from an Energy Conservation Program in China," *American Economic Review*, 115, 408–447.

CHEN, Y. J., P. LI, AND Y. LU (2018): "Career Concerns and Multitasking Local Bureaucrats: Evidence of a Target-Based Performance Evaluation System in China," *Journal of Development Economics*, 133, 84–101.

CHERNIWCHAN, J., B. R. COPELAND, AND M. S. TAYLOR (2017): "Trade and the Environment: New Methods, Measurements, and Results," *Annual Review of Economics*, 9, 59–85.

CHOI, J., J. HYUN, G. KIM, AND Z. PARK (2025): "The Cleanup of US Manufacturing through Pollution Offshoring," *Journal of International Economics*, 154, 104046.

COPELAND, B. R. AND M. S. TAYLOR (2003): "Trade and the Environment: Theory and Evidence," in *Trade and the Environment*, Princeton University Press.

DE LOECKER, J. (2011): "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity," *Econometrica*, 79, 1407–1451.

DEMIRER, M. (2024): "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups," .

DORASZELSKI, U. AND J. JAUMANDREU (2013): "R&D and Productivity: Estimating Endogenous Productivity," *The Review of Economic Studies*, 80, 1338–1383.

——— (2018): "Measuring the Bias of Technological Change," *Journal of Political Economy*, 126, 1027–1084.

DUAN, Y., T. JI, Y. LU, AND S. WANG (2021): "Environmental Regulations and International Trade: A Quantitative Economic Analysis of World Pollution Emissions," *Journal of Public Economics*, 203, 104521.

FAN, H., J. GRAFF ZIVIN, Z. KOU, X. LIU, AND H. WANG (2025): "Going Green in China: Firms' Responses to Stricter Environmental Regulations," *Canadian Journal of Economics/Revue canadienne d'économique*, 58, 385–410.

FORSLID, R., T. OKUBO, AND K. H. ULLTVEIT-MOE (2018): "Why Are Firms That Export Cleaner? International Trade, Abatement and Environmental Emissions," *Journal of Environmental Economics and Management*, 91, 166–183.

FOWLIE, M., S. P. HOLLAND, AND E. T. MANSUR (2012): "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program," *American Economic Review*, 102, 965–993.

GANDHI, A., S. NAVARRO, AND D. A. RIVERS (2020): "On the Identification of Gross Output Production Functions," *Journal of Political Economy*, 128, 2973–3016.

GRETHER, J.-M., N. A. MATHYS, AND J. DE MELO (2009): "Scale, Technique and Composition Effects in Manufacturing SO2 Emissions," *Environmental and Resource Economics*, 43, 257–274.

GRIECO, P. L. E., S. LI, AND H. ZHANG (2016): "Production Function Estimation with Unobserved Input Price Dispersion," *International Economic Review*, 57, 665–690.

GROSSMAN, G. M. AND A. B. KRUEGER (1991): "Environmental Impacts of a North American Free Trade Agreement," *NBER Working Paper*.

GUO, B., H. YU, AND G. JIN (2024): "Urban Green Total Factor Productivity in China: A Generalized Luenberger Productivity Indicator and Its Parametric Decomposition," *Sustainable Cities and Society*, 106, 105365.

HASSLER, J., P. KRUSELL, AND C. OLOVSSON (2021): "Directed Technical Change as a Response to Natural Resource Scarcity," *Journal of Political Economy*, 129, 3039–3072.

HE, G., S. WANG, AND B. ZHANG (2020): "Watering Down Environmental Regulation in China," *The Quarterly Journal of Economics*, 135, 2135–2185.

HOLLADAY, J. S. AND L. D. LAPLUE III (2021): "Decomposing Changes in Establishment-Level Emissions with Entry and Exit," *Canadian Journal of Economics/Revue canadienne d'économique*, 54, 1046–1071.

HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India*," *The Quarterly Journal of Economics*, 124, 1403–1448.

IMBENS, G. W. AND W. K. NEWEY (2009): "Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity," *Econometrica*, 77, 1481–1512.

JO, A. (2024): "Substitution Between Clean and Dirty Energy with Biased Technical Change," *International Economic Review*, n/a.

KWON, O., H. ZHAO, AND M. Q. ZHAO (2021): "The Reallocation Effect of Emissions Cap-and-Trade: Evidence from China," *School of Economics Working Paper Series*.

——— (2023): "Global Firms and Emissions: Investigating the Dual Channels of Emissions Abatement," *Journal of Environmental Economics and Management*, 118, 102772.

LEE, C.-C., Z.-W. HE, AND Z. YUAN (2023): "A Pathway to Sustainable Development: Digitization and Green Productivity," *Energy Economics*, 124, 106772.

LEVINSOHN, J. AND A. PETRIN (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70, 317–341.

LEVINSON, A. (2009): "Technology, International Trade, and Pollution from US Manufacturing," *American Economic Review*, 99, 2177–2192.

——— (2015): "A Direct Estimate of the Technique Effect: Changes in the Pollution Intensity of US Manufacturing, 1990–2008," *Journal of the Association of Environmental and Resource Economists*, 2, 43–56.

——— (2023): "Are Developed Countries Outsourcing Pollution?" *Journal of Economic Perspectives*, 37, 87–110.

LU, Y. AND J. PLESS (2026): "Greening to Grow: Evidence from Environmental Regulation and Industrial Firm Productivity in China," *Journal of Public Economics*, 256, 105596.

LYUBICH, E., J. SHAPIRO, AND R. WALKER (2018): "Regulating Mismeasured Pollution: Implications of Firm Heterogeneity for Environmental Policy," *AEA Papers and Proceedings*, 108, 136–142.

NAJJAR, N. AND J. CHERNIWCHAN (2021): "Environmental Regulations and the Cleanup of Manufacturing: Plant-Level Evidence," *The Review of Economics and Statistics*, 103, 476–491.

OLLEY, G. S. AND A. PAKES (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297.

POPP, D. (2002): "Induced Innovation and Energy Prices," *American Economic Review*, 92, 160–180.

PORTER, M. E. AND C. VAN DER LINDE (1995): "Toward a New Conception of the Environment-Competitiveness Relationship," *Journal of Economic Perspectives*, 9, 97–118.

RAVAL, D. R. (2019): "The Micro Elasticity of Substitution and Non-Neutral Technology," *The RAND Journal of Economics*, 50, 147–167.

RODRIGUE, J., D. SHENG, AND Y. TAN (2022): "Exporting, Abatement, and Firm-Level Emissions: Evidence from China's Accession to the WTO," *The Review of Economics and Statistics*, 1–45.

SHAPIRO, J. S. AND R. WALKER (2018): "Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade," *American Economic Review*, 108, 3814–3854.

SHORROCKS, A. F. (2012): "Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value," *The Journal of Economic Inequality*, 11, 99–126.

TANAKA, S. (2015): "Environmental Regulations on Air Pollution in China and Their Impact on Infant Mortality," *Journal of Health Economics*, 42, 90–103.

WANG, C., J. WU, AND B. ZHANG (2018): "Environmental Regulation, Emissions and Productivity: Evidence from Chinese COD-emitting Manufacturers," *Journal of Environmental Economics and Management*, 92, 54–73.

WANG, H. AND D. WHEELER (2005): "Financial Incentives and Endogenous Enforcement in China's Pollution Levy System," *Journal of Environmental Economics and Management*, 49, 174–196.

WANG, Y., S. C. KUMBHAKAR, AND S. ZHAO (2025): "Dragon's Roar: Unveiling Pollution Abatement Costs in Chinese Manufacturing Sector Using a By-production Approach," .

WANG, Z. (2025): "Non-Neutral Technological Change in Chinese Manufacturing," .

WATABE, Y., R. SOGALLA, AND J. WANNER (2024): "New Trade Models, Same Old Emissions?" *IDE Discussion Papers*.

YANG, D. T., V. W. CHEN, AND R. MONARCH (2010): "RISING WAGES: HAS CHINA LOST ITS GLOBAL LABOR ADVANTAGE?" *Pacific Economic Review*, 15, 482–504.

ZHANG, H. (2019): "Non-Neutral Technology, Firm Heterogeneity, and Labor Demand," *Journal of Development Economics*, 140, 145–168.

Online Appendix

An Anatomy of China's Industrial Cleanup: Biased Technical Change and Emission Cost

by

Ohyun Kwon

Hongsong Zhang

Hao Zhao

Min Qiang Zhao

A Microfoundation for CES Production Function

Consider a firm that jointly optimizes its *production* and *abatement* decisions. The firm chooses the amount of input L and abatement effort π to maximize profits. The firm's net production function is given by

$$Q = \tilde{\omega}_H L \cdot \underbrace{\left[(1-\beta) + \beta \left((1-\pi)^{\frac{1}{\alpha}}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}}_{C(\pi)}, \tag{A.1}$$

where $\tilde{\omega}_H L$ represents gross output and $C(\pi)$ is the *abatement cost* function with respect to the abatement effort π. β is for normalization, σ captures the non-linear response of output with respect to abatement effort, and the role of α is explained in equation (A.2). As the firm increases the abatement share π, emissions Z decrease according to:

$$Z = \tilde{\omega}_H L \cdot \underbrace{\tilde{\omega}_Z^{-1}(1-\pi)^{\frac{1}{\alpha}}}_{E(\pi)}, \tag{A.2}$$

where $E(\pi)$ is the *abatement effectiveness* function in Copeland and Taylor (2003) (CT) except for the additional term $\tilde{\omega}_Z$. Intuitively, $\tilde{\omega}_Z$ can be interpreted as the *abatement efficiency* that is inversely related to the amount of emissions that the firm makes at a given level of gross output and abatement effort and the parameter α governs the sensitivity of

emissions to abatement effort. A lower α implies that emissions are more responsive to changes in abatement effort.[29]

The firm's incentive to lower emissions comes from the emission cost, τ, which enters the firm's profit maximization problem as a cost of emissions. The firm's profit maximization problem can be written as:

$$\max_{L,\pi} \; P(Q)Q - WL - \tau Z, \tag{A.3}$$

where $P(Q)$ is the output price and W is the wage rate. To obtain the CES production function a la equation (6) in the main text, one only needs to make minimal assumptions on the inverse demand function $P(Q)$, only requiring it to be differentiable and downward sloping. To obtain a solution to the optimal π^*, however, we assume $P(Q) \equiv \Phi^{-\frac{1}{\eta}} Q^{\frac{1}{\eta}}$ $(\eta < -1)$. Solving for the optimal level of π^* yields:

$$(1-\pi)^{\frac{1}{\alpha\sigma}} = \frac{\beta}{1-\beta} \frac{\tilde{\omega}_Z}{\tilde{\omega}_H} \frac{W}{\tau}. \tag{A.4}$$

By substituting equation (A.2) into equation (A.1), we can express the production function directly in terms of labor and emissions:

$$Q = \tilde{\omega}_H (1-\beta)^{\frac{\sigma}{\sigma-1}} \left[L^{\frac{\sigma-1}{\sigma}} + \frac{\beta}{(1-\beta)} \left(\frac{\tilde{\omega}_Z}{\tilde{\omega}_H} \right)^{\frac{\sigma-1}{\sigma}} Z^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}. \tag{A.5}$$

Defining the technology parameters

$$\omega_H \equiv \tilde{\omega}_H (1-\beta)^{\frac{\sigma}{\sigma-1}} \tag{A.6}$$

$$\omega_Z \equiv \left(\frac{\beta}{1-\beta} \right)^{\frac{\sigma}{\sigma-1}} \frac{\tilde{\omega}_Z}{\tilde{\omega}_H}, \tag{A.7}$$

we recover the CES production function presented in equation (6) of the main text. As

[29] We can rewrite both loss function (in equation (A.1)) and abatement function (in equation (A.2)) equivalently by redefining $\Theta \equiv (1-\pi)^{\frac{1}{\alpha}}$, but we refrain from doing this simply to maintain comparability of the role of parameter α in CT framework and ours.

equations (A.6) and (A.7) suggest, ω_H is associated with the neutral technology level that scales gross output, while ω_Z is associated with the firm's relative abatement efficiency.

Our system of equations (A.1) and (A.2) closely mimic the framework of Copeland and Taylor (2003) (CT). In the CT model, production is defined as $Q = \tilde{\omega}_H L(1-\pi)$, and by-product emission $Z = \tilde{\omega}_H L(1-\pi)^{\frac{1}{\alpha}}$. Our input-equivalent production function follows a CES form in equation (6), whereas the CT framework assumes a Cobb-Douglas structure $Q = Z^{\alpha}(\tilde{\omega}_H L)^{1-\alpha}$. To demonstrate that our framework generalizes the CT model, we show in Section A.1 that our loss function equation (A.1) can (up to second-order approximation) encapsulate the CT loss function, $(1-\pi)$, as a special case. Specifically, proofs in Section A.1 show that the second-order Taylor expansion of equation (A.1) around $\pi = 0$ yields:

$$Q(\pi) \approx \tilde{\omega}_H L \left[1 - c\pi + \left(\frac{c^2}{\sigma}\left(1 + \beta(\sigma - 1)\right) - c\right)\pi^2\right], \tag{A.8}$$

where $c \equiv \beta/\alpha$. With additional β and c, our framework not only allows the output loss function to have a negative linear term, but its coefficient is not restricted to be -1, allowing for the possibility that the loss function may feature nonlinearity with respect to abatement effort π. Due to this generality, our resulting production function is CES, as is well known, a general form of Cobb-Douglas, which is the resulting input-equivalent production function in CT.

A.1 Generalization of Copeland and Taylor (2003) Framework

We demonstrate that our specification in equation (A.1) encompasses the CT framework as a special case while enabling a more flexible CES structure.

To analyze the properties of this generalization, we define $Q(\pi)$ as the expression in equation (A.1) and perform a second-order Taylor expansion around $\pi = 0$.

First, the first term is simply the potential gross output:

$$Q(0) = \tilde{\omega}_H L. \tag{A.9}$$

Second, the first derivative with respect to π is given by:

$$Q'(\pi) = \frac{dQ}{d\pi} = -\frac{\beta\tilde{\omega}_H L}{\alpha}(1-\pi)^{\frac{\sigma-1-\alpha\sigma}{\alpha\sigma}}\left[(1-\beta)+\beta(1-\pi)^{\frac{\sigma-1}{\alpha\sigma}}\right]^{\frac{1}{\sigma-1}}. \tag{A.10}$$

Evaluating this at $\pi = 0$ yields:

$$Q'(0) = -\frac{\beta\tilde{\omega}_H L}{\alpha}. \tag{A.11}$$

Third, the second derivative is:

$$\begin{aligned} Q''(\pi) = \frac{\partial^2 Q}{\partial \pi^2} = \frac{\beta\tilde{\omega}_H L}{\alpha}\Bigg\{ &\left(\frac{\sigma-1-\alpha\sigma}{\alpha\sigma}\right)(1-\pi)^{\frac{\sigma-1}{\alpha\sigma}-2}\left[(1-\beta)+\beta(1-\pi)^{\frac{\sigma-1}{\alpha\sigma}}\right]^{\frac{1}{\sigma-1}} \\ &+\frac{\beta}{\alpha\sigma}(1-\pi)^{\frac{2(\sigma-1)}{\alpha\sigma}-2}\left[(1-\beta)+\beta(1-\pi)^{\frac{\sigma-1}{\alpha\sigma}}\right]^{\frac{1}{\sigma-1}-1}\Bigg\}. \end{aligned} \tag{A.12}$$

Simplifying and evaluating at $\pi = 0$, we obtain:

$$Q''(0) = \frac{\beta\tilde{\omega}_H L(\sigma-1-\alpha\sigma+\beta)}{\alpha^2\sigma}. \tag{A.13}$$

Combining these terms, we obtain the second-order expansion of equation (A.1):

$$Q(\pi) \approx \tilde{\omega}_H L\left[1-\frac{\beta}{\alpha}\pi+\frac{\beta(\sigma-1-\alpha\sigma+\beta)}{\alpha^2\sigma}\pi^2\right]. \tag{A.14}$$

This expansion demonstrates that our framework is a natural generalization of the Copeland and Taylor (2003) model. Specifically, equation (A.14) relaxes two key constraints of the CT framework: it allows the slope of Q with respect to abatement effort ($dQ/d\pi = -\tilde{\omega}_H L$ in CT) to vary, and it permits the production function to exhibit nonlinear curvature (concavity or convexity) with respect to abatement intensity.

To illustrate this flexibility, we define the parameter $c \equiv \beta/\alpha$. Substituting this into equation (A.14) yields:

$$Q(\pi) \approx \tilde{\omega}_H L\left[1-c\pi+\left(\frac{c^2}{\sigma}(1+\beta(\sigma-1))-c\right)\pi^2\right]. \tag{A.15}$$

This reparameterization highlights two distinct advantages. First, the first-order coefficient c can be set to any value through the appropriate choice of α and β (the standard CT framework corresponds to $c = 1$). Second, the second-order term allows for potential curvature in the abatement function. Crucially, this curvature is not a constraint; it is possible to recover the linear case (zero curvature) by selecting β values that, for any given linear coefficient c, make the coefficient of the quadratic term to become zero. Thus, our model fully encapsulates the CT framework as a special case while providing the necessary microfoundation for the flexible CES production function employed in our main analysis.

A.2 Efficiency and Substitution Effects of Technology Changes

Next, we explore the role of ω_Z in influencing the firm's emissions via two channels: efficiency and substitution effects. Using equation (6), firm-level emission intensity can be expressed as

$$\zeta \equiv \frac{Z}{Q} = \omega_H^{-1} \left(\ell^{\frac{\sigma-1}{\sigma}} + \omega_Z^{\frac{\sigma-1}{\sigma}} \right)^{-\frac{\sigma}{\sigma-1}},$$

where we define $\ell \equiv L/Z$ as the labor-to-emission ratio. Now, let "$\hat{}$" denote the percentage change operator (i.e., $\hat{x} \equiv d \ln x$). We can express the change in emission intensity with respect to changes in ω_Z as

$$\hat{\zeta} = -(\underbrace{(1-\alpha_\ell)\,\hat{\omega}_Z}_{\text{efficiency effect}} + \underbrace{\alpha_\ell \hat{\ell}}_{\text{substitution effect}}), \tag{A.16}$$

where $\alpha_\ell \equiv \ell^{\frac{\sigma-1}{\sigma}} / \left(\ell^{\frac{\sigma-1}{\sigma}} + \omega_Z^{\frac{\sigma-1}{\sigma}} \right)$. Equation (A.16) embodies our definition of *efficiency* and *substitution* effects. The efficiency effect captures the direct impact of a change in dirty technology on emission intensity, holding input ratios constant. In contrast, the substitution effect captures the indirect impact through changes in input ratios induced by the technology change. Clearly, equation (A.16) shows that an increase in ω_Z leads to a decrease in emission intensity via the efficiency effect. An increase in labor-to-emission

ratio ℓ also has a similar effect. To explore the total effect of ω_Z, we divide both sides of equation (A.16) by $\hat{\omega}_Z$ to obtain

$$\begin{aligned} \frac{\hat{\zeta}}{\hat{\omega}_Z} &= -\left((1-\alpha_\ell) + \alpha_\ell \hat{\ell}/\hat{\omega}_Z\right) \\ &= -\left(1-\alpha_\ell\sigma\right), \end{aligned} \tag{A.17}$$

where $\hat{\ell}/\hat{\omega}_Z$ captures the elasticity of labor-to-emission ratio with respect to ω_Z, which by equation (7) is given by $1-\sigma$. Equation (A.17) shows that the total effect of ω_Z on emission intensity depends critically on the substitution elasticity σ and the input ratio weight α_ℓ. When $\sigma > 1/\alpha_\ell$, the total effect becomes positive, indicating that the substitution effect dominates the efficiency effect. In contrast, when $\sigma < 1/\alpha_\ell$, the efficiency effect prevails, leading to a negative total effect. This decomposition provides a clear framework for understanding how changes in dirty technology influence emission intensity through both direct and indirect channels.

B Gross Output Demand and Value-Added Equivalence

This appendix demonstrates how the value-added demand function used in our main theoretical model naturally derives from a gross output demand equation.

Suppose consumers demand the firm's gross output $Q_{Output,jt}$ according to a constant elasticity demand function.

$$Q_{Output,jt} = \tilde{\Phi}_{jt}(P_{Output,jt})^{\eta} \tag{B.1}$$

Here $P_{Output,jt}$ is the gross output price, $\eta < -1$ is the price elasticity of demand, and $\tilde{\Phi}_{jt}$ is a fundamental demand shifter.

Following standard practice in the structural literature (Ackerberg et al., 2015), we assume gross output is a Leontief aggregate of the value-added composite Q_{jt} and intermediate materials M_{jt}. Specifically, we assume $Q_{Output,jt} = \min\{Q_{jt}, M_{jt}/c\}$, where c is the material requirement per unit of output. Profit maximization implies no input waste, yielding the strict equalities $Q_{Output,jt} = Q_{jt}$ and $M_{jt} = cQ_{jt}$.

The implicit price of value added, P_{jt}^{VA}, is defined as the residual revenue left per unit of output after paying for intermediate materials. Let $P_{M,jt}$ denote the price of materials. The value-added price is $P_{jt}^{VA} = P_{Output,jt} - cP_{M,jt}$. To operationalize this, we assume that the material cost share is constant within a given industry. This implies that the value-added price is strictly proportional to the gross output price, such that $P_{jt}^{VA} = \lambda_{\text{ind}(j)} P_{Output,jt}$, where $\lambda_{\text{ind}(j)}$ is the industry-specific value-added share of gross revenue.

We can now derive the demand curve for the value-added composite. Because $Q_{Output,jt} = Q_{jt}$ and $P_{Output,jt} = P_{jt}^{VA}/\lambda_{\text{ind}(j)}$, we substitute these expressions directly into the gross output demand curve in equation (B.1).

$$Q_{jt} = \tilde{\Phi}_{jt}\left(\frac{P_{jt}^{VA}}{\lambda_{\text{ind}(j)}}\right)^{\eta} = \left(\tilde{\Phi}_{jt}\lambda_{\text{ind}(j)}^{-\eta}\right)(P_{jt}^{VA})^{\eta} \tag{B.2}$$

By defining the value-added demand shifter as $\Phi_{jt} = \tilde{\Phi}_{jt}\lambda_{\text{ind}(j)}^{-\eta}$, we recover the exact value-added demand equation presented in our main text.

$$Q_{jt} = \Phi_{jt}(P_{jt}^{VA})^{\eta} \tag{B.3}$$

C Identification with General Production Technology

In the main text, we assume a CES production function to derive closed-form analytical solutions for the unobserved shadow cost τ_{pt} and the relative technology ω_{Zjt}. Here, we demonstrate that our identification strategy that leverages the first-order conditions and the labor share of revenue holds for a general class of production functions, provided they satisfy standard regularity conditions. We first present the proof under a constant markup assumption and then show that identification holds when markup features some flexible structure, e.g., the endogenous markup structure in Atkeson and Burstein (2008).

C.1 Baseline Case: Constant Markup

Model Setup. Consider a general production function for firm j at time t:

$$Q_{jt} = \exp(\omega_{Hjt}) F\left(K_{jt}, L_{jt}, \tilde{Z}_{jt}\right) \tag{C.1}$$

where $\tilde{Z}_{jt} = \exp(\omega_{Zjt}) Z_{jt}$ represents the effective units of the dirty input (emissions). Here, ω_{Hjt} represents Hicks-neutral technology, and ω_{Zjt} represents dirty technology. We assume that $F(\cdot)$ is twice continuously differentiable and exhibits standard neoclassical properties (positive marginal products and diminishing marginal returns).

The firm operates under imperfect competition with a constant elasticity of demand $\eta < -1$, implying a constant markup of $\mu = \frac{\eta}{1+\eta}$. The firm chooses labor (L_{jt}) and emissions (Z_{jt}) to maximize static profits:

$$\max_{L_{jt}, Z_{jt}} P(Q_{jt}) Q_{jt} - W_{jt} L_{jt} - \tau_{pt} Z_{jt} \tag{C.2}$$

where revenue is $R(Q) = P(Q)Q = \Phi_{jt}^{-1/\eta} Q_{jt}^{1+1/\eta}$.

Identification of ω_{Zjt}. The first-order condition with respect to labor equates the marginal revenue product of labor to the wage:

$$\frac{P_{jt}}{\mu}\frac{\partial Q_{jt}}{\partial L_{jt}} = W_{jt} \tag{C.3}$$

Substituting the derivative $\frac{\partial Q_{jt}}{\partial L_{jt}} = \exp(\omega_{Hjt})F_L(K, L, \tilde{Z})$, we can rewrite this in terms of the "true" labor share of revenue, $S^*_{Ljt} \equiv \frac{W_{jt}L_{jt}}{P_{jt}Q_{jt}}$:

$$S^*_{Ljt} = \frac{1}{\mu}\left(\frac{\partial Q_{jt}}{\partial L_{jt}}\frac{L_{jt}}{Q_{jt}}\right) = \frac{1}{\mu}\theta_L(K_{jt}, L_{jt}, \tilde{Z}_{jt}) \tag{C.4}$$

where $\theta_L(\cdot) \equiv \frac{F_L \cdot L}{F}$ is the output elasticity of labor. Note that θ_L generally depends on the effective dirty input $\tilde{Z}_{jt}$.

Condition 1 (Invertibility): Identification requires that the output elasticity of labor, $\theta_L(K, L, \tilde{Z})$, is strictly monotonic in $\tilde{Z}$. This condition implies that the dirty input is not separable from labor in a way that keeps labor's output elasticity constant (as would be the case in a Cobb-Douglas specification where $\sigma = 1$).

Assuming this condition holds (i.e., $\frac{\partial \theta_L}{\partial \tilde{Z}} \neq 0$), we can invert equation (C.4) to solve for the effective dirty input:

$$\tilde{Z}_{jt} = \theta_L^{-1}\left(\mu S^*_{Ljt}, K_{jt}, L_{jt}\right) \tag{C.5}$$

Since Z_{jt} is observed, we uniquely identify the dirty technology as:

$$\omega_{Zjt} = \ln(\tilde{Z}_{jt}) - \ln(Z_{jt}) \tag{C.6}$$

Identification of τ_{pt}. The first-order condition with respect to emissions is:

$$\frac{P_{jt}}{\mu}\frac{\partial Q_{jt}}{\partial Z_{jt}} = \tau_{pt} \tag{C.7}$$

Using the chain rule, $\frac{\partial Q_{jt}}{\partial Z_{jt}} = \exp(\omega_{Hjt})F_{\tilde{Z}}\frac{\partial \tilde{Z}}{\partial Z} = \exp(\omega_{Hjt})F_{\tilde{Z}}\exp(\omega_{Zjt})$. Taking the ratio of the FOCs for emissions and labor eliminates the unobserved price P_{jt}/μ and the Hicks-

neutral term ω_{Hjt}:

$$\frac{\tau_{pt}}{W_{jt}} = \frac{\partial Q/\partial Z}{\partial Q/\partial L} = \frac{F_{\tilde{Z}}(K, L, \tilde{Z}) \exp(\omega_{Zjt})}{F_L(K, L, \tilde{Z})} \tag{C.8}$$

Rearranging for τ_{pt}:

$$\tau_{pt} = W_{jt} \left[\frac{F_{\tilde{Z}}(K_{jt}, L_{jt}, \tilde{Z}_{jt})}{F_L(K_{jt}, L_{jt}, \tilde{Z}_{jt})} \right] \exp(\omega_{Zjt}) \tag{C.9}$$

Since $\tilde{Z}_{jt}$ and ω_{Zjt} are identified in the previous step, and W_{jt}, K_{jt}, L_{jt} are observed, the term in the brackets (the Marginal Rate of Technical Substitution) is fully determined. Thus, τ_{pt} is uniquely identified.

C.2 Extension: Endogenous Markups (Atkeson-Burstein Structure)

We now extend the identification strategy to allow for firm-specific, time-varying markups by adopting the oligopolistic competition framework of Atkeson and Burstein (2008). In this setting, firms compete in quantities (Cournot) or prices (Bertrand) within sectors, generating markups that vary endogenously with the firm's market share.

Market Structure. Assume consumer utility is derived from a nested CES aggregator. There is a continuum of sectors, and within each sector, a finite number of firms compete. Let $\rho > 1$ be the elasticity of substitution between varieties within a sector, and $\eta > 1$ be the elasticity of substitution across sectors.

Under this structure, firm j's perceived elasticity of demand, ε_{jt}, is a weighted average of the within-sector and cross-sector elasticities, weighted by the firm's market share s_{jt} (defined as the firm's share of sectoral revenue):

$$\varepsilon_{jt}(s_{jt}) = \rho(1 - s_{jt}) + \eta s_{jt} \tag{C.10}$$

The firm's profit-maximizing markup μ_{jt} is explicitly determined by its observable market share:

$$\mu_{jt} = \frac{\varepsilon_{jt}(s_{jt})}{\varepsilon_{jt}(s_{jt}) - 1} \tag{C.11}$$

Crucially, because market share s_{jt} is observable, the markup μ_{jt} is structurally identified conditional on the demand parameters (ρ, η).

Identification. We show that the unobserved shadow cost and dirty technology remain identified by distinguishing the markup effect (driven by market share) from the technology effect (driven by input mix).

Recall the first-order condition for labor in terms of revenue shares from equation (C.4), but now with a firm-specific markup:

$$S^*_{Ljt} = \frac{1}{\mu_{jt}} \theta_L(K_{jt}, L_{jt}, \tilde{Z}_{jt}) \tag{C.12}$$

Substituting the structural markup from equation (C.11), we obtain:

$$S^*_{Ljt} \cdot \left(\frac{\rho(1 - s_{jt}) + \eta s_{jt}}{\rho(1 - s_{jt}) + \eta s_{jt} - 1} \right) = \theta_L(K_{jt}, L_{jt}, \tilde{Z}_{jt}) \tag{C.13}$$

The left-hand side of this equation depends only on observables (S^*_{Ljt}, s_{jt}) and demand parameters. Let $\tilde{S}_{Ljt}$ denote this "markup-adjusted" labor share. The identification problem reduces to inverting the relationship:

$$\tilde{S}_{Ljt} = \theta_L(K_{jt}, L_{jt}, \tilde{Z}_{jt}) \tag{C.14}$$

As established in the baseline proof, provided Condition 1 holds (i.e., $\frac{\partial \theta_L}{\partial \tilde{Z}} \neq 0$), the function $\theta_L(\cdot)$ is invertible.

Thus, the identification strategy relies on two distinct sources of data variation:

1. Variation in market share (s_{jt}) identifies the markup μ_{jt}.
2. Variation in the markup-adjusted labor share (conditional on K, L) identifies the effective dirty input $\tilde{Z}_{jt}$ and, subsequently, the dirty technology ω_{Zjt}.

Once ω_{Zjt} is recovered, the identification of the shadow cost τ_{pt} follows exactly the same logic as in the constant markup case, since the markup term cancels out when taking the

ratio of the first-order conditions.

D Additional Deviations and Proofs

D.1 Derivation of the Estimation Equation

This appendix details the derivation of the structural estimation equation (equation (15) in the main text). We proceed in three steps: (1) applying a logarithmic transformation to the CES production function, (2) using the first-order conditions to substitute out the unobserved dirty technology ω_{Zjt}, and (3) using the labor share of revenue to substitute out the unobserved emission cost τ_{pt}.

Step 1: Logarithmic Transformation of the Production Function

Recall the re-parameterized CES production function:

$$Q_{jt} = \exp(\omega_{Hjt}) \left[(X_{jt})^{\frac{\sigma-1}{\sigma}} + (\exp(\omega_{Zjt}) Z_{jt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\nu\sigma}{\sigma-1}}, \tag{D.1}$$

where $X_{jt} = K_{jt}^{\theta} L_{jt}^{1-\theta}$ is the composite clean input. We factor out the clean input term $(X_{jt})^{\frac{\sigma-1}{\sigma}}$ from inside the bracket:

$$Q_{jt} = \exp(\omega_{Hjt}) X_{jt}^{\nu} \left[1 + \left(\frac{\exp(\omega_{Zjt}) Z_{jt}}{X_{jt}} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\nu\sigma}{\sigma-1}}. \tag{D.2}$$

Taking natural logarithms of both sides yields:

$$\ln Q_{jt} = \omega_{Hjt} + \nu \ln X_{jt} + \frac{\nu\sigma}{\sigma-1} \ln \left[1 + \left(\frac{\exp(\omega_{Zjt}) Z_{jt}}{X_{jt}} \right)^{\frac{\sigma-1}{\sigma}} \right]. \tag{D.3}$$

Step 2: Substituting Dirty Technology ω_{Zjt}

We eliminate the unobserved dirty technology term by exploiting the firm's first-order conditions. The firm equates the ratio of marginal products (the marginal rate of technical substitution) to the ratio of factor prices.

First, recall the clean input composite $X_{jt} = K_{jt}^{\theta} L_{jt}^{1-\theta}$. The marginal product of the

composite input X and the dirty input Z are derived from the production function:

$$\frac{\partial Q_{jt}}{\partial X_{jt}} = Q_{jt}^{1-\frac{\sigma-1}{\nu\sigma}} \exp(\omega_{Hjt})^{\frac{\sigma-1}{\nu\sigma}} (X_{jt})^{-\frac{1}{\sigma}}, \tag{D.4}$$

$$\frac{\partial Q_{jt}}{\partial Z_{jt}} = Q_{jt}^{1-\frac{\sigma-1}{\nu\sigma}} \exp(\omega_{Hjt})^{\frac{\sigma-1}{\nu\sigma}} (Z_{jt})^{-\frac{1}{\sigma}} \exp(\omega_{Zjt})^{\frac{\sigma-1}{\sigma}}. \tag{D.5}$$

The marginal product of Labor (L) is found via the chain rule: $\frac{\partial Q}{\partial L} = \frac{\partial Q}{\partial X}\frac{\partial X}{\partial L}$, where $\frac{\partial X}{\partial L} = (1-\theta)\frac{X_{jt}}{L_{jt}}$.

Taking the ratio of the first-order condition for emissions ($\tau_{pt} = P_Q \frac{\partial Q}{\partial Z}$) to the first-order condition for labor ($W_{jt} = P_Q \frac{\partial Q}{\partial L}$), we obtain:

$$\frac{\tau_{pt}}{W_{jt}} = \frac{\partial Q/\partial Z}{\partial Q/\partial L} = \frac{(Z_{jt})^{-\frac{1}{\sigma}} \exp(\omega_{Zjt})^{\frac{\sigma-1}{\sigma}}}{(X_{jt})^{-\frac{1}{\sigma}} (1-\theta)\frac{X_{jt}}{L_{jt}}}. \tag{D.6}$$

Rearranging terms to isolate the unobserved technology component yields:

$$\left(\frac{\exp(\omega_{Zjt})Z_{jt}}{X_{jt}}\right)^{\frac{\sigma-1}{\sigma}} = (1-\theta)\frac{\tau_{pt}}{W_{jt}}\frac{Z_{jt}}{L_{jt}}. \tag{D.7}$$

We substitute the right-hand side of equation (D.7) into the bracket of the logarithmically transformed production function (equation (D.3)):

$$\ln Q_{jt} = \omega_{Hjt} + \nu \ln X_{jt} + \frac{\nu\sigma}{\sigma-1} \ln\left[1 + (1-\theta)\frac{\tau_{pt}Z_{jt}}{W_{jt}L_{jt}}\right]. \tag{D.8}$$

Finally, we expand the clean input term $\nu \ln X_{jt}$ to match the functional form used in the main text. Since $X_{jt} = K_{jt}^{\theta} L_{jt}^{1-\theta}$:

$$\nu \ln X_{jt} = \nu\theta \ln\left(\frac{K_{jt}}{L_{jt}}\right) + \nu \ln L_{jt}. \tag{D.9}$$

Substituting this back yields the intermediate estimation equation:

$$\ln Q_{jt} = \omega_{Hjt} + \nu\theta \ln\left(\frac{K_{jt}}{L_{jt}}\right) + \nu \ln L_{jt} + \frac{\nu\sigma}{\sigma-1} \ln\left[1 + (1-\theta)\frac{\tau_{pt}Z_{jt}}{W_{jt}L_{jt}}\right]. \tag{D.10}$$

Step 3: Substituting Emission Cost τ_{pt}

The equation above still contains the unobserved emission cost τ_{pt}. We recover τ_{pt} using the labor share of revenue. Recall from equation (14) in the main text that the relationship between the shadow cost and the true labor share S^*_{Ljt} is given by:

$$\tau_{pt} = \frac{W_{jt}L_{jt}}{Z_{jt}} \left(\nu \frac{1+\eta}{\eta} \frac{1}{S^*_{Ljt}} - \frac{1}{1-\theta} \right). \tag{D.11}$$

We substitute this expression for τ_{pt} into the bracketed term of equation (D.10). The term inside the logarithm becomes:

$$\text{Term} = 1 + (1-\theta) \frac{Z_{jt}}{W_{jt}L_{jt}} \left[\frac{W_{jt}L_{jt}}{Z_{jt}} \left(\nu \frac{1+\eta}{\eta} \frac{1}{S^*_{Ljt}} - \frac{1}{1-\theta} \right) \right] \tag{D.12}$$

$$= \nu \frac{1+\eta}{\eta} \frac{1-\theta}{S^*_{Ljt}}. \tag{D.13}$$

Finally, we substitute this simplified term back into equation (D.10) and add the output shock e_{jt} (since $\ln Y_{jt} = \ln Q_{jt} + e_{jt}$):

$$\ln Y_{jt} = \omega_{Hjt} + \nu\theta \ln \left(\frac{K_{jt}}{L_{jt}} \right) + \nu \ln L_{jt} + \frac{\nu\sigma}{\sigma - 1} \ln \left(\nu \frac{1+\eta}{\eta} \frac{1-\theta}{S^*_{Ljt}} \right) + e_{jt}. \tag{D.14}$$

Expanding the logarithm of the ratio yields the final estimation equation:

$$\ln Y_{jt} = \omega_{Hjt} + \nu\theta \ln \left(\frac{K_{jt}}{L_{jt}} \right) + \nu \ln L_{jt} + \frac{\nu\sigma}{\sigma - 1} \ln \left(\nu \frac{(1+\eta)(1-\theta)}{\eta} \right) - \frac{\nu\sigma}{\sigma - 1} \ln S^*_{Ljt} + e_{jt}. \tag{D.15}$$

This matches equation (15) in the main text.

D.2 Monotonicity Condition of $\tilde{\phi}(\cdot, e_{jt})$

Proposition: When $\nu > \frac{\sigma-1}{\sigma} > 0$, $\tilde{\phi}(\cdot, e_{jt})$ is strictly decreasing in e_{jt}.

Proof:

The condition $\frac{\sigma-1}{\sigma} > 0$ implies that $\sigma > 1$.

From equation (15), we have

$\frac{\partial\tilde{\phi}(\cdot,e_{jt})}{\partial e_{jt}} = \frac{\partial\omega_{Hjt}}{\partial e_{jt}} + 1 - \frac{\nu\sigma}{\sigma-1}$

Given the data $(k_{jt}, z_{jt}, l_{jt}, W_{jt}, S_{Ljt})$, the investment function $i_{jt} = i(\omega_{Hjt}, \omega_{Zjt}, k_{jt}, W_{jt}, \tau_t)$ defines an implicit function of ω_{Hjt} with respect to e_{jt}. By the implicit function theorem, we get

$\frac{\partial\omega_{Hjt}}{\partial e_{jt}} = -\frac{\partial i_{jt}/\partial e_{jt}}{\partial i_{jt}/\partial\omega_{Hjt}} = -\frac{1}{\partial i_{jt}/\partial\omega_{Hjt}}\left(\frac{\partial i_{jt}}{\partial\omega_{Zjt}}\frac{\partial\omega_{Zjt}}{\partial e_{jt}} + \frac{\partial i_{jt}}{\partial\tau_t}\frac{\partial\tau_t}{\partial e_{jt}}\right)$

Generally, investment increases strictly with Hicks-neutral technology and decreases with factor prices. That is $\partial i_{jt}/\partial\omega_{Hjt} > 0$ and $\partial i_{jt}/\partial\tau_t < 0$.

Given equation (15), ω_{Hjt} can also be expressed as:

$$\omega_{Hjt} = \phi_{jt} - \nu\theta\ln(\frac{K_{jt}}{L_{jt}}) - \nu\ln L_{jt} - \frac{\nu\sigma}{\sigma-1}\ln\left(\nu\frac{(1+\eta)(1-\theta)}{\eta}\right) + \frac{\nu\sigma}{\sigma-1}\ln S^*_{Ljt}, \quad \text{(D.16)}$$

From equation (D.16), we have

$\frac{\partial\omega_{Zjt}}{\partial e_{jt}} = \frac{\sigma}{\sigma-1}\frac{-\nu\frac{1+\eta}{\eta}\left(\frac{1}{S^*_{Ljt}}\right)^2}{\nu\frac{1+\eta}{\eta}\frac{1}{S^*_{Ljt}} - \frac{1}{1-\theta}} < 0$ since $\sigma > 1$ and $\eta < -1$.

From equation (14), we have

$\frac{\partial\tau_t}{\partial e_{jt}} = -\nu\frac{1+\eta}{\eta}\frac{W_{jt}}{Z_{jt}/L_{jt}}\left(\frac{1}{S^*_{Ljt}}\right)^2 < 0$ since $\eta < -1$.

When clean and dirty inputs are gross substitutes ($\sigma > 1$), this implies $\partial i_{jt}/\partial\omega_{Zjt} < 0$. Then, we can easily show that $\partial\omega_{Hjt}/\partial e_{jt} < 0$.

Therefore, if $1 - \frac{\nu\sigma}{\sigma-1} < 0$ and $\sigma > 1$ (that is $\nu > \frac{\sigma-1}{\sigma} > 0$), then $\tilde{\phi}(\cdot, e_{jt})$ is strictly decreasing in e_{jt}.

E Heterogeneity in Technology, Emission Cost, and Emission Intensity

This appendix documents the heterogeneity in technology, emission intensity, and emission costs across firms and regions. We compare the mean values of four key variables: clean technology ($\tilde{\omega}_C$), dirty technology ($\tilde{\omega}_Z$), emission intensity ($\ln(Z/Y)$), and

emission cost ($\ln(\tau)$). We evaluate these variables across three dimensions: (1) firm ownership (SOEs versus non-SOEs); (2) firm size (defined by capital stock, with firms above the median classified as large); and (3) geographic region (Eastern, Central, and Western).[30]

To ensure these comparisons reflect fundamental differences rather than compositional effects, we isolate the residual variation for each variable by partialling out specific fixed effects before calculating group means and 95% confidence intervals. For firm-level comparisons (ownership and size) of clean technology, dirty technology, and emission intensity, we regress each variable on year, four-digit industry, and prefecture fixed effects to isolate within-location-industry variation. For regional comparisons of these same three variables, we absorb year and industry fixed effects to control for macroeconomic shocks and industrial composition. Finally, because emission costs vary strictly at the prefecture-year level, we absorb only year fixed effects for regional evaluations. We plot the residualized values in Figure E.1.

Based on Welch's t-tests with unequal variances, Figure E.1 reveals several striking patterns regarding the distribution of technology and environmental performance across the manufacturing sector. In terms of clean technology, non-SOEs exhibit significantly higher capability levels compared to SOEs. Similarly, large firms possess significantly higher clean technology capabilities than small firms. Regionally, the economically developed Eastern region leads significantly in clean technical efficiency, followed by the Western region, with the Central region lagging behind both. When examining dirty technology, non-SOEs again demonstrate higher baseline levels than SOEs. Interestingly, the difference in dirty technology between large and small firms is negligible and not statistically significant ($p = 0.17$). Regionally, the Central region relies on dirty technology significantly more than the other regions, followed by the Eastern region, while the Western region exhibits the lowest levels.

[30]Eastern Region: Beijing, Fujian, Guangdong, Guangxi, Hainan, Hebei, Jiangsu, Liaoning, Shandong, Shanghai, Tianjin, and Zhejiang; Central Region: Anhui, Heilongjiang, Henan, Hubei, Hunan, Inner Mongolia, Jiangxi, Jilin, and Shanxi; Western Region: Chongqing, Gansu, Guizhou, Ningxia, Qinghai, Shaanxi, Sichuan, Xinjiang, Xizang, and Yunnan.

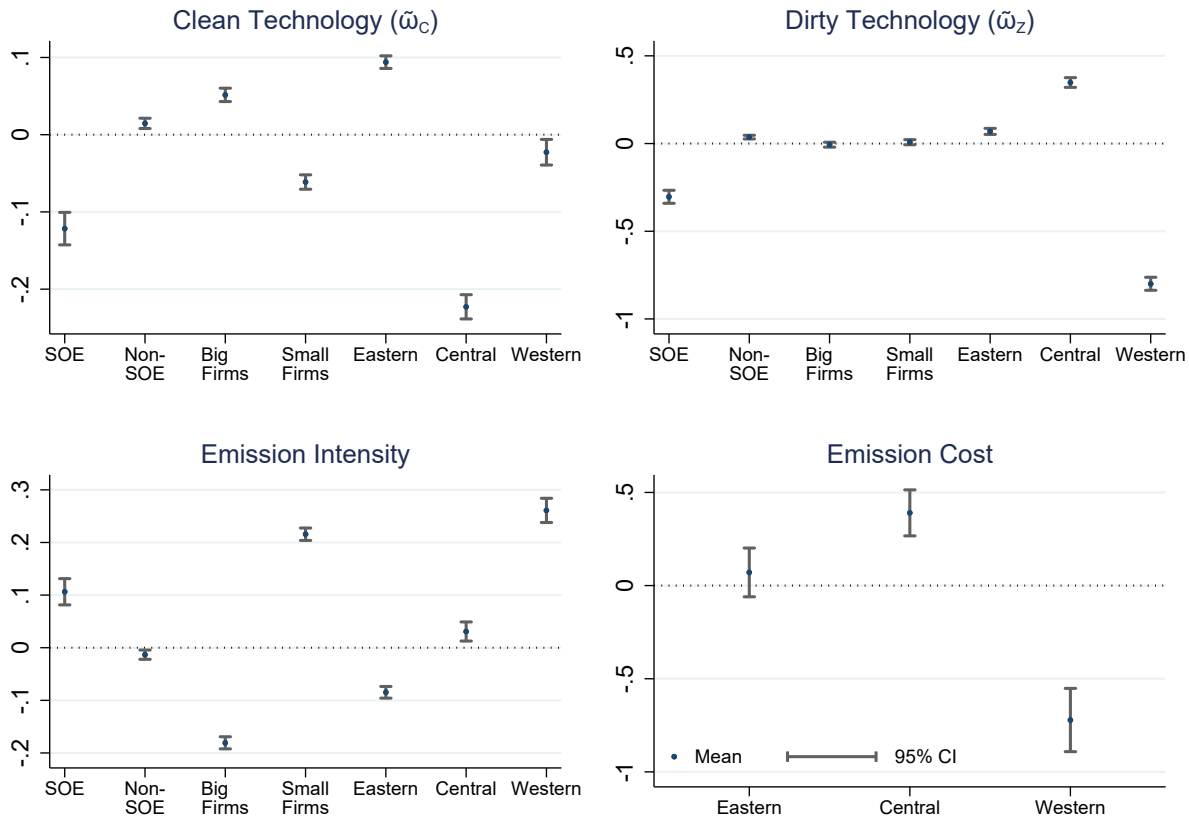


Figure E.1: Heterogeneity in Technology, Emission Intensity, and Emission Costs

Notes: This figure plots the mean values and 95% confidence intervals for clean technology ($\tilde{\omega}_C$), dirty technology ($\tilde{\omega}_Z$), emission intensity ($\ln(Z/Y)$), and emission costs ($\ln(\tau)$) across firm ownership (SOE vs. Non-SOE), firm size (Large vs. Small), and geographic regions (Eastern, Central, Western). Means are calculated from the residuals after partialling out the relevant fixed effects, as described in Appendix E.

Finally, we observe distinct structural divides in environmental outcomes and regulatory burdens. Emission intensity is significantly higher for SOEs than for non-SOEs overall. Furthermore, small firms are significantly more emission-intensive than their larger counterparts. Across geographic areas, the Western region is the most emission-intensive, followed by the Central region, while the Eastern region demonstrates the lowest emission intensity. Correspondingly, emission costs also vary significantly by geography: the Central region faces the highest regulatory costs, followed by the Eastern region, while the Western region experiences significantly lower emission costs than both.

F Further Drivers of Shadow Cost and Technology

F.1 The Impact of Command-and-Control Regulations on Emission Costs

While Section 6.1 demonstrates a strong correlation between our structurally estimated emission costs (τ_{pt}) and price-based discharge fees, we also investigate whether τ_{pt} captures the stringency of major command-and-control policies. Specifically, we examine the "Two Control Zones" (TCZ) policy implemented in 2000 and the SO_2 emission reduction targets set in 2006 under the 11th Five-Year Plan.

Surprisingly, we find little evidence that these nominal policy designations directly elevated firms' perceived shadow cost of emissions. Figure F.1 presents the density distributions of τ across these policy regimes. The distributions for firms in provinces with high versus low 11th Five-Year Plan targets are nearly identical. For the TCZ policy, firms located in designated zones actually exhibit a slightly lower (albeit marginally significant) average shadow cost. This lack of positive correlation is further confirmed by year-by-year difference tests, which remain statistically insignificant throughout the sample period (see Figure F.2).

We formalize these findings in Table F.1. Panel A shows that while the TCZ indicator and SO_2 targets are positively correlated with the shadow cost in univariate regressions, their effect sizes shrink when controlling for the discharge fee (column 4). This pattern persists when examining within-province differences (Panel B). Panel C examines dynamics around the 2006 tightening of SO_2 regulation. Contrary to policy intent, we observe a marginal decline in the shadow cost in provinces with higher reduction goals after their implementation (column 3). While the tightening made the TCZ policy more binding, consistent with Chen et al. (2018), the positive effect on shadow costs is not statistically significant (column 2).

These mixed results contribute to the active debate regarding the efficacy of China's place-based environmental policies. While existing studies document significant real con-

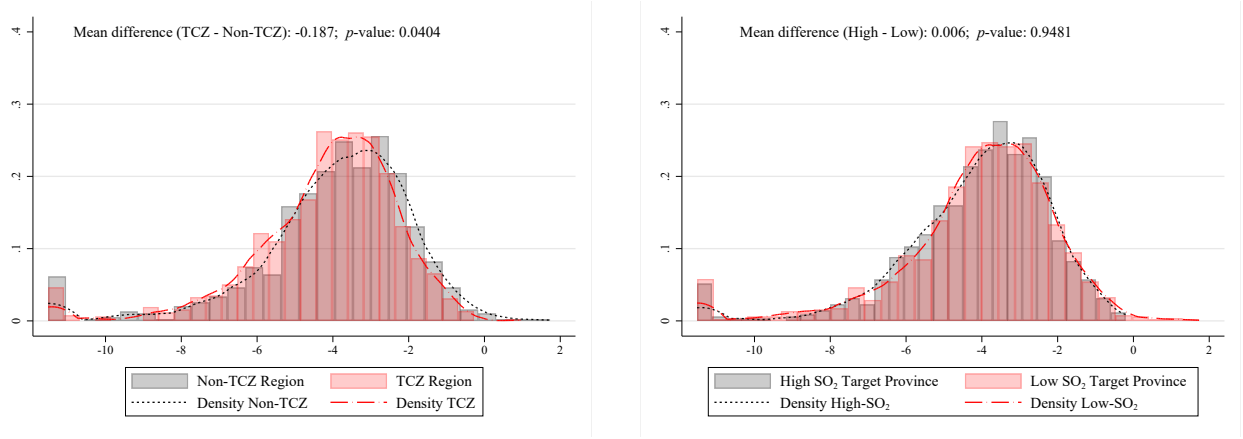
Table F.1: The Impact of Regulations on the Shadow Cost of Emissions

	Dependent Variable: ln(shadow cost)			
	(1)	(2)	(3)	(4)
	Panel A. Year 2004, No Province Fix Effects			
Discharge Fee 2004	0.309***			0.303***
	(0.0782)			(0.0798)
TCZ Region		0.0620		-0.0117
		(0.157)		(0.165)
SO2 reduction target (%)			0.0142	0.00753
			(0.0134)	(0.0149)
Adjusted R^2	0.05	-0.00	-0.00	0.04
	Panel B. Year 2004, with Province Fix Effects			
Discharge Fee 2004	0.345***			0.337***
	(0.0847)			(0.0858)
TCZ Region		0.243		0.200
		(0.154)		(0.147)
Adjusted R^2	0.31	0.26		0.32
	Panel C. All Years, Prefecture and Year Fix Effects			
Discharge Fee 2004 × Post-2005	-0.211**			-0.206**
	(0.0933)			(0.0920)
TCZ Region × Post-2005		0.0200		0.110
		(0.153)		(0.156)
SO2 reduction target (%) × Post-2005			-0.0193*	-0.0176
			(0.0116)	(0.0119)
Adjusted R^2	0.52	0.53	0.53	0.52
No. of Obs.	2043	2080	2080	2043
No. of Prefectures	292	305	305	292

Notes:

(i) This table presents OLS regression results examining the relationship between our estimated log shadow cost of emissions ($\ln \tau$) and various measures of environmental regulation. The unit of observation is the prefecture-year. "TCZ Region" is an indicator variable equal to one for prefectures designated as Two Control Zones. "SO_2 Target" is the provincial-level SO_2 reduction target (%) mandated by the 11th Five-Year Plan. "Discharge Fee 2004" is the log of the prefecture-average emissions fee calculated from the 2004 ASIE. Panel C includes interactions with a post-2005 indicator variable.

(ii) Robust standard error are shown in parentheses in Panels A and B. In Panel C, standard errors clustered at the prefecture level are shown in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$.



(a) Distribution of τ by TCZ

(b) Distribution of τ by SO_2 Target

Figure F.1: Comparison of τ Distributions by TCZ and SO_2 Reduction Target Policies

Notes: This figure plots the kernel density distributions of the log estimated shadow cost of emissions ($\ln \tau$) for firms grouped by two major command-and-control environmental policies. The sample pools all firm-year observations from 2001–2008. Panel (a) compares firms located in prefectures designated as part of the Two Control Zones (TCZ) policy to those in non-TCZ prefectures. Panel (b) compares firms in provinces with high versus low SO_2 reduction targets under the 11th Five-Year Plan.

sequences of the TCZ policy, including health improvements (Tanaka, 2015), shifts in foreign investment (Cai et al., 2016), and heterogeneous firm-level productivity gains (Lu and Pless, 2026), our finding that the policy dummy does not cleanly predict marginal emission costs underscores the limitations of nominal policy indicators. As Lu and Pless (2026) explicitly note, the TCZ designation was not randomly assigned but targeted prefectures that were historically more industrialized and severely polluted. Consequently, these designated zones were precisely the regions with historically weak environmental enforcement and low baseline shadow costs of pollution. This endogenous targeting, combined with lax local enforcement prior to top-down bureaucratic mandates (Chen et al., 2018), explains why a nominal policy dummy does not necessarily correspond to a higher effective marginal cost at the firm level in a cross-sectional comparison. Ultimately, this highlights the critical advantage of our structural approach, which recovers the *de facto* economic cost facing the firm rather than relying on binary policy designations.

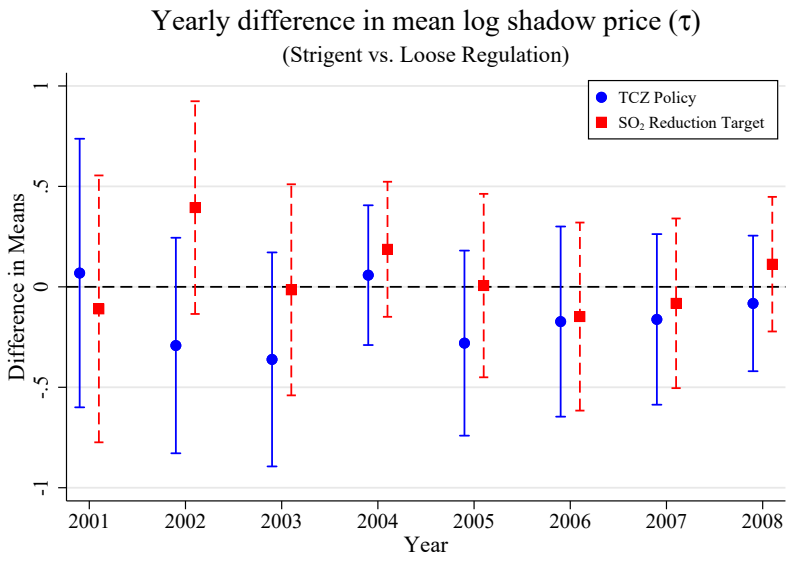


Figure F.2: Differences in Shadow Cost of Emissions by Regions with Stringent vs. Loose Regulations

Notes: The figure shows the estimated year-by-year difference in the mean log shadow cost ($\ln \tau$) for two policy comparisons: TCZ vs. non-TCZ provinces (top panel), and high vs. low SO_2 target provinces (bottom panel). Points are the estimated differences, and bars are 95% confidence intervals. The confidence intervals always include zero, indicating no statistically significant difference in any year.

G Additional Counterfactual Simulations

G.1 Higher Discharge Fee Regime

We first examine the impact of a nationwide tightening of price-based regulation by doubling the prefecture-level emissions discharge fee for all firms in all years. To quantify the effect of this policy, we follow a three-step simulation procedure for each firm-year observation. First, we use the estimated elasticity from our regression analysis (Table F.1) to predict the change in the firm's shadow cost τ_{pt}, resulting from a 100% increase in the fee. Based on an estimate of 0.3, this directly increases the firm's shadow cost of emissions by 30%. Second, we feed the updated shadow cost path into our estimated VAR(1) equations for technology evolution to predict the induced changes in clean-augmenting ($\tilde{\omega}_{Cjt}$) and dirty ($\tilde{\omega}_{Zjt}$) technologies. Finally, holding all other parameters and state variables (e.g., capital, demand shifters) at their baseline levels, we solve each firm's static profit maximization problem under these new counterfactual technology and price paths. This

yields the firm's optimal choices of labor (L_{jt}^{CF}) and emissions (Z_{jt}^{CF}), and the resulting counterfactual output (Q_{jt}^{CF}). By aggregating these new firm-level outcomes, we can assess the policy's impact on aggregate emission intensity, total emissions, and total output.

The results of this simulation, illustrated in Figure 5, show a substantial reduction in aggregate emission intensity. In 2008, the last year of sample period, we find that the doubled emissions fee would lead to a 22.7% reduction in aggregate emission intensity. Our framework allows us to decompose this total effect into a direct cost effect and an indirect effect operating through induced technical change. We find that the reduction in emission intensity is driven almost exclusively by the direct cost effect, which accounts for 93.6% of the decline by inducing firms to substitute away from the now more expensive polluting input. The indirect effect through induced technical change accounts for only 6.4% of net emission intensity change. While the policy does spur progress in dirty technology, its efficiency-enhancing potential is partially offset by a powerful countervailing substitution effect that encourages firms to use the now more productive polluting input. This simulation thus highlights that under this policy, the powerful effect of price-based environmental regulations operates almost entirely by shifting firms' static input choices, rather than through dynamically directing the path of innovation toward a net cleaner outcome.

G.2 Uniform Emission Cost Regime

A central goal of market-based environmental policy, such as a carbon tax or a cap-and-trade system, is to equalize marginal abatement costs across all polluters, thereby achieving emission reductions at the lowest possible aggregate cost. This counterfactual exercise explores the implications of such systems by simulating two uniform emission price regimes. The first is a **national uniform price** regime, in which we solve for a single, nationwide emission cost (τ_t) each year that holds total national emissions constant at their observed baseline level. The second is a **provincial uniform price** regime, in which we solve for a province-specific price (τ_{pt}) each year that holds total provincial emissions

constant. In each scenario, holding firm-level technologies and other state variables fixed at their baseline levels, we re-solve the firm's static profit maximization problem to determine the optimal reallocation of output and emissions.

Table G.1 presents the aggregate outcomes under these three regimes (*status quo*, national uniform price, and provincial uniform price). Because aggregate emissions are held constant by construction, the reported reductions in overall emission intensity directly reflect significant increases in aggregate value added. This represents a classic allocative efficiency gain. In the baseline, firms in high-τ prefectures operate at a high MPE, while firms in low-τ prefectures operate at a low MPE. Implementing a uniform price equalizes the MPE across all firms, inducing a reallocation of production and emissions from low-MPE firms toward high-MPE firms, thus achieving the same environmental outcome while generating more economic output. The national uniform price scheme, by allowing for reallocation across all firms nationwide, generates the largest efficiency gains, driving a 30.8% reduction in aggregate emission intensity relative to the baseline in 2008.

Table G.1: Emission Intensity under Uniform Price Regimes

Year	**Baseline**	**National Uniform Price**	**Provincial Uniform Price**
	Level	Level (% Change)	Level (% Change)
2001	8.50	6.04 (-28.9%)	7.08 (-16.6%)
2008	2.77	1.92 (-30.8%)	2.33 (-16.0%)

Notes: This table compares aggregate emission intensity under the observed baseline ("Status Quo") with two counterfactual uniform pricing regimes: a national uniform price and a provincial uniform price. Both counterfactuals hold total annual emissions constant at their observed baseline levels. Emission intensity is measured in tons per million RMB of value added. The percentage change is calculated relative to the baseline in the same year.

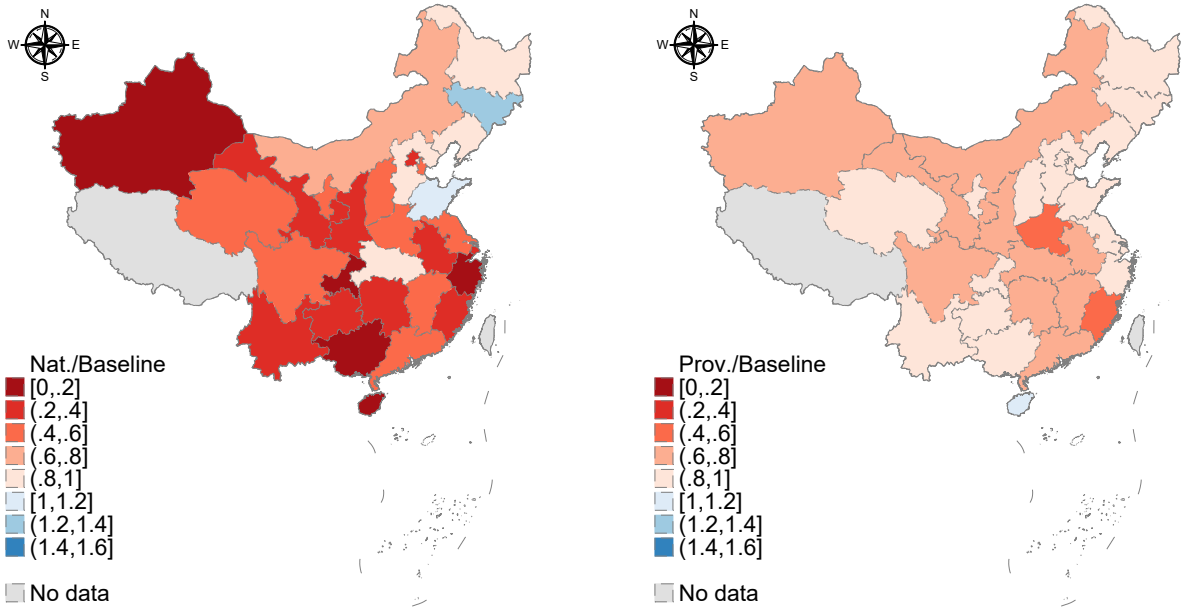
Figure G.1 visualizes the provincial-level consequences of the two uniform pricing regimes. Each panel plots, by province, the ratio of aggregate emission intensity under the counterfactual to the observed baseline. Under the national uniform price (Subfigure G.1(a)), a single nationwide price equalizes the MPE everywhere, generating a redistribution of emissions and output across provincial boundaries. Dirty production flows

toward provinces and firms that can deploy emissions most efficiently. As a result, 28 of the 30 provinces experience substantial reductions in aggregate emission intensity, with the median province seeing emission intensity fall by 56 percent. The cumulative effect of this cross-provincial reallocation is the aggregate gain documented in Table G.1. The two exceptions are Shandong and Jilin, where aggregate emission intensity rises by 20 and 34 percent, respectively. These provinces faced relatively high emission costs in the baseline, so the move to a single nationwide price lowers their effective emission cost. The resulting substitution toward dirty inputs raises local emission intensity by more than the within-province equalization of the MPE can offset.

In contrast, the provincial uniform price (Subfigure G.1(b)) holds total provincial emissions fixed at the baseline while equalizing the MPE across firms within each province. By construction, aggregate emission intensity falls in every province. The induced within-province reallocation raises aggregate provincial output, mechanically lowering emission intensity Z/Q. The median province experiences a 17 percent reduction, with the largest reductions, exceeding 40 percent, concentrated in central and southeastern provinces such as Henan, Fujian, and Guangdong.

G.3 Short- and Long-run Implications of Porter Hypothesis

In this section, we examine the dynamic impact of increasing discharge fees on firms' marginal costs of production. As indicated in Table F.1, a 100% increase in discharge fees is associated with a 30% increase in the shadow price of emissions. This rising shadow price operates through two distinct channels: a *direct channel* affecting relative input costs (static substitution), and an *indirect channel* driven by cost-induced innovation (dynamic technical change). Both channels are structurally captured in our empirical model. We are particularly interested in the latter, as it relates directly to the Porter Hypothesis (Porter and van der Linde, 1995), which posits that well-designed environmental regulations can stimulate innovation, potentially offsetting compliance costs and enhancing competitiveness.



(a) National Uniform Price

(b) Provincial Uniform Price

Figure G.1: Provincial Aggregate Emission Intensity under Uniform Price Regimes

Notes: Each map plots, by province, the ratio of aggregate emission intensity under the counterfactual regime to aggregate emission intensity under the observed baseline, where provincial aggregate emission intensity is computed as total emissions divided by total output summed over the sample period 2001–2008. Panel (a) reports outcomes under a single nationwide emission price; panel (b) reports outcomes under emission prices that vary across provinces but are uniform within each province. Red shading indicates provinces whose aggregate emission intensity falls below baseline (ratio below one), with darker reds corresponding to larger reductions; blue shading indicates provinces whose aggregate emission intensity exceeds baseline (ratio of one or more). All outcomes are aggregated from firm-level simulations.

To isolate the impact of emissions pricing from other confounding factors, we construct a ceteris paribus counterfactual simulation. We utilize data from 2008 to define the base period ($t = 0$). We first compute the firm-specific steady-state values of the innovation terms, ω_{Hj0} and ω_{Zj0}, using the VAR estimates from equations (20) and (21) and setting the initial lagged shadow prices of emissions equal to their observed 2008 levels. In subsequent simulation periods, we hold capital stock, wages, value-added levels, and the components of the productivity evolution process (γ and ξ) fixed at their base-year levels. Under this setup, absent any policy intervention, the marginal cost of production would remain constant over time.

We use the marginal cost of production in the base period (MC_{j0}) as our reference point. In period 1 ($t = 1$), we introduce a permanent, exogenous policy shock: a nationwide 100% increase in discharge fees, which translates to a permanent 30% increase in the shadow price of emissions (τ_{pt}) across all locations. Consequently, in all subsequent years, the lagged values of τ_{pt} utilized in the technology evolution process also reflect this 30% increase.

Crucially, in period 1, the shock operates exclusively through the direct channel. Because the innovation response depends on lagged shadow prices, which remain at pre-policy levels at the start of $t = 1$, technology remains initially unchanged. Thus, the immediate effect is an unambiguous rise in marginal costs as firms face higher input prices without the benefit of improved efficiency. From period 2 onward, the indirect channel activates; firms adjust their innovation strategies in response to the higher shadow prices observed in the previous period, driving changes in ω_{Hjt} and ω_{Zjt}.

To investigate the heterogeneity of this impact, we divide the sample into two subsamples based on the base-year level of ω_{Zj0}. We distinguish between "High ω_Z" firms (those with ω_{Zj0} above the median) and "Low ω_Z" firms (those below).

Figure 4 plots the evolution of the mean normalized marginal cost (MC_{jt}/MC_{j0}) for the full sample and both subsamples from period 1 to 15. The results illustrate a distinct dynamic adjustment path driven by the two channels. In period 1, the direct channel

dominates. For the full sample (solid line), the average marginal cost rises by approximately 13.3%. However, this aggregate figure masks significant heterogeneity. The "High ω_Z" firms (dashed line), which are structurally more reliant on emissions due to the bias of their technology, experience a sharp 23.8% increase in costs. In contrast, the "Low ω_Z" firms (dotted line) are largely insulated, seeing only a modest 2.78% increase.

From period 2 onward, the indirect channel activates. Firms adjust their innovation strategies in response to the higher shadow prices, leading to a decline in marginal costs across the board. This recovery is most pronounced for the "High ω_Z" group, which innovates rapidly to mitigate their high exposure to the regulatory cost. By the new steady state, the gap between the groups narrows, though the initial technical bias continues to dictate the long-run cost burden. In the new steady state, the mean marginal costs for the full sample, "High ω_Z" firms, and "Low ω_Z" firms settle at levels 8.1%, 13.9%, and 2.34% higher than the baseline, respectively.

These results suggest that while regulation imposes an initial cost burden, induced innovation significantly dampens this effect over time. Specifically, for the full sample, the indirect channel offsets approximately 40% of the initial increase in marginal costs caused by the regulation. This finding provides strong empirical support for the mechanism underlying the Porter Hypothesis: firms indeed innovate to mitigate the burden of environmental compliance, though in this specification, the innovation is not sufficient to fully offset the regulatory cost.

H Additional Figures and Tables

Table H.1: Summary Statistics for Variables in the Production Function

Variable	**Mean**	**S.D.**
Value added	42,547.61	86,483.00
Capital	58,713.18	130,580.94
Employment	460.87	657.66
Wage bill	6,726.72	12,728.76
Intermediate inputs	110,637.15	240,581.22
SO_2 Emissions (ton)	226.90	1,506.69
Output price index	1.13	0.33
Price index of intermediate inputs	1.15	0.17
Number of observations	105,356	

Notes: The price indices are normalized to the year 2000 as the base year. The units of value added, capital, wage bill, and intermediate inputs are expressed in thousands of RMB.

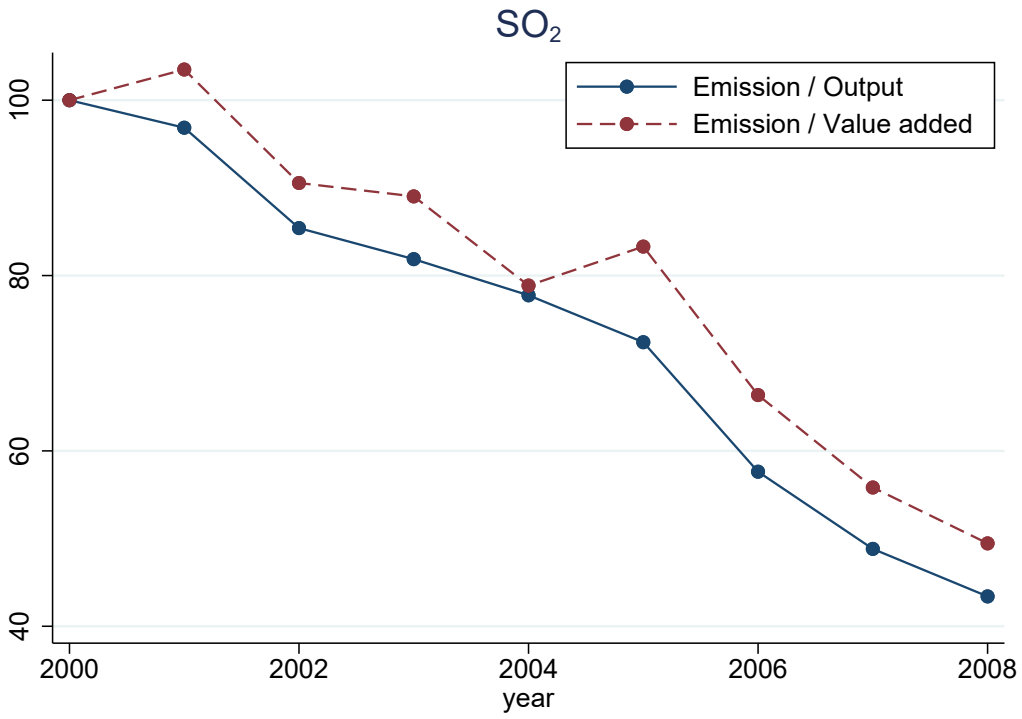


Figure A.1: Industrial SO_2 Emission Intensity in China

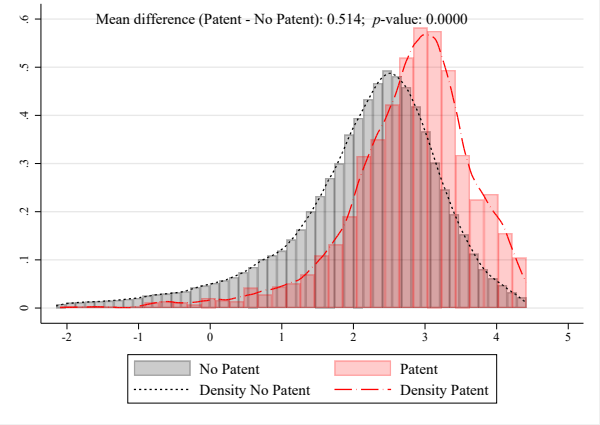
Notes: This figure plots the time trend of two measures of aggregate SO_2 emission intensity (Emissions/Output and Emissions/Value Added) in the Chinese manufacturing sector. Emission data are from China's Environmental Statistical Database (CESD). Real output and value added are from the Annual Survey of Industrial Enterprises (ASIE), deflated using industry-specific price indices from Brandt et al. (2012). All series are indexed to 100 in the year 2000.

To interpret these time trends, we examine the correlation between the three components and firm-level emission intensity. As shown in the first column of Table H.2, all three—clean technology, dirty technology, and the shadow cost—are negatively correlated with emission intensity, suggesting that all are plausible contributors to the observed cleanup in Chinese manufacturing sector. However, these simple correlations are only suggestive. While the correlation of clean technology with emission intensity is more than three times that of dirty technology (-0.576 vs. -0.174), this does not directly imply a proportionally larger impact, as their ultimate effects on output and input choices are mediated through the CES production function. To formally assess the relative importance of each factor, we conduct a full structural decomposition analysis.

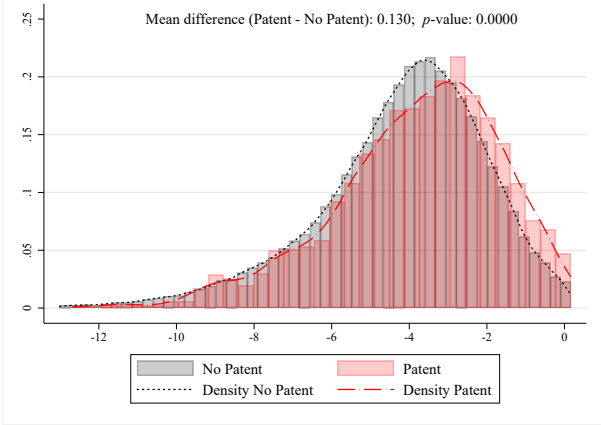
Table H.2: Pair-wise Correlation of Emission Intensity, Technologies, and Emission Cost

Variables	(1) $\ln(Z/Y)$	(2) $\tilde{\omega}_C$	(3) $\tilde{\omega}_Z$	(4) τ
(1) $\ln(Z/Y)$	1.000			
(2) $\tilde{\omega}_C$	-0.576	1.000		
(3) $\tilde{\omega}_Z$	-0.174	0.043	1.000	
(4) τ	-0.217	-0.116	0.459	1.000

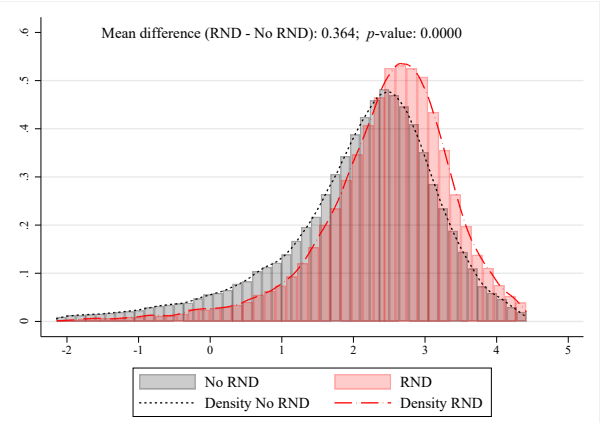
Notes: This table reports pairwise correlation coefficients. The unit of observation is the firm-year for emission intensity ($\ln(Z/Y)$), clean technology ($\tilde{\omega}_C$), and dirty technology ($\tilde{\omega}_Z$). The shadow cost (τ) is at the prefecture-year level and is merged with the firm-level data.



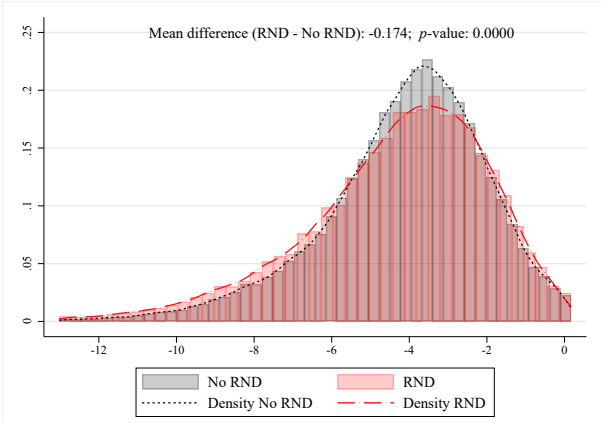
(a) Distribution of $\tilde{\omega}_C$ by Patent Application



(b) Distribution of $\tilde{\omega}_Z$ by Patent Application



(c) Distribution of $\tilde{\omega}_C$ by Positive RD



(d) Distribution of $\tilde{\omega}_Z$ by Positive RD

Figure A.2: Comparison of ω Distributions by Innovation Activities

Notes: This figure plots the kernel density distributions of the estimated firm-level non-neutral technologies, comparing innovative and non-innovative firms. The top panels compare firms based on patenting activity: "Patent" is an indicator for any firm that filed at least one patent during our sample period. The bottom panels compare firms based on R&D: "R&D" is an indicator for a firm reporting any R&D expenditure in a given year. All distributions are pooled across all firm-year observations.

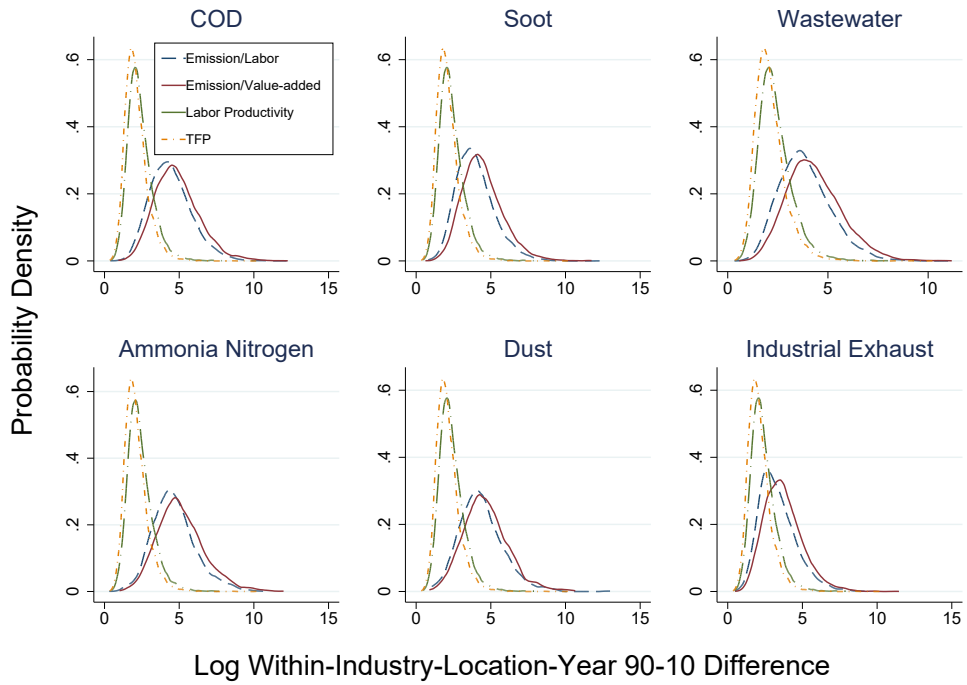


Figure A.3: Dispersion of Environmental and Economic Performance Measures within Industries, Locations, and Years: Other Pollutants

Notes: This figure plots the average dispersion of key firm-level performance metrics. Environmental performance measures are calculated for COD, soot, wastewater, ammonia nitrogen, dust and industrial exhaust. Dispersion is measured as the 90th-10th percentile difference of the log value of each metric, calculated within each four-digit industry-province-year cell. Real value added is computed using industry-specific output deflators from Brandt et al. (2012). Total Factor Productivity (TFP) is estimated following Ackerberg et al. (2015).